Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

FOUNDATION

October 25, 2026 11:43 PM GMT

Robotics | North America

Generalist AI: Can Human Hands Teach a Robot Brain?

Billions in VC capital is pouring into robot foundation models. While peers lean on simulation or internet video, Generalist is venturing that large volumes of real-world human manipulation data, captured with handheld devices, is the path to general-purpose robots.

PrivateCompany

While many are still early-stage, we believe public and private investors should follow the progress of robot foundation model companies as a gauge of how close physical AI is to commercial reality. Generalist is one such company worth watching.

- Robots have been around for decades. The first industrial robot was deployed ~1.5 hours drive from the Morgan Stanley HQ at GM's (now-closed) Fisher Guide Plant outside of Trenton, NJ.
- However, robots were historically engineered for narrow tasks in tightly controlled environments. General-purpose robots that could deliver real-world ROI across a wide range of jobs were a futuristic concept.
- The central challenge with building "Robot Foundation Models", or models that control a robot from high-level reasoning to low-level motor control, has been the lack of training data. Unlike LLMs, there is no internet of physical data that can be readily scraped to build a general purpose robot.
- As a result, each foundation model company has picked a specific data strategy that they believe will build them a sustainable moat and first-mover advantage if their chosen approach is correct.
- The "correct" approach is an open research question. Multiple companies now lean on simulated/generated data (think training in a video game) or internet video data (watching videos of humans doing tasks and then replicating in the physical world) for pretraining due to the cost of collecting data from the real-world.
- **Generalist is taking a different approach, instead venturing that high-fidelity real-world interaction data should transfer most directly to physical manipulation** because it captures real dynamics and avoids the classic sim-to-real gap.
- The common pushback to leveraging physical interaction data is that it's not scalable. Traditionally, robots were trained through "tele-operation" (manually controlling a robot many times until it can repeat the task autonomously), which is expensive, cumbersome, and slow.
- Generalist has developed a handheld device it calls "data hands", hand mounted pincer grippers that capture vision and micro movement data based on an improved version of the Universal Manipulation Interface (UMI) - a

MORGAN STANLEY & CO. LLC

Adam Jonas, CFA
Equity Analyst
Adam.Jonas@morganstanley.com

William Tackett, CFA
Research Associate
William.Tackett@morganstanley.com

Devin Roser
Research Associate
Devin.Roser2@morganstanley.com +1 212 296-4539

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley 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.

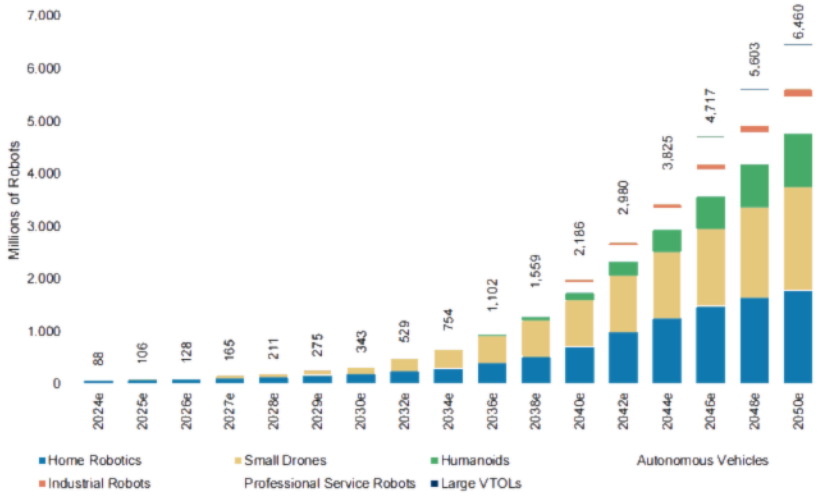
2024 research paper led by Stanford/Columbia/Toyota Research Institute researchers.

