2026

I, ROBOT

Humanoid robots will capture attention; specialized physical AI will capture value



BLUE SHIFT
BY ARTHUR D. LITTLE



“A robot may not injure a human being.”

— Isaac Asimov, *First Law of Robotics*

I, ROBOT

Humanoid robots will capture attention; specialized physical AI will capture value

,

Authors

Dr. Albert Meige, Director of Blue Shift
Rick Eagar, Partner Emeritus
Daniel Granath, Consultant
Sebastian Meyer, Principal
Caroline Mellquist-Lindwall, Consultant
Glen Peres, Manager

Contributors

Lars Thurmann-Moe, Partner
Karim Taga, Partner

COMMITMENT

COMMITMENT

COMMITMENT

COMMITMENT

Executive summary 6

Preamble 10

Chapter 1. Physical AI — A real opportunity dominated by the humanoid narrative 14

Chapter 2. Value will first come from specialized physical AI 34

Chapter 3. Critical uncertainties & future scenarios 46

Chapter 4. What executives should do today 54

Executive summary

There is currently much hype around intelligent humanoid robots. In reality, however, many of the well-publicized demos are either preprogrammed, remote-controlled, or totally AI-generated. At the same time, there is great progress in applying AI to physical machines, with the market growing rapidly. In this report, we explore the latest status of "physical AI" — where it can deliver the most value to business, how it will develop over the coming years, and what executives should do now to prepare for a future of intelligent machines operating alongside humans.

CHAPTER 1: PHYSICAL AI — A REAL OPPORTUNITY DOMINATED BY THE HUMANOID NARRATIVE

Physical AI refers to intelligent and autonomous machines that can perceive, decide, and act in the physical world. Physical AI brings real-world understanding to AI and versatility to robotics. It is defined by embedded intelligence rather than any specific robotic form such as humanoids.

The market is significant and rapidly growing. The current market for physical AI is around US $18 billion, of which humanoids account for only $2–$3 billion. While forecasts differ greatly due to varying category boundaries and assumptions, a likely near-term forecast range is $60–$100 billion by 2030, with a CAGR of more than 20%.

In the longer term, forecasts suggest a range between $4 trillion and $12 trillion by 2050, depending on the extent to which adoption extends beyond industry. Investor commitment is high. Robotics venture capital (VC) reached almost $28 billion in 2025. As an example, humanoid robotics manufacturer Figure AI's valuation rose 1,200% in less than a year.

Value in physical AI will be created across an unevenly mature technology stack and ecosystem comprising three layers: Body, Mind, and Learning Loop. Progress in hardware and perception still outpaces advances in dexterity, reasoning, and learning. The Learning Loop is the least mature layer. The value chain remains unsettled, with new roles emerging around platform integration, application design, and maintenance. A consolidation phase is likely in the near term. China is building advantage through manufacturing scale, supplier density, and coordinated research. The US is strongest in frontier AI, compute, and capital but weaker in the Body layer. Europe and Japan are comparatively stronger in industrial systems automation, systems integration, and applied robotics but weaker in frontier AI models. Geopolitics will play a major role in shaping which players and regions lead the industry.

Physical AI brings real-world understanding to AI and versatility to robotics.

The humanoid narrative is overstating general-purpose readiness. Humanoids make physical AI easy to imagine as a universal labor platform and support the prospect of a bigger market. However, demonstrations are often mistaken for deployment readiness. Deployment depends on much more than a successful demo: a robot must still prove that it can operate safely, repeatedly, and economically in a live setting.

CHAPTER 2: VALUE WILL FIRST COME FROM SPECIALIZED PHYSICAL AI

Strong drivers are accelerating physical AI adoption. Demand-pull factors include long-term structural labor shortages, shifting work preferences, increasing needs for operational resilience, and productivity pressures. Factors pushing adoption include pragmatic workarounds to overcome dexterity limitations, hybrid systems accommodating limited autonomy, cheaper commoditized hardware, new middle-layer players bridging the gap between vendors and buyers, emerging commercial models such as robot-as-a-service, and growing acceptance of integrated robot-human operations.

However, mass adoption at scale remains constrained by technological, economic, and deployment barriers. Lack of human-hand dexterity remains a key factor limiting near-term adoption. Greater autonomy capabilities are still needed in uncontrolled environments. Physical AI requires training data from the physical world, which is difficult to generalize and access. While there are economically attractive use cases, there is not yet a "killer app" with broad applicability and transformative benefits. Physical AI also has the potential to cause harm and damage, while safety regulation lags technical progress. Rules for liability allocation in the case of malfunction are not yet established, and capability control poses unresolved governance challenges. There is a lack of reliable performance standards, and insurability remains uncertain.

Determining the optimal form factor for applications is still immature, and cultural acceptance remains unclear.

The sweet spot for physical AI is in structured environments with some degree of variability. In the short to medium term, wheel-based or arm-only form factors will likely scale first in industrial applications. Sector applications like healthcare, pharmaceuticals, manufacturing, and logistics will evolve toward greater autonomy and improved dexterity. Applications in services, agriculture, construction, mining, and mobility will require further development to operate effectively in less controlled environments. In the longer term, if humanoid robots achieve sufficient improvements in cost, power, autonomy, and dexterity, alongside application-specific software, they could ultimately come to dominate the market.

Advances in autonomy and AI are fundamental to physical AI's trajectory.

CHAPTER 3: CRITICAL UNCERTAINTIES & FUTURE SCENARIOS

Many factors will shape the future trajectory of physical AI. We identified 21 uncertainty factors across six categories: (1) structural and societal changes, (2) autonomy and AI, (3) dexterity and hardware, (4) adoption and use cases, (5) industry structure and dynamics, and (6) economic viability.

The two most critical uncertainties are the pace of technological breakthroughs and the speed of deployment scaling. Advances in autonomy and AI are fundamental to physical AI's trajectory, including improvements in AI architectures, tactile data availability, world models, decision-making systems, and the evolution from deterministic systems to learning-based models, as well as the development of data flywheels and continuous learning loops. Similarly, advances in dexterity and hardware are critical for many applications, requiring progress in actuators, sensors, batteries, and materials. Dexterity remains one of the key bottlenecks limiting expansion into broader use cases.

Economic viability is another major uncertainty affecting adoption. Industry structure and dynamics will also need to evolve to support development and deployment at scale. The breadth and depth of applications and use cases, from industrial to service sectors, will shape how demand develops over time.

These critical uncertainties lead to four plausible scenarios for the future trajectory of physical AI:

1. **Incremental Automation.** Technological progress and deployment scaling remain limited due to persistent technological and institutional barriers.
2. **Industrial Workhorses.** Deployment scales rapidly in targeted industrial applications, but technological breakthroughs remain partial, limiting general-purpose capabilities.
3. **Specialized Intelligence.** Robots achieve high capability levels in specific domains, particularly in tasks requiring complex autonomy or dexterity, but scaling remains constrained by high costs, safety concerns, and integration challenges.

4. **Autonomous Abundance.** Simultaneous breakthroughs in capability and scaling enable widespread deployment, with general-purpose robots becoming viable across industries, including services and potentially consumer applications.

Each scenario carries implications for deployment strategy, priorities, and investment focus. Regardless of which scenarios ultimately prevail, we expect structured industrial environments, hybrid human-robot operating models, and a diversity of robot embodiments — rather than a single dominant humanoid platform — to define the near term.

The decisions made over the next few years will determine which organizations lead and which fall behind.

CHAPTER 4: WHAT EXECUTIVES SHOULD DO TODAY

Executives should act now. Despite the hype, development progress is accelerating, and early movers may gain significant competitive advantage. The challenge for industry leaders is not whether to engage, but how to do so pragmatically and in ways anchored in real operational value rather than technology narratives.

Start by identifying and evaluating use cases. Build a portfolio beginning with scalable, repeatable applications where physical AI is already "good enough." Also, target applications where advances in either autonomy or dexterity could unlock specific high-value tasks. Be creative with new workflows and ways of working.

Initiate foundational "no-regret" moves. Build the foundations for future deployment by strengthening operational data and digital connectivity, engaging with the physical AI ecosystem to build capabilities, and working with governments to support the introduction of appropriate regulation.

Test and pilot in stages. Design pilots to test real-world complexity, beginning with lower-risk, high-return use cases. Establish a sequenced approach with tight scoping to avoid overextension. Consider the full end-to-end economics of deployment at scale, which may differ substantially from the economics of isolated pilots. Adapt strategies as physical AI development scenarios evolve.

Plan for and manage the impact of deployment. Develop clarity on what success looks like, how deployment will affect the organization, and the required approach to transformation. Implement robust change management processes to manage and monitor the transition.

The decisions made over the next few years about where to pilot, what to measure, which vendors to engage, and how to develop the workforce will determine which organizations lead and which fall behind. The winners will not necessarily be those pursuing full autonomy or exclusively humanoid visions, but those focused on solving real operational problems, deploying early, learning quickly, and building internal capabilities and data advantages.

Preamble

Many videos showing hundreds, or even thousands, of humanoids at work are fake, AI-generated, or heavily staged. When the videos are real, the robots are often performing prerecorded choreographies, operating in highly controlled environments, or being teleoperated by humans.

Yet humanoid manufacturers, media outlets, and many consulting firms suggest that a humanoid revolution is imminent, and that robots will soon enter our factories, warehouses, and homes at scale.

Does this sound familiar? Twenty-five years ago, Honda's humanoid robot ASIMO (Advanced Step in Innovative Mobility) carried a similar promise. This is not surprising, and to paraphrase futurist Hans Moravec, "The easy problems are hard."

So are humanoids for real? Yes, meaningful progress has been made. Does this mean humanoids will invade companies in the near future? Probably not. Does this mean we should ignore the humanoid hype? Definitely not.

Humanoids are the moon shot attracting attention, capital, and talent. But the value may emerge elsewhere — in specialized physical AI systems designed for specific, high-value tasks in industrial and operational environments.

For businesses, researchers, and policymakers, physical AI is no longer a futuristic curiosity. It is becoming a technical, economic, regulatory, and strategic reality. This report is built around four core arguments:

Humanoids are the moon shot attracting attention, capital, and talent.

1. **Physical AI is a real opportunity,** currently dominated by the humanoid narrative.
2. **Value will first come from specialized physical AI,** not general-purpose humanoids.
3. **Physical AI remains surrounded** by critical technological, economic, and adoption uncertainties.
4. **Executives should prepare for several possible futures,** rather than bet on a single scenario.

This report explores these questions in depth, drawing on recent scientific progress, industry dynamics, and strategic implications. It also offers a critical perspective on the myths surrounding humanoids, separating realistic opportunities from futuristic illusions.

As a teaser, here is an intriguing anagram:

“AI and robotics”

transforms into

“Ionic bastardo”

Anagrams move in mysterious ways, but this one reminds us that physical AI is not born from a single discipline. It is a hybrid creature, assembled from AI, robotics, sensors, actuators, energy systems, software, data, and industrial know-how.

— Albert Meige, PhD

“The level of hype about … robotics completely distorts people’s understanding of reality.”

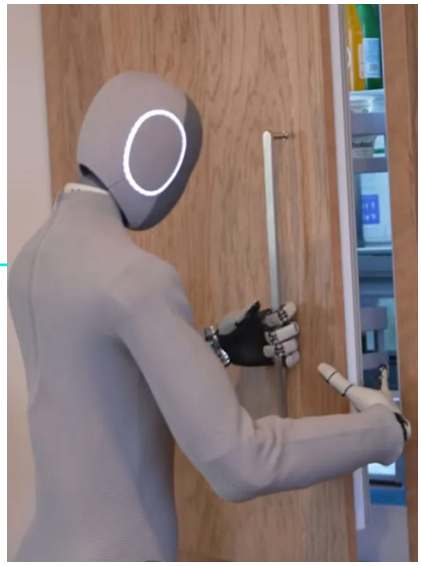
— Rodney Brooks, cofounder, iRobot; founder, Rethink Robotics

THE REAL VALUE OF PHYSICAL AI



Some of the humanoid videos you see are **TOTALLY FAKE** (AI-generated)

Most real videos showcase humanoids that are either **TELEOPERATED**





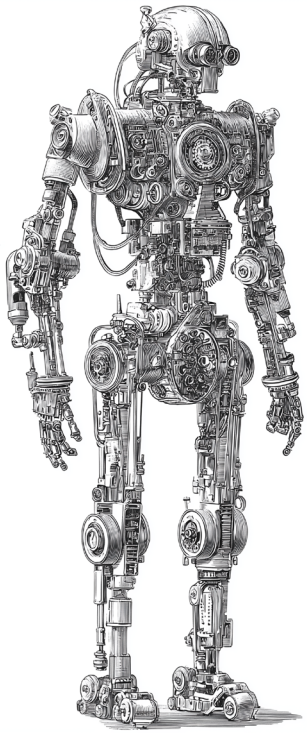


Or prerecorded **CHOREGRAPHIES** without any autonomy

From totally fake to not possible to deploy in real conditions **WHAT IS THE REAL VALUE OF PHYSICAL AI?**

And you may remember humanoids in Japan ... **25 YEARS AGO** Honda Asimo (2000–2022)





CHAPTER 1

PHYSICAL AI — A REAL OPPORTUNITY DOMINATED BY THE HUMANOID NARRATIVE

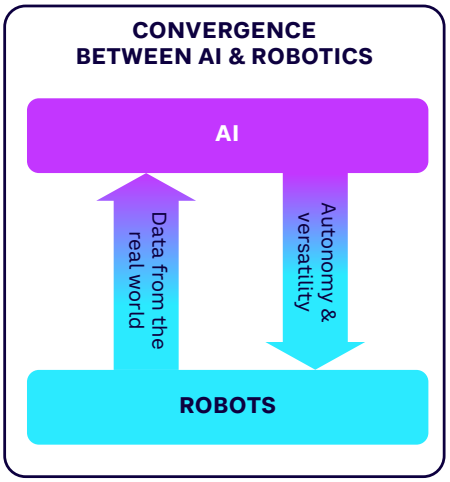
1 PHYSICAL AI — A REAL OPPORTUNITY DOMINATED BY THE HUMANOID NARRATIVE

Physical AI refers to intelligent machines that can perceive, decide, and act in the physical world, combining AI software with robotic hardware across a range of embodiments. It represents a real, large, and fast-growing opportunity. However, the dominance of the humanoid narrative overstates both the current maturity of the value chain and physical AI's readiness for general-purpose deployment.

