Thematic Investing

Physical AI: Robotics Summit Highlights Humanoid Scaling Hurdles

While humanoid robots dominated the Boston Robotics Summit, expectations for scaling fully autonomous, general-purpose humanoids were measured amid several gating factors. We also compiled a humanoid developer landscape.

Thematic Investing

William Thompson
+1 212 526 8641
william.s.thompson@barclays.com
BCI, US

Hannah Greenberg
+1 212 526 2457
hannah.greenberg@barclays.com
BCI, US

Hiral Patel
+44 (0)20 3134 1618
hiral.patel@barclays.com
Barclays, UK

Laia Marin i Sola, CFA
+ 44 (0)20 3134 2531
laia.marin@barclays.com
Barclays, UK

Katherine Ogundiya
+44 (0)20 3134 1391
katherine.a.ogundiya@barclays.com
Barclays, UK

Amy Lian
+44 (0)20 3134 0179
amy.lian@barclays.com
Barclays, UK

Siddharth Gupta
+91 22 6175 2540
siddharth.gupta1@barclays.com
Barclays, UK

Humanoids are coming—but when and at what scale? The recent Boston Robotics Summit focused on the embodiment of AI in humanoid robots. Despite optimism on single/multi-task robots, the tone was more tempered on the near-term scaling of fully autonomous, general-purpose humanoids given: (i) safety, reliability, security & social acceptance; (ii) hardware bottlenecks; (iii) absence of a Physical AI 'GPT moment' for perception; (iv) a training data gap; and (v) competition for compute. While China leads on volume, the US may be better positioned on capability. Our humanoid developer landscape illustrates that many remain in pilot stages, are vertically integrating manufacturing, and plan to leverage Nvidia edge compute. **Barclays is hosting a Physical AI & Robotics Forum in NYC on June 9th. Registration and more details can be found here.**

Completed: 07-Jun-26, 22:46 GMT Released: 08-Jun-26, 08:00 GMT Restricted - External

CONTENTS

Executive Summary

We attended the recent two-day **Robotics Summit & Expo in Boston**, where discussions centered on **Physical AI**, particularly its application in autonomous mobile robots (AMRs) and **humanoid robots**. We participated in various keynote and panel sessions and met with dozens of companies across the robotics ecosystem, including robot developers, suppliers of physical hardware (e.g., actuators, motors), compute hardware (semiconductors, memory, sensors), and providers of training data and simulation platforms. While our colleagues have addressed the broader Physical AI opportunity and macro implications in the recent **Barclays Equity Gilt Study**, our focus was on the **timing and feasibility of scaling Physical AI**, especially general-purpose humanoids.

While there was general optimism around the **near-term deployment of narrow or task-specific robots** in controlled environments, such as welding or logistics robots, the tone was a bit more cautious on the timing of a meaningful inflection toward **fully autonomous, general-purpose humanoids** being deployed in dynamic, human-centric environments, given several challenges:

- **Safety, reliability, security, and social acceptance** were consistently described as non-negotiables. Many speakers drew parallels to the autonomous vehicle industry, which has taken nearly a decade to move from early hype toward broader deployment, largely due to safety concerns and regulatory friction.

- **Hardware remains a bottleneck and presents a chicken-and-egg problem.** There is a lack of off-the-shelf humanoid components, alongside ongoing challenges in locomotion, manipulation, and dexterity. While software development can iterate rapidly, **hardware progress is constrained by supplier lead times, manufacturing complexity, and long feedback cycles**. Until acceptable levels of safety, reliability, security, and public trust are established, scaling manufacturing capacity remains difficult. Interestingly, many humanoid startups are **vertically integrating the design and manufacturing of key components,** including motors and actuators, to ensure cost-effectiveness and supply resiliency with some leveraging auto partnerships (Boston Dynamics/Hyundai, Optimus/Tesla, Mentee Robotics/Mobileye).

- **There is no clear "GPT moment" for human-like perception.** Physical AI still lacks a breakthrough that enables reliable, generalized perception across diverse environments and a long tail of edge cases, limiting progress toward true general-purpose autonomy.

- **A persistent training data gap remains.** While large language and vision models benefit from vast internet-scale datasets, there is no comparable repository of rich, robot-centric data needed to train machines to interact reliably with the physical world. Although the industry is increasingly turning to synthetic data and advanced simulation, the **sim-to-real gap** remains significant. In the near term, progress is expected to come from **narrow deployments with human oversight**, which can accelerate real-world learning loops.

- **Competition for compute is intensifying.** Humanoid development requires compute across simulation, training, and deployment, increasing demand for both **AI data centers and specialized, low-power edge inference hardware** embedded in each robot. Advances in model efficiency, compression, and end-to-end systems integration will be critical to sustaining progress. Interestingly, many humanoid startups are leveraging Nvidia for digital-twin simulation, Physical AI world models and edge compute (Jetson), as well as Google DeepMind for perception & decision-making algorithms.

Other takeaways:

- **While China leads on deployment volume, the U.S. could lead on capability.** China is widely viewed as a high-volume testing ground with greater risk tolerance, but many deployed systems remain scripted, teleoperated, or limited in robustness. It was suggested that U.S. could be better positioned to lead in developing humanoids capable of performing economically useful, industrial-grade work.

Safety, Reliability, Security & Social Acceptance

Safety, reliability, security, and social acceptance were widely cited as non-negotiables. Historically, automated robots have been confined to scripted tasks and controlled environments separated from humans, often by cages. Humanoids, however, are designed to operate in dynamic, human-centric environments. This shift introduces legal, ethical, and societal challenges that closely mirror the experience of autonomous vehicles. Despite early optimism, self-driving technology required more than a decade to move toward broader deployment, largely due to safety concerns and regulatory resistance. Humanoids are likely to follow a similar path, with an interim "human-in-the-loop" phase in which operators supervise and intervene as needed, helping to build trust while systems mature.

That said, failures in most humanoid use cases (e.g., warehouses or factories) generally carry lower consequences than failures on public roads. This lowers the bar for early deployment and increases willingness to roll out imperfect systems to initiate a real-world learning flywheel. Early deployments, even if limited, can generate critical data to improve performance and accelerate iteration. Ongoing advances in AI, including large language and vision-language models, should further compress timelines and support progress toward autonomy.

Even once humanoids demonstrate utility, safety and reliability will remain central to adoption. Real-world environments present a long tail of edge cases, from uneven floors to cluttered layouts, that are difficult to fully simulate. Reliability is also essential to the economic case: failures can cause costly downtime and operational disruption, while undermining worker trust. As a result, many industries will demand high levels of uptime and predictable performance before scaling deployments. As one panelist quipped, even Roomba robots haven't eliminated human vacuuming—they're only marginally effective at cleaning. It was often cited that AI will play a key role in improving reliability from roughly 85% to 95%+.

