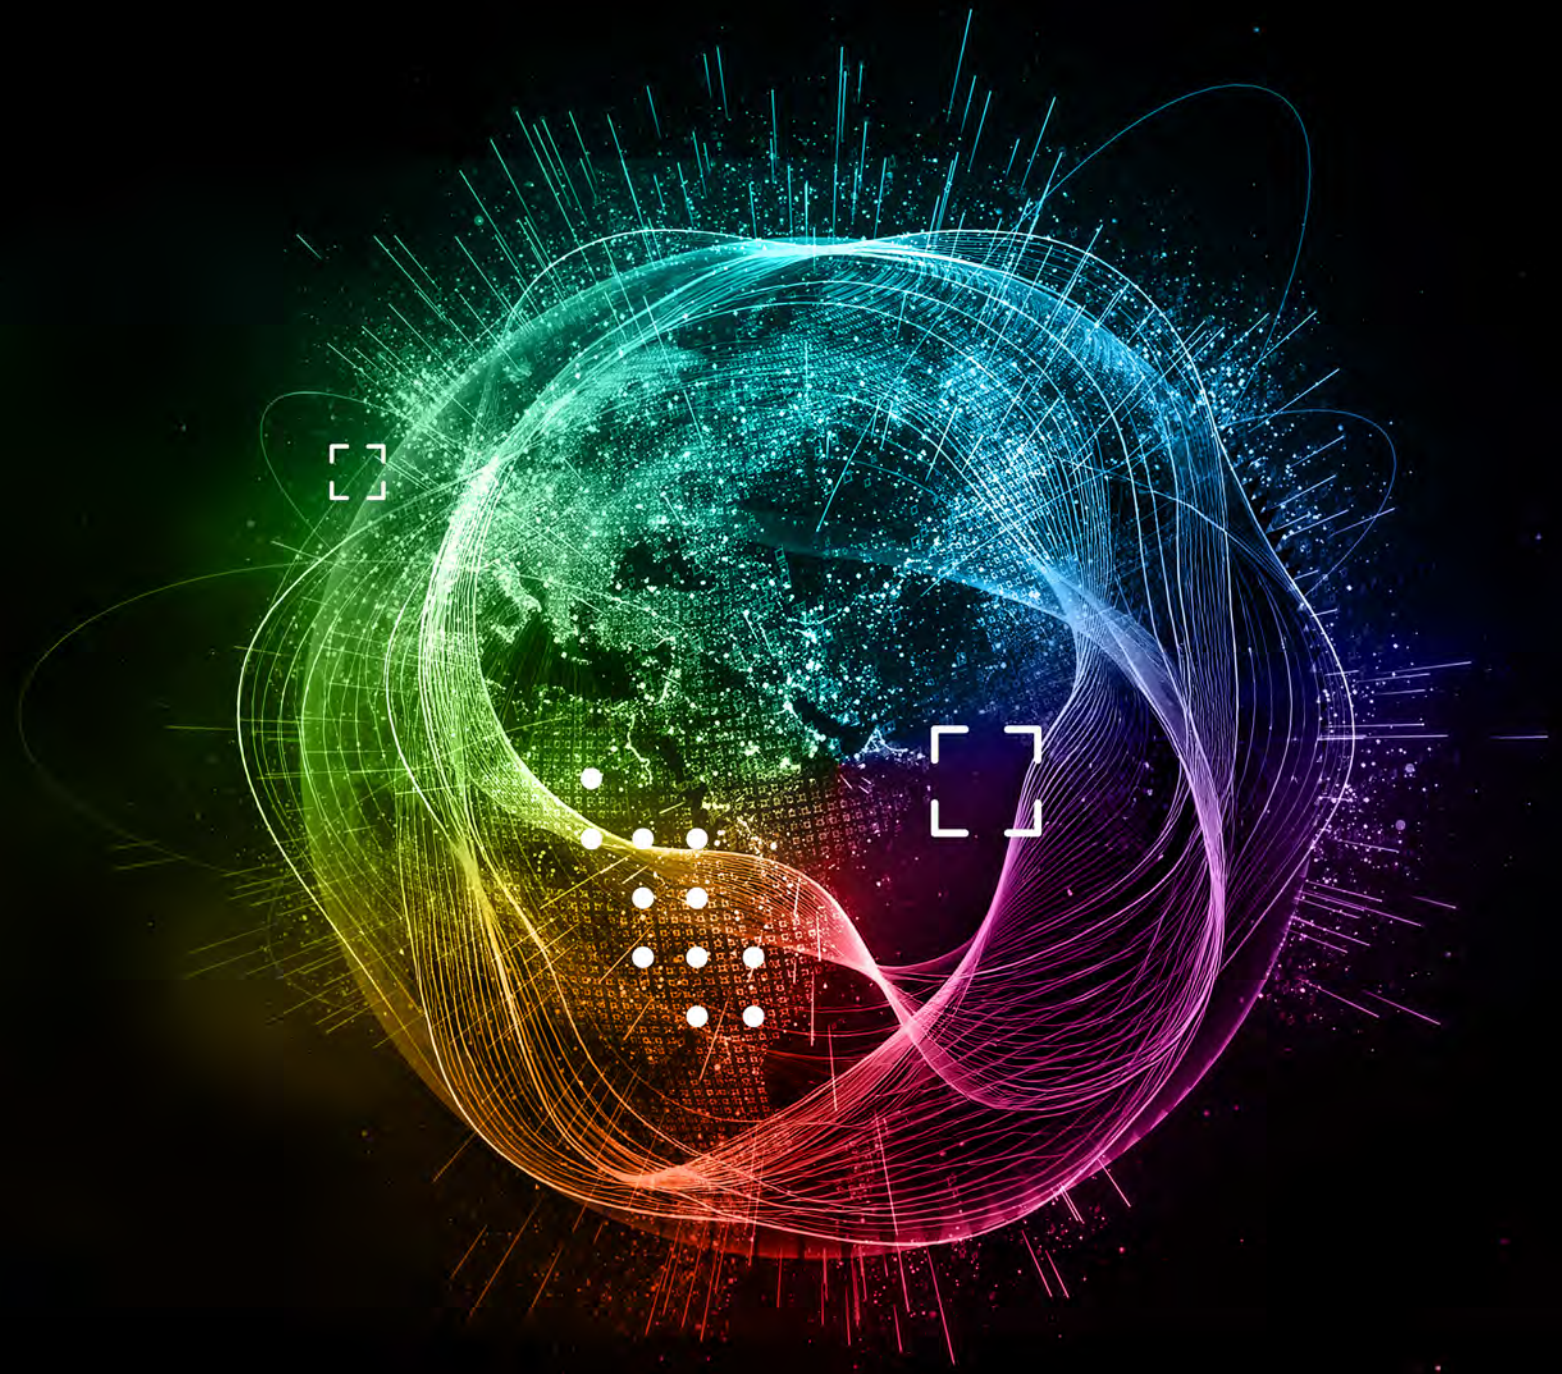
Deloitte.

The AI Dossier

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



The Deloitte AI Institute™

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/AIInstitute

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

What is inside

Consumer
5

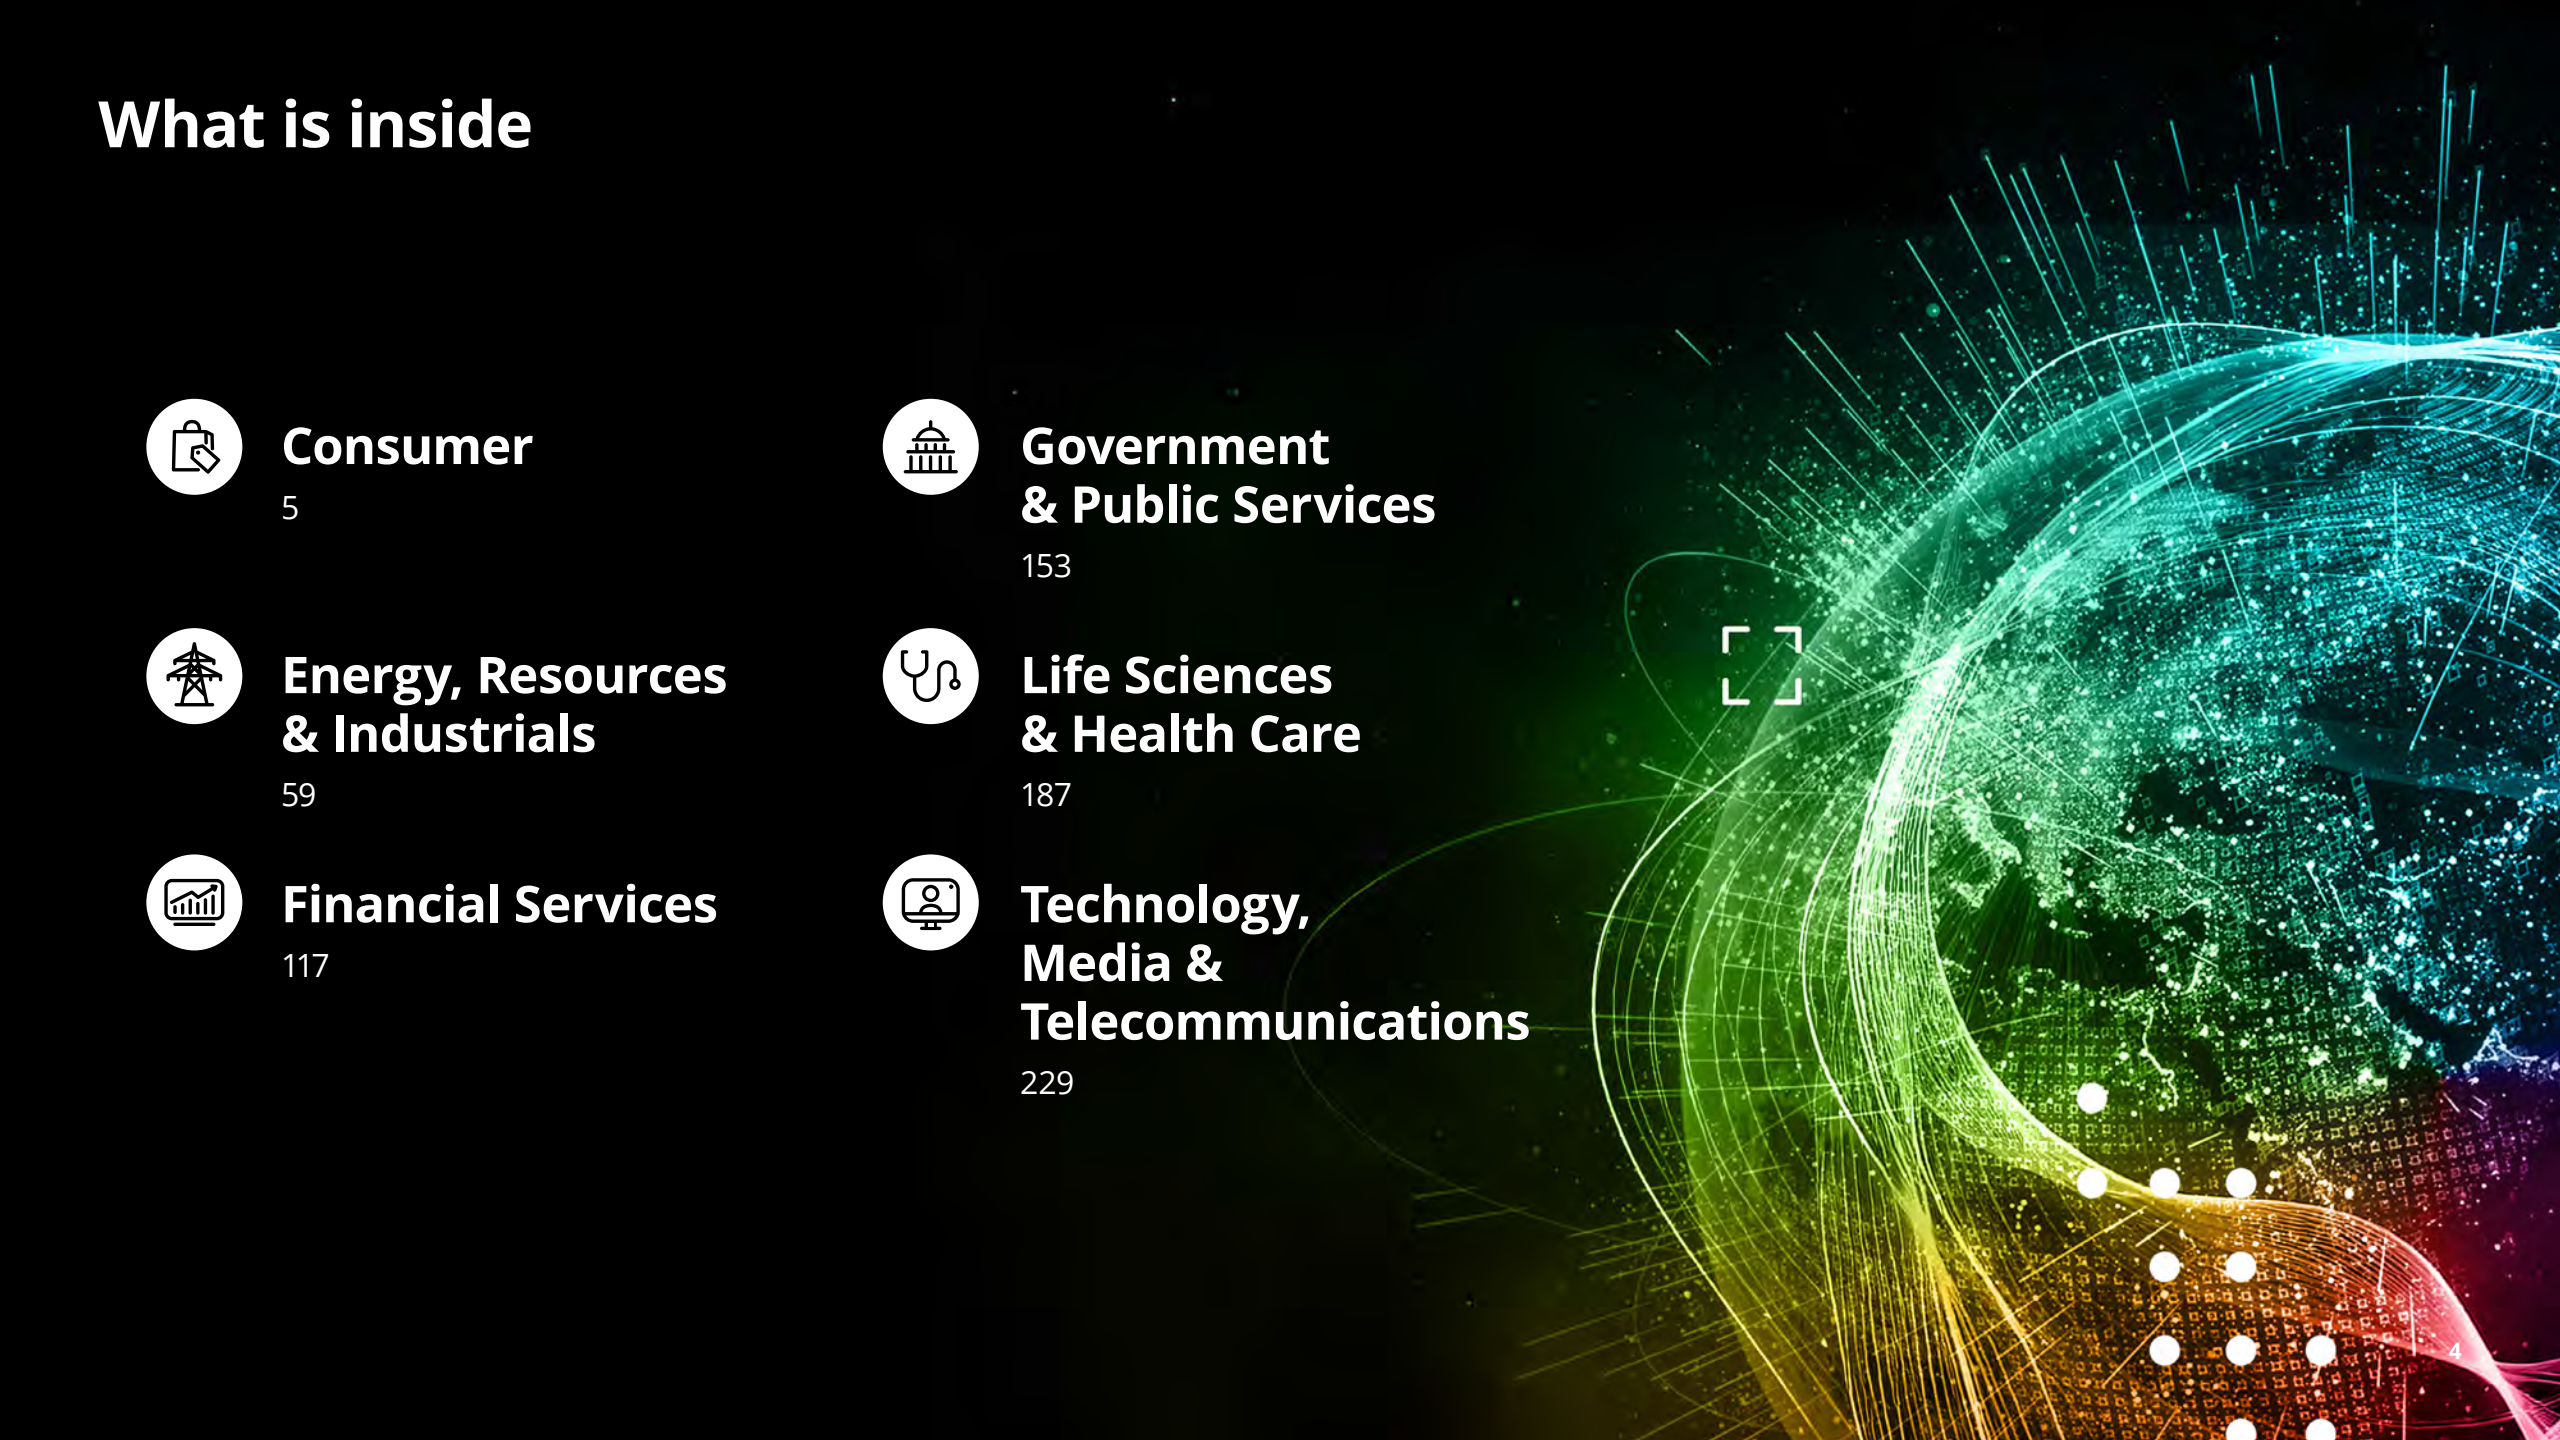
Government & Public Services
153

Energy, Resources & Industrials
59

Life Sciences & Health Care
187

Financial Services
117

Technology, Media & Telecommunications
229



The Consumer AI Dossier



The Consumer AI Dossier

AI has already become an integral part of people's everyday lives, whether they realize it or not. AI innovations are redefining how consumers discover, evaluate, and interact with brands—and compelling new use cases in the consumer space continue to emerge. Generative models are creating troves of rich, personalized content. Autonomous agentic systems are planning and execute tasks on a user's behalf in every part of the business, from product design and pricing to supply chain operations. And emerging physical AI solutions are reshaping how goods move, how stores operate, and how people experience the world around them. These tools are not just enabling better experiences; they are resetting consumer expectations across search, service, commerce, and entertainment.

For consumer companies, this shift creates both opportunity and urgency. AI can drive real-time customer engagement, intelligent automation, and more adaptive decision-making. Large language models, AI agents, and physical AI systems are now capable of handling a wide range business tasks with minimal human intervention, opening new possibilities for operational efficiency and innovation.

However, leveraging these capabilities at scale requires more than technical integration. It demands unified data infrastructure, strong governance, and a willingness to reimagine core business processes. Leading companies are using AI not just to cut costs but to launch new products, redesign customer journeys, and compete on speed, relevance, and personalization.

As regulation evolves and public scrutiny grows, sustainable advantage will come from deploying AI with transparency, oversight, and measurable impact. The winners will not be those who adopt AI the fastest, but those who align it best to strategic goals, operational realities, and consumer trust.

AI innovations are redefining how consumers discover, evaluate, and interact with brands.

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

Dynamic pricing and inventory optimization

Coordinating price and stock decisions in real time

Agentic AI systems can use multiple specialized agents to monitor a wide range of internal and external signals, then dynamically adjust prices, promotions, and inventory to optimize business performance.

ISSUE/OPPORTUNITY

In many retail environments, pricing and inventory decisions are made using fixed rules and periodic adjustments. This approach can leave money on the table when market conditions change quickly. It can also create costly overstocks when demand softens.

Businesses relying on traditional processes can't respond quickly enough to events like a competitor running out of stock, a sudden weather change, or a viral trend shifting demand. Also, by treating pricing and inventory management as distinct processes, many retailers miss opportunities for joint optimization. For example, an item might be discounted without considering replenishment timing, or stock might be held for too long at full price when a strategic promotion could accelerate sell-through.

Agentic AI can unify these activities, with specialized agents continuously collaborating to balance profitability, stock levels, and customer satisfaction.

Tags: **Sales** | **Agentic AI**

HOW AI CAN HELP

Pricing optimization
A pricing agent can continuously learn the price elasticity of each product and track competitor prices, adjusting in real time to capture revenue opportunities, avoid unnecessary markdowns, and react to changing market conditions.

Inventory management
An inventory agent can monitor stock levels across stores and warehouses, factoring in lead times and supply constraints to ensure replenishment decisions align with projected demand and pricing strategies.

Demand forecasting
A demand forecasting agent can analyze signals from internal sales trends, online search patterns, social media, weather forecasts, and local events to anticipate surges or dips in near-term demand.

Promotions and bundling
A promotions agent can design targeted offers and product bundles (e.g., pairing slow-moving items with high-demand products), scheduling them based on real-time sales velocity and inventory.

Collaborative decision-making
All agents in the process share a common situational awareness and negotiate trade-offs. For example, if the demand agent forecasts a surge, the pricing agent might raise prices while the promotions agent delays discounts; conversely, if oversupply is detected, price reductions and targeted promotions might be implemented in specific regions or channels.

Dynamic pricing and inventory optimization

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	Because frequent price changes can be perceived as unfair or arbitrary, dynamic pricing agents should operate within clearly defined policies and thresholds to ensure consistent treatment of customers across channels and regions.
Robust and reliable	Bad data can lead to bad decisions. Agents should be designed with strong data validation and filtering processes to avoid reacting to false signals (such as misinterpreted social trends or inaccurate sales figures).
Transparent and explainable	Dynamic price and promotion changes can confuse customers and internal teams alike. To help address the problem, agents should provide clear reasoning for adjustments, including the data sources and logic used, so pricing and category managers can interpret and communicate the rationale.
Responsible and accountable	Rapid pricing and inventory actions can have strategic and reputational impacts. AI-driven decisions should align with the organization's brand strategy, operating capacity, and regulatory requirements, with final oversight provided by qualified human managers.

POTENTIAL BENEFITS

Increased revenue and margins

Dynamic, coordinated decisions can capture additional profit during high-demand periods and optimize sell-through on slow-moving products.

Reduced waste and overstock

By aligning pricing and replenishment strategies, excess inventory—especially perishables—can be cleared before it becomes unsellable.

Improved customer experience

Timely, relevant promotions increase customer satisfaction and loyalty while maintaining trust in pricing fairness and reducing stockouts.

AI-orchestrated product design

Automated, end-to-end product design powered by AI agents

Agentic AI systems can orchestrate the entire product design lifecycle—from market sensing to concept creation, product development, and iteration—continuously adapting to market changes in real time.

ISSUE/OPPORTUNITY

Traditional product design in the consumer industry is often a linear, stage-gated process that can take months or even years from concept to launch. Also, from hundreds of ideas, often just one or a few options are commercialized. Although this limited approach helps manage complexity, it also slows innovation and impairs an organization's ability to respond quickly to shifting consumer tastes or competitive moves.

The challenge is compounded by siloed functions, with design, sourcing, marketing, and supply chain often operating independently on different data systems and timelines. As a result, valuable insights from sales data, customer feedback, or social trends may not inform product development until it's too late.

Agentic AI can make the entire product design process dynamic and continuously adaptive, reducing time-to-market, unlocking new levels of creativity, and enabling better and faster alignment with what consumers actually want.

Tags: R&D/Product Development | Agentic AI

HOW AI CAN HELP

Market sensing and opportunity identification
A market sensing agent can analyze real-time data from trend reports, social media, consumer sentiment, and competitive intelligence to identify unmet needs and emerging product opportunities.

Concept generation and feasibility analysis
A concept agent can create diverse and innovative product ideas informed by market insights while a feasibility agent evaluates each idea against sourcing options, production cost, manufacturing timelines, and regulatory constraints.

Design development and prototyping
A design agent can produce technical specifications and digital prototypes, enabling rapid iteration and deep product visualization without the time and cost needed to create physical samples.

Validation and dynamic iteration
A validation agent can test designs against historical performance, customer feedback, and simulated market conditions, while a coordination agent can orchestrate updates across product lifecycle management, marketing, and supply chain systems to adjust plans in real time.

AI-orchestrated product design

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Because AI agents might produce concepts that are technically or commercially unviable, outputs should be validated in controlled simulation environments and reviewed through human-in-the-loop processes before advancing to production.

Transparent and explainable

Since design decisions can have major cost, brand, and regulatory implications, agents should provide clear reasoning and evidence for their recommendations, including source data and assumptions.

Responsible and accountable

Products must comply with safety and regulatory requirements. Also, legal questions remain over intellectual property (IP) rights for AI-generated outputs. IP protection and ownership rights can be complex when AI is involved in the creative process. To address such issues, AI-driven design activities should align with brand standards and legal constraints, with final approvals retained by qualified human decision-makers.

POTENTIAL BENEFITS

Faster time-to-market

By enabling rapid concept generation, iterative testing, and digital prototyping, AI can reduce development cycles from months to weeks, allowing brands to respond quickly to market opportunities.

Increased innovation

Generating new and diverse ideas more quickly—in greater volume—expands the creative possibilities for new product development.

Higher product success rates

Innovative design that aligns with real-time market shifts increases the likelihood new products will be a hit with target customers.

Lower development costs

Digital prototyping and early-stage feasibility analysis reduce the need for costly physical samples and late-stage redesigns.

Next-generation store operations

Autonomous in-store coordination to optimize retail execution

Agentic AI systems can coordinate in-store activities by continuously monitoring conditions and taking automated actions to achieve smooth, efficient, customer-responsive operations.

ISSUE/OPPORTUNITY

Running a high-performing retail store involves hundreds of large and small decisions each day: allocating staff to handle peak traffic, restocking shelves when inventory runs low, responding to customer requests, and ensuring that promotional displays are set up correctly. In many cases, these actions are handled reactively, based on direct observation by a manager or sales associate, rather than being driven by data in real time.

This reactive approach can lead to problems such as stock-outs, bottlenecks at checkout, haphazard execution of merchandising plans, and missed sales opportunities—operational frictions that can quickly erode revenue, profitability, and customer satisfaction.

Agentic AI can help stores become highly efficient, semi-autonomous systems—where human associates focus on value-added service and strategic priorities while AI handles routine operational tasks.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Continuous store sensing
A store sensing agent can monitor real-time data streams from cameras, IoT sensors, POS systems, and digital twins to track foot traffic, queue lengths, inventory levels, associate availability, and local events.

Automated compliance monitoring
A compliance agent can use computer vision and sensor data to monitor planogram adherence, promotion execution, and safety hazards, triggering immediate corrective actions as needed.

Dynamic task allocation
A task management agent can reprioritize and assign tasks such as restocking, returns processing, online order pickup, or promotional setup based on current demand and available labor.

Coordinated multi-agent oversight
A store manager agent can oversee all other agents, resolving conflicts, optimizing labor deployment, and coordinating with upstream systems such as workforce management, ERP, and order management platforms.

Next-generation store operations

MANAGING RISK AND PROMOTING TRUST

Robust and reliable
Because retail environments often involve incomplete or questionable sensor data, agents should be designed to function effectively under imperfect conditions, with human escalation protocols in place when AI output is uncertain or fails.

Transparent and explainable
To avoid blind reliance on AI agents, managers and staff should have access to clear explanations of the rationale behind AI-driven decisions.

Responsible and accountable
Store operations must comply with a wide range of standards, including labor laws, safety requirements, and company policies. All AI-driven actions should align with these standards, with ultimate accountability retained by human supervisors.

POTENTIAL BENEFITS

Improved labor productivity
Associates can spend more time on high-value tasks and less time on low-value activities and manual monitoring.

Higher sales conversion and customer satisfaction
Automated store operations can help optimize in-stock rates, checkout lines, and service levels.

Manager bandwidth for strategic decision-making and leadership
Store managers can focus on performance improvement, training, and coaching, rather than spending their days fighting operational fires.

Autonomous supply chain operations

Using AI agents to improve efficiency in global automotive supply chains

Agentic AI systems can improve the efficiency and resilience of automotive supply chains by using specialized agents to forecast demand, optimize planning, detect disruptions, and autonomously adjust operations.

ISSUE/OPPORTUNITY

Automotive supply chains are complex and vulnerable to disruptions from shifting demand, supplier delays, logistics bottlenecks, and external forces such as pandemics, policy changes, and weather. Traditional supply chain processes rely heavily on periodic data reviews and manual adjustments, which often cannot keep pace with sudden changes in demand and supply. These limitations can lead to higher costs, supply delays, and increased operational risk.

With tariffs, global market volatility, and various sustainability pressures (including electrification) reshaping the industry, automakers need supply chains that are dynamic, predictive, and capable of adapting in real time. Agentic AI provides a pathway to autonomous supply chain operations that can be more flexible, efficient, and resilient.

Tags: **Procurement/Sourcing & Supply Chain** | **Agentic AI**

HOW AI CAN HELP

Data readiness and transformation
A data readiness agent can perform quality checks and identify exceptions, while a data generator agent can transform raw inputs into structured data for optimization.

Optimization and demand mapping
A suggestion optimization agent can run AI/ML models to autonomously identify the best-performing options, while a demand mapping agent can align demand signals with the correct product configurations.

Validation and explainability
A validation/explanation agent can review outputs, ensure consistency, and provide transparent reasoning to supply chain managers for greater trust in the system's recommendations.

Autonomous supply chain operations

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Inaccurate recommendations can disrupt production or logistics. AI agents should be validated regularly against real-world outcomes and monitored to provide ongoing reliability.
Transparent and explainable	Because supply chain managers need to understand how recommendations are generated, agents should be designed to provide clear explanations of their optimization logic and how the underlying data was used.
Responsible and accountable	Given the risk of AI agents taking inappropriate or inconsistent action, humans should have the final responsibility for supply chain adjustments.

POTENTIAL BENEFITS

Greater supply chain resilience

Proactive detection of bottlenecks and disruptions—coupled with real-time adjustment—can minimize costly delays and maintain production and supply continuity.

Faster, data-driven decisions

Dynamic demand forecasting and optimization helps supply chains respond quickly to shifting market demand and challenging operational conditions.

Autonomous warranty adjudication

Using AI agents to automate warranty claims processing

Agentic AI systems can streamline the adjudication of automotive warranty claims by using specialized agents to assess claim filings, flag anomalies, generate documentation, and support human adjudicators.

ISSUE/OPPORTUNITY

For automakers, warranty adjudication is a crucial function that directly affects costs, customer satisfaction, dealer relationships, and compliance. Today, the process often involves multiple handoffs, manual reviews, and inconsistent action. This hampers efficiency and speed, increases the risk of undetected fraud, damages relationships with customers and dealers, and drives up administrative costs.

Manual adjudication also makes it difficult to conduct a comprehensive and consistent review of claims. Limited time and resources mean that only a subset of claims can be deeply reviewed. Potential errors, fraudulent claims, and incomplete filings could fall through the cracks. To mitigate such problems, automakers need a more efficient, scalable, and consistent way to manage warranty claims while preserving fairness and transparency.

Tags: **Compliance & Risk** | **Agentic AI**

HOW AI CAN HELP

Data validation and fraud detection
AI agents can review incoming claims for completeness, identify unusual patterns, and flag potential fraud, waste, or abuse before the claim progresses.

Customer and claims history analysis
Agents can cross-check claims with customer and vehicle history, uncovering relevant information that strengthens the adjudication process.

Documentation and reporting
An agent can automatically generate detailed reports for human adjudicators, reducing manual work and providing decision-makers with the necessary context to make informed decisions.

Decision support and denial drafting
Agents can propose denial reasons with clear justifications and then draft denial letters for human approval, helping adjudicators operate more efficiently and improving consistency across claims.

Autonomous warranty adjudication

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Errors in root cause analysis or improper routing can lead to incorrect claim outcomes. Actions of AI agents should be validated against historical claim datasets and continuously monitored for performance accuracy.

Transparent and explainable

Warranty decisions can have a big impact on customers, dealers, and regulatory compliance. As such, AI agents should provide clear reasoning for why a claim was flagged, routed, or recommended for denial.

Responsible and accountable

Given the significant financial and reputational risks associated with warranty adjudication, final decisions should be left to human adjudicators, with AI agents providing decision support rather than operating as fully autonomous systems.

POTENTIAL BENEFITS

Increased efficiency and coverage

Automated warranty adjudication allows for broader and deeper claim analysis—with less time and effort—improving both efficiency and accuracy.

Improved consistency and fairness

Standardized AI-supported processes reduce variability in claim outcomes, helping to make adjudications more consistent, transparent, and fair.

AI assistant for vehicle buying and leasing

Guiding consumers to the right car with personalized, multi-agent assistance

Agentic AI systems can streamline the vehicle buying and leasing process through specialized agents that evaluate numerous purchase options and provide hyper-personalized recommendations to consumers.

ISSUE/OPPORTUNITY

Car buying and leasing is a complex and expensive decision that involves comparing different vehicle models, feature availability, financing structures, total cost of ownership, and dealer inventory. Consumers often find this process intimidating and confusing, which leads to delays, dissatisfaction, or switching to competing brands when their preferred model is unavailable.

For OEMs and dealers, missed sales opportunities during this critical decision window represent lost revenue and weakened customer loyalty. The challenge is compounded by lack of visibility into production pipelines or limited ability to match customer preferences with available inventory, leaving dealers struggling to balance consumer demand with real-world availability.

Tags: **Sales** | **Agentic AI**

HOW AI CAN HELP

Personalized vehicle matching
A central vehicle search and advisor agent helps customers identify models that align with their preferences, budgets, and usage needs, whether buying, leasing, or exploring certified pre-owned (CPO) options.

Inventory and production visibility
An OEM agent analyzes vehicles in production pipelines and offers booking options, helping OEMs capture demand even when current dealer inventory does not meet customer criteria.

Comprehensive financial analysis
A buy agent analyzes total cost of ownership—including loan payments, depreciation, maintenance, and taxes—while a lease agent evaluates lease terms, monthly payments, and conditions to help customers compare financing options with full transparency.

Streamlined communication and support
A communication agent delivers supporting documents, sends summaries, and ensures smooth integration with dealership systems, reducing customer effort and follow-up calls.

AI assistant for vehicle buying and leasing

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable
Because purchase and lease decisions involve major financial commitments, agents need to provide clear explanations of cost breakdowns, assumptions, and trade-offs in their recommendations.

Robust and reliable
Errors in inventory matching or financial analysis can erode customer trust. AI agents should be validated against real-world dealership and OEM data and continuously updated to ensure accuracy.

Responsible and accountable
AI agents can have a significant influence on consumers' car-buying decisions. As such, their outputs should be positioned as guidance tools, with customers and dealer staff retaining final responsibility for understanding and confirming selections.

POTENTIAL BENEFITS

Increased sales conversion
By identifying various inventory and production pipeline options, OEMs can help reduce the number of customers they lose to competitors when the preferred model is unavailable.

Enhanced customer experience
Personalized recommendations and simplified comparisons improve decision-making and reduce the stress of navigating the complexities of financing and leasing.

Reduced dealer workload
Automated handling of routine inquiries reduces a dealership's call volume and required manual effort, allowing sales staff to focus on higher-value interactions with customers.

Marketing content assistant

Content generation

AI can be used to enable the creation of efficient, consistent, and personalized content across a range of modalities.

ISSUE/OPPORTUNITY

Companies face a significant challenge in managing and optimizing marketing content. With hundreds of websites for brand portfolios, each in dozens of languages, companies struggle to allocate enough time and resources to create customer group-specific product descriptions, images, video, and even audio. Enterprises also wrestle with consistency across descriptions, imagery, ads, and other media, and the materials may not always be optimized for the necessary purposes (e.g., product descriptions for search versus e-mail). Companies need a method to provide a seamless and personalized brand experience across different ecosystems and touchpoints.

Tags **Marketing**

HOW AI CAN HELP

Next-gen content generation
With AI, the enterprise can create product descriptions, imagery, video, and more much faster and more consistently than with existing tools and processes.

Personalization at scale
AI models can draw from multimodal data (e.g., text, image, geospatial data) to create personalized and contextually relevant content. The model can be used to catalog content and adapt content and user flow based on language, region, and customer behavior trends.

Assisting compliance
Due to the consistency AI enables across modes, languages, and contextual factors, the enterprise can enhance regulatory compliance for materials across different geographies, cultures, and topics.