PHYSICAL AI IS ABOUT INTELLIGENT MACHINES ACTING IN THE PHYSICAL WORLD

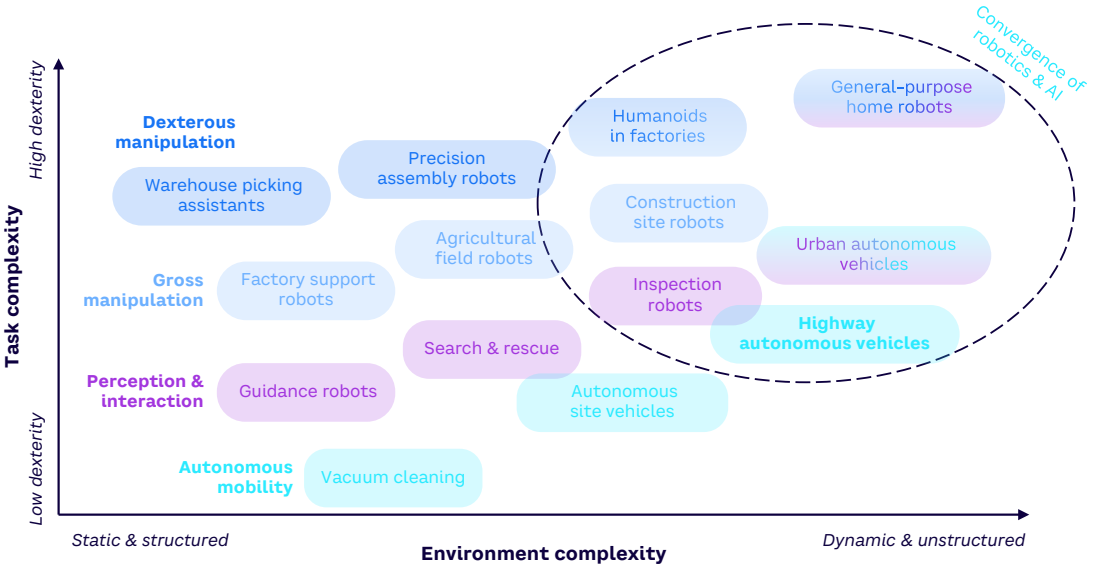
- **Physical AI refers to intelligent and autonomous machines that can perceive, decide, and act in the physical world.** Physical AI sits at the convergence of AI and robotics (see Figure 1) and combines software capabilities such as perception, planning, learning, and decision-making with robotic capabilities such as sensing, manipulation, motion, and control.
- **Physical AI is distinct from both conventional AI and traditional robotics.** Conventional AI has a theoretical understanding of the world. It can classify, predict, recommend, generate, or even decide, but it cannot act in the physical world. Traditional robotics can act in the physical world, but typically in narrow, highly structured, and rigid settings where tasks, sequences, and environments are carefully designed in advance. Physical AI brings real-world understanding to AI and versatility to robotics.
- **Physical AI is defined by embodied intelligence rather than any specific robotic form.** It includes intelligent fixed systems in factories, mobile manipulators in warehouses, autonomous vehicles (AVs) on roads, drones used for inspection and logistics, and robotic systems operating in factories. The defining shift is not the body, but the intelligence embedded within it.
- **Physical AI is a much broader field than humanoids, although humanoids currently dominate headlines.** A humanoid is a robot built to resemble a human, attempting to mimic the advantages of the human form and attributes, such as arms and hands. Figure 2 illustrates four broad categories of physical AI embodiment: (1) dexterous manipulation, (2) gross manipulation, (3) perception and interaction, and (4) autonomous mobility.

Figure 1. What is physical AI?



Source: Arthur D. Little

Figure 2. Four broad categories of physical AI embodiment



Source: Arthur D. Little; Prof. Roland Siegwart, Autonomous Systems Lab, ETH Zurich/EPFL

These categories can be mapped across two dimensions: **task complexity**, which drives dexterity requirements, and **environment complexity**, which drives autonomy needs.

THE MARKET OPPORTUNITY IS SIGNIFICANT & COULD REACH AS MUCH AS $12 TRILLION BY 2050

"Physical AI will dramatically transform manufacturing. It's not a question of if, it's a question of when."

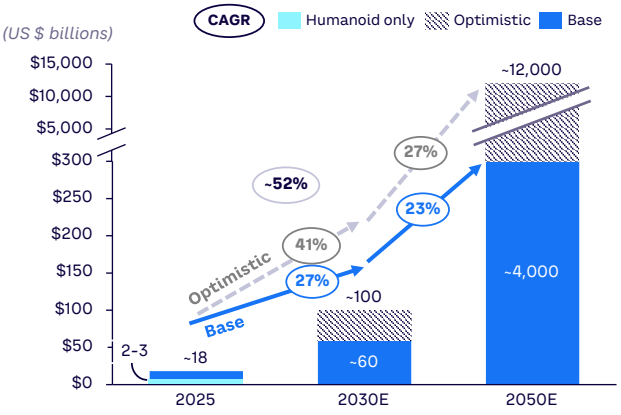
Paul Batsii, CPO, Robotics AI, Siemens

Figure 3 provides approximations of the current and future physical AI market size based on a range of sources.

Current market size is around $18 billion

- **Current market sizes are difficult to estimate due to unclear physical AI scope boundaries.** Available market data from reputable sources typically covers either the broader robotics market, of which only part is AI-driven, or humanoid robotics, which is only a fraction of the total physical AI market. The estimate of $18 billion is based on market size data for "intelligent robotics," expanded to include physical AI software platforms and autonomous mobility markets.
- **The current market for humanoid robotics is estimated at $2–$3 billion.** Humanoids currently account for only around 10% of the total physical AI market, reflecting their relative immaturity.

Figure 3. Global physical AI market size



Wide forecast range driven by differences in market scope, AI penetration assumptions, adoption timing & modeling approaches

Current 2025 market size estimate is around $18 billion,[1] of which humanoid robotics is $2–$3 billion

Near-term 2030 estimate is in the range of $60–$100 billion,[2] depending on different assumptions for overall robotics growth rate & AI penetration

Long-term 2050 estimate is in the range of $4–$12 trillion[3]; high divergence in forecast sources reflects different growth & adoption assumptions; optimistic forecast assumes widespread domestic deployment & broad labor substitution

Note: (1) Based on MarketsandMarkets, expanded 20% to include software platform and autonomous mobility scope; (2) base case assumes physical AI is 54% of ABI 2030 robotics estimate; optimistic case assumes similar % of more optimistic robotics forecasts (Mordor, GMI, etc.); (3) range based on multiple sources, including Barclays, Morgan Stanley, UBS, Deutsche Bank, Bank of America, McKinsey Global Institute, ABI Research, RBC Capital Markets, etc.
Source: Arthur D. Little

Near-term forecasts suggest range of $60–$100 billion by 2030

- **Rapid growth is predicted by nearly all near-term forecasts.** While growth rates and scope definitions differ, most reputable forecasts indicate very strong near-term growth in excess of 20% CAGR.

The opportunity could reach $4–$12 trillion by 2050 if adoption expands beyond industry

- **Most long-term forecasts (to 2050) predict market sizes above $4 trillion.** Although these forecasts are inherently uncertain, most assume rapidly increasing adoption after 2035.
- **The extent to which humanoid adoption expands beyond industry is critical.** Widespread humanoid deployment in domestic and other nonindustrial settings would greatly increase market size. Morgan Stanley estimates that the humanoid-specific market could surpass $5 trillion by 2050, including supply chains, repair and maintenance, and support services, and projects more than 1 billion humanoids in use by then.

Robotics VC investment is already at $28 billion; Figure AI's 1,200% valuation increase shows investor conviction

- **Robotics VC investment reached almost $28 billion,** across more than 1,000 deals in 2025, even as the broader venture market normalized. Private capital is reinforcing market value projections, with investment accelerating rapidly, and major players making significant commitments.
- **Figure AI's valuation rose 1,200% in less than a year.** Some industry voices point to 2026 as the year physical AI enters the investment mainstream. One clear example is Figure AI: the humanoid manufacturer's valuation surged from almost $3 billion to around $39 billion within less than a year, backed by Microsoft, Nvidia, and OpenAI.[1]
- **Investors are pricing physical AI as a future industrial and labor platform,** not merely against current revenues.

Private capital is reinforcing market value projections, with investment accelerating rapidly.

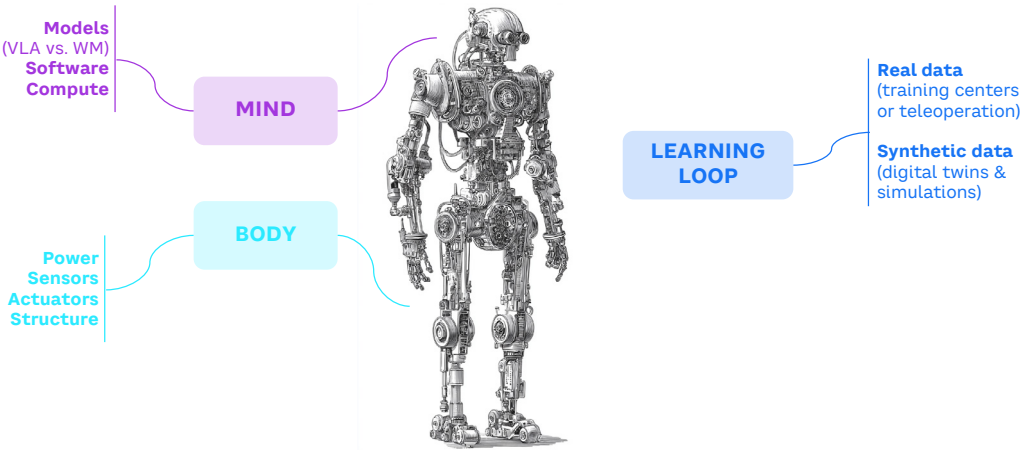
VALUE WILL BE GENERATED ACROSS A LAYERED ECOSYSTEM & TECHNOLOGY STACK WITH UNEVEN MATURITY

Value in physical AI will be generated across an unevenly mature technology stack and ecosystem composed of three layers: Body, Mind, and Learning Loop (see Figures 4 and 5). The value chain structure remains unsettled, with new roles emerging around platform integration, application design, and maintenance. As a result, partnerships, consolidation, and geopolitics will play a major role in shaping which players and regions lead the industry.

Three layers: Body, Mind & Learning Loop

5. **Body — the physical machine that does the job.** This layer determines what robots can physically do and includes actuation, power, sensors, and the overall robot body, defining how well a robot moves, manipulates objects, senses its surroundings, stores energy, as well as how it looks.

Figure 4. Three layers of physical AI's technology stack



Source: Arthur D. Little

Figure 5. Overview of technology readiness

COMPONENT		READINESS	COMMENT
Body	Power		• Battery technology has advanced through EV scale, but runtime, payload, charging & energy use still constrain mobile robots
	Actuators		• Mature component supply exists, but precise, compliant & energy-efficient actuation for dexterous physical interaction remains a core bottleneck
	Sensors		• Cameras, LiDAR, microphones & radar are broadly mature & ready for use, although tactile sensing remains less developed
	Robot structure		• Robot bodies are increasingly buildable, but form-factor fit, stability, weight, robustness & task-specific design remain use-case dependent
Mind	Software		• Data, analytics, control & orchestration software are advancing, but integration across robots, workflows & safety layers is not yet standardized
	Models		• AI models & world models are promising, but broad real-world reasoning is immature; hybrid architectures are more realistic than one giant model
	Compute		• AI compute is powerful, but robots need low-latency onboard reasoning within battery, heat, weight & safety constraints
Learning Loop	Real data		• The weakest layer: physical AI lacks large, reusable action & tactile datasets that link perception, action, force & physical outcome
	Synthetic data		• Simulation is useful for vision, motion & virtual testing but remains weaker for contact, friction & manipulation; sim-to-real is still incomplete

High Medium Low

Source: Arthur D. Little

- **High maturity — movement.** Robots can now move, localize themselves, and avoid obstacles far more effectively than before.
- **Low maturity — dexterity.** Robots still struggle with tasks requiring reliable grasping, touch sensitivity, slip adjustment, and fine force control. This matters because even advanced software will fail if the robot cannot move precisely or sense the environment well enough to operate effectively.

6. **Mind — the intelligence layer that gives robots versatility.** It includes foundational models such as vision-language-action (VLA) and world models (WM), which support perception, scene understanding, physical reasoning, and planning; orchestration and control software connecting models to sensors, robots, workflows, and the physical world; and compute infrastructure, including GPUs, edge compute, accelerators, networking, memory, and inference hardware. For the Mind layer to create value, the robot's other layers must support precise and safe operation in the real world. Different tasks require different forms of intelligence. Deterministic systems provide reliable, repeatable movement and form the base, while learned AI models handle "last mile" actions by adapting the robot's movement when situations change. This hybrid approach combines reliable execution with adaptive flexibility. Although industry experts disagree on the winning architecture, many view this hybrid approach as more promising than relying on "a single giant model doing everything."
   - **High maturity — perception, pattern recognition, and planning.**
   - **Low maturity — reasoning and execution in live environments.**
7. **Learning Loop — the systems through which robots learn from the physical world.** It captures data and uses it to train and improve robots' capabilities over time. Training data can come from either the real world or synthetic simulation:
   - **Real-world data** — training data captured from real-world interactions, such as a film of humans performing a given task.
   - **Synthetic data** — training data generated through simulation of robots operating in virtual environments.
   - **Lowest maturity of the stack** — learning and training remain the least mature parts of the technology stack. (See sidebar for additional discussion of the role of simulation in learning and training.)

Training data can come from either the real world or synthetic simulation.

"The biggest challenge is developing a workable world model … current world models are really focusing on geometry and spatial location, but I'm concerned about physical properties."

Prof. Neville Hogan, MIT

The importance of simulation for physical AI in industry

Insights from Shane McArdle, CEO & President, Falkor (formerly Kongsberg Digital)

For Falkor's CEO, the starting point for physical AI in heavy industry is not robots, but simulation, physics, and operational context. Industrial AI cannot rely on black-box models alone. As McArdle puts it, "If we are 90% accurate in our industry, we are 100% wrong," given the safety and operational consequences.

Successful physical AI requires a hybrid architecture: fast probabilistic AI models working in combination with slower, but trusted, deterministic physics-based simulation. Falkor's role is to provide the industrial platform layer that connects real-time operational data, engineering models, spatial context, simulation, and agentic AI — not as separate systems, but as one coherent environment.

The most credible near-term value lies not in full autonomy, but in transforming workflows and automating tasks.

Its partnership with Nvidia is built around exactly this idea. Using Omniverse libraries alongside live plant data, simulation models, and an agentic AI framework, Falkor's goal is to move decisively from visualization to intelligence, so the behaviors of the real industrial world can be understood. That means not only seeing what a digital twin shows, but also reasoning about what is happening, what could happen, and what actions should be taken.

The most credible near-term value lies not in full autonomy, but in transforming workflows and automating tasks — including through the use of robots. Robots such as drones and underwater AVs already create value because they are purpose-built for inspection and data gathering. In dynamic, high-consequence environments, humans should remain in the loop until trust is genuinely earned. Over time, that model may shift to humans on the loop, managing robots, agents, and tasks by exception — but that transition must be built on demonstrated reliability, not ambition.

The foundation for all of this is the same: understanding how assets are actually operated and maintained. Organizations must build on clean, relevant data linked to real workflows, physics, and trusted operational knowledge. As McArdle says, "You can't put intelligence on top of chaos."

With physical AI advancing this quickly, the question is no longer simply what the technology can do. It is whether organizations truly understand how their assets are being operated and maintained — and if they don't, they shouldn't start with AI.

Value chain takes shape; roles not yet fixed

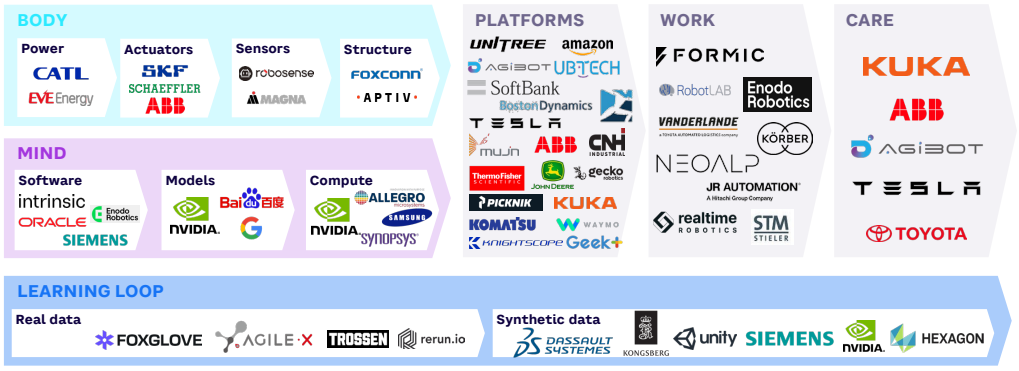
The boundaries of the value chain are still evolving, and players are defining their value-add. The value chain comprises the technology stack (Mind, Body, and Learning Loop), as well as Platform, Work, and Care layers (see Figure 6):

1. **Platform — where players build the robot and deliver the core system.** This layer combines Body and Mind into a working machine, either through in-house integration or supplier ecosystems, and provides the robot's baseline capabilities. In this sense, it delivers the product itself, but not always the full customer solution around it.
2. **Work — where the value chain meets the buyer.** This layer bridges vendors and customers and is enabled by application design players focused on designing use cases, defining routines, and supporting deployment to make machines productive in day-to-day operations. Likely winners in deployment will build platforms that application design players — such as local partners, developers, and integrators — can adapt to the "last mile."
3. **Care — where the system is kept running after deployment.** This layer includes spare parts, maintenance, software updates, troubleshooting, and eventually decommissioning and recycling. In the early market, the platform players and integrators will likely handle this. Over time, however, it may become a more distinct part of the value chain.

