Deloitte.
Together makes progress

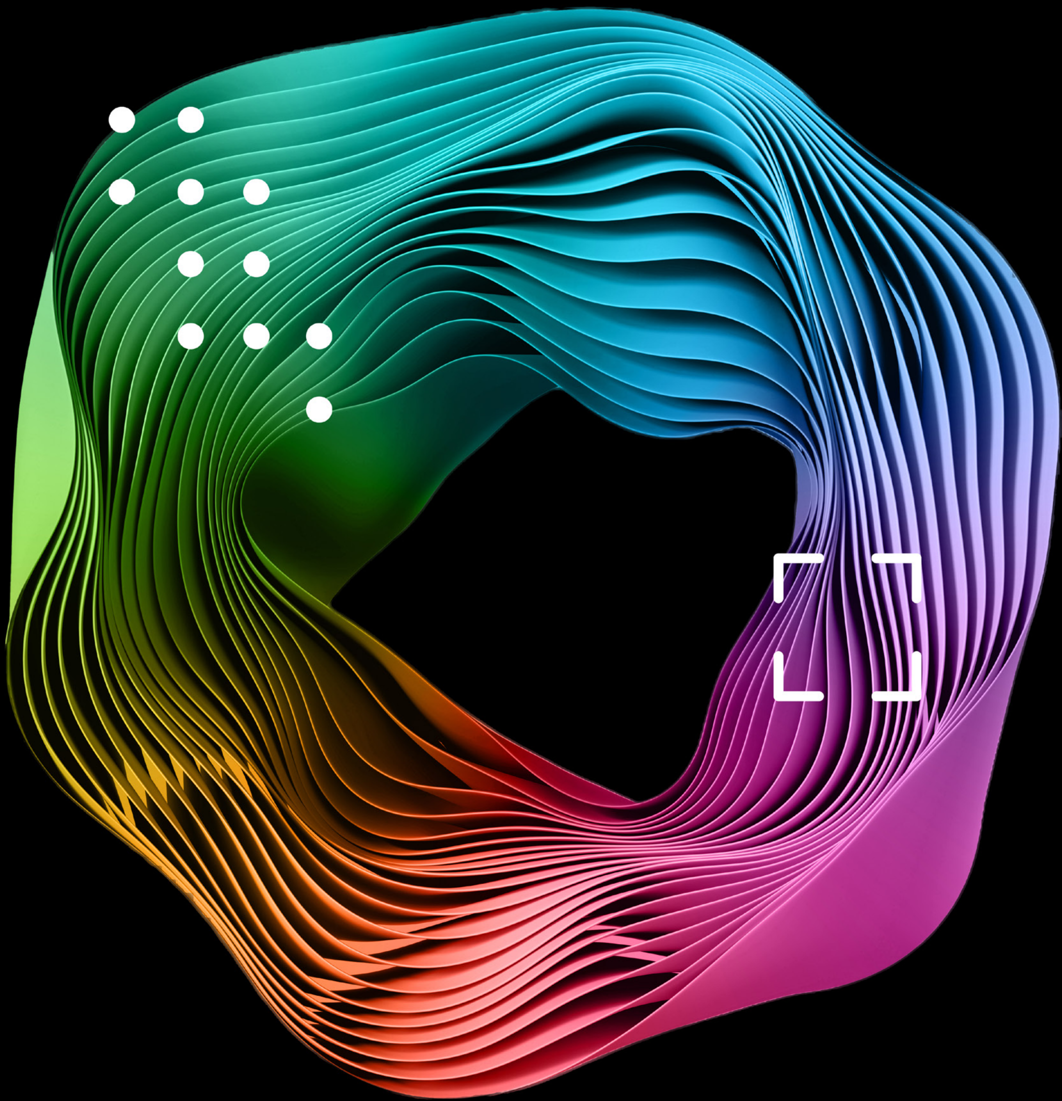

Physical AI:
The moment of acceleration

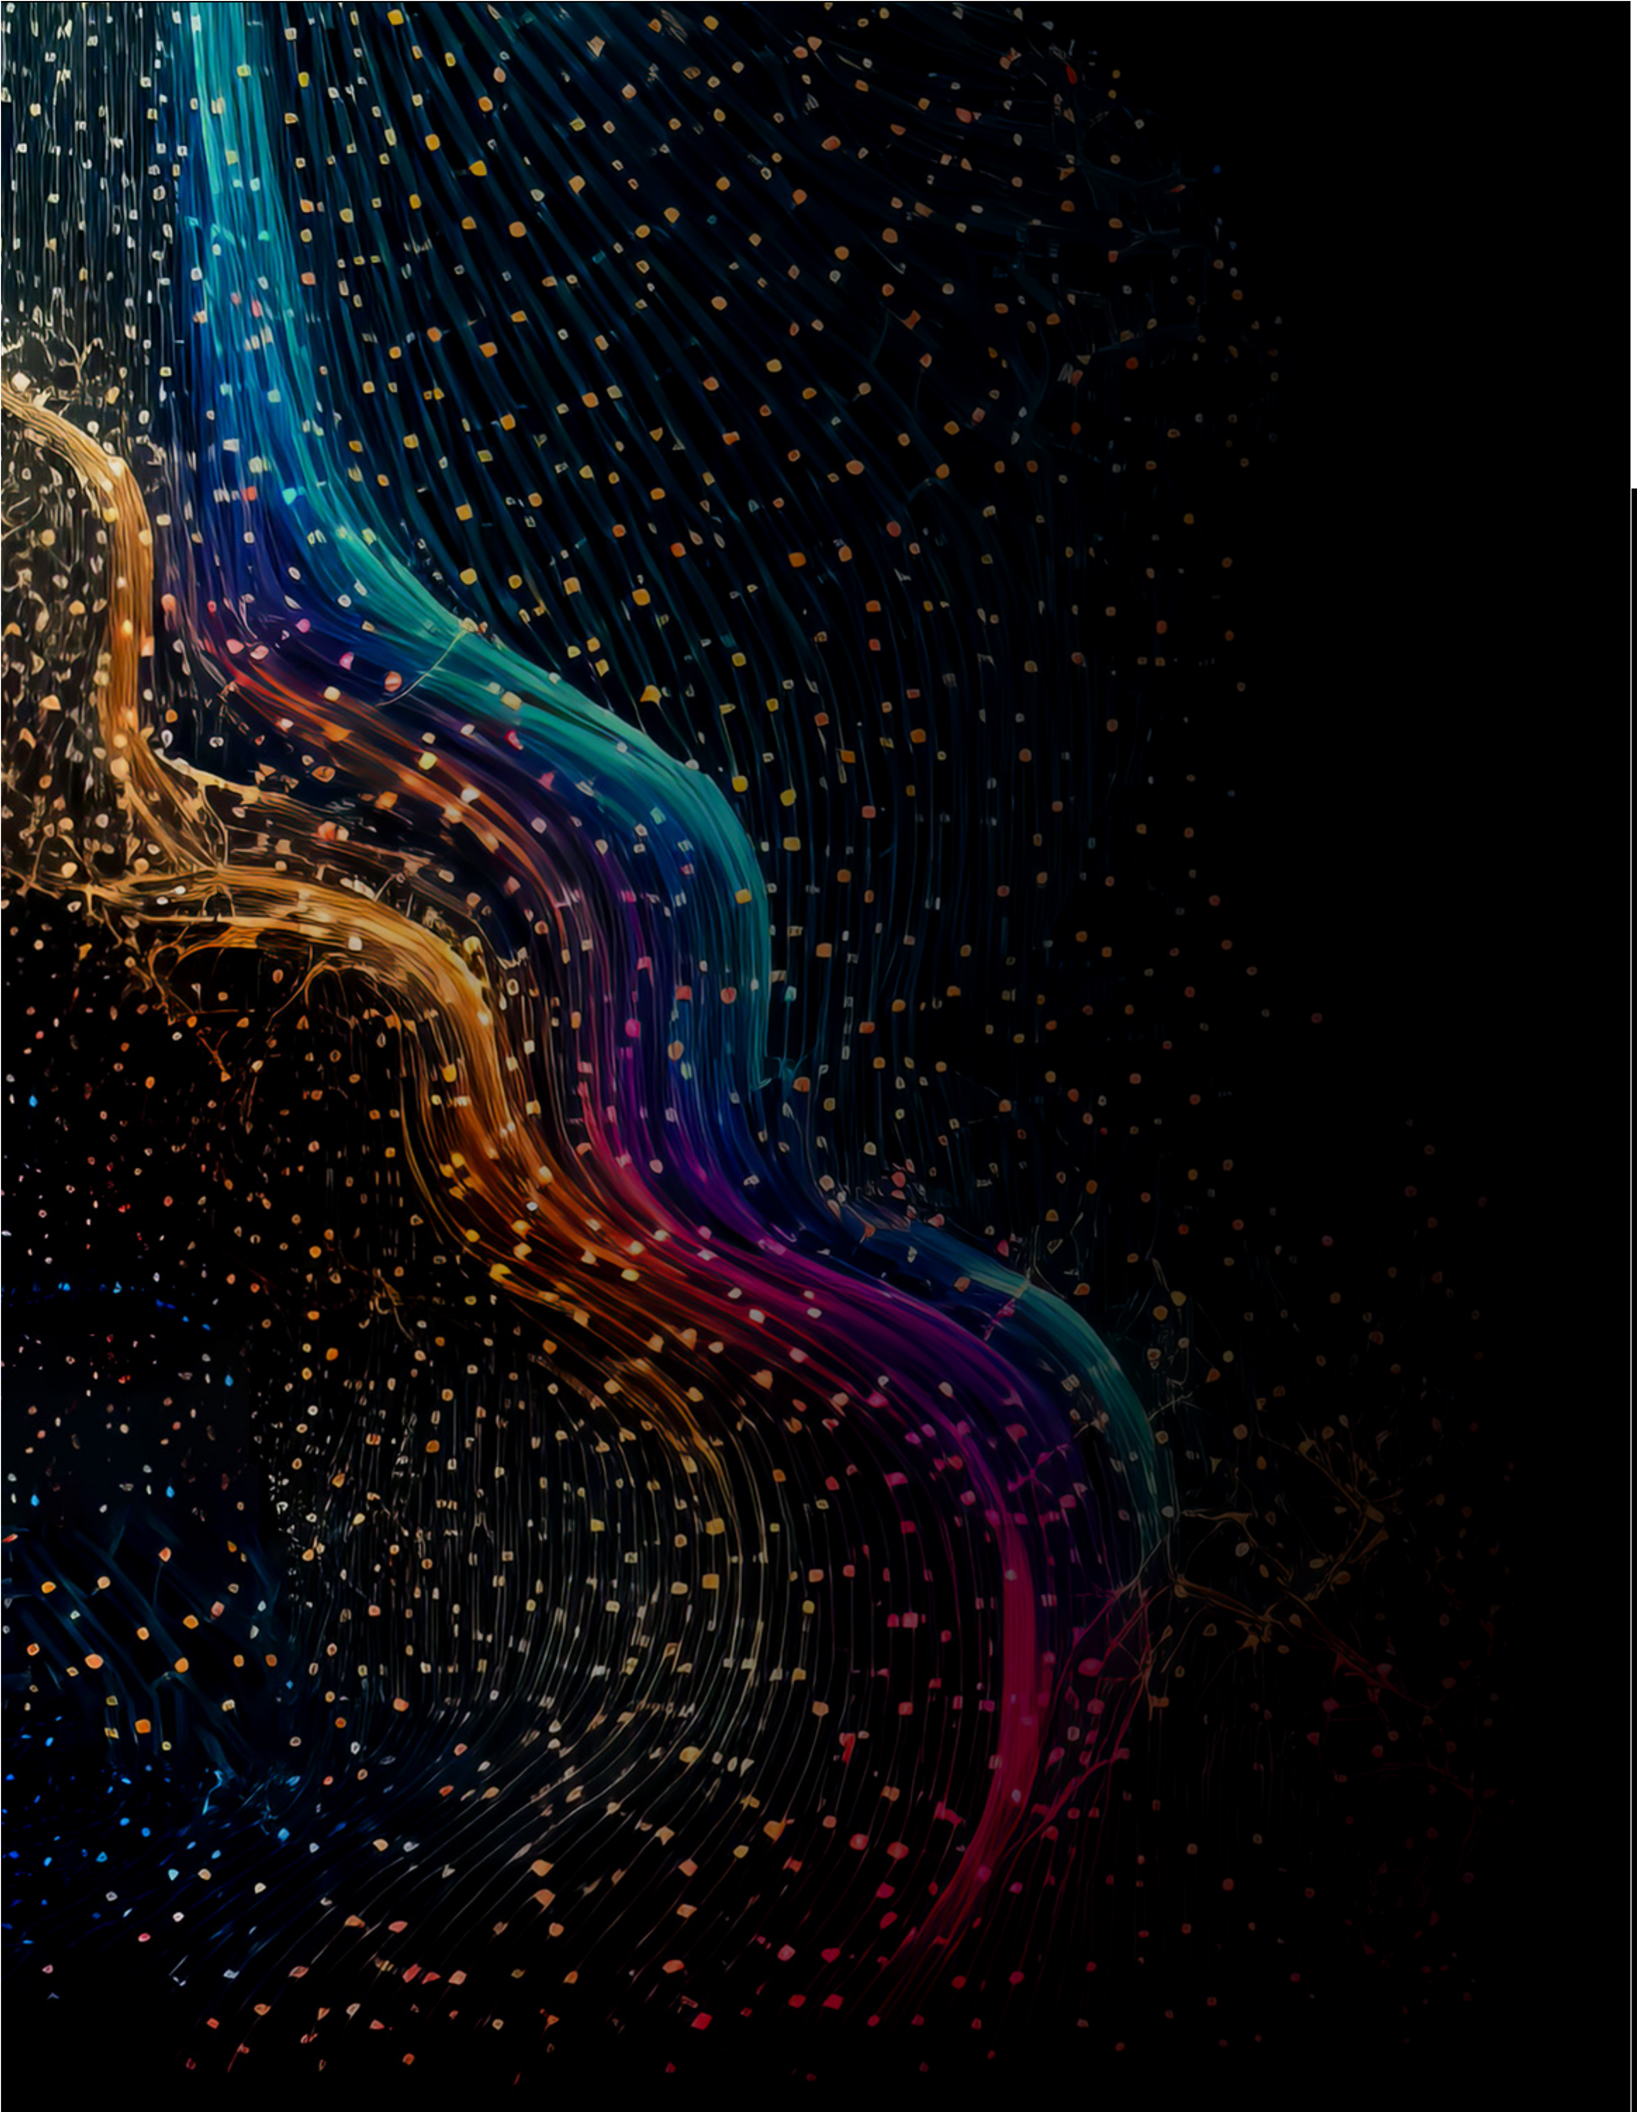

Contents

Building the future: Physical AI comes of age

Around the world, 2025 may prove to be the year when Physical AI (PAI) – the merger of physical systems with AI – definitively moved out of the realms of science fiction and into mainstream business consciousness. What has long been regarded as something for the future is now emerging as a practical reality, driven by cheaper and more capable hardware, and by software that learns how to learn.

Today, just 5 per cent of firms say Physical AI (PAI) is transforming their organisation. Within three years, 41 per cent expect it will. That gap, between current impact and future expectations, defines the story of this paper. And it has real urgency: only 3 per cent of firms have PAI extensively integrated into operations today, yet this is forecast to reach 18 per cent within two years.[1] Those who move now will not just gain an operational edge – they will build the organisational learning that shapes competitive advantage for a decade.

5% ⟶ 41%
share of firms expecting PAI to transform their organisation – today vs. within 3 years

3% ⟶ 18%
firms with PAI extensively integrated today vs. within 2 years

41%
cite cost and resource requirements as the #1 barrier to PAI adoption

This paper does two things. First, it explains why 2025–26 marks a genuine inflection point: why the technology has reached a threshold of practical viability, and why the competitive and governance environment is accelerating adoption. Second, and more importantly, it provides a structured, practical framework for business leaders navigating PAI adoption – where to start, how to sequence investment, and what organisational foundations must be in place for technology to deliver its potential. The focus throughout is on PAI and industrial robotics: the sector where value is being proven today and where the hard lessons of implementation are being learned in real time.

Physical AI: The moment of acceleration

Figure 1: Physical AI explained

Physical AI refers to AI systems embedded in or controlling physical hardware – robots, autonomous vehicles, drones, and smart manufacturing systems – that interact with and act upon the real world. Unlike purely digital AI, PAI perceives its environment through sensors, makes real-time decisions, and executes actions with tangible consequences. It bridges three technology domains to merge digital intelligence and physical reality.

The integration of perception & understanding
(the combination of "eyes" & "brain")

Vision Language Action (VLA) Model
- Vision Language Model (VLM)
- Large Language Model (LLM)
- Transformer → Diffusion → Actions

A closed loop of planning and execution
(The combination of "brain" and "hand")

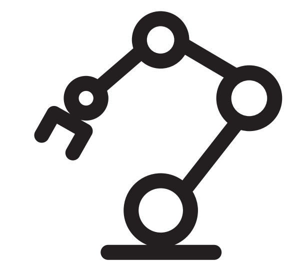

Collaborative and emergent intelligence
(The combination of "individual" and "group")

Multi-Agent Collaboration
