Deutsche Bank
Research

Asia
Industrials

Industry

Humanoid Robot Pulse

Model evolution; China humanoid exports; Tesla production ramp-up (3Q26 edition)

Date
7 October 2026

Iris Zheng, CFA
Research Analyst
+852-2203-5884

Edison Yu
Research Analyst
+1-212-250-7263

Winnie Dong
Research Analyst
+1-212-250-5121

Laura Li
Research Associate
+1-212-250-2266

Xu Liu
Research Associate
+852-2203-5885

Amid the heavy newsflow around humanoid robots, we summarise the key developments in the market globally in this quarterly series.

DB views: Encouraging model evolution; China humanoid robot export to the US rose in August; Tesla ramping up humanoid production

Large models for robots are advancing rapidly: **Figure AI** and **Generalist AI** have both revealed models with zero-shot or low-shot generalization capabilities, whilst robot simulation can now be enhanced by **Atlas from World Labs**. Furthermore, the integration of **OpenAI's GPT-6 Astra** (LLM) with **Physical Intelligence's π0.5** (VLA model) could significantly improve task success rates. Large models are widely considered to be the key bottleneck for the humanoid robots industry, and whilst we await a 'ChatGPT moment' for the sector, **we are encouraged by the rapid progress in this space**.

Separately, we find it surprising that **China's exports of bionic AI robots (mainly robot dogs and humanoid robots) to the US grew further in August**, despite the US restrictions introduced in late July. Meanwhile, several sources indicate that **Tesla is ramping up its Optimus humanoid robot production** at Fremont. We believe this bodes well for supply chain companies, **although challenges remain** for Tesla in areas such as hand manufacturing and chip memory.

Within our APAC Industrials coverage, we prefer Shuanghuan (Buy, closing price RMB38.09), Hengli (Buy, closing price RMB90.31), and Harmonic Drive (Buy, closing price ¥6,600) for the humanoid robot theme. We think Tesla's production ramp-up is also helpful for Sanhua (Hold, closing price RMB35.71). Within our US Auto & Auto Tech coverage, we highlight Tesla (Buy, closing price US$380.68) and continues to monitor Mobileye (Buy, closing price US$7.47) (through Mentee Robotics) as humanoid robotic OEMs.

We summarise the key developments in 3Q26:

- **Model evolution delivering higher success rates and improved generalization capabilities:** The integration of GPT-6 Astra and π0.5 significantly improves robot success rates. World Labs (to be acquired by AMD for US$8.2bn) unveiled Atlas, which can be utilized for robot

Deutsche Bank AG/Hong Kong
IMPORTANT RESEARCH DISCLOSURES AND ANALYST CERTIFICATIONS LOCATED IN APPENDIX 1. Deutsche Bank does and seeks to do business with companies covered in its research reports. Thus, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. THE CONTENT MAY NOT BE DISTRIBUTED IN THE PEOPLE'S REPUBLIC OF CHINA ("THE PRC") (EXCEPT IN COMPLIANCE WITH THE APPLICABLE LAWS AND REGULATIONS OF PRC), EXCLUDING SPECIAL ADMINISTRATIVE REGIONS OF HONG KONG AND MACAU.



simulation. Additionally, models from Figure AI and Generalist AI have achieved zero-shot or low-shot generalization.

- **China humanoid robot exports:** China's bionic AI robot exports totalled 2,079 units in August. Within this, exports to the US grew further to 812 units (up from 623 units in July), reaching a record high.
- **Tesla Optimus:** Tesla is ramping up production of its Optimus robots to hundreds per week with a goal of ~1,000 units per week by year end, although it is facing challenges related to hand assembly and chip memory.
- **Agility Robots:** The company plans production for Digit 5, rising from hundreds units to ~1,000 units, and eventually 10k units which is the planned capacity at its Salem Oregon factory.
- **Large-scale robot deployment plans:** JD.com plans to procure 3mn robots over the next five years, whilst Toyota aims to deploy 400k humanoid robots over the medium term.

We also highlight reports from DB Research in 3Q26:

- Mid-year reality check on the global humanoid robot market, and an comprehensive analysis on ecosystem flywheel, published on 10 August. Link to report.
- Policy update: US ban on foreign-made robotics (including humanoid, quadruped, mobile robots), published on 29 July. Link to report.
- Expert call and site visit takeaways: Chinese robotic reducer expert (link to report); Chinese humanoid robot expert (link to report); Kedali site visit (link to report)

Model evolution: GPT-6 Astra + π0.5 greatly improves success rate of robots

After the launch of GPT-6 Astra by OpenAI in early September, the Galaxy Universal team (under Galbot) from China published a report on GitHub in September where the team has integrated **GPT-6 Atra** with **Physical Intelligence's vision-language-action (VLA) model π0.5** and yielded impressive results, according to 36Kr (link).

The research is conducted under **a simulated environment** and analyses two closed-loop control architectures for dual-arm operation:

- **Direct mode:** Without running π0.5, **GPT-6 Astra directly** reads the three-way camera feeds and robot status, outputs the pose of the dual-arm end effectors and the opening/closing state of the grippers, and **executes 1 to 5 control steps each time**.
- **Mixed mode:** At each decision moment, **π0.5 first generates 50 candidate actions. GPT-6 Astra** observes the same feeds, execution history, and the corresponding dual-arm trajectories of the candidate actions, then makes one of two choices: **either follow the first 1 to 15 steps generated by π0.5, or generate 1 to 5 steps of end-effector pose correction by itself**.