Marketing content assistant

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

While tasked with producing superior marketing materials, AI systems may invent inaccuracies, which will lead to poorer customer engagement and outcomes.

Fair and impartial

Biases in the data (e.g., due to incomplete datasets) could lead to unequal quality of content in the face of different geographical or cultural factors.

POTENTIAL BENEFITS

Catering to the customer

By tailoring content and the user experience based on language, region, and customer preferences, the enterprise can drive customer satisfaction and loyalty.

Revenue growth

Personalized content can promote higher engagement, traffic, and conversions through tailored and relevant marketing experiences.

Cost efficiency

Using AI for content creation allows the enterprise to develop and maintain content at scale without the costs associated with commensurate human labor.

Planning for promotions

Reimagined trade promotions

AI can be used to prepare promotion plans, negotiation materials, pre-works, and pitch-decks.

ISSUE/OPPORTUNITY

When it comes to planning and negotiating trade promotions, Consumer Packaged Goods (CPG) organizations draw from a multitude of data sources and there is often not enough time to filter through all relevant information. What is needed is a way to more rapidly consult data sources to enhance trade pricing negotiations by predicting outcomes, customizing strategies, and tailoring selling stories. At the same time, there is also a challenge in understanding complex transactional data from retailers, which holds valuable insights for the design of successful promotion plans (i.e., what, where, and how to promote).

Tags | Sales

HOW AI CAN HELP

Supporting employees

AI can be used to prepare negotiation materials by combing through older campaigns or deals, sorting the relevant information, and generating suggestions. This helps equip the human employee with materials like pre-works (e.g., consolidated material from prior years) and pitch decks, supporting their negotiations.

Predicting outcomes

AI can help optimize trade shelf spacing and investment allocation by predicting outcomes and conducting scenario building and storytelling. It can also be used to build scenarios with cultural customizations for negotiation processes with retailers.

Optimization support

With AI, users rapidly analyze EPOS data and transactional information to provide insights that help optimize the design of promotional programs, setting the right price points, promotion mechanics, and anticipating sales uplift to inform production processes of the expected demand.

Planning for promotions

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Because price, margin information, and negotiation strategies are consumed by the model, it must be secured to prevent the leakage of sensitive commercial data.

Fair and impartial

The data used to train and fuel the model may be dated, leaving new target groups and small but growing customer segments potentially underrepresented. As a result of this latent bias, the model may be challenged to provide commensurate accuracy for all groups and segments.

POTENTIAL BENEFITS

Driving efficiency

By using AI to augment preparing and sorting materials, the organization promotes efficiency in trade promotion processes.

Trade promotion effectiveness

Leveraging AI can help improve allocation of resources across price, promotion, and negotiation strategies.

Data-driven decision-making

Using AI to create materials for trade negotiations enables human workers to access much more information and make more informed, data driven decisions.

Data access for all

Data-empowered business users

AI can help guide business users to key insights in consumer behaviors by enabling them to combine data from various sources through natural language queries, and by summarizing issues to action without the help of dedicated analysts.

ISSUE/OPPORTUNITY

Everyone in the business should be consumer-focused, but while the marketing function may have access to customer data, business stakeholders in product design, trading, retail operations, supply chain, and other functions may only encounter slices of customer information. Currently, enterprises need dedicated analysts to pull SQL queries and curate data for decision-making, which creates a barrier to customer information and insight. Data is held across different silos, and existing interfaces are only built to answer pre-populated questions. The result is that most business users cannot fully leverage the enterprise's models and data, and cross-functional insights are challenging to achieve.

Tags: **Cross-functional**

HOW AI CAN HELP

Greater access to insights

An AI system can help stakeholders across all business functions better understand the consumer by simplifying data mining and analysis with user friendly interfaces and natural language queries. This allows users to ask questions relevant to their work and extract actionable insights without compromising functionality.

Bringing down data barriers

The system can aggregate data from various sources and domains (e.g., purchasing patterns, customer service, website and browsing data, marketing campaign response) to provide comprehensive insights into consumer behaviors. Reaching across data silos, the system can automatically identify outliers and summarize issues to guide decision-makers to areas requiring attention.

Data access for all

MANAGING RISK AND PROMOTING TRUST

Safe and secure

The AI model is exposed to sensitive and proprietary enterprise data, which creates a risk of potential data leakage. To mitigate this risk, the enterprise may look at restricting data access to the AI provider, as well as carefully determining what consumer data should be exposed to the model.

Robust and reliable

For business users to make confident decisions informed by AI, they need to be able to trust the outputs. To this end, data inputs must be accurate and up to date, and outputs should be validated and monitored.

Transparent and explainable

Business users require sufficient context to interpret consumer data, and while analysis conducted by a data expert inherently contains a level of "human in the loop," when using an AI model, business users need the capacity to understand context and outputs.

POTENTIAL BENEFITS

Agile decision-making

Business users are empowered to make more informed decisions about product launches, sales, and other customer-related initiatives both quickly and efficiently.

Time and resource efficiency

Simplifying data access and analysis for business users can accelerate time to insight without additional burdens on data analysts and the technical workforce.

Seeing is believing

Virtual try-on

AI can be used for style transferring, which allows consumers to see a digital rendering of clothes and other products on their own bodies, in their homes, and elsewhere.

ISSUE/OPPORTUNITY

In the clothing and make-up industry, consumers typically try on products to determine whether they want to purchase and keep it. Yet, this traditional method of selecting products is challenged by online shopping, where the consumer relies on pictures and product descriptions to inform their decision. This can lead to high return rates and affiliated costs to the company, as well as customer dissatisfaction.

Tags | **Customer Service**

HOW AI CAN HELP

Accurate style transferring
By analyzing images or videos of the customer and the desired style, AI can create realistic representations of how the clothing or product would look in the real world.

Greater personalization
By considering factors such as body shape, skin tone, and personal style, AI can suggest suitable products that align with the customer's preferences.

Virtual mix-and-match
AI allows customers to more easily explore a wider range of style options, clothing combinations, and accessories.

Seeing is believing

MANAGING RISK AND PROMOTING TRUST

Private

By working with and augmenting consumer photos and videos, the model is exposed to sensitive or personally identifiable information, which is subject to privacy regulations and standards. Leveraging AI for style transferring requires the enterprise to ensure user data is safely stored, transferred, and used.

Transparent and explainable

When consumers input an image of themselves or their surroundings, they need to understand how that media is used by the enterprise, how consumer-machine interactions are tracked and recorded, and whether there are any privacy risks to the consumer when using the style transferring application.

Fair and impartial

If the training set is unbalanced and therefore biased, renderings for virtual try-ons may be more accurate or realistic for one demographic group over another, potentially impacting customer satisfaction and regulatory compliance.

POTENTIAL BENEFITS

Customization for the customer

Catering to the customer buying experience with a simpler way to explore product offerings promotes customer satisfaction.

Reduced return rates

When customers can better see and imagine how a product looks before making a purchase, it helps reduce the likelihood of mismatched expectations, product dissatisfaction, and returns.

Simpler sales

Making it easier to choose which product to buy by virtue of a simpler method for exploring options can support sales growth.

Trend analysis and insights

AI can be used to analyze data from virtual try-on experiences to gather insights on customer preferences, popular styles, and emerging trends.

Code assist for developers

Augmented developer

AI can be used to supplement the work of software developers by helping create and maintain multiple applications and platforms.

ISSUE/OPPORTUNITY

To give customers a seamless digital experience, enterprises are challenged to develop and maintain applications across different platforms. Yet, developers and other highly skilled professionals are in high demand and short supply. To overcome the talent gap, AI can be used to supplement a developer's effort by automating aspects of code creation and maintenance so the developer can focus on more complex code writing and validating AI outputs.

Tags | **Information Technology**

HOW AI CAN HELP

Offloading lower-level work
AI can augment the completion of repetitive tasks, such as the deployment and maintenance of code across different platforms (e.g., iOS, Android, webapps).

A developer assistant
AI can be used in the development of the code itself, serving as an assistant supporting software developers in writing and maintaining code. It can also promote consistency across platforms and applications, such as by converting functional code to different environments.

Code assist for developers

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Code created with AI may include vulnerabilities that may be difficult to identify during development and even after deployment. Given the importance of cybersecurity, enterprises need to ensure generated code does not introduce security risks.

Robust and reliable

AI is susceptible to errors, and when using it for development tasks, human validation is necessary to mitigate the risk of bugs or vulnerabilities in code as it is created and maintained for multiple applications.

POTENTIAL BENEFITS

Efficient deployments

Using AI can help developers efficiently deploy and maintain code across platforms.

Digital consistency

Using AI helps developers maintain a consistent experience across multiple platforms by ensuring each environment functions at the same level of quality, thanks to automation (e.g., code conversion) that augments developer capacity and capabilities.

Customer support on demand

Customer assistant

AI-enabled virtual agents can improve the customer experience by providing real-time, personalized support and creating new ways of interacting with customers.

ISSUE/OPPORTUNITY

After purchase, customers may seek information or support around a product or service. While traditional call centers have implemented basic AI capabilities to automate responses to customer inquiries, the automation is often limited in its capacity to interpret customer questions and respond in a conversational and helpful way. The need is to accurately and proactively respond to customer inquiries and online trends in an efficient and effective manner.

Tags **Customer Service**

HOW AI CAN HELP

A conversational agent
AI can enable new ways of engaging with customers, using speech-to-text and natural language inputs to generate empathetic and personalized conversations for aftersales support and handling customer complaints.

Better use of human capital
Because generative AI can provide instant, personalized responses to customer queries, offer relevant solutions, and engage in conversations, customers can gain faster response and resolution, and organizations can free up human agents to focus on more complex customer issues.

Customer support on demand

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

The quality and accuracy of customer interactions impact the customer experience and brand impression. If an AI-enabled customer assistant fails to provide accurate and personalized advice or product instructions, it could degrade (rather than enhance) the quality of the customer interaction.

Transparent and explainable

Customers should have the opportunity to gain a clear understanding of what the model can and cannot do. Also, to promote transparency and positive engagements, enterprises should set customer expectations for the virtual assistant.

POTENTIAL BENEFITS

Enhanced customer experience

Personalized and accurate support and troubleshooting contribute to a positive brand reputation and improve customer relationships and loyalty.

Increased efficiency

By using AI to automate various aspects of customer engagement, a higher volume of customer interactions can be accomplished simultaneously, improving response times and creating the capacity to scale with customer demand.

A virtual shopping assistant

Product recommendations

AI can be used to create personalized product recommendations based on customer preferences and behavior.

ISSUE/OPPORTUNITY

Suggesting the right products to customers can increase sales, and hyper-personalized product recommendations are often the most effective at driving a sale. Data-based product recommendations are already possible today, but they often lack a conversational, natural language tone.

What is more, recommendations may lack a hyper-personalized quality as they are based on broader customer segments and purchase history, as opposed to individual customer search criteria and feedback.

Tags: **Sales**

HOW AI CAN HELP

Hyper-personalized recommendations
Based on customer input and preferences, AI can generate tailored recommendations, making the buying process more personalized and convenient. In addition, the interactive and iterative approach to product recommendations that AI enables can yield more targeted suggestions than current search engine capabilities.

Image as input/output
Consumers can enter an image of preferred styles (e.g., a celebrity in a designer outfit), and the AI model can identify products and make recommendations based on the image.

A virtual shopping assistant

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

Latent bias in training and testing data may lead the model to express a preference toward some products or product combinations when making recommendations. Ongoing monitoring, data updates, and human validation can contribute to continuous improvement and bias mitigation.

Private

The model may be exposed to customer data throughout the course of an interaction, and that personal information may be subject to regulatory protections. Important considerations include how the customer data is stored, transferred, and used, as well as how the data is consumed and used by the model itself.

POTENTIAL BENEFITS

Enhanced customer experience

Delivering personalized and accurate support, guidance, and troubleshooting helps create a positive brand reputation and improves customer relationships and loyalty.

Increased efficiency

Using AI to automate selected customer engagement activities can improve efficiency and scalability while improving customer satisfaction.

Next-level market intelligence

Market research

By harnessing AI's capacity to read and summarize vast amounts of relevant material, companies can expedite market research and gain concise insights for effective decision-making in new markets.

ISSUE/OPPORTUNITY

When researching entry possibilities in new markets or customer groups and identifying new target segments, enterprises face a variety of challenges. Things like a lack of market data, unfamiliar customer preferences, cultural and economic differences, competitive analysis difficulties, regulatory complexities, high market entry costs, potential brand perception challenges, and uncertainties about demand and market acceptance all impact the speed and quality of market research.

Tags: **Marketing**

HOW AI CAN HELP

Market intelligence
AI can help simulate market scenarios, generate synthetic data to fill data gaps, predict customer preferences based on existing patterns, offer cross cultural insights, aid in competitor analysis, suggest compliance strategies, optimize market entry costs, simulate brand perception scenarios, and provide demand forecasting to reduce uncertainties.

Information synthesis
AI enables rapid market research by efficiently reading and summarizing extensive volumes of pertinent material, presenting the information in a readily understandable format for market research teams.

Novel market segmentation
AI generated data may reveal new and previously unidentified market segments within the target market. This can open up additional opportunities for niche marketing and product customization.

Richer personas
Rather than relying on basic surveys and focus groups for understanding consumer likes and dislikes, AI can identify specific customer preferences and create detailed profiles. Using AI, market research teams can even create fictional yet plausible customer personas based on the market's unique characteristics, helping the company better understand their potential customers' behavior and preferences.

Next-level market intelligence

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	AI models may learn from biased datasets, leading to biased outputs that do not accurately represent the actual market.
Robust and reliable	Given AI's potential to hallucinate and produce inaccurate outputs, AI-generated insights should be verified with real-world data and traditional research methods to ensure accuracy and reliability.
Responsible and accountable	While AI can complement market research, it should not replace traditional research entirely, as it may miss qualitative nuances and human expertise.
Transparent and explainable	To trust AI outputs, users require the ability to understand which samples and research methods were used to generate recommendations and insights.

POTENTIAL BENEFITS

Cost-effective research

AI can reduce the costs associated with traditional market research methods by generating large datasets and simulating scenarios.

Risk mitigation

By simulating market responses, CPG companies can identify potential risks and challenges in the new market before making substantial investments. This helps reduce the chances of product failure and financial losses.

Integrated business planning

AI consolidation of forecasting and planning across the enterprise

AI can help an organization consolidate real-time sales, demand, and supply data across all functions, creating a single source of truth to drive faster, more strategic decisions in finance, supply chain, marketing, and sales.

ISSUE/OPPORTUNITY

Today's companies have a wide variety of systems for planning and forecasting. However, the individual outputs from those disparate systems often conflict with each other and don't provide a unified view of what's really going on. Different teams—finance, supply chain, marketing, and sales—create their own forecasts using siloed data and inconsistent approaches. The potential results? Mismatched projections, inefficiencies, delayed decision-making, and significant operational waste.

Tags: **Operations**

HOW AI CAN HELP

Real-time consolidation
AI can consolidate real time inputs from sales, inventory, marketing trends, and supply chain metrics to produce dynamic forecasts.

Sophisticated analysis
AI enables trend recognition, historical pattern analysis, and early alerting on supply demand gaps, while also facilitating scenario planning and pricing strategy refinement all through a unified dashboard.

Actionable insights
The system can continuously update itself as new data flows in, signaling demand shifts or regional product affinities and providing decision-makers with actionable insights.

Integrated business planning

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Given the system's critical business impact, resilience is key. AI models should be capable of updating in real time and integrating multiple data streams reliably and accurately. Extensive pilot testing can help fine-tune model accuracy before scaling.

Transparent and explainable

A user-facing dashboard that clearly shows inputs, trends, and recommendations can help business leaders understand how forecasts are generated, what assumptions are at play, and what real-world data is influencing outputs— reducing blind reliance on the system and promoting human-AI collaboration.

Safe and secure

To help mitigate security breaches and operational disruptions, robust security protocols should be embedded in both the technology infrastructure and data flows, with IT overseeing access controls, data encryption, and integration with existing ERP systems.

POTENTIAL BENEFITS

Unified forecasting with less redundancy

AI can help minimize conflicting forecasts across departments, creating a single source of truth for the entire enterprise.

Improved collaboration and decision-making

Cross-functional teams are able to operate from the same real-time data set, improving alignment. Also, leaders spend less time on data consolidation and cross-checking, enabling them to make better-informed decisions more quickly.

Greater supply chain efficiency

Integrated business planning powered by AI enables better inventory and warehouse management, which can reduce supply disruptions, shortages, and waste.

Social media content generation

Automated, multimodal content creation that is trend-aware and always on

AI is now being used to autonomously produce social media content—text, images, hashtags, and videos—that aligns with brand identity and capitalizes on viral trends in real time.

ISSUE/OPPORTUNITY

Social media is a key channel for communicating with customers and shaping brand perceptions, and an important driver for awareness, engagement, and sales conversion. But creating personalized, high-quality content at speed and scale—while maintaining brand consistency and legal compliance—is a difficult balance.

Large enterprises often rely on global agencies to support social media content across dozens of brands and channels. This approach can be very costly, time-consuming, and limited by human working hours. Also, in a media environment where trends can shift in an instant, traditional methods likely cannot scale or respond quickly enough to keep pace with opportunities in real time.

Tags: **Marketing**

HOW AI CAN HELP

Detecting and analyzing trends and events
AI can help detect and analyze influencer trends and brand affinity across a wide range of social media platforms 24/7, identifying opportunities to shape consumer expectations in real time. Retrieval augmented generation (RAG) capabilities provide real time access to social data, such as trending hashtags, viral video clips, and current events.

Model-agnostic orchestration
Content creation tasks can be dynamically routed to the most cost effective or best performing AI models, optimizing output while reducing compute costs.

Generating multimodal creative content
AI offers the ability to autonomously generate creative content across modalities, while remaining contextually and culturally aware. Key capabilities include: (1) LLMs for generating social media copy, product descriptions, captions, and hashtags; (2) multimodal image models for visual asset generation, including pack shots, brand imagery, and marketing visuals; and (3) short form video generation.

Social media content generation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Retrieval-augmented generation can reduce hallucinations and improve model performance over time. Fallback models and safety nets can mitigate failures or inappropriate content generation under unpredictable conditions.

Fair and impartial

The content generation pipeline should be evaluated regularly for potential cultural, social, or representational biases. Human oversight can ensure that outputs reflect brand values.

Private

No personal user data should be used in the generation process; models should be trained and tuned on anonymized or public datasets. Data residency and usage should comply with regional regulations, including the EU AI Act.

POTENTIAL BENEFITS

Always-on, real-time responsiveness

Traditional content workflows often require long lead times for ideation, approval, and execution. With AI, brands can respond almost instantly to real-time events, seasonal trends, or cultural moments by generating relevant content in minutes, enabling more agile and timely brand engagement.

Scalable content production at low marginal cost

AI enables brands to produce high volumes of personalized, platform-specific content—text, images, and video—without requiring a linear increase in resources. The system can support hundreds of product lines and campaigns with minimal incremental effort, greatly improving operational leverage.

Cost reduction through automation and budget reallocation

Augmenting external creative agencies and internal content teams with AI-generated outputs can reduce operational costs. It can also free up budget that can be reallocated toward more strategic initiatives, such as paid media, analytics, or customer experience improvement.

Data-driven personalization and targeting

AI systems can tailor content for audience segments based on behavior, geography, platform norms, or product affinity. This micro-personalization allows brands to deliver relevant content to niche audiences, increasing engagement and conversion potential.

Improved productivity

By automating repetitive or time-intensive content generation tasks, marketing and creative professionals can focus more on high-value work such as strategy, brand storytelling, or campaign optimization. This reallocation of effort can lead to improved job satisfaction and better use of talent.

Consistent brand voice and visual identity

With proper tuning and governance, AI-generated content can more reliably align with predefined brand guidelines, helping to ensure a unified voice across markets, languages, and touchpoints. The system can learn and reinforce tone, terminology, and aesthetic standards consistently.

Semi-autonomous warehouse loading and unloading robotics

Human-supervised physical AI for dynamic material handling

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

ISSUE/OPPORTUNITY

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

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

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

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

Tags: **Distribution & Logistics** | **Physical AI**

HOW AI CAN HELP

Perception of unstructured physical environments
Vision and sensor fusion allow systems to identify objects, assess orientation, detect obstacles, and understand spatial constraints in real time.

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

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

Fallback continuity mechanisms
Manual override capability ensures uninterrupted operations during technical interruptions, allowing workers to complete tasks manually if robots experience downtime.

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

Semi-autonomous warehouse loading and unloading robotics

MANAGING RISK AND PROMOTING TRUST

Safe and secure

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

Responsible and accountable

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

Transparent and explainable

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

POTENTIAL BENEFITS

Reduced unloading time

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

Lower ergonomic strain

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

Operational resilience through fallback mechanisms

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

Improved space utilization

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

Cross industry applicability

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

Autonomous transport for urban mobility services

AI-driven mobility in unpredictable urban environments

Autonomous vehicles provide ride-hailing, mobility services, and delivery of goods without human drivers, bringing physical AI directly into daily consumer transportation.

ISSUE/OPPORTUNITY

Urban mobility systems may face driver shortages, rising costs, and inconsistent service availability. Traditional transport models struggle to scale efficiently. Cities can experience mobility gaps in underserved neighborhoods where ride share driver availability is limited and in off-peak hours when driver supply drops sharply despite continued passenger need.

The opportunity is to deploy autonomous passenger services that improve accessibility while reducing reliance on human drivers, enabling consistent service across hours and locations within regulatory frameworks that help enable safety and public acceptance.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

Perception and navigation
AI interprets complex urban environments including traffic patterns, pedestrian behavior, construction zones, and road conditions to navigate safely through city streets.

Regulatory-compliant design
Systems align with approval requirements including data reporting, safety certifications, and operational restrictions mandated by local transportation authorities.

Safety-constrained autonomy
Operations remain supervised and comply with strict safety requirements through speed limits, restricted operating zones, and conservative decision-making that prioritizes passenger and public safety.

Human fallback mechanisms
Escalation paths exist for edge cases where remote operators can provide guidance or take control when the autonomous system encounters situations outside its operational design domain.

Passenger interaction systems
Vehicles communicate with users directly through voice interfaces, in-vehicle displays, and mobile apps to coordinate pickups, provide route information, and address passenger requests.

Fleet-level optimization
Vehicles with fleet telemetry are deployed based on demand patterns, positioning cars near areas with expected pickup requests to minimize wait times and improve service coverage.

Autonomous transport for urban mobility services

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Autonomous vehicles are high-risk AI systems, making cybersecurity a fundamental design requirement. Systems must be hardened against sensor spoofing, adversarial attacks, and unauthorized access, with clear incident-response protocols and regular third-party security assessments occurring before and throughout commercial operation.

Responsible and accountable

When an autonomous vehicle is involved in a public incident, accountability cannot be ambiguous. Responsibility frameworks must be established *before* deployment begins. Detailed operational logs and safety event records should be maintained to support incident investigation, regulatory reporting, and iterative safety improvement.

Fair and impartial

The promise of autonomous mobility—expanding access to underserved neighborhoods and off-peak hours—can only be realized if fleet deployment algorithms are actively designed for equity, not just efficiency. Routing, availability, and pricing models should be regularly audited to ensure they do not systematically disadvantage riders based on location, income, or inability to access digital payment methods.

POTENTIAL BENEFITS

Broader access

Mobility access improves for a broader population as service can be provided in areas and at times when driver availability dips, expanding to new locations and times.

Improved scalability

Service scalability increases as fleets can be sized to match demand without recruiting and retaining drivers.

Better service

Passenger experience becomes more consistent and predictable as vehicle behavior, routing, and service quality follow standardized protocols, subject to regulatory approval and public acceptance.

Lower costs

Operating costs decrease as driver dependence is reduced, potentially enabling lower fares and expanded geographic coverage.

Multipurpose household service robots

Reasoning-enabled service robots for home environments

Physical AI-enabled service robots that use reasoning models to perform household support tasks (e.g., cleaning assistance, item retrieval, setup, basic monitoring) in dynamic home environments, within defined safety and autonomy boundaries.

ISSUE/OPPORTUNITY

Current household robots require detailed task programming for every specific action, limiting their usefulness to narrow, pre-defined activities. Users must explicitly instruct robots on each step of a task—where to go, what to pick up, how to handle objects, and when to stop—making deployment time-consuming and limiting robots to repetitive, identical tasks. Household environments change constantly with objects moved, furniture rearranged, and new items introduced, causing pre-programmed instructions to quickly become outdated and require manual updates.

The barrier to adoption is the programming burden rather than hardware capability—households need robots that can reason about their environment and infer appropriate actions based on context, goals, and safety constraints rather than following rigid scripts.

The opportunity is to shift from scripted automation to reasoning based physical AI that can interpret context, infer appropriate actions, and operate reliably in unstructured home environments—dramatically reducing setup effort while expanding practical value.

Tags: **Customer Experience** | **Physical AI**

HOW AI CAN HELP

Contextual reasoning
AI infers appropriate actions by understanding the current situation, user goals, and environmental context rather than requiring explicit step-by-step instructions for every task variation.

Human interaction
Systems respond naturally to conversational requests and environmental cues, allowing users to communicate intent at a high level rather than specifying detailed procedures.

Reduced task programming
Explicit instructions are minimized as systems learn to generalize across similar tasks and adapt to environmental changes without requiring manual reprogramming.

Safety-constrained autonomy
Actions remain bounded within defined safety envelopes that prevent damage to property, ensure human safety, and avoid behaviors outside approved operational limits.

Multipurpose household service robots

MANAGING RISK AND PROMOTING TRUST

Private

Household service robots observe some of the most intimate spaces in people's lives—continuously capturing audio, video, spatial maps, and the behavioral patterns of every occupant, including children. Data collection should be limited strictly to what task execution requires, processed on-device wherever possible, and never shared with third parties without explicit, informed consent from all household members.

Safe and secure

A robot that physically manipulates objects inside a home—around children, elderly occupants, and pets—presents harm potential that is immediate and concrete. Safety boundaries must be rigorously defined and tested well beyond lab conditions, and network-connected systems must be secured against unauthorized access that could allow external parties to remotely observe or control devices operating inside private residences.

Responsible and accountable

When a robot causes harm or property damage through an autonomously reasoned action, it can be difficult to determine who is responsible—the hardware manufacturer, the AI developer, or the platform operator. Allocation of responsibility must be defined contractually before deployment, not resolved after an incident. Operational logs should be sufficient to reconstruct exactly what the system perceived, inferred, and did at the time.

POTENTIAL BENEFITS

Ease of use

Less instruction required as users can communicate goals at a high level and allow systems to determine implementation details based on environmental reasoning.

Improved interaction

Systems feel more intuitive as they respond to natural language and contextual cues rather than requiring users to learn specialized programming interfaces or command structures.

Long-term scalability

Automation expands gradually as reasoning capabilities improve and systems learn to handle increasingly complex household tasks through experience and model updates.

Broader task coverage

More activities automated as systems can handle variations and novel situations without explicit programming for each specific scenario encountered.

Vision-enabled store operations

Real-time retail execution through vision

Vision enabled store operations leverage in store computer vision and edge analytics to track shelf execution and planogram adherence, enabling timely adjustments to product placement based on real-time conditions.

ISSUE/OPPORTUNITY

Retail execution and shelf compliance are traditionally validated through manual audits, which are time-consuming, inconsistent, and reactive. Field representatives visually inspect product placement, stock levels, and promotional displays across distributed retail environments, traveling from store to store to compare physical shelf arrangements against planogram specifications. Each inspection requires the representative to mentally compare what they see against ideal layouts, estimate spacing and facings, and document deviations for later follow-up.

Manual validation limits coverage and slows corrective action, as representatives can visit a fraction of locations each week, and by the time audit reports reach stores, shelf conditions may have changed. Inconsistent shelf execution can reduce sales performance and brand visibility, as products placed in wrong locations receive less customer attention, out-of-stock situations go undetected, and promotional displays fail to meet brand standards.

The opportunity lies in automating visual validation through computer vision, enabling faster identification of misplacement, out-of-stock risk, or suboptimal layout. However, accuracy should remain high across varying lighting conditions, store formats, and device types to help enable trust and usability at scale.

Tags: **Sales** | **Physical AI**

HOW AI CAN HELP

Edge vision in the aisle
Computer vision models analyze shelf images to detect product placement and spacing, identifying individual SKUs, counting facings, and recognizing when items are incorrectly positioned.

Real-time feedback loops
Field users receive immediate guidance on corrective actions, with visual overlays showing which products need adjustment.

Context aware planogram reasoning
AI compares observed layout, with ERP-integrated reconciliation, against expected configuration templates, highlighting deviations and prioritizing corrections.

Scalable validation coverage
Automated analysis increases inspection frequency without increasing labor, enabling daily checks rather than weekly or monthly manual audits.

Human in the loop execution
Field staff remain in control; AI provides recommendations and visual overlays, to help enable fast, informed corrections without autonomous physical action.

Vision-enabled store operations

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Accuracy across the full range of real-world retail conditions is a prerequisite for this use case at scale—and those conditions are demanding: variable store lighting, shelf clutter, inconsistent device camera quality, and thousands of SKU variations across product ranges and regional markets. Models validated on a narrow set of stores can struggle in new environments, so performance monitoring must be conducted continuously across the live deployment footprint, not as a one-time pre-launch exercise.

Responsible and accountable

When AI-generated compliance assessments inform commercial consequences (e.g., supplier penalties, disputed promotional fees, or contract disputes), the basis for those assessments must be defensible, not just directionally accurate. Organizations should maintain clear audit trails of what the system detected, how it was interpreted, and who acted on the output, with human oversight required before automated findings trigger any consequential commercial action.

Transparent and explainable

The human-in-the-loop design this system depends on only works if field representatives can understand why a deviation has been flagged. Recommendations should be accompanied by visual evidence, a plain-language explanation of expected versus observed shelf layout, and a confidence indicator—giving staff what they need to exercise true judgment rather than following outputs they cannot question or challenge.

POTENTIAL BENEFITS

Faster compliance validation

Immediate detection reduces correction lag, enabling same-visit fixes rather than waiting for audit reports to be processed and communicated.

Improved merchandising effectiveness

Optimized placement increases sales performance, as products consistently appear in planned positions that maximize visibility and align with promotional campaigns.

Reduced manual auditing effort

Automation lowers time spent on inspections, freeing representatives to focus on relationship building and strategic merchandising improvements.

Higher execution consistency

Standardized validation improves brand reliability across locations, ensuring consistent shelf presentation across store formats and markets.

Fleet telemetry and route optimization

Adaptive logistics driven by edge intelligence

Physical AI systems embed intelligence directly into delivery vehicles, using onboard sensors, edge computing, and real-time connectivity to continuously perceive operating conditions and adapt routing, driving behavior, and delivery execution while vehicles are in motion.

ISSUE/OPPORTUNITY

Delivery networks operate across diverse traffic conditions, distribution constraints, and customer delivery windows. Static route planning fails to adapt dynamically to real-time conditions, leading to delays and inefficiencies. Multiple tracking vendors and inconsistent telemetry standards create integration challenges that slow scaling. Logistics inefficiencies increase fuel consumption, delay deliveries, and reduce retailer service levels.

The opportunity is to deploy physical AI-enabled fleets where vehicles themselves become intelligent actors—continuously sensing conditions, adjusting execution in real time, and coordinating with fleet systems to maintain service reliability at scale. However, interoperability, data standardization, and secure integration with manufacturing and warehouse systems are prerequisites for system-wide orchestration.

Tags: **Distribution & Logistics** | **Physical AI**

HOW AI CAN HELP

Real-time route optimization
AI models analyze location, traffic, and delivery progress to dynamically adjust routing in response to changing conditions, accidents, or unexpected delays.

Load sequencing optimization
Algorithms optimize delivery sequencing and product mixing to reduce turnaround times, ensuring products are loaded in the order they'll be delivered and minimizing time spent searching for items at each stop.

Driver behavior analytics
Telemetry supports identification of inefficient or unsafe driving patterns including harsh braking, excessive idling, or suboptimal speed management that increases fuel consumption and risk.

Cross-system data integration
Fleet data aligns with production schedules and warehouse dispatch systems for coordinated outbound logistics, ensuring vehicles depart when orders are ready and arrive when receiving docks are available.

Predictive delay modeling
Edge-deployed AI anticipates disruptions and recommends proactive rerouting before delays impact delivery schedules, accounting for historical traffic patterns, weather forecasts, and known construction zones.

Fleet telemetry and route optimization

MANAGING RISK AND PROMOTING TRUST

Private

Continuous fleet monitoring generates substantial personal data about drivers: precise location history, behavioral patterns, working hours, and biometric data (when driver-facing cameras are used). To address emerging legal requirements on biometric data collection, AI-enabled fleet solutions will need explicit driver consent, clear retention limits, and regular vendor security audits before and after deployment.

Responsible and accountable

AI-generated driver behavior scores can have direct employment consequences, including disciplinary action, retraining requirements, or dismissal. Organizations must establish clear governance over how these scores inform HR decisions, with humans retaining authority for any consequential employment action. Also, drivers must be allowed to access their own data and to formally contest assessments they believe to be inaccurate.

Fair and impartial

Driver scoring models trained on historical fleet data risk penalizing drivers systematically for factors outside their control, such as operating in high-congestion urban areas, covering more demanding routes, or driving older vehicles with lower-quality sensors. AI models should be regularly audited to confirm they reflect genuine driving behavior rather than route difficulty or equipment variability, so that performance assessments are genuinely comparable across the fleet.

POTENTIAL BENEFITS

Reduced time-to-retailer

Dynamic routing improves service-level performance by avoiding delays, minimizing wait times at delivery locations, and ensuring on-time arrivals within promised delivery windows.

Improved safety monitoring

Behavior analytics reduce operational risk by identifying drivers who need additional training, detecting patterns that predict accidents, and enabling proactive interventions before incidents occur.

Higher delivery reliability

Real-time adjustments mitigate disruption impact, allowing logistics managers to communicate accurate arrival times to retailers and maintain service commitments despite unexpected obstacles.

Lower transportation cost

Fuel efficiency and idle-time reduction decrease expenses through optimized routes, reduced unnecessary mileage, and improved driver behavior that eliminates wasteful practices.

Edge–cloud architecture for consumer mobility

Distributed intelligence enabling real-time physical action in vehicles

An edge–cloud physical AI architecture distributes intelligence between vehicles and centralized platforms to enable real time perception and control at the edge, while supporting fleet wide learning, data management, and continuous improvement in the cloud. This approach balances ultra low latency safety requirements with scalable model training and deployment across geographically distributed mobility fleets.

ISSUE/OPPORTUNITY

Mobility vehicles generate massive volumes of sensor data from cameras, lidar, radar, and other onboard systems while operating in dynamic, safety-critical environments. Many driving decisions must be made within milliseconds, making it impractical and unsafe to rely solely on cloud-based processing due to network latency, bandwidth constraints, and connectivity variability. At the same time, fully localized intelligence limits the ability to learn from fleet-wide experiences, slowing improvement of perception and control models and preventing vehicles from benefiting from rare or geographically distributed edge cases.

Managing, transferring, labeling, and reusing raw sensor data at scale is costly and creates development bottlenecks. Infrastructure limits on onboard compute, storage, and network capacity further constrain how much data can be processed or transmitted. The opportunity is a unified edge–cloud architecture that enables real-time local execution while coordinating centralized learning, data management, and deployment to accelerate autonomous driving development without violating real-world system constraints.

Tags: **Information Technology** | **Physical AI**

HOW AI CAN HELP

Real time physical intelligence at the edge

Vehicles locally process camera, LiDAR, radar, and telemetry data to perceive surroundings and execute physical actions immediately. Edge AI ensures millisecond level response for safety critical maneuvers even during connectivity loss.

Fleet level learning in the cloud

Edge–cloud data platforms curate, prioritize, and replay real world driving scenarios (including rare edge cases) to continuously improve perception and control models. Simulation and synthetic data augment real world data to accelerate learning without increasing on road risk.

Continuous closed loop improvement

Edge systems infer component health from real world behavior, reducing sensor dependence, while fleet data and simulation refine models in the cloud. Validated updates are pushed back over the air, steadily improving safety, performance, and reliability across vehicles without hardware changes.

Edge–cloud architecture for consumer mobility

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

This architecture requires that edge AI be consistently capable of making safety-critical decisions—emergency braking, hazard avoidance, collision detection—within milliseconds. A model that performs well in testing but degrades unpredictably under real-world conditions of sensor noise, adverse weather, or hardware variation across vehicle generations creates direct safety risk. Continuous validation across the live fleet is not optional; it is the foundation on which everything else depends.

Safe and secure

The over-the-air update pipeline, which pushes new AI models simultaneously to an entire deployed fleet, is both this architecture's greatest strength and its most acute vulnerability. A compromised or insufficiently validated update could affect thousands of vehicles at once. Securing the full update lifecycle, from model training through cryptographic signing and staged deployment, should be treated as a critical fleet-wide safety requirement, not an IT governance checkbox.

Responsible and accountable

Continuous cloud-based model updates create an accountability challenge that is structurally unique to this architecture. When an AI-driven safety event occurs, determining which model version was running on which vehicle at that moment is not straightforward when models are being updated fleet-wide on an ongoing basis. Operators must maintain version-controlled deployment records that are precise enough to reconstruct the exact system state at the time of any safety-related event, and governance processes must ensure independent validation before any update reaches production vehicles.

