Deloitte.

Together makes progress

The Physical AI Dossier

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

GLOBAL AI & EMERGING MARKETS

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 Financial Services Physical AI Dossier

Summary: The Financial Services Physical AI Dossier

For an industry where reliability is foundational and regulatory scrutiny is constant, Physical AI offers powerful capabilities—but requires equally rigorous governance

Financial services may not appear, at first glance, to be a heavily physical industry. But beneath the digital surface, financial institutions operate extensive and consequential physical infrastructure (ATM networks, data centers, server farms, trading terminals, and branch technology) on which customers, markets, and regulatory obligations directly depend. Unlike most physical assets in other sectors, this infrastructure does not degrade gradually and visibly; it can fail suddenly, with immediate and highly visible consequences for customers and counterparties alike.

This characteristic makes financial services infrastructure a particularly compelling environment for Physical AI's predictive capabilities. The value of anticipating a failure before it occurs—rather than responding after the fact—is asymmetric: prevention is dramatically cheaper than remediation, and the reputational and regulatory cost of an outage far exceeds the cost of the maintenance that would have prevented it.

At the same time, financial services is an industry where the governance dimensions of any AI deployment are unusually consequential. These institutions operate under some of the most rigorous regulatory oversight of each sector, with strict requirements around system reliability, auditability, data residency, and accountability[5]. A Physical AI system that cannot explain its predictions, trace its reasoning, or demonstrate consistent performance to a regulatory standard is not deployable in this environment—regardless of its technical capability.

This creates a discipline that, while demanding, ultimately serves the industry well. Financial institutions that build Physical AI deployments to the governance standards their regulators require should have built systems that are also more trustworthy, more auditable, and more resilient than those deployed under less demanding conditions.

Solution: Edge + Robotics | Task-specific

Predictive maintenance for IT and ATM infrastructure (1/2)

Predictive monitoring enabled by edge integration of physical machines

DESCRIPTION

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

ISSUE/OPPORTUNITY

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

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

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

HOW PHYSICAL AI CAN HELP

Telemetry-driven failure prediction

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

Asset-specific health modeling

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

Preparation of remediation actions

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

Risk-weighted assessment

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

Human-in-the-loop control

All actions remain subject to human approval, to help ensurei accountability, auditability, and regulatory compliance.

Edge-compatible inference

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

POTENTIAL BENEFITS

Reduced downtime

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

Lower operational risk

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

Maintenance efficiency

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

Predictive maintenance for IT and ATM infrastructure (2/2)

