Deloitte.

Together makes progress

The Physical AI Dossier

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

GLOBAL AI & EMERGING MARKETS

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

Foreword

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

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

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

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

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

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

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

Nitin Mittal
Global AI Leader
Deloitte LLP

Jim Rowan
Deloitte US, Head of AI
Deloitte Consulting LLP

Sam Roddick
Chief Strategy Officer, Global Consulting Services
Deloitte Global

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

The Consumer Physical AI Dossier

Summary: The Consumer Physical AI Dossier

Physical AI is reshaping how goods move, how stores operate, and how consumers experience the world around them

Consumer companies operate across one of the broadest and most varied physical footprints of each industry—supply chains spanning continents, distribution networks measured in thousands of locations, retail environments that should be staffed and maintained at scale, and an end point that is the most unpredictable physical environment of all: the home. This breadth has historically meant high labor dependency, significant operational variability, and limited ability to maintain consistency across touchpoints.

Physical AI addresses each of these structural characteristics directly. Where consumer operations are labor-intensive and repetitive, autonomous systems can take on physical tasks more reliably and at greater scale. Where quality and compliance depend on human observation across hundreds of locations, AI-powered vision can provide continuous and consistent oversight that manual auditing cannot. Where consumers expect faster, more responsive service, AI-coordinated physical systems can help companies meet those expectations without proportional increases in cost.

The consumer sector also presents some of Physical AI's most demanding deployment conditions. Consumer environments are dynamic and unpredictable: stores rearrange, delivery routes change constantly, and homes are entirely unique. Unlike industrial settings where Physical AI can be deployed in controlled, structured conditions, consumer applications must function reliably in human-scale environments. This raises the bar significantly on robustness and adaptability.

Consumer trust adds another layer. As Physical AI becomes more visible to end consumers (in stores, in delivery interactions, and increasingly in the home) the standards for transparency and privacy are set not just by regulators but by consumers themselves. Companies that earn that trust may find Physical AI to be a source of lasting competitive advantage.

Autonomous transport for urban mobility services (1/2)

AI-driven mobility in unpredictable urban environments

DESCRIPTION

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.

HOW PHYSICAL 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.

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.

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.

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.

Regulatory-compliant design

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

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.

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.

Lower costs

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

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.

Autonomous transport for urban mobility services (2/2)

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.