- Coordination / Orchestration Layer
- Types
- Structures
- Strategies

AI-driven machines bridge digital intelligence and formats to bring physical intelligence to bear in a wide-range of different capabilities and environments

How can we use AI to enable robotics?

Robot
Robots can be programmed to carry out physical actions, navigate environments and manipulate objects in real world

Edge Services
Hardware at the network's entry points that locally process data reducing latency

Advanced Artificial Intelligence

Agentic AI
Agentic AI systems can fast-track decision-making without requiring constant human input

Advanced Artificial Intelligence (Cont'd)

Generative AI
Create new and original content based on algorithms trained on large datasets

Computer Vision
Key to enabling real-time perception, train of robots/machines by observing human behaviours, "map" physical spaces for simulations etc.

Machine Learning
Algorithm enable computers to learn and improve from experience

Deep learning
Neural networks model complex data patterns to detect anomalies

AI-Enabled Robot
AI-enabled robot combines physical AI with robotics and is a powerful tool with cross-industry applications

Key Form Factors

Humanoids
Humanoids robot replicate human characteristics and fucntionality to augment the workforce and increase output

Task Specific
Task specific robots have a less prescriptive form factor and are more used for repeatable process and tasks

Quadrupeds
Quadrupeds harness stability, agility, and adaptability to navigate challenging environments

Autonomous Mobile Robots (AMRs)
AMRs utilize intellignet navigation to freely move through environments, operating across diverse settings

Autonomous Vehicles
Autonomous vehicles are similar to other robots in both design and intelligence

Drones
Drones are flying machines that can enable autonomous modes through physical AI and other technologies

Why now? The convergence that changes everything

PAI's move from theoretical potential to commercial reality is the result of several simultaneous advances that have reached a critical threshold together.

Hardware cost and capability: The cost of multimodal sensors has dropped materially in recent years while precision has improved 60 per cent – expanding the range of tasks that can be automated.[2] Improved force control is enabling more complex tasks to be automated, with flexible grasping now achieving success rates above 95 per cent. Edge computing costs have followed the same commoditisation curve as industrial cameras before them.

Software that learns from simulation: AI models now train inside high-fidelity digital twins – virtual replicas of factories – and transfer learned behaviours reliably to the physical world. New open-sourced physics engines have made these training environments both more realistic and dramatically cheaper.[3] A tripartite computing architecture, spanning simulation for training, cloud for iteration, edge for execution, has emerged as the industry standard.