POTENTIAL BENEFITS

Operational resilience

Vehicles remain functional during network disruptions due to edge autonomy.

Broader scenario coverage and model robustness

Fleet-wide data aggregation combined with synthetic data generation improves performance across diverse weather, traffic, and road conditions.

Faster innovation cycles

Fleet wide learning accelerates improvement without manual recalibration.

Lower system cost

Selective data transmission and local inference reduce bandwidth and cloud compute costs.

Consistent experience at scale

Ride quality, braking behavior, and navigation improve uniformly across fleets.

Safer consumer mobility

Real time, on vehicle decision making reduces accident risk in dynamic environments.

Robotic stowing and picking system

Shelf-based picking and stowing in warehouses

Robotic systems automate stowing and picking at warehouse shelf interfaces and delivery stations, using computer vision to identify items in cluttered slots, spatial modeling to track shelf occupancy, and force-sensitive manipulation to handle products in tight clearances.

ISSUE/OPPORTUNITY

Shelf-based warehouse operations require workers to repeatedly reach into densely packed slots, bend to low shelves, lift items overhead, and manipulate products with varying fragility and weight in minimal clearance spaces. These repetitive motions create ergonomic risks—back injuries, shoulder strain, repetitive stress injuries—that drive workers' compensation costs and turnover.

Traditional rigid automation cannot handle obstructed items in cluttered slots, navigate tight clearances without damaging adjacent inventory, or grasp items of varying shapes without crushing or dropping them.

AI-enabled robotic systems can perform shelf-based picking and stowing for a substantial portion of SKUs, reducing ergonomic risk while expanding automation scope.

Tags: **Distribution & Logistics** | **Physical AI**

HOW AI CAN HELP

Accurate perception in dense shelf slots

Vision models can identify items despite partial occlusion, varied shelf lighting, and tight spacing that creates ambiguity about item boundaries and grasp points.

Item-level automation decisions

AI evaluates each SKU's physical characteristics to determine which items can be reliably handled robotically versus which should route to human workers, optimizing labor division based on actual system capabilities.

Fine manipulation with force feedback

Force sensors provide real-time feedback during grasping, adjusting grip pressure based on item rigidity and detecting contact with shelf edges to abort unsafe motions before damage occurs.

Footprint-aware system design

AI supports layout optimization for limited space.

World modeling of shelf geometry

AI tracks which slots are occupied, how items are positioned, and available clearances, enabling motion planning that avoids collisions with shelves and neighboring inventory.

Robotic stowing and picking system

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

The commercial case for this system rests on AI vision and force-sensitive manipulation being dependable enough in challenging real-world situations—irregular packaging, partial occlusion, varying weight and fragility—to be trusted with a substantial share of shelf operations. Models that perform well on common products but degrade on unfamiliar or awkwardly packaged items don't just reduce efficiency; they actively disrupt the human-robot division of labor the whole system depends on.

Safe and secure

Robotic arms operating near human workers pose immediate physical safety risks if force control or object detection fails. Safety boundaries must be validated across the full range of real operating conditions—not just controlled testing scenarios—with reliable fallback behaviors when the system encounters situations, objects, or worker proximity outside its defined operational envelope.

Responsible and accountable

A framework for allocating responsibility when problems occur must be clearly defined before deployment—not negotiated after an incident. Operational logs capturing what the system perceived, inferred, and physically executed should be maintained at sufficient fidelity to support incident investigation, insurance claims, and liability determination.

POTENTIAL BENEFITS

Ergonomic risk reduction

Reduced repetitive reaching, bending, and lifting for human workers by offloading physically demanding shelf interactions, particularly for heavy items and awkward positions.

Operational consistency

Reduced performance variability across shifts, sites, and seasonal workforce fluctuations, with robotic systems maintaining consistent throughput.

Labor cost savings

Lower manual handling effort allows facilities to maintain throughput with fewer workers (or redeploy workers to tasks requiring human judgment).

Expanded automation coverage

Current systems can handle approximately 75% of SKUs based on physical characteristics, compared to much lower coverage with traditional fixed automation.

Improved order accuracy

Standardized and consistent picking logic minimizes human error, ensuring the right items are picked every time and significantly reducing mis-picks, rework, and customer returns.

Vision-enabled robotic induction for high-variability consumer logistics

Handling SKU variability at industrial throughput

Vision-enabled robotic systems use advanced computer vision, perception, and machine-learning models to identify, orient, grasp, and transfer a wide variety of items across inbound logistics flows. These systems operate across conveyor-based induction as well as floor-loaded and palletized trailer unloading, handling high SKU diversity, reflective or damaged packaging, inconsistent presentation, and unstructured environments at industrial throughput.

ISSUE/OPPORTUNITY

Conveyor induction in distribution centers involves extreme product variability—including potentially thousands of SKUs with different geometries, weights, packaging materials, and labeling—making manual induction a physically demanding, error-prone bottleneck. Traditional rule-based automation often fails when handling reflective surfaces, damaged packaging, inconsistent item presentation, or unlabeled products because these systems depend on rigid templates and known item geometries.

The opportunity is to deploy physical AI systems that can reason about physical objects in motion and adapt manipulation behavior in real time—without reconfiguration—while operating at industrial throughput.

Tags: **Distribution & Logistics** | **Physical AI**

HOW AI CAN HELP

Tolerance of variability

AI models can identify and classify items despite significant differences in geometry, surface reflectivity, label placement, and packaging condition, eliminating the need for pre-configured templates for each SKU.

Real-time classification and routing

Vision models process items continuously as they arrive, supporting immediate routing decisions in high-speed conveyor environments where delays create bottlenecks.

Closed loop learning from physical outcomes

Execution results (drops, misfeeds, successful placements) feed back into model behavior, improving robustness across new SKUs and packaging variations.

Edge-based execution

AI inference runs on local computing hardware positioned near the robot to meet the low-latency requirements needed for continuous industrial throughput.

Non-safety-critical deployment context

Systems operate in zones isolated from human workers, reducing the safety certification and liability burden compared to collaborative robot applications.

Adaptive manipulation in the physical loop

AI adjusts robotic grasp points, placement force, and release timing dynamically based on observed item characteristics, reducing jams, misfeeds, and dropped items.

Vision-enabled robotic induction for high-variability consumer logistics

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

High throughput industrial environments leave little room for error. A vision model that misclassifies reflective packaging, damaged labels, or unfamiliar SKUs doesn't just slow the line; it creates misrouting, exceptions, and bottlenecks that ripple across the entire distribution center. Performance must be validated continuously across the full range of real-world items that the system will encounter in production, not just the SKUs present at the time of design and testing.

Responsible and accountable

The system's closed-loop learning design means model behavior changes continuously based on real-world outcomes. This creates a unique accountability challenge: the system that caused a misrouting event today may not be the same system that was used last month. Operational logs must capture classification decisions, manipulation outcomes, and model versions with enough detail and accuracy to support root cause analysis, contractual dispute resolution, and continuous improvement.

Transparent and explainable

Continuous model updates based on operational feedback require a higher level of governance discipline than static automation systems. Operators need to understand when model behavior has changed, what drove the change, and whether performance improvements on newly encountered SKUs have introduced degradation elsewhere. Model updates should be documented, tested against all SKUs, and communicated to operations teams before being deployed to production lines.

POTENTIAL BENEFITS

Higher throughput

Increased items inducted per hour by removing manual handling bottlenecks at inbound stations.

Operational consistency

Stable performance across SKU changes, seasonal product variations, and peak demand periods without system recalibration.

Lower error rates

Reduced misrouting, mislabeling, and downstream exception handling through more consistent item identification and placement.

Labor productivity

Reduced reliance on repetitive manual induction tasks, allowing workers to focus on exception handling and quality verification.

Reduced ergonomic strain and injury risk

Automation removes repetitive lifting and manual handling from high-risk inbound tasks.

Autonomous material movement in consumer fulfillment environments

Physical AI-driven logistics in dynamic, human-shared facilities

Autonomous mobile robots (AMRs) safely transport materials across warehouses and factory floors shared with human workers. Using AI based perception and edge autonomy, robots detect people, equipment, and obstacles in real time, dynamically adjusting routes and speed. Fleet level orchestration coordinates multiple robots to reduce congestion, improve throughput, and maintain safe, flexible operations without fixed infrastructure.

ISSUE/OPPORTUNITY

Internal material transport in warehouses and manufacturing facilities relies heavily on manual labor using forklifts, pallet jacks, and hand carts. Traditional automated guided vehicles (AGVs) require fixed guide tracks, magnetic tape paths, or segregated operational zones that isolate them from human workers. As facility layouts and workflow patterns change frequently to accommodate seasonal demand, new product lines, or process improvements, fixed-path automation becomes a constraint that limits operational flexibility.

Mobile automation that can safely operate in dynamic environments alongside human workers without requiring permanent infrastructure modifications enables facilities to efficiently reconfigure layouts and processes while maintaining automated material flow.

As robotic fleets grow, local autonomy alone creates systemic bottlenecks: traffic jams at high-use corridors, task queuing at popular workstations, and cascading delays when disruptions occur. The opportunity is to deploy physical AI systems that can safely reason and act in motion, enabling flexible material transport that adapts continuously to real world conditions while operating alongside people.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

Human-object discrimination

Perception models using computer vision and machine learning differentiate humans from static objects like pallets, storage racks, carts, and structural obstacles, enabling the robot to apply different behavioral rules depending on what it detects in its path.

Human aware safety envelopes

Robots enforce dynamic speed limits, stopping distances, and approach behaviors tuned to local safety standards, facility zones, and regulatory requirements.

Adaptive speed control

Robots automatically reduce speed or stop when humans are detected within defined proximity zones, with behavior adjusted based on approach angle, human movement patterns, and local safety requirements.

Edge based decision execution

Edge-based autonomy enables immediate responses to sudden obstacles or human movements without requiring communication with centralized traffic control systems, reducing latency and maintaining safe operation even during network disruptions.

Fleet-level Orchestration

AI based fleet orchestration optimizes task allocation, path planning, and workload balancing in real time, enabling coordinated multi robot operations while reducing bottlenecks and idle time.

Dynamic navigation

AI continuously recomputes optimal paths based on real-time observations of congestion patterns, temporary obstacles, floor conditions, and human activity, avoiding the need for pre-programmed routes that become obsolete when layouts change.

Autonomous material movement in consumer fulfillment environments

MANAGING RISK AND PROMOTING TRUST

Safe and secure

The success of this system hinges on robots safely sharing space with workers. Human-object discrimination must perform reliably across all live operating conditions: poor lighting, crowded peak-period aisles, workers in non-standard positions, and edge cases not well-represented in training data. Failures here are not just performance shortfalls; they could lead to human injury or death.

Responsible and accountable

When an AMR is involved in a collision or near-miss with a worker, fault attribution between the AI developer, robot manufacturer, systems integrator, and facility operator can be a serious and complex challenge. Accountability frameworks must be established before deployment. Also, operational logs capturing robot perception, decision-making, and motion at the time of any safety event must be maintained with sufficient accuracy and detail to support regulatory reporting, insurance claims, and liability determination.

Transparent and explainable

Workers sharing a facility with autonomous robots have a vested interest in understanding how those robots will behave around them: what triggers a slowdown or stop, how they should act when a robot approaches, and what to do when behavior seems unexpected. Clear communication about robot behavior rules is not just an ethical obligation; it's an operational requirement. Workers who don't understand or trust robot behavior create unsafe interactions and workarounds that can undermine the human-robot collaboration the system depends on.

POTENTIAL BENEFITS

Reduced transport labor

Lower dependence on manual material movement for repetitive routes, allowing workers to focus on tasks requiring judgment and dexterity.

Scalable deployment

Reduced infrastructure requirements enable faster rollout across multiple sites without extensive facility modifications or downtime.

Improved safety

Reduced collision risk in shared human-robot spaces through consistent detection and predictable, conservative robot behavior around people.

Operational flexibility

Facility layouts, storage locations, and workflow patterns can be modified without reengineering robot paths or installing new guidance infrastructure.

Higher asset utilization

Reduced robot idle time through better task sequencing and routing, allowing facilities to handle higher workloads with existing fleets rather than purchasing additional robots.

Programmable and general-purpose robots for consumer operations

Adaptive physical AI systems operating across dynamic consumer environments

General purpose physical AI robots, including humanoid and mobile platforms, are designed to perform multiple tasks across dynamic consumer environments. Powered by unified vision language action models, these systems can adapt to inspection, material handling, basic maintenance, and support tasks through software updates, enabling flexible deployment, human supervised autonomy, and reuse of the same hardware as operational needs evolve.

ISSUE/OPPORTUNITY

Most industrial robots are designed for narrowly defined physical tasks, limiting flexibility when products, layouts, or processes change. A welding robot cannot easily be repurposed for material handling, and a picking robot cannot perform quality inspection without significant hardware modification or replacement.

This specialization creates substantial retooling costs and long deployment timelines whenever operational needs evolve, forcing organizations to maintain large fleets of single-purpose machines that sit idle when their specific task is not needed. Traditional automation fails in these environments because it depends on fixed layouts, rigid programming, and narrow task definitions. Reconfiguring automation when workflows change is slow and capital intensive. The opportunity is physical AI systems that can perceive, reason, and act in real time, allowing the same robotic platform to adapt to new tasks, environments, and workflows without physical retooling—while operating safely alongside human workers.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

Real time environmental perception

Robots continuously perceive shelves, products, people, tools, and obstacles using vision and sensor fusion, maintaining an up to date world model rather than relying on static maps.

Safety-constrained behavior

AI limits actions to defined safety envelopes, ensuring robots operate within speed, force, and proximity constraints appropriate for shared human-robot workspaces.

Shared learning across tasks

Experience and training from one application transfers to others, as skills learned for inspection (e.g., object recognition) support maintenance tasks (e.g., part identification).

Human-supervised autonomy

Robots operate under controlled conditions with human oversight, performing routine tasks autonomously while escalating complex or ambiguous situations to human operators.

Software-driven capability expansion

New tasks are added through software updates and model training without hardware changes, allowing the same robot platform to take on additional tasks over time.

Multi-task adaptability

Robots dynamically transition between picking, material transport, inspection, and support tasks, allowing a single platform to serve multiple operational roles.

Programmable and general-purpose robots for consumer operations

MANAGING RISK AND PROMOTING TRUST

Safe and secure

General-purpose robots that acquire new capabilities through software updates without hardware changes create a safety certification challenge that single-purpose automation doesn't face. Specifically, safety validation completed at commissioning may be invalidated by a subsequent update. As such, each meaningful software-driven capability expansion should trigger a fresh risk assessment—not be treated as a routine update covered by existing certification.

Responsible and accountable

When a software update adds a new physical capability and the robot subsequently causes harm or damage during that task, liability allocation can be difficult to determine. Contractual accountability frameworks must be established before each new capability is deployed, not negotiated after an incident occurs.

Transparent and explainable

Although human-supervised autonomy is a core design requirement, it is only possible if workers and supervisors know what the robot is currently configured to do, what constraints govern it, and what has changed since the last update. Effective human oversight requires clear, accessible documentation of current capabilities and limitations that is updated every time a new task or model is deployed.

POTENTIAL BENEFITS

Improved flexibility

Single platforms support multiple tasks, enabling organizations to redeploy robots as operational priorities shift without purchasing specialized equipment for each new application.

Extended hardware lifespan

Software updates extend platform usefulness by adding capabilities and adapting to new tasks, reducing the frequency of hardware replacement and improving return on capital.

Reduced retooling effort

Less task-specific automation investment required when processes change, since generalized robots can be reprogrammed rather than replaced or extensively reengineered.

Future-proofed automation path

Generalized, software-defined platforms scale with evolving processes and increasing SKU variability, reducing the need for repeated reengineering as operations change.

The Energy, Resources & Industrials AI Dossier



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.

Tags: **Sales** | **Agentic AI**

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.

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.

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.

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.

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.

Tags: **Manufacturing & Quality**

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.

Tags: **Field Services**

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.

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.

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.

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.

Higher overall productivity

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

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.

Tags: **Compliance & Risk** | **Physical AI**

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

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.

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.

Remote human supervision

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

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 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.

Tags: **Operations** | **Physical AI**

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.

Tags: Field Services | Physical AI

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.

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.

Tags: **Manufacturing & Quality** | **Physical AI**

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.

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.

The Financial Services AI Dossier



The Financial Services AI Dossier

As AI capabilities mature and regulatory clarity improves, AI is moving from experimental to essential—powering financial services businesses that are smarter, faster, and more responsible.

In an environment where trust, precision, and speed are paramount, AI is enabling financial services firms to detect risk earlier, serve customers more effectively, and compete with greater agility. Whether it's fraud prevention in real time, hyper-personalized marketing, predictive trading, or automated claim reporting, AI is shifting how institutions create value and manage complexity across banking, capital markets, and insurance.

AI is already having a transformative impact on knowledge work, helping people do their jobs more efficiently and effectively and supporting human decision-making with faster, deeper analysis of complex data patterns. It is also helping behind the scenes to triage and manage cyber threats, develop and debug new computer programs, and even generate synthetic data to train smarter AI models.

Agentic systems go even further, enabling AI to handle tasks such as algorithmic trading while keeping workers in the loop for oversight. Also, while financial institutions might not seem like obvious candidates for physical AI, beneath the surface they rely heavily on a critical infrastructure of physical assets—ATM networks, data centers, trading terminals, branch systems—that can fail suddenly in highly visible ways. This makes financial services a particularly compelling environment for physical AI's predictive capabilities.

AI innovations can help financial institutions address rising pressure to modernize legacy systems, reduce costs, and respond to evolving customer expectations. However, they also bring heightened scrutiny around fairness, transparency, and accountability. Firms that want to harness the full power of AI will therefore need to invest not only in improved model performance but also in foundational elements such as governance, data quality, and organizational readiness.

AI's powerful capabilities come at an opportune time to help financial institutions address rising pressure to modernize legacy systems, reduce costs, and respond to evolving customer expectations.

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-powered risk management and regulatory compliance

Using AI agents to provide 24/7 risk and compliance monitoring

Agentic AI systems can act as an always-on compliance and risk management team, with specialized agents continuously monitoring risks across a financial institution's operations.

ISSUE/OPPORTUNITY

Financial services firms must continuously monitor for risks—ranging from fraudulent transactions to compliance violations—across a vast landscape of operations. At scale, with millions of customers and complex regulatory demands, risk management can be costly, labor intensive, and subject to error. Missed risks can lead to massive penalties, reputational damage, and systemic vulnerabilities. Yet, human teams are limited in their ability to keep pace with the overwhelming volume of transactions, communications, and rule changes that require their oversight.

Traditional monitoring approaches can be siloed and reactive, identifying issues only after damage has been done. Multi-agent AI offers a proactive, comprehensive solution that can scan for risks, interpret new regulations, and test institutional resilience under different scenarios.

Tags: **Compliance & Risk** | **Agentic AI**

HOW AI CAN HELP

Risk detection and monitoring
Agents that specialize in different risk areas (such as credit card fraud, anti-money laundering (AML), or cybersecurity) can automatically detect and monitor risks. They can also work together, sharing alerts across domains. For example, if a cybersecurity agent detects a breached account, it can warn a transaction monitoring agent to increase scrutiny on that account.

Regulatory change monitoring
Other specialized agents can parse new regulations using natural language processing, interpret rules and guidelines, map new requirements against company controls to identify compliance gaps, and then draft reports or suggest control changes for humans to implement.

Stress testing and scenario analysis
A stress testing agent can simulate scenarios such as interest rate jumps or credit crunches, coordinating with agents representing different departments (e.g., loans, trading, deposits) to assess institutional safety.

AI-powered risk management and regulatory compliance

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Financial firms face massive penalties and reputational damage for compliance failures. As such, multi-agent systems should be validated against historical fraud cases, past audits, and synthetic stress scenarios to ensure their outputs are dependable and comprehensive.

Safe and secure

Pulling data from various databases to feed AI agents requires significant IT work and raises potential security concerns. Given the sensitive nature of risk management data, agentic systems must be secured against unauthorized access or accidental disclosure of private customer information.

Transparent and explainable

Risk management is overseen by humans who are personally liable and may be wary of trusting "black box" agents. Explainability is key. Agents need to provide clear audit trails (for example, explaining why a transaction was flagged, and including references to rules or past cases). This clarity makes AI outputs easier for human officers and regulators to trust and validate.

Fair and impartial

Biases in data sources or training models can result in unfair or inaccurate risk assessments, raising concerns for customer trust, brand reputation, and regulatory compliance.

POTENTIAL BENEFITS

Higher assurance and lower compliance costs

Proactive and comprehensive monitoring reduces the likelihood of missed issues, helping firms avoid massive penalties while reducing the high labor costs associated with traditional compliance processes.

Reduced risk exposure and faster detection

Always-on agentic systems can sift through transactions, emails, and news at speed and scale, reducing a firm's exposure by identifying risks much sooner.

Improved audit and regulatory relations

AI agents specializing in compliance can speed up audits, doing in hours what might otherwise take human teams weeks to complete.

Ultra-personalized financial advice and wealth management

Using AI agents to deliver highly personalized financial planning

Agentic AI systems can serve as an automated financial advisor, providing tailored wealth management by analyzing an individual's full financial picture and adapting plans to changing needs and market conditions.

ISSUE/OPPORTUNITY

Traditionally, financial advisors segment clients into broad categories—by age, income, or risk tolerance—and provide standardized advice. This limits personalization and overlooks the nuances of an individual's financial situation. At the same time, wealth management firms often reserve high-quality, customized advice for affluent clients because it is expensive and resource-intensive to deliver at scale.

As client expectations shift toward greater personalization and cost-efficiency, firms are seeking ways to deliver more customized service at lower operational cost—without compromising regulatory compliance or fiduciary responsibility.

Tags: **Customer Experience** | **Agentic AI**

HOW AI CAN HELP

Market monitoring
A market scout agent can track real-time developments such as interest rate changes, stock prices, and macroeconomic indicators to identify relevant opportunities and risks.

Personal financial analysis
A client profile analyst agent can aggregate financial data (e.g., bank accounts, credit card spending, loans) and behavioral patterns (e.g., risk-taking versus saving) to build a detailed and dynamic personal profile for each client.

Planning and strategy generation
A planning agent can synthesize market and client inputs to recommend bespoke financial strategies—for example, rebalancing investment portfolios in response to market changes or personal life events such as a home purchase or retirement.

Compliance integration
A compliance agent can help ensure recommendations adhere to fiduciary standards and regulatory guidelines, making the AI outputs safe and trustworthy for both clients and advisors.

Continuous fine-tuning
Unlike static robo-advisors, multi-agent AI systems can provide adaptive guidance, refining recommendations as circumstances evolve, rather than offering static one-time plans. They can identify portfolio risks in real time, flag anomalies for review, and automate activities such as threshold-based rebalancing, factor-based investing, and tax loss harvesting.

Ultra-personalized financial advice and wealth management

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Because errors could cause significant client losses, agents should be validated against historical financial data and tested under varied market and life-event scenarios to ensure they generate sound, resilient advice.
Transparent and explainable	AI-driven decisions—such as rebalancing actions or investment recommendations—should be accompanied by clear rationales and audit trails, enabling accountability and regulatory review.
Fair and impartial	Advisory models should be monitored to ensure they do not steer clients toward biased product selections, specific asset classes, or vendor-preferred solutions. Advice must be applied equitably across client segments.
Private	Because AI agents handle highly sensitive financial data, strong safeguards must be in place to protect information, including anonymization and strict access controls.

POTENTIAL BENEFITS

Personalization at scale

Multi-agent AI makes high-quality, personalized wealth management affordable for mid-market clients, expanding access beyond the traditionally affluent segment.

Better outcomes

Custom investment strategies ultra-personalized by AI to fit an individual's life stages, goals, and preferences can improve customer satisfaction and financial outcomes in a cost-effective way, boosting retention and conversion.

Strategic differentiation

AI-driven wealth management can help a firm establish a premium position in a commoditized market by offering a high-touch yet affordable client experience that is difficult to achieve in other ways.

Greater scalability and efficiency

Automation of portfolio management tasks enables firms to serve more clients at lower cost, freeing human advisors for complex decisions.

AI agents for algorithmic trading and market simulation

Enhancing trading strategies and insights with multi-agent collaboration

Agentic AI systems enable trading firms to use specialized agents to simulate artificial markets and to execute diverse trading strategies, enabling smarter, faster trades and richer insights into market dynamics.

ISSUE/OPPORTUNITY

Financial firms have long been at the forefront of harnessing intelligent technologies such as algorithmic trading and AI to achieve a winning edge. Now, multi-agent AI is taking the game to a whole new level.

Traditional algorithmic trading transformed financial markets. However, behind closed doors, multi-agent systems are already in production and demonstrating better adaptability to market changes—outperforming single-strategy approaches in various timeframes. In fact, it is now common for different algorithmic strategies (i.e., agents) to be running concurrently and even "competing" for capital allocation based on performance.

The latest innovation is to leverage more explicit agent frameworks and inter-agent communication, with multiple and varied trading or simulation agents acting independently yet sharing information to optimize outcomes. Looking ahead, we expect to see greater adoption of multi-agent reinforcement learning for strategy development. We may also see exchanges using agent-based AI to monitor market stability or simulate the impact of rule changes. Without such tools, firms risk being outpaced in highly competitive markets (and regulators may risk missing early warning signs of instability). However, in many cases, the biggest innovations could remain hidden from view for competitive reasons.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Market simulation

In artificial market simulations, agents representing different trader archetypes (e.g., retail investors, institutional traders, market makers) can enable researchers, investment firms, and regulators to observe emergent phenomena and test different strategies and scenarios before they go live.

Specialized trading agents

A firm can simultaneously deploy live agents that apply different strategies in the market—such as short-term arbitrage, medium-term trend following, or options hedging—sharing signals to avoid conflicts and improving portfolio resilience.

Coordinated action

Agents can operate semi-autonomously under the supervision of a coordinator agent or human risk manager, ensuring portfolio alignment and preventing overexposure. In some cases, different AI agents might even collude or negotiate with each other to create a competitive advantage.

Reinforcement learning and communication

Advanced systems can use multi-agent reinforcement learning, where agents learn from each other through trial and error in simulated environments. Some of these systems might leverage LLMs to enable agents to communicate and explain their reasoning, making the simulation results easier to interpret.

AI agents for algorithmic trading and market simulation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Because unpredictable agent interactions can create market instability, systems should be stress-tested against diverse scenarios, including flash-crash conditions, to ensure agents behave reliably under pressure.
Fair and impartial	Agents should be monitored to prevent behaviors that could unfairly favor certain counterparties or drift into market manipulation. Simulations should be designed to reflect diverse trader types, avoiding skewed outcomes.
Private	Trading data and strategies are highly sensitive and commercially valuable. Systems should ensure strict data security and access controls to protect intellectual property and prevent leakage of proprietary trading logic.
Safe and secure	Given the adversarial nature of markets, safeguards should be established to prevent agents from being deceived or manipulated by other firms' AI. Strong cybersecurity and resilience measures are essential to maintain safe operations.

POTENTIAL BENEFITS

Diversified trading performance

Multi-agent systems can function like an "AI investment committee," combining different strategies to improve returns across market regimes.

Faster reaction times

Agents can operate in milliseconds, dynamically hedging exposures and adjusting positions to manage risk in volatile conditions. This increased speed can help a firm boost its trading volumes, validate analyses in real time, and potentially drive greater profitability while mitigating risks.

Systemic insight

Agent-based simulations could provide regulators and institutions with cutting-edge tools to identify systemic risks or opportunities, enabling better preparation for disruptive events.

AI agents for credit underwriting

Adaptive, data-driven underwriting through multi-agent collaboration

Agentic AI systems can transform credit underwriting through specialized agents that analyze applicant data, monitor market context, assess risk, and maintain compliance—creating highly personalized lending decisions.

ISSUE/OPPORTUNITY

Traditional credit underwriting often relies on generic segmentation (e.g., age, income, or credit history) and static risk models. This approach can be slow, inflexible, and exclusionary, frequently locking out those with thin credit files or non-traditional income sources. At the same time, lenders struggle with outdated workflows that fail to adapt dynamically to changes in markets or individual circumstances.

As lenders work to expand their businesses responsibly, they need more agile underwriting solutions capable of assessing varied data and real-time conditions and delivering accurate decisions across a broad customer base.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Multisource data aggregation

A data gathering agent can compile structured and unstructured information (e.g., bank statements, tax filings, e-commerce history, location signals) to build an accurate and comprehensive borrower profile.

Dynamic scoring and simulations

Another agent can apply adaptive scoring models and run repayment scenario simulations, guiding approval, denial, or escalation recommendations based on risk capacity and context.

Regulatory alignment

A compliance-focused agent can confirm all assessments meet suitability requirements, fiduciary standards, and traceability benchmarks. This solution layer supports audit readiness and decision traceability.

Workflow orchestration

Agents working in concert—each handling discrete steps in the credit workflow, from data gathering to scoring and reporting—can provide continuous fine-tuning of underwriting decisions as new data or market conditions emerge.

AI agents for credit underwriting

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Credit decisions can have life-changing implications. Agents must be validated continuously against historical performance and varied economic conditions. Simulated stress tests should regularly challenge outcomes to ensure resilience across credit cycles

Transparent and explainable

Regulators and borrowers require clarity on underwriting decisions. Agents should log their reasoning (e.g., which income streams or risk patterns influenced a score), presenting audit-ready explanations that demystify the model's logic.

POTENTIAL BENEFITS

Faster, smarter approvals at lower cost

Multiple data sources, specialized agents, and adaptive models enable faster, more precise underwriting without increasing manual workloads.

Financial inclusion

By incorporating non-traditional signals and actively monitoring for bias, AI agents can extend credit responsibly to underserved populations.

Operational scalability and quality

Multi-agent systems can boost lender efficiency and client satisfaction by enabling scalable, high-volume credit processing with consistent quality and regulatory compliance.

Intraday liquidity optimization

Reconciling breaks in real time to free trapped cash and lower buffers

Agentic AI systems can automatically reconcile breaks and timing mismatches as they occur, so banks can unlock trapped cash faster and operate with a smaller intraday cash buffer.

ISSUE/OPPORTUNITY

Across payment systems, securities depositories, and internal subledgers, breaks and timing mismatches trap cash until exceptions are resolved. The shift to faster settlement cycles (e.g., T+1 in U.S. markets) compresses post-trade windows and raises the cost of delays or late funding. At the same time, supervisory agencies expect better monitoring of intraday flows and positions, even as instant payment operating hours expand the volume and timing of movements that must be continuously managed. Taken together, these trends make same-day reconciliation and proactive liquidity steering a necessity, not a nice-to-have.

Tags | Operations | Agentic AI

HOW AI CAN HELP

Autonomous, event-level matching
Multiple AI agents continuously match payment confirmations, Nostro statements, ledger entries, and settlement messages (ISO 20022 rich data) to resolve breaks as they arise—requesting clarifications, attaching evidence, and escalating only true exceptions. This releases cash sooner.

Intraday forecasting and liquidity routing
A forecasting agent aggregates real-time statuses (e.g., cross-border tracking, queued RTGS payments) and predicts cash peaks and troughs. A treasury agent then recommends actions—sweeps, intraday credit, collateral moves, or CLS pay-ins—to minimize buffers while meeting payment obligations.

T+1 and instant-payments readiness
Exception agents accelerate break resolution and funding decisions under compressed timelines, using enriched ISO 20022 data and end-to-end payment visibility to cut manual follow-ups and reduce same-day funding surprises.

Intraday liquidity optimization

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Because reconciliation errors can trigger fails and liquidity penalties, agents should be tested on historical breaks, simulated queues, and stress days (month-end, index rebalancing). Agents should be designed to fall back to deterministic matching rules and human review for ambiguous cases.

Transparent and explainable

Since treasury, operations, and auditors must understand every break closure, agents should be designed to produce line-item audit trails (source message IDs, fields compared, confidence scores) and plain-language rationales for funding recommendations and payment releases.

Fair and impartial

Because intraday funding decisions affect business lines and clients differently, agents should be designed to apply consistent prioritization policies (e.g., client SLAs, regulatory cutoffs) and be monitored to avoid systematically disadvantaging lower-volume corridors or smaller counterparties.

Safe and secure

Given the high value of payment instructions, agents should be designed with strong cybersecurity controls (segregated runtimes, signed message handling, anomaly detection) and resilience testing to prevent spoofed confirmations or malicious instruction replay.

POTENTIAL BENEFITS

Less trapped cash and lower buffers

Faster break closure and proactive sequencing reduce precautionary intraday buffers and daylight overdraft usage, improving liquidity efficiency and lowering carry costs.

Fewer fails and fees

Real-time matching and RTGS-aware scheduling reduce late payments and fail charges while improving counterparty confidence and SLA performance.

Audit-ready compliance

Automated evidence trails support supervisory expectations for intraday monitoring and reporting, strengthening an organization's control posture with less manual effort.

Transformation with speed and confidence

Code assistant for digital transformation

AI can enable banks to increase digitization at a faster pace through code assistants.

ISSUE/OPPORTUNITY

Many financial institutions are pursuing cloud and data transformations, which are essential steps in preparing the organization for using AI tools of many kinds. In some cases, legacy hardware is retired as data is curated and shifted to the cloud, freeing up humans for more valuable work while bringing down the costs associated with on-premise infrastructure. Yet, these kinds of transformations are significant undertakings that can bring long lead times and high costs. There is also a risk of failure and error.

Tags | **Information Technology**

HOW AI CAN HELP

Supercharge your human capital
AI can be used as a component of cloud and data transformations to empower developers working across the enterprise on applications, data engineering, machine learning, and front-end development.

A helping hand in code development
As organizations explore new digital and cloud capabilities, development teams can accelerate and simplify their work by using AI as a force multiplier when writing, debugging, and documenting code, as well as translating ideas to code.

A shorter path to software
Part of success in transformation hinges on how quickly new enabling software can be deployed. There are opportunities to use AI in software development to shorten the lifecycle and more quickly reach a stable and deployable version, such as by helping rapidly write APIs, ETL, data pipelines, or even front-end code.

Transformation with speed and confidence

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Partial automation of programming-related tasks requires the system to be reliably available and accurate. If availability cannot be guaranteed to an acceptable extent, the benefits of automation must be weighed against the risk of erroneous or buggy code.

Responsible and accountable

The training data for foundation models may create legal risks related to intellectual property or copyright infringement. If the training data contains copyrighted material, the organization deploying the model needs to evaluate whether the presence of intellectual property in the training set could lead to legal challenges against the enterprise. Also, while the use of AI can accelerate the work of developers, without a human in the loop (e.g., to validate and debug code), critical failures may occur. Shoring up accountability may involve documenting and communicating standards and expectations for employees using AI.

Safe and secure

By using an AI system, proprietary code bases may be exposed to third parties, raising questions around the security of the data and controlled access to it. An inadvertent breach of confidential intellectual property could have significant enterprise impacts.

POTENTIAL BENEFITS

Lower transformation costs

By shortening the software development lifecycle, the organization can reduce overall costs for digital and cloud transformation.

Lower the bar to digital entry

Using AI opens the door for financial services organizations of all sizes, capabilities, and technology maturities to digitize and move to the cloud in a way that was previously out of reach for many.

Business intelligence at your fingertips

Enterprise-wide data search and access

Make business intelligence via enterprise data search accessible to all through natural language interfaces.

ISSUE/OPPORTUNITY

For many banking and insurance companies, data is stored in a multitude of locations, from local hardware to cloud storage solutions. This makes it difficult to effectively query different databases and retrieve relevant information quickly and efficiently. Complicating the matter, multiple mergers and acquisitions over time may have compounded the diversity of data locations and databases, which hinders data mining for insights. Financial services organizations are information-intensive enterprises, and without the capacity to easily query all data, the result is poorer or incomplete insights that can increase enterprise risk while threatening customer satisfaction.

Tags: **Cross-functional**

HOW AI CAN HELP

An enabling interface
AI capabilities can be built on top of an existing solution to facilitate the communication of queries from the user to the search layer. This serves as the interface between search layers and databases, allowing users to easily mine all enterprise data, as well as generate structured analytics reports.

Speed to insight
With the ability to query and analyze disparate data sources using AI as an interface, the enterprise can move past traditional business intelligence techniques and dramatically reduce the time required to generate insights while increasing the workforce accessibility to business intelligence.

Business intelligence at your fingertips

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

When it comes to governance and control, granting more data access to a wider segment of the workforce can create a more complex challenge of restricting who in the organization is permitted to access sensitive business data.

Robust and reliable

Given the known challenges with model reliability and the potential for hallucination, banking and financial services organizations face the risk of inaccurate or false AI-derived insights influencing decision-making and leading to negative ramifications even at the market level.

Private

When dealing with sensitive and proprietary information, the organization must contend with securing the data, removing or obscuring it in training and testing sets, and evaluating the model to determine whether it could "leak" protected information, either due to faulty function or a targeted attack.

POTENTIAL BENEFITS

Lower technical hurdles

AI as an interface between search and data enables business users to query databases and obtain tailored results without in-depth programming experience. This gives more of the workforce access to business intelligence without additional burdens on IT and data science teams.

A new level of data-driven decisions

Real-time access to all of an enterprise's data can help organizations become even more insight driven, which supports improved growth prospects through access to the right insight at the right time.