After completing this segment, it observes the new frame and makes a new decision.

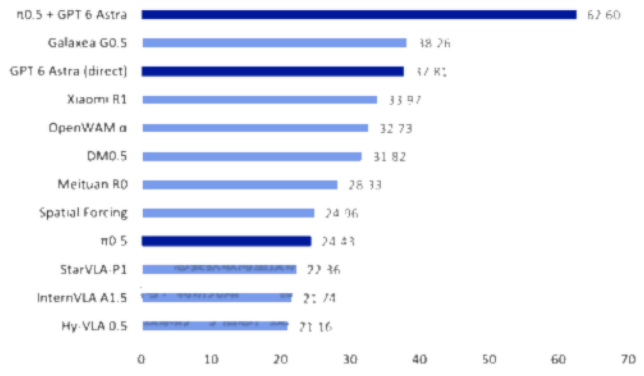
Firstly, the team conducted the **RoboDojo experiment** and selected 10 types of complex robot operation tasks. The results show that **the mixed mode achieved a 48% success rate and an average score of 62.60**, which is even 64% higher than the score of the second-place entry at 38.26.

More notably, **Astra only modified 14.4% of all executed control steps**, meaning that most of the actions are still completed by π0.5, and Astra is mainly responsible for understanding tasks, judging the current status and deciding whether corrections are needed at key nodes. The experiments show that Astra is best at semantic understanding, visual recognition, spatial relationship judgment and task planning. However, large models suffer from shortcomings when it comes to tasks that require precise contact, stable grasping and continuous control, such as building towers and drawing mahjong tiles. **This suggests that large language models like Astra could help with robot policy, but still lacks the understanding of the physical world and the operation of the physical world are still two distinct capabilities.** A typical problem revealed by the experiments for Astra is that understanding the goal does not mean being able to complete the action stably, and this is where π0.5 generates value.

Furthermore, the team conducted the **RoboLab experiment** with another 10 types of tasks. The results demonstrate **Astra's advantages in examining the understanding of general robot policies** on vision, object relationships, task instructions and spatial structures. The experiments also demonstrate the **zero-shot generalization** as Astra was not additionally trained for these specific tasks.

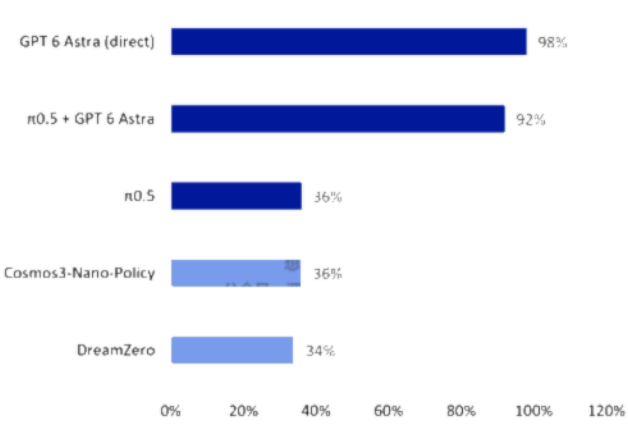
Interestingly, in the RoboDojo experiment, π0.5+Astra consumed a total of about **625mn tokens**, while Astra Direct exceeded **1.13bn tokens**.

Figure 1: RoboDojo-Sim: Average score across 10 tasks



Source: Company data, Deutsche Bank Research; Note: Average score across the ten selected tasks. The ten tasks were selected and stratified based on $\pi_{0.5}$'s publicly reported success rates, including tasks with low success rates. As the official model results do not provide aggregated data for the same task subset under a standardized randomized testing setup, the figures differ from those in the official overall leaderboard covering all tasks.

Figure 2: RoboLab: Success rate of 10 tasks



Source: Company data, Deutsche Bank Research; Note: Success rates for the ten tasks selected by RoboLab. Each method was tested 50 times, with an equal number of trials for each task. GPT 6 Astra (direct) achieved 49/50, $\pi_{0.5}$ + GPT 6 Astra achieved 46/50, $\pi_{0.5}$ and Cosmos3-Nano-Policy both achieved 18/50, and DreamZero achieved 17/50. The horizontal axis represents the success rate (0–100%).

Figure 3: RoboDojo-Sim: Per-task score (0–100)

No	Task	DM0.5	GalaxeaVLA (G0.5)	Xiaomi Robotics-1	OpenWAM-α	π0.5	π0.5 + GPT 6 Astra	GPT 6 Astra (direct)
1	Organize table (整理桌面)	44.00	46.33	57.67	62.50	23.33	60.00	30.00
2	Classify objects by language (按语言分类物品)	0.47	1.07	2.00	1.33	0.60	38.00	60.00
3	Imitate sorting sequence (模仿排序顺序)	1.80	1.67	2.50	2.90	1.60	55.00	0.00
4	Arrange the largest number (排出最大数字)	7.85	4.11	8.56	4.36	2.29	50.00	57.00
5	Pack objects into box (物品装箱)	14.72	17.12	18.69	20.83	18.36	50.00	50.00
6	Classify objects (物体分类)	26.83	10.33	17.20	5.53	24.67	71.00	100.00
7	Build a tower (搭塔)	55.20	82.93	52.60	52.53	37.73	64.00	12.00
8	Draw mahjong tiles (麻将杠牌)	56.67	90.00	41.33	32.00	26.67	40.00	0.00
9	Fold clothes (叠衣服)	28.96	32.75	41.49	51.31	29.12	100.00	40.00
10	Put bottles into bin (瓶子入桶)	81.70	96.30	97.70	94.03	79.93	100.00	36.00
	Overall	**31.82**	**38.26**	**33.97**	**32.73**	**24.43**	**62.60**	**37.81**

Source: Deutsche Bank Research, 36Kr; Note: Average score across the ten selected tasks. The ten tasks were selected and stratified based on $\pi_{0.5}$'s publicly reported success rates, including tasks with low success rates. As the official model results do not provide aggregated data for the same task subset under a standardized randomized testing setup, the figures differ from those in the official overall leaderboard covering all tasks.