Open ecosystems accelerating innovation: Hugging Face reached one million users in 2025 with a $299 desktop robot, open-sourcing all hardware designs and software.[4] A growing number of open reasoning vision-language-action (VLA) models for autonomous systems are becoming available. Capability previously locked behind proprietary systems, is now available to researchers worldwide. At the same time, breakthroughs in cross-embodiment learning are advancing forward robot training and flexibility: enabling skills acquired on one robot platform to transfer to entirely different robot morphologies without retraining.[5]

These converging forces have produced measurable results at scale. Business leaders around the world are looking seriously at how PAI can be integrated into their operations. Over 500,000 industrial robots were deployed in 2024, with annual installations forecast to reach 700,000 by 2028. Collaborative industrial robots comprise a growing share, reaching almost 65,000 installations in 2024.[6] According to a Citi GPS report, there are currently around 405 million robots of all kinds in production globally – a figure projected to reach 1.3 billion by 2035.[7] And increasingly numbers of these robots will be augmented with some form of PAI.

Physical AI: The moment of acceleration

A global race: With governance catching up

Beyond technological advances, PAI is accelerating because the surrounding ecosystem is falling into place. Capital investment, hyperscale platforms, national industrial strategies and emerging standards are reinforcing one another to create sustained momentum. As with cloud and digital platforms before it, early leaders will not simply adopt AI faster, they will help shape the architectures, norms and supply chains that others will later conform to.

Capital is flowing at unprecedented scale. In the US alone, robotics-related startups raised over $10.3 billion in the first 11 months of 2025 – a year-on-year increase of 61 per cent. The US launched a $500 billion infrastructure initiative for AI computing centres; the EU committed €200 billion through InvestAI; China's 2025 Government Work Report officially designated 'embodied intelligence' as a future-oriented strategic industry.

Governance is moving in parallel. The EU AI Act comes into force in 2026, with a voluntary code of practice already encouraging early compliance. China is promoting a national top-level design, while encouraging local governments to develop differentiated industrial clusters. The ISO and IEC released the first international AI standards in 2025.[8] This shift matters. Regulation is emerging because PAI is moving from experimentation to deployment – a signal of maturity, not constraint.

PAI in the real world

For the public, perceptions of PAI have been shaped less by factories than by fiction. More recently, everyday encounters with semi-autonomous robotic vacuum cleaners and lawnmowers alongside news stories about autonomous cars and humanoid robots are shaping expectations. For leaders in manufacturing and industry, intelligence embedded into physical systems is becoming more familiar rather than exotic. This normalisation changes behaviour. It reduces organisational resistance, makes risks feel more manageable and accelerates willingness to experiment. This means, PAI is increasingly viewed as a practical extension of digital transformation into the physical world.

PAI's impact is already visible across different sectors and applications (Figure 2). In manufacturing, BYD's Xi'an plant operates at approximately 97 per cent automation – AI-guided robots and AMRs responding to live physical signals in a facility that is, in essence, a robotic factory.[9] In warehousing and logistics, PAI-enabled robots are being deployed for pick-and-pack at scale and increasingly for delivery.[10] Across China and the United States, autonomous driving and robotaxi services are now operating on public roads, with pilots and smaller scale trials in many other regions. In healthcare, robots deliver medications in hospitals, freeing nursing staff for higher-value tasks; a University of York prototype AI diagnostic robot assists with patient triage.[11]

For most business leaders, PAI's current impact remains limited. In Deloitte's State of AI Survey 2026, just 5 per cent of firms report PAI transforming their industry today – compared with 45 per cent for traditional AI and machine learning.[12] But the trajectory is unambiguous: 10 per cent of businesses expect PAI's transformational impact within the next year, and 31 per cent within three years. Uptake will be highest in consumer and life sciences and healthcare sectors (both 22 per cent), technology, media and telecommunications (18 per cent), and energy, resources and industrials (16 per cent). Right now, industrial robotics is where PAI is proving its value most concretely – and it is where the lessons that will inform all subsequent PAI adoption are being written.

Figure 2: Physical AI and robotics are performing a range of functions

Across different industries and sectors physical AI robots can deliver a wide-range of use cases based around a core set of functional capabilities

Create

What: Robotics used for the construction of manufacturing physical goods.

Example use case:
Car manufacturers using robotic arms to assemble a car

Handle

What: Interact with the world through manipulating objects or using tools.

Example use case:
A humanoid robot moves totes in a warehouse

Transport

What: Robotic vehicles to move and deliver people or goods.

Example use case:
An autonomous vehicle gives a person a ride to the airport

Observe

What: Robotics to observe, survey, or view an area.

Example use case:
Drones scan bridges and infrastructure for stability

Protect

What: Robotics for physical protection or security.

Example use case:
Drones serve as a first responders for emergencies

Interact

What: Robotics to supplement public human interaction.

Example use case:
Humanoid robots that are used in customer service or care environments