Security represents an additional, and often underappreciated, hurdle. Humanoids are fundamentally networked, software-defined systems that integrate sensors, actuators, AI models, and continuous connectivity. This exposes them to cyber risks ranging from unauthorized access and model tampering to data corruption and operational disruption. A compromised robot is not just an IT issue but a potential physical and operational liability. Robust security architectures, disciplined update regimes, and built-in fail-safes are therefore prerequisites for both enterprise adoption and public acceptance. Ultimately, sustained confidence in humanoids will depend as much on cybersecurity and system resilience as on physical safety and task performance.

Absence of a Physical AI 'GPT Moment' for Perception

Physical AI still lacks a "GPT moment" needed to unlock reliable, general-purpose perception and action. A major challenge for Physical AI is enabling robots to execute plans with human-like perception—perceiving their environment, planning actions, and acting autonomously—while developing general-purpose skills that remain reliable and safe across diverse environments, including rare edge cases. This requires tightly integrated hardware, perception, and motor control so robots can interpret their surroundings and make real-time decisions despite shifting lighting, moving people and objects, unpredictable obstacles, and non-standard layouts.

This gap between human and machine capability is captured by Moravec's paradox. In the 1980s, roboticist Hans Moravec observed that tasks humans find effortless—perception, navigation, and basic dexterity—are extraordinarily difficult to automate, while tasks humans find challenging, such as advanced logic or chess, can often be handled effectively by algorithms.

OpenAI's (private; not covered) 2020 release of GPT-3 is widely cited as an inflection point for large language models (LLMs), but it was enabled by a prior breakthrough: Google's 2017 transformer architecture and self-attention mechanism. Earlier models processed text sequentially, limiting speed and long-range context. Self-attention allowed models to process all words in parallel and weight their relationships, enabling much larger, more capable, and more efficient LLMs.

Seeking analogous gains in robotics, researchers often draw on Daniel Kahneman's *Thinking, Fast and Slow*, which distinguishes between System 1 (fast, intuitive action) and System 2 (slow, deliberate reasoning). A similar dual-system approach is emerging in humanoid robotics: fast, low-latency controllers handle reflexive responses (high-frequency actions), while slower, more powerful models perform planning and long-horizon reasoning (high-level task planning). Reinforcement learning plays a complementary role by enabling robots to iteratively improve these behaviors through trial-and-error interaction, particularly in learning control policies and adapting to new environments.

A central element of this strategy is Vision Language Action (VLA) models, which unify perception, language understanding, and motor control. Rather than separate vision, planning, and control pipelines, VLAs take visual observations and language instructions as input and directly output actions. By closing the loop between understanding and doing, VLAs allow robots to interpret high-level goals—such as "pick up the red cup"—and execute them physically, making them a key bridge toward general-purpose Physical AI.

FIGURE 1. **Basic illustration of a potential Vision Language Action (VLA) model architecture**

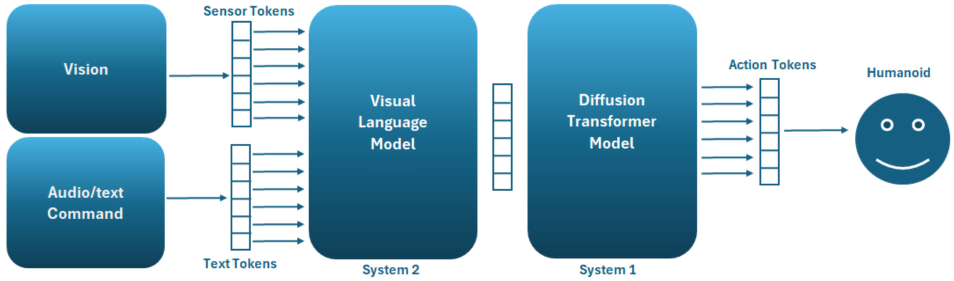

Source: Nvidia & Barclays Research

The long-term ambition is foundational world models for robotics: single systems that can perceive, plan, and act across many tasks and environments, and even transfer across different robot bodies (cross-embodiment) to reduce training costs and improve skill transfer. This would mirror what GPT-style models did for language by creating a scalable platform that improves with data and adapts to new tasks without being rebuilt from scratch.

So why hasn't Physical AI had its "GPT moment" yet? Because physical embodiment adds layers of difficulty that text-only AI doesn't face. High-quality real-world robot data is scarce and expensive, as discussed below in Persistent Physical AI Training Data Gap, simulation only partially closes the gap, and models must run with ultra-low latency and high reliability on constrained hardware. Progress is therefore harder and more incremental—but ongoing advances in VLAs, dual-system designs, data collection, and simulation suggest a breakthrough remains plausible.

Persistent Physical AI Training Data Gap

Efforts are underway to accelerate the data flywheel and close the sim-to-real gap. Data was characterized as the "fuel" that powers AI. While large language models (LLMs) and vision-language models (VLMs) have benefited from vast amounts of text and video data from the internet, there is no comparable repository of rich, diverse data capturing the world from a robot's perspective. Millions of YouTube videos capture human activity, but none of them come with the needed kinematic context—joint motions, actuator commands, sensor feedback, etc.—to directly train robots.

Meanwhile, autonomous vehicles enjoy a unique advantage: millions of cars gathering data on public roads, which general-purpose humanoids cannot currently replicate. Real-world humanoid data collection is inherently slow, expensive, and risky. Even with teleoperation (a human remotely controlling the robot), each robot can only run a limited number of hours per day, and any significant mishap can result in expensive hardware damage or downtime. Teleoperated data also tends to be narrow and task-specific, captured under controlled conditions with a particular robot and setup, lacking exposure to the broad variability and rare edge cases robots will face in the wild.

It's not just about quantity—quality and diversity of data are also critical. Huge volumes of low-quality or overly uniform data can yield "garbage in, garbage out" results: models that work in the lab but fail in the real world. The most valuable training experiences are often the hardest to acquire: unusual, unstructured, hard-to-anticipate scenarios that truly test a robot's abilities and can involve human intervention, yet by definition are rare.

To address the data gap, roboticists are increasingly turning to synthetic data via advanced simulations and digital twins of robots and their environments. Modern physics-based simulators enable thousands of virtual robots to practice in parallel under varied conditions, generating orders of magnitude more data than physical experiments could produce. While simulated data is usually less rich than real data, it offers an 80/20 advantage: allowing developers to explore a wide range of scenarios quickly and safely, then focus limited real-world tests on the toughest remaining cases. Simulations serve as effective digital training grounds, enabling humanoids to learn basic skills before encountering the real world.

Yet even this strategy has limits due a "sim-to-real" gap. Robots trained in simulation often require further fine-tuning with real data to handle the unpredictability of physical environments. A hybrid approach is emerging. For example, Tesla (TSLA, covered by Dan Levy) is leveraging simulation know-how from self-driving cars to train its Optimus humanoid. Elon Musk has described plans for an "Optimus Academy", where tens of thousands of physical robots will train via self-play in a controlled facility, alongside millions of simulated robots running in a high-fidelity physics environment. The goal is to amass experience at unprecedented scale, then use real-world testing to calibrate and validate what's learned virtually—gradually closing the sim-to-real gap.