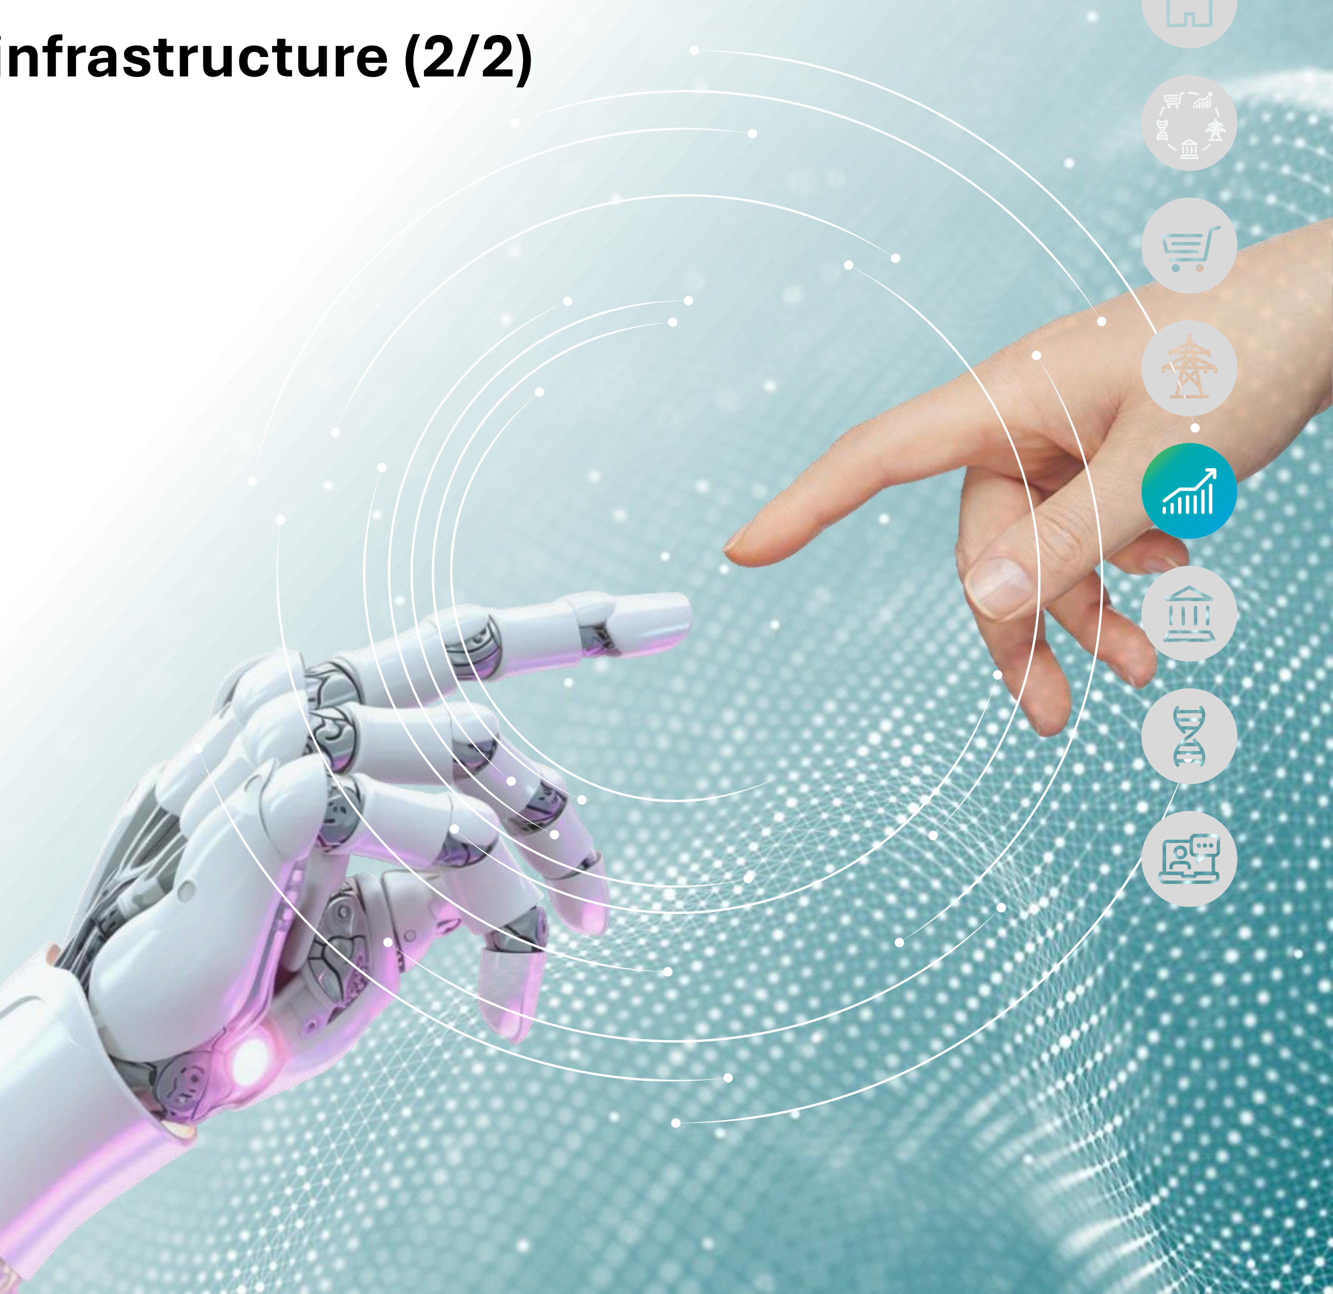

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

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

Responsible and accountable

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

Transparent and explainable

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

Solution: Edge + Robotics | Humanoid robots

Humanoid robots for branch operations (1/2)