"The value chain is still fluid because the technical architecture itself is not settled."

Thibauld Jongen, Cofounder, Common Sense Robotics

Figure 6. Physical AI value chain, including player examples



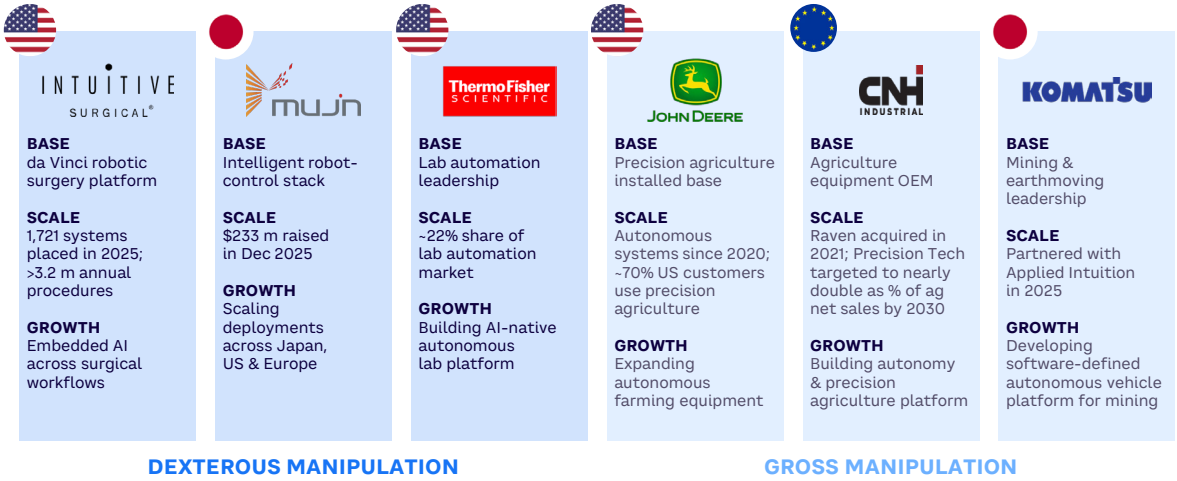
Note: The list of players is illustrative and not exhaustive
Source: Arthur D. Little

A phase of collaboration between traditional automation players and AI experts is likely. Partnerships, acquisitions, and hybrid business models between traditional automation players and AI experts are rational paths forward. Traditional automation companies bring engineering and safety expertise, while newer AI players contribute models, cloud infrastructure, and software capabilities. Few companies are strong in both areas at the same time. As a result, the next phase may reward companies that can build partnerships combining complementary strengths quickly enough to achieve market relevance.

A consolidation phase is likely in the near future. The high number of companies reflects the amount of capital and expectation that has flowed into the space rather than the number of players likely to secure a durable place in it. In China alone, there are currently around 300 automation companies, most of them focused on humanoid robots. This level of fragmentation is unlikely to be sustainable once commercial realities begin to outweigh market enthusiasm and investment-driven momentum. As deployment realities become clearer, the market will likely consolidate around players that can combine technical capability, real customer value, and a credible path to scale.

Manipulation is one of the first major physical AI value pools, as automation leaders extend existing platforms into higher-value tasks. In the dexterous and gross manipulation categories, including precision manufacturing, medical, logistics, agricultural, and construction robotics, physical AI represents a major growth platform for both established players and startups, often building on many years of automation development. Figure 7 provides selected examples for illustration.

Figure 7. In manipulation categories, automation leaders are extending into higher-value AI-driven platforms



Note: The list of players is illustrative and not exhaustive
Source: Arthur D. Little

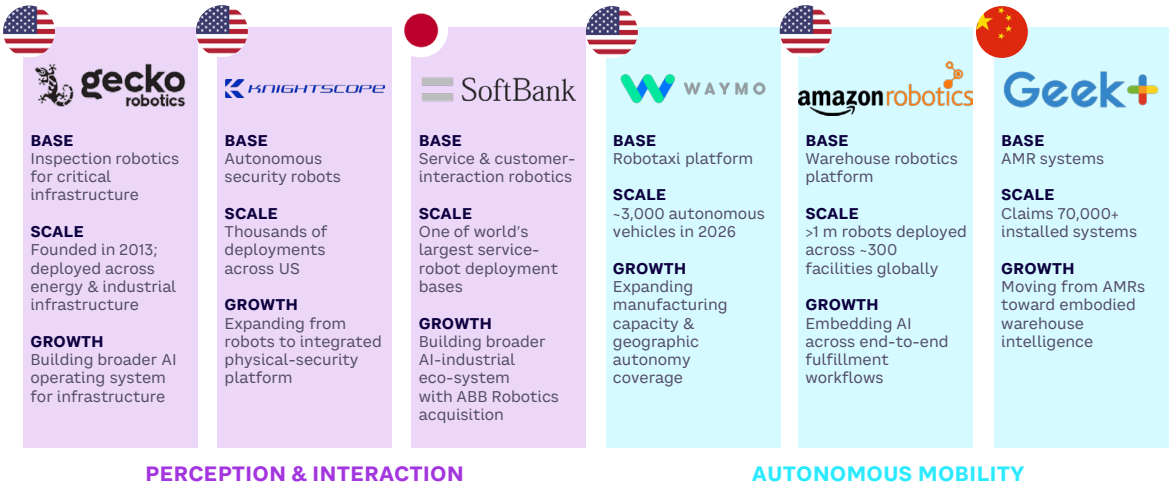
Perception and mobility players are evolving from robots into autonomous operating platforms. In the perception and interaction and autonomous mobility categories — including inspection, service, security, delivery, cleaning, autonomous mobile robots (AMRs), warehouse mobility, robotaxis, and autonomous delivery vehicles — many players are shifting from selling robots to offering autonomous end-to-end operating platforms (see Figure 8).

Humanoid vendors are moving from prototypes to early production, but industrial-scale deployment remains largely unproven. Although detailed production plans are generally unavailable, there are growing signs that leading humanoid players such as UBTECH, Unitree, AgiBot, Tesla, Boston Dynamics, and Figure AI are accelerating toward mass production within the next one to two years (see Figure 9).

"Humanoids sold last year are still largely used for demos, entertainment, and POCs, with still a long way to go."

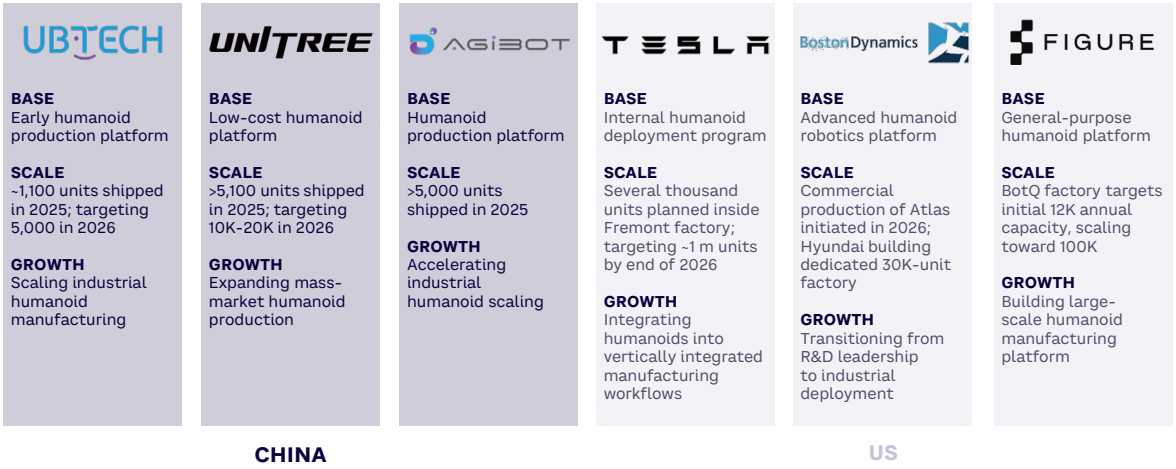
William Shi, CEO of Europe & Americas, AgiBot

Figure 8. Physical AI is a major growth platform for perception and interaction/autonomous mobility categories



Note: The list of players is illustrative and not exhaustive
Source: Arthur D. Little

Figure 9. Mass production of humanoid robots is commencing



Note: The list of players is illustrative and not exhaustive
Source: Arthur D. Little

GEOPOLITICS SHAPES PHYSICAL AI THROUGH DIFFERENT REGIONAL STRENGTHS, WEAKNESSES & STRATEGIES

Physical AI is increasingly viewed as a strategic macroeconomic capability by states and regions. The market is therefore driven not only by customer demand and technical progress, but also by policy, national ambition, and fear of falling behind. Key states and regions are seeking to attract capital, research, and deployment activity. However, they possess different strengths, weaknesses, and strategies for gaining advantage, with China currently leading in several dimensions (see Figure 10).

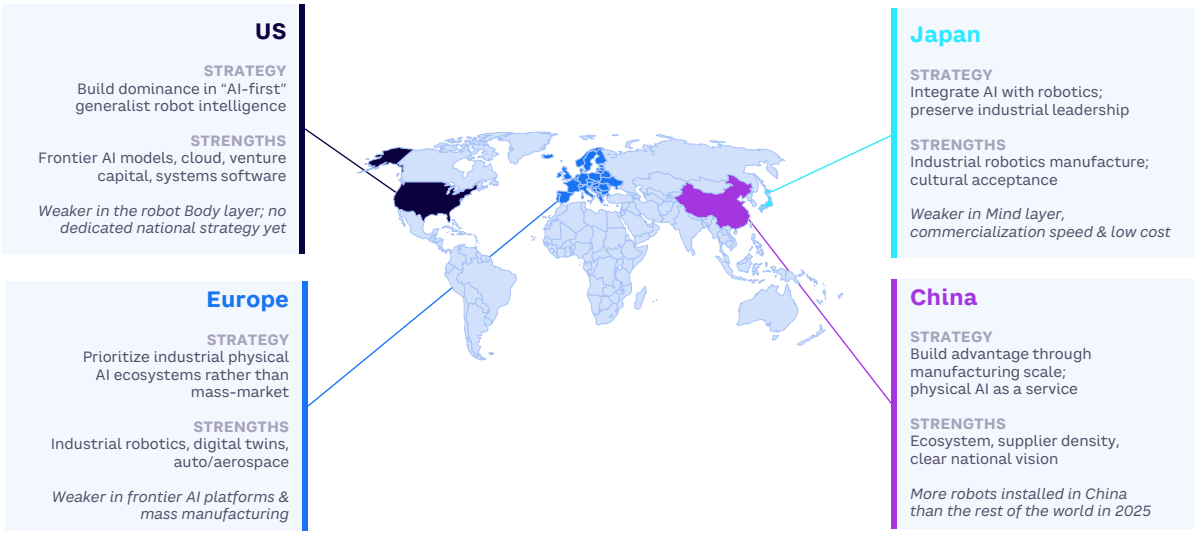
US is strongest in Mind layer but weaker than China in Body & manufacturing at scale

Key strengths in the US include frontier AI models, cloud/compute infrastructure, VC, systems software, and high-value robot integrators. Large players including Nvidia, Tesla, Figure AI, Boston Dynamics, Agility Robotics, and Apptronik act as anchors in an extensive ecosystem. The US is comparatively weaker, especially relative to China, in the Body of the robot, such as cost-effective mass manufacturing, actuators, reducers/gearing, motors, sensors, batteries, and dense supplier ecosystems.

The US strategy focuses on "AI first" over explicitly physical AI:

- **The US aims to lead in "generalist" robotics intelligence.** The White House 2025 AI Action Plan seeks to "win the AI race" by reducing barriers, building domestic AI infrastructure, and establishing leadership in AI standards and security. Meta, Nvidia, Google DeepMind, and multiple startups are investing heavily in robotics foundation models.

Figure 10. Regional strengths and strategies for gaining physical AI advantage



Source: Arthur D. Little

- **Robotics manufacturing is also an investment target, although there is still no dedicated national robotics strategy.** The Action Plan highlights measures intended to position the US as a world-class manufacturer of AI-related technologies, including autonomous drones, self-driving cars, robotics, and related systems. Industry groups such as the Association for Advancing Automation continue to advocate for a dedicated national robotics strategy.
- **Skills shortages are increasingly viewed as a constraint.** The National Robotics Association estimates that open robotics-related positions exceeded 340,000 by the end of 2025, a 42% increase over 2024 levels.[2] Specialized skills in robot operating systems, computer vision, and embedded systems are identified as the most acute shortages.

Physical AI is increasingly viewed as a strategic macroeconomic capability by states and regions.

The US trails China in deployment at scale, with near-term focus on manufacturing and logistics. Use cases and deployment examples in the US include:

- **In 2024, the US installed roughly 34,200 industrial robots,** compared with approximately 295,000 installed in China.[3] Morgan Stanley estimates that Chinese companies account for 63% of the global humanoid robot supply chain.[4]
- **According to data from The Robot Report, the largest near-term deployment growth in the US will likely be in AMRs.**[5] These are generally wheeled rather than humanoid and may include dual-arm manipulators and collaborative robot (cobot) systems. This growth builds on the already extensive robotic footprint across warehouses, fulfillment centers, hospitals, and factories.
- **Logistics remains a key deployment area.** For example, Agility Robotics's Digit humanoid robot is already operating at scale in GXO Logistics facilities following a two-year trial and is also being tested by Amazon.
- **Manufacturing handling operations are another priority.** Humanoids from Agility Robotics, Figure AI, Tesla, and Boston Dynamics are being piloted for repetitive material handling, palletizing, tote movement, parts delivery, and factory support tasks.
- **Infrastructure maintenance is also important.** Companies such as Gecko Robotics offer AI-enhanced robots for inspection and maintenance monitoring in the oil & gas and utility sectors.
- **Healthcare is another priority area.** AI-enhanced robotic systems are being trialed to anticipate and adjust surgical interventions in real time, including systems such as Intuitive's da Vinci.

China builds industrial advantage through manufacturing scale, supplier density & coordinated research

More robots were installed in China during 2025 than in the rest of the world combined.[6] China is moving rapidly from "AI in the cloud" to "AI in steel." Many in the West have yet to fully grasp the speed at which China is developing physical AI. China has established a clear national vision for industrial and physical AI, aimed at scale and strategic autonomy. More than 80% of humanoid robot installations in 2025 came from Chinese firms, led by UBTECH, Unitree, and AgiBot.[7] The China strategy builds AI and physical AI as a service:

- **AI and physical AI are framed as strategic industrial infrastructure,** not merely as technologies: China's AI ambition is not limited to chatbots or generic models; it focuses on industrial productivity, acceleration, automation, and robotics.
- **A shrinking population is a main driver.** Robots have long been a focus for the Chinese Communist Party, partly due to the country's shrinking population and a growing need for factory workers.
- **Humanoids act as a moon shot narrative to mobilize talent, capital, and attention.** Humanoid robots appear less as an immediately deployable industrial solution than as a flagship ambition used to structure the broader physical AI ecosystem.
- **The endgame is scale and strategic autonomy.** China's model aims to reduce dependency, build domestic capabilities, and industrialize AI across strategic sectors.

Many in the West have yet to fully grasp the speed at which physical AI is developing in China.

Physical AI in China covers the whole spectrum of use cases, ranging from accelerating molecule discovery to humanoids:

- **XtalPi** — coupling AI and physical AI to discover new pharmaceutical molecules through industrialized experimentation: fewer human errors, faster cycles, robotic labs, and platformized discovery capacity.
- **AgiBot** — humanoids remain largely in demos and proofs of concept, with real deployment still constrained by cost, balance, data, and use-case economics.

China has established a clear path from "0 to 100" through deployment discipline:

- **Everything is being packaged "as a service."** Compute, AI, automation, manufacturing, and marketing are being turned into scalable service platforms.
- **Deployment starts with data factories.** Data is treated as an industrial asset. China claims to have hundreds of training centers/data factories to collect, clean, structure, label, and package data from the physical world for robots.
- **China's key advantage is its ecosystem and cluster approach, combined with a strong public-private research loop.** Manufacturing scale, dense supplier networks, supply chain overlaps with EV factories, and the ability to quickly generate product upgrades together allow even "good enough" robots to improve faster than Western competitors expect.