Additionally, early deployments of incomplete systems in supervised, low-stakes roles can further kick-start learning. Putting a small number of constrained humanoids to work under human oversight can uncover unanticipated failure modes and provide valuable data, which feeds back into the training loop. Over time, this real-world learning flywheel can dramatically improve performance and reduce reliance on simulation alone, much as early autonomous vehicles improved through limited on-road testing.

Physical AI Will Intensify Competition for Compute

Physical AI will drive demand for AI data centers and low-power edge compute. During the conference, there were several mentions of the need for more compute to support humanoid simulation, training, and deployment. It was noted that the perception stack can cost ~$20,000 per humanoid alone. NVIDIA has often described this through a three-tier computing architecture:

- **Simulation compute.** Training humanoids in large-scale, physics-accurate simulations and digital twins requires substantial AI data center capacity and energy. Running thousands of parallel simulations for synthetic data generation and reinforcement learning represents a distinctly new compute-intensive load that traditional software-only AI doesn't face.

- **Foundation model training ("the brain builder")**. Vision Language Action (VLA) models that fuse vision, language, and sensor inputs into action plans are trained in AI data centers. These models can reach 10–20 billion parameters, meaning heavy GPU use, long training cycles, and high infrastructure cost. As humanoid development accelerates, this training burden competes with other AI workloads for scarce compute capacity.

- **On-robot (edge) compute.** Once deployed, humanoids must run inference locally under tight latency and power constraints. Perception-to-action feedback loops often demand responses in just tens of milliseconds to maintain balance and safety, making it impractical to run large foundation models on the robot itself. Instead, models must be compressed, distilled, or re-architected for battery-powered hardware. NVIDIA's open VLA model GR00T N1.6 (~3B parameters) shows the push toward smaller, more deployable AI brains.

This three-tier scheme highlights a fundamental tension: models big enough to learn rich general behavior are expensive to train, while models small enough to deploy on a robot must be highly optimized. Thus, physical AI fuels demand for both more cloud GPU computing (for simulation and training) and specialized, low-power edge hardware (for onboard control). Future progress will rely not only on scaling up compute but also on breakthroughs in model efficiency, compression, and system architecture to deliver state-of-the-art AI within robots' strict power and latency constraints.

Physical AI's Chicken-and-egg Hardware Problem

Hardware presents a scaling challenge. It was emphasized that hardware may be the key bottleneck limiting humanoid scalability. Improving software is much faster than improving hardware, which requires slower development cycles with suppliers. As a result, it's unrealistic to mass produce humanoids until safety, reliability, security, and social acceptance are achieved, posing a classic chicken-and-egg problem for scaling.

Several companies stressed the need to simplify hardware by design: eliminate unnecessary parts, remove cables in moving joints (to prevent failure), embrace symmetry to minimize unique parts, and allow quick field swapping of modules. The industry needs to standardize and commoditize key hardware to drive down costs (targeting ~$20k per humanoid, taking cues from automotive). It was noted that there remains a shortage of off-the-shelf humanoid components as we saw a lot of 3D-printed parts, which are great for prototypes but generally not cost-competitive at scale (as detailed in Additive Manufacturing Reshaping Manufacturing, Layer by Layer).

Enhanced locomotive, manipulation, and dexterity remain critical focus areas. Precise, real-time motor control at each joint is vital and requires tiny, highly integrated chips with built-in current sensing and power management. Enhanced dexterity is especially critical in the hands with leading designs targeting up to ~22 degrees of freedom per hand; one humanoid hand we saw at the conference had limited relative dexterity and still cost ~$2,000 per hand.

Actuators (motors with gear boxes) are another major cost and complexity driver. A single humanoid typically requires 30–60 actuators, and while competition among actuator suppliers is increasing, differentiation is shifting toward deeper integration. Several vendors noted efforts to embed firmware, sensors, and safety features directly into actuators to improve precision motor control, torque performance, and fault detection. The challenge is delivering this sophistication while maintaining reliability, manufacturability, and cost discipline—requirements that will ultimately determine whether humanoids can move beyond prototypes into scalable products.

Sensors are another critical constraint on humanoid scalability. Humanoids require dense, multi-modal sensing across vision, force, torque, touch, and balance to operate safely in unstructured environments. High-performance tactile sensors in the hands, force-torque sensing at joints, and robust proprioceptive feedback (i.e., ability to sense its own position, movement, and force in space) are essential for dexterous manipulation and safe human interaction, but they add cost, complexity, and integration risk. Several companies noted that today's sensor stacks are still too fragile, expensive, or difficult to manufacture at scale. As with actuators, progress will depend on tighter integration—combining sensing, compute, and control into fewer, more reliable modules—while ensuring sensors remain durable, calibrated, and replaceable in real-world deployments.

One supplier suggested that battery performance remains a key issue. If humanoids are running through their battery capacity, backup humanoids will be required for operations to run 24x7 operations, adding additional costs. That being said, some new humanoid designs incorporate hot-swappable battery systems that minimize downtime by allowing quick battery replacement. Notable examples include Boston Dynamics' Atlas, Mentee Robotics' upcoming Mobileye humanoid, Unitree's G1/H1 models, and AgiBot's Expedition series – all of which are built to support on-demand battery swapping for continuous operation.

Interestingly, many humanoid robotics companies are vertically integrating hardware design and manufacturing to mitigate supply constraints, accelerate iteration, and better control costs. For example, 1X (private, not covered) builds its proprietary tendon-driven motors entirely in-house—from copper winding through final actuator assembly—and reports that it has refined these motor designs since 2015, producing roughly 17,000 motors since its California facility became operational. Apptronik (private, not covered) is similarly developing proprietary high-torque actuators with an emphasis on a simplified bill of materials to support scalable production of its Apollo humanoid, complemented by a pilot and strategic manufacturing collaboration with Jabil (JBL, covered by Tim Long) to build Apollo units and deploy them within selected Jabil manufacturing operations. Boston Dynamics (80% owned by Hyundai Motor, 005380.KS, not covered) plans to leverage standardized components from Hyundai's automotive supply chain to improve the reliability and manufacturability of its Atlas humanoids as production scales. Tesla (TSLA, covered by Dan Levy) follows a comparable strategy, repurposing electric vehicle-grade motors, power electronics, and its in-house FSD (Full Self-Driving) computing platform for Optimus, with the longer-term objective of achieving automotive-like production volumes and costs—on the order of tens of thousands of units annually and a target unit cost around $20k over time.

China Leads on Volume, but US May Lead on Capability

The physical AI race seems to reflect diverging approaches. While China clearly leads in volume of humanoid deployments, as detailed in AI gets physical: The decade of the robot belongs to China, one industry expert quipped that many of the Chinese models are almost like "toy robots" that lack the strength, capabilities, and safety for industrial use. Furthermore, they are often either (i) pre-programmed for demonstrations (e.g., dancing humanoids) or (ii) teleoperated by a human. They're not fully autonomous, not general-purpose, not able to do more complex tasks, and not able to operate in rugged environments. One equipment supplier indicated that they view China as a high-volume test bed that lets them experiment and innovate, since the product life cycle in China is much faster with greater risk tolerance, but they believe the US will lead China on developing humanoids that get deployed to do useful work for industrial applications.