Model evolution: Atlas from World Labs for spatial intelligence; AMD to acquire World Labs for US$8.2 billion

World Labs introduced its **omni world model (全模态世界模型) Atlas**, on 1 September 2026 (*link*). According to the company, Atlas is a multimodal autoregressive diffusion transformer (多模态自回归扩散变换器): **all inputs are combined into a shared spatial context.** Atlas uses that context to generate what comes next, staying consistent in 3D with everything it has seen and imagining what lies beyond it.

Atlas can perform a broad range of tasks spanning **world generation, reconstruction, and simulation:**

- **Camera-Controlled Generation:** Atlas generates images and videos from one or more images with pixel-perfect camera control, outputting up to 1 minute of video at 1440p.
- **Spatial Reconstruction:** Atlas reconstructs real world scenes from one to dozens of input images. It generates both image frames from novel views and explicit 3D outputs, outperforming state-of-the-art models specialized for 3D reconstruction.
- **Space-Time Simulation:** Atlas models space and time from input videos, reframing videos for dramatic visual effects and **enabling Real-to-Sim workflows for robotics.**



- **Image Generation:** Atlas generates images and 360 panoramas from text; it can follow complex prompts, render text, and generate a wide variety of visual styles.

For robotics, Atlas can help scale Real-to-Sim for both navigation and manipulation. Atlas can reconstruct a space in explicit 3D from a few images. Then, as a simulated robot moves through space, Atlas also generates the RGB and depth data its sensors would observe along the way. The world and the robot's view of it come from the same model.

- **Robot navigation examples:** The company captured two large environments with a cell phone video, using 24 frames each for reconstruction. Scanning spaces like these traditionally requires elaborate and expensive equipment. The company then simulates different kinds of robots navigating different paths and use Atlas to generate images from the perspective of the robot's body-mounted cameras.
- **Robot manipulation example:** From a few casual recordings, Atlas aids in building a simulation that also captures how objects move and interact. Once a task is simulated, users can vary it easily: change the objects, their positions, the robot's motion, the lighting, and the background. The result is diverse training data and testing environments for robotics at scale.

Figure 4: The evolution roadmap of World Labs' model

Time	Evolution	Type	Core capability
Dec 2024	Generating World	AI system	Persistent, navigable 3D world generation
Sep 2025	**Marble**	Multimodal World Model	Create 3D worlds from text, images, video, or coarse 3D (3D粗模) layouts; interactively edit, expand, and combine worlds 3D worlds can be exported as Gaussian splats (高斯泼溅), meshes (网格), or videos
Oct 2025	**RTFM (Real-time Frame Model)**	World Model	Autoregressive diffusion transformer operating on sequences of frames Learned rendering without relying on explicit 3D representations
Sep 2026	**Atlas**	Omni World Model	Multimodal autoregressive diffusion transformer: all inputs are combined into a shared spatial context Camera-contolled world generation, spatial reconstruction, and space-time simulation

Source: Company data, Deutsche Bank

On 28 September 2026, **AMD announced that it has signed a definitive agreement to acquire World Labs**, the spatial intelligence startup co-founded by AI pioneer Dr. Li Feifei (*link*). According to AMD, this acquisition will help the company "develop the hardware, software and systems that will power the next generation of AI and strengthen the open AI ecosystem". Dr. Li will join AMD as Executive Vice President and Chief Scientist, continuing to lead World Labs team alongside its other two co-founders. The all-stock transaction is valued at approximately **US$8.2 billion**, making it the second-largest acquisition in AMD's history, behind only its US$35 billion acquisition of Xilinx. AMD expects the deal to close by the end of 2026.

Model evolution: Figure AI released Helix 2.5 for zero-shot generalization

On 17 September 2026, Figure AI introduced **Helix 2.5**, demonstrated as its most advanced neural network to address the question: whether a humanoid can enter an unfamiliar home and immediately perform whole-body tasks autonomously (*link*).

Helix 2.5 was pretrained on Index, Figure AI's global-scale robot training dataset unveiled in August, and learned three specific behaviors: tidying rooms, folding towels, and making beds. During testing, humanoid robots were placed in 30 unseen homes and involved 420 household tasks. Result shows that, Helix 2.5 achieved **a 56% zero-shot success rate**, namely without any prior data collection or fine-tuning. Figure AI claims this marks the first demonstration of zero-shot whole-body generalization on a humanoid robot.

The Index pretraining proved to be the key differentiator on this generalization capability: models trained without Index recorded a mere 8% success rate, whereas those trained with Index surged to 56%. Additionally, Helix 2.5 introduced robust self-correction capabilities. When encountering errors, the robots demonstrated the ability to autonomously step back, adjust their stance, and reposition themselves to rectify the mistake before proceeding tasks.

