


China takes the lead in robots



JPR
Jon Peddie Research

4 Saint Gabrielle Court
Tiburon, CA, 94920-1619
+1 415 435 9368
www.jonpeddie.com

WHITE PAPER

China Takes the Lead in Robotics

An Analysis of China's Strategic Advance in Robotics and Automation

JON PEDDIE RESEARCH

Tiburon, California



Copyright Jon Peddie Research 2026. All rights reserved.

Reproduction in whole or in part is prohibited without written permission
from Jon Peddie Research.

This report is the property of Jon Peddie Research (JPR) and made available to a restricted number of clients only upon these terms and conditions.

Agreement not to copy or disclose. This report and all future reports or other materials provided by JPR pursuant to this subscription (collectively, "Reports") are protected by: (i) federal copyright, pursuant to the Copyright Act of 1976; and (ii) the nondisclosure provisions set forth immediately following.

License, exclusive use, and agreement not to disclose. Reports are the trade secret property exclusively of JPR and are made available to a restricted number of clients, for their exclusive use and only upon the following terms and conditions. JPR grants site-wide license to read and utilize the information in the Reports, exclusively to the initial subscriber to the Reports, its subsidiaries, divisions, and employees (collectively, "Subscriber"). The Reports shall, at all times, be treated by Subscriber as proprietary and confidential documents, for internal use only. Subscriber agrees that it will not reproduce for or share any of the material in the Reports ("Material") with any entity or individual other than Subscriber ("Shared Third Party") (collectively, "Share" or "Sharing"), without the advance written permission of JPR.

No material from JPR may be used for AI training. The term "AI training" refers to the process of teaching a machine learning model to make predictions or decisions based on data. That includes but is not limited to: 1. Supervised Learning: Training with labeled data (input-output pairs), such as image classification, spam detection, sentiment analysis. 2. Unsupervised Learning: Training with unlabeled data to find patterns or groupings, such as clustering (e.g. K-means), dimensionality reduction (e.g. PCA). 3. Self-Supervised Learning: A hybrid where the system generates labels from the data itself, such as contrastive learning, masked language modeling (used in GPT, BERT). 4. Reinforcement Learning (RL): Agents learn by taking actions and receiving rewards or penalties, such as game-playing AIs (e.g. AlphaGo), robotics, navigation. 5. Unsupervised Reinforcement Learning: Agents explore environments without a defined task or reward, such as Pretraining agents for general skills before fine-tuning. 6. Multi-Task Learning: Training one model to handle multiple tasks simultaneously, such as a model that can translate and summarize text. 7. Transfer Learning: Using a pre-trained model on a new, related task with minimal retraining, such as fine-tuning GPT or BERT on a specific domain. 8. Few-Shot / One-Shot / Zero-Shot Learning. Few-shot: Trained on a very small number of examples. One-shot: Trained on just one example. Zero-shot: Performs tasks without any task-specific training. 9. Meta Learning ("Learning to Learn"): Training models that can adapt to new tasks with very little data, such as Model-Agnostic Meta-Learning (MAML). 10. Curriculum Learning: Training models by gradually increasing task difficulty. 11. Online Learning: Model updates continuously as new data arrives, such as stock prediction systems. 12. Federated Learning: Training across decentralized devices without sharing raw data, such as mobile keyboard prediction models. 13. Contrastive Learning: Learns by comparing positive and negative pairs of data. Used in: Self-supervised vision models like SimCLR, CLIP. Specialized AI Training Terms, Fine-tuning: Adjusting a pre-trained model on a new dataset. Pretraining: Initial broad training, usually on large datasets (e.g. GPT pretraining). Prompt Tuning / Instruction Tuning: Adjusting models using instructions or task-specific prompts. RLHF (Reinforcement Learning from Human Feedback): Used in GPT models to align behavior with human intent. Distillation: Training a smaller model to mimic a larger model. Active Learning: Selecting the most informative data points to label next. Semi-supervised Learning: Mix of labeled and unlabeled data for training.