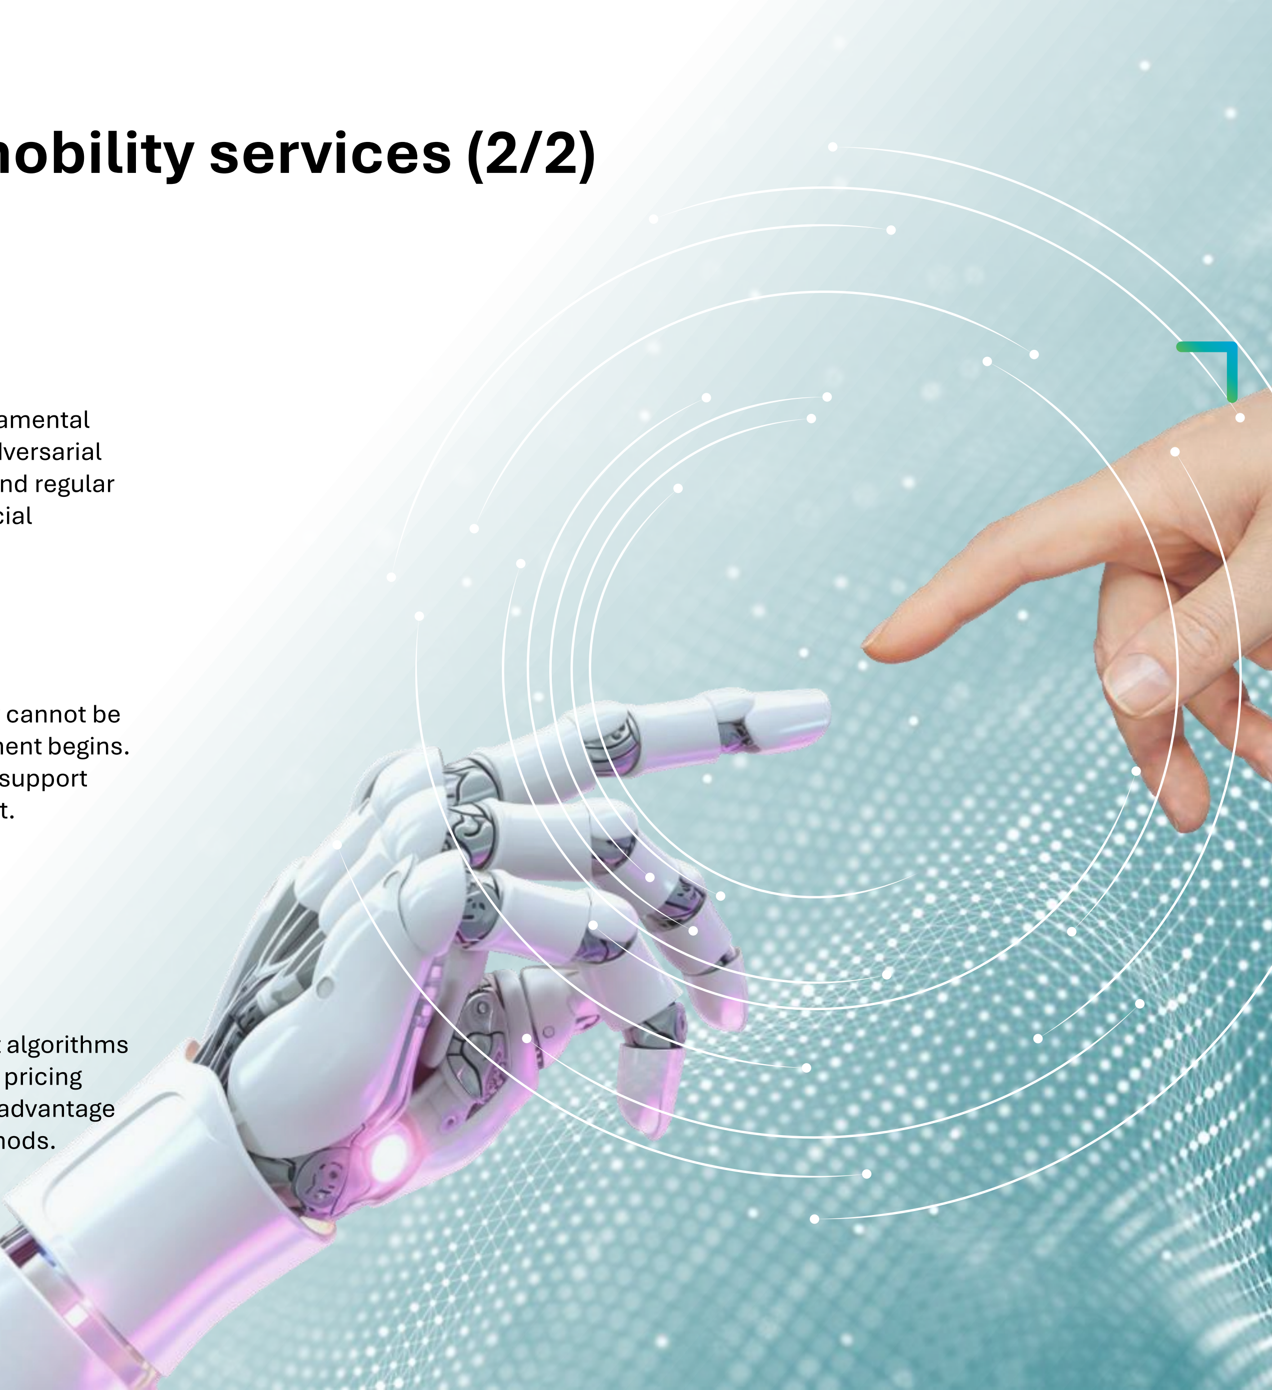

Solution: Edge + Robotics | Humanoid robots

Multipurpose household service robots (1/2)

Reasoning-enabled service robots for home environments

DESCRIPTION

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 each specific action, limiting their usefulness to narrow, pre-defined activities. Users should 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.

HOW PHYSICAL 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 each task variation.

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

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.

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

POTENTIAL BENEFITS

Ease of use

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

Broader task coverage

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

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.

Multipurpose household service robots (2/2)

MANAGING RISK AND PROMOTING TRUST

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.

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.

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.

Solution: Computer Vision

Vision-enabled store operations (1/2)

Real-time retail execution through vision

DESCRIPTION

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.

HOW PHYSICAL 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.

Context-aware planogram reasoning

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

Real-time feedback loops

Field users receive immediate guidance on corrective actions, with visual overlays showing which products need adjustment.

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.

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.

Vision-enabled store operations (2/2)

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, such as, 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.