Automating front-of-house experiences with governance

DESCRIPTION

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

ISSUE/OPPORTUNITY

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

HOW PHYSICAL AI CAN HELP

Spatial perception & crowd sensing

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

Conversational triage

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

Edge-based execution for resilience

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

Branch-aware behavior

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

Embodied escalation control

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

POTENTIAL BENEFITS

Reduced wait times

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

Lower exception load on human staff

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

Operational consistency across branches

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

Measurable physical KPIs

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

Humanoid robots for branch operations (2/2)

MANAGING RISK AND PROMOTING TRUST

Private

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

Responsible and accountable

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

Fair and impartial

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

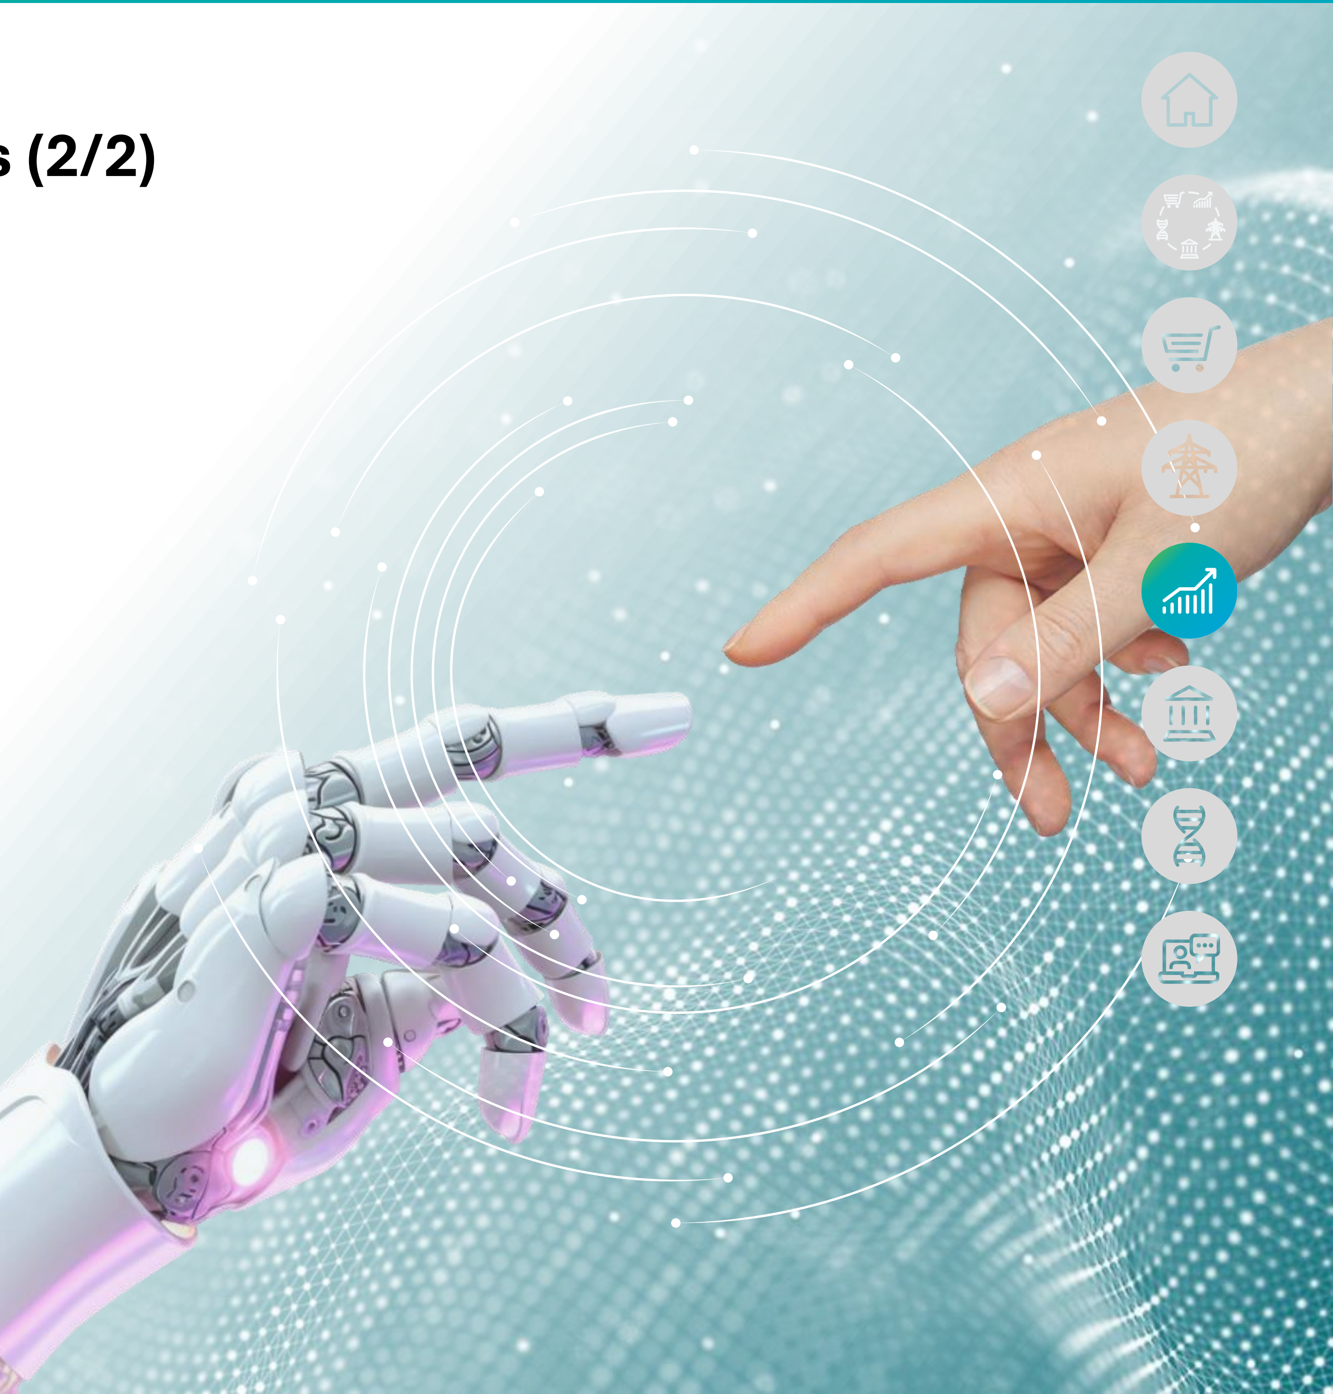

Solution: Edge + Robotics | Task-specific

ATM cash forecasting, replenishment and autonomous cash logistics (1/2)

Predictive cash availability through edge intelligence

DESCRIPTION

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

ISSUE/OPPORTUNITY

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

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

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

HOW PHYSICAL AI CAN HELP

Edge-enabled cash demand forecasting

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

Contextual signal integration

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

Dynamic replenishment recommendations

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

Risk-based prioritization

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

Logistics planning support

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

Human-controlled execution

Replenishment and logistics decisions remain under human oversight, to help ensure accountability for security-sensitive operations.