Subscriber shall be liable for any breach of this agreement and shall be subject to cancellation of its subscription to Reports. Without limiting this liability, Subscriber shall be liable for any damages suffered by JPR as a result of any Sharing of any Material, without advance written permission of JPR.

Important Note related to Sharing of Material: If Subscriber wishes to Share any Material with any Shared Third Party, it shall contact JPR and request permission to do so. JPR shall forthwith prepare an appropriate Sharing agreement, which shall set forth the name and contact information of the Shared Third Party, the terms and conditions of such Sharing, the compensation to JPR for such Sharing, and the agreement by Subscriber and the Shared Third Party that the Shared Third Party shall not, under any circumstances, Share any of the Material with any other entity or person.

Disclaimers. JPR Reports contain "reviews" of various products and information. The content of these Reports represents the interpretation and analysis of statistics and information that are either generally available to the public or released by entities or individuals deemed responsible. The information in the Reports is believed to be reliable but is not guaranteed as to its accuracy or completeness. The Reports do not endorse or attempt to sell any products and should not be interpreted differently. JPR assumes no responsibility for the correct or incorrect usage of any trademarks or service marks.

China's robot boom outpaces its chip independence

Deployment leads the world; chips still lag.

China has taken a clear lead in humanoid robots, home to more than 140 companies building them today. Electric vehicle makers are rushing into the sector: Xpeng plans mass production this year, and BYD unveils its first humanoid this month. That scale now doubles as a live test of China's ambitions in physical AI, the race to build machines that see, think, and move on their own. The deployment numbers, the chip supply chain, and the tug-of-war over materials all tell one connected story.

China's humanoid robot boom has grown into one of the most closely watched proxies for the country's AI chip self-sufficiency push. More than 140 domestic humanoid robot companies now compete for share, drawing in electric vehicle makers eager to apply their manufacturing scale, supplier networks, and battery expertise to a new product category. XPeng plans mass production of its Iron humanoid this year, running three of its own Turing chips per unit for 2,250 TOPS of effective compute, a vertical-integration move that skips third-party silicon entirely. Cumulative Turing shipments already exceed 200,000 units, with roughly 1 million targeted for 2026. BYD enters the race this month with its Xiao Di humanoid and has not disclosed its chip supplier; one trade report claims BYD runs perception and control on chips designed by its own semiconductor unit, a claim BYD itself has not confirmed.

#	Supplier	Chip / Product	Confirmed Chinese humanoid customers
1	Nvidia (US)	Jetson AGX Thor, Jetson T4000, Jetson AGX Orin	Unitree H2, UBTech, Galbot G1, Zhiyuan/AgiBot, Zhongqing, EngineAI, UniX AI — Nvidia's Orin still powers most Chinese humanoid start-ups (TechCrunch); the Unitree partnership was announced June 1, 2026
2	D-Robotics/ Horizon Robotics (China)	Xuri S600 SoC (560 TOPS INT8, integrated "brain + cerebellum")	Beijing Innovation Center's Tiangong 3.0, Linglong 2.0, TARS A3, UBTech, Fourier, Booster, Robotera, X Square Robot — "more than 20 leading clients," backed by a \$150M Series B2 in April 2026
3	Huawei (China)	Ascend 310B/310P NPU + Kunpeng CPUs, CloudRobo platform	UBTech (comprehensive cooperation agreement, May 2025); Leju Robotics' Kuavo runs on Huawei Cloud's Pangu embodied model; CloudRobo launched at WAIC July 2026
4	XPeng (in-house)	Turing AI chip — 750 TOPS effective per chip	XPeng's own Iron humanoid uses 3 Turing chips (2,250 TOPS effective); cumulative Turing shipments exceed 200,000 units, with ~1 million targeted in 2026 (XPeng)
5	Rockchip (China)	RK3588 SoC + RK1820/1828 AI coprocessors (6–26 TOPS)	UBTech's mass-market U1 consumer humanoid, unveiled July 2026 at 119,800–990,000 yuan — RK3588 is called "the mainstream choice for edge computing... for humanoid robots in China" (SCMP)
6	Qualcomm (US)	Dragonwing IQ10 (up to 700 TOPS), TurboX IRB10	Stardust Intelligence/Atribot's humanoid at WAIC 2026, the first product-level IQ10 robotics deployment (Thundercomm); reference design unveiled at Computex June 2026
7	Black Sesame Technologies (China)	Huashan A2000 ("brain," up to 1,000 TOPS) + Wudang C1236 ("cerebellum")	Tianwen humanoid (with Wuhan University); Rovar X3 and a Fourier dexterous hand demoed on its SesameX platform at CES 2026
8	MetaX (China)	Purpose-built embodied-AI edge chips (specs not yet disclosed)	50/50 joint venture with UBTech, signed June 12, 2026, targeting sample production in H2 2027
9	Cambricon (China)	Siyuan 590/690 AI accelerators	Named among 10 domestic chipmakers with Day-0 adaptation for the Wise Kaiwu 3.0 world model; role is largely training/adaptation rather than confirmed onboard modules (Embodied Global); targeting ~500,000 accelerators shipped in 2026 (Tom's Hardware)
10	AMD (US)	Ryzen 7 8845HS APU (host CPU, not the AI accelerator)	Fourier Intelligence's GR-3 humanoid uses it as the general-purpose "body controller," typically paired with a Jetson Orin/Thor for AI compute (Fourier)