Humanoid Developer Landscape

FIGURE 2. Humanoid Developer Landscape (excluding China)

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
1X (Private, 2014, US/ Norway)	NEO – home/consumer humanoid for safe, autonomous household tasks; can switch to remote human teleoperation via 1X "Expert" for complex chores.	Pre-orders opened late 2025 (US$200 deposit; US$20k unit price). First deliveries expected in 2026. Vertically integrated NEO Factory in Hayward, CA ramping to ~10,000 units/year; a second facility (San Carlos, CA) slated for late 2026, enabling capacity ~100,000+ units/ year by end-2027 with added automation.	$100M Series B (Jan 2024); $23.5M Series A2 (Mar 2023). OpenAI Startup Fund (investor).	Does extensive in-house development of critical components (custom "Revo2" actuators, batteries, transmissions, structures, sensor integration) to control costs and mitigate supply risks. Reports it has built ~17k Revo2 motors since production ramp-up with design continuously refined since 2015.	NVIDIA Jetson AGX Thor platform serves as each robot's on-board computer. Runs 1X's video-pretrained world model (1XWM) as the core AI policy for perception and motion planning.
Agility Robotics (Private, 2015, USA)	Digit – humanoid for automating repetitive, labor-intensive tasks in warehouses and factories (bulk tote handling, machine tending, etc.).	Opened RoboFab factory (Salem, OR) in late 2023, with initial output in the low hundreds per year, scalable to >10,000 units/year at full capacity. First Digits delivered to select partners in 2024 (Agility Partner Program); broad commercial availability targeted for 2025.	$150M Series B (Apr 2022). Investors/pilot partners include Amazon, DCVC, Playground Global; strategic partnerships with Toyota (RaaS pilot), GXO (logistics deployment), Schaeffler (investor & manufacturing collaborator), and NVIDIA (AI/simulation partner).	Contract manufacturing & supply partnerships (e.g., forging supply of robotics parts through Schaeffler). Built a new Oregon facility to centralize assembly of Digit's sophisticated leg and arm mechanisms.	On-board AI (vision & control) runs on Nvidia Jetson (plans to upgrade to Jetson Thor for next-gen Digit) and uses NVIDIA's Isaac Sim for virtual training at scale, integrated with Agility's Arc cloud platform for remote management and software updates.
Apptronik, Inc. (Private, 2016, USA)	Apollo – general-purpose humanoid designed to safely perform repetitive tasks alongside humans in factories & warehouses (parts transport, packing, assembly support). Next-gen Apollo prototype scheduled debut in late 2026.	Pilot stage – delivered pilot Apollo units to Mercedes-Benz (auto assembly line deployment, late 2024). Manufacturing partnership with Jabil to build Apollo at scale starting 2025. Volume production to commence after next-gen launch (expected late 2026).	$935M Series A (Feb 2026) led by B Capital. Investors include Google, Mercedes-Benz (pilot/strategic partners); John Deere, Qatar Investment Authority. Jabil serves as key contract manufacturing partner. Also collaborating with Google DeepMind (AI research partner).	Emphasizes mass manufacturability – in-house developed high-torque actuators (reducing SKUs and hardware cost) and simplified component bill of materials for efficient assembly. Partnering with Jabil provides global supply chain access and high-quality volume production processes ("robots building robots" concept with Apollo assisting on its own line).	NVIDIA Jetson AGX Orin (current prototype compute; planned upgrade to Jetson Thor). Extensive use of NVIDIA Omniverse for digital twin simulation in design and test; working with Google DeepMind on specialized AI algorithms for learning and control to accelerate capabilities.
Beyond Imagination (Private, 2018, USA)	Beyond Bot – a teleoperated humanoid (mobile base with dual robotic arms) extending skilled human labor via telepresence. Targeted at factory, pharma lab, and manufacturing support with remote human-in-the-loop control (via VR suit) and autonomous assist.	Pre-commercial – prototypes tested in controlled settings. Aimed to commence first enterprise pilot deployments by 2025–26, initially focusing on environments like pharmaceutical and biotech labs, manufacturing cells, etc., where remote human dexterity is needed.	$100M Series B (May 2025).	Uses off-the-shelf industrial robotic arms and sensors interfaced with proprietary control systems to accelerate prototyping. Focus on teleoperation reduces need for complex advanced autonomy hardware; prioritizes reliable connectivity and human safety features.	"Aura" OS – a unified platform enabling seamless integration of human teleoperators and AI modules. Real-time remote control is combined with assistive autonomy for routine tasks. Plans to incorporate large language and vision models for improved semi-autonomy over time.