Physical AI: The moment of acceleration

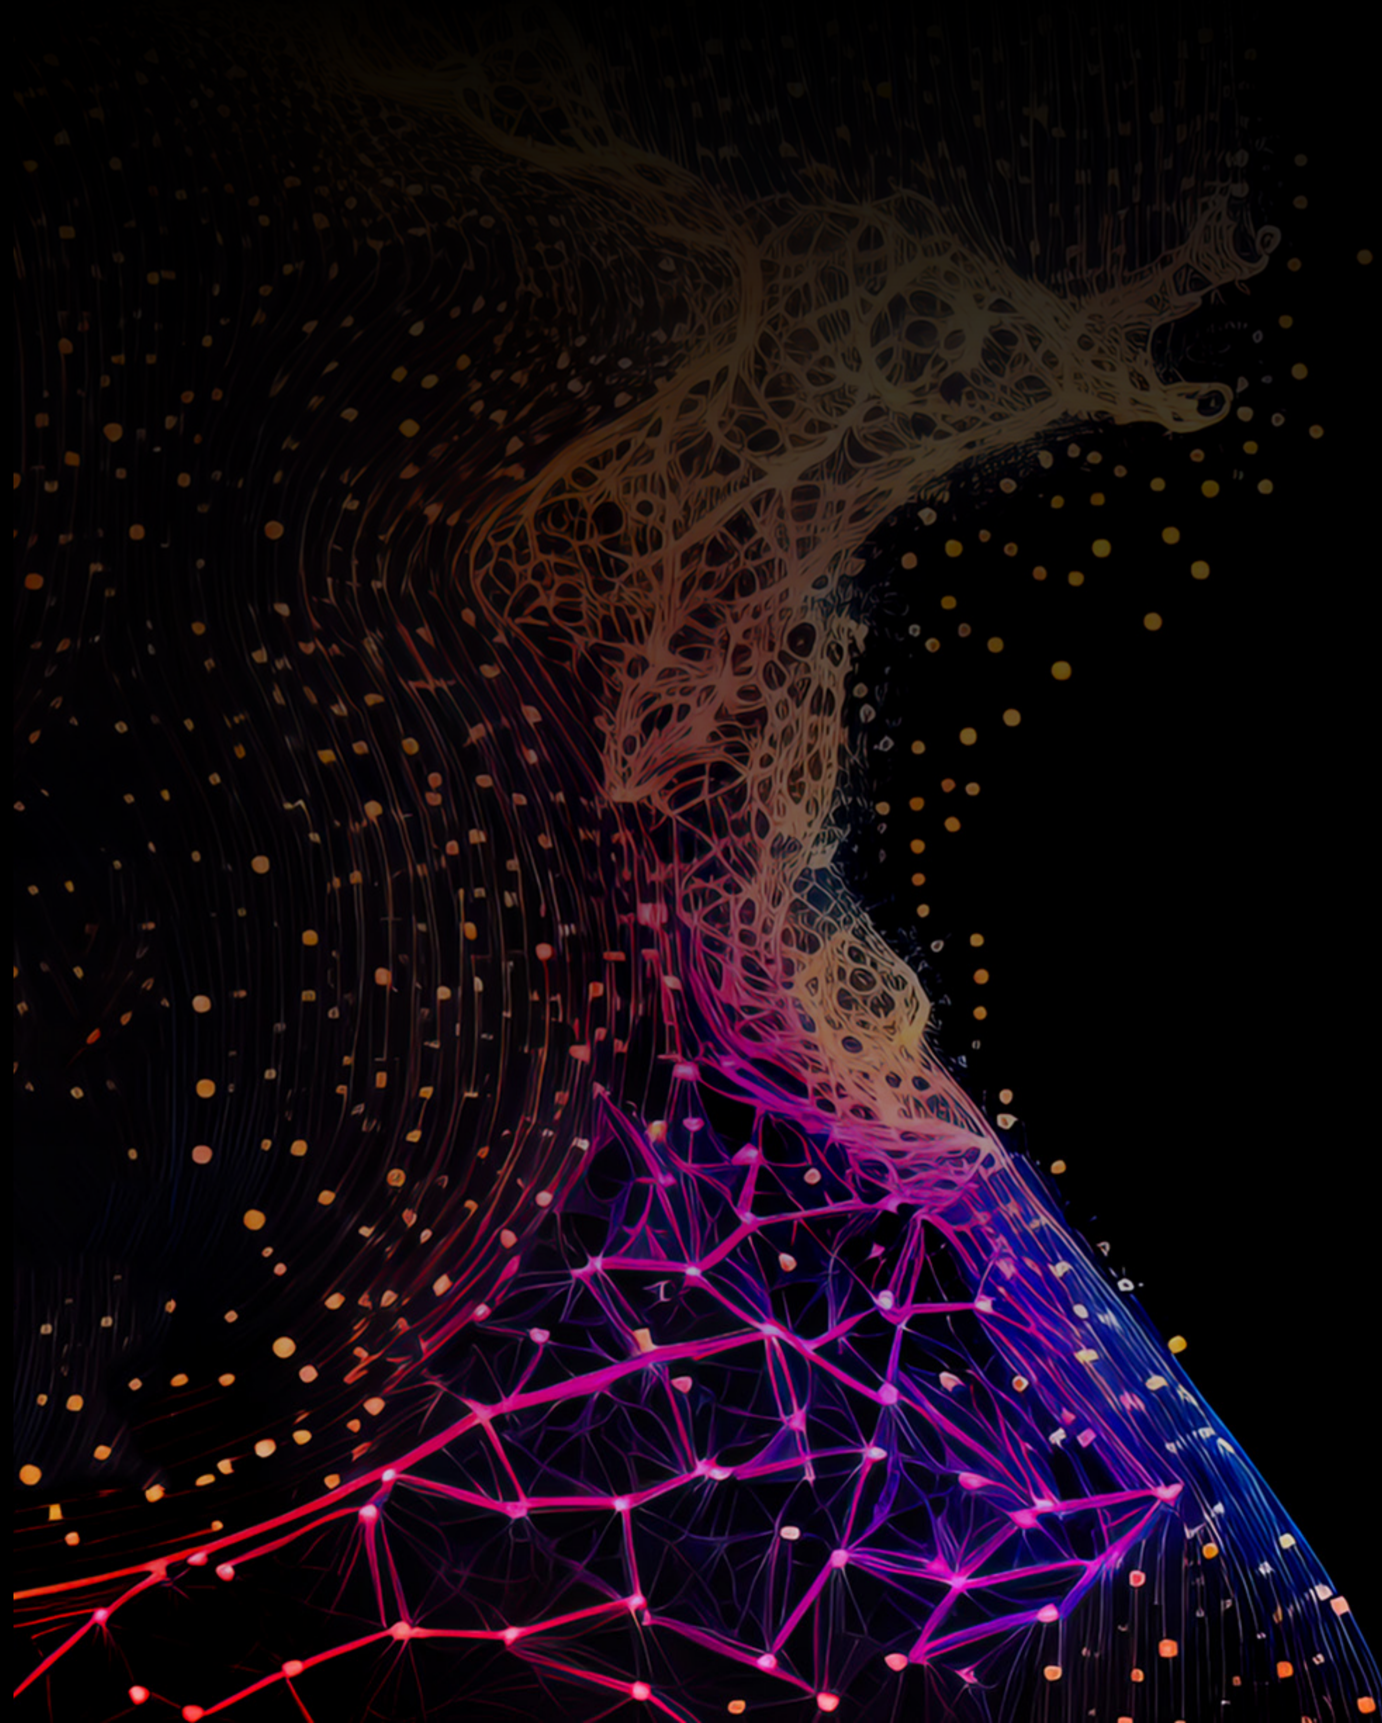

Where PAI scales value: Industrial robotics

Industrial robotics is PAI's proving ground – the environment where it must earn credibility by delivering measurable ROI under real operating conditions. Successfully deploying PAI here, where systems must learn to perceive, decide and act reliably, creates a hardened technological and operational core. However, important this foundation is, the value captured within manufacturing operations is not the end point. Instead, it provides the kinetic energy for broader transformation. What begins on the factory floor establishes a blueprint for integrating digital intelligence with physical assets, enabling systematic scaling across the enterprise and extended through the value chain.

The value radiates outward in three distinct layers (Figure 3):

Operational value: focusing on the manufacturing core

Unlike traditional robots, which rely on fixed programming and excel in tightly controlled, repetitive tasks, PAI-enabled robots do not tire, lose focus, or require manual reconfiguration. They autonomously handle predictable variation. The result: near elimination of human error in precision tasks, significant reductions in downtime, optimised resource utilisation, a new level of complex automation and a factory floor that self-optimises.

Value spillover: integration across the value chain

In R&D, AI agents prototype and test designs in digital twins, shortening new product introduction (NPI) cycles. In logistics, AI-driven robots coupled with real-time data can enable agile supply chains and dynamic inventory management. This unprecedented synchronisation from design to delivery reduces waste and human errors, enhances responsiveness, and accelerates time to market for new products.

Disruptive innovation: enabling new business models

PAI enables new production models, such as Factory-as-a-Service (FaaS), where on-demand, reconfigurable production lines can be deployed for short-run or custom manufacturing, and Operation-as-a-Service (OaaS), allowing companies to outsource the operational intelligence of a plant. These models will shift competitive advantage from physical assets to algorithmic capability. Advanced manufacturing, once capital-locked, becomes accessible on demand.

Physical AI: The moment of acceleration

Figure 3: Physical AI from operational improvement to business model evolution

Physical AI impacts will go beyond manufacturing optimisation and innovation to reshape how value flows across the supply chain and enable new business models.

3-layer business value of PAI deployment

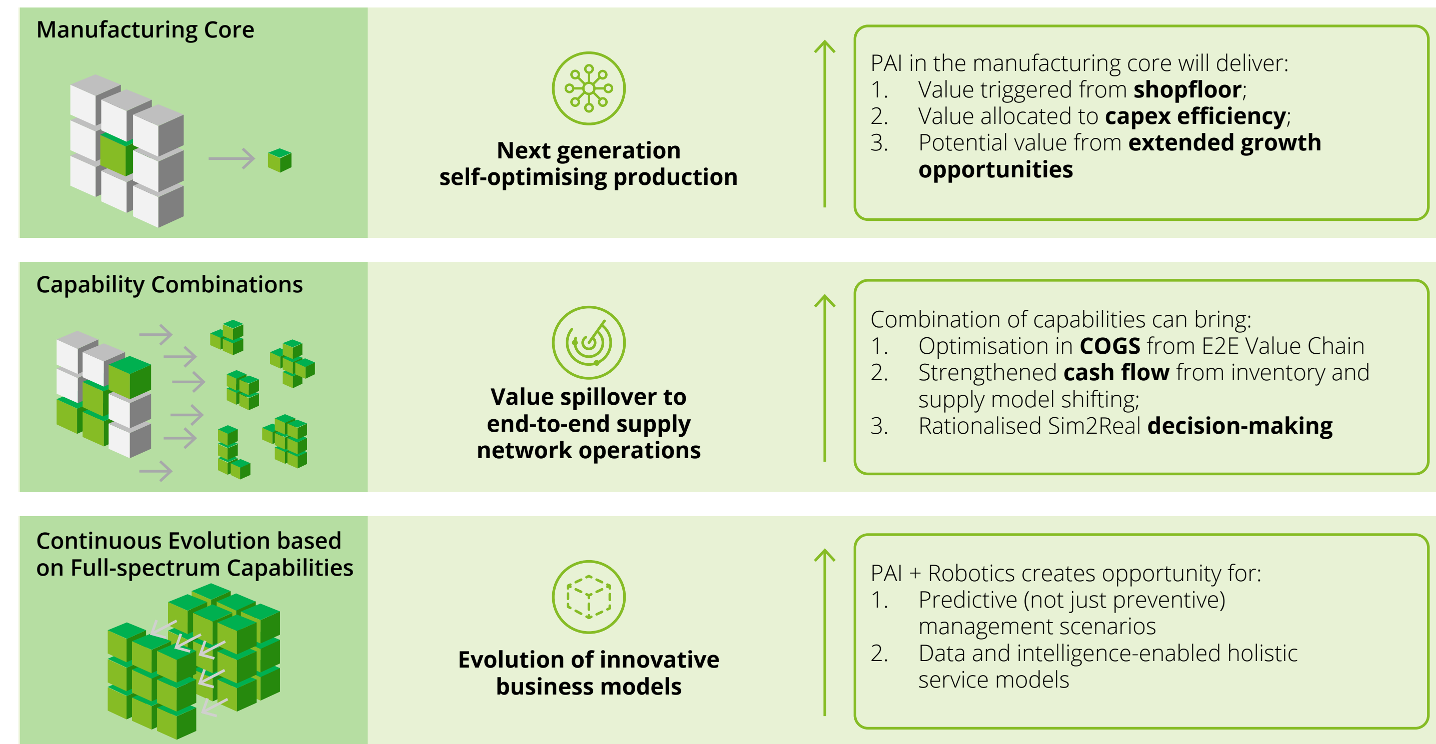

The PAI technology stack

PAI and industrial robotics are not a single technology but a deeply integrated stack (Figure 4). Understanding the six layers – and their interdependencies – is essential for sound investment decisions and avoiding costly gaps:

- **Platform:** The digital twin environment where AI models are trained and workflows are simulated before physical deployment.
- **AI and Cognition:** the 'brain' responsible for reasoning and decision-making, controlling adaptive interaction with the physical world, from environmental understanding to task and motion planning.
- **Sensing and Perception:** The system's 'eyes, ears and skin', comprising cameras, force-torque sensors, and sound-wave detectors that fuse sensor data to create a real-time model of the physical world.
- **Mechanical Structure:** The robot 'body' – articulated arms, humanoids, AMRs – defines its physical capabilities and degrees of freedom. Modular design and advanced planning and scheduling (APS) enable flexible production and rapid reconfiguration between tasks.
- **Compute and Control:** Edge computing connects the AI brain to the mechanical body. Low-latency processing enables smooth, precise and safe operation, especially important in human-collaborative environments.
- **Power and Actuation:** An often-underestimated constraint. PAI is energy-intensive, and scaling it requires robust battery management and resilient power supply.

The power of this stack comes from the synergy between layers. Taken together, the stack mirrors a biological system: sensing provides the eyes and touch, cognition forms the brain, mechanical systems act as the body, and compute and control serve as the nervous system that binds perception to action.

Mastering it is what transforms effective but isolated automation into intelligent, adaptable, self-learning physical systems. Gaps in any one – inadequate digital twin fidelity, insufficient edge compute, unstandardised sensing data – will constrain what the AI layer can achieve, regardless of how advanced the models themselves are.