Figure AI revealed that **Index is currently generating 35 minutes of new human experience every second**.

Figure 5: The evolution roadmap of Helix – Helix 02 – Helix 2.5

Time	Evolution	Type	Core capability
Feb 2025	**Helix**	VLA model "System 1, System 2"	Full-upper-body control S2: "Think slow", a pretrained VLM for scene understanding and language comprehension S1: "Think fast", a latent-conditional visuomotor transformer, handles low-level control
Sep 2025	Go-Big	Pretraining Project	Direct transfer from 100% egocentric human video data to robot behavior
Jan 2026	**Helix 02**	VLA model Add "System 0"	Whole body control, loco-manipulation S0: a foundation model, learned prior over how people move while maintaining stability
Aug 2026	Index	Robot training dataset	Figure-exclusive system for sourcing real-world physical human behaviour data at scale
Sep 2026	**Helix 2.5**	VLA model	Zero-shot whole body generalization of Index-pretrained tasks in unseen environment Self-correction over long horizons

Source: Company data, Deutsche Bank Research

Model evolution: Generalist AI launched GEN-1.5 with one-shot/few-shot learning capability

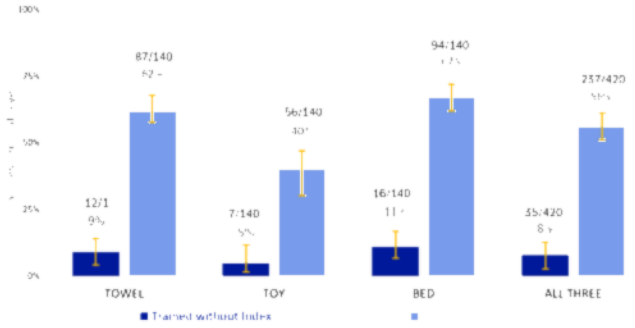
On 19 August, Generalist AI unveiled **GEN-1.5**, their latest robot foundation model that exhibits broad one-shot and few-shot learning from demonstration capabilities, as well as zero-shot generalization. GEN-1.5 is a large multimodal that processes video input (30s of memory, alongside other sensor, language, and proprioceptive inputs) and produces 100Hz action trajectories.

GEN-1.5's capabilities include the below. Generalist AI notes that these capabilities appear to emerge directly from pretraining on large amounts of physical interaction data.

- **One-shot learning via in-context prompting:** The model learns new tasks in seconds when prompted with 3-12s of a single demonstration, no training required.
- **Few-shot** adaptation via gradient descent: The model can be fine-tuned to a new task in 1-10 gradient steps on 1-5min of data (~10-15 demonstrations).
- **Improvising new strategies and tool use:** The model generalizes at the level of behavioral strategies: forming entirely new trajectories to reach a goal, using unseen tools (e.g. brush, dustpan, etc.) to create new solutions to tasks demonstrated with other tools, and working ambidextrously even when prompted or fine-tuned to perform the task with a specific hand.

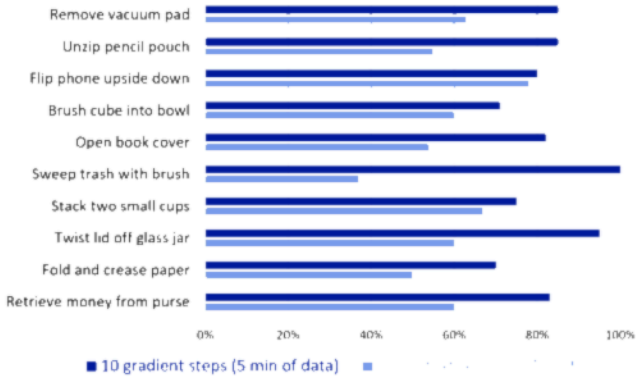
Experiments across 10 tasks by Generalist AI show **59% (±10% std. dev.) average success with one-shot in-context prompting**, straight from the pretrained model. With few-shot learning, **performance rises to 83% (±9% std. dev.) via 10 gradient steps on 5min of data per tasks** (~50 demonstrations).

Figure 6: Helix 2.5 index pretraining drives zero-shot generalization



Source: Company data, Deutsche Bank Research; Note: Zero-Shot Success in 30 Homes: Success rates of models trained from scratch vs. models trained with Index. Index accounts for the majority of Helix 2.5's zero-shot capability.

Figure 7: GEN-1.5 exhibits general ability from just one-shot or few-shot demonstrations



Source: Company data, Deutsche Bank Research; Note: GEN-1.5 can learn short-horizon atomic manipulation tasks at 83% average success rate after 10 steps of gradient descent on 5 minutes of data per task, and 59% with zero gradient updates and 3 to 12 seconds of demonstration data using emergent in-context learning.

Figure AI: Four-stage roadmap for building humanoid robots

Separately, in a live interview with RoboStrategy in September, Figure AI's CEO Brett Adcock shared his **"four stages of maturity" framework for building humanoid robots** (*link*). The four stages include:

1) **Building great hardware:** with the required torque, speed, runtime, dexterity, and range of motion, combined at an acceptable weight and cost.

2) **Whole-body AI architecture:** getting the system to run well end-to-end, from input commands to output long-horizon autonomy.