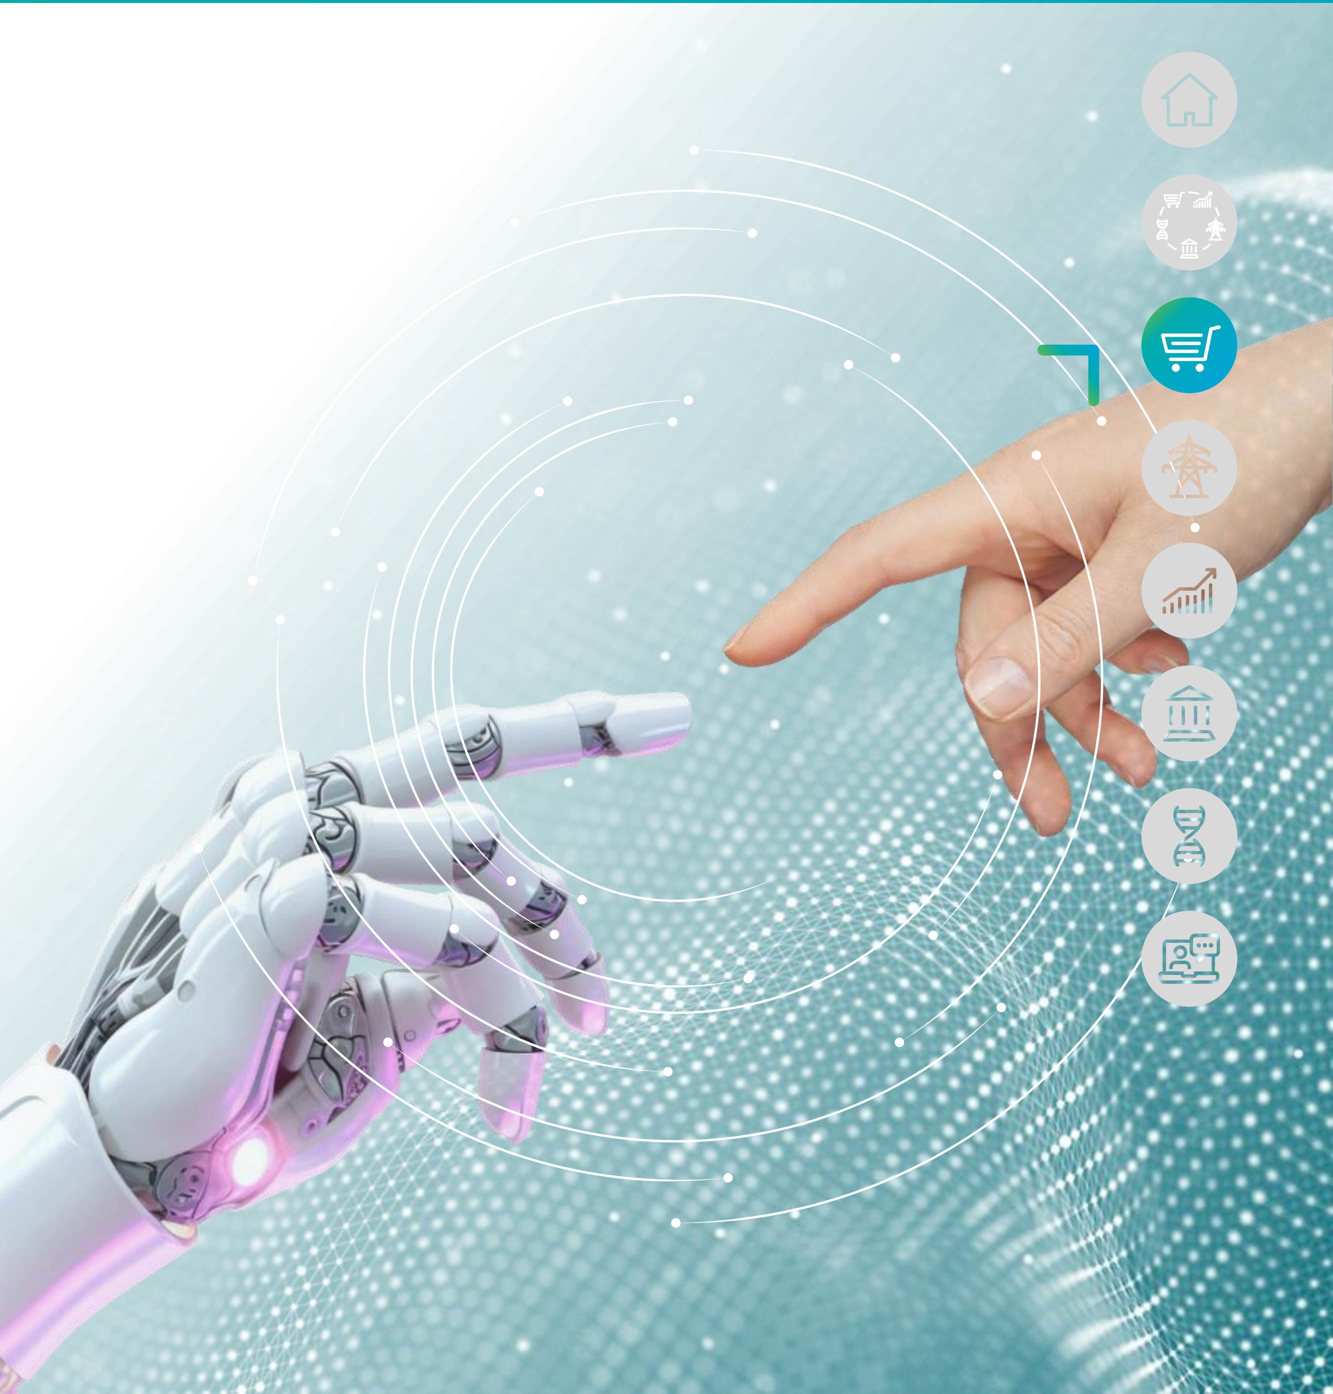

Solution: Edge + Robotics | Autonomous vehicles

Fleet telemetry and route optimization (1/2)

Adaptive logistics driven by edge intelligence

DESCRIPTION

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.

HOW PHYSICAL 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.

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.

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.

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.

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.

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.

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.

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, to help enable logistics managers to communicate accurate arrival times to retailers and maintain service commitments despite unexpected obstacles.

Fleet telemetry and route optimization (2/2)

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.