Figure 4: The technology building blocks of physical AI robotics

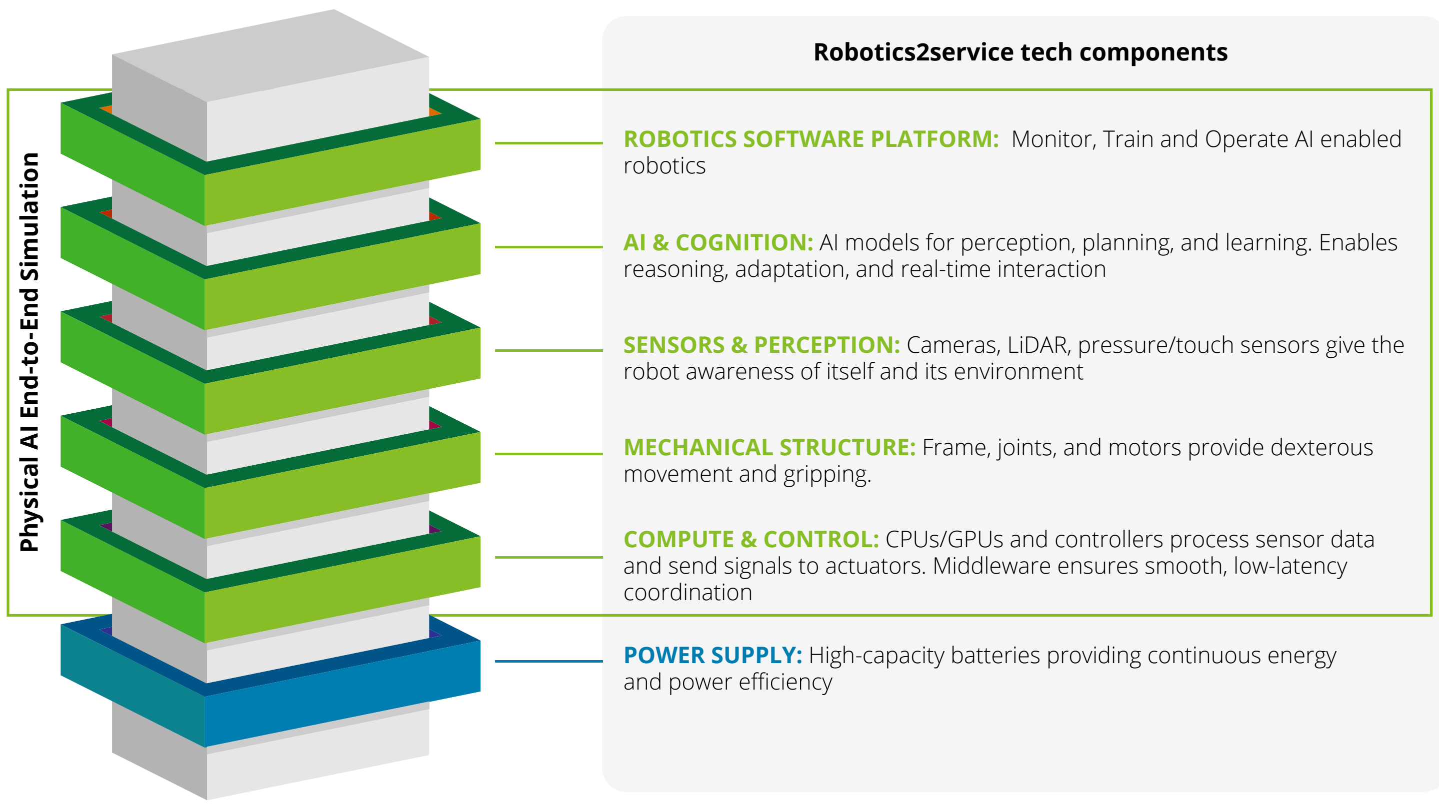

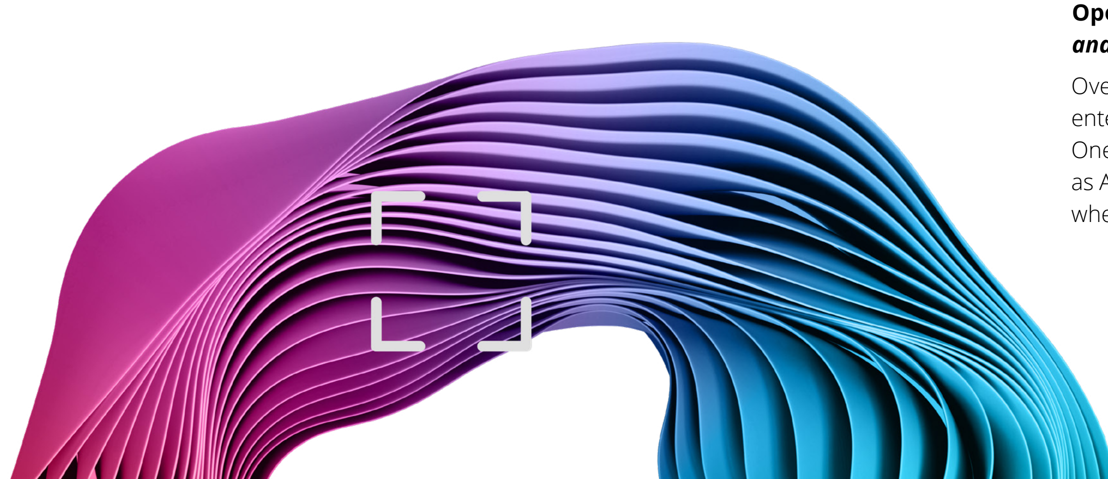

Physical AI: The moment of acceleration

Getting past the bottlenecks

While the potential of PAI is clear, moving from today's robotics-enabled plants to full autonomous, end-to-end operations will require organisations to overcome a set of significant technological and operational bottlenecks.

In Deloitte's State of AI 2026 survey, the most commonly reported barriers to PAI adoption are: cost and resource requirements (41 per cent), challenges identifying use cases (36 per cent), talent and skills gaps (33 per cent), and technology or data availability (31 per cent).[13] These map onto two distinct categories of challenge that require fundamentally different responses.

> "There are still fundamental challenges in the evolution of the PAI stack itself over which business leaders have little or no control. Our view is that the best approach for many industrial and manufacturing leaders is to be an informed observer and strategic partner in research and pilots."

Technological bottlenecks *(largely outside most organisations' control)*:

Progress is driven by technology companies and academia. Key constraints include AI's limited accumulation of physical 'common sense', or its intuitive understanding of how objects behave in novel situations.

They also include the simulation-to-reality gap, where models trained in digital twins don't always transfer cleanly to real-world environments. Additional challenges arise from the complexity of safety certification for autonomous systems in shared human spaces, and from the cost of advanced sensors and edge compute, which remains prohibitive for many applications.

For most leaders, the appropriate response is to act as an informed observer and strategic partner in pilots, rather than to attempt to solve these constraints independently.

Operational bottlenecks *(squarely in organisations' control – and the deciding factor for most organisations)*:

Overcoming operational bottlenecks requires coordinated, enterprise-wide action spanning process, technology and people. One challenge is the adequacy of the automation foundation itself, as AI struggles in low-automation, high-variation environments where lean and standardised processes are not yet in place.

A second constraint is digital infrastructure readiness, including whether Programmable Logic Controllers (PLCs), Supervisory Control and Data Acquisition (SCADA) systems, and Manufacturing Execution Systems (MES) can integrate with AI orchestration platforms. Without Industrial Internet of Things (IIoT) connectivity Physical AI cannot operate at scale.

Talent is a further bottleneck, as hybrid professionals who speak both operations and data science are scarce and take years to develop. At scale, compute and power supply also move beyond pure engineering considerations: competition for high-performance chips and resilient power infrastructure can shape who can deploy and operate PAI effectively.

Unlike technological constraints, these bottlenecks can be resolved through deliberate leadership action.

> "There are internal readiness factors that determine how effectively you can deploy and scale any PAI solution, today or tomorrow – and critically, they are in your control. Successful PAI implementation is as much about adapting as it is adopting."

The critical insight: technological bottlenecks will be resolved by the technology industry, operational bottlenecks sit within the control of individual organisations. In practice, capturing value from PAI depends on both deploying new technology and reshaping processes, skills, and operating models around it.

Realising the value: The dual maturity lens

PAI does not deliver value by default. The journey will be punctuated by rapid advances in some areas and delays while organisation or technology catches up. Success requires two distinct maturity curves to advance in step (Figure 5).

Advanced technology delivers value only when organisations are ready to integrate it effectively, just as operational maturity must be paired with modern capabilities to avoid leaving value untapped.

This is because PAI is not simply a solution to be deployed, but a capability that needs to be cultivated through operational discipline and organisational learning. Leadership in the next decade will come from organisations that can systematically align operational maturity with the accelerating technology frontier, rather than relying on advanced research capability alone. The path forward is to build the PAI-ready environment incrementally—one standardised workstation, one upskilled operator, one digital twin at a time.

Figure 5: The dual maturity lens

Technology Application Maturity	Stage 1: Automation	Stage 2: Collaborative Digitalisation	Stage 3: Digital Twin	Stage 4: Physical AI
Critical Path To Next Maturity Level - Illustrative -	• Industry 4.0 / Physical AI strategy • Digital core systems (ERP, WMS, MES, PLM/ PDM, DCS) • Manufacturing automation with robotics • Data foundation and governance • Automated 'Human-Machine Collaborative Process' design • Supply chain automation	• Process SOP standardisation • Digital lean manufacturing • Material flow planning and visibility • 'Digital Thread' integrated architecture design • Digital core and robotics collaborative technology design	• Robot twin, vision twin, product twin build up, and digital twin application and integration • Industrial knowledge graph setup • Integrated supply chain and manufacture simulation • Organizational transformation upon new technology	• Physical AI enabled robotics deployment • VLA model fine tuning • Physical AI synthetic data • Physical AI agentic application

Operational Maturity and Transformation Readiness	Standardisation	Lean Operation	Organisation and Talent Structure
Key Features Along the Critical Path - Illustrative -	• Human-disciplined SOP execution • Quantified parameters encoded in machine settings • Digital modelling in integrated systems; interfaces standardised • Self-describing standards for autonomous system configuration	• Reactive problem-solving; basic waste observation • Cellular flow, pull systems; Andon for real-time response • Predictive lean with digital simulation and model-driven Kaizen • Autonomous self-healing operations with AI-driven anomaly resolution	• Skilled hands for manual execution; supervisors manage people • Operators as equipment users; managers steward human-machine systems • Hybrid "translator" roles emerge; integrated OT/IT teams • Workforce co-innovates with AI; governance for AI lifecycle