3) **Scaling intelligence:** the primary bottlenecks here are data and compute. Adcock believes humanoid robotics will ultimately require far more of both than LLMs do.

4) **Scaling manufacturing:** once a robot is sufficiently intelligent to do more tasks, faster, and with a higher success rate, the focus shifts to integrating humanoids into the economy at massive scale.

Adcock argues **the first two - hardware and AI autonomy - can't simply be solved by investing money**; they're long-duration, deep-engineering problems, which he compared to "orbital rockets".

By contrast, **the last two - intelligence and manufacturing - can be accelerated with capital**. Adcock believes doing these final stages correctly will eventually require hundreds of billions of dollars.

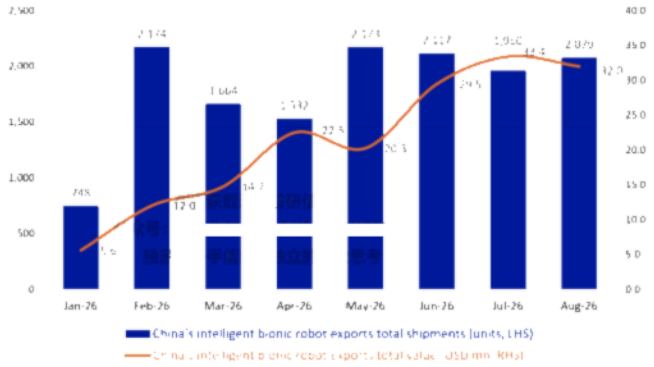
China humanoid robot exports to the US still rose in August despite restrictions

Since January 2026, China Customs has disclosed exports of **"intelligent bionic robots"** as a standalone category. This category covers AI-powered autonomous robots that resemble human or animal forms or structures and can independently perform tasks. We believe "intelligent bionic robots" primarily consist of quadruped robots and humanoid robots.

In total, export volume reached close to 15,000 units in 8M26, remaining at around 2,000 units per month from May to August. Meanwhile, export value totaled US$170mn, with an ASP of ~US$11,700 in 8M26. **In August, export volume edged up 6% MoM to 2,079 units**, whilst export value declined by -4.5% MoM as ASP fell -10% from July's series high. This ASP decline was mainly attributable to lower export ASPs in the US (down 35% MoM) and Asia (down 14% MoM). Despite the pullback, the total export ASP in August remained nearly double the 1Q26 level.

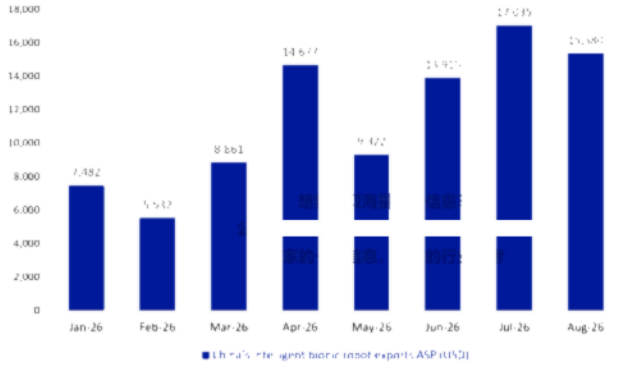
What surprised us was the resilience of exports to the US: China's intelligent bionic robot **exports to the US grew further in August to 812 units**, representing **a 30% MoM increase from 623 units in July**. This occurred despite the US FCC's ban since late July on new models of foreign-made robots (including humanoid, quadruped, and mobile robots) (*link to report*). That said, the ASP of exports to the US dropped significantly to ~US$8,500 in August, down from over US$13,000 in July.

Figure 8: China's intelligent bionic robot exports total shipments (units) and total value (USD'mn)



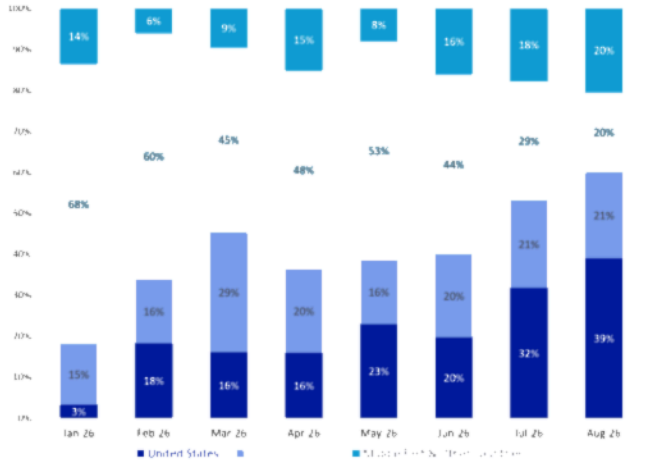
Source: China Customs, Deutsche Bank Research

Figure 9: China's intelligent bionic robot exports ASP (USD)



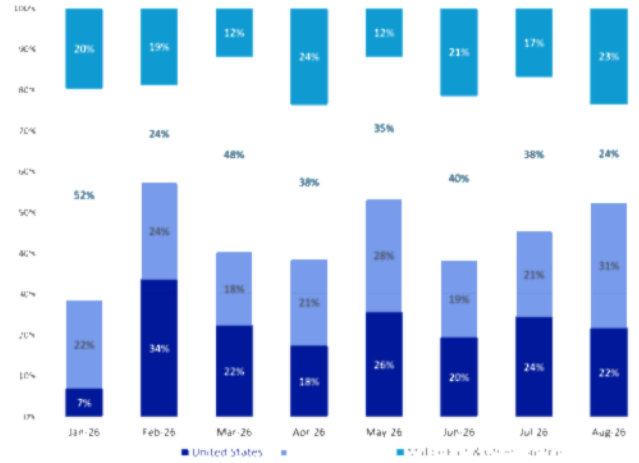
Source: China Customs, Deutsche Bank Research