Fixing the missing data issue

Synthetic data generation

Generate synthetic data for model training, anomaly detection, and identifying cyber and deception attacks.

ISSUE/OPPORTUNITY

Missing data is a significant challenge for financial services organizations. Datasets may be incomplete, data transfers restricted, and potential anomalies underrepresented in the data. Using synthetic data can help overcome these challenges. In cloud transformation, data transfers may be delayed due to risks and regulations around data governance, and use of synthetic data can enable a smoother and more efficient transformation. Meanwhile, machine learning anomaly detection systems (such as those for identifying fraud, waste, and abuse) must be trained on data from previous events. Yet, the dearth of data from these rare events can make anomalies harder to assess.

Tags | **Information Technology**

HOW AI CAN HELP

Improve model training
Generative AI can be used to quickly create synthetic data to supplement machine learning model training data, which is then used to aid and accelerate digital and cloud transformations. In this way, generative AI complements the enterprise's wider AI initiatives, fueling (rather than replacing) other AI deployments.

Amplify anomaly event detection
The rarity of anomaly events can make it difficult to train machine learning systems to detect instances of fraud, waste, and abuse. Creating synthetic data with generative AI gives ML systems a larger suite of examples that lead to a greater capacity to find patterns and anomalies in the data.

Harden the organization's cyber posture
Just as synthetic fraud data can be used to train models to identify fraudulent activity, synthetic adversarial data can be used to train models to detect and mitigate cybersecurity risks and deceptive behaviors targeting virtual assistants.

Fixing the missing data issue

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

A significant risk when generating synthetic data is that historic biases can creep into the generated data, perpetuating those biases (e.g., in the case of certain communities or socio-economic groups being underrepresented in the data because those groups have had lower levels of banking activity in the past). This bias is not necessarily intentional but is still harmful.

Robust and reliable

Synthetic data created with generative AI can be limited in its scope and scale, and it should not be presumed to be accurate or perfectly reflective of real-world data. An over-reliance on synthetic data may inject problems with data reliability, which can hamper the validity and usefulness of outputs and model training.

POTENTIAL BENEFITS

Faster path to the cloud

Generative AI-created synthetic data can accelerate digital and cloud transformations by making the transition smoother and more efficient.

Tackling fraud

Use synthetic data to train machine learning systems on rare or unknown events, such as novel types of fraud.

Security confidence

Adversarial synthetic data contributes to the enterprise-wide imperative to fortify digital assets against cyber threats.

Getting to know your customer

Research-based report generation

Generate and summarize reports on new customers to inform employee decisions for customer onboarding.

ISSUE/OPPORTUNITY

New customers are the lifeblood of a growing financial services enterprise, but onboarding customers can be a highly manual and time-consuming process. Know Your Customer (KYC) standards and rules require institutions to develop meta reports on customers based on economic outlook, equity research, adverse media, and new prospect due diligence. These are high human involvement tasks, with desk-based meta-research consuming valuable time and resources.

Tags | **Customer Service**

HOW AI CAN HELP

Condensing results for easier consumption

AI can be used to summarize and filter results from existing search engines to inform meta reports, as well as to summarize information for the customer relationship manager.

Research and analysis to inform reports

Generative AI, alongside other machine learning models, can be used to conduct preliminary data searches and meta-analysis, potentially accelerating the KYC process.

Getting to know your customer

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

When using AI to perform research and analysis, there is a risk that it could skip or misconstrue highly relevant information, which could skew the conclusions in the meta-analysis and hamper sound decision-making. If a new customer is erroneously assessed to be higher risk and a relationship manager declines the opportunity to engage the customer, the consequences are missed revenue and diminished customer engagement.

Private

When dealing with a customer's financial or personally identifiable information, the enterprise faces legal and regulatory standards for data privacy. When using AI, the organization should take steps to ensure sensitive information does not inadvertently leak through model outputs, as well as govern who has access to the model, the underlying data, and the customer data it references.

POTENTIAL BENEFITS

Timely insights

Faster and more efficient search and analysis can give decision-makers more up-to-date information and insights that enable better, more timely decision-making around customer onboarding.

Cost reduction

By streamlining and simplifying the report generation process, costly labor hours can be redirected to more valuable work.

Efficiency

The end user saves time and effort by more easily accessing and consuming relevant information.

Enhanced AI support for customers

Financial guardian

A personal virtual assistant powered by AI can cater to daily needs of customers.

ISSUE/OPPORTUNITY

Customer service and engagement is vital for financial services organizations, particularly as they transform from a product-focused to a customer-focused business. As a part of that, there is a pressing need for the enterprise to rapidly and accurately answer both common and complex customer queries and do so with the timeliness customers expect. Yet, increased digitization in financial services has reduced access to representatives who can answer customer questions; this is at a time when financial services customers seek a hyper-personalized experience.

Tags: **Customer Service**

HOW AI CAN HELP

An empathetic, personalized interface

Generative AI can be paired with other models to create a customer interface that delivers a hyper-personalized experience, such as by training the generative AI model to provide answers or insights with empathy. It also overcomes some of the challenges around more traditional chatbots that can lack an empathetic tone.

A more capable digital agent

Using generative AI can improve the usefulness and accessibility of a chatbot interface. With the capabilities of large language models (LLMs), a personal digital assistant can summarize contracts and answer nuanced questions, and the customer may enjoy a range of interface options, including text, audio, and imagery.

Enhanced AI support for customers

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

While a generative AI-enabled solution may provide valuable answers and recommendations most of the time, there is a risk that too much confidence may be placed in the validity of the outputs, both by the organization and the customer. Generative AI is not an infallible oracle, and an overreliance on the AI solution may have a detrimental impact on customer actions, which can in turn increase financial risk.

Private

When confidential or personally identifiable information is inputted via the digital interface by the customer, the financial institution is obligated to follow the laws and rules that dictate how that sensitive information can be transmitted, stored, and accessed. Failing to do so could raise legal peril and potentially subject the enterprise to greater cyber risks.

Transparent and explainable

End users require a clear understanding of how their information will be processed and that they are interacting with a machine. At the same time, the enterprise needs to be able to interpret outputs and understand how and why the generative AI model created a given output.

POTENTIAL BENEFITS

Customer satisfaction

Greater accessibility and more timely answers can lead to a more personalized and satisfying customer experience. This can drive an increase in net promoter score, reflecting increased customer retention and loyalty.

A cycle of efficiency and growth

While customer loyalty and brand reputation fuel business growth, integrating generative AI into virtual assistants can further enhance operational efficiency. AI-enabled systems can serve more customers at scale while the human workforce focuses on resolving the most complex issues or attracting new customers.

Financial inclusivity through hyper-personalization

When customers feel that banking applications are delivering a valuable, personalized experience, it encourages greater customer engagement and interest in service offerings, supporting both the customer's financial wellbeing, as well as that of the enterprise.

Customized marketing for the individual

Hyper-personalized sales and marketing assistant

Regulatory-compliant marketing material generation across different geographies.

ISSUE/OPPORTUNITY

Financial services marketing operations are increasingly coming under regulatory scrutiny for issues such as mis-selling and misinformation. Part of the challenge for multinational organizations is that cultural differences as well as varying levels of customer understanding about products may create regulatory risk for enterprises in a given geography. To overcome this, organizations are investing significant manual labor to maintain a compliant marketing function, which is both time consuming and costly.

Tags **Marketing**

HOW AI CAN HELP

Customized materials for different audiences
AI can be used to create marketing materials that contain the appropriate tone, language, and cultural references, while also supporting consumer understanding of the product to maintain regulatory compliance.

Personalized sales at scale
With AI, financial services organizations have the ability to create marketing materials that are customized to individual customers—and to do so at scale.

Customized marketing for the individual

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

For AI-generated marketing to be valuable, organizations must be able to rely on the validity of the output. AI can be prone to hallucinations, and when tasked with creating marketing that touts one product over another, there is a risk the model will return false statements. This injects potential regulatory violations that could result in fines and other penalties. To shore up reliability with AI, outputs should be validated by a human with the subject matter understanding to do so.

Fair and impartial

Datasets may contain latent bias of which the organization is unaware. This could be due to how the data was acquired, recorded, and curated, and the challenge is compounded when operating in multiple geographies. Organizations should consider the datasets used to train and fuel generative AI systems and whether unknown bias could lead to marketing materials that fall short because they fail to reflect important geographical and cultural differences.

POTENTIAL BENEFITS

Individual-level marketing

Hyper-personalized marketing takes the enterprise to a new stage of customer engagement and enticement that is infeasible through manual effort alone.

Confident compliance

Using AI for marketing development can help ensure the content remains in line with regulatory expectations across many geographies, thereby reducing regulatory risk.

Driving marketing ROI

Personalized marketing can support new sales, strengthen the customer relationship, and reduce the cost of marketing operations while also improving the timeliness of outreach and engagement.

Ensuring the integrity of claims

Automated claims reporting

AI can be used to automatically generate reports based on descriptions or pictures of the relevant subject.

ISSUE/OPPORTUNITY

During claims processing for property and casualty insurance, claims agents must decide whether a destructive event was insured, and if so, the amount and cost of the damage. These processes are complex and can be time consuming, and claims agents have few tools to support their decision-making.

Tags: **Operations**

HOW AI CAN HELP

Virtual damage rendering
AI can be used to help visualize the damage by replicating it virtually. The basis for the replication and visualization can be customer conversations, damage documents, photos, official reports, and other relevant media. In visualizing the data in this way, the claims agent is empowered to make better decisions when assessing the degree and cost of damage.

Automated claims reporting
With AI, claims reports can be generated automatically based on photographic evidence.

Ensuring the integrity of claims

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Damage visualization requires a high degree of accuracy, and erroneous AI outputs could lead to claims being paid incorrectly, potentially leading to overpayment (a detriment to the organization) or underpayment (a detriment to the customer).

Transparent and explainable

If claims agents use AI to automate aspects of claims processing but are unable to articulate to customers how the AI model derived its outputs or contributed to the cost and damage assessment, customers may not accept the outcome of the claims process.

POTENTIAL BENEFITS

Cost reduction

Faster claims processing and more accurate damage assessments can reduce labor costs and claims payments.

Customer satisfaction

By processing claims faster and with less administrative burden, customer satisfaction improves due to quicker, more streamlined adjudication and payment.

Identifying fraud

More expedient claims processing has the advantage of identifying potential fraud more quickly, helping to ensure the integrity of claims and payments.

A virtual bank experience

VR-enabled retail banking centers

Virtual reality customer agents powered by generative AI can change the retail banking experience and interactions.

ISSUE/OPPORTUNITY

Banks seek to provide customers with multiple methods of interacting with their accounts, services, and offerings. As digital services grow, customers may prefer to conduct banking transactions remotely through virtual transactions. However, this approach still requires a human agent, and employees are the most expensive cost component in customer service—particularly when three levels of customer support are necessary. Chatbots can help automate virtual transactions, but existing chat tools are limited to specific, pre-programmed dialogue and options.

HOW AI CAN HELP

Bring the bank to the customer
With a generative AI-enabled virtual space, customers can use a VR headset to conduct business with the financial institution and interact with a service representative from the comfort of their own home in a way that is convenient for the customer.

Hyper-personalized service
A generative AI agent can provide conversational, tailored responses to questions about customer accounts and financial needs. This helps the enterprise cater to its customers' desires for a personalized experience while also avoiding the costs that come with adding more human customer service workers.

Speed and quality of service
In a virtual space, customer data can be accessed in real time by the conversational agent to provide faster, higher quality service and offerings.

Tags: **Customer Experience**

A virtual bank experience

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable

Given a generative AI-enabled chatbot's capacity for conversational outputs, there is a risk that customers may not realize they are interacting with a machine. To promote trust, customers should not only be informed that the chatbot is not human, but they should also understand how their inputs and information are stored, accessed, and used.

Fair and impartial

The datasets used to train and inform the chatbot may contain latent biases, such as underrepresented customer groups or semantic deficiencies in some languages but not others. As a result, the model may simply not work as well for some customers, leading to a variety of negative customer impressions and complaints.

Responsible and accountable

With issues around reliability and data quality, if the chatbot outputs erroneous data or recommendations, a human stakeholder needs to be accountable for the outcome. Organizations can promote accountability by keeping a human in the loop and documenting roles and responsibilities.

POTENTIAL BENEFITS

Customer-focused experience

A generative AI-enabled virtual banking experience connects customers with tailored digital representatives that can converse in the customer's preferred language in a timelier manner.

Liberated human capital

With conversational chatbots satisfying most customer inquiries, the human workforce can focus on the most complex or value-driving customer service activities.

Cost avoidance

Virtual agents allow the bank to serve more customers without expanding the human workforce, which helps limit customer service costs.

Focused cyber

AI-driven vulnerability management that focuses on real threats

AI can help filter, analyze, and prioritize security alerts, allowing organizations to focus their resources on true vulnerabilities—which can help reduce their cybersecurity risk.

ISSUE/OPPORTUNITY

Financial institutions and other large enterprises rely on multiple security tools that generate millions of alerts daily—many of which are false positives or non-actionable.

Security teams and developers spend time triaging and remediating vulnerabilities that pose little or no real risk. This inefficiency delays responses to real threats, creates operational fatigue, and increases the likelihood of breaches—threatening business continuity and brand reputation.

Tags **Compliance & Risk**

HOW AI CAN HELP

Evaluating alerts
AI models can evaluate alerts by calculating breachability (likelihood the vulnerability can be tapped in a real-world scenario) and exploitability (likelihood it can be actively taken advantage of to inflict harm).

Setting priorities
An AI-driven system can use enterprise-specific context such as workload risk profiles, application architecture, and network exposure to separate high-risk from low-risk vulnerabilities. AI automatically identifies which threats demand attention and which can be deprioritized, providing a data driven, risk-based prioritization process.

Focused cyber

MANAGING RISK AND PROMOTING TRUST

Robust and reliable
The vulnerability management system should be continuously retrained and tested against live alert data and historical incidents to validate its prioritization logic. A human-in-the-loop approach provides important oversight on critical decisions, especially during early-stage deployment.

Safe and secure
The models should be hosted in secure environments (e.g., on-prem or restricted cloud instances), with strict access controls and audit trails. Systems must be designed to prevent data leakage and be regularly tested for robustness.

Transparent and explainable
AI outputs should include the rationale for prioritization with traceable decision paths that help analysts understand and trust the recommendations.

POTENTIAL BENEFITS

Less noise
Moving from thousands of daily alerts to a handful of real threats enables security and engineering teams to focus their time and effort more effectively.

Improved security and reputation
Decreasing security incidents strengthens a company's standing with customers, investors, and regulators.

Cost savings and improved efficiency
Better resource allocation allows security and engineering teams to spend less time and money protecting the organization from cyberthreats.

Enhanced productivity and morale
Minimizing time spent on false positives reduces friction between the security and engineering teams, improving productivity and morale.

Predictive maintenance and autonomous operations for IT infrastructure & ATMs

Predictive monitoring enabled by edge integration of physical machines

AI models continuously sense, interpret, and act on real-world signals from distributed physical assets (ATMs, servers, cooling systems), enabling autonomous, edge-based detection and intervention to prevent failures, maintain uptime, and ensure operational continuity even in low-connectivity environments.

ISSUE/OPPORTUNITY

Failures of physical IT/OT infrastructure and ATMs directly affect customer access, trading operations, and regulatory obligations. When ATMs fail, customers lose access to cash services—potentially at critical moments—damaging trust and satisfaction, cash handling component wear and sensor faults are common physical failure modes. Server, network, power, and cooling failures can disrupt market access and channels, creating missed opportunities and potential reporting gaps. Traditional reactive maintenance leads to outages, emergency interventions, and elevated operational risk.

The core gap is lack of edge integration across heterogeneous machines—without it, you can't collect high frequency signals reliably, act locally during connectivity loss, or enforce consistent device identity/security—so models remain "blind" or delayed.

The opportunity is to anticipate physical asset degradation earlier and reduce incidents through AI-powered predictive maintenance, while cautiously progressing to limited, controlled action under strict governance and human oversight that satisfies regulatory requirements for accountability.

Tags: Information Technology | Physical AI

HOW AI CAN HELP

Telemetry-driven failure prediction
AI models analyze historical and real-time telemetry sourced through edge gateways/ agents to identify degradation patterns that typically precede failure events, enabling earlier intervention than rule-based monitoring.

Preparation of remediation actions
For low-risk scenarios, AI systems prepare potential remediation steps—such as restart, patching, or component replacement—without executing them autonomously.

Asset-specific health modeling
Different classes of physical assets are modeled separately, accounting for hardware type, age, usage intensity, and operating conditions rather than applying generic thresholds.

Human-in-the-loop control
All actions remain subject to human approval, ensuring accountability, auditability, and regulatory compliance.

Edge-compatible inference
Inference can be executed close to physical assets to meet latency, reliability, and data-residency requirements.

Risk-weighted assessment
Predicted failures are evaluated in terms of customer impact, operational criticality, and regulatory sensitivity, helping teams prioritize responses.

Predictive maintenance and autonomous operations for IT infrastructure & ATMs

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

The main gap being addressed is unreliable edge integration that leaves models blind to degradation signals. A predictive system that misses failures in older or non-standard assets creates exactly the reactive maintenance cycle it is designed to prevent—with direct consequences for customer access, operational continuity, and the business case for deployment.

Responsible and accountable

Human approval for all remediation actions is the mechanism through which the system maintains regulatory compliance. This requires documented records of what the AI predicted, what action was prepared, who approved it, and what followed. Without this audit trail, the human-in-the-loop design provides regulatory assurance in principle but not in practice.

Transparent and explainable

Financial institutions cannot deploy predictive systems they can't clearly explain to regulators. Models must be able to show which signals indicated degradation, how failure probability was scored, and why certain assets were prioritized, to a standard that satisfies regulatory examination.

POTENTIAL BENEFITS

Reduced downtime

Fewer customer-facing outages as predictive maintenance prevents ATM and infrastructure before service impact and edge-local detection reduces time-to-detect and time to triage.

Maintenance efficiency

Shifting from reactive to predictive maintenance allows technicians to address issues during scheduled windows rather than responding to urgent failures.

Lower operational risk

Earlier, planned interventions reduce the likelihood of emergency situations that could trigger regulatory scrutiny or customer complaints.

Humanoid robots for branch operations

Automating front-of-house experiences with governance

Humanoid robots act as a "digital receptionist" in bank branches, sensing customer presence, movement, and queue dynamics to autonomously manage check in, triage, and routing. By combining embodied interaction, on device perception, and governed human oversight, they optimize front of branch operations in regulated financial environments.

ISSUE/OPPORTUNITY

Branch lobbies still rely heavily on human reception and ad-hoc queue handling, which can create inconsistent experiences, longer waits, and avoidable staff load, especially during peak traffic. Customers often need quick routing (who to see, where to go, what to bring) rather than a full teller interaction, but traditional processes don't scale efficiently. Banks also need ways to reduce friction at the front desk while maintaining service continuity and reinforcing in-branch experience standards. The opportunity is to deploy humanoid robots to digitalize branch reception, reducing delays and improving access to information via automated check-in/out, appointment booking, and real-time wait-time updates, while preserving human staff for higher-value interactions.

Tags: Customer Service | Physical AI

HOW AI CAN HELP

Spatial perception & crowd sensing

Robots detect queue length, congestion, and customer flow using vision and proximity sensors—adjusting triage behavior dynamically without manual supervision.

Edge based execution for resilience

Core sensing and routing logic runs locally, ensuring continued branch operation during network latency or partial outages.

Conversational triage

Answers common questions and routes to the right staff/team. Captures intent fast (e.g., "replace card," "open account") to enable the right first handoff.

Embodied escalation control

Physical presence enables controlled handoff—robot can physically guide customers to service points or staff, reducing misrouting and confusion.

Branch aware behavior

Robots adapt interaction patterns based on branch layout, time of day, and live staffing levels—something screen based systems cannot do.

Humanoid robots for branch operations

MANAGING RISK AND PROMOTING TRUST

Private

A humanoid robot continuously sensing customer presence, movement, and behavior in a bank branch collects biometric and behavioral data about identifiable individuals who have not necessarily consented to automated observation. Organizations must be transparent about what is captured, apply strict retention limits, and ensure that data is not repurposed beyond the immediate branch interaction it was collected to support.

Responsible and accountable

Financial services regulators hold institutions accountable for customer outcomes regardless of whether interactions are delivered by humans or automated systems. Accountability cannot be passed to the technology vendor when something goes wrong in a regulated customer-facing environment.

Fair and impartial

A triage system that works well for the majority but underperforms for customers with different languages, communication styles, or accessibility needs creates unequal access to branch services. In a regulated financial environment serving a diverse public, this is both a reputational and a compliance concern. Performance must be validated across the full customer population.

POTENTIAL BENEFITS

Reduced wait times

Faster intake and better queue flow. Fewer re-queues and faster routing improve throughput during peak periods.

Operational consistency across branches

Embodied behavior enforces standardized intake and routing regardless of staff mix, branch size, or local practices.

Measurable physical KPIs

Enables tracking of queue density, dwell time, handoff accuracy, and physical congestion—not just digital metrics.

Lower exception load on human staff

Staff engage only when physical AI flags ambiguity, VIP handling, or regulatory exceptions—preserving human judgment where it matters.

ATM cash forecasting, replenishment and autonomous cash logistics

Predictive cash availability through edge intelligence

AI models running on edge-equipped ATMs predict cash depletion using physical machine signals and contextual factors (day-of-week, seasonality, local events). Cash forecasting and intelligent replenishment are being deployed with the goal of autonomous coordination of cash logistics.

ISSUE/OPPORTUNITY

Out-of-cash ATMs erode customer trust and drive-up emergency servicing costs. When customers hit empty ATMs, frustration and inconvenience reduce satisfaction and loyalty and can push them to competitors with more reliable cash access. Static replenishment thresholds break under variable demand because fixed rules don't adapt to changing withdrawal patterns from holidays, local events, weather disruptions, or neighborhood shifts.

Manual planning also struggles to incorporate irregular but predictable spikes (concerts, sporting events, payday cycles). Without edge integration, real time visibility into machine health, cash levels, and local demand signals remains limited, delaying detection and response.

The opportunity is to equip ATMs with edge intelligence that continuously interprets physical machine signals and contextual demand drivers, enabling accurate cash forecasting, risk based prioritization, and optimized replenishment planning—reducing outages and logistics costs while keeping execution under strict human and security controls.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

Edge-enabled cash demand forecasting

AI models run on edge-connected ATMs, combining cash-level telemetry and device signals with usage patterns to predict depletion timelines more accurately than fixed thresholds.

Dynamic replenishment recommendations

AI adjusts replenishment timing and quantities based on predicted demand rather than relying on static refill schedules.

Contextual signal integration

Models incorporate temporal effects, seasonal patterns, and local events influencing withdrawal behavior to anticipate demand changes before they occur.

Logistics planning support

In the future, AI could assist in coordinating routes and schedules for cash handling teams to optimize operational efficiency.

Human-controlled execution

All replenishment and logistics decisions remain under human oversight, ensuring accountability for security-sensitive operations.

Risk-based prioritization

ATMs with the highest likelihood of customer impact are prioritized for intervention based on location importance and depletion probability.

ATM cash forecasting, replenishment and autonomous cash logistics

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Inaccurate forecasting directly produces the out-of-cash events and unnecessary logistics costs the system exists to eliminate. Models must be validated across the full range of demand conditions they will encounter in production, including irregular spikes from local events and neighborhood shifts that fixed thresholds and historical averages cannot anticipate.

Safe and secure

AI forecasting outputs that reveal replenishment timing and ATM cash status are operationally sensitive intelligence that, if compromised, could directly enable criminal exploitation of logistics operations. This is not a generic cybersecurity concern; it is specific to this use case. Forecasting outputs and machine-level cash data must be secured with protections commensurate with the physical security risks of cash handling.

Fair and impartial

Risk-based prioritization that favors high-traffic or commercially significant ATMs may systematically result in lower cash availability in lower-income or lower-density areas where customers have the fewest alternatives. Cash availability is a basic financial access need, and for institutions with public service obligations, AI-driven prioritization that disadvantages underserved communities is both a regulatory and a reputational concern.

POTENTIAL BENEFITS

Fewer out-of-cash events

Higher ATM availability as predictive forecasting prevents cash depletion before customers are affected.

Improved customer experience

More reliable access to cash strengthens customer trust and satisfaction by ensuring ATMs are available when needed.

Lower logistics costs

Fewer emergency replenishment trips and better route density driven by real-time edge visibility into depletion risk and machine status.

The Government & Public Services AI Dossier



The Government & Public Services AI Dossier

AI is becoming a critical tool for driving mission outcomes in the government and public services industry. From frontline service delivery to policymaking and infrastructure management, AI systems are reshaping how government and public services organizations operate, deliver services, and engage with constituents.

In an environment often constrained by legacy systems, budget pressures, and complex regulatory mandates, AI offers the ability to modernize processes, improve responsiveness, and unlock new levels of efficiency and insight.

AI is already being deployed for a wide range of use cases across the sector, such as drafting policy documents, analyzing and summarizing global legislation and policies, simulating urban planning scenarios, hyper-personalizing education, and engaging with the public through virtual public servants.

Agentic AI enables more autonomous support across many areas—including regulatory oversight, permitting, benefits, intelligence, and health documentation—augmenting, not replacing, human expertise. Similarly, physical AI offers governments powerful tools to protect people, deliver services, and manage enormous physical domains (roads, utilities, public buildings, transit systems, public safety infrastructure, etc.) that are chronically under-resourced given the maintenance and oversight they require.

AI innovations, when paired with secure data integration and worker oversight, can reduce administrative burdens, streamline operational workflows, and enhance the precision and speed of government decision-making.

As government and public sector use cases mature and the regulatory landscape evolves, AI is poised to become not just a tool for modernization, but a catalyst for more adaptiveand data-driven governance. However, the adoption of AI in government and public services also brings unique challenges.

Public institutions must meet high standards for transparency, fairness, and accountability—often in environments where privacy concerns, procurement complexity, and uneven digital maturity can limit deployment. As such, success depends not just on technical innovation, but on cross-agency collaboration, trust-building with the public, and alignment with legal and ethical frameworks.

AI systems are reshaping how government and public services organizations operate, deliver services, and engage with constituents.

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-supported regulatory examinations and inspections

Augmenting regulators with multi-agent systems for fast, more accurate review

Agentic AI systems can support regulators with specialized agents that analyze data, review large volumes of documents, and uncover insights to streamline examinations and inspections.

ISSUE/OPPORTUNITY

Regulatory examinations and inspections are essential to ensuring compliance in industries such as financial services, food safety, and transportation. These processes require regulators to review vast amounts of historical and current market data, pull from multiple document sources, and interpret complex regulations. The work is labor-intensive, time-consuming, and prone to bottlenecks, which can delay the detection of compliance risks.

As regulatory responsibilities expand and the business environment becomes increasingly complex, regulators need tools that can streamline and accelerate analysis without sacrificing quality or rigor. Agentic AI offers the ability to automate routine examination tasks, highlight anomalies, and provide real-time support to examiners.

Tags: **Compliance & Risk** | **Agentic AI**

HOW AI CAN HELP

Quantitative analysis
AI agents can extract formulas, run calculations, and compare data against regulatory benchmarks to identify areas of potential concern.

Document review
The systems can process large volumes of reports, filings, and other documents, uncovering key insights and flagging areas that warrant attention from human examiners.

Dynamic examiner support
A conversational interface allows regulators to interact with the system, request clarifications, and guide the direction of analysis, ensuring a human stays in the loop.

Cross-domain applicability
The framework currently being used for financial services examinations can be adapted to other types of regulatory inspections, including food safety, tax compliance, and transportation.

AI-supported regulatory examinations and inspections

MANAGING RISK AND PROMOTING TRUST

Private	Because examinations can involve sensitive institutional and personal data, agentic systems should employ strict data protection protocols and comply with all relevant privacy regulations.
Robust and reliable	Flawed analyses can lead to missed risks or incorrect enforcement actions. To avoid such problems, AI agents should be validated against known regulatory outcomes and tested continuously for accuracy.
Responsible and accountable	Since regulatory judgments can have serious legal and economic consequences, AI agents should be used primarily as decision-support tools, with human regulators retaining final authority over important decisions and actions
Fair and impartial	To ensure equitable oversight, agents should be designed to apply standards consistently across institutions and sectors, minimizing the risk of bias in how examinations are prioritized or conducted.

POTENTIAL BENEFITS

Faster examinations at lower cost

Automated analysis and document review can shorten the time required to complete inspections, helping regulators do their jobs efficiently and effectively.

Improved risk detection

By examining vast amounts of quantitative and qualitative data, AI agents are more likely to detect problems or anomalies that could go unnoticed.

Scalable oversight

With AI agents handling routine analysis, regulators can expand their oversight capacity without a corresponding increase in staff, helping to maintain compliance in complex and high-volume domains.

AI-driven permitting

Accelerating and simplifying government permitting processes

Agentic AI systems can streamline permitting workflows by using specialized agents to scan applications, extract key data, check compliance, and provide real-time feedback to applicants and reviewers.

ISSUE/OPPORTUNITY

Permitting is a critical function for government agencies at the federal, state, and local level, covering areas such as land use, energy, waterways, and housing. These processes can be slow, opaque, and labor-intensive, requiring reviewers to manually assess large volumes of documentation against complex and sometimes overlapping regulations. Small errors or missing information in applications can create lengthy delays, increasing costs for both applicants and agencies. With government entities facing larger workloads and limited staffing, applicants and regulators alike are seeking opportunities for efficiencies. New approaches can help accelerate permitting while ensuring compliance with local, state, and federal requirements.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Automated document review
AI agents can scan and extract relevant data from submitted applications, checking for completeness and identifying missing or inconsistent information.

Compliance validation
Specialized agents can evaluate applications against permitting requirements, standard operating procedures (SOPs), and applicable regulations, flagging potential compliance issues early in the process.

Applicant feedback and support
Using AI and other advanced technologies, agents can provide real-time feedback to applicants, helping them correct errors or supply missing documentation before delays occur.

Reviewer assistance
AI agents can support human reviewers by summarizing key findings, verifying that evaluations align with regulatory standards, and highlighting areas requiring human attention.

AI-driven permitting

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	To maintain public trust, agents should be designed to apply requirements consistently across all applicants.
Robust and reliable	Errors in interpreting regulations can lead to incorrect approvals or major delays. To improve consistency and reliability, agents should be trained on verified SOPs and continuously updated as requirements evolve.
Private	Since applications often contain sensitive personal and corporate data, agentic AI systems should be designed with strong data protection measures and full compliance with privacy regulations at all three levels of government.
Responsible and accountable	Permitting decisions can have important legal and financial ramifications. To avoid errors in accuracy or judgment, agentic AI systems should be primarily used for decision support, with final approval authority retained by human reviewers.

POTENTIAL BENEFITS

Reduced costs and permitting timelines

By automating document review and compliance checks, processing times can be reduced from months to days, lowering costs and accelerating the permitting process.

Improved consistency and fairness

Standardized evaluations conducted or supported by AI agents can help all applicants be assessed against the same criteria, reducing variability across reviewers.

Reduced applicant frustration

Real-time feedback improves transparency and helps applicants correct issues earlier, reducing costly delays and improving trust in the process.

AI-enhanced benefits eligibility

Improving access to government benefits programs

AI systems can streamline the application and determination processes for government benefits programs and public assistance.

ISSUE/OPPORTUNITY

Determining eligibility for public benefits such as health services, food stamps, unemployment insurance, workers comp, and childcare can be a time-consuming process for applicants and government agencies alike. Traditional public assistance portals may have redundant data entry, unclear requirements, and complex navigation, which can frustrate applicants (especially those with limited digital literacy) and delay access to benefits.

Meanwhile, government caseworkers manage lengthy manual intake processes and may have a lack of support in interview preparation and execution. These challenges increase agency workloads and can lead to incomplete applications, inconsistent reviews, and a higher likelihood of errors.

Tags: **Customer Service** | **Agentic AI**

HOW AI CAN HELP

Conversational intake
Conversational agents can replace static application forms with guided, human-like interactions that clarify questions and adapt dynamically to applicant responses.

Automated document intelligence
Document analysis agents can extract and validate data from submitted records, reducing manual effort and maintaining accuracy in eligibility assessments.

Interview assistance
AI agents can support human caseworkers by preparing interview outlines, validating applicant responses in real time, and highlighting missing or inconsistent information.

Onboarding and on-the-job assistance
AI can accelerate new caseworker onboarding by generating tailored training materials and insights from historical cases, helping agencies address workforce challenges more efficiently. AI agents can also support human caseworkers by preparing interview outlines, validating applicant responses in real time, and highlighting missing or inconsistent information.

End-to-end applicant support
The specialized agents can work together across the intake and review processes to provide applicants with a seamless experience that is less intimidating and confusing, helping to ensure applications are complete and compliant.

AI-enhanced benefits eligibility

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable

In order for applicants and caseworkers to trust AI-driven recommendations, agents should provide clear explanations of eligibility determinations, including references to applicable rules and policies.

Robust and reliable

Conversational agents might misunderstand user inputs, especially in cases of language barriers or low digital literacy. Agents should be trained on multiple languages and continuously monitored for performance.

POTENTIAL BENEFITS

Faster access to benefits

Automating intake and eligibility can reduce processing times, helping applicants receive support more quickly.

Cost savings and improved productivity

Using AI agents to streamline processes and automate manual tasks can reduce costs and enable caseworkers to focus more attention on complex cases and applicant support.

Improved applicant experience

Guided interactions, real-time feedback, and simplified workflows can reduce frustration and confusion, especially for applicants with limited digital literacy.

Smarter intelligence

Improving intelligence reporting via AI agents and multimodal data integration

Agentic AI systems can autonomously collect, synthesize, and cross-reference intelligence data from various sources, producing faster, more comprehensive insights.

ISSUE/OPPORTUNITY

Intelligence analysts, including those working with open-source intelligence (OSINT), face the daunting task of searching, collecting, and interpreting vast amounts of multimodal data—from social media and public records to satellite imagery, communications, and financial reports. This manual process is time-consuming and limited by the capacity of human analysts to track patterns across large datasets.

As threats evolve and the volume of public and classified data continues to grow at breakneck speed, agencies need tools that can help accelerate the production of intelligence reports while improving accuracy, adaptability, and depth of insight. Agentic AI offers the potential to automate core intelligence gathering activities so human analysts can focus on validation and higher-level reasoning.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Automated data discovery
AI agents can continuously scan for and retrieve information from multiple structured and unstructured data sources, both classified and public—including briefings, news media, and other reports—reducing the need for manual searches and data collection.

Contextual linking of insights
If analysis of an old report produces new leads—such as a location or personal association—AI agents can automatically launch related queries to build a more complete picture of personal networks and activities.

Multimodal data integration
Specialized AI agents can process different types of inputs (e.g., text, imagery, sensor data) and combine the results into comprehensive, integrated insights.

Adaptive intelligence reporting
As the various agents collaboratively draft and refine reports, they can highlight key findings and update them dynamically as new data becomes available.

Smarter intelligence

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Automated agents might misinterpret ambiguous or misleading data, potentially affecting the accuracy of intelligence outputs. This is especially true for open-source information, which can be intentionally misleading or fake. AI agent actions should be validated against verified sources and continuously monitored to reduce errors.

Safe and secure

Given the sensitive nature of intelligence queries, special care must be taken to prevent adversaries from influencing the model (or gathering their own intelligence from what is queried).

Transparent and explainable

Providing adequate transparency in autonomous information exchange can pose significant technical and operational hurdles, especially given the complexity of integrating data from multiple sources. AI agents should document their reasoning and data sources so the outputs are easier to trust and audit.

Responsible and accountable

Given the life-and-death consequences of intelligence information, AI agents should be viewed as decision-support tools, with human analysts and leaders continuing to have final responsibility for decisions and actions.

POTENTIAL BENEFITS

Faster examinations at lower cost

Automating data collection and integration accelerates intelligence reporting and reduces manual effort, which can enable agencies to respond more quickly and efficiently to emerging threats.

Deeper insights

Cross-modal, autonomous querying and analysis can uncover hidden connections that might be missed with traditional manual processes, providing a more comprehensive and dynamic understanding of complex domains.

Improved analyst capacity

By offloading routine search and synthesis tasks, AI agents allow human analysts to focus on higher-level reasoning, strategic decision-making, and collaboration with colleagues.

Automated clinical documentation

Reducing provider burden by automating patient visit notes

Agentic AI systems can automatically generate structured clinical documentation from patient visits, reducing providers' workload, improving record quality, and enabling them to spend more time with patients.

ISSUE/OPPORTUNITY

For providers in large government health care systems, documenting patient visits can be an administrative burden. Traditional processes require physicians to spend hours after each visit manually writing SOAP (Subjective, Objective, Assessment, Plan) notes and updating electronic health records (EHRs). This not only disrupts the flow of care but often requires after-hours "pajama time" spent on paperwork, which contributes to provider burnout. Also, incomplete or inconsistent documentation creates risks for care quality, regulatory compliance, and continuity across providers.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Automated transcription and generation
A generator agent listens to the audio of patient visits, converting it into structured SOAP notes and EHR-ready data, which minimizes the need for manual entry.