Physical AI: The moment of acceleration

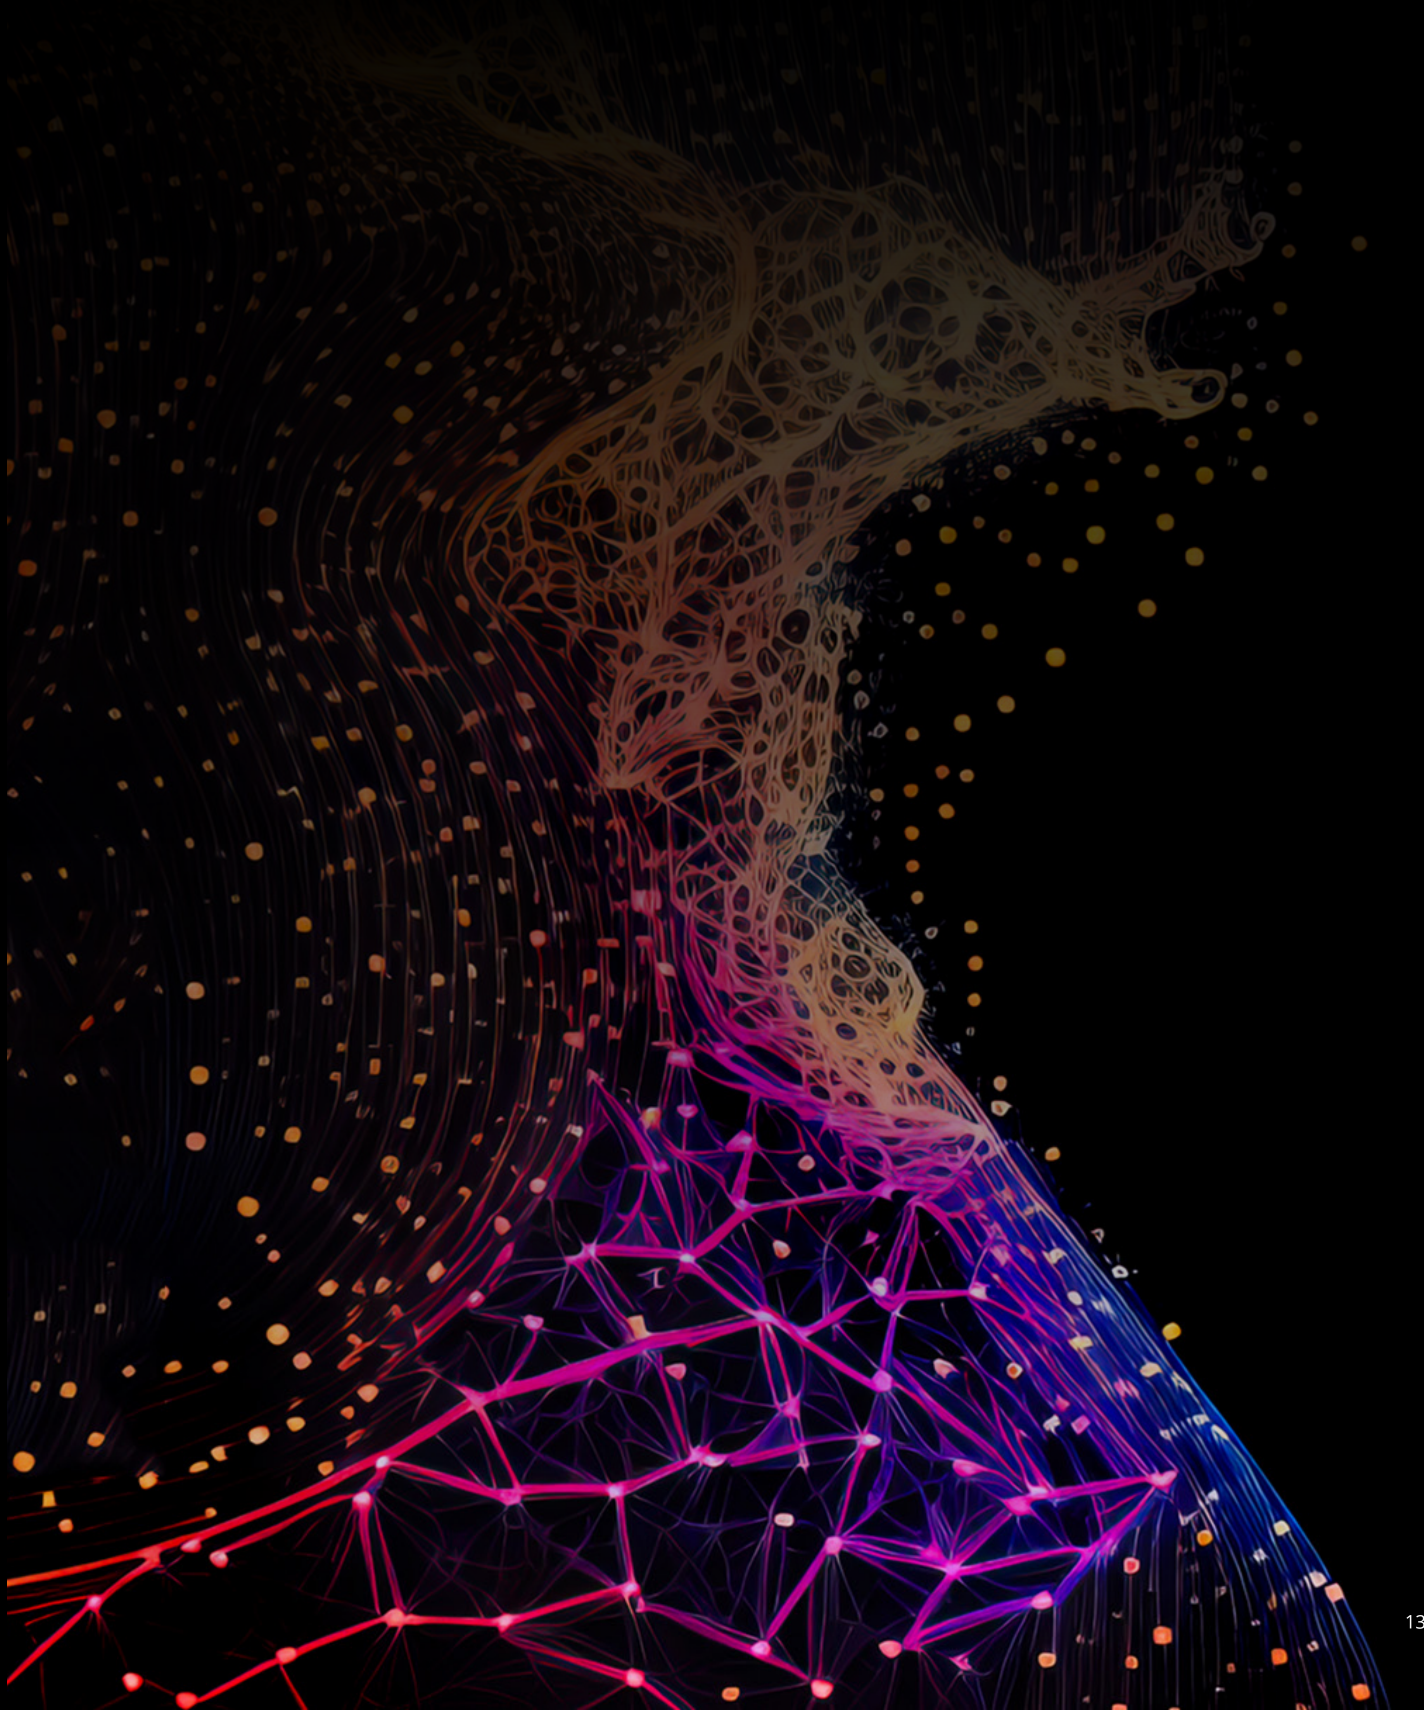

Technology application maturity: Four stages

The evolution from rigid machinery to truly autonomous, adaptive physical systems follows a structured ascent through four stages. Each has a defined end-state, a critical path, and prerequisite technologies. Each stage marks a clear break in what machines can sense, decide, and act on in the physical world. Understanding this architecture is the basis for strategic sequencing – knowing what to build, when to build it, and why.

Stage 1 – Automation. Machines execute predefined sequences with precision, speed, and repeatability. The system is blind to its environment but performs flawlessly under controlled conditions. Human intervention is required for setup, changeover, and exception handling. Mature automation is a capital-intensive and engineering-driven. It is not about machine intelligence, but mechanical precision and control logic. Key prerequisites: Programmable Logic Controllers (PLCs), servo motors, dedicated tooling, safety systems, basic Human-Machine Interfaces (HMI). Many organisations still operate Stage 1 with significant pockets of manual work. Completing Stage 1 across the value stream is the prerequisite for everything that follows – the temptation to skip ahead is one of the most common and costly mistakes.

Stage 2 – Collaborative Digitalisation. Machines become aware of human presence and environmental context. Physical guarding gives way to collaborative safety architectures. Machines and humans share workspace fluidly. Data begins to flow upward from the shop floor. This stage is driven by sensor integration and evolving control software. Existing machines are augmented with new capabilities – new sensors, operating systems, software upgrades – and new forms of collaborative robots are introduced. Key prerequisites: collaborative robots (cobots), 2D/3D vision cameras, force-torque sensors, IIoT gateways, edge computing nodes. Stage 2 is where operations become visible: the data generated here (cycle times, error rates, operator interactions) becomes the fuel for Stage 3.

Stage 3 – Digital Twin. A persistent, bi-directional synchronisation between physical operations and their virtual representations. Every machine, workpiece and process has a digital shadow reflecting its real-time state. Simulation evolves from an offline design tool into an operational runtime environment for prediction, optimisation and commissioning. This is the most demanding lift in the entire technology maturity model because it requires the convergence of operating technology and IT, alongside data infrastructure investment and organisational buy-in for simulation fidelity. Key prerequisites: digital twin platforms, physics engines, high-bandwidth low-latency connectivity, time-series databases, MES and ERP integration layers. Stage 3 is the gateway to PAI: without a validated digital twin, AI models trained in simulation will fail in reality.

Stage 4 – PAI. Physical systems that perceive unstructured environments, reason about optimal actions, and execute with autonomy. At this stage, the separation between the brain (generalised intelligence), the eyes (perception), and the body (physical embodiment) becomes explicit rather than implicit.

PAI generalises across tasks, learns from both real and simulated experience, and improves continuously. The human role shifts from operator to supervisor, exception handler and co-innovator.