Europe is strong in industrial systems but weak in frontier AI platforms & mass manufacturing

- **Europe lags the US and China in scale but can leverage strengths in industrial systems.** Compared to the US, Europe lacks world-leading capabilities in frontier AI models, hyperscalers, and late-stage VC. It also lacks China's low-cost manufacturing ecosystems. However, it has strengths in industrial robotics, cobots, industrial software, digital twins, automotive automation, aerospace systems, high-precision manufacturing, and applied AI for factories. Industry experts see an opportunity for Europe to transition legacy industries, such as automotive and manufacturing, into the physical AI market.
- **EU policy focus on AI risks is now shifting toward industrial support and regulatory harmonization.** Recently, the European Commission launched a €20 billion (~US $23 billion) plan for AI gigafactories with advanced supercomputing capacity. The EU's broader AI Continent strategy aims to accelerate AI adoption across manufacturing, healthcare, automotive, and public services while reducing dependence on US and Chinese platforms.
- **An EU AI-powered robotics strategy is planned for this year,** intended to align member states on robotics deployment and standards.

Industry experts see an opportunity for Europe to transition legacy industries.

Europe prioritizes industrial physical AI ecosystems rather than mass-market robotics in use cases and deployment:

- **The near-term priority is AI-assisted flexible factories** where robots can handle repetitive, physically demanding, or variable tasks. These may be wheeled or hybrid rather than humanoid. For example, Siemens and startup Humanoid announced successful deployment trials of the HMND 01 Alpha robot in Siemens's Erlangen Electronics Factory, performing autonomous logistics tasks using Nvidia's physical AI stack.
- **Industrial manufacturing, especially automotive, is key.** For example, BMW and Hexagon Robotics are deploying the AEON humanoid robot to support automotive production, battery assembly, and component manufacturing.
- **Another major deployment area is warehouse logistics and intralogistics.** For example, DHL announced in 2025 that it would deploy more than 1,000 additional robots across UK operations to support e-commerce and healthcare logistics growth.
- **Healthcare, elderly care, and infrastructure/energy systems are also priority areas.** European robotics strategies consistently identify healthcare, rehabilitation, assistive robotics, and social robotics as major growth categories. Startups such as Wandercraft are extending robotic mobility systems into industrial humanoids and physical assistance platforms. AI-enhanced robots for energy, rail, and other infrastructure are also priority areas.

Similar to Europe, Japan is strongest in the Body layer & industrial systems

Japan is still dominant in industrial robotics manufacturing but weaker in the Mind layer, commercialization speed, and low-cost manufacturing. Through firms such as Fanuc, Yaskawa, Kawasaki Robotics, Omron, and Mitsubishi Electric, Japan still holds around 70% of the world's industrial robot market share. However, it is weaker than the US in frontier AI models, hyperscale compute, cloud platforms, and AI software ecosystems and increasingly faces pressure from China in low-cost robotics manufacturing and commercialization speed. Given its culture and recent history, Japan's strength could lie in acting as a proving ground for physical AI adoption in a symbiotic society where physical AI, teleoperated robots, and humans work and interact together.

AI robotics in Japan is seen as essential national infrastructure:

- **AI robotics has a renewed urgency.** In light of labor shortages and aging demographics, Japan declared in March 2026 that it aims to build a domestic physical AI sector and capture a 30% share of the global market by 2040.
- **The Japanese strategy emphasizes three priorities:** (1) integrating generative AI with robotics, (2) accelerating deployment in labor-short sectors, and (3) preserving Japan's industrial leadership.

AI robotics in Japan is seen as essential national infrastructure.

Japan focuses on logistics, manufacturing, retail, transport, healthcare, and others. Several examples illustrate Japan's broad long-term focus:

- **Logistics.** Mujin offers a software-defined AI robotics platform for warehouses and factories, including picking, palletizing, and warehouse automation.
- **Manufacturing.** Nvidia and Fujitsu announced a 2025 partnership to build Japanese AI infrastructure for healthcare, manufacturing, the environment, next-generation computing, and customer services.
- **Retail.** Seven-Eleven Japan and Telexistence announced a 2025 partnership to develop Astra, a humanoid robot using a vision-language-action foundation model, targeting store deployment in 2029.
- **Transport.** Beginning May 2026, Japan Airlines is piloting humanoid robots at Haneda Airport for baggage/cargo loading, unloading, and cabin cleaning.
- **Healthcare.** METI (Ministry of Economy, Trade and Industry) lists long-term-care robot technology as a formal policy area.

HUMANOID NARRATIVE OVERSTATES GENERAL–PURPOSE READINESS

"The spring festival gala robots and the Unitree kung fu style performances show progress and ambition indeed ... but these systems are still preprogrammed or tightly scripted."

Prof. Marina Zhang, University of Technology Sydney

Humanoids dominate the narrative because they make physical AI look like a universal labor platform

- **Humanoids help tell the physical AI story.** Humanoids make the story easier to imagine, market, and fund. A robot shaped like a person suggests a machine that can work where people work, use what people use, and eventually augment or replace a wide range of human labor. However, the storytelling power is stronger than the current commercial proof.
- **Humanoids support the prospect of a bigger market.** Many actors reinforce the humanoid story because they benefit from a larger market narrative. Investors and large tech companies prefer the universal robot story because it supports larger funding rounds, stronger valuations, and the idea that robotics is the next major AI platform.
- **Hype accelerates but distorts learning.** This hype is useful because it attracts talent, capital, executive attention, and government interest into a field with long development cycles. But it can also distort learning by shifting attention away from the less visible work needed for adoption: field performance, safety, site integration, operating logic, and evidence of business value.

"Legs create attention, arms earn money."

Georg Stieler, Head of Robotics & Automation, Stieler Technology

Humanoids make the story easier to imagine, market, and fund.

Demonstrations can be mistaken for deployment readiness

- **Demonstrations are overread as readiness.** Visible progress is too quickly interpreted as proof that a new baseline capability has arrived. Robot videos online often mix scripting, teleoperation, human triggering, and heavily controlled conditions while being interpreted as evidence of broad autonomy. Adoption depends on much more than a successful demo: a robot still must prove that it can operate safely, repeatedly, and economically in a live setting.
- **Structured tasks are ready first.** Deployment is strongest where tasks, flows, and constraints are well-defined. In the foreseeable future, the market will focus first on structured environments such as factories, logistics sites, warehouses, mobility systems, retail stores, and selected service settings, where tasks and constraints are predictable and trained staff are near. Structured task execution is already producing credible value.

"Humanoids sold last year [were] largely used for demos, entertainment, and POCs, with still a long way to go."

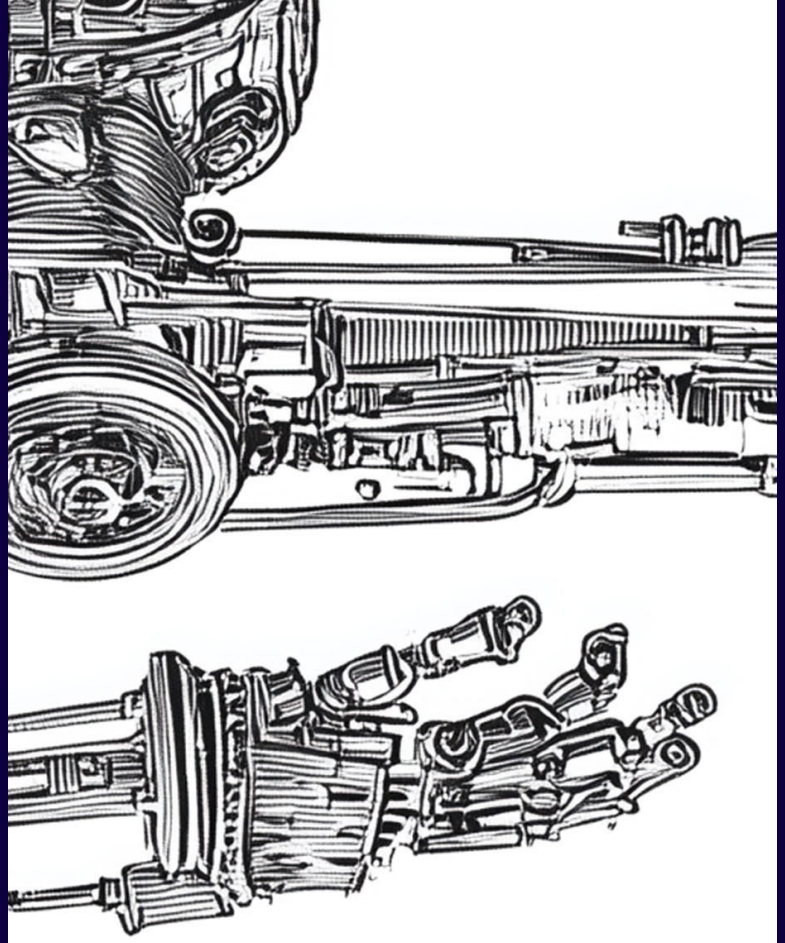
William Shi, CEO of Europe & Americas, AgiBot

Visible progress is too quickly interpreted as proof that a new baseline capability has arrived.

“Of course you were programmed, by nature or nurture or both.”

– Ex Machina

2



CHAPTER

VALUE WILL FIRST COME FROM SPECIALIZED PHYSICAL AI

2

VALUE WILL FIRST COME FROM SPECIALIZED PHYSICAL AI

Several strong drivers, such as labor shortages, shifting work preferences, and resilience needs, are accelerating demand for physical AI. However, key technological, economic, and deployment barriers are currently constraining rapid scale-up. The sweet spot for initial deployment is in structured environments with a degree of variability, favoring specialized applications in manufacturing, logistics, agriculture, and mobility. Applications in service and support, construction, mining, healthcare, and pharmaceuticals are also promising.

STRONG DRIVERS ACCELERATE PHYSICAL AI ADOPTION

Adoption is being propelled by both demand-pull and supply-push dynamics (see Figure 11).

Demand–pull factors

- **Labor shortages are becoming a structural issue.** In many major countries, labor shortages will likely become a long-term structural, rather than cyclical, problem. For example, the OECD (Organisation for Economic Co-operation and Development) forecasts an average drop of 8% in the working-age population across its members by 2060, including more than 20% for Germany and 30% for Japan.[8] Falling birth rates, aging populations, and continued downward pressure on labor availability are key drivers of the shortage. Physical AI could ultimately become a key solution to offset the shortfall.
- **Work preferences are shifting away from repetitive roles.** In industries such as manufacturing, logistics, and care, many companies are struggling to recruit and retain people for roles that are repetitive, physically demanding, low-status, or have night shifts, even with higher pay. For example, in the UK in 2025, the highest worker shortages were in transportation and storage, health and social work, and accommodation and food services.[9] Around 450,000 manufacturing jobs in the US remained unfilled in 2025. Surveys show younger Americans increasingly view factory jobs as unattractive and physically demanding. Many of these jobs are in the sweet spot where physical AI can create value, especially in tasks such as heavy lifting or cleaning. Despite understandable fears from employees, many employers see the next generation of robotics as enabling humans to focus on other tasks through job redesign rather than being substituted entirely.

"Young people do not want to work in an industry that mostly operates at night and in the evening."

Christian Oestergaard, Lead Visionary & Senior Group Strategist of IT-Production/IoT/AI, PostNord

- **Operational resilience is rising in importance.** Multiple factors increase the need for operational resilience in large companies, including geopolitical fragility, supply chain chokepoints, economic volatility, and digital dependence. Recent surveys show that between 60%–80% of business leaders believe volatility and risk have increased. Even without these macro pressures, many operations must deal with last-minute staffing gaps, seasonal spikes, fluctuating demand, supply variability, and tighter service expectations. Responding effectively to sudden disruptions is often constrained by limited human resources. Physical AI can help operations function effectively in times of crisis. For example, physical AI is now being trialed in logistics to improve operational stability.

Figure 11. Adoption drivers for physical AI



Source: Arthur D. Little

- **Productivity pressures are increasing.** Rising labor costs put pressure on company margins and challenge existing ways of working. To stay competitive, companies must improve productivity and automate a much wider range of tasks. This pushes businesses to use labor more efficiently and focus employees on higher-value activities.

Supply-push factors

- **Task-specific pragmatic workarounds for dexterity are being developed.** Instead of trying to replicate the full versatility of the human hand, developers are building workarounds using simplified grippers and interchangeable end effectors that are reliable enough for specific tasks. This pragmatic approach expands the range of feasible applications, even though broader dexterity challenges remain.

"When we're designing the machine for a task, why would we limit it to mimic the human form factor...? We can optimize the end effector for a particular job."

Dave Grant, CEO, PickNik Robotics

Hardware is becoming more standardized, lower-cost, and easier to connect.

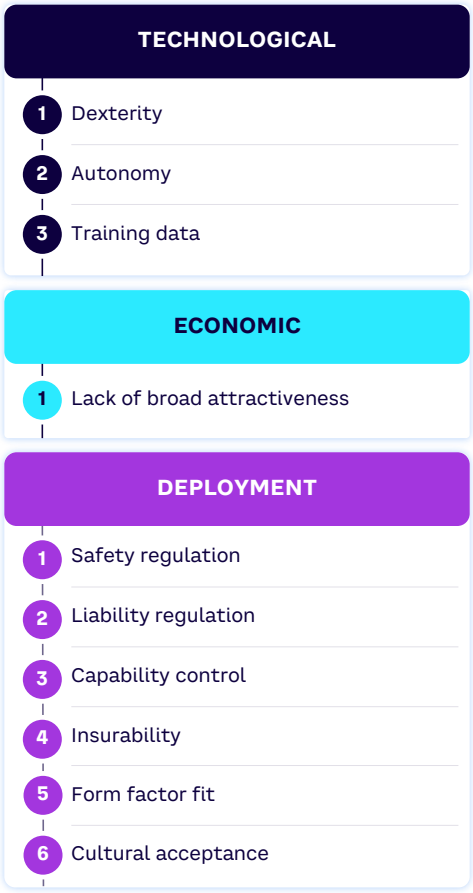
- **Hybrid systems are improving autonomy.** Similarly, autonomy capabilities, while still limited in their understanding of real-world surroundings, are improving enough to begin deployment in structured environments. Hybrid architectures promise a way forward: they combine classical rule-based deterministic approaches for repeatable tasks with trained AI models that enable adaptability to new situations. For example, a robot can be programmed with rules for navigating a factory floor while using AI to analyze the objects it handles, enabling greater flexibility in task execution.
- **Commoditization of hardware is lowering adoption thresholds.** Hardware is becoming more standardized, lower-cost, and easier to connect, enabling manufacturers and other industry players to use hardware more as plug-and-play solutions. Robotic arms, cameras, and mobile platforms are no longer rare or ultra-expensive, making it easier to assemble task-specific systems that fit each worksite's process at a reasonable cost. For example, light cobot arms weighing around 5 kg and costing $25,000–$35,000 a decade ago can now be bought from Chinese suppliers for less than $10,000. This not only lowers cost; it enables more companies to experiment with physical AI and helps address technical barriers such as autonomy and dexterity.
- **An emerging middle tier is bridging the gap between producers and buyers.** New players (e.g., NEOALP, EndoRobotics, Realtime Robotics) are emerging at the application design layer in the ecosystem, providing a bridge between robot manufacturers, software providers, and buyers. These players can offer workflow integration, fleet orchestration, and context-specific application knowledge, lowering adoption risk for buyers and accelerating time to value.

"The biggest challenge right now is we see a lot of customers that are extremely motivated to start with humanoid robotics. But it's totally unclear to them how to do this."

