Morgan Stanley | RESEARCH

IDEA

September 16, 2026 09:00 PM GMT

Robotics | North America

RoboFest California Preview: Bots by the Bay

Adam Jonas, CFA
Equity Analyst
Adam.Jonas@morganstanley.com +1 212 761-1726

William Tackett, CFA
Research Associate
William.Tackett@morganstanley.com +1 212 761-6028

Devin Roser
Research Associate
Devin.Roser@morganstanley.com +1 212 296-4539

While bipedal humanoids doing complex work may still be a few years away, robotics-focused AI labs have shown some real progress over the last few months across multiple different model approaches. We believe we are nearing a point where the manufacturing & integration is the limiter to adoption.

PrivateCompany

Private companies highlighted: Generalist AI, Physical Intelligence, Genesis AI, Applied Intuition, Skild AI, AMI Labs, 1X, Figure AI, Dyna Robotics

On October 12-13 ahead of our annual private company Spark Conference (Laguna Beach, CA), we are visiting Generalist AI, Physical Intelligence, Genesis AI, Applied Intuition, and others in the SF Bay Area as part of our 'RoboFest' Bus Tour. Space for RoboFest CA is extremely limited - if you miss this one we hope you can join one of our future RoboFests in the months ahead (Austin, Seattle, Zurich next on the calendar).

What's New & What's Next in Robotics AI Development?

Robotics AI labs starting to generate real revenue. Startup Skild AI announced it had hit $100mn in ARR this month, just ten months after its first commercial deployment (revenue was reportedly ~$30mn in 2025). Customer count also reportedly rose from 8 in Jan to more than 60, with hundreds of robots now operating across warehouses and factories. Of course, one company doesn't establish a trend, but this is the first revenue figure of real size anyone in the category has publicly disclosed to our knowledge.

Multiple labs claiming 'one-shot' learning. Previously, post-training a model for a new task could take thousands of demos over weeks. However, recently, Generalist's GEN-1.5 reported a 59% first-try success after a demonstration lasting 3-12 seconds, with the core model itself left untouched. Skild's S1, released in August, describes much the same result from a single human video.

World models are hot, VLAs less so. As of 12-18 months ago, VLA Models, or models that translate what a robot sees (vision) and what it is told (language) into motor commands (action), were all the rage. However, VLAs struggle with data hunger and general brittleness in dynamic settings. For example, a model could work well on what resembles its training data, but quickly degrade when the environment changes, resulting in a new task or environment needing fresh data collection. This year, a number of players including NVIDIA, Yann LeCun's AMI Labs, Physical Intelligence, Genesis AI, and others have brought attention to new

investors should be aware that the firm may have a conflict of

Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

foundation models that leverage "world models" or "world action models" that envision how a scene could change during a task, and choose the right action with the best envisioned outcome.

Data strategies differ sharply, and more than one seems to be working. Many companies are showing progress despite having near opposite approaches to robot learning. For example, Skild AI's model leans hard on simulation and learning from 3rd party videos (highly available, but lower fidelity). Meanwhile, Generalist's model leans heavily on hundreds of thousands of real-world interaction data. Both claim to have achieved some level of one-shot learning.

Locomotion and manipulation folded into one whole-body model. For humanoids specifically, many robots have generally run separate models for each. However, this inadvertently confined them to work done standing in place. Lifting something shifts the center of gravity/balance, so keeping the two systems apart would cap what the robot can attempt. Models from companies such as Google DeepMind, 1X, and Figure AI have recently combined both into one model.

Early evidence suggests robot models may follow LLM scaling laws. We would argue that the rate at which model performance rises as training data grows is critical for the path of robotics penetration, especially as robots begin being deployed and collecting data in large numbers in the real world. NVIDIA's EgoScale paper, published in February, offered arguably the first reasonably solid evidence, finding a log-linear relationship between training data volume and model performance across 20,854 hours of human video. Dyna Robotics reported a similar pattern in August (DYNA-2) at roughly 50 times the scale, saying performance improved smoothly from 1,000 to 1,000,000 hours of training video with no sign of plateauing.

What's standing in the way of physical-world RSI? In our view, primarily manufacturing + systems integration. The LLM takeoff happened after 2 primary phenomena occurred: 1) models reached a capability threshold where they could be genuinely useful; 2) millions of people adopted it in their daily lives creating the necessary flywheel where ROI attracted additional investment and data flywheels spurred ongoing model development. We would argue robotics models today are now capable of being genuinely useful in many (albeit highly structured) settings, but there are still barriers holding back mass adoption preventing the same flywheel as with LLMs. To be specific, manufacturing scale remains relatively nascent and capable robots cost six figures (or more). ***Simply put, there are not enough 'data probes' and the probes that do exist are prohibitively expensive for most risk-averse and capital-constrained physical world customers.*** Real-world robot data is often perceived as the least scalable and most expensive today, but could quickly become the predominant training source as robots become deployed in much greater numbers across diverse settings.

We recently published on how the US could reshore actuators, which can represent >50% of the robotics Bill-of-Materials: Robotics: Actuators: Robot Muscle Needs American Hustle (13 Sep 2026).

For our thoughts on how physical AI/robotics can help rebuild the US manufacturing base and address AI's 'PR problem' please see our report Prometheus Shrugged.

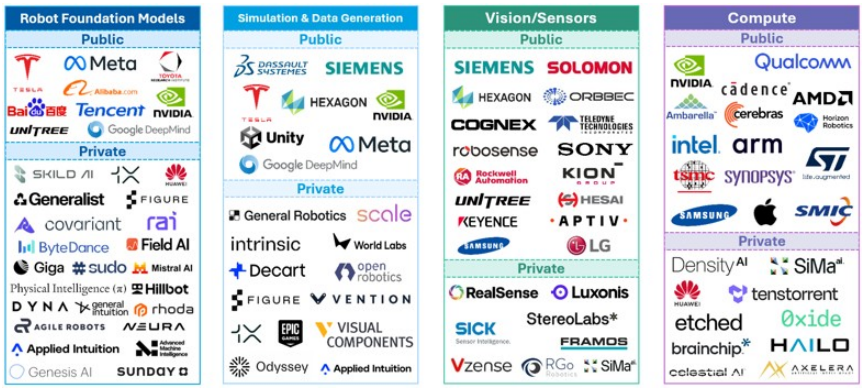


The content addressing private companies is being provided for informational purposes only and does not constitute a solicitation or imply future research coverage if the company goes public. Content is based on unaudited information. No investment recommendation is provided as there is limited public information available for private companies. Investors should conduct their own due diligence and be aware that additional or different information may be provided by the company.

Exhibit 1: Notable Robot Foundation Model Companies by Primary Training Data Source (Not all Inclusive)



Exhibit 2: Notable Robotics AI Enablers: Public & Private



Note: May not capture 100% of players. Does not include autonomous driving-specific companies (see separate section).

For Further Reading:

- Robotics: Actuators: Robot Muscle Needs American Hustle (13 Sep 2026)
- SpaceX: 20 Topics in 20 Minutes (20 Aug 2026)
- Robotics: Humanoid Horizons: Do Robots Have a PR Problem? (28 Jul 2026)
-
- SpaceX: The Rowdy Robot: What's the Catch? Bots on the Bayou (31 Aug 2026)
- China Industrials: Robotics: US Protectionism - Competition Rising (3 Aug 2026)
- Robotics: Thoughts on US Robot Protectionism (30 Jul 2026)
- SpaceX: Prometheus Shrugged (11 Sep 2026)
- SpaceX: AI's Final Frontier; Initiate at Overweight, PT $300 (7 Jul 2026)