A defining shift at this stage is the decoupling of intelligence from embodiment. Instead of programming skills robot by robot, PAI systems increasingly separate generalised intelligence, trained in simulation and refined through fleet learning, from specific physical forms. Capabilities learned on one machine can transfer across different robots and tasks with minimal retraining, dramatically reducing deployment time and expanding reuse. This ability to generalise transforms robotics from bespoke automation into scalable physical intelligence.

This is the frontier. The path is not yet fully proven, but the contours are emerging from leading adopters. Key prerequisites: AI foundation models adapted for embodied tasks, reinforcement learning frameworks, fleet orchestration platforms, edge AI accelerators, automated data labelling pipelines. Stage 4 is not a destination but a new mode of operation: instead of a factory of individual robots, it becomes a learning production system, continuously improving as it operates.

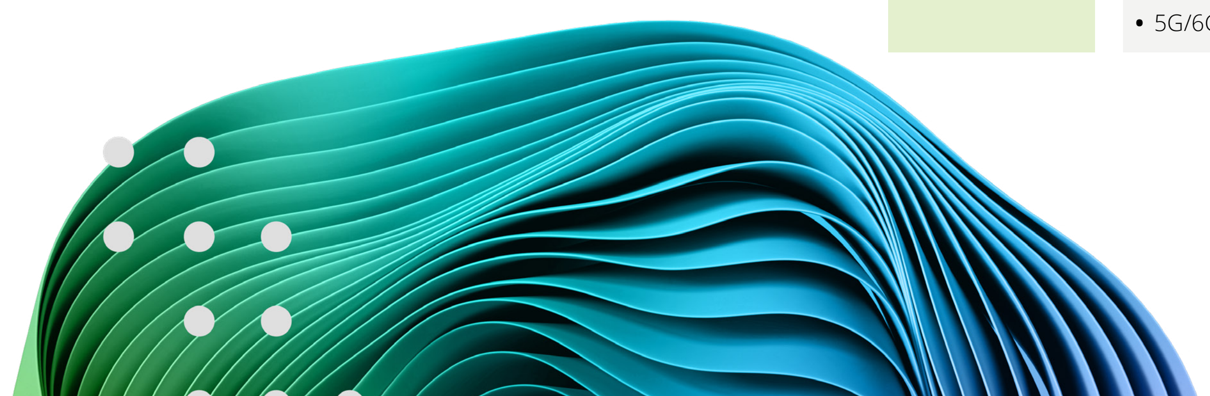

Physical AI: The moment of acceleration

Figure 6: Physical AI maturity roadmap

	80 % of enterprises today Opportunity: to get organisations "PAI ready" for the future			**20 % of enterprises today** Opportunity: to implement PAI solutions at scale
	Stage 1: Automation	**Stage 2: Collaborative Digitalisation**	**Stage 3: Digital Twin**	**Stage 4: Physical AI**
End-state Goal	• **High automation adoption:** target greater than 70%, high penetration rate of pr Woduct manufacturing processes with significant labor reduction. • Pursuing **predictable output, lower operating costs**, and **higher quality stability**.	• **Implementation of a set of collaborative digital platforms**, achieve system-level standardisation of: – Robot task assignment and work content. – Work environment – Operational execution (incl. optimisation) – Data language and interaction protocols – Technology application	• **High Sim2Real models coverage**. • Accelerate new product **time-to-market**, achieve **deep flexible manufacturing**, and **improve productivity** through digital twin technology. • Achieve **proactive risk management**, **knowledge asset accumulation**, and **low-cost innovation**.	• Continuous and **high-density model training**. • A factory with **no human involvement** and **AI autonomous decision-making**.
Key Technologies - Illustrative -	• Robotic platforms, ERP systems, Warehouse mgmt systems Mfg. execution systems, Product lifecycle mgmt, Product data mgmt, SCADA • MDM/Data lake/ Low-code/No-code development platform • Motion control and actuating mechanisms. • Dedicated industrial robots, smart machine, and automated systems • Basic sensors, Computer vision, • RPA/Business Process Management platform • 5G/6G/IoT	• Low-Code/No-Code platforms and workflow engines • Value stream estimation and simulation • Advanced analytics tools • IoT and automatic identification technology • Supply chain control tower • Unified data model, master data, and metadata management • Cloud services, edge computing hardware and frameworks	• Multi-physics modeling and simulation engine, digital twin platform • Machine learning and deep learning platform • Multimodal Generative AI • Natural language processing and information extraction, graph databases • Synthetic data and physical simulation technology • Data model prediction and scene planning	• Physical AI deployment • VLA model training • Physical AI synthetic data simulation training in digital twin • 'Autonomous Decision-Making and Collaboration' agent development

Physical AI: The moment of acceleration

Operational maturity and transformation readiness

Technology application maturity determines what is possible. Operational maturity determines what is executable. The most sophisticated PAI system will deliver little value if deployed into an organisation lacking the foundational discipline, flow, and human architecture to absorb it. Three interdependent dimensions must evolve in parallel with the technology stages.

Dimension 1: Standardisation – the language of scalability

PAI thrives on consistency. Variability in parts, processes, or interfaces is noise that confuses both classical automation and AI models. Standardisation does not mean rigidity – it means controlled, documented variation that can be taught to a system. At lower maturity, process knowledge lives in operators' heads, quality depends on the least-skilled person on the shift. As maturity advances, documented standards enable predictable environments for collaborative robots and vision systems. At the digital twin stage, standards become machine-enforced – equipment communicates through common protocols with unified data schemas, virtual replicas synchronise with physical assets at scale. At the PAI frontier, standards become adaptive and self-describing: equipment auto-configures on connection, workpieces carry their own process recipes.

Dimension 2: Lean operation – waste elimination as a precondition for intelligence

Automating a wasteful process merely bakes in the waste. Lean principles such as value-stream mapping, continuous and visible flow, decoupled processes, root-cause problem-solving create healthy processes. PAI then acts as a force multiplier on an already optimised system. At lower maturity, automation yields isolated speed gains that are lost to waiting and transport. As lean maturity deepens, pull systems create predictable, visible value streams – the prerequisites for sensor deployment and collaborative automation. At the digital twin level, lean becomes predictive and model driven. At the PAI frontier, operations become autonomous: AI agents detect anomalies, diagnose root causes, and implement countermeasures through simulation-validated policies. The kaizen culture of continuous incremental improvement is also the cultural bridge that helps workforces view AI as problem-solving augmentation, not job replacement.

Dimension 3: Organisation and talent transformation – the hardest and most neglected dimension

The introduction of PAI requires new roles, new skills, new decision rights, and new career pathways. The questions every leader should be asking now: Where are the hybrid professionals who speak both operations and data science? How are silos between engineering and IT broken down? Who manages joint ownership of digital twins and robot fleets? Is the workforce ready to shift from manual execution to supervision and continuous improvement of AI systems?

Each technology era has demanded a new type of human workforce. Automation moved the standard from operator memory to machine settings – a torque value stopped being 'tighten until it feels right' and became 45 Newton-metres programmed into a spindle controller. Operators shifted from performing value-add to enabling it.

Today, digital twins are moving process knowledge from manuals to systems; a new breed of 'translators' are emerging who understood both the physics of the process and the logic of the code. Now, PAI demands the next evolution: operators become supervisors of autonomous systems; managers orchestrate fleets of intelligent agents; the translators of today become the architects of tomorrow's human-machine workflows. Discovery and development of use cases will be led from the frontline. The PAI-enabled organisation will no longer be a hierarchy of people, but– a network of human and artificial intelligence, learning and improving together.

Governance: The emerging dimension

The three operational dimensions are not static checklists – they are organisational states that must continuously evolve alongside the technology maturity stages. A governance dimension has also become critical. As PAI reaches the factory floor, a new class of risks emerge: physical health and safety, system autonomy, accountability, and operational control – often arising at machine speed and beyond direct human intervention. Leading companies are establishing corporate-level AI governance: rules for data use and AI applications, measurable approaches to explainability, and clear accountability frameworks. The current trend is restrictive – defining the circumstances under which AI should not be used. This will evolve as high-risk scenarios are protected by robust rules and confidence grows. Over time, the challenge will shift from restricting autonomy to defining where and how it can be safely granted within clearly governed boundaries and guardrails.

3.4 Navigating the journey: Three failure modes to avoid

Even organisations with sound strategy consistently encounter the same failure modes. Across the entire PAI deployment lifecycle, the aim should be to compress the transformation timeline while safeguarding the business case by aligning technology, process, and talent around a singular outcome: real-world, sustained business impact.

Operational instability during deployment

Every PAI deployment disrupts production schedules, workforce routines, and the equilibrium of a value stream. Without simulation-first validation and phased rollout, instability triggers resistance and extends the path to stability. The fix: rigorous change management embedded within the technical rollout to safeguard throughput and quality at every inflection.

Extended break-even

The interval between first investment and first value is the period of greatest organisational vulnerability. By decomposing the PAI transformation into modular, self-contained value increments, it is possible to capture cash-visible returns within quarters, not years. Early wins build conviction among stakeholders, fund the next phase of investment, and transform the initiative from a cost centre into a self-sustaining value engine.

Failure to achieve take-off after go-live

Many organisations that invest heavily in PAI (as they have in previous digitalisation initiatives) find themselves unable to scale beyond the pilot. Alongside the robots and code, organisations must embed the blueprints, training curricula and governance models needed to operate, adapt and extend the capability. Technology must become a repeatable, scalable engine that accelerates with each successive implementation.

Figure 7: Increasing confidence in physical AI value delivery

Commercial application of physical AI requires companies to build strong organisational and technological foundations.

