Deloitte.

The Consumer AI Dossier

A selection of high-impact use cases

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

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.

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.