Christian Tauber, Cofounder & CEO, NEOALP

- **New commercial models are lowering the purchase threshold.** Emerging commercial models such as robot as a service and expanded after-sales offerings are becoming important features in the industry's rollout phase. These models aim to reduce up-front capital expenditure and commercial risk by shifting the value proposition from selling a machine to delivering a working service. The likely winners may therefore be defined not only by robot performance but also by who can wrap the most attractive and practical service, financing, deployment, and maintenance model around it.
- **Integrated human/robot operations are unlocking faster adoption.** There is growing recognition that physical AI can create value in hybrid environments where robots handle structured physical tasks while humans focus more on supervision, exception handling, and judgment. A robot does not need to replace a worker end to end; it can create value as soon as it reduces operational strain, improves continuity, or takes over more dangerous tasks. This means buyers do not need to wait for fully autonomous robots before adoption while also supporting a more positive narrative for employees.

Figure 12. Adoption barriers for physical AI



Source: Arthur D. Little

ADOPTION AT SCALE REMAINS CONSTRAINED BY TECHNOLOGICAL, ECONOMIC & DEPLOYMENT BARRIERS

Despite rapid progress, mass adoption of physical AI remains constrained by technological, economic, and deployment-related barriers (see Figure 12).

Technological barriers include insufficient dexterity, autonomy & training data

- **Human-hand dexterity is still unmatched.** The complexity of a physical interaction determines the required level of dexterity. Tasks requiring fine manipulation, force control, and handling different surfaces and frictions demand higher dexterity. Dexterity can be measured in degrees of freedom, with the most advanced devices on the market at the beginning of 2026 having around 20–22 degrees.[10] Although this is impressive, the human hand has 27 degrees of freedom and far greater adaptability than today's robotic hands. As a result, fine-manipulation tasks involving grip, touch, force control, or small adjustments remain difficult. Many applications, including most household tasks, require these capabilities.

"To create a robot that can make a cup of tea in any household, filling any kettle from any tap, extracting a teabag from any container, getting a cup from any cupboard and milk from any fridge, is a very difficult problem."

Stefano Bensi, General Manager, Softbank Robotics EMEA

The predictability of the environment shapes the required level of autonomous decision-making.

- **Human-level autonomy is still unmatched.** The predictability of the environment shapes the required level of autonomous decision-making. Limited autonomy can be sufficient in structured settings such as factories with clear partitions, flat floors, controlled processes, known object locations, and mostly repetitive tasks, while greater autonomy is needed in less predictable environments such as construction sites, roads, households, mines, or agricultural settings. Current VLA and transformer-based approaches can learn useful correlations from demonstrations, but higher autonomy will require stronger world models, causal reasoning, and recovery logic to handle edge cases beyond the training data. Autonomy requirements increase when robots must reason about context, force, friction, object behavior, and physical interaction.
- **Robot training data is still hard to generalize.** Unlike large language models and other AI systems that can access abundant online training data, physical AI requires training data about the physical world and how to interact with it, which is much more limited and harder to access. This makes learning slower, more complicated, and harder to generalize across tasks and environments. Humans can teach robots how to interact through teleoperation and simulations based on previously generated and digitized action data. However, each deployment site has unique characteristics that require adjustment of the robot's control and task orientation.

"Synthetic data already works very well for perception. The challenge starts when contact matters: friction, tactile sensing, grasping, and surface interaction are still difficult to simulate with enough quality."

Brian Gerkey, CTO, Intrinsic, and Board Chair, Open Robotics

Economic use cases for physical AI are not yet attractive at broad scale

- **The "killer app" is not yet apparent.** Robotics is still relatively expensive, and uptime remains uncertain, so each viable use case must deliver significant efficiency, productivity, and/or cost-reduction benefits to achieve an attractive ROI. The current physical AI market lacks a true killer application with broad applicability and compelling economics that would justify wide-scale rollout. Adoption rates are therefore limited by the speed at which individual use cases can be economically justified.

"All we need for robotics is a killer app. It really is the thing that makes the technology worth paying for."

Prof. Neville Hogan, MIT

Physical AI has the potential to cause harm to humans and damage to the environment.

Deployment barriers, particularly governance & integration, are slowing adoption

- **Safety regulation lags behind technical progress.** Physical AI has the potential to cause harm to humans and damage to the environment and therefore requires dedicated safety regulation. The stakes are extremely high, as even a single fatality or serious injury caused by a robot could significantly delay deployment or halt it altogether, and would likely be viewed as unacceptable by society. Deployments are already moving into uncontrolled environments before safety rules and regulations are fully established. In most regions, existing regulations mainly cover conventional workplace and product safety. The EU's 2024 AI Act has an indirect impact by imposing stricter requirements on higher-risk AI systems. In addition, anticipated 2027 EU machinery regulations will introduce requirements for CE (Conformité Européenne) marking, risk assessments, and robot safety compliance.[11] The US does not yet have a federal robotics law. China, by contrast, introduced the HEIS (Humanoid Robotics and Embodied Intelligence Standard System) 2026 regulations in February 2026, covering AI models, hardware, safety, ethics, and lifecycle management.

"For humanoids, safety governance is lagging deployment: there is still no clear guidance for how to assess risk before these robots move further into public and industrial environments."

Elena Messina, Principal, Prospicience LLC

- **Liability frameworks are not yet established.** Clear rules do not yet exist to allocate responsibility for malfunctions across AI model vendors, robot manufacturers, application developers, system integrators, and operators. Cybersecurity risks are also higher for physical AI than for other digital systems that do not interact directly with the physical environment.
- **Capability control poses unresolved governance challenges.** Defining and controlling what physical AI systems are permitted to do raises difficult governance issues. Physical AI robots may gain new capabilities over time or be redeployed into settings beyond those originally certified as safe. Abnormal or unexpected situations may also cause robots to behave in unpredictable ways, and they may not always understand when they are acting outside their competence range. Demonstrating safety across all operating conditions remains challenging.

- **Reliable performance standards are lacking.** Buyers still lack widely accepted independent standards for validating a robot's promised performance. Instead, buyers must rely on vendor claims regarding factors like battery life, precision, tolerance, and reliability. Organizations such as ISO (International Organization for Standardization) are working on performance criteria and testing methods, but standards for intelligent and humanoid robots remain at an early stage. These frameworks will become increasingly important for buyers wishing to invest in physical AI fleets at scale.
- **Insurability is not yet well established.** The absence of mature standards and regulation around safety, performance, and liability creates challenges for insurers. Humanoids are especially difficult to insure because there is limited operational history, and even once certification frameworks emerge, acceptance in uncontrolled settings may take time. In Japan, some robot insurance products already exist, but they remain limited in scope and expensive. In the meantime, buyers may need to accept greater operational risk and focus on mitigating it through design and operational controls.
- **Establishing the best-fit form factor is still immature.** More time is needed for developers and industrial end users to determine the best robot morphology, or form factor, for different deployment settings. Humanoids may suit certain general-purpose tasks, such as warehousing or facility support, but in many settings, simpler form factors, including wheeled robots, multi-arm systems, or upgraded industrial robots, may be more suitable.
- **Cultural acceptance remains unclear.** Humanoids and other physical AI devices are experienced by humans as being more than just neutral machines. People naturally project attributes such as personality, emotion, and social behavior onto intelligent robots. Western cultural narratives frequently portray robots as being potentially threatening, whereas some Asian narratives, such as Japan's robotic cat Doraemon, frame them as companions or helpers. Physical AI developers and users still need a better understanding of how to create conditions that enable robots and humans to work harmoniously across different roles and settings.

Humanoids and other physical AI devices are experienced by humans as being more than just neutral machines.

The sweet spot for physical AI is in structured environments with a degree of variability

The emerging sweet spot for physical AI is structured work with enough variability to benefit from intelligence, but not so much complexity that deployment becomes impractical. Applications that meet these criteria are near-term priorities for development (see Figure 13).

Each application for physical AI reflects a different combination of autonomy and dexterity. For example, robotaxis can navigate highly dynamic environments but have limited manipulation abilities, while surgical robots can perform highly precise manipulation tasks but operate with tightly constrained autonomy.

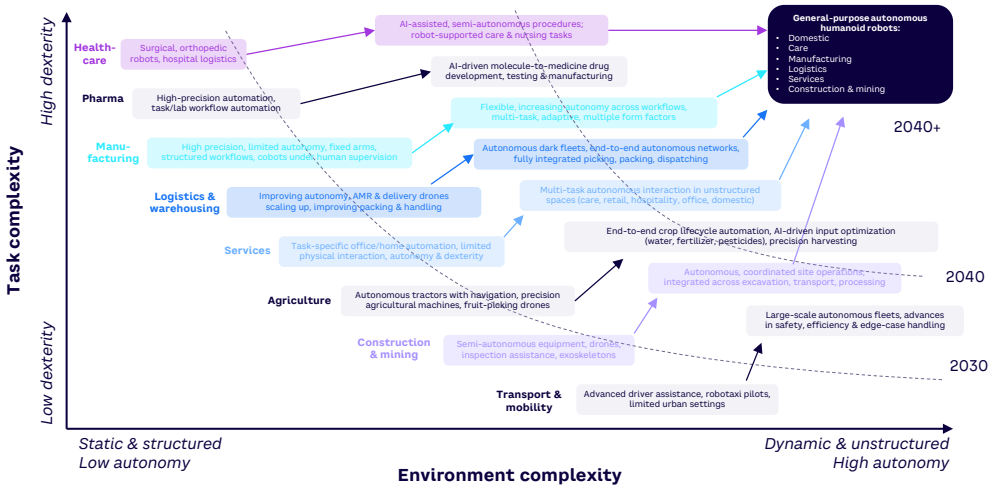
"Factories and logistics are much more likely than homes in the near term, as these environments are more structured. Objects are more known. The task flow is more predictable."

Prof. Roland Siegwart, ETH Zurich

Structured sectors will move toward greater autonomy & dexterity

- **Healthcare intelligent robotics has strong potential but remains high risk.** Robotics is already widely used in surgical applications. Physical AI also has potential in hospital logistics and, over the longer term, as an assistant in care and nursing tasks. However, autonomy will likely remain tightly controlled because of the level of exposure to vulnerable humans. French company Wandercraft, for example, shows how specialized physical AI can create value today. Its self-balancing exoskeleton Atalante X supports gait rehabilitation in clinical settings, illustrating how embodied AI succeeds first in constrained, high-supervision environments.

Figure 13. Developing sector applications will reflect converging advances in dexterity and autonomy

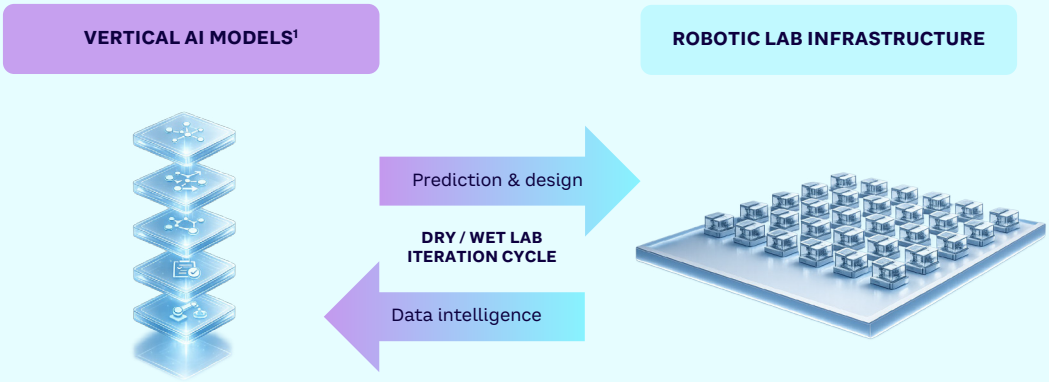


Source: Arthur D. Little; Prof. Roland Siegwart, Autonomous Systems Lab, ETH Zurich/EPFL

Case study: XtalPi & the autonomous pharmaceutical lab

XtalPi is building an AI-driven "autonomous lab" platform that combines robotics, AI models, quantum physics, and cloud computing to accelerate pharmaceutical R&D (see Figure A). Its approach integrates AI "dry labs" with robotic "wet labs" in a closed-loop design-make-test-analyze workflow, enabling faster and more automated drug discovery. The company orchestrates hundreds of robotic workstations for synthesis, crystallization, dispensing, testing, and sample handling, operating continuously with minimal supervision. XtalPi claims this infrastructure can improve experimental efficiency by 5x–8x while significantly increasing data quality and scale. The platform is designed to shift pharmaceutical R&D from labor-intensive experimentation toward AI- and automation-intensive discovery workflows.

Figure A. Example autonomous lab platform for pharma R&D



Note: (1) Structural prediction model, activity prediction model, molecule generation + screening model, pipeline application model, viability + physicochemical properties model, solid-state model, antibody model, automation model
Source: Arthur D. Little, XtalPi

- **Pharmaceutical applications in drug discovery could be transformative.** Applications for physical AI in pharmaceuticals focus on accelerating discovery and development by using robotics to close the loop between AI-driven experimentation and physical execution. The long-term vision is a fully autonomous "self-driving" drug discovery lab (see case study). Significant barriers remain, however, including biological complexity and variability, data quality and fragmentation, and regulatory constraints.
- **Manufacturing applications will move from precision automation toward adaptive multi-task workflows.** Typical manufacturing applications initially include high-precision, lower-autonomy roles and cobots supporting multi-task workflows. Over the next five to 15 years, this will likely extend toward greater dexterity, adaptability, and multi-task capability across multiple form factors, including humanoids. Traditional production line automation is already well served by conventional robotics, meaning future value creation is more likely to come from flexible and adaptive workflows.
- **Logistics and warehousing developments will initially focus on optimization rather than replacement.** Near-term developments include AMRs deployed at scale with improved dexterity and autonomy, eventually extending to autonomous "dark" fleets capable of fulfilling end-to-end orders.

Applications in less structured sectors will require further advances in autonomy & robustness

- **Service and support applications have large potential, but domestic robots remain distant.** Physical AI could be especially suitable for use cases in settings with many people, repeatable tasks, and onsite maintenance staff. This includes hospitals, hotels, airports, shops, laundromats, museums, and care facilities, with use cases such as delivering objects, guiding, and cleaning. The dream of domestic home robots remains distant, however, for three main reasons:
  1. **Homes combine open-ended layouts** with large numbers of messy and unstructured objects.
  2. **Significant safety risks** must be resolved before robots can move freely around people in homes.
  3. **Costs would need to fall sufficiently** for mass-market consumer adoption.

Capable domestic robots are unlikely to become widely available for at least another 15–20 years.

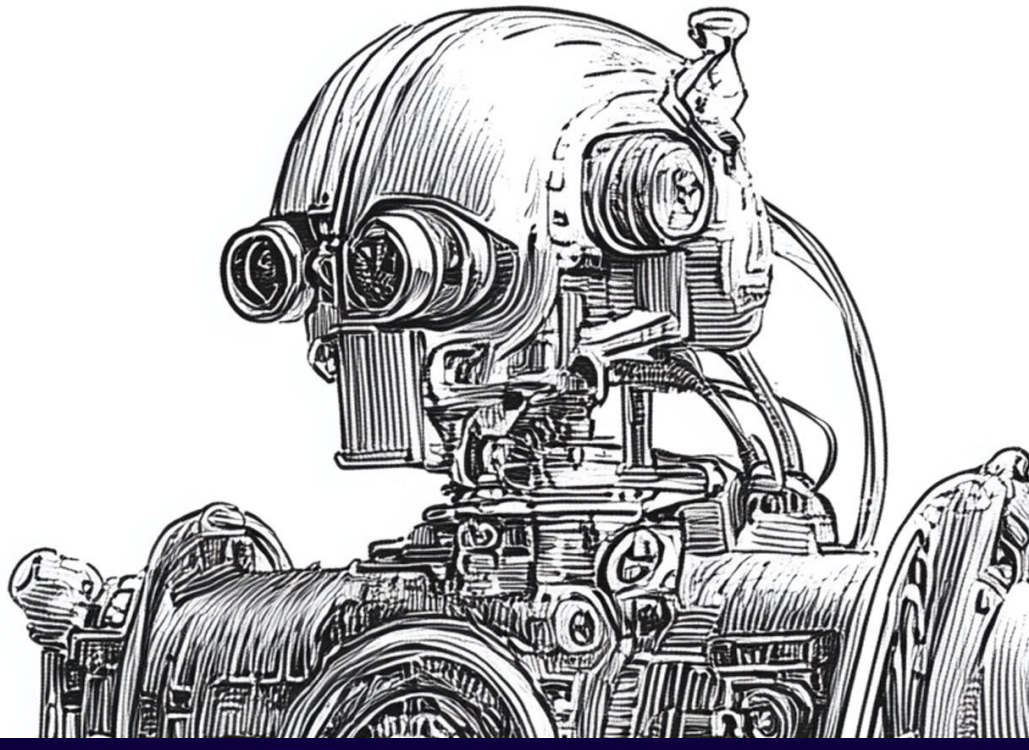
Experts believe that capable domestic robots are unlikely to become widely available for at least another 15–20 years:

- **Agriculture will likely move toward greater end-to-end automation and precision.** In agriculture, AVs already operate with limited human intervention as tractors, sprayers, and harvesters, while AI supports precision crop management. Over the next five to 15 years, this will likely extend toward end-to-end crop lifecycle automation, more precise weeding and harvesting, and greater livestock automation.
- **Construction and mining could benefit significantly if reliability and economics improve.** Robotics is already deployed selectively for heavy or dangerous tasks. Over time, this could extend to more integrated autonomous operations that reduce labor intensity. However, the variability of these environments and the economics of the industries mean that large-scale deployment will depend on cost-effectiveness and proven reliability.
- **Transport and mobility systems will likely see large-scale autonomous rollout.** In transport, the next five to 15 years should bring further progress in safety and efficiency with broader rollout of large-scale autonomous mobility systems.