- *For more context, we highly recommend NVIDIA Researcher Jim Fan's talk at Sequoia AI Ascent discussing the "death" of tele-operation to train robots, and rise of wearables and human-captured ("egocentric") video.*

- Since its founding in 2024, the company has shown material progress with each of its 3 model releases (GEN-0, GEN-1, and GEN-1.5), growing its dataset from roughly 270,000 to more than 500,000 hours, lifting reported average task success from about 64% to 99% on select tasks, and most recently demonstrating emergent one-shot learning of new tasks from a single short demonstration.

TLDR on Generalist AI Robot Foundation Models: Most robots today are programmed for one narrow job. Which means teaching them a new task can take days of expert setup, demonstrations, and training. A robot foundation model is a single AI model trained on large amounts of physical data allowing it to control robots that 'generalize' across many tasks, much as LLMs handle many kinds of text-based tasks. Generalist AI collects that data with handheld "data hands" that let people perform ordinary tasks in homes, warehouses, and workplaces, with no robot required, and the company says its dataset exceeds 500,000 hours. With a model trained at this scale, the company has claimed it can learn certain new tasks with only a few seconds of demonstration.

Exhibit 1: Morgan Stanley Global Robot Model (Reach Out for a Copy): Projecting 6.5bn Robots Globally by 2050



Source: Morgan Stanley Research estimates

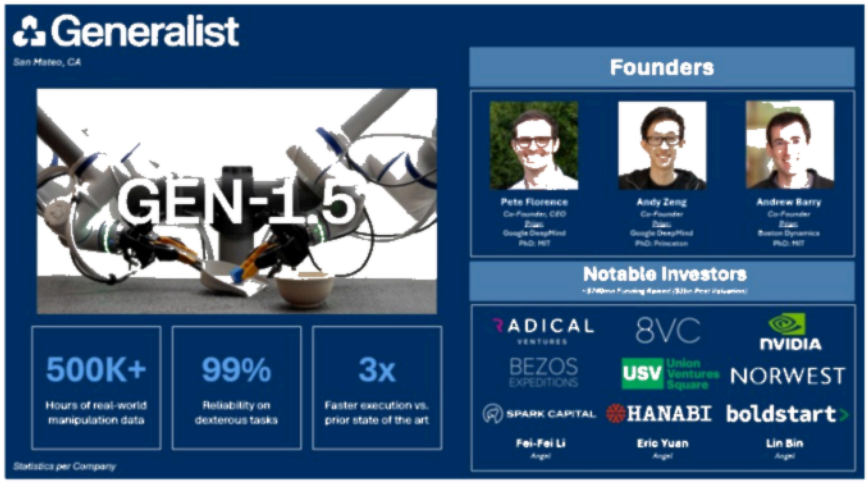
What Made Generalist Interesting To Us?

1. **Claims of one shot learning.** GEN-1.5 claims to learn new manipulation tasks from a single 3-12 second demonstration with no training step, averaging 59% success across 10 everyday tasks, including unzipping a pencil

pouch, twisting the lid off a glass jar, folding and creasing paper, and retrieving money from a purse. While only on simple, short-horizon tasks thus far, it appears to be one step in seriously cutting the cost of teaching robots new jobs and eliminating days of data collection.