Validation and compliance
A validator agent reviews the generated notes for clinical accuracy, completeness, and regulatory compliance, confirming that records meet organizational and legal standards.

Workflow orchestration
An orchestrator agent manages the entire documentation process, routing tasks among the specialized agents and providing smooth, efficient workflows.

Seamless integration
The agentic AI system formats the validated documentation into common formats (e.g., DOCX, HTML) for direct integration into EHR systems, reducing friction and manual effort for providers and staff.

Automated clinical documentation

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	Documentation agents should be designed and tested to work effectively across varied patient populations and care contexts.
Robust and reliable	Because errors in transcription or orchestration could compromise documentation quality, AI agents should be validated regularly and continuously updated to reflect evolving medical and compliance requirements.
Transparent and explainable	For providers to trust that AI-generated documentation accurately reflects patient encounters, AI agents must be able to explain how notes were created and why specific items were flagged for correction.
Private	Some patients might be uncomfortable having their medical visits recorded and then analyzed by AI. Every patient should be given the choice to opt out. Also, AI agents must strictly comply with HIPAA and other relevant regulations, encrypt audio and text records, and minimize retention of personally identifiable information.

POTENTIAL BENEFITS

Lower costs and reduced provider burnout

Automating documentation can reduce clinical staff workload, reducing costs and freeing providers from after-hours paperwork, which can improve their work-life balance.

Improved record quality and compliance

Standardized, AI-validated notes can improve the completeness and consistency of patient records, enhancing care continuity and compliance with regulations.

Data-driven insights

Structured documentation enables advanced analytics and quality improvement initiatives, helping government health systems improve care delivery at scale.

Virtual public servant

Citizen engagement

AI can enable virtual assistants that provide personalized responses to citizen questions about public services.

ISSUE/OPPORTUNITY

Government organizations perform a range of functions, from supporting public health to promoting tourism. Data about government and public services, however, is often stored in a variety of formats and locations (e.g., on-prem, cloud), challenging interoperability. When citizens contact agencies to inquire about services and resources, human agents are challenged to rapidly access and summarize information to satisfy citizen questions. This is a time-consuming, labor-intensive endeavor for the organization, and it may not meet citizen expectations for fruitful engagement.

Tags: **Customer Service**

HOW AI CAN HELP

A digital agent for engagement
An AI-enabled virtual assistant can serve as the interface between citizens and government information, helping with questions and transactions via empathetic, natural language.

Reaching across datasets
The virtual assistant can distill and summarize information from myriad sources on a variety of topics to answer questions in a multitude of languages regarding service requirements and appointment options.

Virtual public servant

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

While virtual assistants may be valuable for providing information, they may not be suited to providing true insight and advice. Agencies need to guard against over-reliance on AI solutions and the potential for citizens to take some action based on faulty or improper AI output.

Safe and secure

A model tasked with providing accurate information may be a target for cyber criminals seeking to access sensitive information or manipulate the model and its underlying data. Many government agencies contend with cybersecurity regulations and standards, making model security a priority.

Robust and reliable

Model accuracy and timeliness depends in part on the data sources it can access. If information is outdated or incorrect, it creates a risk of erroneous outputs. Human stakeholders responsible for updating information have a direct impact on model reliability and user trust.

POTENTIAL BENEFITS

Promoting citizen engagement

When public services are more accessible due to more efficient and robust technology, user engagement and citizen satisfaction in government offerings increase.

Increasing accessibility

A virtual assistant powered by AI can interact with citizens in their preferred language and ultimately help bring down social barriers to engaging public services.

Citizen satisfaction

Government agencies operate in service to the public, and providing fast access to information about services promotes a positive public perception of government function.

Insights for all

Knowledge management

AI can serve as an interface to help public sector organizations become insight-driven by making data more accessible.

ISSUE/OPPORTUNITY

From the census to transportation and procurement, government agencies collect and release huge amounts of open datasets. By encouraging the use, reuse, and distribution of open datasets, government organizations can promote data-driven innovation and citizen-centric services if combined with an agency's internal datasets. For public industry stakeholders to become truly insight-driven, they require the means to interrogate all relevant data, even if they lack a technical background in data science or related fields.

Tags: **Cross-functional**

HOW AI CAN HELP

Greater accessibility
AI can provide a natural language interface that allows non-technical users to access and understand data that might otherwise only be accessible to technical users.

Democratizing insights
Rather than placing all of the burden for data analysis, interpretation, and visualization on a technical team, an AI interface reduces the effort by allowing more stakeholders to work with the data and derive their own insights.

Insights for all

MANAGING RISK AND PROMOTING TRUST

Safe and Secure

An AI model that taps into a variety of datasets can make it difficult to control which data is accessed by which stakeholders in which organization, raising important considerations for model security and governance.

Private

When dealing with sensitive and proprietary information that is subject to varying laws and regulations across jurisdictions, organizations are called on to ensure the AI model does not leak, inadvertently divulge, or inappropriately access sensitive or restricted data.

Transparent and explainable

To accurately interpret data and AI outputs, the end user needs to understand which data was referenced for the output, which could not be accessed, and the potential biases in the available data.

POTENTIAL BENEFITS

Scaling data access

An AI solution that can access a variety of datasets and data types allows public servants to draw conclusions from a broader set of knowledge and information.

Fostering collaboration

When more public servants can access insights and knowledge, it promotes insight-driven action across agencies, helping to fuel greater collaboration between a larger set of stakeholders.

Faster insights

AI can help accelerate the process of identifying and consuming relevant information, driving speed and efficiency.

Simulating urban planning scenarios

Urban planning/future of cities

AI can be used to help urban planners in the ideation and design of novel urban concepts.

ISSUE/OPPORTUNITY

More than 56 percent of the world's population—4.4 billion people—lives in cities.[1] By 2050, the urban population is expected to double, with upwards of 70 percent of people living in cities. The scale and speed of urbanization brings a host of challenges, such as lack of affordable housing, overburdened transportation systems, traffic congestion, lack of drinking water, sanitization issues, and environmental quality. The challenge for city officials and urban planners is to imagine the future of cities by overcoming creative hurdles and developing city designs that are resilient, sustainable, and human centric.

Tags | R&D/Product Development

HOW AI CAN HELP

Generating 3D city models
Using AI, thousands of 3D images can be rapidly created to help guide and refine a city design. Such 3D images form part of the design brief for urban planners and the master city plan.

Simulate natural disasters
AI can simulate natural disasters like earthquakes, floods, or hurricanes to evaluate the vulnerability of city infrastructure and plan for resilient urban infrastructure.

Planning for the future
By simulating population growth and demographic trends, AI can develop scenarios for urban expansion and plan for adequate infrastructure, housing, transportation, and public services that accommodate urban growth.

Simulating urban planning scenarios

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

While AI models may create interesting or attractive designs, they require human review and validation to ensure they meet urban planning requirements and can be feasibly built in the real world.

Transparent and explainable

A lack of contextual knowledge of urban planning may lead AI to develop improbable scenarios, and analysts need to be able to understand how and why the model produced an output in order to confirm and validate it.

POTENTIAL BENEFITS

Super-charge creativity

Using AI to rapidly create a plethora of designs and scenarios helps city officials imagine the future of cities and plan for upcoming challenges.

Faster ideation and iteration

With a faster method to create design iterations, urban planners can accelerate the design and decision-making processes.

Improved decision-making

Using AI in city planning enables decision-makers to model various scenarios and optimize urban designs for better resource utilization, sustainability, and quality of life for residents.

Education 2.0

Hyper-personalized education

AI can be used to hyper-personalize digital teachers that can adapt to student learning needs and curricula.

ISSUE/OPPORTUNITY

The demand for schoolteachers can often exceed supply. While the available teachers contend with larger class sizes, they also need to accommodate students with different learning styles and educational needs. Yet, because of the one-to-many nature of traditional schools, teachers are challenged to deliver the kind of personalized learning support and instruction that students need to be successful.

HOW AI CAN HELP

A digital, adaptive teacher
AI can serve as a virtual instructor, drawing from resources and lesson plans to hyper-personalize the learning experience. The model can check the student's work and comprehension and adapt lessons and learning strategies according to the student's individual weaknesses, strengths, and preferences.

A force multiplier for teachers
When personalized digital teachers can work with students one-on-one to master new skills and knowledge, the human instructor can focus on higher level planning, interacting with students, evaluation, and student support.

Tags: **Learning & Development**

Education 2.0

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

While digital teachers can offer valuable advantages in adaptive learning, the model should not be expected to satisfy all of the important lessons teachers impart, such as social lessons around collaboration, conflict resolution, and empathy. The human element in teaching is essential, and educational institutions need to take a responsible approach to integrating AI-enabled teachers.

Robust and reliable

Because AI is susceptible to outputting inaccuracies and hallucinations, there is a risk that a virtual teacher could teach incorrect facts or produce poor learning strategies.

Private

Student data is subject to education regulations, making model security and data privacy a priority when deploying digital teachers.

POTENTIAL BENEFITS

Catering to the student

Employing adaptive learning with AI can promote knowledge retention and understanding by tailoring teaching approaches to the student's learning style.

Remedy the talent gap

Leveraging AI helps overcome teacher shortages, allowing more students to access quality education.

Removing barriers

An AI-enabled teacher is not restricted to a physical classroom. With online access, digital teachers could be accessible to students in any environment or geography, helping to bring down barriers to attending school.

Digitizing policymaking

Policy creation assistant

AI can be used to search large volumes of policy documents and output natural language responses to user queries in complex policy environments.

ISSUE/OPPORTUNITY

Because the data relevant to government agencies is stored in different locations and formats, it can be difficult for analysts and policymakers to effectively query datasets and retrieve relevant information in a timely manner. With nomenclature issues, it can also be challenging to identify associated data topics and types. The result is a diminished ability to digitize policymaking and discussion, as well as increasingly complicated interactions around policy matters.

Tags: **Operations**

HOW AI CAN HELP

AI assistant
AI can identify data related to common themes and topics and then summarize that information in response to user queries, helping to identify policy differences, conflicts, and gaps.

Citizen engagement in policymaking
Using AI, governments can create interactive platforms and chatbots that encourage citizens to participate in policymaking discussions. The AI-driven interface can gather public opinions and feedback on policies, making it easier for citizens to voice their views.

Digitizing policymaking

MANAGING RISK AND PROMOTING TRUST

Private

Some of the data relevant to policy issues may be sensitive or restricted, and the AI model may require controls to limit which users can access which datasets.

Fair and impartial

Various stakeholders aim to influence policymaking. AI might be biased in giving higher weight to comments and input coming from some sources over others. This has the potential to produce biased policies that favor certain businesses or segments of society.

POTENTIAL BENEFITS

Data query at scale

By reviewing large volumes of policy documents, AI can help users accelerate information gathering and increase their capacity and efficiency in querying disparate datasets.

Participatory policymaking

Using AI to better identify and incorporate a varied views and stakeholders supports more robust and complete representation in policy matters.

Drafting contracts and SoWs

Procurement

AI can analyze offerings from existing vendors, match them to organizational needs, generate requests for proposals, and analyze the responses.

ISSUE/OPPORTUNITY

Governments procure billions of dollars in goods and services annually.[2] Traditionally, government procurement requires significant volumes of paperwork, which can lead to delays. Many government procurement contracts are highly detailed and often incorporate a range of clauses and requirements from payment terms to export controls to wage and workforce requirements. Drafting requests for proposals (RFPs) and contracts and then generating statements of work (SoWs) requires significant time and resource investments.

HOW AI CAN HELP

Automated drafting
AI can automate the RFP and SoW writing processes by generating the initial drafts based on templates, historical documents, or specific prompts provided by procurement officials.

Extracting information
AI's advanced Natural Language Processing (NLP) capabilities can help extract relevant clauses and requirements from existing contracts, SoWs, and legal documents. Such information can be used to either create new contracts or assess the risks posed by existing contracts.

Tags | **Procurement/Sourcing & Supply Chain**

Drafting contracts and SoWs

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable

AI may not be able to explain why certain clauses are added to a contract while others are excluded, which is vital information for the human user validating the outputs.

Private

Ingesting existing and historical contract data may pose data privacy and legal hurdles. Model governance is necessary to ensure the AI model, as well as the underlying data, meet privacy rules, regulations, and standards.

POTENTIAL BENEFITS

Time savings

Creating initial document drafts with AI can expedite the writing process and lead to significant time savings, compared to manually creating each RFP or SoW from scratch.

Improved consistency

AI can develop drafts while adhering to predefined guidelines in prompts, which supports a greater level of consistency across report writing.

Multilingual citizen services

Service delivery

AI can help with language translation to support the delivery of services to citizens.

ISSUE/OPPORTUNITY

Many governments around the world serve populations with varying language proficiency and linguistic backgrounds. This challenges agencies to develop multilingual websites, translate official documents, and support frontline workers with translation tools so they can better communicate with all citizens.

Tags: **Customer Service**

HOW AI CAN HELP

Aiding frontline workers
AI can be used to create real-time audio and text messages in different languages as frontline workers interact with residents around a variety of services, such as social care, health care, and emergency response.

Translating official documents
Government agencies often handle the publication of official documents, laws, regulations, and policies. AI can streamline the translation process and help produce accurate and consistent translations.

Announcement and website translation
Government websites and public information (e.g., health and travel advisories) can be translated quicky to make essential information more accessible to the population.

Multilingual citizen services

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

The data used to train an AI model for use in translation may not be consistently accurate or robust across all languages, which could in turn lead to poorer translations and less access to citizen services for some language speakers than for others.

Private

The translating model may be exposed to sensitive information, necessitating steps to ensure the model does not mishandle or inappropriately divulge protected data and thus violate data privacy regulations.

POTENTIAL BENEFITS

Real-time translation

When audio or text can be translated into a multitude of languages in real time, it enables more seamless and conversational interactions with different language speakers.

Translation at scale

AI can handle large volumes of document translation, giving an agency capacity to ensure government information and services are accessible to a greater audience.

Summarizing legislative documents

Legislative administration

AI can help legislative staff more rapidly transcribe and summarize hearings, legislation, documents, and official announcements.

ISSUE/OPPORTUNITY

Legislative offices are expected to hold hearings on important topics, respond to constituents, and make public announcements in the form of press releases. Manually transcribing hearings and meetings is a time-consuming task.

Further, developing new legislation (where staff play a pivotal role in research) requires sifting through voluminous policy proposals and research published by experts.

Tags: **Operations**

HOW AI CAN HELP

Summarizing official documents
Auto-generating transcripts of hours-long committee hearings and summarizing important bills and hearings can significantly reduce the administrative burden on staffers.

Process and summarize policy proposals and research
Legislative staff review a large volume of policy proposals and recommendations published by experts. AI can quickly summarize the documents for them, so staffers can spend more time on higher level policy analysis and decision-making.

Summarizing legislative documents

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

AI may perpetuate latent biases based on its training set and generate skewed summaries that are partisan and favor certain ideologies.

Private

Ingesting internal policy proposals can expose sensitive information, requiring organizations to take measures that protect the confidentiality of internal documents.

POTENTIAL BENEFITS

Reducing burdens

Generating summaries of official hearings can reduce administrative burdens on legislative staff so they can focus on more complex tasks.

Saving time

AI can quickly retrieve information and summarize it, saving legislators and staff time when reviewing lengthy, complex, or detailed documents.

Global policy tracking

Automated tracking and analysis of public policy

Organizations in the public and private sector can use AI to monitor, interpret, and analyze public policy developments in real time across hundreds of countries.

ISSUE/OPPORTUNITY

Tracking national policy developments on a global scale is a resource-intensive and highly fragmented process. Policy documents vary in language, structure, formatting, and accessibility, making it difficult for international organizations, governments, and advocacy groups to maintain a coherent and timely view of global policy trends.

HOW AI CAN HELP

Data collection and analysis
AI can automatically gather, structure, and analyze vast volumes of policy documents from government websites and public sources around the world. The technology performs multilingual data extraction, applies natural language processing to categorize and summarize policies, and synthesizes insights into structured outputs that can be validated by subject-matter experts.

Tags: **Compliance & Risk**

Global policy tracking

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

The AI model should be designed to avoid reinforcing systemic biases. Human experts from different backgrounds and regions are embedded in the feedback loop to validate model outputs, helping to ensure representation across geographies and policy contexts.

Robust and reliable

Automated systems should undergo rigorous, iterative testing to help ensure the reliability of outputs. Policy insights are continuously benchmarked against human analysis and real-world policy documents to maintain a high level of accuracy and dependability, especially in politically sensitive or under-reported regions.

POTENTIAL BENEFITS

Greater accuracy, scalability and knowledge sharing

AI can enable ongoing, real-time monitoring and analysis of thousands of policies across hundreds of countries without an exponential increase in manual effort. It can also enable local entities and other stakeholders to identify global policy trends, compare regional approaches, and uncover best practices.

Improved efficiency

Combining data collection and policy analysis with AI-powered automation can save organizations significant time and resources.

Improved decision-making

Use of AI for policy tracking and analysis can provide organizations with more timely, structured, and reliable data to support effective decision-making and planning.

Autonomous patrolling and threat detection

Robotic patrolling for government facilities

Autonomous mobile robots (AMRs) patrol buildings and facilities, continuously monitoring environments, detecting potential threats, and escalating incidents based on severity. Systems interpret scenes locally and alert security personnel when intervention is required.

ISSUE/OPPORTUNITY

Continuous physical security coverage is costly and difficult to staff with humans alone. Fixed cameras lack mobility and context. Security guards patrol facilities on fixed routes or remain stationed at checkpoints, creating predictable gaps in coverage that leave areas unmonitored for extended periods.

Fatigue and attention degradation affect human guards during overnight shifts or long patrol routes, reducing vigilance when monitoring repetitive environments where incidents are infrequent. Guards cannot be in multiple locations simultaneously, forcing facilities to choose between comprehensive coverage requiring large security teams or accepting blind spots during routine patrols. Fixed cameras provide static views but cannot investigate suspicious activity, follow subjects through facilities, or provide contextual assessment of situations that require mobile perspective.

Physical AI robots can provide persistent, mobile monitoring while reserving human attention for verified incidents requiring judgment, de-escalation, or physical intervention.

Tags: Compliance & Risk | Physical AI

HOW AI CAN HELP

Autonomous navigation

Robots move through buildings independently, following patrol routes, navigating around obstacles, using elevators, and adapting paths when areas are blocked or access is restricted.

Continuous operation

Robots patrol without fatigue, maintaining consistent coverage during overnight hours, weekends, and holidays when staffing human security teams is most expensive and challenging.

Threat escalation logic

Events are classified by severity, with minor anomalies logged for review while urgent threats trigger immediate alerts with live video feeds to security personnel.

Human alerting

Security teams are notified when needed, receiving situational context, robot location, and video evidence to inform their response without requiring constant monitoring of robot feeds.

Visual scene interpretation

AI identifies abnormal situations such as unauthorized personnel in restricted areas, doors left open after hours, unattended packages, or environmental hazards like water leaks or smoke.

Autonomous patrolling and threat detection

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Robots that misclassify routine activity as threats create alert fatigue that causes security teams to discount warnings; those that miss genuine incidents defeat the purpose of deployment. Both failure modes are serious. Reliability across varying lighting, occupancy levels, and facility configurations is essential for the system to deliver sustained operational value.

Private

Autonomous robots continuously recording video as they patrol occupied facilities capture detailed data about individuals' movements and behaviors throughout the working day. Organizations must define clear limits on what is recorded, how long footage is retained, who can access it, and what prevents patrol data from being used for employee monitoring beyond legitimate security purposes.

Fair and impartial

Threat classification models trained on specific facility types or populations may generate disproportionate false positives for individuals whose appearance or behavior differs from training data norms. In government facilities serving the public, biased threat detection creates civil rights exposure and undermines the institutional trust on which legitimate security operations depend.

POTENTIAL BENEFITS

Improved coverage

Facilities are monitored continuously as robots provide round-the-clock patrols without breaks, shift changes, or attention lapses that create security gaps.

Faster detection

Threats are identified earlier through constant mobile monitoring that catches developing situations before they escalate, rather than discovering incidents only during periodic guard rounds.

Security consistency

Coverage does not degrade over time as robots maintain identical patrol patterns and alertness levels regardless of hour, day, or duration of deployment.

Lower labor dependence

Fewer routine guard tasks as robots handle repetitive patrols, allowing security personnel to focus on response, investigation, and situations requiring human judgment.

Audit-ready physical decision trails

Sensor data, movement paths, and escalation decisions can be logged for post-incident review, public accountability, and legal scrutiny.

Smart city operations and urban infrastructure modernization

City-scale intelligence for monitoring, service delivery, and resilient operations

Physical AI combines urban sensors, drones, and autonomous robots to monitor and manage city infrastructure in real time. Edge intelligence detects events, sending data to cloud platforms for analysis and decision-making, while robots perform tasks like inspection and maintenance, linking sensing and action across the city.

ISSUE/OPPORTUNITY

Smart city initiatives aim to deliver dense, sustainable urban environments with high service quality and operational efficiency. In practice, however, both city infrastructure and large public complexes still rely heavily on fragmented data, manual inspections, and labor intensive operations. Urban systems—transportation, utilities, public safety, and environment—operate in silos with limited real time, street level visibility, making it difficult to detect issues early, coordinate responses, or prioritize maintenance. At the same time, large residential and commercial developments depend on sizable human workforces for cleaning, maintenance, and deliveries, resulting in high operating costs, uneven service quality, and limited scalability as demand grows.

The opportunity is to deploy physical AI as foundational urban infrastructure. Edge integrated sensors, drones, and autonomous mobile robots enable continuous, real time monitoring and verification across cities and facilities. By detecting hazards at the edge, streaming curated events to the cloud, and unifying insights across systems, physical AI enables coordinated dispatch, proactive intervention, and demand driven resource allocation—delivering more consistent services, improved resident experience, and more resilient urban operations without relying on static schedules or manual oversight.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

City-scale sensor and edge integration

AI aggregates data from transportation, utilities, and environmental sensors with edge compute to detect anomalies locally (e.g., hazards, leaks, congestion) and publish event alerts with minimal latency.

Robotic service execution

AMRs handle cleaning, waste handling, internal logistics, and routine maintenance across public buildings, transit hubs, campuses, and residential complexes. Service schedules adapt dynamically based on actual usage patterns.

Infrastructure health monitoring

Detects stress and failure in urban systems by analyzing diverse data such as bridge vibration patterns, water flow irregularities, or grid fluctuations, identifying maintenance needs before failures disrupt service.

Decision support for operators

AI informs planners without exerting autonomous control, providing recommendations that human operators review and approve before implementation.

Human governance and oversight

Public authorities retain control over policy decisions and operational changes, with AI systems providing analysis and recommendations.

Mobility and crowd analysis

Improves traffic and pedestrian flow by identifying congestion patterns, predicting rush-hour bottlenecks, monitoring crowd density at public events, or detecting unusual movement patterns.

Smart city operations and urban infrastructure modernization

MANAGING RISK AND PROMOTING TRUST

Private

City-scale sensors, drones, and robots operating continuously across public spaces and residential areas create a mass surveillance capability covering entire urban populations that have no practical ability to opt out. Governments must define strict, publicly documented limits on what is collected, how long it is retained, and what prevents repurposing beyond the infrastructure management purposes that justified deployment.

Responsible and accountable

The use case preserves human governance by design; however, it is easy for AI recommendations to become de facto decisions as alert volumes grow and operator capacity is stretched. AI recommendations must be clearly distinguished from human decisions in operational records, and accountability frameworks must be actively maintained, not assumed to hold just because that was the original design intent.

Fair and impartial

Governments have fairness obligations that private sector entities do not. AI-driven urban resource allocation that systematically directs maintenance and services toward higher-density or more commercially active areas is not just an operational shortfall—it is a failure of the public duty that justifies the deployment. Monitoring for equitable outcomes across the full jurisdiction must be built into governance from the outset.

POTENTIAL BENEFITS

Reduced operating costs

Replaces large human teams performing repetitive tasks with autonomous systems, significantly lowering facility management expenses while enabling smart city developments to achieve financial sustainability.

24/7 facility operations

Enables round-the-clock service delivery—cleaning, deliveries, maintenance—without shift-based labor costs or coordination complexity, supporting the always-on nature of smart city infrastructure.

Greater urban resilience

Cities respond more effectively to stress and incidents, as cross-system visibility enables faster detection of cascading failures, coordinated response across departments, and proactive mitigation of predictable disruptions.

Consistent service quality

Eliminates variability in service delivery based on staffing availability, ensuring all residents receive uniform maintenance and support as envisioned in smart city livability standards.

Better-informed planning

Decision-makers gain clearer visibility into city operations, enabling evidence-based infrastructure investments, more accurate service demand forecasts, and better understanding of how proposed changes might affect interconnected systems.

The Life Sciences & Health Care AI Dossier



The Life Sciences & Health Care AI Dossier

AI is accelerating transformation across life sciences and health care—sectors where precision, timeliness, and trust can mean the difference between life and death. From research labs to hospital wards, AI innovations (including LLMs, AI agents, and physical AI) are enabling faster discovery, more targeted treatments, and more efficient and effective delivery of care.

In life sciences, AI is streamlining drug discovery and development—helping to identify promising new molecules, optimizing clinical trial design, and identifying patient cohorts for precision therapies. AI is also being used to automate regulatory compliance and to optimize supply chain performance through improved supply and demand forecasting.

In health care, AI is improving diagnostic accuracy, personalizing treatment pathways, and easing administrative burdens. Increasingly, agentic AI and physical AI are automating complex workflows—supporting clinical decision-making, coding and billing, and even orchestrating and executing elements of care delivery with minimal human intervention.

With growing volumes of clinical, genomic, and real-world data, AI is helping organizations extract insights at a speed and scale that traditional methods cannot match. And as the technology matures, AI will not just enhance isolated functions—it will likely become integral to advancing scientific discovery, improving population health, and building more resilient, responsive health systems with ongoing human oversight.

That said, the use of AI in regulated, high-stakes environments demands rigorous validation, explainability, and alignment with ethical standards. Data privacy, equity in outcomes, and trust in automation are paramount. To succeed, companies will need to effectively integrate AI within existing care and research frameworks while navigating evolving regulatory requirements and public expectations.

AI is helping organizations extract insights at a speed and scale that traditional methods cannot match.

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

Multi-modal diagnosis and clinical decision support

Collaborating across data types to improve diagnostic accuracy

Agentic AI systems can act as a clinical co-pilot, with specialized agents analyzing different data sources—such as imaging, lab results, and patient records—and then combining insights to support more accurate and timely diagnoses.

ISSUE/OPPORTUNITY

Diagnosing complex conditions often requires synthesizing data from multiple modalities: radiology scans, lab tests, patient histories, genetic profiles, and physician notes. This process traditionally depends on coordination across multiple human specialists, each interpreting one slice of the evidence. Limited time, high patient volumes, and fragmented systems make it difficult to integrate all relevant data into a cohesive picture, which can delay treatment or lead to missed diagnoses.

As health systems face mounting pressure from clinician shortages and rising patient loads, the need for decision support that is both comprehensive and efficient is urgent. Multi-agent AI provides an opportunity to automate routine analysis across diverse data sources, highlight patterns that humans might overlook, and generate integrated recommendations that help physicians make faster and more informed decisions.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Specialized data analysis
Specialized AI agents can apply deep expertise in focused areas. For example, imaging agents can flag abnormalities on scans, NLP agents can parse physician notes and pathology reports, and other agents can interpret lab results and genetic markers.

Continuous learning from records
Agents can draw on the electronic health record as common context and shared memory, adapting recommendations over time based on a patient's history and treatment responses (in addition to population-level outcomes).

Contextual integration
A planning or coordination agent can link evidence across modalities—for example, connecting a lung nodule on a CT scan with smoking history and lab findings—before proposing next steps such as a biopsy or targeted treatment.

Workflow automation
By triaging scans, pre-annotating records, and uncovering key insights, agents can reduce the time physicians spend on routine analysis, allowing them to focus on patient interactions and issues requiring complex human judgment.

Multi-modal diagnosis and clinical decision support

MANAGING RISK AND PROMOTING TRUST

Safe and secure
Errors in diagnosis can be life-threatening. As such, agents should be deployed with strong safeguards, including rigorous validation of outputs and controlled simulation testing before integration into live workflows.

Transparent and explainable
Because clinical decisions carry major consequences, agents should provide transparent reasoning, including references to underlying data. Validator agents can be used to audit and cross-check other agents' conclusions.

Private
Medical data is highly sensitive, which means that agent systems must follow strict access controls, comply with HIPAA, GDPR, and similar regulations, and use privacy-preserving techniques to minimize unnecessary data exposure.

Responsible and accountable
Since liability in clinical settings is complex and risky, agent recommendations should initially be deployed under physician supervision, with final accountability for diagnosis and treatment decisions retained by human clinicians.

POTENTIAL BENEFITS

Faster, more accurate diagnoses
Automated integration of multimodal data helps generate insights earlier and reduces diagnostic errors, enabling quicker, more effective interventions.

Reduced clinician workload
By offloading burdensome, data-heavy tasks such as image triage and report synthesis, agents give physicians more time for patient care and complex decision-making.

Improved patient outcomes
Earlier detection and personalized treatment plans can lead to better survival rates, reduced complications, and higher overall quality of care.

Hyper-personalized health care

Delivering tailored, continuous care through multi-agent collaboration

Agentic AI systems can act as a 24/7 virtual care team, with specialized agents monitoring patient data, providing personalized coaching, and coordinating interventions to deliver hyper-personalized health care at scale.

ISSUE/OPPORTUNITY

Managing chronic conditions such as diabetes, hypertension, or heart disease requires continuous monitoring, timely adjustments, and patient engagement outside of clinical visits. Human care managers and clinicians often have limited bandwidth to track individual patients in real time, leaving care gaps where small issues can escalate into serious and costly hospitalizations or complications. Patients, meanwhile, often struggle to follow generalized care plans that don't adapt to their lifestyles, personal preferences, or daily health fluctuations. The result is higher readmission rates, preventable emergency visits, and reduced patient satisfaction and health outcomes. Multi-agent AI offers a way to make health care more proactive and personalized, addressing issues early and potentially scaling up to provide individualized support to large populations

Tags | Customer Experience | Agentic AI

HOW AI CAN HELP

Continuous patient monitoring
A monitoring agent can analyze real-time data from wearable devices, glucometers, or other sensors, detecting subtle patterns such as recurring glucose spikes or irregular heart rhythms.

Personalized coaching
A coaching agent can engage patients directly through apps or chatbots, offering tailored advice, timely reminders, and personal encouragement that reflects their specific health needs and behaviors.

Evidence-based insights
A knowledge agent can stay updated with medical guidelines and population-level data, interpreting broader patient patterns and suggesting personalized, evidence-based adjustments to care.

Automated, coordinated care
Agents can schedule follow-ups, suggest medication adjustments (with clinician approval), offer healthy recipes, and even auto-order groceries—all orchestrated through a shared patient profile.

Hyper-personalized health care

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	Biased training data could lead to less effective guidance for certain demographic groups. To minimize problems and avoid reinforcing existing disparities, AI agents should be trained and validated on diverse patient datasets.
Private	Because agentic AI systems continuously collect and analyze sensitive patient information, they must provide strict privacy protections and fully comply with data regulations such as HIPAA and GDPR. This includes obtaining clear patient consent.
Transparent and explainable	Patients and clinicians could initially be skeptical of following AI-driven guidance. To help build trust, AI agents should provide understandable explanations for their recommendations, including specific references to the data or guidelines that informed them.
Responsible and accountable	Clinical decisions can have major safety implications. As such, AI agent outputs should be reviewed by care teams where appropriate, with final accountability for treatments and adjustments retained by licensed clinicians.

POTENTIAL BENEFITS

Improved patient outcomes at lower cost

Timely, personalized interventions can help improve patient adherence to treatments, reduce complications, prevent costly emergency visits, and enable earlier corrections when managing chronic diseases. More broadly, agentic AI systems can boost operational efficiency and clinician productivity.

Scalable personalized care

AI agents can enable clinicians to provide individualized care to thousands of patients simultaneously, increasing reach and helping to address workforce shortages.

Higher patient engagement and satisfaction

AI can help patients feel supported between clinical visits, improving adherence to treatment plans and fostering trust in their care providers.

Automated customer service

Improving health plan member interactions through multi-agent systems

Agentic AI systems can handle a wide variety of customer service tasks—such as benefit inquiries, claims status updates, and enrollment support—seamlessly escalating complex cases to human representatives.

ISSUE/OPPORTUNITY

Health plans face large volumes of inquiries from members and prospective members across multiple channels, ranging from routine coverage questions to complex claims issues. Traditional call centers and support teams often struggle with long wait times, inconsistent service quality, and high staffing costs, all of which can frustrate members and erode customer satisfaction.

In today's competitive marketplace, a poor service experience can increase churn, while inefficient processes can drive up operating costs through the roof. Health plans need a way to deliver fast, accurate, and personalized support without continuously scaling human staffing levels, particularly during high-volume periods such as open enrollment.

Tags: **Customer Service** | **Agentic AI**

HOW AI CAN HELP

Specialized AI agents
Individual agents can handle specific tasks such as answering benefit questions, processing claims inquiries, and assisting with new member onboarding, providing responses that are timely, accurate, and consistent.

Seamless collaboration
AI agents can share context across tasks, improving their service quality and allowing for smooth escalation to human representatives when needed, without requiring members to repeat their information.

Dynamic prioritization
Requests can be triaged based on urgency, member status, or complexity, ensuring that high-priority cases receive timely attention.

Proactive engagement
AI agents can initiate reminders or follow-ups, such as prompting prospective members with enrollment information or reminding existing members of required documentation.

Automated customer service

MANAGING RISK AND PROMOTING TRUST

Private
Health plan support requires handling sensitive personal and medical information. Agentic AI systems must be designed with strict data protection protocols and full compliance with regulations like HIPAA.

Robust and reliable
Since members rely on accurate information to make critical financial and health decisions, AI-generated responses should be validated against authoritative data sources and continuously monitored for accuracy.

Responsible and accountable
Health-related interactions often require a strong dose of empathy and humanity. AI agents should be implemented as a complement to human staff, with clear escalation paths to ensure that members always have access to human support when needed.

POTENTIAL BENEFITS

Improved member experience
Faster, more accurate, and more consistent responses across channels help improve the overall customer experience, build member trust and satisfaction—and reduce churn.

Increased operational efficiency
Automating routine tasks frees up human staff to focus on complex or high-priority cases, boosting effectiveness while at the same time improving productivity and efficiency.

Enhanced scalability
Agentic AI can help health plans maintain consistent levels of service and responsiveness in times of fluctuating inquiry volumes (e.g., during open enrollment) without proportionally increasing staffing.

Smarter clinical trials

Improving trial design, recruitment, and monitoring with multi-agent systems

Agentic AI systems can optimize clinical trials by designing protocols, selecting sites, identifying participants, and monitoring progress in real time, helping trials run faster, safer, and more efficiently.

ISSUE/OPPORTUNITY

Clinical trials are one of the most expensive and resource-intensive stages of drug and therapy development, often slowed by challenges such as inefficient site selection, poor patient recruitment from narrow patient pools that don't fully reflect the broader population, and frequent protocol amendments. These inefficiencies can lead to delays, higher costs, and a higher risk of trial failure—ultimately impacting the speed-to-market and efficacy of new treatments.

In the face of growing therapeutic pipelines and increasing pressure to reduce time-to-market, life sciences organizations need tools that can improve trial design and execution while maintaining compliance with strict regulatory standards.

Tags | R&D/Product Development | Agentic AI

HOW AI CAN HELP

Optimized site and protocol design

AI agents can analyze hospital capabilities, historical trial data, and patient access to recommend optimal trial sites, appropriate sample sizes, and robust protocols that balance scientific rigor with operational feasibility.

Data integration and harmonization

By combining data from electronic health records, registries, and wearable devices, agents can create a holistic, real-time view of trial operations, ensuring decisions are made with the most current and complete information.

Adaptive trial management

Agents can continuously monitor trial progress and outcomes, dynamically adjusting resource allocation or protocol details to keep trials on track and compliant.

Patient stratification and recruitment

Specialized agents can identify and stratify eligible patient populations, improving recruitment efficiency while reducing delays caused by insufficient enrollment.

Smarter clinical trials

MANAGING RISK AND PROMOTING TRUST

Fair and impartial	AI models should be trained and validated using demographically broad datasets to avoid systemic bias in recruitment strategies. Representation audits can help ensure no population segments are inadvertently favored or excluded.
Private	Since clinical trials depend on sensitive patient data, agents must be designed with strong privacy protections, data governance controls, and full compliance with regulations such as HIPAA and GDPR.
Robust and reliable	Trial outcomes directly impact patient safety and therapeutic approval. AI-generated recommendations need to be validated against clinical expertise and continuously refined through real-world feedback.
Responsible and accountable	Over-reliance on automation can lead humans to become complacent and miss important clinical nuances. AI agents should be treated as decision-support tools with physicians and trial managers maintaining final accountability.

POTENTIAL BENEFITS

Lower trial costs and improved trial success rates

Improved trial design and adaptive management reduce delays and inefficiencies, driving down the high costs associated with clinical development. Also, real-time recruitment monitoring increases the likelihood of trial completion.