Humanoid form factors could still dominate in the long term

In the short to medium term, wheel-based and arm-only form factors will likely scale first in industrial applications because they are cheaper, more energy efficient, and easier to tailor to specific tasks. Bipedal form factors should become widespread first in entertainment and interaction-oriented applications. However, they could eventually become valuable for tasks requiring maneuverability, stair navigation, and the ability to cope with work targets with diverse degrees of freedom. If humanoid cost, power efficiency, autonomy, dexterity, and application-specific software improve sufficiently, humanoid robots could still become scalable over the long term and ultimately dominate parts of the market.

CHAPTER 3



CRITICAL UNCERTAINTIES & FUTURE SCENARIOS

3 CRITICAL UNCERTAINTIES & FUTURE SCENARIOS

The future trajectory of physical AI is subject to many uncertainties. The most critical uncertainties are the speed at which breakthroughs will be made in autonomy and dexterity and the pace and breadth of deployment scaling. Together, these determine whether the future evolves toward one of four plausible scenarios: Autonomous Abundance, Specialized Intelligence, Industrial Workhorses, or merely Incremental Automation.

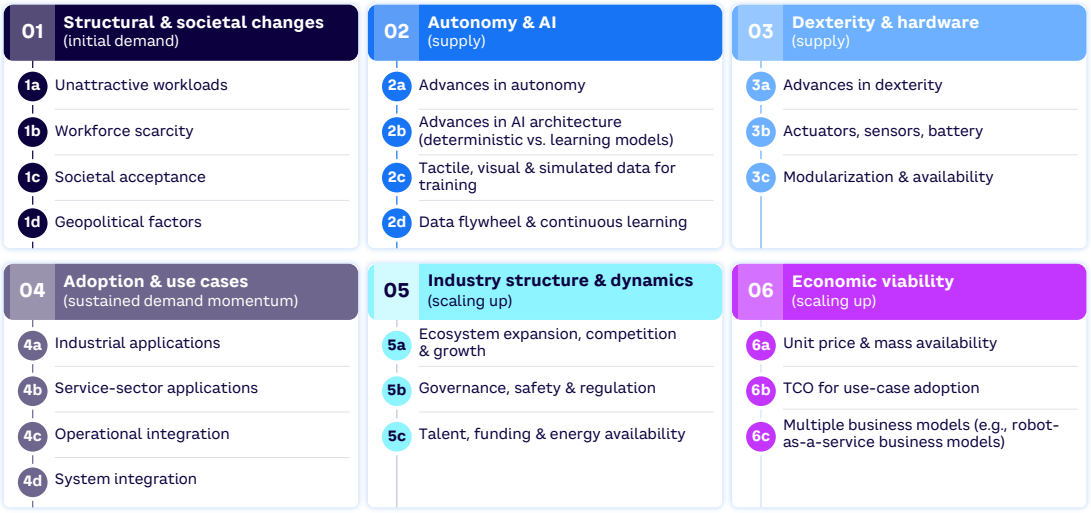
SIX SETS OF FACTORS SHAPE THE TRAJECTORY

While physical AI will likely be deployed at scale across industry and society over time, the trajectory remains uncertain. Understanding these uncertainties is useful for strategic planning. Six sets of shaping factors will influence future development: structural and societal changes, autonomy and AI, dexterity and hardware, adoption and use cases, industry structure and dynamics, and economic viability (see Figure 14).

TWO CRITICAL UNCERTAINTIES DOMINATE: TECHNOLOGICAL BREAKTHROUGHS & DEPLOYMENT SCALING

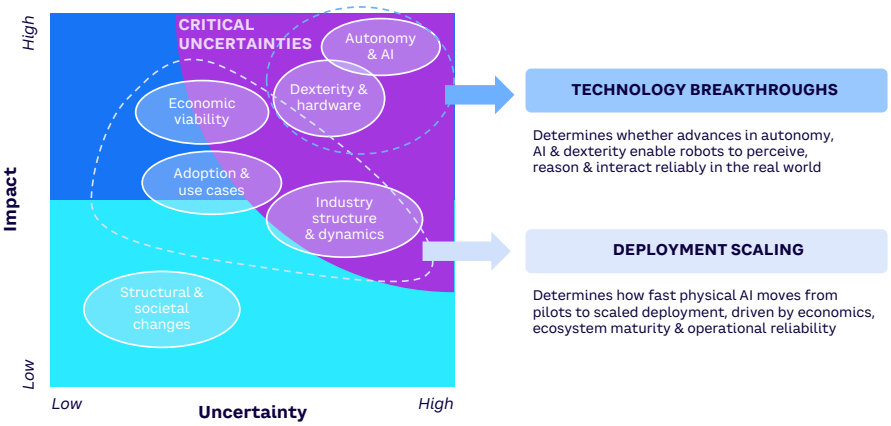
Each set of factors can be assessed according to its impact on physical AI development (low to high) and the degree of uncertainty associated with it (low to high). This produces the ranking shown in Figure 15. Pulling these shaping factors together, the two most critical uncertainties that emerge are the pace of technological progress and the speed of deployment scaling.

Figure 14. Six sets of factors will shape future development



Source: Arthur D. Little

Figure 15. Shaping factors ranking and critical uncertainties



Source: Arthur D. Little

Autonomy/AI & dexterity/hardware are the most consequential uncertainties

As shown in Figure 15, the shaping factors with the highest impact and uncertainty are autonomy and AI and dexterity and hardware:

- **Advances in autonomy and AI are fundamental to physical AI's trajectory.** This includes advances in AI architectures, tactile data availability, world models, and decision-making systems, as well as the evolution from deterministic control systems toward learning-based models. It also includes the development of data flywheels and continuous learning loops. There is still a high degree of uncertainty around how rapidly these technical challenges can be overcome.
- **Advances in dexterity and hardware are equally critical.** Progress in actuators, sensors, batteries, materials, and robot design defines the physical limits of what robots can do. Dexterity, particularly manipulation and tactile interaction, remains one of the most important bottlenecks preventing broader deployment beyond structured environments.

Societal acceptance remains important but is likely to adapt gradually.

Economic viability, industry structure/dynamics & adoption/use cases have the next-highest impact & uncertainty

These three sets of factors largely determine the achievable pace of scale-up:

- **Economic viability will determine whether adoption moves beyond pilots.** Factors such as unit cost, TCO, and business model alignment (e.g., robot as a service) will determine whether physical AI solutions can scale commercially.
- **Industry structure and dynamics will need to develop further.** Competition, partnerships, and the emergence of integrators will determine how quickly solutions can be developed, deployed, and maintained at scale.
- **Adoption and use cases must become sufficiently broad.** The breadth and depth of applications across industrial and service sectors will shape how quickly demand develops. Integration into real workflows and operating environments is essential for moving from demonstrations and pilots to scaled deployment and sustaining adoption momentum.

Structural & societal changes, while important, have relatively lower impact & uncertainty

Factors such as workforce scarcity and job attractiveness are already well understood. Although geopolitics will influence competitive positioning and regional leadership, it is less likely than technological progress or deployment economics to fundamentally change the overall pace of physical AI development. Societal acceptance remains important but is likely to adapt gradually as physical AI demonstrates safety, reliability, and economic value.

CRITICAL UNCERTAINTIES LEAD TO FOUR PLAUSIBLE SCENARIOS, EACH WITH STRATEGIC IMPLICATIONS

Based on these two dimensions, four plausible scenarios emerge for the future trajectory of physical AI (see Figure 16).

Scenario 1: Incremental Automation (low technology breakthroughs, slow scaling)

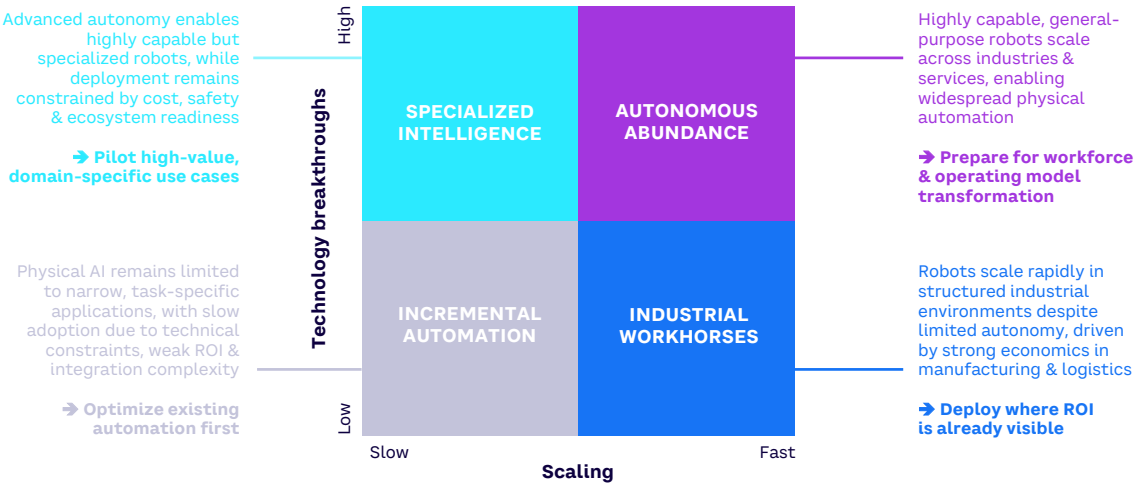
Incremental improvement; current expectations prove overly optimistic. In this scenario, both technological progress and scaling remain limited. Despite incremental improvements, fundamental bottlenecks in autonomy, dexterity, and learning persist. Deployment remains concentrated in well-defined tasks and highly structured environments, with robots acting primarily as assistive tools rather than independent workers. Human oversight remains essential, and consumer applications remain largely out of reach.

This scenario reflects a world in which current expectations about physical AI prove overly optimistic.

Strategic implications:

- Prioritize optimization of existing automation assets.
- Avoid overcommitting to immature platforms.
- Target narrow, high-ROI use cases where economics are already proven (e.g., warehouse AMRs, cobots, quality control).
- Maintain workforce-centric operating models.

Figure 16. Four future scenarios for physical AI



Source: Arthur D. Little

The future trajectory of physical AI is subject to many uncertainties.

Scenario 2: Industrial Workhorses (low technology breakthroughs, fast scaling)

Pragmatic scaling; value is created through targeted industrial applications. In this scenario, deployment scales rapidly, but technological breakthroughs remain limited. Robots do not achieve general-purpose capability but become reliable and cost-effective enough to scale across structured industrial environments. Adoption accelerates in manufacturing, logistics, and warehousing, where tasks are repetitive but variable enough to benefit from flexible automation. Systems remain specialized, with limited dexterity but high reliability within defined workflows. Strong ROI, improving commercial models, and ecosystem development support widespread deployment. Hybrid operating models remain common, with humans handling exceptions and oversight.

This scenario reflects a pragmatic evolution of physical AI and will likely be the most commercially viable path over the next decade.

Strategic implications:

- Aggressively automate structured environments with visible ROI.
- Redesign operations around robotic economics, focusing on throughput, utilization, and error reduction.
- Invest heavily in orchestration and systems integration (e.g., example workflows, fleet management).
- Design greenfield factories and warehouses around robotic operation from the outset.

Scenario 3: Specialized Intelligence (high technology breakthroughs, slow scaling)

Breakthrough capabilities emerge in specific domains, but adoption remains limited. In this scenario, significant technological breakthroughs occur, but scaling remains limited. Robots achieve high levels of autonomy or dexterity in selected domains, but high costs, safety concerns, and integration challenges prevent broad deployment.

Physical AI systems excel in areas such as mobility, advanced manufacturing, defense, and specialized industrial services but remain concentrated in high-value applications and controlled environments. Deployment is characterized by advanced pilots and niche applications rather than broad adoption. Industries such as transport, defense, or high-end manufacturing benefit from these capabilities, but widespread workforce transformation does not occur.

Strategic implications:

- Pilot high-value, domain-specific applications where labor is scarce or risks are high.
- Treat physical AI as strategic capability infrastructure that can create competitive advantage.
- Build proprietary physical AI ecosystems.
- Expect slower adoption rates and selective deployment.

Scenario 4: Autonomous Abundance (high technology breakthroughs, fast scaling)

Breakthrough capabilities combine with rapid scaling to enable widespread, general-purpose deployment. In this scenario, advances in autonomy and dexterity are matched by favorable economics, mature ecosystems, and rapid deployment. General-purpose robots become viable across a wide range of industries, including services and potentially consumer applications.

High levels of autonomy reduce the need for human intervention, and physical AI becomes integrated into daily life and work. Productivity gains are substantial, and labor markets undergo significant transformation. Many tasks traditionally performed by humans become automated, while new roles emerge around supervision, coordination, and system design.

This scenario also introduces new challenges, including workforce displacement, regulatory complexity, and broader societal implications.

A diversity of robot embodiments will persist rather than a single dominant humanoid platform.