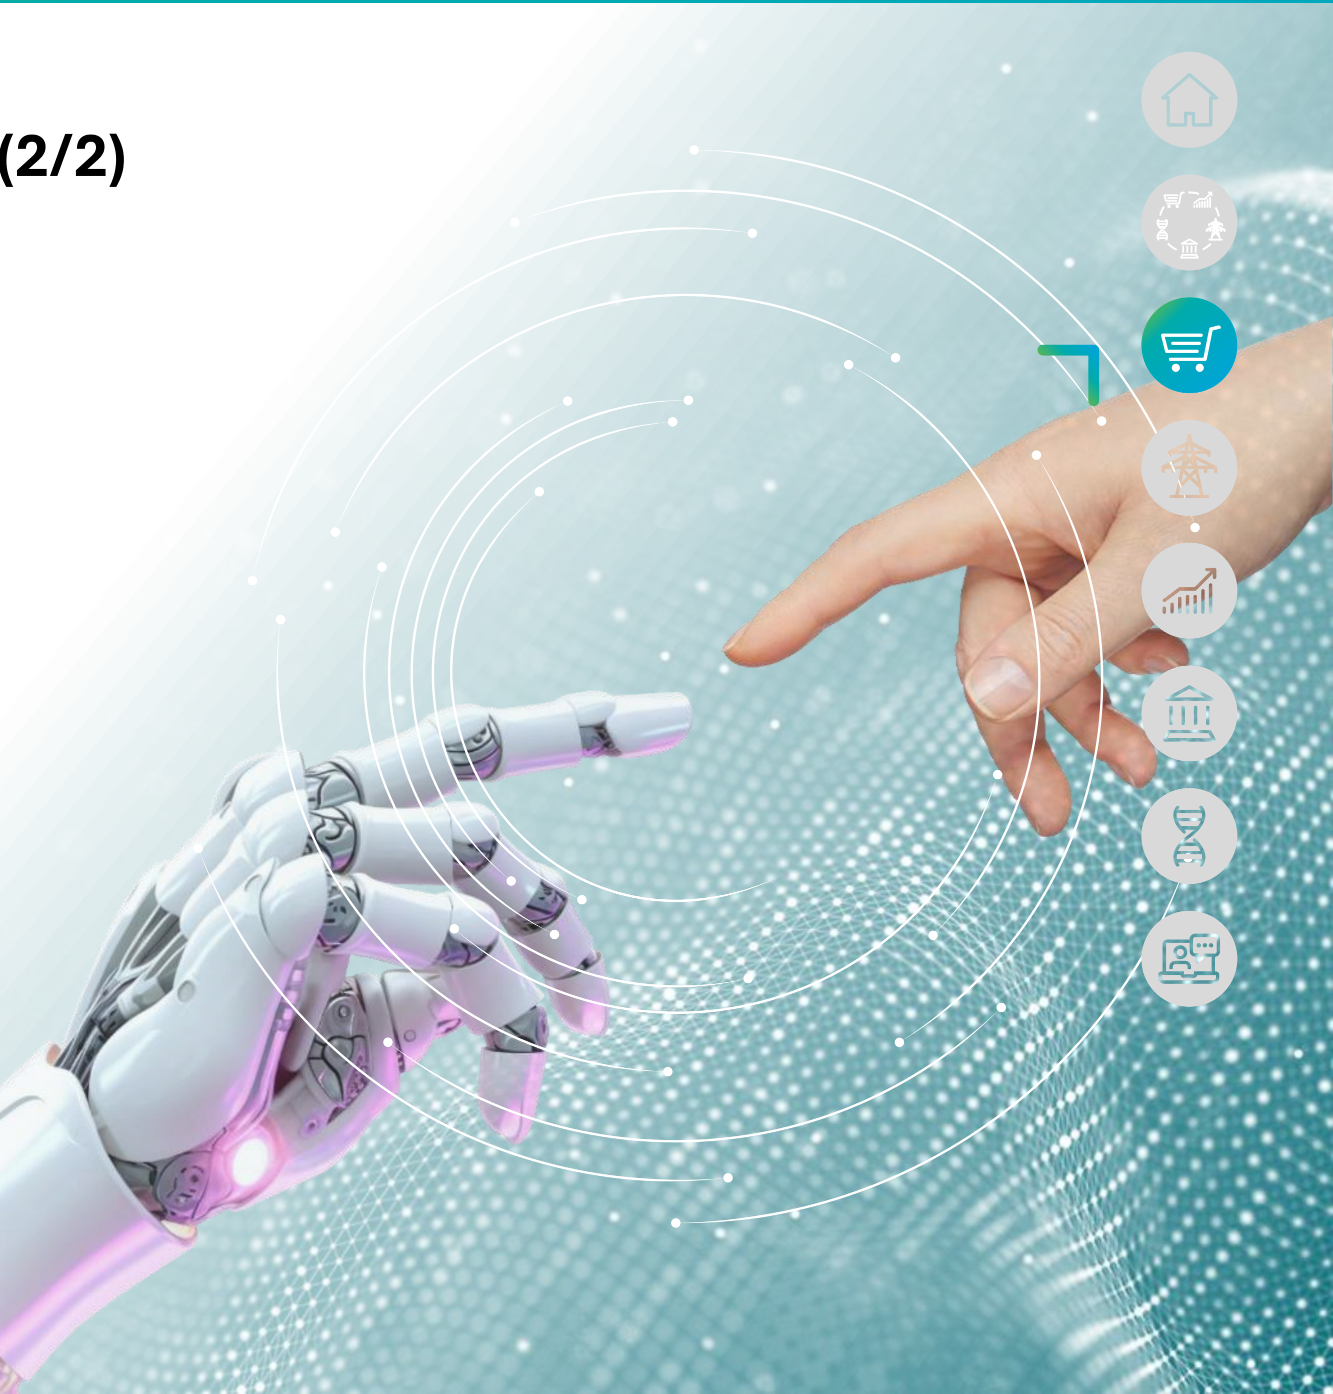

Edge–cloud architecture for consumer mobility (1/2)

Distributed intelligence enabling real-time physical action in vehicles

DESCRIPTION

An edge–cloud Physical AI architecture distributes intelligence between vehicles and centralized platforms to help 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 volumes of sensor data from cameras, LiDAR, radar, and other onboard systems while operating in dynamic, safety-critical environments. Many driving decisions should 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.

HOW PHYSICAL 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.

POTENTIAL BENEFITS

Operational resilience

Vehicles remain functional during network disruptions due to edge autonomy.Easier rollout across regions.

Lower system cost

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

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.

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.

Edge–cloud architecture for consumer mobility (2/2)

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.

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.

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.