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
Boston Dynamics (80%-owned by Hyundai Motor Group (005380.KS), 1992, USA)	Atlas – humanoid for a broad range of heavy industrial tasks (material handling, assembly) designed to operate safely around people. Also produces non-humanoid robots: Spot quadruped, Stretch logistic robot.	Industrial prototype phase – “productized” Atlas unveiled Jan 2026; undergoing trial deployments in Hyundai factories (e.g., parts kitting and sequencing tasks) under oversight. Plans to integrate Atlas into select Hyundai production lines by ~2028 and ultimately expand to complex assembly tasks by ~2030. Hyundai constructing a U.S. robot plant (target ~30k units/year by 2028) to facilitate eventual volume production.	80% stake acquired by Hyundai Motor Group (2021). Key partners: Hyundai Mobis (supplying drivetrain & actuators), Google DeepMind (integrating foundation AI models), NVIDIA (providing advanced hardware & simulation).	Leverages Hyundai’s automotive supply chain: standardized EV-grade motors, sensors, and structural components for high reliability and manufacturability. Designed to minimize unique parts count and maximize robustness in factory conditions.	Atlas’s on-board computing is powered by NVIDIA Jetson AGX Thor, enabling real-time multi-sensor fusion and complex AI-driven control. Integrated with Boston Dynamics’ Orbit fleet management software (currently used for Spot) for logging, updates, & coordination. Atlas is being enhanced with DeepMind’s multi-modal AI models for advanced decision-making and dexterous skills.
Engineered Arts (Private, 2004, UK)	Ameca & Mesmer – ultra-realistic social humanoids for interactive experiences (museums, exhibitions, marketing). Focused on human-like expressions and conversational ability.	Limited production / rentals – ~200 life-like robots delivered worldwide, mostly to museums, theme parks, exhibitions. Building capacity gradually (not yet mass-producing). R&D on next-gen walking/dexterous capabilities (e.g., improved locomotion & AI conversation) ongoing via a Redwood City (US) lab.	$10M Series A (Dec 2024) led by Helium-3 Ventures (Matt Bellamy). Notable installations: MSG Sphere (Las Vegas), GSK (pharma expo events), Computer History Museum (CA).	In-house development of expressive servo actuators and realistic silicone facial/skin systems (“Mesmer” platform) for lifelike appearance. Low-volume craft manufacturing and customization per client use case.	Tritium AI platform integrates third-party large language models and multimodal AI, enabling naturalistic conversations and interactions. Units remain Internet-connected for continuous learning and content updates delivered remotely.
Figure (formerly AI Inc.) (Private, 2022, USA)	Figure 01/02/03 – working towards a full-size general-purpose humanoid for warehouse, logistics & light manufacturing tasks; current prototypes (Figure 01/02) used in test environments.	Pre-commercial – prototypes have completed over 1,000 hours in pilot testing (part-handling trial at a BMW facility in 2024). Scaling engineering & manufacturing — building a pilot production facility in California (BotQ) to produce first saleable units by ~2025–26.	$675M Series B (Mar 2024, co-led by Microsoft & OpenAI). >$1B Series C (Sep 2025) includes Brookfield, NVIDIA, and Intel Capital.	Vertical integration strategy: developing many sub-systems in-house. Negotiating supply agreements for actuators and sensors to support anticipated high volume of production. Partial reliance on contract manufacturers (for complex chassis and transmissions) until its own facilities scale.	After first integrating OpenAI’s GPT-based models, Figure developed its own “Helix” AI for onboard robot learning and control. Relies on extensive simulation (Microsoft Azure cloud + NVIDIA GPU clusters) to train diverse skills before deploying to physical units.
Google DeepMind (Alphabet subsidiary (GOOG), 2010, UK & USA)	Robotics AI R&D – AI research lab developing general-purpose foundation models and advanced AI algorithms for robotics (e.g., multi-modal large models like “Gemini”). Provides software, but not physical robots.	R&D only – no commercial hardware. Contributes to robotics through partnerships. Examples: integrating DeepMind’s AI models into Boston Dynamics’ Atlas (to improve cognition), collaborating with startups (Apptronik, Figure) on their robot autonomy.	A division of Alphabet (acquired in 2014). Partners include Boston Dynamics, Apptronik, Toyota Research, and others on robotics AI.	Not a hardware maker; focuses on cloud compute, data, and software. Robotic partners include Boston Dynamics (integrating AI into Atlas) and Apptronik (Apollo’s AI).	Operates large-scale cloud training on Google’s TPU/GPU data centers to develop cutting-edge robotics algorithms. Develops high-level cognitive control and perception models that can be integrated into partner companies’ robots (e.g., vision-language models, reinforcement learning frameworks) to enhance their autonomy and versatility.
Kawasaki Heavy Industries (7012.JP, 1896, Japan)	Kaleido (RHP) – long-running in-house humanoid R&D platform (latest prototype Kaleido 9). Aimed at proving heavy-lift, rough-terrain humanoid capabilities for disaster response and industrial support.	R&D only – prototypes (Kaleido series) regularly demonstrated (e.g., carrying 18 kg payload in 2025 trials) but no commercialization plans yet. Focus on incremental technology improvements; emphasis on robotics research over near-term productization.	Corporate-funded R&D (KHI). Partnerships: co-founded a “Physical AI center” in Silicon Valley in 2026 with NVIDIA, Microsoft, Fujitsu to explore advanced AI integration.	Applies KHI’s heavy machinery expertise: custom hydraulic/pneumatic actuators, industrial controllers, robust exoskeleton-based designs. Prototypes built at KHI labs.	Incorporating cutting-edge AI from partners (e.g., Nvidia Isaac Sim used for simulation, Microsoft’s cloud AI for control prototypes). Focus on stability and precision under extreme conditions; uses classical control algorithms complemented by emerging cognitive modules.

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
Mentee Robotics (Acquired by Mobileye (MBLY), 2019, Israel)	Third-generation, fully integrated humanoid (in development; no brand name public yet) built for autonomous operation in industrial and logistics tasks, with in-house actuators, dexterous hands, and advanced perception.	Pre-commercial – accelerated development under Mobileye (since acquisition in 2026). Prototypes in testing; proof-of-concept deployments expected ~2026 in controlled environments (e.g., Mobileye's facilities). First limited production run targeted ~2028, likely leveraging Mobileye/Intel's manufacturing resources.	Acquired by Mobileye (Intel) in Jan 2026 (~$900M deal). Previously bootstrapped.	Integrates Mobileye's automotive-grade vision sensors and high-precision drive systems into humanoid hardware. Developing proprietary actuators (high torque-to-weight) and tactile hands in-house to maintain tight integration and cost control.	Leverages Mobileye's computer vision and RSS safety frameworks from autonomous driving; employs few-shot learning pipelines to allow robots to rapidly learn new tasks with minimal human demonstrations.
Mind Robotics (Private - Rivian spinout, 2025, USA)	Factory automation robotics R&D – not a humanoid form factor (focuses on mobile manipulators & multi-axis arms). Aims to create AI-driven hardware for complex EV manufacturing tasks (assembly, material handling).	R&D stage – prototypes tested within Rivian's EV production lines; no commercial product offerings yet. Using Rivian's facilities as testbeds during 2025–26. Targeting external pilot systems in automotive & aerospace by 2027–28, if R&D milestones met.	$500M Series A (March 2026) co-led by Accel & Andreessen Horowitz; spin-out from Rivian Automotive (major shareholder and development partner).	Initially repurposes off-the-shelf industrial robot components (arms, mobile bases) for quick prototyping; long-term, hardware manufacturing would be outsourced to established automotive/ robotic partners if scaled.	Emphasizes foundation model-driven planning – training large AI models on Rivian's production data. Heavy use of simulation on Rivian's high-performance computing, combined with on-robot GPUs for real-time control and adaptation.
NEURA Robotics (Private, 2019, Germany)	4NE-1 – full-size "cognitive" humanoid in development for heavy-duty industrial use, aiming for high strength (arms targeting ~100 kg lift).	Pilot trials – prototypes are undergoing evaluation in select industrial settings; no public sales yet. A strategic partnership with Schaeffler (2025) will supply actuators and refine design, with plans to deploy thousands of NEURA units across factories by the mid-2030s. A manufacturing facility in Germany is being established to support future production.	~€8 M seed/angel (2019–21). Schaeffler (strategic manufacturing & integration partner); support from local government programs in Germany.	Emphasizes high-performance hardware (strong servo motors, heavy-payload gearboxes). Pursues a European-centric supply chain for quality control and autonomy. R&D and assembly operations co-located in Germany; iterative design influenced by pilot feedback from industrial partners.	Focused on proprietary "Industrial Cognitive Engine" – combining large language models with symbolic AI reasoning – to enable advanced decision-making.
PAL Robotics (Private - sub of Royal Group (UAE), 2004, Spain)	TALOS – a human-size force-controlled bipedal robotics platform for R&D and custom enterprise solutions (e.g., assembly assistance). Also sells TIAGo mobile manipulators and custom robots for specific use cases.	Low-volume R&D sales – a handful of TALOS units sold to research institutions (e.g., used by Airbus for assembly process trials). No mass production plans; primary business is project-based development of custom robotic solutions (TALOS, TIAGo, etc.).	Self-funded plus EU grants. Collaborations: Airbus (R&D pilot for airplane assembly tasks using TALOS); EU Horizon research programs.	Uses standardized actuators and ROS-based control frameworks, allowing high customization per client use case. Low-volume, in-house assembly of robots.	Primarily open-architecture control – supports ROS-based frameworks and custom sensor integration. No dedicated in-house AI system provided.
Rainbow Robotics (277810.HK, 2011, South Korea)	HUBO & RB-H – humanoid R&D platforms (building on KAIST's Hubo series) for research and tech development; also sells collaborative robot arms (cobots) and mobile robots (current revenue generators).	Humanoids in R&D – no commercial humanoid product yet; focus on advancing biped locomotion & manipulation. Uses HUBO derivatives as testbeds while generating revenue via cobots and mobile robots for industrial clients.	Public listing in 2021. Samsung Electronics acquired ~35% stake by 2025.	Leverages KAIST's decades of humanoid research (Hubo labs) and reuses components from existing robot products (cobot actuators, controllers) for humanoid prototypes. Samsung partnership opens access to high-efficiency manufacturing and supply chain resources if commercialization proceeds.	Current prototypes utilize ROS and Rainbow's custom motion control algorithms.