Strategic implications:

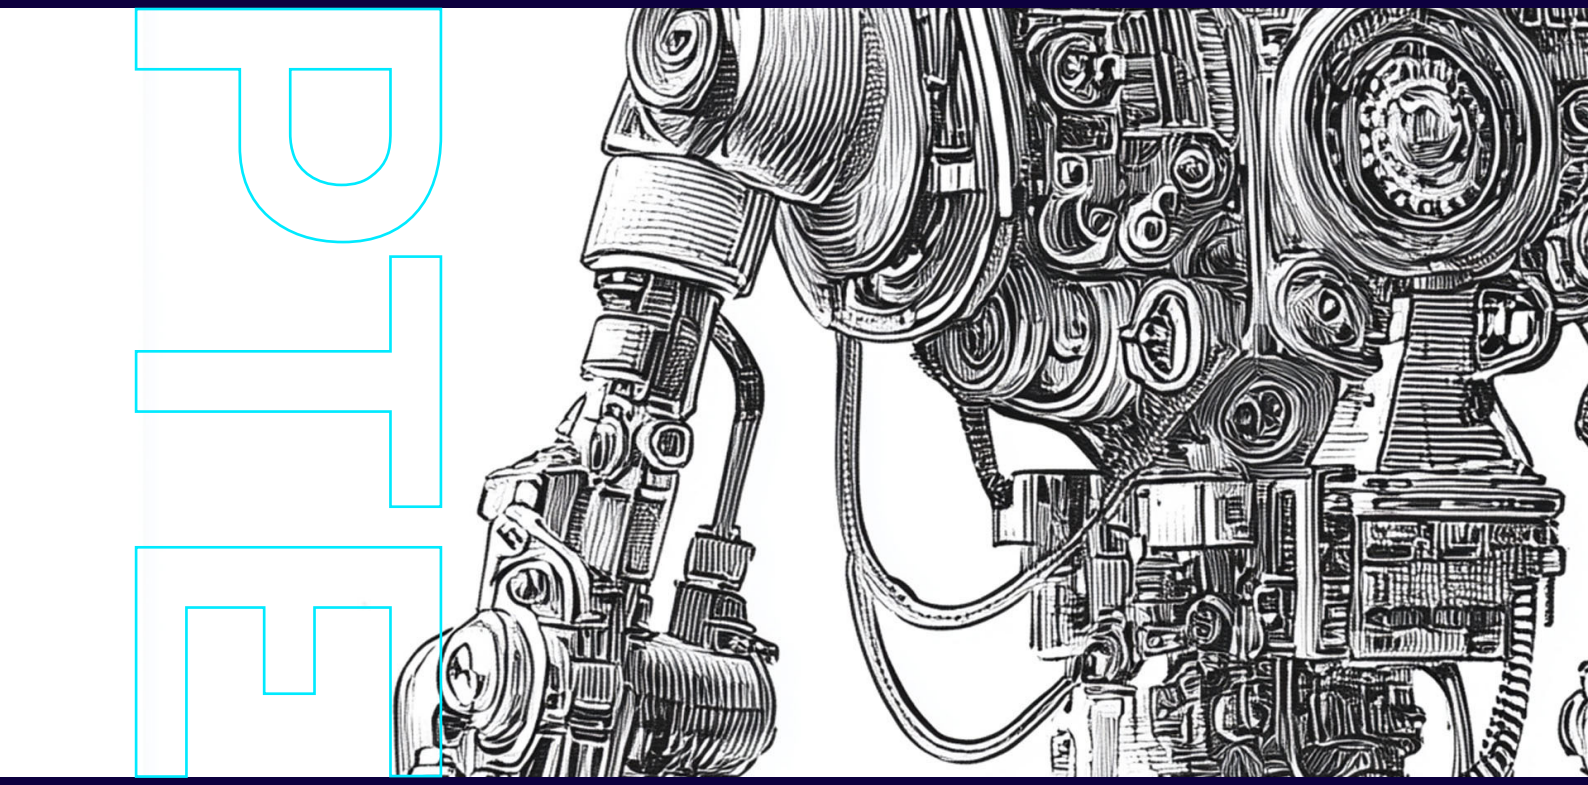
- Prepare for large-scale workforce and operating model disruption.
- Move quickly to avoid structural cost disadvantages.
- Increase control over strategic AI and robotics infrastructure.
- Redesign workforce strategies around human-machine collaboration.
- Prepare for ecosystem and regulatory disruption related to employment, safety, liability, sovereignty, and governance.

In practice, the future trajectory of physical AI could embody elements of several or all of these scenarios at different stages of development. A key takeaway is that Autonomous Abundance should not be treated as the default outcome.

As discussed in Chapter 2, regardless of which scenario ultimately prevails, over the next five to 10 years, we can expect that:

- Structured industrial environments will dominate.
- Hybrid human-robot operating models will prevail initially.
- A diversity of robot embodiments will persist rather than a single dominant humanoid platform.

CHAPTER 4



WHAT EXECUTIVES SHOULD DO TODAY

4 WHAT EXECUTIVES SHOULD DO TODAY

As physical AI development accelerates, executives need to act now to secure competitive advantage. The initial focus should be on tightly scoped pilots for scalable use cases where physical AI autonomy and dexterity are already "good enough," while also considering new workflows and ways of working. Companies should make no-regret moves today to strengthen operational data and digital infrastructure and engage with ecosystem players. Deployment will require robust change management capabilities.

"Every time I visit China, I'm struck by the speed and scale at which robotics and physical AI are being deployed. This is something global business leaders need to take very seriously."

Marc Fleischmann, Chief Software & AI Officer, KUKA Robotics

EXECUTIVES SHOULD ACT NOW

- **Development progress is accelerating.** While physical AI is not yet mature enough for universal deployment, it is already strategically significant. There are strong indications that the industry is approaching an inflection point where progress in physical AI — not just humanoid, but also other form factors — is accelerating rapidly.
- **Early movers may enjoy significant competitive advantage.** The cost of inaction rises as capabilities improve and early adopters build operational and data advantages. The challenge for industry leaders is not whether to engage, but how to do so pragmatically, anchored in real operational value rather than technology narratives. Decisions made over the next few years will determine whether organizations build a learning advantage or remain dependent on external providers once the technology matures.

"Don't be late to the game.... Start to build up in-house knowledge; how these technologies can be used and shall be used."

Pascal Gohl, Head of Robotics, Hexagon Robotics

Companies should begin by identifying a portfolio of potential use cases.

START WITH IDENTIFICATION & EVALUATION OF USE CASES

- **Build a portfolio, starting with broadly scalable use cases.** Companies should begin by identifying a portfolio of potential use cases. Initially, value is most likely to be created through broadly scalable, repeatable use cases where physical AI already provides sufficient levels of autonomy and dexterity. For a typical manufacturing company, this points toward labor-intensive, nonspecialized operations, including goods receiving, warehousing, logistics, assembly, testing, packing, and deliveries. Existing pain points, such as labor shortages, safety risks, or operational bottlenecks, are often the best candidates. Additional upside may emerge from targeted applications enabled by advances in autonomy or dexterity, such as inspection, maintenance, and autonomous mobility.
- **Target niche, high-value applications as well.** Consider applications where advances in autonomy or dexterity unlock specific high-value tasks. These may not support broad deployment but can deliver strong ROI within clearly defined boundaries. Examples include inspection and maintenance of distributed assets in remote or hazardous environments.

These use cases do not require general-purpose intelligence, but they do require sufficient reliability to operate under real-world conditions. Scaling in these sectors is often more selective due to integration complexity and regulatory constraints, but the economic value per use case can still be significant.

- **Reimagine workflows and ways of working.** Avoid focusing solely on direct substitution of human labor. Instead, consider how physical AI could enable redesigned workflows and new operating models that combine human and machine capabilities. These changes may unlock value beyond traditional efficiency and productivity improvements.

"It is very useful to understand the problems you actually want to solve.... Where in my production do I use manual labor today? Where do I have my largest costs or work-environment problems?"

Klas Kronander, Cofounder, Enodo Robotics

Organizations should actively engage with emerging ecosystem players.

INITIATE NO–REGRET MOVES & BUILD FOUNDATIONAL CAPABILITIES

- **Strengthen operational data and digital infrastructure.** Once promising use cases have been identified, organizations should begin building the foundations required for future deployment. These are no-regret because they are necessary in virtually all physical scenarios. The most important include strengthening operational data quality and digital connectivity, including data capture, data cleanup, connectivity, and edge computing. Often, workflows are poorly understood at a detailed level, and actual operations frequently differ from documented procedures. As a result, operational data is often fragmented and inconsistent, creating major obstacles to deployment and making it difficult to assess costs and benefits. Early action is important because these improvements can take significant time.

"Think about and understand your processes; that is the important first step ... and get internet connection on the shop floor."

Dr. André Scholz, Head of Innovation Department Autonomous Factory, Siemens

- **Engage with the physical AI ecosystem to build capability.** Organizations should actively engage with emerging ecosystem players to build internal knowledge and position themselves advantageously. This includes participating in industry networks, understanding key player capabilities, and establishing partnerships and alliances for studies and pilots. Early pilots with multiple vendors help organizations evaluate potential solutions while building familiarity with the technology landscape. Engagement should extend beyond robotics vendors to include application developers, systems integrators, and other ecosystem participants.

In some cases, participation in academic studies may support first-mover advantage and improve internal knowledge and understanding. Investors should likewise recognize that opportunities extend beyond robot manufacturers to the broader ecosystem, including components, software, integration, maintenance, and other enabling services. Service provider roles (such as maintainers, insurers, and integrators) will likely expand.

- **Monitor and help shape regulation.** Governments play a critical role in enabling adoption, as regulatory and safety barriers remain significant. Industry players can help shape practical regulatory frameworks through active engagement with policymakers. Regulatory sandboxes can accelerate learning by enabling real-world experimentation while avoiding premature ineffective regulation. In particular, safety certification and governance frameworks will be critical to reducing deployment risk at scale.

Companies should focus on targeted pilots designed to test real-world performance.

TEST & PILOT IN STAGES

- **Pilot real-world use cases.** Companies should focus on targeted pilots designed to test real-world performance. Many systems that perform well in controlled demonstrations struggle when exposed to operational variability, edge cases, and reliability challenges. Even seemingly simple environments such as logistics operations can present significant complexity.
- **Start with lower-risk/higher-return use cases.** Initial pilots should focus on applications that have broad scalability potential and relatively low technological complexity. These are more likely to generate attractive returns. Companies should consider not only humanoids, but also simpler task-optimized solutions such as robotic arms, wheeled robots, and mobile manipulators. Back-office or industrial applications may present lower risk than customer-facing deployments.
- **Sequence deployments with tight scoping and clear objectives.** Physical AI initiatives should be sequenced, with clear decision points before wider deployment to avoid high costs and failures. One of the biggest risks is overextension. Pilots should therefore be tightly scoped, with clearly defined success metrics linked to cost, productivity, and reliability. Broad experimentation without clear objectives often leads to inconclusive results.
- **Evaluate economics pragmatically.** Consider the full end-to-end economics of deployment at scale, which may differ substantially from pilot economics. Performance should be validated under real operating conditions. Evaluation should include technical reliability, workflow integration, maintenance requirements, service support, and organizational readiness. This is particularly important in distributed industries such as telecommunications and energy, where integration complexity is high and failures can be costly. Maintaining flexibility is critical. Organizations should also avoid premature vendor lock-in, retain access to operational data, and continue building internal capabilities alongside external partnerships.

- **Adapt strategy as the market evolves.** As discussed in Chapter 3, organizations should continuously monitor market developments and adjust their deployment strategies to suit. For example, an Incremental Automation scenario favors optimization of existing automation assets and selective deployment of high-ROI applications. An Industrial Workhorses scenario favors aggressive automation of structured environments and significant investment in systems integration and workflow redesign. A Specialized Intelligence scenario favors development of proprietary high-value, domain-specific capabilities. And an Autonomous Abundance scenario requires rapid operating model transformation and workforce redesign.

The potential of physical AI to transform industry, and ultimately society, is substantial.

PLAN & MANAGE DEPLOYMENT IMPACTS

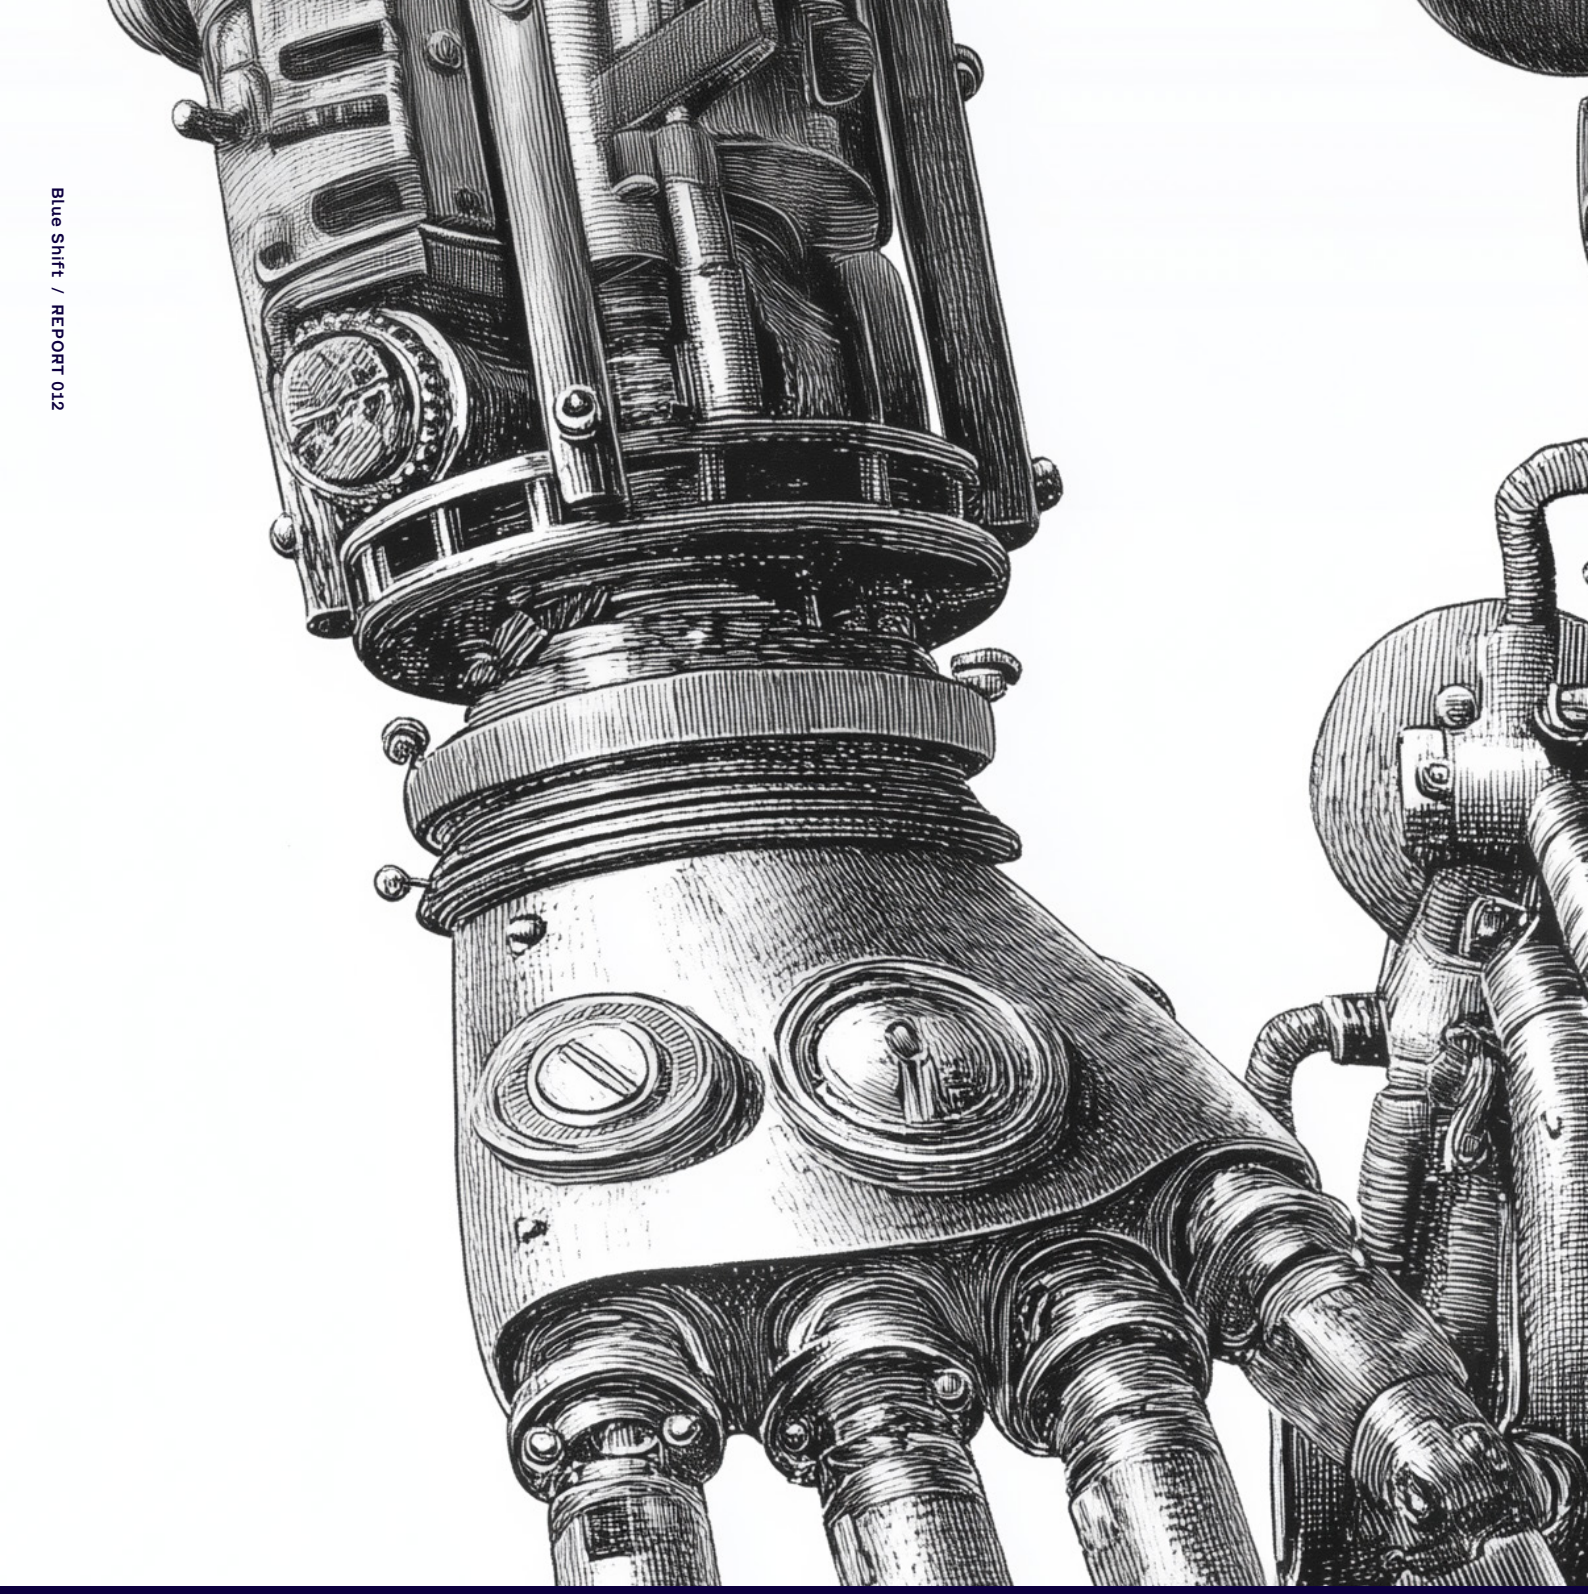
- **Define success, impact, and approach.** Organizations need clear definitions of what success looks like, what the impact will be on the organization, and how the transformation will be managed. This includes understanding who will be affected, how roles will change, and what implementation roadmap, risks, and mitigation measures are required. Transparent communication with employees is essential.
- **Manage and monitor change carefully.** Successful deployment depends as much on organizational adoption as on technical performance. Robust change management approaches should therefore be embedded from the outset. Priorities include building awareness, encouraging engagement, developing skills and knowledge, sustaining adoption over time, and monitoring and measuring progress against defined objectives.