More accurate trial results

By including a broader and more demographically representative participant pool, clinical trials are more likely to produce results that reflect the real-world effectiveness of a drug across different populations.

Accelerated time-to-market

Smarter site selection, faster recruitment, and fewer protocol amendments shorten trial timelines, enabling faster delivery of new therapies to the marketplace.

Better commercial outcomes

Drugs that are tested and proven effective across a wide range of demographic groups are more likely to gain broad clinical adoption, which increases their commercial potential.

Dynamic inventory management

Using coordinated AI agents to optimize supply chains and inventory

Agentic AI systems can transform supply chain and inventory management through specialized agents that help optimize inventory levels and supply chain performance while reducing costs.

ISSUE/OPPORTUNITY

Managing inventory in life sciences and pharmaceuticals is particularly challenging due to complex supply chains, strict regulatory requirements, and unpredictable demand for critical products. Traditional inventory management often relies on fixed rules and periodic actions, which can lead to costly overstocks, stockouts that disrupt patient care, and inefficiencies in production and procurement.

In addition to regulatory and operational complexity, companies must also navigate factors such as geographic variations in demand, perishability of products, and external disruptions like transport delays or weather events. These challenges can result in shortages with serious health implications, or excess inventory that drives up costs and waste.To remain competitive and compliant, organizations need inventory management systems that can respond dynamically to real-time changes in demand, supplier performance, and production capacity—minimizing waste and ensuring that essential therapies and products are always available.

Tags | Procurement/Sourcing & Supply Chain | Agentic AI

HOW AI CAN HELP

Demand forecasting agents
AI agents can analyze historical sales data, market trends, and external signals to help predict demand fluctuations and proactively adjust inventory levels and plans. AI can also integrate local geographical characteristics, disease prevalence, and socioeconomic factors to generate highly accurate, micro-market-specific demand forecasts.

Supplier and procurement agents
Agents can monitor supplier performance, lead times, and procurement activities, ensuring purchasing decisions align with both short-term needs and longer-term forecasts.

Production and scheduling agents
AI can optimize production schedules, identifying bottlenecks and aligning output with available resources and expected demand. By leveraging cross-silo data and advanced analytics, AI can identify patterns and trends that traditional methods might miss, helping to mitigate product shortages.

Distribution and logistics agents
Agents can monitor distribution flows, detect inefficiencies, and propose adjustments in real time to help ensure inventory is positioned where it is most needed.

Coordinated orchestration
An orchestration layer confirms that actions across procurement, production, and logistics remain aligned, enabling holistic, end-to-end inventory optimization.

Dynamic inventory management

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Poor forecasts or system errors can trigger significant supply chain disruptions. AI agents should be validated against human planners and real-world performance, with continuous monitoring to ensure accuracy and stability.

Transparent and explainable

Inventory decisions can have a major impact on costs, compliance, and patient access. As such, agent-driven decisions should be supported by clear rationales that supply chain professionals can understand and verify.

Responsible and accountable

Automated systems can overlook subtle operational factors that a human expert might consider crucial. Conversely, over-reliance on automation can make human experts complacent. Human oversight of AI is essential, with supply chain professionals empowered to override or adjust agent recommendations as needed.

POTENTIAL BENEFITS

Reduced costs and improved financial performance

Dynamic, data-driven inventory planning powered by agentic AI can lower inventory holding costs, reduce excess stock, and minimize spoilage of perishable products.

Improved service levels

Real-time monitoring and coordinated decision-making reduces the risk of stockouts, ensuring patients and providers receive products when needed.

Greater supply chain agility

Organizations can respond more quickly to disruptions or demand spikes, maintaining resilience in the face of changing market or regulatory conditions.

End-to-end autonomous drug discovery and development

Accelerating drug discovery through collaborative AI agents

Agentic AI systems can drive the drug discovery and development process—from molecular design to clinical trials—using specialized agents that work together to accelerate innovation and reduce costs.

ISSUE/OPPORTUNITY

Drug discovery and development is one of the most expensive and time-consuming challenges in life sciences, often costing billions of dollars and spanning many years or even decades. Researchers must manually generate candidate molecules, predict their properties and efficacy, assess toxicity, and evaluate manufacturability (often across disconnected systems and siloed teams). This fragmented approach slows innovation, drives up costs, and increases the risk of late-stage failures.

With growing demand for faster therapeutic innovation and mounting R&D pressures, life sciences organizations need ways to both increase the throughput of candidate discovery and improve confidence in early-stage selections. Multi-agent AI offers an opportunity to streamline molecule design, integrate diverse data sources, and automate iterative refinement of promising compounds.

Tags: R&D/Product Development | Agentic AI

HOW AI CAN HELP

Agent-driven design of novel compounds
Specialized AI agents can generate novel chemical structures, simulate compound interactions, and predict pharmacological properties, providing a broad and high-quality candidate pool. AI can also propose modifications to known molecules in early-stage drug development, rapidly evaluating their therapeutic potential and feasibility.

Multi-modal integration
AI systems can incorporate data from different domains—including chemical, biological, and safety data—creating a holistic view of each compound's potential and reducing the risk of pursuing weak candidates.

Collaborative refinement
Toxicity, manufacturability, and regulatory-focused agents can evaluate compounds in parallel, sharing data with molecule generation and simulation agents to iteratively improve designs. An orchestration agent can oversee and manage the specialized agents, integrate outputs, and prioritize workflows, facilitating seamless progression from initial ideation to candidate selection.

Clinical trials and regulatory submissions
Agents can design preclinical studies and trial protocols based on predicted outcomes. Also, agents can prepare submission-ready documentation, track compliance requirements across regions, and adapt protocols to evolving regulatory standards.

End-to-end autonomous drug discovery and development

MANAGING RISK AND PROMOTING TRUST

Robust and reliable
Early design errors can lead to costly downstream failures. AI outputs should be validated against human benchmarks and continuously monitored.

Transparent and explainable
Regulatory approval requires meeting strict standards for safety, efficacy, and traceability. Agent reasoning and decision paths should support explainable outputs, including transparency into the data and assumptions used.

Responsible and accountable
Drug discovery is a complex challenge with life-or-death implications. AI agents should be deployed as co-pilots to human scientists, with humans retaining accountability for final go/no-go decisions.

POTENTIAL BENEFITS

Accelerated time-to-market
Automated generation and evaluation of compounds can significantly shorten early-stage development timelines—from years to months—enabling rapid responses to emerging diseases and other unmet needs.

Lower R&D costs
By rapidly identifying and prioritizing candidates that meet key property constraints, AI can reduce development costs and improve the quality of compounds advancing to later stages.

Higher likelihood of success
Comprehensive, multi-modal analysis and real-time adaptation improve the chances that selected candidates will advance successfully through clinical trials. The results could include reduce late-stage attrition, improved ROI, and timely delivery of safer, more effective therapies.

A co-writer for appeals

Denial appeal letters

AI can be used to draft denial appeal letters, drawing from patient records and medical policies and guidelines in a faster, more cost-effective way.

ISSUE/OPPORTUNITY

When a medical insurance claim is denied, hospital billing staff face a costly and lengthy process of reviewing patient records and medical policies to create an appeal letter. For US hospitals, appeals-related administrative costs are measured in billions of dollars. Part of the challenge is the amount of time required for staff to compile an appeal. While more than 60% of denied claims are recoverable, vague reasons for denial and limited hospital billing resources result in only 0.2% of in-network claims being appealed, with millions of dollars written off as uncollectible loss each year.[3]

Tags: **Operations**

HOW AI CAN HELP

Retrieving policies and guidelines
An AI-based retrieval model can reach across large volumes of medical policies and member plans to identify the necessary information for a claims appeal.

Extracting patient data
Using extractive algorithms, the organization can rapidly consult unstructured medical notes, medications, lab results, and other electronic health records.

Writing the appeal
With the necessary information gathered by AI, an LLM can be used to generate an appeal letter.

A co-writer for appeals

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable

When consulting highly detailed guidelines, policies, and records to appeal a claim denied for vague reasons, AI models may misinterpret the denial or the records, leading to an unsuccessful appeal. Ultimately, a human needs to be accountable for validating appeal letters.

Private

By drawing from electronic health records, the model is consuming health information whose protection is subject to laws and regulations. Organizations must ensure that the data ingested and information outputted aligns with data protection and patient privacy expectations.

POTENTIAL BENEFITS

Reclaim revenue

Automating the denial appeal process can supplement hospital billing resources, leading to more denial appeals filed and potentially more revenue recovered.

Efficiency improvement

Compared to traditional manual methods, AI technologies can enhance the speed and efficiency of appeals activities, such as substantiating claims and drafting appeals. Their potential to streamline processes for both simple and complex cases can make the legal workflow more timely and cost-effective.

Faster admin for payers, providers, and patients

Accelerated prior authorization

Using AI to analyze medical policies, guidelines, and provider-submitted information about underlying issues, patient needs, and medical history, an organization can automate a Prior Authorization submission (Provider) or generate a Prior Authorization approval or denial (Payer).

ISSUE/OPPORTUNITY

The Prior Authorization process is manual and labor-intensive for both health care payers and providers. The process requires the input of coders who understand the intent of a payer's Prior Authorization policies, as well as the need for medically necessary care management plans. The time required to analyze medical records and policies to make determinations on Prior Authorization submission, approval, or denial can lead to a long administrative process between the payer and provider, which can negatively impact patient satisfaction and the customer experience.

Tags: Operations

HOW AI CAN HELP

Supporting the provider
For providers, AI can help prepare a Prior Authorization submission by analyzing submission requirements and guidelines and cross-referencing with a patient's medical records to ensure necessary requirements are met. AI can then aid in submission to the payer and continually learn which best practices tend to lead to Prior Authorization approvals.

Supporting the payer
For payers, AI can help reduce the time required to make a Prior Authorization decision, impacting the patient experience. It also helps mitigate fraud by determining if there are anomalies in a provider's coding practices and supports compliance by analyzing submitted Prior Authorization requests and records against the payer's policies and procedures.

More efficient operations
For both payers and providers, using AI can reduce work burdens and streamlines the ability to handle Prior Authorizations, which can reduce costs while improving patient experiences.

Faster admin for payers, providers, and patients

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Prior Authorization requires the provider and payer to communicate sensitive patient data, such as protected health information (PHI) and personally identifiable information (PII), etc., which means this data is exposed to the model. Risks include unauthorized third-party access, as well as AI systems inadvertently revealing sensitive information during the generation process, thus compromising patient data confidentiality.

Fair and impartial

The process for submitting and responding to a Prior Authorization involves a standard set of rules and the patient's medical history, which introduces the potential for bias in AI models. This bias might arise from the historical data used to train the model (e.g., disparities in health care treatment or outcomes), which could lead the AI model to inadvertently perpetuate and even amplify such biases by making its own biased decisions or recommendations. The use of standardized authorization rules and patient-specific medical history, alongside continuous monitoring and careful evaluation, helps mitigate this risk and promotes fairer outcomes.

Robust and reliable

While the process for submitting and responding to a Prior Authorization revolves around a standard set of rules and the patient's medical history, there is a risk the model will misinterpret nuanced medical conditions of underrepresented populations that were not in the training dataset, and thus falsely deny the need.

POTENTIAL BENEFITS

Speed and efficiency

With AI, providers and payers may require less time to understand policies, research patient medical records for compliance, and generate, approve, or deny a Prior Authorization request.

Continuous learning

An AI feedback loop refers to the cyclical process whereby the AI model's output is presented to users or evaluators for feedback, which is then used to iteratively update and refine the model. This enhances the consistency and quality of outputs, enables providers to gain a deeper understanding of payer policies, streamlines decision-making processes, and helps payers optimize their procedures.

Improved patient experience

As the Prior Authorization process becomes more efficient, patients can receive the care management they need without needless waiting for administrative processes to conclude. This can increase patient satisfaction by improving the administrative and patient experience.

Simplifying claims submission

Medical coding

AI can be used to create codes for a claims department to categorize incoming claims and billing for medical services and procedures, which can improve the accuracy, efficiency, and speed of the claims process.

ISSUE/OPPORTUNITY

The claims submission process in the medical industry can be laborious and error-prone, involving the manual categorization of a large volume of incoming claims with complex medical codes. This time-consuming task leads to backlogs, delays, and potential payment issues for health care providers.

Tags: **Operations**

HOW AI CAN HELP

Transformed claims processing
Using AI to help categorize incoming claims and analyze and assign accurate codes can improve the overall accuracy, efficiency, and speed of claims processing. This can result in faster reimbursements for providers and a streamlined experience for both the claims department and patients.

Reduced labor burden
By leveraging AI, human effort in the claims submission process can be redirected to higher value tasks, which could result in administrative cost savings for the payer.

Simplifying claims submission

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

A Large Language Model (LLM) used in medical billing may be susceptible to bias from skewed training data, incorrect labels, and under-represented cases, potentially leading to incorrect claim categorization. To mitigate these issues, careful data collection, diverse model testing, and continuous monitoring and adjustment are vital for ensuring fair and accurate performance.

Private

To assess coding accuracy, the LLM compares the billed codes with the patient medical history, which exposes the patient's data to the underlying model and creates potential privacy risks that need to be mitigated.

Robust and reliable

Medical coding is highly regulated with strict penalties for over/under coding. The accuracy and reliability of LLM outputs in this regard is essential, as mistakes could have serious consequences. Reliability may be challenging in part because patient medical history may contain multiple modalities (e.g., text, images, and video).

POTENTIAL BENEFITS

Accuracy to limit revenue loss

Leveraging an LLM can help reduce the risk of coding errors. This can increase billing accuracy and decrease revenue loss due to errors.

Time efficiency

Automating the review of medical records can save valuable time for health care practitioners, enabling them to focus on more meaningful work.

A physician's message manager

Provider in-basket management

An LLM can be used to process messages in a health care provider's in-basket, accelerating responses while enabling physicians to focus on patient-facing care.

ISSUE/OPPORTUNITY

The amount of time required for primary care providers (PCPs) to accomplish both administrative and patient care responsibilities can exceed what is possible in a day. In some cases, upwards of two-thirds of their time is spent on administrative, non-patient facing work.

The 21st Century Cures Act encourages electronic medical records (EMR) in-basket usage, which led to a dramatic increase in in-basket messages during the COVID-19 pandemic. The result is a significant burden on PCPs, which is contributing to physician burnout.

Tags: **Customer Service**

HOW AI CAN HELP

Triaging the in-basket
AI can be used to review routine messages (e.g., Rx refills, scheduling) and delegate simpler tasks to automation.

Message assistant
PCPs can leverage AI to summarize complex clinical messages for review and use the model to draft replies for provider input and response. AI models consult prior in-basket replies and EHR data when creating drafts.

Insights at scale
By using AI-enabled in-basket message systems at scale, organizations can identify issues related to patient negativity in their messages. The insights into complaints, expressions of dissatisfaction, frustration, confusion, or concern about care can inform interventions that may improve the patient experience.

A physician's message manager

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Use of AI for in-basket systems involves collecting, processing, and storing large amounts of sensitive patient data, such as medical history, diagnoses, and treatment plans. This data is subject to strict privacy laws, and any unauthorized third-party access could result in legal and financial consequences for health care providers.

Responsible and accountable

If messages are composed or summarized with inaccurate information, it could lead the PCP to erroneous decision-making or poor patient engagement, which can have significant consequences for patient health, trust in the health care provider, and the reputation of the organization.

POTENTIAL BENEFITS

Physician support

By using an AI-enabled in-basket system, the PCP's time-consuming administrative tasks are reduced, permitting more patient-facing work and mitigating one cause of physician burnout.

Timely responses

A more efficient process for working through in-basket messages can lead to faster responses to patient needs, contributing not only to a better patient experience but potentially also better health outcomes.

Patient sentiment

By identifying and tracking signals of negativity at scale, health care providers can gain insights into common pain points in the patient experience. This could help them proactively address these issues, whether by adjusting their practices, improving communication, or implementing other interventions to enhance patient satisfaction.

Democratizing AI model creation

Knowledge domain model development

Generative AI can be used to improve existing AI models by removing user interface (UI) hurdles through reinforcement learning (RL) without the need for technical staff.

ISSUE/OPPORTUNITY

Developing novel AI models for life sciences and health care organizations continues to demand a high degree of technical proficiency to perform data exploration, feature engineering, model training, and evaluation. Frequently, the steps involved in model training lack a user-friendly interface, posing accessibility challenges for health care professionals and domain experts who may not possess extensive technical backgrounds. Simultaneously, the quality and relevance of model outputs hinges significantly on domain expertise and practical experience. Overcoming this divide between technical acumen and domain knowledge remains an obstacle in harnessing the complete capabilities of AI within the field of life sciences and health care.

Tags | R&D/Product Development

HOW AI CAN HELP

Empowering professionals

With its capacity for learning from and adapting to iterative feedback, generative AI can act as an enabler for professionals across various sectors. It offers the opportunity to continually refine existing domain-specific AI models by adding new training data. This iterative enhancement increases the model's accuracy, utility, and relevance to the user's specific professional needs. In this way, generative AI can empower professionals by providing them with tailored, precision AI tools that evolve with their work.

Streamlining health care model development

Generative AI can help simplify model development in the complex and highly regulated health care industry. By focusing on intuitive user interface designs and automated processes, generative AI minimizes UI obstacles, making it more accessible for professionals to refine and improve their existing models.

Improving alignment

Generative AI leverages reinforcement learning (RL) techniques, a type of machine learning where an AI system learns to make decisions by trial and error, to validate and improve its own outputs. This process assists in mitigating prevalent AI challenges, including hallucinations or confabulations, ambiguity, and colloquialism misuse. As a result, it bolsters AI's reliability and furnishes professionals with more precise models and predictions, thus aligning AI capabilities more closely with user requirements.

Democratizing AI model creation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Hallucinations or confabulation could lead to the execution of incorrect procedures or use of suboptimal reagents and equipment, causing inaccurate experiments and inefficient use of resources. Particularly in medical or pharmaceutical labs, inaccurate information could even lead to compliance or regulatory issues.

Transparent and explainable

The generative AI system incorporates tools that offer transparency into data engineering pipelines, including data preparation stages. This inherent transparency facilitates an understanding of the AI model's functioning within the organization, fostering trust in the accuracy and reliability of the AI system's outputs. It is a crucial component of the AI use case, demonstrating the system's accountability and promoting its acceptance across the organization.

POTENTIAL BENEFITS

Enhance institutional knowledge access

AI can help reduce institutional knowledge loss due to employee exits and enable on-demand access to domain-specific knowledge across the organization.

Increase development throughput

Domain area experts can drive more self-sufficient model experimentation and development by utilizing natural language model outputs and synthesizing insights about optimal procedures, reagents, equipment, and techniques into a comprehensive and accessible format.

Cost management

This democratic approach to AI model development empowers employees to take part in model experimentation, reducing costs associated with machine learning operations and technical specialists.

Optimizing lab procedures

Experimental design

AI can be used to create procedural templates and recommendations on best practices (e.g., reagents, equipment, techniques).

ISSUE/OPPORTUNITY

Laboratory personnel, including researchers, technicians, and managers, often face challenges in maintaining up-to-date procedural templates and maintaining the consistent application of best practices, especially as scientific knowledge evolves rapidly. These challenges can lead to inefficiencies, errors, and inconsistency in experiments or analyses. Additionally, without a central source of curated recommendations, time and resources may be wasted sourcing and comparing various reagents, equipment, and techniques. These pain points present an opportunity for AI to streamline and enhance laboratory processes.

Tags | **R&D/Product Development**

HOW AI CAN HELP

Generation of novel processes
Leveraging historical data and scientific principles, an AI model could suggest novel experimental designs, more efficient processes, or alternate uses of reagents and equipment, stimulating innovation in laboratory procedures.

Data analysis and interpretation
AI uses a large language model (LLM) to analyze data from lab protocols, equipment specifications, previous experimental designs, reagent usage, and techniques, providing a holistic understanding of laboratory procedures and principles.

Optimizing lab procedures

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

The integration of multimodal text and images of complex structures and processes in experimental design presents complexity. This can heighten the risk of unworkable, unfeasible, or inefficient designs, as interpreting and accurately representing this diverse and intricate data can be challenging. These challenges could potentially lead to errors in the design and execution of experiments, resulting in failed or less reliable outcomes and unnecessary time and resource expenditure.

Responsible and accountable

In the event of erroneous design recommendations, accountability may be an issue. Determining who bears the responsibility for incorrect designs and their potential consequences is important. The roles of human oversight and system validation need to be clearly defined.

Transparent and explainable

With the application of AI in experimental design, there may be challenges related to explainability. If scientific or academic papers are to be published based on the results, authors need to be able to adequately explain the methodology behind the AI-recommended designs, which can be inherently complex due to the black-box nature of some AI models.

POTENTIAL BENEFITS

Efficiency

LLMs can reduce the time and effort needed for experimental design by streamlining and accelerating data analysis and procedure consolidation, and by providing best practice recommendations.

Lower cost

With less time required for experimental design, organizations can reduce the overall operational costs of experiments while also increasing throughput.

Revealing the rules

Automated regulatory compliance

AI can be used to support compliance by processing large amounts of regulatory documents from multiple geographies.

ISSUE/OPPORTUNITY

Compliance with ever-changing regulations in every geography is a costly, time-consuming process for pharmaceutical companies. Even with significant investment in legal help, regulatory compliance can be hard to achieve.

HOW AI CAN HELP

Text processing
AI can be used to extract regulations for a specific purpose from thousands of pages of regulatory texts, enabling compliance.

Transforming the legal support ecosystem
As AI handles the laborious, detail-oriented process of regulatory text processing, it can also lead to a commensurate decrease in the need for third-party legal and compliance support.

Mitigating financial risk
By employing AI in regulatory compliance, the potential financial risk associated with non-compliance can be reduced.

Tags: **Compliance & Risk**

Revealing the rules

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable	AI models may produce outputs that are hard to interpret, making it difficult to validate them and explain the reasoning to regulatory authorities.
Private	While regulatory authorities may vary, data privacy around personal health information remains a priority, and data that is not anonymized first may leak and become inappropriately disclosed.
Robust and reliable	An AI model trained to extract compliance factors from regulatory documents may be susceptible to outputting information that looks accurate but is a hallucination, making human validation an important element for mitigating risks around reliability.

POTENTIAL BENEFITS

Cost reduction

Using AI to process regulatory documents reduces the need for humans to perform time-consuming tasks, potentially lowering the cost of compliance.

Fuel for growth

When regulatory compliance becomes tractable across geographies because of AI processing capabilities, it helps the organization confidently expand business operations globally.

20/20 impurity detection

AI-driven visual inspection for particulate matter in IV fluids

Computer vision powered by AI can be used to detect particulate contamination in IV bags, reducing product waste and improving patient safety in life sciences manufacturing.

ISSUE/OPPORTUNITY

Pharmaceutical manufacturers, particularly those producing IV fluids and life-saving therapies, face a persistent and costly challenge: detecting particulate matter in sterile products. Despite sterile manufacturing environments, small particles—such as plastic, dust, or other foreign materials—can still enter IV bags, posing serious health risks to patients.

This issue is not new—dating back to the 1940s—and despite ongoing improvements, a scalable, reliable, and cost-efficient solution has remained elusive. Historically, detection has relied on manual inspection, often using contingent labor, leading to inconsistent results, high labor costs, and significant product waste. Also, every incident of contamination risks brand reputation, regulatory scrutiny, and potential product recalls.

Tags | **Manufacturing & Quality**

HOW AI CAN HELP

Automated, real-time inspection
By combining AI vision capabilities with generative AI models trained on synthetic and real-world data, manufacturers can automate the inspection process at scale. High-resolution imaging and computer vision detect anomalies in fluid packaging with greater precision and consistency than the human eye.

Continuous learning
AI enhances the system by learning from historical defect data, adapting to new defect types, and identifying potential causes through pattern recognition across vast datasets. The AI system not only flags potential contamination in real time but also enables traceability—helping identify root causes by analyzing patterns across manufacturing lines, geographies, or specific production lots. This insight enables proactive correction and long-term process improvements.

20/20 impurity detection

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

AI models should undergo rigorous testing across multiple manufacturing lines and environments to help ensure high accuracy and minimal false positives/negatives. Redundancy checks, human-in-the-loop validation, and performance monitoring help ensure reliable operation even under variable lighting or packaging conditions.

Transparent and explainable

AI-based contamination detection can provide clear, trackable results, allowing manufacturers to understand why a product passes or fails inspection. Detailed imaging and reporting help ensure accountability and regulatory compliance while enhancing confidence in quality control.

POTENTIAL BENEFITS

Improved patient safety and confidence

Consistent detection of contaminants before they enter the supply chain increases patient safety, and fewer quality incidents enhance brand trust with hospitals, regulators, and patients.

Operational efficiency and scalability

Replacing manual inspection with AI reduces reliance on contingent labor and speeds up quality control processes. Also, AI-based solutions can be deployed at scale across multiple products, manufacturing lines, and facilities worldwide.

Improved ROI and reduced costs/waste

AI helps drive measurable ROI through improved efficiency, reduced labor costs, and avoidance of costly recalls. Also, fewer discarded bags due to false positives or late-stage detection leads to significant material savings.

Quality and traceability

Root cause analysis helps address upstream issues in the manufacturing line, improving overall process quality.

Transforming pharmaceutical research & development

Physical AI smart labs for drug discovery

Physical AI-powered "smart labs" autonomously execute the entire design-make-test-analyze (DMTA) cycle in drug discovery. These systems design molecules using AI-driven computational models ("in silico"), generate synthesis plans, and directly execute them through integrated robotic lab platforms.

These systems orchestrate execution, physically synthesize molecules, transfer samples via automated handling systems, and run assays using connected lab instruments, analyze results, and initiate the next iteration—with minimal human intervention beyond oversight and exception handling.

ISSUE/OPPORTUNITY

Pharmaceutical discovery is constrained by slow, manual DMTA cycles that take 8–12 weeks per iteration, limiting the number of molecules that can be tested and delaying time-to-market for new therapies. Traditional discovery requires manual setup of experiments, physical sample movement, and fragmented instrument workflows—creating bottlenecks and increasing error risk in repetitive tasks. Capital-constrained biotech companies face existential time pressure but lack the workforce to scale discovery operations. Physical AI and robotics offer an opportunity to automate up to 70–80% of standardized processes, accelerating cycle times and enabling faster kill decisions on failing candidates while progressing winners. Autonomous systems can work continuously, handle greater experimental loads, and improve portfolio NPV by compressing discovery timelines.

Tags: R&D/Product Development | Physical AI

HOW AI CAN HELP

AI-driven molecule design & synthesis

AI algorithms design novel molecules in silico, predict synthetic routes, and generate execution plans including reagent plating, liquid handling sequences, and robotic orchestration.

AI-optimized experimental design

Machine-learning models predict which experiments will yield the most information, reducing unnecessary data generation by up to 50% and maximizing throughput from existing lab capacity.

Autonomous sample transport

AMRs transport samples between synthesis, purification, and testing stations. Automated assay platforms run efficacy and safety screens, upload results directly to AI systems, and close the loop for next-cycle decision-making.

Real-time data analysis & adaptive iteration

AI monitors experimental outcomes in real time, evaluates performance against objectives, identifies failures early, and autonomously designs the next DMTA cycle—compressing iteration time from weeks to days.

Robotic synthesis & purification

Robotic arms and automated synthesis platforms execute AI-generated protocols autonomously, performing liquid handling, reaction setup, purification, and quality control (QC) without human intervention, operating 24/7 when facilities permit.

Transforming pharmaceutical research & development

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Silent error propagation is the most distinctive reliability risk here. Errors in robotic synthesis, sample handling, or AI iteration decisions can compound across multiple cycles before detection, advancing flawed candidates and discarding viable ones. This corrupts the scientific foundation on which regulatory submissions rest. Reliability must be validated across the full automated workflow, not just individual steps.

Responsible and accountable

Regulatory authorities expect documented human accountability for scientific decisions. A system operating with minimal human intervention must have an exceptionally robust audit trail to satisfy that expectation. Organizations must clearly document the boundary between AI-generated recommendations and human-validated judgments at every stage, making full traceability of experimental decisions a deployment prerequisite, not a backward-looking governance addition.

Transparent and explainable

Regulatory accountability requires that qualified scientists can assess and stand behind the AI's scientific reasoning—not just validate its outputs. When AI makes decisions to drop candidates or advances molecules based on predicted properties, the reasoning must be traceable to a standard that satisfies both internal scientific review and external regulatory examination.

POTENTIAL BENEFITS

Radical cycle time compression
Autonomous smart labs reduce DMTA cycle times from 8–12 weeks to under 2 weeks, enabling faster progression of viable drug candidates, earlier termination of failures, and significant improvement in portfolio Net Present Value (NPV).

Scalable innovation infrastructure
Pharma companies can scale discovery operations by adding parallel robotic workstations, creating a path to continuous, factory-style drug discovery without workforce constraints.

Increased discovery capacity
By automating 70–80% of standardized workflows, labs can handle larger portfolios with the same physical infrastructure and workforce, focusing human effort on the 20% of artisanal, non-standardizable tasks.

Capital efficiency for biotech
Small biotech companies can deploy smart labs to replace large discovery teams, reducing headcount requirements while achieving higher throughput—critical for one-product companies racing against cash burn.

Enhanced reliability & reproducibility
Robotic systems eliminate human error in repetitive tasks, ensure consistent execution of protocols, and capture tacit operator knowledge that would otherwise be lost, improving data quality and experimental reproducibility.

Automated drug dispensing systems

AI-powered robotic systems for accurate and scalable medication management

AI-powered robotic pharmacy systems automate medication dispensing using physical AI and robotics. Prescription dispensing robots handle counting, packaging, and labeling, while prescription filling robots accurately fill vials or blister packs. Automated storage and retrieval systems organize inventory and enable rapid access. Combined with AI algorithms and computer vision, these systems ensure accurate verification, dosage calculation, and low-error, continuous pharmacy operations with minimal manual intervention.

ISSUE/OPPORTUNITY

Manual pharmacy workflows rely on repetitive tasks such as counting, filling, and inventory handling, which increase errors and slow down operations. These processes are difficult to scale while maintaining accuracy and consistency, leading to delays and higher workload for pharmacists. Physical AI-driven robotic systems offer an opportunity to automate dispensing, filling, and storage, reducing errors, improving efficiency, and enabling scalable, high-throughput pharmacy operations.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

AI-driven prescription processing

AI validates prescriptions, ensures accurate data verification, and reduces manual review effort to speed up processing.

Computer vision verification

Ensure correct medication identification and labeling while detecting discrepancies to prevent dispensing errors.

Robotic dispensing systems

Automate counting, packaging, and labeling of medications while ensuring precision and consistency in repetitive tasks.

Automated storage & retrieval

Manage inventory and enable fast medication access while optimizing stock organization and reducing retrieval time.

Real-time safety checks

Detect drug interactions and patient risks while supporting informed decision-making for safer dispensing.

Robotic filling systems

Fill vials and blister packs with high precision, reducing variability and improving dosage accuracy.

Automated drug dispensing systems

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Medication dispensing errors cause direct patient harm—and a system that introduces its own error modes in place of human ones does not improve patient safety; it substitutes one risk for another. Reliability must be validated continuously across the full range of medications and dosage forms handled, not just the products well-represented during design and testing.

Safe and secure

A compromised dispensing system that alters dosage instructions or suppresses drug interaction checks is not an IT problem—it is literally a matter of life and death. Security must therefore be treated as a patient safety issue, not an IT governance issue, with protections reflecting the direct clinical consequences of system compromise or manipulation of prescription data and dispensing logic.

Responsible and accountable

Pharmacist accountability for final dispensing decisions is a regulatory and professional requirement that AI assistance cannot displace. The governance design—including which checks are AI-assisted and which require human sign-off—must be explicit, enforced, and documented. This is the dimension that ultimately determines whether the system is deployable in a regulated healthcare environment.

POTENTIAL BENEFITS

Reduced errors & improved safety
Minimize manual errors in dispensing and filling while enhancing patient safety and reducing adverse events.

Higher workforce efficiency
Reduce repetitive workload for pharmacy staff while enabling greater focus on patient care activities.

Scalable 24/7 operations
Support continuous, high-volume dispensing while enabling pharmacies to handle increased demand efficiently.

Faster processing time
Accelerate prescription handling and delivery while reducing patient wait times and improving service levels.

Consistent & accurate operations
Ensure standardized and reliable processes while improving quality control and reproducibility.

Optimized inventory management
Improve stock organization and retrieval speed while reducing shortages and enhancing inventory visibility.

Surgical robotics for microsurgery

AI-assisted precision surgery

High-precision surgical robots/robotic hands support microsurgical procedures under direct human supervision. These systems feature extremely low latency, high accuracy, and tight integration between AI perception and physical actuation during live operations.

ISSUE/OPPORTUNITY

Microsurgery demands precision beyond normal human motor capability, while safety requirements leave no tolerance for error. Manual procedures limit consistency and accessibility. Delicate operations such as nerve repair, vascular reconstruction, or ophthalmic surgery require movements measured in fractions of millimeters, performed while the surgeon's hands naturally experience tremor and fatigue. Surgical outcomes vary based on individual surgeon dexterity, experience level, and physical condition, creating inconsistency in patient results for identical procedures.

Complex microsurgical techniques remain accessible only to specialists at major medical centers, limiting patient access based on geography and surgeon availability.

The opportunity is to use physical AI to enhance surgical precision while keeping clinicians fully in control. Leveraging robotic hands enables direct translation of surgeon intent into controlled end effector motions at the point of contact, reducing tremor at the source and enabling precise force control.

Tags: **Field Services** | **Physical AI**

HOW AI CAN HELP

High-resolution perception

AI interprets visual and sensor data during surgery, enhancing the surgeon's view with magnification, filtering, and highlighting of critical anatomical structures that guide precise intervention.

Latency-constrained execution

Decisions and actions occur in real time with minimal delay between surgeon input and robotic response, maintaining the natural feel of direct tissue interaction critical for surgical judgment.

Precision motion control

Robotic hand executes movements at sub-millimeter accuracy, filtering out hand tremor and scaling surgeon hand movements to finer instrument movements for delicate tissue manipulation.

Safety-critical integration

AI operates within strict procedural constraints including movement boundaries, force limits, and anatomical safety zones that prevent inadvertent damage to surrounding structures.

Human-controlled operation

Surgeons retain authority over all actions, with robotic systems translating their commands into precise physical movements rather than making autonomous surgical decisions.

Surgical robotics for microsurgery

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

A latency spike or motion control failure during nerve repair or vascular reconstruction is not a recoverable error—it's a matter of life and death. Validation must reflect actual surgical conditions, not laboratory benchmarks, and must demonstrate performance across the full range of procedures, tissue types, and edge cases the system might encounter in live use.

Safe and secure

The safety envelope—movement boundaries, force limits, and anatomical safety zones—is the primary mechanism protecting patients during live procedures. Any failure that causes the system to act outside that envelope, whether from software error, sensor malfunction, or external compromise, has direct and potentially irreversible consequences. Fail-safe design is not a feature; it is the fundamental condition for clinical deployment.

Responsible and accountable

Surgeon accountability for patient outcomes is both a clinical and a legal requirement that the system's design must actively support. Logs of all commands, system states, and force feedback throughout each procedure are the evidentiary record on which that accountability rests. Without a complete and reliable audit trail, the governance principle ("surgeon retains authority") applies in theory but cannot be demonstrated in practice.

POTENTIAL BENEFITS

Improved surgical outcomes
The robotic hand enables finer tissue handling and more accurate suture placement, reducing trauma and complication rates and improving healing times for microsurgical cases.

Consistency
End effector level actuation and sensing from the robotic hand deliver uniform precision regardless of surgeon fatigue or tremor, reducing variability in outcomes.

Patient benefit
Recovery outcomes improve through less tissue trauma, reduced scarring, shorter hospital stays, and lower complication rates that translate into better long-term function.

Expanded treatment capability
More procedures become feasible as robotic precision enables operations that exceed manual capabilities, bringing advanced microsurgical techniques to cases previously considered too complex.

Service robots for hospital operations

Autonomous execution of routine logistics

Robots perform repetitive service tasks in hospitals, such as transporting medications, meals, linens, and supplies. These systems navigate human environments and interact with staff and patients. Robots physically move items through corridors, elevators, and clinical areas while making real time route and task decisions.

ISSUE/OPPORTUNITY

Hospitals face labor shortages and high turnover in repetitive service roles. Manual task execution consumes skilled staff time. Nurses and porters spend significant portions of shifts transporting medications, lab samples, linens, and supplies between nursing stations, pharmacies, laboratories, and patient rooms.

The opportunity is to deploy AMRs to automate routine point to point logistics, freeing clinical staff for direct patient care. On robot physical AI enables resilient local decision making (obstacle avoidance, dynamic re routing, queued handoffs) that preserves service continuity during peak demand or partial network outages.

Tags: **Operations** | **Physical AI**

HOW AI CAN HELP

Indoor navigation

Robots move safely in human spaces, navigating hallways, elevators, and doorways while avoiding obstacles like patients, visitors, equipment, and staff moving through dynamic hospital environments.

Human interaction handling

Robots operate around people, yielding right-of-way, communicating arrival through alerts or displays, and responding appropriately when staff or patients need to access items or clear pathways.