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
Sanctuary AI (Private, 2018, Canada)	Phoenix – full-size general-purpose humanoid designed for "labor-as-a-service" in workplaces (retail, warehouses, manufacturing). Gen 7 (2023) introduced 20-DoF humanlike hands (25 kg lift capacity).	Pilot phase – completed successful retail pilot (Mark's shop, 2023: 110+ tasks performed by Phoenix via mix of autonomy & telecontrol). Iterating Gen 8 prototypes (2024) to improve dexterity and reliability; expanding pilot programs into logistics and manufacturing (multi-week trials) through 2026. Aiming for late-decade productization.	C$75+ million raised (2018–2023), including C$58.5M Series A (Mar 2023) and C$30M federal SIF funding. Investors: Bell Canada, Verizon, Export Development Canada, Workday Ventures. Partners: Apptronik (hardware collaboration on joint designs); Cycorp (cognitive AI research).	Early prototypes assembled in-house using off-the-shelf actuators and sensors; now exploring contract manufacturing for initial production scale-up. Key hardware components (e.g., actuators, batteries) currently sourced from established robotics suppliers, with plans for custom designs as volumes grow.	Cognitive architecture ("Carbon") – multi-module AI (vision, language, memory, motor control) that combines large language models with logical reasoning for general intelligence. Training approach blends human teleoperation demonstrations with reinforcement learning to teach skills.
Tesla (TSLA, 2003, USA)	Optimus ("Tesla Bot") – general-purpose humanoid targeting unsafe, repetitive, or mundane tasks in factories (long-term vision includes consumer use). Built using Tesla's EV components (electric actuators, high-density battery packs, FSD computer).	Prototype phase – multiple Optimus units (Gen 2/3) built since 2021; tested internally in Tesla's automotive plants. As of 2024, two bots performing simple autonomous tasks (material transport, basic assembly) at a Tesla factory. Tesla aims to achieve large-scale production in the late 2020s (long-term goal of tens of thousands of units once safety and performance are proven).	Self-funded via Tesla's internal R&D.	Leverages Tesla's existing EV supply chain (in-house motors, power electronics, FSD computer chips) to reuse proven components and minimize costs. Plans for mass production involve expanding Optimus assembly at Giga Texas using Tesla's Gigafactory infrastructure. Tesla designs its AI chips (FSD, Dojo D1) internally but outsources fabrication to chip foundries (e.g., TSMC, Samsung). Elon Musk has also expressed interest in establishing a dedicated chip fabrication facility (a 'Terafab').	Autopilot-derived AI: Optimus leverages Tesla's vision-first perception approach, neural-network tooling, and training infrastructure originally developed for Autopilot/FSD, with models trained on Tesla's Dojo supercomputing platform. Current prototypes employ a hybrid control approach, combining learned perception and motion primitives with scripted task logic and supervised testing, resulting in limited but improving autonomy.
Toyota Motor Corp. (7203.JT, 1937, Japan)	T-HR3 – advanced teleoperated humanoid (~1.5 m) for assistive and remote presence tasks (e.g., healthcare, construction, disaster relief). Toyota also produces a smaller Human Support Robot (HSR) – a wheeled assistive service robot.	R&D focus – T-HR3 prototypes have been publicly demonstrated since 2017 (remote-controlled via master arm exosuits) but no current commercialization plans. Toyota's robotics projects remain long-term research efforts (some limited pilot tests in internal facilities and at "Woven City" living lab).	Internally funded by Toyota (Toyota Research Institute). Collaborations with academic partners (University of Tokyo, MIT) for robotics and AI research.	Harnesses in-house automotive-quality servos, sensors, and high-reliability control hardware. Prototype production performed at Toyota's R&D groups; if commercialized, would leverage Toyota's manufacturing and supply chain infrastructure.	Teleoperation-centric – T-HR3 is primarily controlled by a human operator via Toyota's full-body Master Maneuvering System exosuit with visual and force feedback. The robot itself has minimal autonomous perception or decision-making.