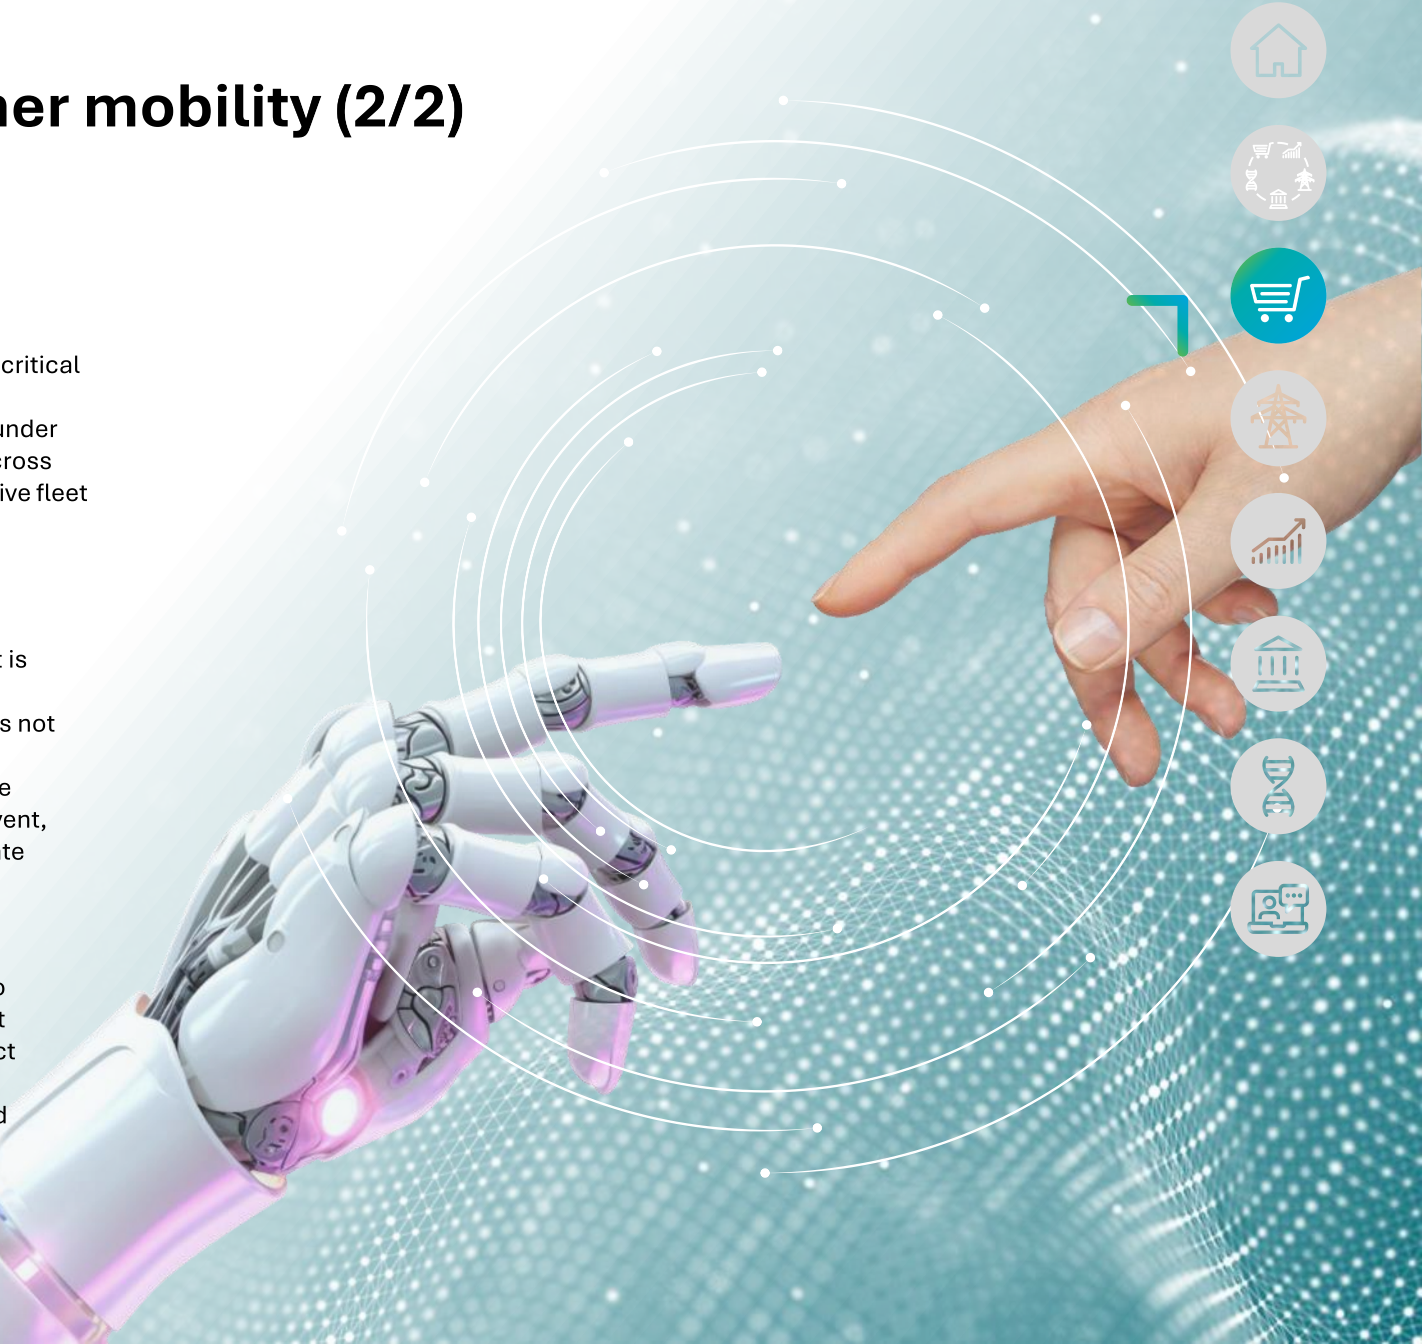

Solution: Edge + Robotics | Task-specific

Robotic stowing and picking system (1/2)

Shelf based picking and stowing in warehouses

DESCRIPTION

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 (Stock Keeping Unit), reducing ergonomic risk while expanding automation scope.

HOW PHYSICAL 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.

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.

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.

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.

Footprint-aware system design

AI supports layout optimization for limited space.

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.

Expanded automation coverage

Current systems can handle approximately 75% of SKUs based on physical characteristics, compared to much lower coverage with traditional fixed automation.[2]

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).

Improved order accuracy

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

Robotic stowing and picking system (2/2)

MANAGING RISK AND PROMOTING TRUST

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.

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.