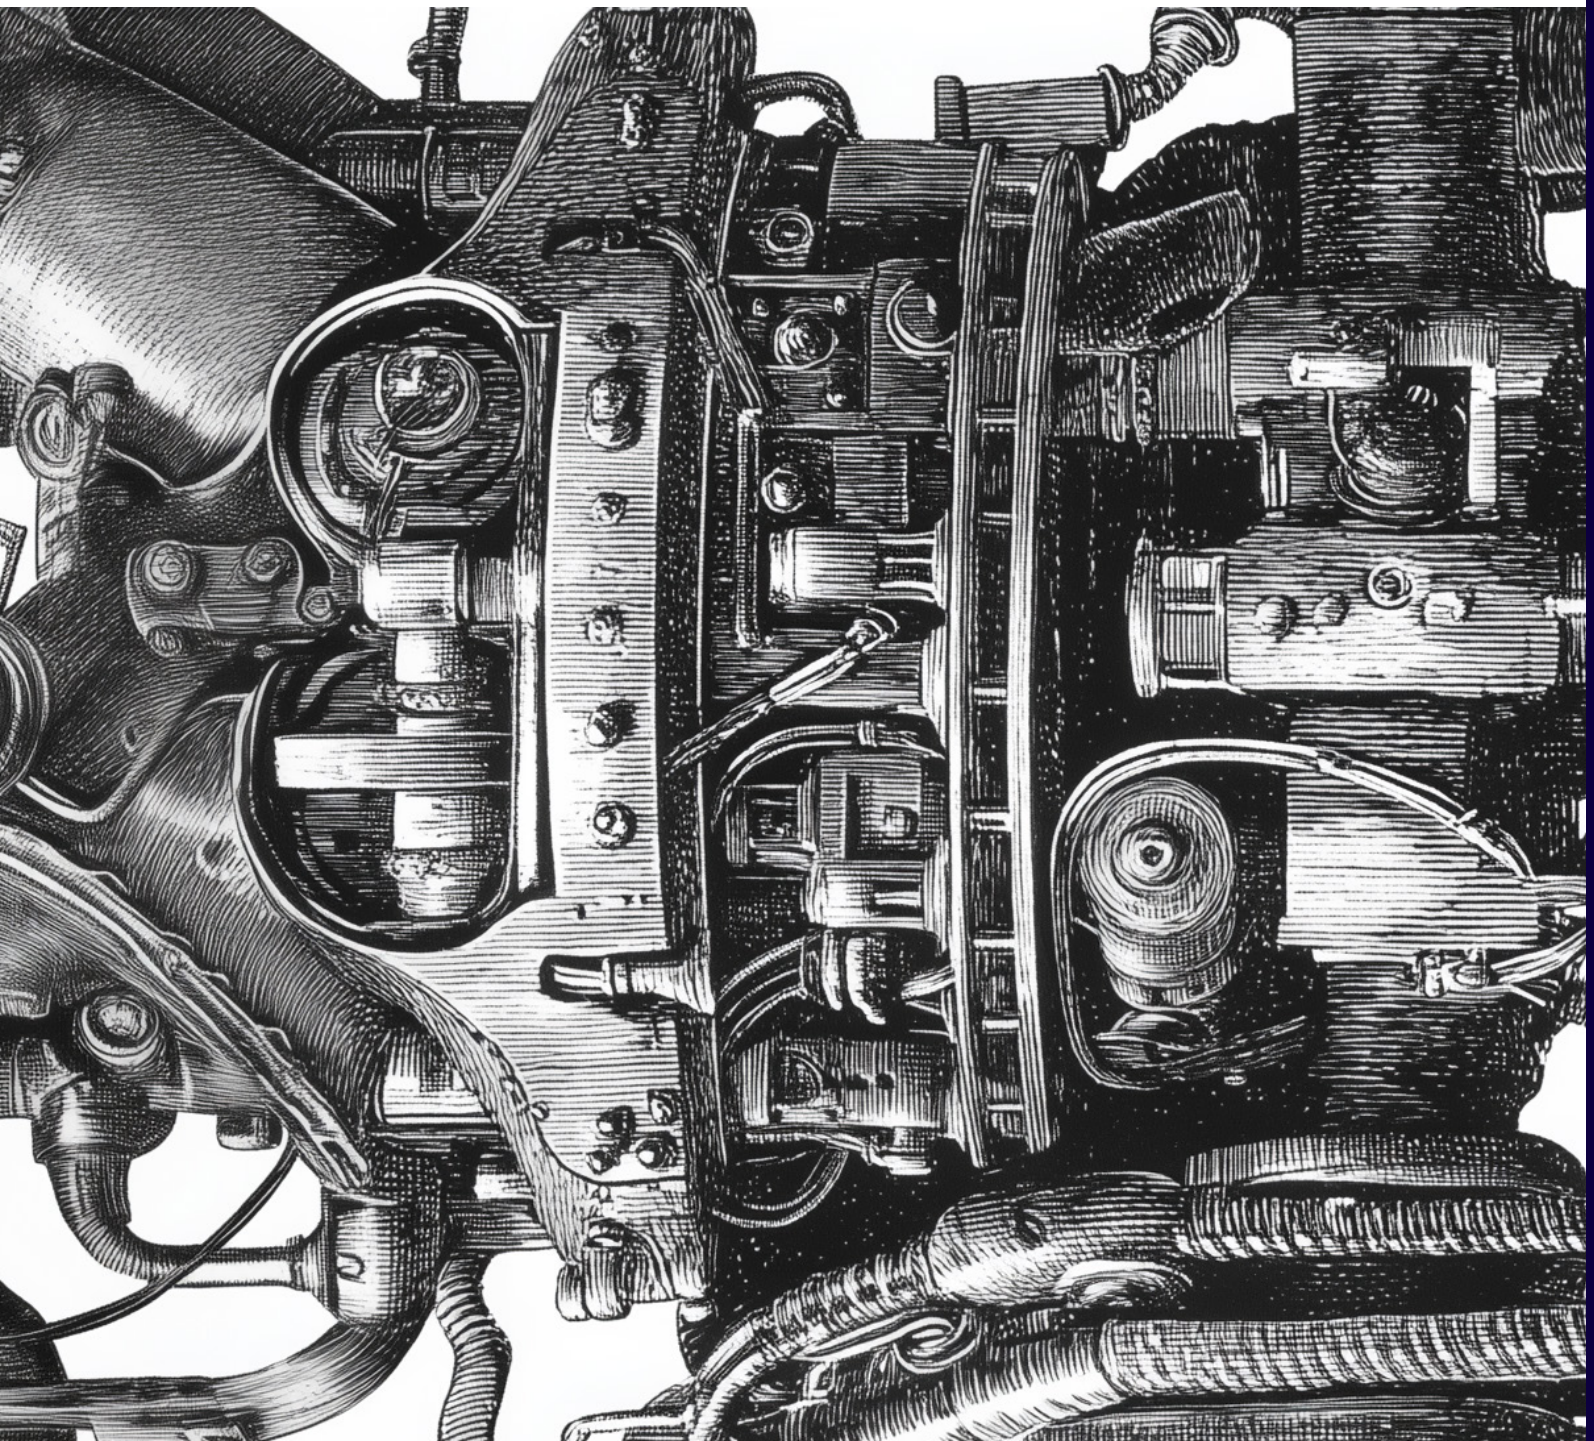
CONCLUSION

Physical AI will not transform society overnight. However, decisions made over the next few years about where to pilot, what to measure, which partners to engage with, and how to prepare your workforce will help determine which organizations lead and which fall behind. The winners will not necessarily be those chasing full autonomy or humanoid visions alone, but those that focus on solving real operational problems, deploying early, learning quickly, and building internal capabilities and data advantages. While the current market contains a significant degree of hype, the potential of physical AI to transform industry, and ultimately society, is substantial, and the pace of development is accelerating. Executives in industries with relevant use cases should engage now.

Endnotes

1 "Figure Exceeds $1B in Series C Funding at $39B Post-Money Valuation." Figure, 16 September 2025.
2 "State of Robotics 2026: A Sector on the Cusp of AGI-Driven Transformation." National Robotics Association, 10 February 2026.
3 "World Robotics Report 2025." International Federation of Robotics, 2025.
4 "Humanoids: A $5 Trillion Market." Morgan Stanley, 14 May 2025.
5 Crowe, Steve. "2026 State of the Robotics Industry Report." The Robot Report, 26 January 2026.
6 Liljas, Per Elinder. "Human Record Shattered by Bipedal Robot." SvD Naringsliv, 19 April 2026.
7 "Global Humanoid Robot Installations Reach 16,000 Units in 2025 as Mass Production Picks Up Pace." Counterpoint, 13 January 2026.
8 "OECD Employment Outlook." OECD, 9 July 2025.
9 "Business Insights and Impact on the UK Economy: 22 May 2025." UK Office for National Statistics, 23 October 2025.
10 "Sanctuary AI Demonstrates In-Hand Manipulation Capabilities for Improved General Purpose Robot Dexterity." Press release, Sanctuary AI, 12 December 2025.
11 European Union. "Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on Machinery and Repealing Directive 2006/42/EC of the European Parliament and of the Council and Council Directive 73/361/EEC (Text with EEA Relevance)." EUR-Lex, 2023.





Acknowledgments

We would like to express our deepest gratitude to the esteemed experts worldwide who generously shared their knowledge and insights during the interviews conducted for this Report. Their expertise and contributions have been invaluable in shaping the understanding and depth of our research. In particular, we would like to acknowledge the following experts:

Ali Alessandro Ayach, Partner, Healthcare & Life Sciences (HCLS), ADL
Toshihiro Ariki, Partner, Automotive & Manufactured Goods (AMG), ADL
Paul Batsii, CPO, Robotics AI, Siemens
Stefano Bensi, General Manager, SoftBank Robotics EMEA
Rajesh Duneja, Partner, Telecommunications, IT, Media & Electronics (TIME), ADL
Martijn Eikelenboom, Partner, Sustainability, ADL
Juan Farré, President & CEO, Danish Technological Institute (DTI)
Marc Fleischmann, Chief Software & AI Officer, KUKA Robotics
Michael Fischer, Senior Advisor, ADL
Denis Gardin, Director of Innovation & Future Technology, MBDA
Brian Gerkey, PhD, CTO, Intrinsic, and Board Chair, Open Robotics
Pascal Gohl, Head of Robotics, Hexagon Robotics
Dave Grant, CEO, PickNik Robotics
Claus Grubern, Partner, AMG, ADL
Neville Hogan, Professor, Mechanical Engineering and Brain & Cognitive Sciences, MIT
Peter Howard, President & CEO, Realtime Robotics
Hiroshi Ishiguro, Professor, Osaka University
Shirish Jain, Partner, Financial Services, ADL
Thibauld Jongen, PhD, Cofounder, Common Sense Robotics
Arnaud Jouron, Managing Partner, Performance, ADL
Josef Karbassi, EVP Business Development, Vacuum Automation, Piab Group
Petter Kilefors, Managing Partner, Growth, ADL
George Konidaris, Cofounder & Chief Scientific Advisor, Realtime Robotics;
Associate Professor of Computer Science, Brown University
Michael Kolk, Managing Partner, Global Practice Leader, Innovation, ADL
Klas Kronander, PhD, Cofounder, Enodo Robotics
Goetz Kuras, Partner, Financial Services, ADL
Emilio Lapiello, Partner, Innovation, ADL
Chui Lee, Managing Partner, Transformation, ADL
Katlyn Lewicke, Head of Global Automation Strategy & Intel, GXO Logistics
Lorenzo Masia, Professor, Intelligent Bio-Robotic Systems; Executive Director,
Munich Institute for Robotics & Machine Intelligence (MIRMI)
Shane McArdle, President & CEO, Falkor (formerly Kongsberg Digital)
Elena Messina, Principal, Prospicience LLC
Jérôme Monceaux, Founder & CEO, Enchanted Tools
Christian Oestergaard, Lead Visionary & Senior Group Strategist of IT-Production/IoT/AI, PostNord
Enguerran Ripert, Principal, Growth, ADL
André Scholz, PhD, Head of Innovation Department Autonomous Factory, Siemens
William Shi, CEO of Europe & Americas, AgiBot
Roland Siegwart, Professor, Mechanical & Process Engineering, ETH Zurich
Arjun Singh, Partner, Financial Services, ADL
Georg Stieler, Head of Robotics & Automation, Stieler Technology
Christian Tauber, Cofounder & CEO, NEOALP
Janahan Tharmaratnam, Partner, HCLS, ADL
Jocelyne Troccaz, Emeritus Research Director, CNRS
Garrett Whitmore, Partner, Innovation, ADL
Kai-Oliver Zander, Partner, Performance, ADL
Marina Zhang, Associate Professor, University of Technology Sydney (UTS)

Dr. Albert Meige, Director of Blue Shift, Arthur D. Little
meige.albert@adlittle.com

Disclaimer

The AI field is still evolving rapidly. This report reflects the best of our knowledge as of June 2026.

Blue Shift, by Arthur D. Little, explores the impact of technologies on business, society, and humans. The Blue Shift Report covers these topics in depth, inviting guest authors, academics, and artists to contribute to the conversation.

Want to join the discussion? Please address your comments to our editorial office at Arthur D. Little, New Fetter Place West, 2nd floor, 55 Fetter Lane, London EC4A 1AA, United Kingdom — Telephone: +44 7710 536 471.

For further information, visit: www.adlittle.com.

Copyright 2026, Arthur D. Little. All rights reserved.

BLUE SHIFT

BY ARTHUR D. LITTLE

ARTHUR D LITTLE