Source: Company reports & Barclays Research

FIGURE 3. **Humanoid Developer Landscape (China)**

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
AgiBot (Zhiyuan Robotics) (Private, 2023, China)	Expedition A2 – full-size general-purpose humanoid designed for multi-sector deployment (manufacturing, logistics, retail, security)	Early mass production – delivered 5,000 units by late 2025, with management targeting continued scaling through 2026. Deployed across multiple pilot verticals; still refining autonomy and robustness.	Undisclosed private funding; founded by ex-Huawei executives.	Vertically integrated manufacturing centered in Shanghai. In-house design and production of motors, actuators, batteries, and soft materials; tightly coupled domestic supplier network enabling rapid cost-down and fast production ramp.	Proprietary embodied-AI stack (GO-series foundation models) integrating vision, language, and control. Extensive use of simulation and large-scale real-world data collection; on-board NVIDIA Jetson-class edge compute
AI2 Robotics (Private, 2023, China)	AlphaBot series – full-size general-purpose humanoids for diverse industrial tasks (auto, factories).	Pre-commercial / pilot – initial prototypes deployed in late 2025 (e.g., trial at a Dongfeng auto plant). Aggressive scaling plan (1,000→10,000 units in 2026) reliant on internal ramp-up (not yet validated).	Raised ~CNY 1 B (~US$150 M) across multiple quick rounds in 2025 (state-owned industry funds, private VC, and internet giants).	Heavy vertical integration – developing core hardware (actuators, sensors) and end-to-end manufacturing capacity internally; targeting vehicle-scale production efficiency (AI-driven assembly lines) by 2026.	Vision-Language-Action “AlphaBrain” – a powerful embodied AI model guiding planning & control. Extensive synthetic data pre-training and simulation to enable generalizable skills with minimal real-world fine-tuning; large-scale cloud training.
Booster Robotics (Private, 2023, China)	T1 & K1 – medium (T1) and compact (K1) humanoids for R&D, education, and competitions.	Commercial – launched K1 globally via online platforms in late 2025 (1,000+ unit orders; 700+ units delivered to 200+ clients across 20+ countries). Booster-powered teams won RoboCup 2025 (Adult & Kid sizes) using these platforms.	$28 M raised (2023–25).	Standardized design & manufacturing – leverages modular, affordable hardware (mass-produced actuators, controllers) to achieve scale.	Open SDK & multi-sensor stack – built on NVIDIA Orin & Isaac SDK. Integrates vision & LiDAR for perception; developer-friendly API for custom behaviors. Emphasizes stability and ease-of-use over advanced autonomy (for learning & competitive programming).
DEEP Robotics (Private, 2017, China)	DR02 – full-size industrial humanoid, also Lynx series wheel-leg robots and Jueying quadrupeds.	Prototype / R&D – unveiled DR02 (Oct 2025). Focusing on extensive field testing for harsh environments (power grid inspections, emergency response). Scaling plans undisclosed (emphasis on technology demonstration and partnerships).	CNY 500 M Series C (~US$70 M, Dec 2025).	Leverages existing supply chain from quadrupeds for motors and perception systems; modular “quick-detach” component design for faster maintenance.	Embodied AI platform – multi-sensor (vision + thermal + LiDAR) fusion and robust autonomy optimized for hazardous scenarios. Developing “DeepVLA” navigation & multi-modal control system to enable stable autonomous operation across varied environments (e.g., outdoors, extreme weather)
ENGINEAI (Private, 2023, China)	T800 – full-size industrial humanoid for inspection, manufacturing, and heavy-duty tasks; complemented by PM01/ SA01 platforms.	Initial mass production ramp in 2026 – in 2026 with a target throughput of one T800 unit per 15 minutes (~10k per year). First T800 batches delivered; second production base planned in Zhengzhou to expand capacity.	$200 M Series B (Apr 2026)	Highly automated, closed-loop factory covering component testing, assembly, and quality control. Full-stack in-house R&D across joints, actuators, and control electronics; multi-regional manufacturing footprint planned.	Multi-modal embodied-AI architecture trained heavily in simulation; reinforcement learning for locomotion and manipulation. NVIDIA Jetson-based edge compute for on-robot inference.

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
EX Robots (Private, 2013, China)	EX humanoid series – hyper-realistic bionic humanoids (lifelike androids) for entertainment, hospitality, and R&D.	Limited production – small-scale build-to-order manufacturing (roughly 2–4 units per month) for museums, expos, and special projects; pricing ~CNY 1.5–2 M per unit. No large deployments (intended for exhibit and demonstration scenarios).	Not available.	In-house fabrication of lifelike components – internal production of silicone skin, robotic animatronics, and custom micro-motors. Focus on bespoke assembly rather than mass manufacturing; single-site factory in Dalian.	Expressive AI – proprietary software emphasizes facial expression recognition and emotional interaction. A multi-modal interaction system (vision + voice + expression mapping) drives human-like facial movements and reactions; high actuator density in head/face for fine emotional cues.
Flexiv (Private, 2016, China/USA)	Rizon series – adaptive 7-DoF robotic arms (upper-body robots) with force control for flexible manufacturing tasks.	Commercial – shipped 100+ Rizon units by 2020; serving industrial customers (auto, electronics) with repurchase orders and initial paid deployments. Expanding into new industries (logistics, agriculture, healthcare) with upcoming dual-arm and mobile platforms (e.g., MICO, Enlight).	>US$122 M total raised. Strategic partnerships with major Chinese manufacturers (for integration into assembly lines and distribution channels).	Hybrid production model – core hardware (force-control actuators, sensors, and control boards) and algorithms developed in-house; final assembly at dedicated factories in China (Shenzhen, Foshan) for cost efficiency. Strategic supply agreements for key components.	Adaptive control & force-sensing AI – advanced closed-loop force feedback and computer vision allow fine manipulation and assembly tasks. Combining proprietary force control algorithms with conventional machine vision and deep learning for flexible automation in dynamic environments (not LLM-driven).
Fourier (Fourier Intelligence) (Private, 2015, China)	GR-1 / GR-2 – electric, full-size general-purpose humanoids evolved from Fourier's exoskeleton platform.	Small-batch production / early commercial – GR-1 units shipped to research labs and pilot customers from 2024. GR-2 launched in 2024 and remains in testing; broader commercialization expected after design stabilization post-2025.	¥400 M Series D (~$60 M, Jan 2022) led by SoftBank Vision Fund; Prosperity7 (Aramco) also investor.	Core proprietary electric actuators ("Fourier Smart Actuators") enable non-hydraulic architecture optimized for manufacturability. Mix of in-house assembly and domestic component suppliers.	Vision-language-action stack supported by heavy simulation training. Teleoperation and task learning supported through multi-modal perception and ROS-compatible tooling.
GAC (Guangzhou Auto) (2238.HK, 2005, China)	GoMate – third-generation embodied-intelligence humanoid with hybrid wheel–biped mobility (38 DoF).	R&D / internal pilots – unveiled December 2024. Internal trials within GAC factories during 2025; component-level batch production planned first, followed by small-batch whole-robot production targeted for 2026.	Corporate-funded R&D by GAC Group.	Leverages automotive-grade components and GAC's EV supply chain, including solid-state battery technology (~6-hour runtime). Variable wheel–leg chassis improves stability and energy efficiency.	Vision-centric navigation and decision stack adapted from autonomous driving research; fusion of vision, inertial, and haptic sensing.
Galbot (Private, 2023, China)	Galbot S1 & G1 – bipedal service humanoids for retail and logistics support (unmanned stores, pharmacy robots).	Initial commercial deployments – ~10 humanoids implemented in Beijing pharmacies (2025) performing autonomous order retrieval and delivery; planning to roll out 100 robot-run stores nationwide by late 2025.	~CNY 2.4 B (~US$335 M) raised (mid-2025) from strategic investors (e.g., CATL) and local government funds.	Vertical stack – designs own hardware (servo motors, control boards) and custom electronics. Partners with Bosch for industrial-grade components. Emphasizes cost-effective fabrication and durability (target <¥600k per unit) via AI-optimized design workflows.	Sim2Real AI paradigm – heavy reliance on synthetic data pre-training to reduce real-world data needs. Combines high-fidelity simulation with minimal targeted real-world fine-tuning. Focuses on robust vision and dexterous object handling with minimal human supervision.