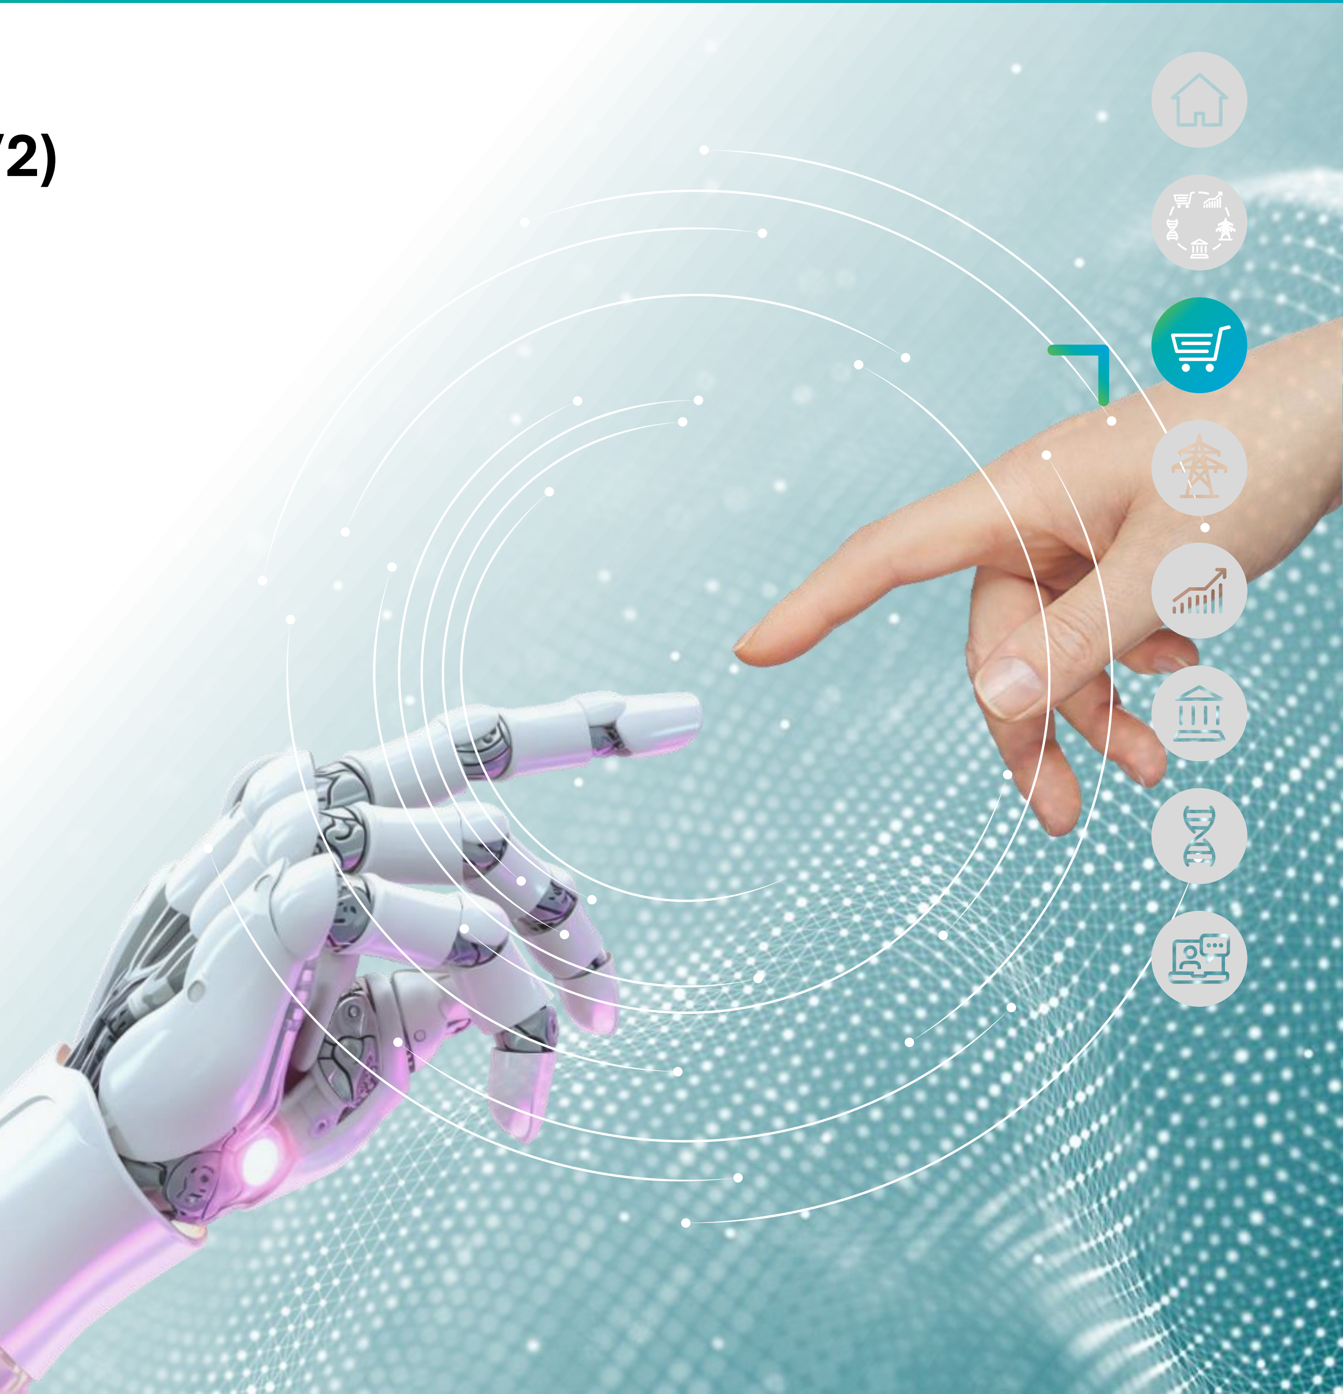

Solution: Edge + Robotics | Task-specific

Vision-enabled robotic induction for high-variability consumer logistics (1/2)

Handling SKU variability at industrial throughput

DESCRIPTION

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 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.

HOW PHYSICAL 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.

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.

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.

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.

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.

POTENTIAL BENEFITS

Higher throughput

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

Labor productivity

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

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.

Reduced ergonomic strain and injury risk

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

Vision-enabled robotic induction for high-variability consumer logistics (2/2)

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.