Table 1. Top 10 robot AI processor suppliers.

Nvidia still supplies the compute that most of the field runs on. The company holds roughly 92% of AI training accelerators used in humanoid robotics worldwide, and Chinese firms depend on it for about 87% of pre-training compute. Jetson Thor modules, rated at 2,070 FP4 TFLOPS and drawing 130 W, now sit inside Unitree's H2 research humanoid, UBTech's platforms, Galbot's G1, Zhiyuan/AgiBot units, and EngineAI hardware, making Orin and Thor the default building block for most Chinese humanoid start-ups today. Cost drives the push toward domestic alternatives as strongly as export controls do: UBTech's chairman has said foreign chips make up roughly a third of hardware bill-of-materials cost. Core-component localization across Chinese humanoid programs now exceeds 70%, and a July 2026 US export rule further tightened the performance-density threshold for chip exports to China. For ISVs, that software stack matters as much as the silicon: Teams that build on Nvidia's CUDA tooling today gain a shorter path to production, and switching to a domestic accelerator later means re-validating models and inference pipelines on a different compute stack.

Domestic silicon is closing the gap fast, and ISVs and silicon teams tracking this market should watch a handful of names. D-Robotics, spun out of Horizon Robotics, packs its Xuri S600 chip with an integrated "brain and cerebellum" design rated at 560 TOPS INT8, and has signed more than 20 clients, including the Beijing Innovation Center's Tiangong 3.0 program, which targets mass delivery in the second half of 2026, plus Linglong 2.0, Tars' A3, Fourier, Booster Robotics, and Robotera. Huawei pairs its Ascend NPUs and Kunpeng CPUs with a CloudRobo software platform, and signed a comprehensive cooperation agreement with UBTech in May 2025 to build on that combination; UBTech's U1 emotional AI model trained on that same Ascend framework. Rockchip's RK3588 SoC now runs the emotional AI model inside UBTech's mass-market U1 consumer humanoid, priced from 119,800 to 990,000 yuan, and the company calls RK3588 the mainstream choice for edge compute in Chinese humanoid robots. Black Sesame Technologies pairs its Huashan A2000 "brain" chip with a Wudang C1236 "cerebellum" chip inside the Tianwen humanoid built with Wuhan University, rating the A2000 family at up to 1,000 TOPS. Cambricon and MetaX round out the domestic bench: Cambricon's Siyuan accelerators handle training and model adaptation work for a coalition of 10 domestic chipmakers targeting roughly 500,000 units shipped in 2026, and MetaX signed a joint venture with UBTech in June 2026 to design purpose-built embodied-AI edge chips, with sample production targeted for the second half of 2027. AgiBot has already shifted to domestic microcontrollers for its robot "cerebellums" and is evaluating additional Chinese chipmakers for its next-generation A3 platform.