Task execution

Deliveries are completed autonomously as robots transport items between designated locations, confirm delivery through digital handoffs, and return for the next assignment without human intervention.

Operational reliability

Tasks are executed consistently with predictable delivery times, proper handling of temperature-sensitive medications or fragile items, and automated documentation of completed deliveries.

Fleet coordination

Multiple robots operate together, sharing elevator access, coordinating routes to avoid congestion, and balancing workload distribution to maintain service levels during peak demand periods.

Service robots for hospital operations

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Hospital service robots transporting medications, lab samples, and temperature-sensitive supplies must perform reliably across a live clinical environment that includes crowded corridors, elevator queues, and potential network outages. Delivery failures or unexpected stops can disrupt care workflows, delay medication administration, and create congestion in spaces where patient safety depends on clear access.

Safe and secure

Autonomous robots in hospitals must respond safely to unpredictable human behavior, including patients with impaired mobility, children, and individuals in distress. Physical collisions in clinical spaces could cause injury or disrupt care. Safety boundaries must be validated across the full diversity of people and conditions encountered in live hospital environments, not just controlled testing environments.

Responsible and accountable

When a robot fails to deliver medication on time, delivers to the wrong location, or causes a care disruption, accountability must be clearly allocated between the AI developer, manufacturer, integrator, and hospital operator. Delivery logs and operational records must be sufficient to support investigations and demonstrate compliance with hospital governance and patient care requirements.

POTENTIAL BENEFITS

Labor relief

Staff focus on higher-value work; for example, nurses can spend more time on patient assessment and care rather than transport and other routine service tasks.

Service consistency

Tasks are completed reliably with uniform timing and quality regardless of shift staffing levels, employee experience, or competing demands on human workers' attention.

Scalability

Operations expand without staffing increases as additional robots handle growing delivery volumes during facility expansions and workload peaks without proportional labor costs.

Operational efficiency

Routine tasks are automated, eliminating delays from staff unavailability, reducing delivery errors from miscommunication, and maintaining consistent service during peak periods.

Eldercare and memory care support systems

Addressing critical care gaps through embodied AI

Humanoid robots deployed in memory care and eldercare settings provide companionship, medication management, and emotional support for individuals with Alzheimer's, dementia, and age-related care needs.

ISSUE/OPPORTUNITY

Aging populations are pushing care needs beyond available caregiving capacity, leaving facilities and families struggling with staffing shortages, medication adherence, and meaningful engagement for patients with cognitive decline. Overburdened caregivers—both professional and familial—cannot provide continuous supervision or companionship, while isolation and medication errors accelerate health risks.

Humanoid AI assistants offer a path forward, they provide consistent, always available support for routine monitoring, medication reminders, and engagement—while keeping caregivers firmly in control of all clinical and judgment intensive decisions. By extending care continuity without replacing human empathy or accountability, Physical AI helps stabilize care quality, reduce caregiver burden, and improve patient well being at scale.

Tags | Customer Experience | Physical AI

HOW AI CAN HELP

Emotional connection & companionship
Humanoid form factor enables elderly individuals to personify and connect with the AI assistant, leveraging the innate human tendency to bond with physical objects. The robot engages in conversations on topics the patient raises, providing stimulation and reducing isolation.

Medication compliance & health monitoring
AI monitors medication schedules and reminds patients to take prescribed medications at appropriate times. System tracks compliance and can alert caregivers to missed doses or concerning patterns.

Governed intervention
Escalation to caregivers when thresholds are crossed. No autonomous clinical decision making.

Consistent, reliable presence
Unlike human caregivers managing multiple patients or working shifts, the AI assistant provides continuous availability. Patients benefit from predictable, patient interaction without fatigue or irritability.

Eldercare and memory care support systems

MANAGING RISK AND PROMOTING TRUST

Private

Robots monitoring elderly patients with cognitive decline collect highly sensitive data—medication adherence, behavioral patterns, and indicators of deterioration—about individuals who may lack capacity to consent. Strict governance is essential, with meaningful consent processes involving family members or legal guardians and clear limits on data access, retention, and secondary use.

Safe and secure

Humanoid robots in close proximity to elderly patients with cognitive decline and impaired mobility must behave predictably under all conditions—including when patients become confused, agitated, or interact physically in unexpected ways. Safety boundaries must be validated specifically for memory care populations, whose responses cannot be assumed to mirror those of other user groups.

Responsible and accountable

Caregivers must retain full clinical accountability. In practice, this requires reliable escalation protocols that alert human caregivers when thresholds are crossed, and governance frameworks ensuring families and care facilities understand precisely what the robot does autonomously versus what triggers human involvement.

Transparent and explainable

Families and care facilities deploying robots with cognitively impaired patients must understand what the system does, what data it collects, and how it escalates concerns—since patients themselves may lack capacity to assess these things. Transparency obligations fall on deploying organizations, not just on individual users, given the vulnerability of the population being served.

POTENTIAL BENEFITS

Addresses eldercare labor shortages
Augments overstretched caregiving workforce by handling routine monitoring, companionship, and medication management tasks. Enables human caregivers to focus on tasks requiring judgment and empathy.

Reduces caregiver burden
Family members and professional caregivers gain peace of mind from continuous patient monitoring. System handles routine reminders and engagement, reducing stress on human caregivers.

Improved medication adherence
Timely reminders and compliance tracking reduce missed doses and escalation risk, with alerts routed to caregivers when needed.

Improves patient outcomes
Consistent monitoring and medication compliance support reduce adverse health events. Emotional engagement and mental stimulation may slow cognitive decline for dementia patients.

Cleanroom manufacturing automation

Precision operations in sterile environments

Robotic systems perform material handling, assembly, inspection, and localized environmental monitoring in pharmaceutical and biologics cleanroom spaces, reducing human presence in sterile zones while maintaining strict contamination control and validated quality standards.

ISSUE/OPPORTUNITY

Pharmaceutical and biologics manufacturing requires sterile production environments where contamination control is critical and quality failures carry massive costs. Cleanroom operations are labor-intensive, requiring workers to don extensive protective equipment, follow rigorous gowning procedures, and work in uncomfortable controlled environments with restricted airflow and temperature regulation. Human presence introduces contamination risk despite protective protocols, as people generate particulates through movement, breathing, and skin shedding—even when fully covered.

Quality failures in sterile manufacturing trigger costly batch rejections that can reach millions of dollars, regulatory investigations that threaten facility operating licenses, and potential patient safety issues that create liability exposure.

Current physical AI systems lack cleanroom certifications and ISO collaborative robotics standards for sterile environments, with updated standards expected in Q3 2028.

These automated cleanroom applications—especially post-certification—can deliver substantial value in high-margin sterile production settings where quality and consistency are paramount and where reducing human presence directly reduces contamination risk while maintaining production capacity.

Tags: **Manufacturing & Quality** | **Physical AI**

HOW AI CAN HELP

Certified cleanroom robotics

Purpose-designed systems that meet cleanroom classification standards perform material handling, assembly, and inspection tasks in sterile zones without introducing particulate or biological contamination, using specialized materials and designs validated for controlled environments.

AI-driven quality monitoring

Advanced vision systems with closed-loop process control detect anomalies in real-time and automatically adjust parameters to maintain sterile processing conditions and product quality, identifying deviations before they cause batch failures.

Collaborative sterile operations

Physical AI in cleanroom manufacturing is optimized for minimal human presence to reduce contamination risk. Over time, evolving safety and cleanroom standards may enable limited, governed human robot collaboration for exception handling and complex tasks—but this remains secondary to automation first sterile operations.

Cleanroom manufacturing automation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Robotic systems in pharmaceutical cleanrooms must perform reliably within validated sterile environments where any deviation—a dropped component, a missed quality check, or an uncontrolled movement—can contaminate a batch and trigger regulatory investigation. Reliability standards must meet pharmaceutical validation requirements, not just general industrial robotics benchmarks.

Responsible and accountable

Pharmaceutical manufacturing is one of the most heavily regulated production environments in the world. The audit trail requirement is not just good governance—it is a regulatory obligation. When AI contributes to a quality failure, the evidentiary record must be sufficient to satisfy FDA or equivalent scrutiny.

Transparent and explainable

Quality assurance teams and regulatory inspectors reviewing AI-driven inspection decisions need to understand what anomalies were detected, how severity was assessed, and what actions were taken—both to exercise oversight and demonstrate compliance. Opaque AI quality decisions that cannot be explained to a regulatory inspector are not compatible with pharmaceutical manufacturing standards.

POTENTIAL BENEFITS

Contamination risk reduction
Minimizing human presence in sterile zones reduces particulate generation and biological contamination while maintaining consistent environmental controls and product quality.

Regulatory compliance support
Automated documentation, process traceability, and consistent adherence to validated procedures strengthen an operation's regulatory compliance posture in heavily scrutinized pharmaceutical manufacturing.

Enhanced product consistency
Robotic systems eliminate human variability in repetitive tasks, delivering more consistent results in high-value sterile production where quality failures trigger costly batch rejections and regulatory scrutiny.

The Technology, Media & Telecommunications AI Dossier



The Technology, Media & Telecommunications AI Dossier

AI is reshaping the technology, media, and telecommunications sectors from the inside out. For businesses providing the digital foundation of the modern economy, AI technologies—including generative AI, AI agents, and physical AI—are driving both product innovation and operational reinvention. From code generation and content creation to network optimization and audience analytics, AI is accelerating time to market, enhancing user engagement, and unlocking new revenue models.

In the technology sector, AI is transforming how software is designed, built, and maintained. Developers are using AI-powered tools to write, test, document and debug code faster, while IT operations teams are deploying intelligent agents to monitor systems, predict outages, and automate resolution. In hardware, AI is enabling smarter design of chips and devices, shortening development cycles and improving performance.

Media organizations are leveraging generative AI to create hyper-personalized content, automate editorial workflows, and synthesize massive volumes of user data into actionable insights. Increasingly, agentic AI systems are moving beyond task automation—autonomously managing workflows, supporting technical sales, and orchestrating elements of content management across the industry landscape. Meanwhile, synthetic voice, image, and video capabilities are opening new formats for storytelling and user interaction.

In telecommunications, AI is enhancing technician support, streamlining network maintenance, and optimizing operations—helping providers deliver more reliable service and respond faster to customer needs. As 5G expands, AI will be critical to orchestrating complex, distributed infrastructure while enabling real-time services—from autonomous vehicles to immersive media.

Yet as AI becomes embedded in core platforms and content, the stakes rise around governance, transparency, and responsible use. Issues such as IP ownership, misinformation, algorithmic bias, and deep fakes present real risks. To stay competitive and credible, companies must balance speed of innovation with safeguards that build trust.

To stay competitive and credible, companies must balance speed of innovation with safeguards that build trust.

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

Tags | Primary business function | Agentic AI | Physical AI

AI-powered technical sales

Streamlining the end-to-end sales process with role-based AI agents

Agentic AI systems can augment sales teams by using specialized agents that mirror traditional sales roles, automate research and CRM tasks, and provide real-time support to boost sales performance and efficiency.

ISSUE/OPPORTUNITY

Technology sales are increasingly complex, requiring teams to manage long sales cycles, multiple stakeholders, and extensive product knowledge. Sales representatives often spend large amounts of time on administrative work—researching prospects, updating CRM systems, and preparing materials—leaving less time for relationship building and deal closure. These inefficiencies can reduce responsiveness to customers, lead to inconsistent engagement, and limit organizations' ability to scale sales operations effectively.

Tags: Sales | Agentic AI

HOW AI CAN HELP

Prospecting and research support

An AI-based business development representative (BDR) agent can research prospects, prepare initial pitch materials, and draft outreach communications, helping to provide consistent and timely engagement.

Lead qualification and CRM automation

Specialized AI agents can automatically update CRM systems, validate customer data, and handle routine record-keeping, reducing errors and freeing sales staff for higher-value activities.

Account executive (AE) enablement

An account executive AI agent can support the next stage of the sales process by preparing detailed materials, assisting with live customer inquiries, and coordinating documents to support the close of the deal.

Workflow integration

The AI agents can connect with sector-specific sales tools and platforms, align processes with industry best practices, and collaborate to provide seamless workflows across the sales organization.

AI-powered technical sales

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Incorrect data entry or poor workflow execution could result in errors that undermine credibility and damage customer relationships. To minimize such problems, AI agents should be continuously monitored, validated, and tested against live sales processes.

Private

Because sales processes involve handling sensitive customer information, agents should comply with all relevant data protection regulations (such as GDPR and CCPA), minimize unnecessary data exposure, and use encryption to secure communications and CRM updates.

Responsible and accountable

While AI can support sales teams with research and recommendations, ultimate responsibility for customer relationships, pricing decisions, and contractual commitments should remain with human sales professionals and business leaders.

Safe and secure

Sales platforms and CRM systems are prime targets for cyberattack. AI agents should be integrated with enterprise-grade cybersecurity safeguards to prevent breaches, account takeovers, or data leaks that could damage trust with customers.

POTENTIAL BENEFITS

Increased sales productivity and lower costs

Automating research, CRM updates, and other sales support tasks can reduce costs while enabling human sales teams to spend more time engaging customers and closing deals.

Improved customer experience

Faster responses and more consistent communications improve customer satisfaction and help strengthen relationships and trust throughout the buying journey.

Scalable sales operations

AI agents can increase sales capacity without proportional increases in headcount, enabling organizations to expand their business in new and existing markets.

AI agents for customer success

Using agentic AI to improve post-sale support and customer success

Agentic AI systems can improve customer success and maximize long-term customer value by using AI to handle numerous post-sale support activities, allowing humans to focus on relationship building.

ISSUE/OPPORTUNITY

For technology companies, post-sale customer success is increasingly central to revenue growth, with a customer's long-term value to the business hinging on renewals, upsells, and cross-sells. However, customer success managers (CSMs) are often stretched thin, spending significant time on administrative work such as preparing success plans, drafting strategic review materials, and managing renewals in CRM systems. This leaves less capacity for relationship building and strategic conversations that secure customer loyalty.

As organizations scale and expand their account portfolios, they have a critical need to automate routine post-sales activities while delivering high-quality, personalized support. Agentic AI offers a way to maintain high-touch engagement—even for down-market or high-volume accounts—without proportional increases in staffing.

Tags: **Customer Service** | **Agentic AI**

HOW AI CAN HELP

Customer success planning
An overarching customer success agent reviews CRM and sales data to draft initial success plans and then continuously updates them with new inputs about usage, deployments, and engagement.

Ongoing monitoring and insights
Specialized agents analyze customer health scores, product usage, and support interactions, identifying risks and opportunities for proactive intervention by CSMs.

Strategic review preparation
Another AI agent prepares first drafts of materials for quarterly or strategic business reviews, providing customer-facing teams with timely, relevant insights.

Renewal and upsell support
A renewal agent drafts a renewal playbook, updates the CRM system, prepares quotes, and recommends product or pricing mixes, then cycles information back to the customer success agent for onboarding or cross-sell support.

Conversational support
Agentic AI-powered virtual assistants orchestrate advanced language models to deliver natural, adaptive conversations—handling a wide range of customer queries in real time and multiple languages. By combining orchestration with generative AI capabilities, these agents provide fast, accurate, and consistent responses, while automatically documenting interactions for future reference. This enables scalable, personalized support at lower cost.

AI agents for customer success

MANAGING RISK AND PROMOTING TRUST

Robust and reliable
Since inaccurate data or flawed insights could undermine customer trust and loyalty, AI agents should be validated against actual customer usage patterns and regularly updated with new business rules.

Transparent and explainable
For customer-facing teams to trust and act on AI-driven recommendations, agents need to provide clear rationales for their outputs and highlight supporting data.

Responsible and accountable
Strong customer relationships are built on personal empathy and trust. As such, humans will always be a critical part of the sales and support process (pre- and post-sale) and have ultimate responsibility for building and maintaining relationships that last.

POTENTIAL BENEFITS

Higher customer retention
By proactively monitoring customer health and addressing problems early, AI agents can help organizations maintain strong relationships and reduce churn.

Increased sales growth
Agents can identify upsell and cross-sell opportunities that drive account expansion and long-term revenue growth.

Lower costs and improved productivity
Automation of routine planning, reporting, and CRM updates reduces costs and frees CSMs to focus on relationship building and other high-impact activities that boost customer satisfaction and value.

AI agents for software engineering

Automating the software lifecycle with multi-agent collaboration

Agentic AI systems can accelerate and improve software development by using specialized agents that generate, test, debug, and deploy code, enabling faster delivery of higher-quality applications at lower cost.

ISSUE/OPPORTUNITY

Software engineering is an increasingly complex challenge, with organizations under pressure to quickly deliver new features and products while maintaining high levels of reliability, security, and compliance. Traditional development approaches require large teams of software engineers to execute repetitive tasks such as writing boilerplate code, testing, debugging, and managing deployments. These activities are time-consuming and costly. Also, they can impair innovation and development by diverting skilled engineers away from strategic design and problem-solving.

Tags: R&D/Product Development | Agentic AI

HOW AI CAN HELP

Specialized coding agents
AI agents can generate code from natural language requirements, translating business or product needs directly into executable code.

Automated testing agents
Testing agents can validate functionality, identify defects, and ensure that new code integrates seamlessly with existing systems.

Debugging and troubleshooting agents
When errors are found, agents can automatically propose fixes or apply patches, quickly resolving problems and accelerating development timelines.

Deployment and operations agents
Specialized agents can handle integration, deployment, and monitoring, facilitating smoother releases and more resilient production environments.

Code summarization and documentation
Agentic AI can automatically generate clear, human-readable documentation alongside code, improving maintainability and accelerating future development—without burdening engineers with manual write-ups.

AI agents for software engineering

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

In software development, code quality and integrity are paramount. AI agents should be validated against test suites, subjected to rigorous regression testing, and monitored in production to help ensure reliability at scale.

Transparent and explainable

Since developers must understand, maintain, and extend AI-written code, agents should provide clear documentation and traceability of their logic and coding decisions. This includes explaining why specific code structures, libraries, or fixes were chosen.

Private

Agents must avoid exposing sensitive business logic, proprietary algorithms, or customer data during development and deployment. Data privacy protections and secure handling of intellectual property are essential, particularly when using shared or cloud-based training data.

Responsible and accountable

Although AI agents can perform coding tasks well, ultimate responsibility for design decisions, compliance, and final code deployment should rest with human engineers and engineering managers. Clear governance frameworks should ensure that accountability is not abdicated to autonomous systems.

POTENTIAL BENEFITS

Faster development cycles

Automation of coding, testing, and debugging shortens release timelines, enabling quicker delivery of features and products.

Improved code quality

Continuous AI-driven testing and refinement can reduce defects and lead to more reliable, secure software.

Greater productivity and efficiency

By offloading repetitive software development tasks, engineers can focus on higher-value activities such as strategic design, innovation, and problem-solving.

AI agents for service lifecycle management

Automating technology and telecom service activities

Agentic AI can manage key service lifecycle functions—including customer support, network operations, and billing—operating with minimal human intervention to improve efficiency and scalability.

ISSUE/OPPORTUNITY

Telecom and technology providers face rising demand for responsive, always-on service delivery. Traditional service lifecycle processes—such as handling customer inquiries, managing network performance, and processing billing—often rely on large teams performing repetitive tasks across multiple systems. These manual approaches can create delays, inconsistencies, and higher operating costs, while limiting the ability to scale operations effectively.

Tags: **Operations** | **Agentic AI**

HOW AI CAN HELP

Specialized functional agents
AI agents can be tailored for specific functions such as customer service, network monitoring, billing management, and field support, each handling tasks autonomously with high speed and consistency.

Customer support automation
In support environments, AI agents can triage tickets, resolve common issues, and escalate complex cases as needed, whether through voice or text-based channels.

Operational integration
Integrated into existing enterprise platforms, agents can update records, monitor systems, and provide real-time recommendations, seamlessly aligning with current workflows.

Scalable orchestration
Multi-agent systems can coordinate activities across service lifecycle stages, enabling efficient end-to-end management with minimal human oversight.

On-site troubleshooting support
AI-powered tools equip field technicians with access to troubleshooting guides and generate clear, step-by-step resolution plans, improving on-site problem solving and network performance.

AI agents for service lifecycle management

MANAGING RISK AND PROMOTING TRUST

Robust and reliable	Maintaining data integrity and ensuring accurate ticket resolution are critical, especially when agents operate autonomously. AI agents should be continuously validated for accuracy and stability.
Safe and secure	Given the risks of cyberattacks targeting critical telecom and technology infrastructure, agent systems should be hardened against intrusion and designed with robust cybersecurity measures to protect networks from sophisticated attacks.
Private	Since agents often handle sensitive customer and operational data, strong security protocols and strict compliance with privacy regulations are essential.
Responsible and accountable	Agentic AI can make mistakes or overlook nuances in complex service cases. Human supervisors need to closely monitor AI agents and retain final responsibility for actions and decisions, intervening whenever necessary—especially in situations that involve sensitive customer information or regulatory compliance.

POTENTIAL BENEFITS

Greater efficiency and lower costs

Automating routine service tasks can improve the productivity of support and operations teams, reducing both the time and cost required to address customer inquiries and technical issues.

Improved customer satisfaction

AI-driven service that is faster, more consistent, and more reliable enhances customer satisfaction and loyalty.

Scalable operations

Organizations can expand support capacity and manage higher service volumes without proportional increases in headcount. Also, existing service staff can shift their focus to activities that are more complex and strategic.

Efficient marketing spend

Using AI agents to improve the efficiency of marketing spend

Agentic AI systems can help media and entertainment companies get more value from their marketing spend by forecasting financial outcomes, evaluating ROI, and running budget scenarios to optimize growth, retention, and profitability.

ISSUE/OPPORTUNITY

Media and entertainment companies invest heavily in marketing to drive subscriptions, viewership, and engagement; however, measuring and optimizing financial impact remains a complex challenge. Marketers must balance budgets across multiple channels, predict outcomes such as customer acquisition and retention, and spend their money as efficiently and effectively as possible—particularly during high-profile events, which have outsized financial implications.

Traditional forecasting and ROI analysis often rely on manual processes and disconnected tools, making it difficult to produce timely, accurate insights. Without more adaptive, data-driven approaches, marketers risk overspending, missing revenue opportunities, or failing to adjust strategies quickly enough to reflect changing consumer behavior and market conditions.

Tags: **Marketing** | **Agentic AI**

HOW AI CAN HELP

Financial forecasting
Marketers need tools that can continuously analyze performance, forecast outcomes, and support decision-making on where best to invest their resources. A financial forecasting agent can analyze marketing budget inputs and business assumptions to generate full-year projections for revenue, spend, and profitability.

Marketing effectiveness analysis
A marketing effectiveness agent can analyze historical and current performance data to provide recurring insights on ROI, ROAS, and customer lifetime value.

Scenario planning
A scenario planning agent can run simulations of different budget strategies to model impacts on subscription starts, retention, and overall profitability.

Coordinated decision support
By sharing insights across forecasting, effectiveness, and scenario planning, AI agents can collaborate to provide a comprehensive view of marketing's financial contribution.

Efficient marketing spend

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Inaccurate forecasts and simulations can lead to poor budget decisions. To improve reliability, AI agents should be continuously tested and refined to reflect real-world market conditions.

Private

Because customer-level data is used to measure marketing effectiveness, strong data governance and compliance with privacy regulations are essential.

Responsible and accountable

AI agents are fallible and cannot perfectly anticipate all external factors or sudden market shifts. Also, AI hallucinations remain a concern. Ultimately, humans need to be responsible and accountable for all decisions and actions.

POTENTIAL BENEFITS

Higher revenue

By automating forecasting and scenario planning, agentic AI helps drive higher revenue through more accurate bundle sign-up predictions and targeted marketing efforts.

Improved efficiency and conversion rates

Greater efficiency in owned and paid channels—including both traditional linear TV and connected TV (CTV)—can lead to better conversion rates for subscriptions.

Faster, more agile decisions

Agentic AI enables a business to respond more quickly to market changes and optimize marketing spend during high-impact events.

AI-supported budget allocation

Using AI agents to optimize how marketing budgets are allocated

Agentic AI systems can inform and optimize global, regional, and national marketing budget allocations across cost centers, brands, content types, and performance channels.

ISSUE/OPPORTUNITY

Global media and entertainment companies operate across multiple regions, products, and channels, with significant marketing investments needed to drive subscriptions and engagement. Traditional budgeting processes often rely on static planning cycles and siloed data, making it difficult to dynamically reallocate spend across cost centers, brands, content types, and performance channels.

Without better budgeting tools, organizations risk under-investing in high-performing areas, overspending in less effective channels, and creating global and regional strategies that are out of line with local market realities.

Tags: Finance | Agentic AI

HOW AI CAN HELP

Cost center analysis
AI agents analyze marketing spend across cost centers and recommend reallocations that optimize key performance indicators (KPIs) such as sign-ups, churn, and viewership.

Product, channel, and vendor optimization
Other agents assess product, channel, and vendor mix within each market, using predictive models to maximize subscription starts and profitability.

Regional and country-level scenarios
Scenario planning agents run simulations of budget allocations at regional and country levels, evaluating the impact of incremental or cost-saving changes on subscriptions, retention, and profitability.

Coordinated recommendations
Agents collaborate to provide a unified view of marketing performance and deliver actionable recommendations for smarter, more flexible budget allocation.

AI-supported budget allocation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable
Bad marketing budget decisions can have serious consequences for the broader business. Agents should be validated regularly against historical budgets and real-world outcomes, using continuous performance tracking to help ensure accuracy and minimize the risk of misallocation.

Transparent and explainable
Agents should be able to provide clear explanations of their reasoning—e.g., how reallocating $X from channel A to channel B affects ROI—so marketing leaders can understand and trust their recommendations.

Responsible and accountable
AI agents should be used as decision-support tools, with final responsibility reserved for marketing and finance leaders who understand strategic business contexts.

Safe and secure
Given the sensitivity of financial plans, systems must be built with enterprise-grade security to help ensure budget forecasts and proprietary data are protected from breaches or misuse.

POTENTIAL BENEFITS

Improved financial outcomes
Optimizing the product, channel, and vendor mix within each country can maximize subscription starts, customer retention, and overall profitability.

Greater spend efficiency
Efficiency improvements in both CTV and owned marketing channels can drive higher conversion rates and reduce wasted marketing spend.

More flexible decision-making
Real-time scenario planning enables marketing teams to adjust budgets dynamically, aligning strategies with rapidly changing market conditions.

AI for gamers

Game content development

Developers can leverage AI—including generative models— to maintain and update their game with new assets and content in line with user community requests and interests.

ISSUE/OPPORTUNITY

Game development requires a massive up-front investment in time, resources, and capital. AAA games can cost tens of millions of dollars to develop and take years to complete. These costs will only rise as players increasingly demand more complex games, more post-release support, and more frequent content updates. AI provides the gaming industry with an opportunity to bend the cost curve through enhanced development efficiency, while also simultaneously meeting player demands.

Tags | R&D/Product Development

HOW AI CAN HELP

Ongoing content development

Post release, developers can rapidly generate and deploy new gaming assets as expansions or microtransactions, such as seasonal or downloadable content (e.g., new characters and skins). Developers can use text prompts to generate new content in line with the existing game and requests from the community, then upload those assets to the game.

AI for gamers

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable	Generated content resulting from a model trained with proprietary third-party data may lead to copyright claims if it is deemed to be too similar (without substantial variation).
Safe and secure	Players' personally identifiable information (PII) could be fed into the models as they interact within the game, which raises risks around cybersecurity and regulatory compliance. The collection of PII, even inadvertently, places an obligation on the organization to secure the data as it is accessed, transferred, and stored.
Fair and impartial	Generated assets may over-index on player segments providing feedback or residing in specific regions. This uneven sampling of the input data could lead to bias in what assets are generated, and it may lead to missed opportunity and revenue as some customers are ignored.

POTENTIAL BENEFITS

Greater efficiency for greater creativity

By automating the process of creating game content, developers have more capacity to work on creative game designs and explore new, innovative ideas.

Cater to gamers

More immersive, controllable, responsive, engaging, and unique experiences for gamers (based on community requests and existing popular assets) has a direct impact on player lifetime value.

Drive new revenue

When add-on content can be generated with minimal human involvement, it creates new revenue streams with minimal investment.

Content creation with AI

AI-enabled creative tools

Content creation can be facilitated and enhanced with AI tools that minimize the need for manual editing and time-consuming content management.

ISSUE/OPPORTUNITY

Content creators and managers are faced with large volumes of data that require considerable time to generate, edit, and oversee. Creators also face tight deadlines that require high levels of efficiency for content management and editing.

Significant time and resource investments are needed for video and image editing, and the volume of content creates challenges around data management and finding the right content at the right time.

Tags | R&D/Product Development

HOW AI CAN HELP

Creative assistant tool
AI can be used to create imagery and apply edits using descriptive commands. Features like conversational editing, text-to-template, and text-to-image allow users to expedite the editing phase of the content creation process.

Picture editorial
Producers can automate footage management with video-to-text AI to evaluate and create tags for scenes and content. Text-to-video commands (e.g., "add more rain to this scene") can be used to enhance and accelerate the editing process.

AI "reshoots"
Content creators can use scripts and 3D scans of actors to generate new content, alter footage to create more realistic special effects, and allow studios to make edits without the need for reshoots.

Content creation with AI

MANAGING RISK AND PROMOTING TRUST

Responsible and accountable	AI tools may be trained with large databases of media and content, some of which may be copyright protected. As a result, the model outputs may include aspects of a creator's or studio's work or style that are not attributed to them, which raises legal risks.
Robust and reliable	Noticeable changes in style and brand quality due to AI content creation and editing may erode consumer trust in the brand and content.
Private	If bad actors access the underlying models or applications, it could contribute to the spread of fake content on behalf of the organization, leading to misinformation. Model owners should ensure strong privacy and access controls to mitigate this risk.

POTENTIAL BENEFITS

Greater efficiency

Content management stakeholders can gain efficiencies by leveraging creative tools to facilitate work and even create net-new content across the production lifecycle.

Improved content quality

Generating novel content can supplement the human creative process and potentially lead to a highe-quality product.

Content tailored to the audience

With AI, creators can hyper-personalize content with prompts driven by consumer trends and interests.

Marketing content multiplier

On-brand publishing

Using AI, marketing content generation can be quicker and more cost effective, while still preserving the company's brand identity.

ISSUE/OPPORTUNITY

When multiple authors are contributing to a piece of marketing or business content, there can be quality and consistency issues with tone and brand values. Authors are challenged to consistently balance product promotion with thought leadership and insight. As such, on-brand publishing is a significant time and cost investment that requires a long-term commitment to generating content that establishes the organization's (or its leaders') subject matter authority. However, the return on investment for on-brand publishing can be difficult to measure because the impact itself is complex and challenging to quantify.

HOW AI CAN HELP

Cohesive content generation
AI systems can be trained with on-brand content to mimic the style of company marketing materials and generate new, high-quality content rapidly and on demand.

Tailored, personalized messaging
With AI, organizations can easily create multiple versions of the same on-brand marketing tailored to different customers and audiences.

Ideation with generation
Marketing departments can leverage AI to quickly create multiple versions of content in various styles to identify the most compelling and persuasive option.

Tags: Marketing

Marketing content multiplier

MANAGING RISK AND PROMOTING TRUST

Transparent and explainable

Personalized advertisements may be customized based on data collected or purchased from individuals. This may be off-putting to consumers who realize the organization has such broad access to their data, potentially harming the enterprise's brand reputation and undermining consumer trust. One way to mitigate this outcome is to ensure data collection and usage policies are transparent and communicated meaningfully to the consumer.

Responsible and accountable

Content produced by AI systems may not be subject to the same protections as human-generated content. Companies need to be wary of infringing on copyrighted material used to train AI systems.

Safe and secure

When brand data is used to train AI, there is a risk of data leaks that could result in sensitive information or IP being divulged to competitors. Companies need to ensure that their proprietary information is safely stored, transferred, and used, as well as monitor model outputs to validate that protected information is not being revealed.

POTENTIAL BENEFITS

Instant marketing

Companies can create content better tailored to their brand and customers, iterating through multiple drafts as needed.

Time savings

As AI systems instantly generate content, human staff can shift to an editorial role.

Broader marketing range

With the ability to easily create content across various formats, styles, and topics, companies enjoy greater flexibility in how they reach their customers. This also allows companies to more rapidly adapt to marketing trends.

Language translation at scale

Content localization

AI can be used to quickly and easily scale content across regions by translating and converting text and audio into regional languages.

ISSUE/OPPORTUNITY

The ability to create and translate content at scale can be a competitive differentiator for multinational enterprises, but it can also command significant time and resources, and rapid, on-demand translation may be difficult to achieve.

HOW AI CAN HELP

Tools for custom localization and quality assurance
AI can be used to help organize and manage complex file type, analyze content before translation to optimize localization, and integrate glossaries, term bases, and language tools into workflows.

Content personalization across industries
AI-powered content personalization can supercharge localization efforts by helping to improve engagement, build brand loyalty, and increase conversions.

Speech recognition during translation
AI can be leveraged to enable voice user interfaces (VUI), transcribe video and audio content into text, and simultaneously translate spoken content into the target language.

Tags: **Operations**

Language translation at scale

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

Bias in the data used for content personalization could lead to unequal and unfair recommendations for certain groups of customers. In addition, AI applications are often trained on datasets from major languages, which means LLMs may have lower accuracy rates for less common languages and alternative dialects.

Transparent and explainable

Messaging and tone may change with language translation, which may negatively impact the text or audio being generated and the overall quality of the content. Localization should be audited to make sure that the messaging remains consistent with the original intent.

POTENTIAL BENEFITS

Enhancing translation

Translation processes using AI can lead to improved speed, accuracy, and scalability.

Improving the customer experience

A wider availability of language resources with the quality and speed enabled by AI promotes a higher-quality user experience.

Ensuring quality

Organizations can leverage AI to automate quality assurance for the localization of digital assets by providing more accurate natural language processing.

Enhancing chip innovation

Semiconductor chip design & manufacturing

AI can be used to iterate chip designs by having designs "compete" across a set of performance dimensions.

ISSUE/OPPORTUNITY

With demand for ever more powerful semiconductor chips, design complexity is rising. As semiconductor sizes continue to shrink, density scaling becomes a challenge, since upgraded features are required to fit on perpetually smaller chips.

HOW AI CAN HELP

Iterative chip design
AI can generate and iterate chip designs and improve the outputs by having chip designs "compete" across a set of performance dimensions. At each new iteration, chip parameters are tweaked based on learnings from the best-performing designs in past iterations. These models are trained on existing layouts to learn patterns and constraints and generate new layouts that meet specific design requirements.

Tags: **R&D/Product Development**

Enhancing chip innovation

MANAGING RISK AND PROMOTING TRUST

Safe and secure

With each new generation of novel designs, there is a risk of IP leakage and data breaches for proprietary chip designs and technical specifications generated by the LLM that could severely damage the organization's competitive advantage. There should be rigorous security protocols in place to protect against this.

Responsible and accountable

When using AI for design, the organization needs to consider how to secure copyrights or patents and protect the IP of chip designs that are moved into production.

Transparent and explainable

For complex simulation processes, the organization needs the capacity to understand how and why the model determined a scenario or design to be optimal. Design validation requires users and stakeholders to be able to understand the reason for the outputs.

POTENTIAL BENEFITS

Cost and time

By shortening the development lifecycle, the enterprise can reduce total development costs.

Create new ideas

AI can help improve designs or discover entirely novel designs that optimize performance based on specific criteria, such as power consumption, performance, location, and manufacturability.

Tech specs on demand

Technical sales, operations, and field staff knowledge management

AI can help sales, operations, and field staff quickly find and translate technical specifications to enable faster knowledge retrieval.

ISSUE/OPPORTUNITY

Technology offerings require deep technical understanding and the ability to find the right technical specifications in a timely manner. When it comes to translating technical specs and responding to customer technical questions, sales, operations, and field staff can be challenged to translate the information and effectively communicate it to the customer.

A big part of the problem is the time-consuming process of scouring through vast amounts of unstructured information and knowledge documents that contain the specifications and answers customers are seeking.

Tags: **Sales**

HOW AI CAN HELP

Spec summarization and search
AI can be used to create summaries of technical specifications based on targeted text-based queries to help understand which products meet customer requirements. It can suggest features and integrations that align with the customer's existing technology stack and vendors, as well as provide links to articles or an internal knowledge base for future reference.

Knowledge management update
Sales case histories and other current documents can be used to update knowledge management so similar technical inquiries in the future can be rapidly addressed with previous resolution steps and summarizations.

Automated technical demos
AI can be used to automate the creation of software demonstrations tailored to specific clients and use cases. This is achieved by training on demo scripts and sample interactions to generate demonstrations showcasing a solution's key features and benefits.

Tech specs on demand

MANAGING RISK AND PROMOTING TRUST

Private

Customer data (e.g., sales case history, customer tech stack/ vendors) must be processed by the model, making it necessary to continuously monitor model outputs and safeguard customer data to mitigate privacy risks.

Robust and reliable

AI models are susceptible to hallucinations, or factual inaccuracies, making human validation essential to establishing trust in the outputs and the decisions they inform. What is needed is a verification process to ensure the accuracy and reliability of information derived from the model (e.g., spec summarization, demos). This has a direct impact on answering customer questions, and by extension, customer satisfaction and sales.