POTENTIAL BENEFITS

Fewer out-of-cash events

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

Lower logistics costs

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

Improved customer experience

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

ATM cash forecasting, replenishment and autonomous cash logistics (2/2)

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

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

Safe and secure

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

Fair and impartial

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

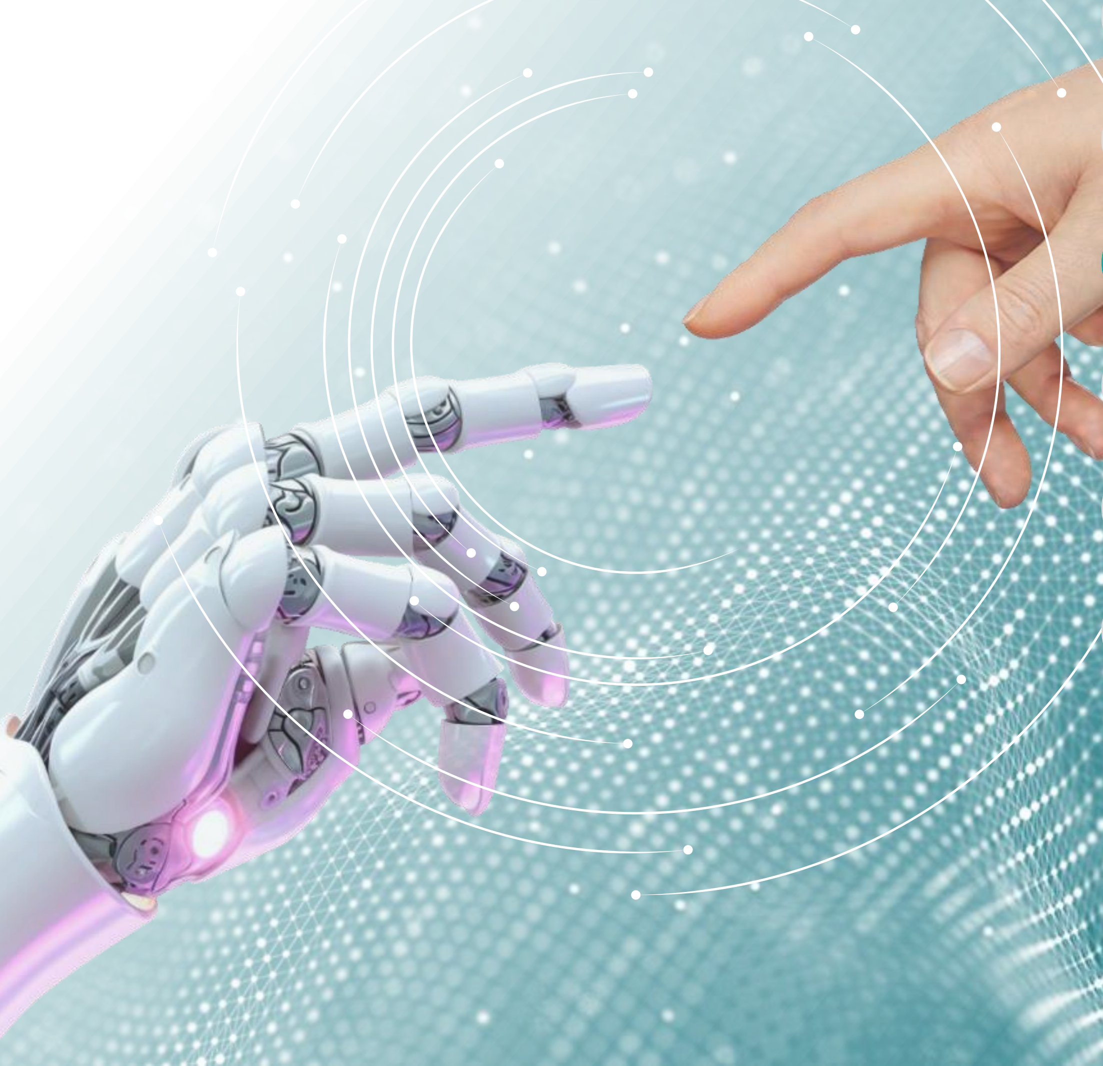

End Notes

5. Managing explanations: how regulators can address AI explainability