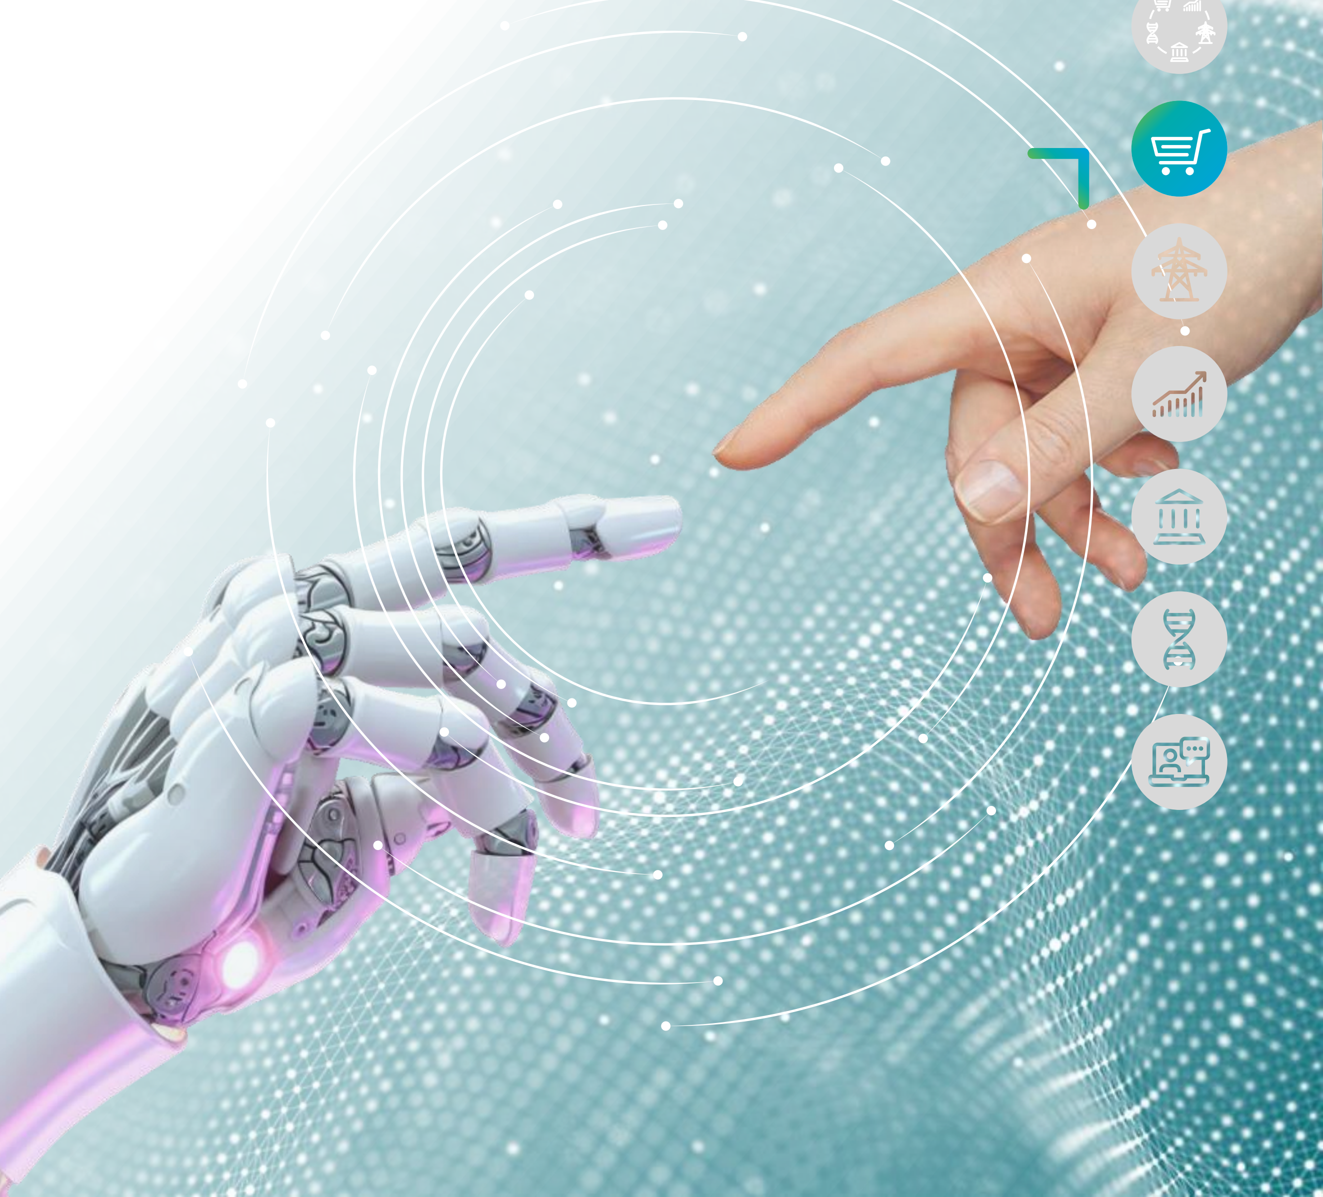

Autonomous material movement in consumer fulfillment environments (1/2)

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

DESCRIPTION

Autonomous mobile robots 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.

HOW PHYSICAL 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.

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.

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 helps enable 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.

Human-aware safety envelopes

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

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.

POTENTIAL BENEFITS

Reduced transport labor

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

Operational flexibility

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

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.

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.

Autonomous material movement in consumer fulfillment environments (2/2)

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.

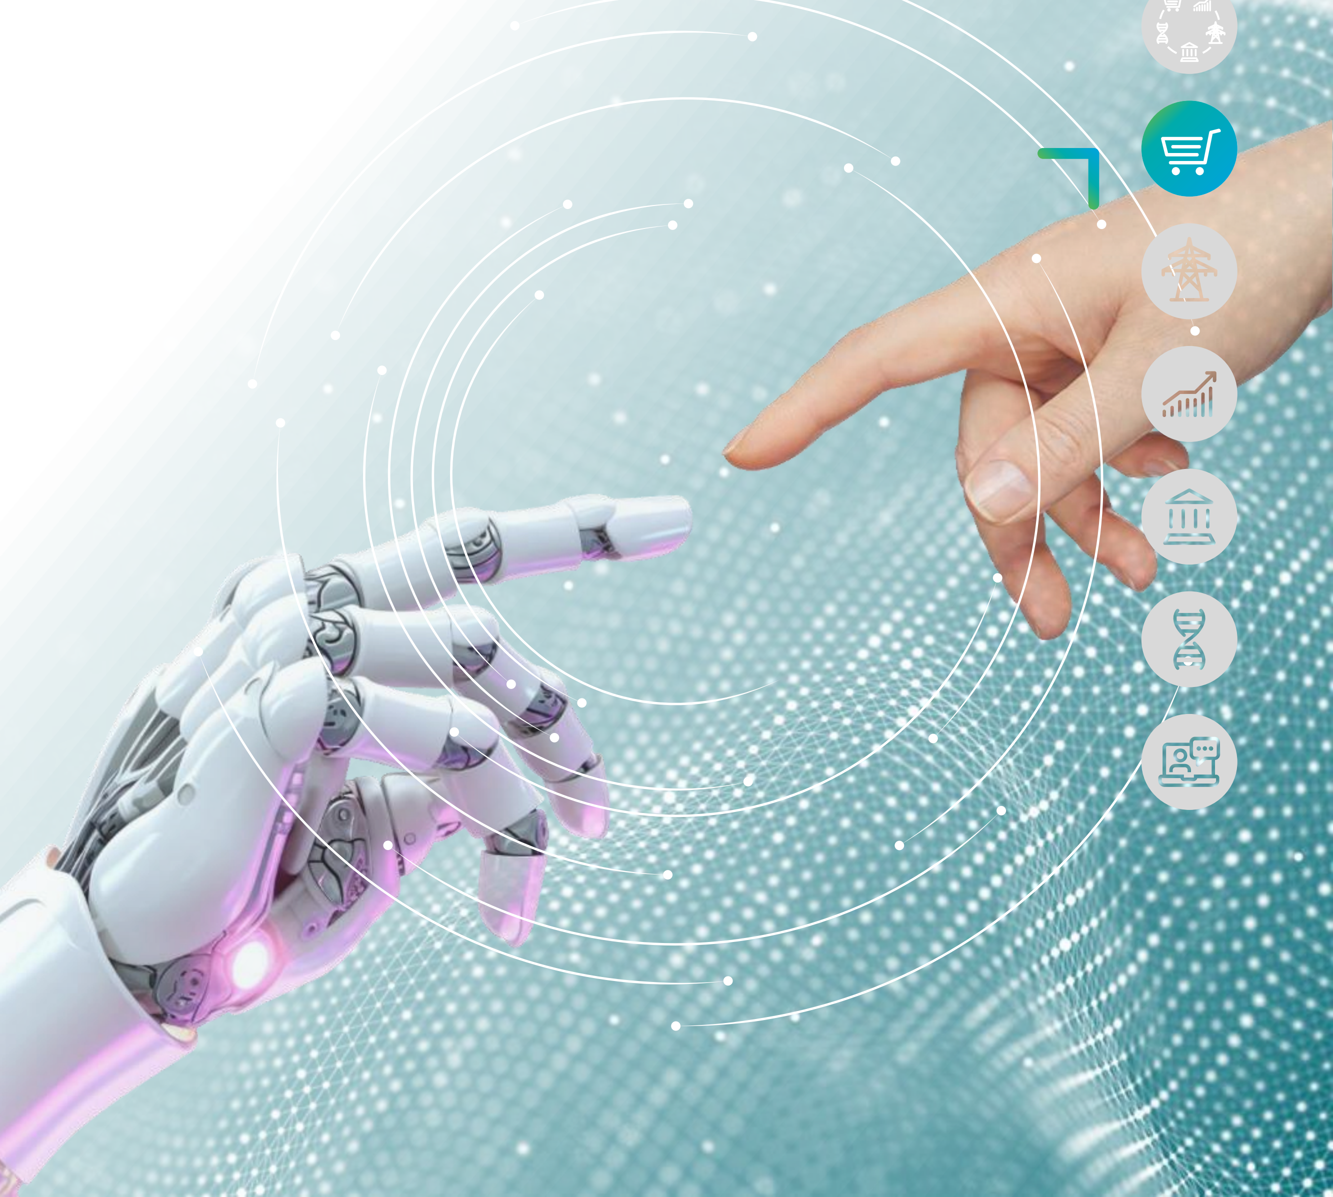

Solution: Edge + Robotics | Humanoid robots

Programmable and general-purpose robots for consumer operations (1/2)

Adaptive Physical AI systems operating across dynamic consumer environments

Description

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, enabling the same robotic platform to adapt to new tasks, environments, and workflows without physical retooling—while operating safely alongside human workers.