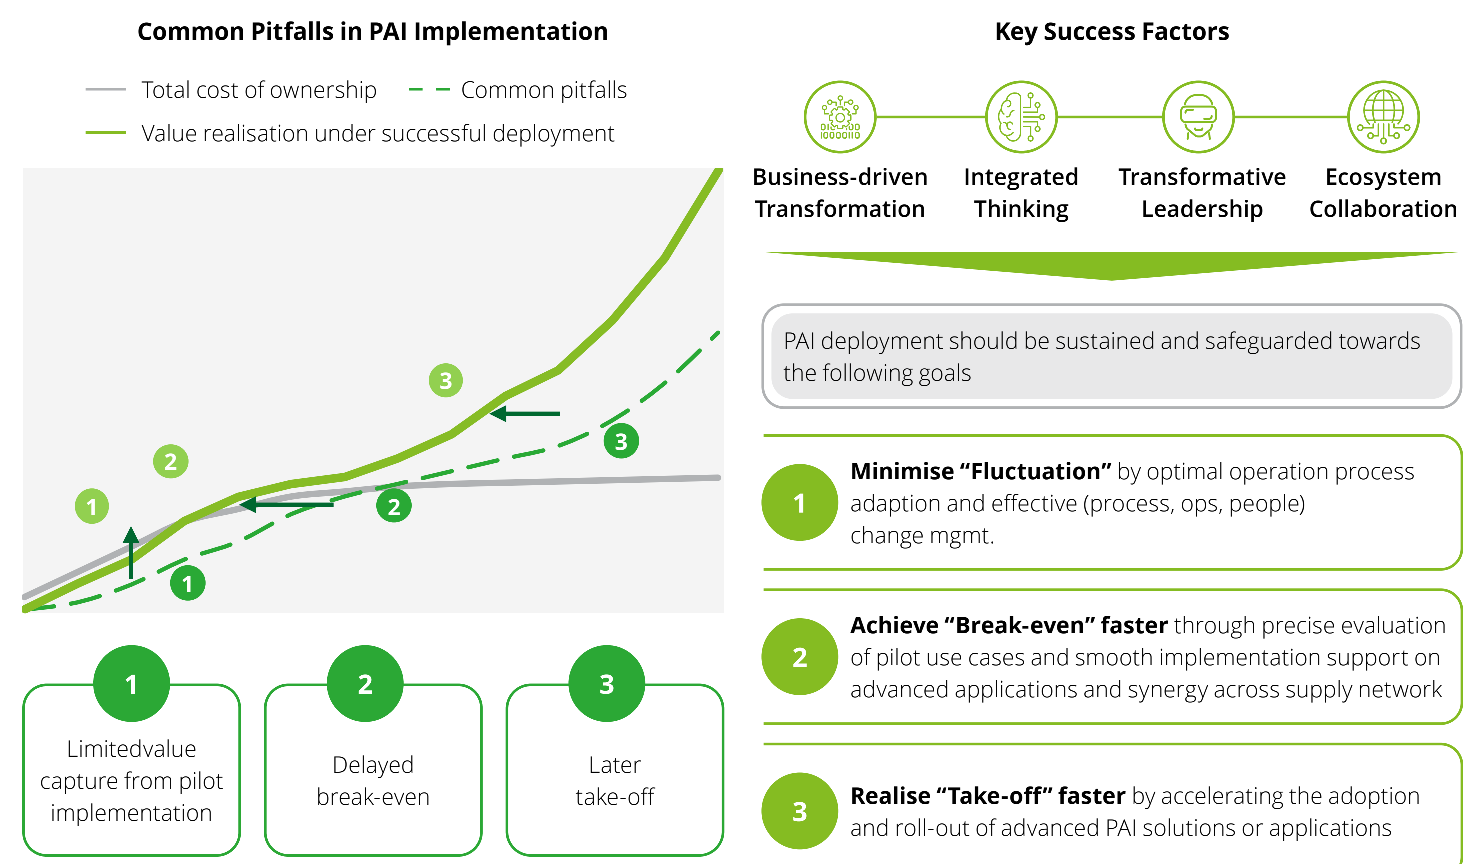

The three questions every operations leader should be asking now

PAI is no longer a question of 'if' or 'when' – it is a question of readiness. With 41 per cent of business leaders expecting transformational impact within three years and PAI integration set to grow six-fold in two years, the window to build the operational foundation is narrowing.[14]

As with any rapidly evolving technology, many of PAI's implications are still to play out. Even so, clear and increasingly structured pathways are emerging for the journey ahead. As PAI becomes more widely adopted, more visible and, critically, as its value becomes more tangible, its application will expand rapidly across sectors and value chains. The cost of moving too slowly is not just missed efficiency, but the loss of the organisational learning that comes from being an early mover.

For industrial and manufacturing leaders, three questions should drive the agenda:

Where are you on the technology maturity ladder?

Stage 1 completion – consistent, lean, digitally accessible automation – is the prerequisite for everything that follows. If your shop floor still runs on tribal knowledge and uncontrolled variation, PAI investment is premature regardless of what the technology can offer.

Are your operational fundamentals PAI-ready?

Assess the three dimensions honestly: standardisation, lean process maturity, and talent architecture. The organisations that will capture PAI's value are not those with the flashiest pilots, they are those with the most disciplined foundations. Technology deployed into operational immaturity will fail to deliver.

Are you building the right human architecture now?

The talent bottleneck is real and long-lead. Hybrid professionals who speak both operations and data science act as a bridge between the plant floor and the server room and take years to develop. The organisations assembling cross-functional teams and building frontline-led use case discovery now will move significantly faster as the technology scales.

The pioneers learning PAI on the factory floor today are writing the competitive playbook for the next decade. The frontier has shifted. The future is being built right now.

Authors

Scott Xiaozhen Sun
Asia Pacific Physical AI Centre of Excellence leader
Deloitte China
scottsun@deloittecn.com.cn

Joe Zhuo Deng
Supply Chain & Network Operations, Physical AI Centre of Excellence, manager
Deloitte China
jdeng@deloittecn.com.cn

Key Contacts

Dr. Amanda Williamson
Director
Deloitte New Zealand AI Institute
amawilliamson@deloitte.co.nz

Paulo Osorio
Partner
Deloitte New Zealand
paosorio@deloitte.co.nz

Ruby Osborne
Physical AI specialist
Deloitte New Zealand AI Institute
ruosborne@deloitte.co.nz

Glossary

AMR – Autonomous Mobile Robots	Self-navigating robots that use sensors, mapping, and onboard intelligence to move materials and perform tasks safely in dynamic environments without fixed paths.
Batch-size-one-pulling	A demand-driven production model where each individual unit is produced in response to a specific customer order, enabling mass customization with minimal inventory and waste.
Edge Devices	Industrial hardware that processes data locally at or near machines (rather than in the cloud), enabling low-latency control, real-time analytics, and autonomous decision-making
ERP – Enterprise Resource Planning	Integrated enterprise software that coordinates core business processes – such as finance, supply chain, manufacturing, and procurement – using a shared system of record.
LLM – Large Language Model	A deep-learning model trained on large volumes of text to understand, generate, and reason with human language, forming the cognitive backbone for many modern AI and multimodal systems
MES – Manufacturing Execution System	Software that monitors, manages, and optimizes production operations on the factory floor in real time, bridging enterprise planning systems and physical manufacturing.
NPI – New Product Introduction	The structured process of taking a product from concept through design, validation, and ramp up into full scale manufacturing and market launch.
PLM – Product Lifecycle Management	A digital framework for managing product data, processes, and decisions across the full lifecycle – from ideation and design through production, service, and retirement.
SCADA – Supervisory Control and Data Acquisition	The software and hardware layer that monitors and controls industrial operations in real time.
VLA – Vision Language Action model	A multimodal AI model that combines vision and natural language understanding to directly generate executable physical actions, enabling robots to perceive, reason, and act in the real world from high level instructions
VLM – Vision Language Model	A multimodal AI model that integrates visual perception and language understanding, allowing systems to interpret images or video and reason about them using natural language.

Endnotes

1. Deloitte, State of AI in the Enterprise, January 2026
2. https://www.deloitte.com/cz-sk/en/issues/generative-ai/state-of-ai-in-enterprise.html
3. NVIDIA, Train a Quadruped Locomotion Policy and Simulate Cloth Manipulation with NVIDIA Isaac Lab and Newton, 29 September 2025, Train a Quadruped Locomotion Policy and Simulate Cloth Manipulation with NVIDIA Isaac Lab and Newton | NVIDIA Technical Blog
4. Arxiv.org, OpenVLA: An Open-Source Vision-Language-Action Model, 5 September 2024, https://arxiv.org/abs/2406.09246;
5. NVIDIA Physical AI at NeurIPS: Alpamayo-R1 and the Cosmos Ecosystem, 12 December 2025, https://introl.com/blog/nvidia-physical-ai-alpamayo-cosmos-neurips-2025
6. TechCrunch, Hugging Face opens up orders for its Reachy Mini desktop robots, 9 July 2025, https://techcrunch.com/2025/07/09/hugging-face-opens-up-orders-for-its-reachy-mini-desktop-robots/
7. Gemini Robotics 1.5 brings AI agents into the physical world, 25 September 2025, https://deepmind.google/blog/gemini-robotics-15-brings-ai-agents-into-the-physical-world/
8. IFR, World Robotics 2025, September 2025
9. CitiGPS, The Rise of AI Robots: Physical AI is Coming for You https://www.citifirst.com.hk/home/upload/citi_research/rsch_pdf_30297368.pdf
10. AI Impact Assessment Guideline (ISO/IEC FDIS 42005:2025); AI Explainability Guideline (ISO/IEC TS 6254:2025); AI Risk Management Guideline (ISO/IEC 23894)
11. Sustainability Magazine, How BYD is Harnessing AI to Accelerate Global EV Production, 9 October 2025, https://sustainabilitymag.com/news/how-chinas-byd-is-using-ai-to-scale-global-ev-manufacturing
12. Forbes, The Rise of AI-Powered Robotics, And The Future of Work, 15 April 2025, https://www.forbes.com/sites/jackkelly/2025/04/15/the-rise-of-ai-powered-robotics-and-the-future-of-work/
13. https://www.york.ac.uk/news-and-events/news/2024/research/ai-robotics-emergency-room-waiting-time/
14. Deloitte, State of AI in the Enterprise, January 2026
15. Deloitte, State of AI in the Enterprise, January 2026
16. Deloitte, State of AI in the Enterprise, January 2026

Deloitte.

Together makes progress