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
Hiwonder (Private, 2015, China)	Tonybot & DIY Kits – small-scale open-source humanoids and robotics kits (e.g., 17-DoF Tonybot for STEM education).	Commercial – ongoing low-volume product sales since ~2017. Thousands of programmable kits and mini-humanoids sold globally via e-commerce (Amazon, AliExpress) for hobbyist, school, and STEM use.	Self-funded (no major VC rounds). Profitable from kit sales; partnerships with electronics distributors to reach overseas customers (AliExpress, Amazon).	Modular manufacturing – leverages off-the-shelf microcontrollers (Arduino/ESP32/Jetson) and standard servo components; designs custom servo lines (LX series) for higher performance. Assembly in Shenzhen; kits emphasize user-friendly build & customization.	Educational AI – supports accessible programming (Scratch, Arduino, Python) with integrated vision and speech modules for demonstrations. Provides multi-modal interaction on advanced models (Tonybot uses small language models for voice & vision interplay) to teach coding and AI concepts.
Kepler Robotics (Private, 2023, China)	Forerunner K2 – general-purpose humanoid emphasizing hybrid wheel–leg mobility for energy efficiency.	Pilot / small-batch production – initial K2 units (dozens) delivered to partners for co-development. Limited production underway; broader rollout dependent on pilot outcomes.	Undisclosed early-stage venture funding.	Rapid-iteration manufacturing using domestic suppliers; hybrid mobility reduces actuator count and battery load, simplifying production.	Standard vision and LiDAR perception combined with proprietary motion planning; cloud-connected telemetry for continuous software improvement.
Leju Robotics (Private, 2016, China)	AELOS / ROBAN / KUAVO series – humanoids from compact educational models to full-size **Kuavo** general-purpose robots (industrial & service).	Pre-production / early manufacturing – established a high-volume humanoid production line (Guangdong) in early 2026 with target capacity ~10,000 units/year. Small-scale deliveries (Aelos/Roban kits) since 2018; scaling up mass production of Kuavo full-size robots through 2026.	~¥1.5 B (~US$200 M) raised in 2025. Early backers (2019–21) included Tencent AI Accelerator, Microsoft Accelerator, Lenovo Capital, etc. Partnering with Schaeffler (actuator supply).	Ecosystem approach – vertical integration of core components, plus global partnerships (with Huawei, Alibaba, Haier) to leverage supply chain depth. Automated assembly lines (one robot build ~30 min) planned to support widespread adoption.	Open-source & large-model integration – built on ROS with proprietary gait-control algorithms. Integrating large models (Huawei's Pangu) for high-level cognition. Targets multi-industry deployment (from retail guides to manufacturing co-workers) via adaptable AI.
Noetix (Private, 2023, China)	Bumi – compact consumer humanoid (94 cm, 12 kg) for family-friendly companionship and education; also larger N2 (118 cm) in development.	Pre-commercial – Bumi launched Oct 2025 (~¥9,988 on JD.com); thousands of pre-orders across national shopping festivals (Singles' Day, 12.12). N2 prototype ran a half-marathon demonstration in Apr 2025, garnering 2,500 orders (targeting shipments mid-2026).	No formal VC funding disclosed (likely seed-funded by founders & local grants; part of Beijing humanoid innovation hub).	Low-cost design – repurposes smartphone-grade components (batteries, sensors, SoCs) for affordability. Lean assembly in Beijing and partner EMS facilities to achieve consumer-level pricing (~US$1,400).	Consumer-grade AI & control – proprietary motion control system for smooth walking and interactive gestures; open programming interfaces for educational use. Emphasizes limited autonomy (1–2 hr battery, supervised mode) with basic vision & voice modules for family-friendly interactions.
RobotEra (Private, 2023, China)	STAR1 – high-performance humanoid optimized for speed and outdoor mobility (record ~3.6 m/s).	R&D / pilot – prototypes demonstrated record-setting runs in 2024. No commercial deployments announced; focus remains on robustness and stability.	Angel and early institutional funding (undisclosed).	Custom lightweight structures and shock-absorbing joints; low-volume prototype manufacturing using advanced machining and rapid prototyping.	Vision- and IMU-based perception with GPU-accelerated control optimized for high-speed locomotion and balance.
UBTECH Robotics (9880.HK, 2012, China)	Walker S1 / S2 – full-size humanoids for factory and commercial service roles.	Early commercial deployments – Walker units delivered to automakers and logistics firms; 500+ pre-orders reported as of 2024 (per IPO filings). Production scaling following 2023 IPO.	Hong Kong IPO (December 2023). Pre-IPO strategic investor Tencent; key partners include Foxconn and BYD.	Designs core servos and controllers in-house; Foxconn partnership supports scalable assembly. Supply chain largely China-centric with automotive and electronics overlap.	Vision- and SLAM-based navigation with sensor fusion (cameras and LiDAR); cloud-edge coordination over 5G for fleet management and updates.

Company (Status / Founded / HQ)	Robot Product / Program	Commercial & Production Status	Funding / Key Investors & Partners	Supply Chain	Perception Stack
Unitree Robotics (Private, 2016, China)	G1 – compact humanoid developer platform (~1.3 m) aimed at R&D and education; complements Unitree's quadruped lineup.	Early shipments – G1 kits delivered to labs and developers starting late 2024. Focus remains on developer and research markets rather than enterprise deployment.	Bootstrapped via product sales (no major outside VC disclosed).	Cost-optimized design using standardized parts; manufacturing scale and supplier relationships inherited from high-volume quadruped robot production.	NVIDIA Jetson Orin NX with ROS-based open software stack enabling developer experimentation and customization.
Xiaomi Corp. (1810.HK, 2010, China)	CyberOne – full-size humanoid concept showcasing Xiaomi's AI and hardware capabilities.	R&D only – internal pilots in Xiaomi EV factory during 2025–26. No commercialization roadmap disclosed.	Internally funded by Xiaomi.	Repurposes Xiaomi smartphone and IoT supply chain components (cameras, sensors, SoCs) for rapid prototyping; internal actuator and controller development.	Xiaomi in-house vision, speech, and emotion-recognition AI; largely teleoperated with limited autonomy, focused on human–robot interaction.
XPENG Robotics (XPeng Inc. subsidiary (9868.HK), 2020, China)	IRON – highly anthropomorphic humanoid for service and industrial scenarios.	Pre-production – next-generation IRON unveiled November 2025. Mass production targeted by end-2026, with first commercial deployments (retail and enterprise) expected in 2027.	Funded by XPeng Motors via internal R&D; Alibaba Cloud acts as a technology partner for mapping and cloud infrastructure.	Full-stack in-house development drawing on automotive supply chain (actuators, hands, batteries, proprietary Turing AI chips). Dedicated humanoid robot factory under construction in Guangzhou.	XPeng "Physical AI" stack combining vision-language-action foundation models with high-TOPS on-board compute derived from autonomous driving systems.

Source: Company reports & Barclays Research

Related Research

- Thematic FICC Research - AI gets physical: Ten things to know about humanoids
- Thematic FICC Research- AI gets physical: The decade of the robot belongs to China
- FICC Research - Global economic drivers: How humanoid robotics matters for macro
- Global Macro - From macro to markets: Embodied AI: Wealth creation or economic displacement?
- Thematic FICC Research - The future of work: AI gets physical
- Thematic FICC Research - AI gets physical: The decade of the robot