Figure 10: China's intelligent bionic robot exports shipments (units) by region mix



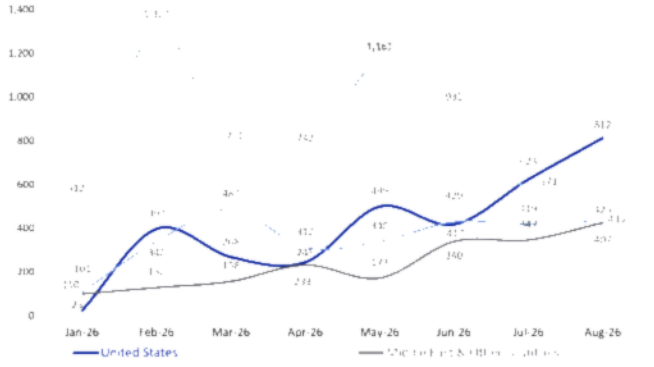
Source: China Customs, Deutsche Bank Research

Figure 11: China's intelligent bionic robot exports value (USD) by region mix



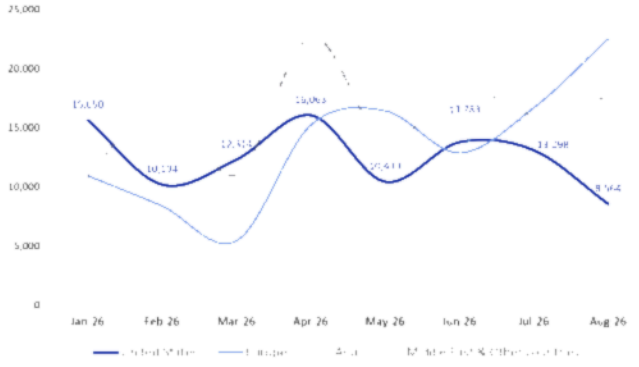
Source: China Customs, Deutsche Bank Research

Figure 12: China's intelligent bionic robot exports shipments to key regions (units)



Source: China Customs, Deutsche Bank Research

Figure 13: China's intelligent bionic robot exports ASP by key regions (USD)



Source: China Customs, Deutsche Bank Research

Tesla ramping up production of Optimus humanoid robots amid challenges

The Information *(link)* reported on 25 September that Tesla has increased its production of Optimus by ~tenfold in recent months, **rising to several hundred units per week in August** at Fremont compared with only a few dozen in 2Q26. Tesla is reportedly targeting a continuous automated line capable of **more than 1,000 robots a week by end of the year**, and **an eventual target of ~20,000 a week**. Meanwhile, Tesla is facing challenges and key bottlenecks include:

- **Complex hand assembly:** Each hand and forearm reportedly contains more than 100 screws and other small components, many of which require manual assembly.
- **Production-line precision:** Equipment used to assemble the robots' hands and joints does not always align components accurately, resulting in higher rework requirements.
- **Durability and sensor issues:** The hands are not sufficiently durable for sustained operation, while some touch sensors have experienced reliability problems. Tesla has developed a replaceable sensor-equipped "sensing glove," which is expected to be incorporated into **next year's model**.
- **Limited task generalization:** Optimus still requires job-specific programming and training. Even basic new tasks can take several days to learn.
- **Controlled deployment:** Most units are currently used internally for testing, AI training and data collection. Factory deployments remain limited to closely supervised and tightly controlled environments.

Separately, per Yahoo Finance *(link)*, Elon Musk posted on 1 October that **Tesla has cut the AI5 chip's RAM memory in half to 72GB and the AI6 chip's memory by a third to 144GB**, calling it "the only way to get enough volume for Optimus production," which "greatly" reduces cost. He added that the cuts should have a "negligible effect on Optimus performance," as memory bandwidth is a bigger limiting factor than total memory capacity. Simultaneously, **Micron**'s CEO comments on the earnings call on 30 September that **humanoid robots are expected to need more than 200GB of memory and multiple terabytes of storage per unit**, similar to autonomous vehicles.

Additionally, The Paper in China *(link)* reported on 21 September that **Tesla's robotics team arrived in Ningbo** on 16 September to commence a new round of **factory audits at key suppliers**, including **Tuopu Group** (601689.SS), **Sanhua** (002050.SZ), and **Joyson Electronics** (600699.SS). Prior to these audits, all of these companies had already secured orders for Tesla's humanoid robot, though allocation volumes may vary. The news also clarified that the audits are not restricted to Ningbo; rather, Tesla's robotics team is visiting supply chain partners across the broader Yangtze River Delta region, including Hangzhou and Shanghai. Most of these partners are existing suppliers that already collaborated with Tesla on its automotive business. Meanwhile, Sina news (*link*) reported on



17 September that Tesla recently placed **a new order for approximately 5,000 Optimus** with its suppliers. Following previous pre-production orders of several hundred units, this marks the first mass-production order at the thousand-unit scale.