HOW PHYSICAL 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.

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.

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.

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.

Multi-task adaptability

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

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.

Programmable and general-purpose robots for consumer operations (2/2)

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.

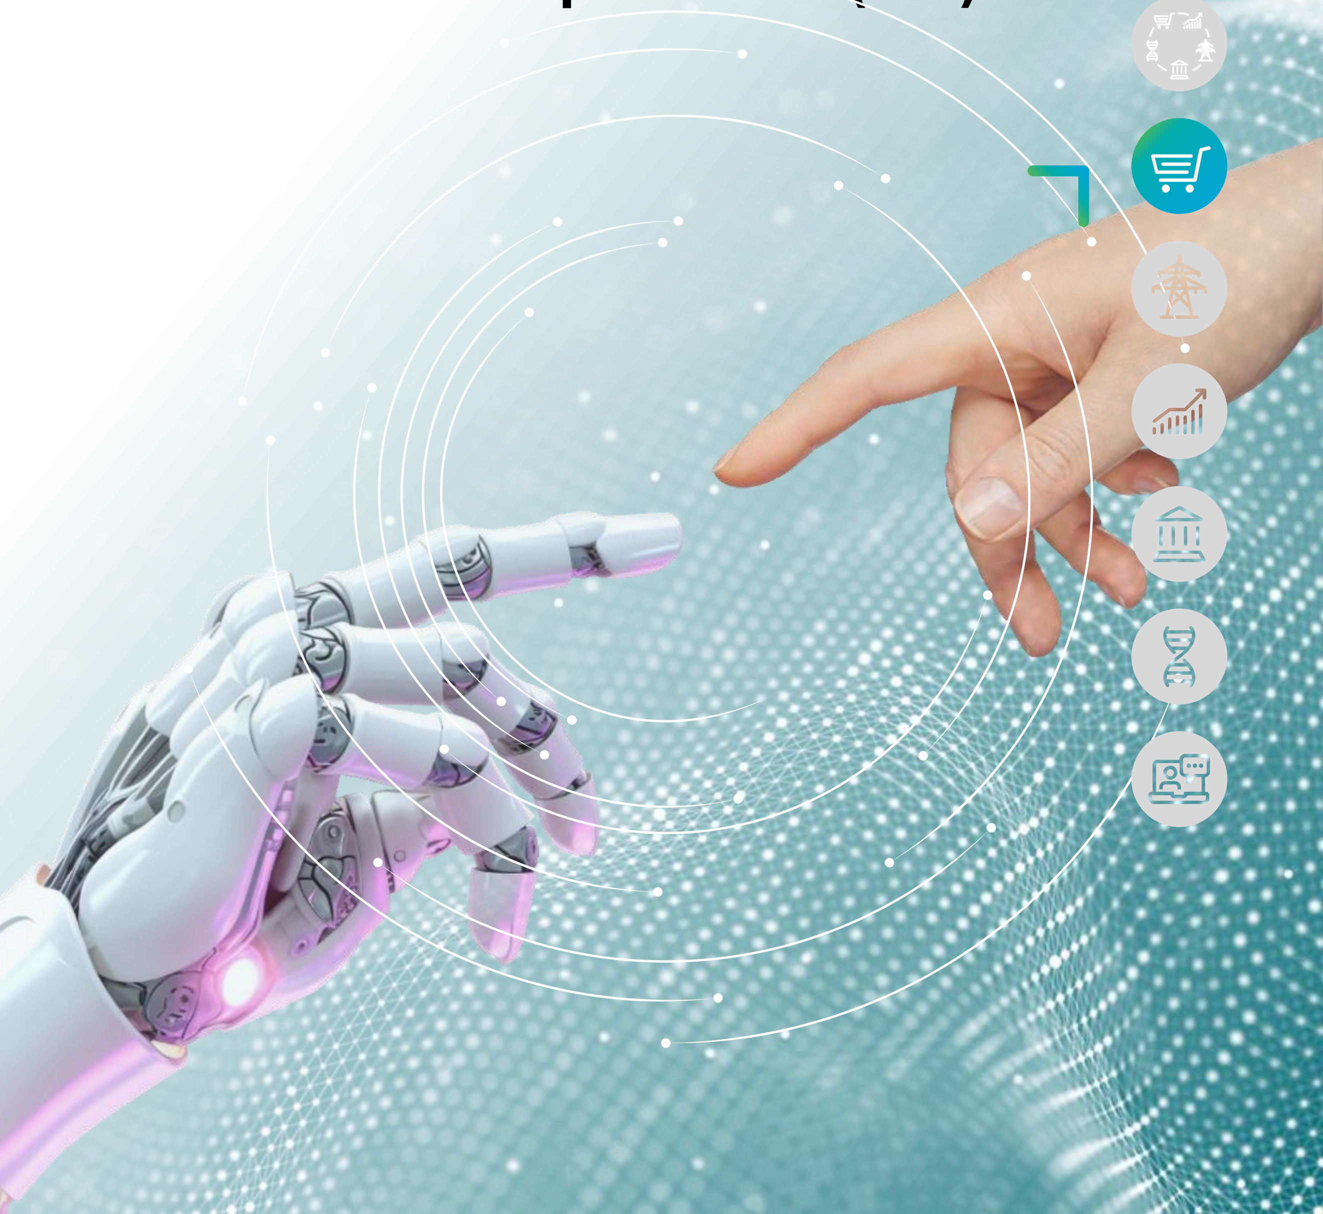

About Deloitte

End Notes

2. Physical AI in retail supply chains: From the dashboard to the warehouse floor - Automated Warehouse