CHINA HUMANOID ROBOT

Asia Leaders Conference 2026 Takeaways: Panel Session: Full-Stack Embodied AI & Commercialization Roadmap

Asia Leaders Conference 2026
Aug 31-Sep 2, 2026
Explore >



Jacqueline Du
+852-2978-1785 ;
jacqueline.du@gs.com
Goldman Sachs (Asia) L.L.C.

At our Asia Leaders Conference 2026 in Hong Kong, we hosted a panel discussion with Robin Han, Co-Founder and CEO of Sudo, Cewu Lu, Co-Founder and CEO of NoeMatrix.ai, and Tianqi Luo, Co-Founder and CFO of Galaxea, focused on Full-Stack Embodied AI & Commercialization Roadmap. **Key discussion points:**

(1) Panelists broadly see embodied AI moving beyond pure research/demo excitement into early commercialization validation, **with the bottleneck shifting from isolated task success to whether reliability, generalization, and deployment economics can improve simultaneously.**

(2) **Data strategy should not be framed as a binary choice between real-world data or simulation; the key metric is "intelligence efficiency per dollar,"** or how much model capability can be generated per unit of effective data, compute, talent, and capital.

(3) **Panelists debated whether world models will dominate pure VLA or fuse with it to form World Action Models, but agreed that commercial success ultimately depends on an architecture's ability to translate understanding and reasoning into stable, reliable real-world manipulation.**

(4) **Initial commercialization is expected in logistics, select retail, and industrial sub-tasks, prioritizing narrower applications and the maturation of reusable skills, as successful deployment fundamentally relies on matching current technological capabilities in speed, precision, and generalization with specific commercial scenarios**.

Detailed Takeaways

- **Commercialization bottleneck: The current stage of embodied AI has moved past pure research/demo excitement into the early commercialization**

Goldman Sachs does and seeks to do business with companies covered in its research reports. As a result, 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. For Reg AC certification and other important disclosures, see the Disclosure Appendix, or go to www.gs.com/research/hedge.html. Analysts employed by non-US affiliates are not registered/qualified as research analysts with FINRA in the U.S.

phase, with the **key bottleneck shifting from isolated task success to whether reliability, generalization, and deployment economics can improve simultaneously**. Robin Han, Co-Founder and CEO of Sudo, emphasized that customers will not lower reliability thresholds simply because the task is hard: moving from 50%-60% success to 85% may be technically meaningful, but remains insufficient when commercial requirements are close to 100%, especially in select industrial use-cases. Cewu Lu, Co-Founder and CEO of NoeMatrix.ai, framed deployment cost as both an engineering optimization problem and a scientific problem. He noted that model-design know-how may still be under-optimized, with pi-style training hours (around 10,000 hours) suggesting the gap is not just brute-force scaling. Looking ahead, Galaxea's Tianqi Luo suggested the inflection point could arrive faster than expected if the data pathway is solved: once scaling begins, data accumulation could become exponential, for example collecting 1 million hours may take one year, while collecting 10 million hours may also be achievable within a similar time frame. The current bottleneck is more about data than an inherent ceiling in model capability; if data collection continues to improve, embodied AI's "intelligence level" could move from roughly "three-to-four years old" to "ten-plus years old" in less than five to ten years. Mr. Luo noted that because compute is the larger cost driver (vs. data itself), improving data quality and productivity is critical to avoid wasting compute.

- **Foundation/world model end-game: Panelists debated whether world models will dominate pure VLA or fuse with it to form World Action Models, but agreed that commercial success ultimately depends on an architecture's ability to translate understanding and reasoning into stable, reliable real-world manipulation.** NoeMatrix.ai's Lu argued that world models are meaningfully superior to pure VLA and will likely dominate in the near-to-medium term because they add stronger memory, spatial reasoning, analytical capability, and understanding of physical-state changes, though current architectures remain far from optimal for more complex robotic tasks. Galaxea's Luo framed the likely path as fusion rather than replacement: VLA maps vision/language to action, world models predict environment dynamics and action consequences, and World Action Models combine "imagining the future" with executable action generation. Sudo's Han emphasized the ultimate commercialization test: whichever architecture (VLA, WAM, or something else) can translate understanding and reasoning into stable, reliable manipulation in real-world settings will matter most.
- **Data strategy: Embodied AI data strategy should not be framed as a binary choice between real-world data or simulation; the more important question is which company can convert each unit of effective data into the steepest intelligence improvement, i.e. "intelligence efficiency per dollar."** Galaxea's Luo argued that simply labeling companies as "real-data" or "high-simulation" players is too overly simplified, as the real differentiation lies in the full system is collected, processed, evaluated, trained on, and closed-looped, a how much model capability is generated per unit of data, compute, capital. Mr. Han of Sudo highlighted that simulation must play a role real-robot data is constrained by collection speed, cost, and efficien

人在草木间

significant sim-2-real gap, the question is not whether simulation has scale advantages, but how effectively it can be used. Prof. Lu of NoeMatrix.ai similarly framed UMI, Ego, or any current data format as unlikely to be the final answer, and potentially even outdated in a few years; however, their near-term value is that they help build reusable infrastructure across data, software, training methods, distribution understanding, and evaluation, which should remain useful even if the winning data modality changes. In this sense, simulation and real-world data are better viewed as variables in a broader optimization function rather than mutually exclusive technology routes.

- **Initial commercialization for embodied AI is expected in logistics, select retail, and industrial sub-tasks, prioritizing narrower applications and the maturation of reusable skills, as successful deployment fundamentally relies on matching current technological capabilities in speed, precision, and generalization with specific commercial scenarios**. Within industrial, however, the panelists noted the earliest opportunities are likely to be narrower and more bounded tasks, while use cases requiring 99%+ reliability are likely to come later given the higher bar for model capability, skill reliability, and robot hardware performance. The more relevant lens is not just "which vertical comes first," but which reusable skills mature first — e.g. sorting, packing, folding, and flexible-object handling — because once one model-driven workflow starts to run, follow-on scenarios may scale faster through recombination of the same underlying capabilities. From a commercial logic perspective, customer scenarios typically demand three key attributes: speed, precision, and generalization. These three requirements combine to form an infinite number of matrices. Crucially, there will always be specific scenarios where the current technological capabilities align perfectly with the required speed, precision, and generalization. These are precisely the scenarios that should be prioritized for initial commercialization. Therefore, successful deployment is ultimately the result of the effective intersection between technological capability and commercial application.