Qualcomm entered this race later than its domestic and foreign rivals and already landed its first Chinese humanoid design win, pairing its Dragonwing IQ10 platform, rated up to 700 TOPS, with Stardust Intelligence's Atribot humanoid at the World AI Conference (WAIC) in 2026. AMD holds a narrower role in this field: Fourier Intelligence's GR-3 humanoid runs a Ryzen 7 8845HS chip purely as a general-purpose body controller, with a separate Nvidia module handling the AI compute alongside it. Among the domestic contenders, D-Robotics and Horizon Robotics rank as the fastest riser, moving from unveiling the Xuri S600 in November 2025 to more than 20 named clients and two mass-production humanoid programs within nine months, backed by a \$150 million funding round in April 2026. That pace gives CIOs evaluating Chinese robot vendors a fast-moving supplier list to track alongside each vendor's own production timeline.

The deployment figures underneath this chip race are stark. Global humanoid shipments reached roughly 13,000 units in 2025, and Chinese firms captured more than 84% of that volume, according to Omdia data cited by 36Kr. Unitree alone shipped over 5,500 humanoids in 2025 and now targets 10,000 to 20,000 units for 2026. AgiBot claims the top global spot for 2025, with 5,100 humanoid shipments and a 39% market share, and it has deployed more than 10,000 robots across multiple form factors as of mid-2026.

Trade data confirms the scale of this shift. China's robot exports hit ¥113.2 billion (\$15.7 billion) in the first quarter of 2026, and humanoid robot exports surged 210% year over year in that same period. Customs officials created a new tariff code, "intelligent bionic robots," to track the category separately for the first time. That move signals something larger: Humanoid robots have crossed from experimental hardware into a line item that customs authorities now track with the same rigor as autos or electronics. China's manufacturing base, supply chain depth, and willingness to ship robots into commercial pilots ahead of full reliability explain much of this lead, and volume, cost, and export momentum sit squarely with Chinese builders today.

A new assessment from DSET, a Taipei, Taiwan-based think tank, frames China's robotics push as a whole-of-nation strategy that borrows the playbook it used in electric vehicles and drones, the same approach that handed BYD and DJI roughly 70% global share in those markets. Nathanael Cheng, the report's lead researcher, told Nikkei Asia the strategy carries real potential to win the global physical AI race, and China has not reached that stage. Smart Analytics Global puts China's 2025 share of humanoid and robot-dog shipments at 85%, built on a large domestic market, foundation models, manufacturing scale, real-world data, and batteries. The report also finds real gaps: premium reducers, motors, machine tools, and sensors come mainly from Japan and Europe today, and memory chips and printed circuit boards could become the next bottleneck, since a single humanoid needs around 150 PCBs and the US PCB share has fallen from 30% to roughly 4% since the early 2000s. China's rare-earth magnets and battery materials hand it separate leverage over robot makers abroad. Washington answered in late July, blocking new import approvals for foreign-made humanoid, quadruped, and even robotic-vacuum models, and it grandfathered products already on the market.

What do we think?

China's humanoid robot volume now outpaces its chip independence. Domestic suppliers gain design wins fast; Nvidia still anchors most training and much on-device compute across the field. ISVs building for this market should plan for dual-sourced silicon, and CIOs evaluating Chinese robot vendors should ask directly which chips run inside the unit and how that dependency shifts under new export rules.

Inflection point. Humanoid robots mark an inflection point in AI, moving models out of data centers and onto machines that must sense, decide, and act in real time under strict power and latency limits. China's shipment volume, export growth, and new customs tracking show a market that has moved past prototypes into production. Silicon teams that solve on-device inference at this scale first will set the compute baseline for the next generation of embodied AI, and chip sourcing will decide who keeps up.