We have visited several component suppliers for Tesla Optimus. The most shared production ramp-up at Tesla by the suppliers is **1,000-1,500 units per week by September/October**, and **over 2,000 units per week by end of 2026**, with exponential growth in 2027. The comments are directionally consistent with the news reports.

Jason Calacanis forecast Optimus fleet to hit one billion units by 2036 after a private Optimus demonstration

On his "This Week in Startups" podcast in late August (*link*), the American tech entrepreneur and venture capitalist Jason Calacanis highly praised Tesla's Optimus, stating that, "Optimus is going to be the greatest product ever made by humanity". According to his forecast, there will be **one billion Optimus worldwide by 2036**. **After watching a private Optimus demonstration recently**, Jason asserted that the brain of Optimus is "five generations" beyond those robots that were showcased by Chinese domestic players at the World Humanoid Robot Games in Beijing.

Agility Robotics: Updates from prospectus; Digit 5 launch

On 4 September 2026, Agility Robotics and Churchill Capital Corp XI (SPAC) jointly filed Form S-4, the registration statement for the proposed business combination. Key takeaways from this document include:

- **Valuation:** A pre-money equity value of US$2.5bn, with expected gross proceeds of over US$620mn – comprising US$421mn from CCXI's trust account (assuming no redemption rights exercised) plus US$201mn from PIPE financing led by Foxconn at US$10 per share.
- **Financials:** Agility Robotics delivered **revenue of US$1.78mn for FY2025** (including US$1.13mn from related parties), representing a more than **fivefold increase** over US$310,301 recorded in FY2024. Following substantial SG&A and R&D expenses, operating loss reached US$140mn in FY2025, up from US$70mn in FY2024. Loss per share stood at US$40.29 and US$22.03 in FY2025 and FY2024 respectively. Robot sales totaled US$1.55 million in FY2025, which included **a purchase of five robots by Amazon totaling approximately US$1 million**.
- **Operations:** Agility Robotics operates under two business models: Robot-as-a-Service (RaaS) model and Ownership model. In RaaS model, customers pay one-time deployment fee of US$25,000 for an individual robot onsite and a recurring **monthly subscription fee of US$8,500** to get access to software and maintenance services. In Ownership model, customer pay **an upfront purchase cost of US$200,000**, a one-time deployment fee of US$20,000, and an annual Arc software subscription fee and maintenance fee of US$36,000.
- **Industry application:** Digit v4 robots have already been deployed across warehouse, logistics and manufacturing customers including Amazon,



GXO Logistics, Schaeffler, Toyota Motor Manufacturing Canada, and Mercado Libre. Current commercial deployments are concentrated in material-handling tasks such as moving bins and totes. Additionally, Agility's revenue exhibits **high concentration** among few customers. For instance, a single customer has committed to a US$300mn order for 1,000 Digit v5 robots under a three-year RaaS contract, yielding an estimated ASP of US$300,000 per unit.

- **Capacity:** Agility Robotics opened a 70,000 sq ft manufacturing facility "RoboFab", in Salem, Oregon, with claimed **annual capacity up to 10,000 units**. In July 2026, the company opened new Fremont facility with size of 60,000 sq ft, for software, AI training and R&D. Agility Robotics assumes a weighted adoption curve for Digit 5, starting at ~800 units in 2027, scaling to ~7,000 units in 2030 and reaching up to ~25,000 units in 2035. The ultimate production will depend on pace of deployment.

Separately, on 15 September 2026, Agility Robotics unveiled **Digit 5**, engineered specifically for safe human-robot collaboration at scale. Agility is the first launch partner for NVIDIA's Halos safety framework, enabling Digit 5 to collaborate safely in factories and warehouses without traditional physical barriers.

Building on three years of commercial deployment experience from Digit 4, Digit 5 introduces massive hardware upgrades:

- **40% higher payload** – Repeatedly lifts loads up to 50lbs (~23kg), a nearly 40% increase from Digit 4's 35lbs, powered by cycloidal actuators integrated into leg design.
- **Faster charging** – Achieves 10:1 run-to-charge ratio (90-minute runtime with 9-minute charge), significantly up from Digit 4's 2:1 ratio.
- **Swappable end-effectors** – Allows Digit to dynamically switch grippers when transitioning between different tasks, such as tote handling and machine tending.
- **Extended reach** - Digit 5 stands 5' 11" (1.81 m) and weighs 284lbs (129 kg). It extends its maximum reach height to 7.2 feet (2.2 m), up from Digit 4's 5.5 ft.

Agility said it also plans to target global deployment for Digit 5, **beginning with the EU and United Kingdom**. Early access to Digit 5 is expected to **begin in the first half of 2027 in N. America first**, with general availability for manufacturing, warehouse and distribution operators expected **by the end of 2027**. The company also most recently announced that it signed a MOU with Fort Robotics, a safety platform building the Trust Layer for physical AI, on a partnership that further advances the safety infrastructure of Digit 5.

Large scale robot deployment plan by JD Logistics and Toyota

JD Logistics aims for the world's largest embodied robot fleet

At JDDiscovery 2026 (the JD Global Tech Explorers Conference) held on 9 September 2026, **JD.com** unveiled its **Physical AI Acceleration Plan**, aimed at driving embodied AI from real-world validation to commercial deployments (*link*).