POTENTIAL BENEFITS

Efficiency with automation

AI's ability to quickly consult and summarize technical specifications greatly reduces manual effort for sales, operations, and frontline staff when responding to technical sales inquiries.

Tailored to the customer

Greater personalization in responses and demonstrations improves the customer sales experience and increases the chances for conversion.

Supporting sales staff and other stakeholders

With AI, staff can rapidly create content to support the sales and marketing process, and to address specific questions from customers and partners.

AI-powered RFP and knowledge assistant

Automated proposal generation and sales knowledge management

AI can produce RFP responses automatically and help sales teams prepare for pitches by providing easy access to internal knowledge resources through smart chatbots.

ISSUE/OPPORTUNITY

Sales processes are often constrained by how quickly teams can access institutional knowledge and respond to Requests for Proposals (RFPs). Many sales teams have only days to coordinate across multiple departments and deliver detailed technical and commercial responses. Their ability to respond can be slowed by manual processes, fragmented internal documentation (e.g., playbooks and product briefs), inconsistent proposal quality and knowledge reuse across teams, and limited tools to extract and synthesize key information.

AI-powered tools can accelerate sales professionals' ability to retrieve, understand, and reframe information for client needs— without requiring technical expertise or deep coordination across departments.

Tags: **Sales**

HOW AI CAN HELP

Providing easy access to internal knowledge through chatbots
Salespeople can converse with AI-powered chatbots to quickly and easily retrieve sales playbooks, technical specs, competitive positioning, and customer references directly from internal documentation repositories.

Automatically drafting RFPs
AI models can produce high-quality, tailored RFP responses by finding and summarizing relevant content from existing sales documents, aligning answers with internal knowledge bases, and incorporating reusable proposal components.

Providing individualized sales support with little or no coding
Non-technical users, including sales reps and subject matter experts, can generate summaries, extract insights, and draft proposals through a simple user interface—no prompts or coding required.

Enabling customized sales processes and tools
Technical users can integrate AI tools directly into other internal systems, workflows, or dashboards to build more personalized applications.

AI-powered RFP and knowledge assistant

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

AI-assisted workflows can deliver higher-quality outputs—in less time— than traditional approaches. Users can flag incorrect responses or incomplete information; these are logged and reviewed in recurring QA cycles. Also, fallback mechanisms should exist to ensure consistent availability if problems arise with the AI models.

Transparent and explainable

Documentation should be provided for both business users and developers to explain how the system processes inputs and generates outputs. The chatbot interface includes citation tracing, where users can see which source documents were utilized to generate responses. Proposal-generation tools can allow users to edit and review outputs before submission, promoting human-in-the-loop oversight and transparency.

Safe and secure

All data and model interactions should occur within a secured internal environment, with no calls to third-party APIs unless vetted and approved. Systems should support audit logging for all user interactions to help ensure traceability and compliance. Role-based access controls can ensure only authorized personnel are able to view or generate sensitive proposal content.

Private

The system should not log personally identifiable information (PII) unless required by specific business rules and protected under internal data governance protocols. Feedback mechanisms should be anonymized where appropriate, helping to ensure user privacy while supporting continuous improvement. RFPs and customer documents processed in the system should be stored temporarily and purged according to data retention policies.

POTENTIAL BENEFITS

Faster deal cycles

Sales teams can respond to RFPs and prepare sales pitches/collateral much more quickly than before, accelerating the sales cycle.

Higher win rates

With centralized, AI-assisted knowledge access, sales teams can produce responses that are more consistent and comprehensive—reducing errors and potentially improving win rates (especially for opportunities with time-sensitive budget windows).

Increased sales rep productivity

Salespeople can search for materials or draft proposals more quickly, freeing them to focus on sales strategy, client relationships, and personal follow-ups.

Path to commercialization

Once validated internally, AI-powered sales tools have the potential to be offered to external customers, turning an internal efficiency driver into a revenue-generating product.

Automated test case generation

AI-powered test case generation and automation in chip development

As chip designs become more complex and product cycles accelerate, engineering teams are leveraging AI to automate test case generation and validation.

ISSUE/OPPORTUNITY

Chip development demands exhaustive testing and validation due to increasing functional complexity and the high cost of post-release defects. Human testers may struggle to keep pace with the volume and sophistication of required test cases, leading to potential quality issues, slower development cycles, and growing verification costs. Yet, security vulnerabilities or missed bugs can result in major product delays, public backlash, and brand damage, prompting chip manufacturers to add even more layers of testing.

Tags | R&D/Product Development

HOW AI CAN HELP

Automating test creation
AI tools can be used to create new test cases from product requirement documents, bug histories, and structured datasets. These tools can assist engineers by proposing a wider set of test scenarios including ones not previously considered and by automating portions of test implementation through code generation.

Identifying test gaps
AI systems can also help identify gaps in testing coverage and can prioritize high-risk areas based on historical failure data, although integration with structured data and internal governance systems remains an ongoing challenge.

Automated test case generation

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Generated test cases can be validated against known test results and manually vetted to help ensure they hold up under real-world complexity. Also, systems can be stress-tested with increasingly complex product requirement documents to assess scalability and robustness across chip generations.

Transparent and explainable

AI-generated test cases can be accompanied by natural language summaries or rationales explaining why certain logic or edge conditions were selected. Engineers can trace outputs back to source inputs (e.g., PRD sections, bug databases), enabling better understanding and debugging of the AI system itself.

Safe and secure

The development and inference processes can occur in sandboxed environments with strict access controls to prevent accidental leakage of proprietary information. Integration with external AI services should be carefully managed to ensure no sensitive IP or design data is exposed to third-party systems.

POTENTIAL BENEFITS

Increased test coverage and enhanced product quality

AI can enable the generation of more comprehensive test cases than previously possible with human effort alone, allowing for earlier defect detection. Also, by identifying edge cases and potential failure modes, AI can reduce the risk of catastrophic bugs slipping into production.

Faster time-to-market

Automation accelerates the validation process, allowing development teams to keep up with faster chip release timelines and feature rollouts.

Operational efficiency and cost control

AI helps teams do more with less, reducing reliance on manual testers and mitigating the need to grow headcount to handle increasing workloads.

Improved development process

As the test tools mature, there is potential for deeper integration with the design and verification phases, improving end-to-end development flow across decentralized teams.

AI-powered source separation for music remastering

Separating mixed audio tracks into their component parts using AI

AI can separate vocals or instruments from mixed audio tracks even when the original files are not available, opening up new opportunities for licensing, remixing, archival restoration, and monetization.

ISSUE/OPPORTUNITY

Many recordings in music labels' back catalogs were produced at a time when multitrack preservation practices were inconsistent, and, in many cases, the original recordings have been lost, damaged, or never existed in isolated formats. This limits the ability to fulfill requests for custom edits—such as instrumentals, a cappella songs, or remixes— thereby stalling lucrative licensing deals, particularly for synchronization (music in film, television, and advertising) and derivative content creation. Manual audio reconstruction is costly, time-consuming, and often technically infeasible at scale. Yet demand for high-quality, tailored audio continues to grow, especially with the global expansion of streaming and sync opportunities.

Tags | R&D/Product Development

HOW AI CAN HELP

Separating music into its component parts

AI, particularly deep learning-based source separation models, can analyze a fully mixed audio file and isolate its constituent elements— vocals, guitar, bass, drums, ambient noise, etc.—into discrete audio tracks with high fidelity. These models have matured significantly in recent years and can now perform at a level sufficient for commercial use in many scenarios. Rather than depending on traditional DSP (digital signal processing) or manual studio methods, the AI learns from large datasets of music to "de-mix" the sound using learned patterns of frequency and structure. .

Leveraging Software-as-a-Service

Most deployments today use AI-powered SaaS platforms that allow internal teams to process catalog tracks quickly and securely. Internal quality control—along with artist or management approval—is then layered on to ensure that the extracted stems meet the creative and technical expectations of the project.

AI-powered source separation for music remastering

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

All outputs from AI models are subject to expert human review. Because source separation can introduce artifacts, tracks should be assessed case-by-case to determine if the fidelity is suitable for commercial or creative use. Teams should be trained to identify when alternative methods or manual interventions may be more appropriate.

Transparent and explainable

Processes for using AI in audio separation should be clearly defined internally and communicated externally as needed. Stakeholders—including sync partners, artists, and producers—should be informed when AI-generated stems are used, and how those stems were derived from the source material.

Responsible and accountable

All source separation use should be logged, and responsibility for approving commercial use should rest with both label- and artist-facing teams. If stems are to be reused, remixed, or publicly released, the appropriate clearance workflows—including licensing and revenue-sharing— must be followed.

POTENTIAL BENEFITS

Commercial monetization of back catalogs

AI-powered source separation can make more recordings available for synchronization deals, remixing projects, or global reissues in alternate formats.

Accelerated time-to-license

The speed and efficiency of AI can minimize delays associated with locating or recreating stems, enabling a faster turnaround for time-sensitive media productions. .

Cost-efficient alternative to studio sessions

AI offers a high-quality yet faster and less expensive alternative to manual isolation or re-recording, which are both time-consuming and expensive.

Artist-led remix and reimagination projects

Using AI to extract source elements, artists can revisit and reinterpret their own work or collaborate across genres. Even in less creative scenarios, artists and labels can maintain full control over what gets extracted and used, ensuring all usage aligns with legal, creative, and ethical standards.

Operational scalability

AI can systematically process large volumes of tracks, with human review reserved for final quality control, increasing throughput without compromising quality.

AI-powered archive access and extraction

Transforming historical news content into a valuable asset

AI enables news organizations to recover legacy content lost to system or format issues—turning dormant information into a usable, searchable, and monetizable asset.

ISSUE/OPPORTUNITY

News archives hold cultural, journalistic, and commercial potential. But over time, many of the most significant stories—especially interactive long-form journalism, investigative pieces, and special coverage—have become inaccessible due to technological evolution, changes in content management systems (CMS), format obsolescence, and a lack of centralized archives.

Reporters and editors often cannot locate stories they know exist, especially from the early digital era (late 1990s to early 2010s). Multimedia components such as photos, graphics, and maps have not always been retained or migrated, rendering even recovered content incomplete.

Tags: **Operations**

HOW AI CAN HELP

Document extraction and digitization
AI models can process and extract structured information from legacy formats such as PDFs, microfilm scans, and outdated HTML, even when metadata is missing or inconsistent.

Content reconstruction
AI tools can intelligently identify article structure (headlines, subheads, body text, captions, bylines), reconstruct layout context, and reassemble fragmented articles into coherent, readable documents.

Semantic indexing and search
Large Language Models (LLMs) enable content to be semantically tagged and categorized, improving discoverability across themes, time periods, people, and places—even when specific keywords are not used.

Metadata enrichment and linking of multimodal assets
AI can supplement missing or corrupted metadata (e.g., publication date, author, topic) by analyzing linguistic and contextual clues. Also, the technology can cross-reference and re-link associated images, graphics, or videos from various archives where files may have been separated during prior migrations.

Improved access
AI can provide improved interfaces—such as chat-style queries or timeline exploration—to help users engage intuitively with the archive.

AI-powered archive access and extraction

MANAGING RISK AND PROMOTING TRUST

Fair and impartial

Systems are designed to ensure access to historical content across different eras and communities. Bias mitigation strategies are incorporated into model training and metadata tagging to avoid skewed representation of topics, regions, or individuals.

Robust and reliable

Extraction and structuring workflows are tested across various content types and legacy formats to help ensure consistent quality. Human oversight is embedded throughout the process to validate the accuracy and fidelity of reconstructed articles.

Transparent and explainable

A clear audit trail should be maintained for all AI-processed content, including logs of when and how specific items were extracted, tagged, and categorized. Explanatory overlays and metadata annotations help end-users understand the origin and limitations of AI-reconstructed documents.

POTENTIAL BENEFITS

Editorial improvements

Journalists can rediscover and repurpose historic reporting, improving storytelling quality and institutional memory. The AI-powered solution speeds up research for retrospective or investigative reporting by eliminating the need to manually dig through archives.

Monetization

AI can enable news organizations to expand their relationships with libraries, educational institutions, and content platforms while providing the foundation for new archive-based products, such as nostalgia-based newsletters and historical collections. What's more, it positions news organizations to negotiate more effectively with AI companies looking for premium training data by presenting them with a curated, high-quality proprietary dataset.

Improved operational efficiency

The solution can reduce ad-hoc archive retrieval and reduces the need for internal technical support to help recover content. Also, it strengthens the organization's institutional capabilities for structured knowledge management.

Network sensing, assurance and autonomous recovery

AI-supported reliability for national networks

Machine learning and agentic AI systems analyze telemetry from network assets—such as switches, routers, radio units, antennas, and fiber—to detect faults, localize root causes, and assess severity. This is complemented by drone based inspections that capture visual and sensor data from hard to reach infrastructure, with AI correlating physical observations and network signals to improve diagnosis. Human operators validate insights and execute remediation actions, ensuring control and compliance.

ISSUE/OPPORTUNITY

Telecommunications networks are large, complex physical systems where localized faults can quickly cascade into widespread service disruption. Manual monitoring struggles to keep pace with the scale and velocity of telemetry data, resulting in delayed detection and prolonged outages—conditions made less tolerable by strict regulatory requirements and high customer expectations for uptime. Network operations centers receive thousands of alerts daily from distributed infrastructure spanning cell towers, fiber nodes, data centers, and customer premises equipment. Operators must distinguish genuine failures from routine fluctuations and trace root causes amid cascading alarms, where a single physical issue—such as damaged fiber or a failing router—can trigger hundreds of downstream alerts. Drone based inspections are used to survey towers, antennas, and fiber routes, capturing visual and sensor data from hard to reach assets. Physical AI correlates this drone inspection data with network telemetry to identify incidents earlier, pinpoint root causes more clearly, and support faster resolution—while keeping humans in control and ensuring regulatory compliance.

Tags: **Information Technology** | **Physical AI**

HOW PHYSICAL AI CAN HELP

Telemetry-based fault detection

AI systems continuously monitor telemetry from physical network equipment to detect abnormal behavior indicating faults before they cause widespread service impact.

Severity and impact assessment

AI helps classify incidents based on scale and potential customer impact, enabling operators to prioritize responses to the most critical failures first.

Cross-network correlation

Signals from multiple network components are analyzed together to identify where issues originate, filtering out cascading alarms that merely reflect downstream effects.

Remediation recommendation

AI proposes corrective actions for operator review such as rerouting traffic, restarting equipment, or dispatching field technicians to specific locations.

Human-in-the-loop execution

All actions are validated and executed by operations teams, maintaining accountability and compliance with regulatory requirements for network changes.

Root-cause identification support

Likely causes are suggested based on observed network behavior and correlated with drone based visual inspections of physical assets, accelerating diagnosis by highlighting probable failure points for targeted operator investigation.

Network sensing, assurance and autonomous recovery

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

AI fault detection systems must distinguish genuine failures from routine fluctuations reliably. False positives waste operator attention, while false negatives allow faults to cascade into widespread outages. Reliable operation across a diverse set of network equipment, traffic patterns, and failure modes is essential.

Safe and secure

Telecommunications networks are critical national infrastructure, and AI systems that recommend traffic rerouting or equipment interventions are a high-value adversarial target. A compromised detection system could suppress genuine alerts or trigger unnecessary interventions, causing service disruption at scale. Security protections must reflect the critical infrastructure context and the public consequences of system compromise.

Responsible and accountable

All remediation actions require human validation and execution—maintaining accountability and regulatory compliance for network changes. This demands clear documentation of what the AI detected, what was recommended, who approved execution, and what outcome occurred, creating an audit trail supporting regulatory reporting, SLA dispute resolution, and post-incident investigation.

POTENTIAL BENEFITS

Reduced outage duration
Earlier detection and diagnosis shorten incidents by enabling faster operator response and reducing time spent on manual root cause analysis.

Operational focus
Teams spend less time triaging alerts and correlating symptoms, as AI combines network telemetry with drone based inspection data from physical assets, allowing network engineers to focus on resolution rather than diagnosis.

Faster fault isolation
Operators can identify and localize issues more quickly, reducing mean time to repair and minimizing the duration of service outages when failures do occur.

SLA protection
Improved compliance lowers penalty risk by preventing service level agreement violations through proactive incident management and faster restoration.

Lower field maintenance costs
Proactive maintenance reduces emergency repair visits by preventing urgent failures, reducing overtime labor and expedited parts shipments.

AMR-enabled physical AI for predictive quality control

Predictive maintenance insights for operators

Physical AI systems analyze trends across the end-to-end infrastructure value chain—from component quality and deployment conditions to network and equipment performance—in manufacturing facilities, to predict and prevent service-impacting issues before they occur. Autonomous Mobile Robots (AMRs) act as mobile inspection and sensing platforms within data centers and network operations environments, capturing high-frequency visual, environmental, and asset-health data, reasoning over in-process signals in real time, and triggering corrective actions before defects propagate downstream.

ISSUE/OPPORTUNITY

Quality control today is largely reactive, identifying defects only after materials, machine time, and labor have already been consumed—driving scrap, rework, delays, and margin erosion. The core limitation is the lack of real time visibility into upstream indicators such as material variation, equipment drift, and environmental change. Physical AI addresses this by continuously analyzing patterns across materials, processes, equipment, and environment, with Autonomous Mobile Robots serving as mobile, in line sensing and inspection platforms across the factory floor. Equipped with vision systems, sensors, and edge AI, AMRs capture high frequency, in process quality and equipment data that is fused with production and material telemetry, enabling closed loop AI systems to predict quality risks early and stabilize production before defects occur.

Tags: **Manufacturing & Quality** | **Physical AI**

HOW PHYSICAL AI CAN HELP

Predictive analytics across supply chain

AI systems continuously monitor material quality at intake, tracking variations and correlating them with downstream production outcomes to predict potential defects before materials enter production.

Real-time process optimization

Machine-learning models analyze equipment performance data, environmental conditions, and production parameters to detect early signs of process drift and automatically trigger corrective adjustments.

Automated robot reprogramming

Based on predictive insights, AI systems autonomously adjust robot parameters, tooling settings, or process flows to compensate for predicted variations, eliminating manual intervention cycles.

AMR-enabled physical AI for predictive quality control

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Predictive quality systems must perform reliably across continuous production variability. A system that misses early drift indicators allows defects to propagate through downstream processes, while a system that generates excessive false alerts disrupts production unnecessarily. Both failure modes erode the operational trust on which closed-loop automation depends.

Responsible and accountable

If AI can autonomously adjust robot parameters and process flows, organizations must define clear governance boundaries specifying which adjustments can be made autonomously and which require human validation before execution. This helps prevent unreviewed AI decisions from affecting product quality at scale.

Transparent and explainable

To validate the model's reasoning and detect drift, process engineers need to understand what upstream signals drove a quality risk prediction and what corrective adjustment was applied. Closed-loop systems that adjust processes without explainable reasoning make it difficult to identify systematic errors before they affect product quality or propagate downstream.

POTENTIAL BENEFITS

Defect prevention and scrap reduction

Predictive AI, supported by AMR-based in-line inspection, detects quality issues early in TMT manufacturing lines, preventing defects from spreading and reducing scrap, rework, and wasted materials.

Optimized production efficiency

Automated process adjustments eliminate production interruptions for manual corrections, maintaining continuous flow and maximizing throughput while ensuring consistent quality standards.

Supply chain intelligence integration

End-to-end visibility enables manufacturers to provide feedback to suppliers about material quality trends, improving upstream quality and creating closed-loop optimization across the entire supply chain.

Sustainable network operations

Energy-aware control of physical networks

Physical AI systems sense real-time conditions across distributed physical assets—radios, power units, cooling systems, and edge equipment—and autonomously adjust their operating states at the edge. Decisions are executed locally within safety bounds, enabling energy-efficient operation of large-scale physical networks without compromising reliability.

ISSUE/OPPORTUNITY

Telecom networks consume significant energy, driving both operating costs and carbon exposure. Static energy management wastes power during periods of low demand, as cellular base stations, data centers, and network equipment often run at full capacity regardless of real time traffic. Edge-deployed IoT sensors, vision-enabled monitoring systems, and wearable devices used by field technicians provide granular visibility into equipment utilization, environmental conditions, and on-site activity. Network operators face growing pressure to meet regulatory sustainability targets while controlling energy costs that represent a major share of operational spend. Physical AI leverages data from edge devices, IoT sensors, and vision-based systems to dynamically adjust network operations in real time—scaling power usage up or down based on actual demand patterns—while maintaining service quality and regulatory compliance.

Tags: Operations | Physical AI

HOW PHYSICAL AI CAN HELP

Edge-level physical arbitration
Physical AI resolves conflicts between competing objectives (energy, performance, safety) locally—e.g., deciding whether to throttle, sleep, or reroute power when conditions degrade.

Physical degradation-aware control
AI adapts energy behavior based on equipment age, wear, and thermal stress, reducing long-term physical damage—not just short-term energy use.

Sensor-driven micro-actuation
Continuous feedback from temperature, vibration, and load sensors enables fine-grained physical adjustments (fan speeds, power draw, cooling flow) rather than coarse system-wide controls.

Human supervision
Operations remain controlled by network engineers who set policies, review AI recommendations, and maintain override authority to ensure service commitments are met.

Sustainable network operations

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Physical AI that autonomously adjusts power states, cooling flows, and operating parameters across distributed network infrastructure must perform reliably under all conditions—including equipment aging, unexpected load spikes, and sensor degradation. Incorrect energy adjustments can compromise service quality or accelerate physical wear on the very assets the system is designed to protect.

Responsible and accountable

Network engineers theoretically retain policy-setting and override authority; however, as AI handles increasing volumes of continuous adjustments, the practical ability to review individual decisions is diminished. Organizations must define clear boundaries between autonomous edge actions and decisions requiring engineer review, maintaining clear audit trails to support accountability when service or equipment outcomes are disputed.

Transparent and explainable

Network engineers setting energy policies need to understand how physical AI balances competing objectives—energy efficiency, service quality, and equipment protection—and why specific adjustments were made. Without this visibility, engineers cannot calibrate policies effectively, identify misconfigured optimization objectives, or maintain true oversight of systems making continuous physical changes.

POTENTIAL BENEFITS

Extended asset lifespan

Energy decisions informed by physical stress signals reduce premature component failure across equipment fleets.

Reduced human intervention in physical tuning

Engineers shift from manual parameter tuning to policy-level oversight as AI handles continuous physical adjustments.

Reputational protection

Environmental risk is mitigated as operators demonstrate measurable progress on sustainability goals, protecting brand reputation among environmentally conscious customers and stakeholders.

Scalable cross-site consistency

The same physical control logic applies across factories, networks, data centers, or plants—despite different layouts and equipment mixes.

Data center operations

Anticipatory control of data center environments

Physical AI systems sense real-time conditions across distributed physical assets—Physical AI systems may forecast hardware degradation and thermal stress in data centers and coordinate maintenance actions with workload migration to reduce failure risk.

ISSUE/OPPORTUNITY

In the TMT sector, data centers are complex physical systems where localized issues can quickly cascade across compute, power, and cooling infrastructure. Traditional planning and monitoring struggle to capture the dynamic interplay between workload intensity, thermal behavior, and hardware aging, often reacting only after hotspots, performance degradation, or failures impact service availability. Static thresholds frequently miss early stress signals as servers age, cooling systems lag demand, and power components experience uneven loads.

Data center digital twins overlay facilities design, power and cooling systems, and network architecture to simulate "what if" scenarios, enabling operators to anticipate thermal, capacity, and performance impacts of build out or configuration changes before making physical interventions.

Physical AI enables a shift to predictive, infrastructure aware operations by continuously analyzing physical signals to anticipate risk earlier. This helps enable allows operators to proactively migrate workloads, rebalance thermal loads, or schedule targeted maintenance—under strict governance models that keep humans in control—improving resilience, uptime, and energy efficiency at scale.

Tags | Information Technology | Physical AI

HOW PHYSICAL AI CAN HELP

Thermal and stress prediction

AI models forecast temperature and load conditions that increase failure risk by analyzing patterns in cooling system performance, compute workload distribution, and environmental factors that create thermal stress on physical infrastructure.

Maintenance coordination support

AI assists in aligning physical interventions with operational constraints, helping planners schedule hardware replacement, cooling system maintenance, and infrastructure upgrades during periods that minimize business impact.

Human escalation pathways

Important decisions remain with data-center operators, helping to ensure that knowledge and accountability stay with qualified personnel who understand business context and can assess trade-offs that AI systems cannot fully evaluate.

Hardware degradation forecasting

Signals from physical components are used to anticipate end-of-life events, identifying at-risk servers, storage devices, and network equipment. AI assists operators rather than acting autonomously, providing recommendations and analysis while humans retain authority over all actions that could affect availability, performance, or risk exposure.

Guardrailed decision support

AI assists operators rather than acting autonomously.

Scenario evaluation

Potential actions are evaluated for risk and impact before execution, modeling the consequences of workload shifts, maintenance windows, or equipment changes to not create new problems while addressing identified risks.

Data center operations

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Physical AI that forecasts hardware degradation and thermal stress must perform reliably across the dynamic interplay of workload intensity, cooling behavior, and equipment aging. A system that misses early stress signals may allow conditions to develop until failures force emergency shutdowns, negating the predictive value that justified deployment.

Safe and secure

Data centers are critical infrastructure whose AI-assisted operations span compute, power, and cooling systems—a broad attack surface where compromised control logic could trigger cascading failures. Manipulated thermal predictions or maintenance scheduling could cause outages affecting dependent services. Security protections must match the operational criticality and customer service obligations of these facilities.

Responsible and accountable

Humans retain control of all consequential actions—workload migrations, maintenance scheduling, and infrastructure changes—with AI providing recommendations for operator review. However, this human-in-the-loop approach must be adhered to in practice. Consistent documentation of AI recommendations, human decisions, and outcomes is essential when service disruptions or equipment failures require investigation.

POTENTIAL BENEFITS

Reduced unplanned outages

Higher availability through earlier identification and proactive intervention before hardware failures or thermal events force emergency shutdowns or trigger cascading failures across interconnected infrastructure.

Asset longevity

Improved lifecycle management by reducing exposure to stress conditions that accelerate hardware degradation, extending useful equipment life and reducing premature replacement costs.

Operational resilience

Better preparedness for incidents through predictive analysis that helps enable advance planning, resource staging, and coordinated response strategies rather than reactive crisis management when failures occur.

Semiconductor manufacturing orchestration

Fleet-level coordination of fab robots

Physical AI systems coordinate collaborative robot fleets across semiconductor fabrication facilities and adjacent downstream operations (e.g. testing, packing, and quality inspection), dynamically assigning tasks such as transport, inspection, and handling based on production needs.

ISSUE/OPPORTUNITY

Semiconductor fabrication facilities operate with hundreds of robots performing important tasks, yet traditional systems assign each robot to fixed functions regardless of real-time production needs. Transport robots move wafers between processing stations, inspection robots check for defects, and handling robots load and unload equipment, each following predetermined routes and schedules. When production priorities shift or equipment becomes available ahead of schedule, fixed robot assignments create bottlenecks as idle robots in one area cannot assist with backlogs elsewhere. Facilities purchase additional robots to help enable sufficient capacity for peak demands in each function, driving capital costs higher than necessary if robots could be reassigned dynamically.

Quality requirements demand strict contamination control and precise handling, constraining how robots can be redeployed without risking yield loss. In downstream test/pack areas, throughput can fluctuate with shift coverage and attendance, creating avoidable output dips even when equipment capacity exists. Robots and humans work side-by-side at 50/50 ratios during training phases. Humans teach procedures during day shifts; robots maintain operations independently during night shifts and workforce gaps.

The opportunity is fleet-level orchestration that dynamically assigns robots to the highest-priority tasks while respecting cleanroom constraints and quality standards, maximizing throughput with fewer total robots.

Tags: **Manufacturing & Quality** | **Physical AI**

HOW PHYSICAL AI CAN HELP

Fleet-level reasoning
AI assigns tasks dynamically across the robot fleet, directing available robots to the highest-priority activities based on current production status and bottleneck locations.

Reduced fixed roles
Robots are not task-locked, enabling transport robots to assist with inspection or handling when those functions become bottlenecks, increasing overall fleet utilization.

Throughput-aware coordination
Actions align with production goals, prioritizing movements that accelerate wafers through rate-limiting process steps rather than following rigid predetermined schedules.

Quality preservation
AI respects fab constraints including contamination zones, handling protocols, and equipment compatibility requirements that protect semiconductor yield and reliability.

Human governance
Engineers supervise orchestration, setting production priorities, defining quality constraints, and maintaining oversight of robot assignments to help enable fab safety and yield targets.

Remote enablement
AR-based assistants deliver step-by-step, hands-free fab tool maintenance guidance, enable real-time multilingual troubleshooting using live equipment context, and connect engineers to remote experts.

Semiconductor manufacturing orchestration

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Fleet-level AI orchestration of robots must perform reliably under real-world production conditions, including equipment changes, yield excursions, and unpredictable robot availability. Orchestration failures that send robots to incorrect zones, create handling conflicts, or violate contamination protocols can directly affect wafer yield in a manufacturing environment where quality failures are extremely costly.

Safe and secure

Integration of AI robot orchestration into semiconductor production and contamination control systems creates a large attack surface where compromised logic could direct robots to violate contamination zones or bypass quality checks. Security protections must reflect the yield and commercial consequences of fabrication integrity being compromised through unauthorized access or manipulation.

Responsible and accountable

Engineers supervise orchestration by setting priorities and defining constraints. However, dynamic fleet orchestration generates too many individual assignment decisions for engineers to review each one. Organizations must clearly define which assignment decisions can be executed autonomously within pre-approved boundaries and which require human validation before execution to maintain meaningful oversight.

POTENTIAL BENEFITS

Lower capital intensity

Fewer robots required as dynamic assignment enables smaller fleets to handle the same production volumes by eliminating idle capacity in underutilized functional areas.

Operational efficiency

Resources are better utilized as robots spend more time performing value-adding activities and less time idle, increasing return on expensive automation investments.

Quality protection

Standards are maintained through AI enforcement of contamination protocols, handling procedures, and equipment compatibility rules that prevent quality excursions from dynamic reassignments.

Higher throughput

Bottlenecks are reduced through intelligent resource allocation that concentrates robot capacity where production flow is currently constrained, improving wafer cycle times.

Ultra-low-latency network protection

Enabling safety-critical instant response through device intelligence

In telecom networks, next generation physical AI embeds GPU based intelligence directly at the edge—within base stations, network equipment, cameras, and sensors—enabling real time threat detection without relying on centralized processing. By analyzing signals locally in milliseconds, these systems can identify physical and network threats such as site tampering, unauthorized access, equipment anomalies, or service impacting attacks, and autonomously take immediate protective actions—isolating affected assets, triggering failover, or adjusting operating parameters—before issues propagate across the network, while keeping human operators in the oversight loop.

ISSUE/OPPORTUNITY

In telecom environments, critical safety and security decisions must be made in milliseconds to prevent physical damage, service disruption, or harm to personnel. For example, AI enabled wearables worn by field technicians at cell sites can continuously monitor proximity to live equipment, restricted zones, or hazardous conditions and trigger immediate alerts or automatic equipment shutdowns if unsafe behavior is detected—without relying on network connectivity. Extending this model to the network itself, Physical AI embedded directly into base stations, edge devices, and security sensors can autonomously detect threats such as site tampering, unauthorized access, or abnormal signal behavior and take instant protective actions, including isolating affected network elements, denying access, or activating failover paths. By eliminating transmission latency and enabling on device decision making, telecom operators can contain threats at the edge before they cascade across interconnected networks, while retaining human oversight for escalation and compliance.

Tags | Information Technology

HOW PHYSICAL AI CAN HELP

On-camera AI processing
Lightweight AI models run directly on smart cameras, processing video streams locally and making threat detection or safety decisions without any network transmission, achieving response times under 10 milliseconds—enabling intervention that is essentially instantaneous.

Instantaneous alert and action systems
On-device processing enables immediate triggering of alerts, alarms, or automated safety responses (e.g., crane shutdowns, access denials) at machine-speed rather than network-limited speeds, ensuring protective actions occur fast enough to prevent incidents.

Device-embedded threat recognition
In telecom networks, AI embedded in edge devices such as site cameras, base stations, and access sensors detects physical and network threats in real time. By processing data locally, these systems can autonomously isolate affected assets, deny access, or trigger failover—containing threats instantly without waiting for centralized analysis, while retaining human oversight.

Ultra-low-latency network protection

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

AI making autonomous protective decisions within milliseconds must perform with extremely high reliability. A false positive that shuts down a cell site or denies access to a field technician creates operational disruption; a false negative that misses a genuine threat allows the incident it was designed to prevent. Both failure modes carry immediate operational consequences.

Responsible and accountable

Autonomous protective actions executed at machine speed leave no time for human review before execution, making pre-deployment governance critical. Organizations must define precisely which actions the system can take autonomously, under what conditions, and with what constraints—with human operators retaining authority over escalation, investigation, and recovery following an autonomous intervention.

Transparent and explaianable

Human operators in the oversight loop must understand what triggered an autonomous protective action—what threat was detected, what evidence supported it, and what the system did—to assess appropriateness, investigate false positives, and refine detection policies. Systems acting without explainable reasoning cannot be meaningfully supervised even when humans nominally retain oversight authority.

POTENTIAL BENEFITS

Life-saving response times

Eliminating transmission latency enables AI systems to detect and respond to threats or safety violations fast enough to prevent incidents, potentially saving lives in security and industrial safety applications where milliseconds matter.

Scalable safety infrastructure

Distributed intelligence on individual devices avoids the bandwidth and processing bottlenecks of centralized systems, enabling organizations to deploy comprehensive safety monitoring across large facilities without infrastructure constraints or centralized processing limitations.

Network-independent operation

On-device processing ensures critical safety systems function even during network outages or connectivity issues, maintaining protection under all conditions without dependence on network infrastructure availability.

Robotic quadrupeds for stadium operations and sports broadcasting

Agile, AI-driven mobility for live events and complex venues

Physical AI-powered robotic quadrupeds are being deployed in sports venues to support stadium operations and live broadcasting. These mobile robots combine sensors, cameras, and remote-control capabilities to navigate physical environments, assist security teams with monitoring, and capture stable, dynamic visual content for live sports coverage without disrupting on-ground activities.

ISSUE/OPPORTUNITY

Large sports venues face operational challenges related to real-time monitoring, safety oversight, and immersive fan engagement. Security teams often need early visibility into crowded or hard-to-access areas without increasing human risk.

Simultaneously, broadcasters seek innovative ways to capture engaging visuals while maintaining stability and reliability in dynamic environments.

Physical AI systems present an opportunity to extend human capabilities through mobile robotic platforms that operate on the ground, interact safely with surroundings, and deliver real-time visual intelligence and coverage.

Tags: Operations | Physical AI

HOW PHYSICAL AI CAN HELP

Intelligent mobile surveillance
AI-enabled perception allows robotic quadrupeds to navigate venue perimeters, monitor movement patterns, and provide continuous visual feedback, supporting preventive monitoring and operational awareness across large physical spaces.

Computer vision–based situational detection
Computer vision models process live video feeds to identify unusual activity or objects, enabling faster alerts and informed decision-making while maintaining human oversight of critical actions.

Remote operation and assisted autonomy
On-device processing enables AI-supported control systems to help operators guide robots through complex terrain, relay audio-visual information, and perform initial assessments before human teams intervene.

Edge based execution for reliability
Core perception and mobility decisions run locally on the robot, enabling continued operation even under network congestion or partial connectivity loss.

Robotic quadrupeds for stadium operations and sports broadcasting

MANAGING RISK AND PROMOTING TRUST

Private

Quadrupeds with cameras and sensors in public sports venues continuously capture video covering large numbers of spectators who have not consented to robotic surveillance. Organizations must define clear policies on what is recorded, how long footage is retained, what prevents repurposing for individual tracking, and how this monitoring is disclosed to event attendees.

Safe and secure

Quadruped robots navigating densely crowded stadium environments alongside spectators, staff, and athletes must operate safely at all times. Navigation failures or loss of control in a crowd could cause serious injury. Safety boundaries and emergency stop capabilities must be validated under live event conditions, including post-match crowd surges and unexpected human-robot contact.

Fair and partial

Computer vision models trained on specific crowd types or venue configurations may perform less accurately in different circumstances—generating disproportionate false positives for individuals whose appearance differs from training data norms. In security monitoring applications, biased detections create reputational risk for venues and event operators.

POTENTIAL BENEFITS

Enhanced operational visibility

Mobile robots extend monitoring reach across stadium environments without increasing direct human exposure.

Richer audience experience

Dynamic, ground-level visuals add engaging perspectives while maintaining broadcast stability and reliability.

Improved safety support

Early situational awareness enables informed responses before deploying personnel into uncertain conditions.

Contacts



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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



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

www.deloitte.com/us/AIInstitute

Deloitte.

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.

Copyright © 2026 Deloitte Development LLC. All rights reserved.

Endnotes

1. "Overview," Urban Development, The World Bank, accessed September 7, 2023.
2. US Government Accountability Office (GAO) analysis of Federal Procurement Data System, June 2025.
3. Karen Pollitz, Justin Lo, Rayna Wallace, Salam Mengistu, "Claims Denials and Appeals in ACA Marketplace Plans in 2021," KFF, February 2023.