Detailed initiatives outlined across each field include:

- **Logistics**: Over the next five years, JD Logistics plans to procure **3 million robots**, 1 million autonomous vehicles, and 100,000 drones, committing to establish the world's largest embodied robot fleet. This massive deployment aims to advance the fully automated end-to-end logistics workflows.
- **Data collection**: JD Cloud plans to establish the world's largest data collection center for embodied AI, targeting the acquisition of over **10 million hours of real-world human video data** within the next two years.
- **Large model**: JD launched JoyAI-Echo WM, a real-time interactive world model at this event. JD has established a comprehensive JoyAI foundation model matrix - spanning multimodal, world, and embodied models - to provide the core model capabilities required to power physical AI deployments.
- **Computing power**: In collaboration with domestic partners, JD Cloud plans to construct a 100,000-GPU high-density cluster. The initiative aims to build an autonomous, highly reliable, and low-latency intelligent computing infrastructure to support the training, simulation, and scale deployment of its JoyAI large models.

Toyota to install 400,000 self-developed humanoid robots

On 18 September 2026, **Toyota Motor** announced that the Group would **invest JPY1 trillion (approx. US$6.4 billion) per annum starting from 2028** to accelerate factory automation and robot deployment (*link*). Under this plan, Toyota will introduce **400,000 in-house-developed humanoid robots**, deploying 150,000 units across its own plants, and the remaining 250,000 across group affiliates worldwide. The Executive Vice President of Toyota stated that their goal is to create a world where robots coexist with humans, rather than replacing them.

Named "ELEY" (Embodied Learning Robot for Enhanced Yield), the 50-kg humanoid robot moves on a wheeled mobile base. ELEY is capable of performing delicate tasks such as grasping parts, transferring components from one bin to another, and folding fabric. Currently, workers wear specialized grippers modeled after ELEY's hands to repeat routine tasks, allowing the robots to learn directly from real-world human demonstrations.

Key catalysts for the humanoid robot space



Figure 14: Humanoid robot events in next several months

Humanoid robot events		
12-15 Oct 2026	AI	2026 OCP Global Summit in California
12-16 Oct 2026	Robot	2026 CIIF (中国国际工业博览会) in Shanghai
22/10/2026	Tesla	3Q26 results (BBG estimate)
Late 2026	Tesla	To unveil Optimus V3
Late 2026	Tesla	To start Optimus V3 mass production
Late 2026	XPeng	To start mass producing humanoid robots
2026	Li Auto (理想)	Wheeled humanoid robot to be released in 2026
2026	DEEP Robotics (云深处)	In A-share IPO process (Completed response to the first round of IPO inquiries on 7 August 2026)
2026	AGIBOT (智元)	In H-share IPO process (Confirmed on 24 July 2026)
2026	DOBOT (越疆)	In A-share IPO process (IPO application approved by Shenzhen Stock Exchange on 22 July 2026)
2026	Leju Robotics (乐聚)	In A-share IPO process (IPO application accepted by Shenzhen Stock Exhange's ChiNext Board (创业板) on 19 May 2026)
2026	Xiaomi	New humanoid robot for mass production to be released in 2026
2026	Galbot (银河通用)	Reportedly in IPO process
19/02/2027	Unitree Robotics (宇树)	Last day of the lock-up period for certain offline investors (一部分网下投资者限售期到期) — 0.56% of total shares
2027	GAC (广汽)	To start mass producing humanoid robots

Source: Public information, Deutsche Bank

Links to previous reports on Humanoid Robots

- **Rise of Humanoid Robots expert call (V):** Takeaways from a Chinese robotic reducer expert and our channel checks (V) - 17 Sep 2026
- **Humanoid robot reducer:** Kedali (not covered) site visit takeaways - 17 Sep 2026
- **Humanoid Robot:** Mid-year reality check; Ecosystem flywheel spinning - 10 Aug 2026
- **Humanoid Robot:** US bans foreign-made robotics (humanoid, quadruped, mobile robots) - 29 Jul 2026
- **Rise of Humanoid Robots expert call (IV):** Takeaways from a Chinese humanoid robot expert - 28 Jul 2026
- **Humanoid Robot Pulse (2Q26 edition):** US activities stepping up; China scaling fast; Supportive policies - 16 Jul 2026
- **Humanoid Robot:** Comparing Unitree, UBTECH, DEEP, Dobot and Leju: Opportunities and risks coexist - 15 Jun 2026
- **Mobileye:** Mentee humanoid deep dive - 9 Mar 2026
- **Humanoid Robot (III):** Six visions for 2026 – Scaling, iterating and diversifying - 6 Mar 2026
- **Humanoid Robots Pulse (4Q25 edition):** Emerging electronics and Japanese players; Humanoid policy - 18 Dec 2025
- **Rise of Humanoid Robots expert call (III):** Takeaways from expert call with Humanoid - 4 Dec 2025
- **Newsletter:** Tesla, Sanhua, AgiBot, Unitree, Figure update - 23 Oct 2025
- **Newsletter:** Tesla Optimus; Unitree IPO; Nvidia and Hon Hai - 12 Sep 2025
- **Rise of Humanoid Robots expert call (II):** Takeaways from an expert at a leading Chinese humanoid robot start-up - 17 Jun 2025
- **China Macro:** The upcoming AI Boom in China - 26 Mar 2025



- **Rise of Humanoid Robots expert call (I):** Takeaways from expert call with Agility Robotics - 20 Feb 2025
- **Humanoid Robot (II):** Go big and go home - 14 Feb 2025