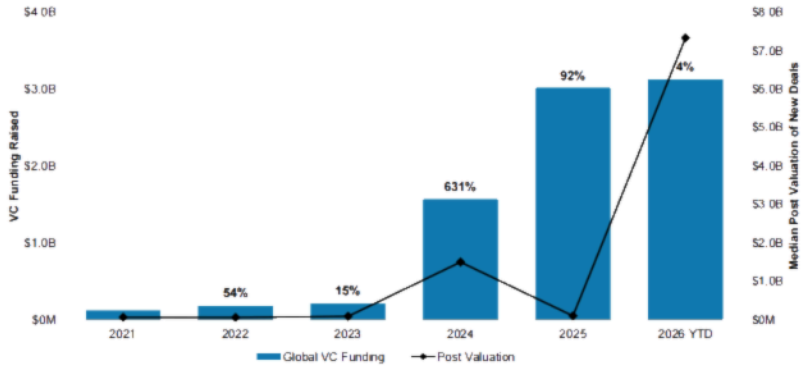
2. **Claims GEN-1.5 may follow LLM scaling laws.** Generalist reports that its models improve predictably as pretraining data grows, similar to LLMs. If true, it could have wide implications for robotics, since it would allow companies to forecast returns on their collected data, better optimize compute, and would shift the competitive question toward who can collect the most diverse data. Whether this dynamic is only true for models trained on physical interaction data is an open question to be proved out.
3. **Opposite approach to others means results have an interesting read-across to industry.** Many peers lean on simulation (Genesis, NVIDIA), internet video (Rhoda, Skild), or smaller teleoperated datasets, while Generalist is venturing that very large volumes of real-world human interaction data are the key to a general-purpose model. This makes Generalist, at the minimum, a useful test case. If the company's approach works, real world data collection at scale could prove to be a serious moat for companies, and if it fails, the field's current bootstrapping approaches look more correct.

Exhibit 2: Generalist AI Company Profile



Source: Generalist AI, PitchBook, Morgan Stanley Research

Exhibit 3: VC Funding for Robotics Foundation Models (PitchBook)



Source PitchBook (as of 10/7/2026), Morgan Stanley Research

For Further Reading:

- Physical AI: The Morgan Stanley Robot Almanac (7 Dec 2025)
- Space: Starcloud: Data Centers in Space - AI Goes Orbital (27 Jan 2026)
- Video: Autonomy at Altitude: Joby Aviation (22 Jan 2026)
- Defense Tech: Anduril: a NeoPrime(r) (12 Jan 2026)
- Embodied AI: Saronic: Robots Ahoy! (We're Gonna Need a Better Boat) (2 Nov 2025)
- Robotics: I Just Pre-Ordered a Humanoid Robot (4 Nov 2025)
- Embodied AI: Humanoid Disruptors: Unitree - Humanoids Starting at $16K (28 Jan 2025)
- Embodied AI: AI Robotics Disruptors: Physical Intelligence (19 Dec 2024)

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. 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.

Thoughts on the Current State of Robotics AI

What is holding back an AI robotics take-off? *We'd divide it into 2 areas: 1) Model Capability; 2) Physical World Complexities.*

Model Capability Hurdles

Lack of training data (core problem, now generally understood): By far the most pointed to problem with building robotics foundation models, although for good reason. To put it simply, there is no internet scale of data to train robot models, so companies need to collect or generate their own. There are 3 broad areas companies have focused on: real-world collected data, synthetic/simulated data, and existing video. Some companies have taken an explicit approach on one end of the spectrum: such as Generalist with real-world physical data or Skild AI with internet-scale video. Some have taken a broader approach, such as Physical Intelligence or Google DeepMind.

- Because no one yet knows which mix of real, simulated, and video data actually works best, the field hasn't been able to establish the reliable scaling laws that enabled the frontier LLMs of today, nor has robotics even proven to be a field that observes LLM-like scaling laws (though Generalist has shown promising results, *see below*).

Evaluation: How do we rank models? With LLMs, companies can run a model on a fixed dataset in a closed environment and obtain a result anyone can replicate. With robotics, every company has different hardware in different environments performing different tasks. Even the same robot arm could have slightly different actuator wear that results in different results. As a result, there's no clean way to rank models against each other to tell what is or is not working. There are simulated benchmarks, but the "sim-to-real" gap, especially with contact-heavy tasks, is very real even today. When progress can't be measured cleanly, it's challenging for researchers, financiers, and adopters to efficiently direct their effort and capital to the best solution.

Cross embodiment problem: Robots differ enormously in their bodies, including arm geometry, number of joints, grippers, sensors, and camera placements, so data collected on one robot doesn't transfer cleanly to another. This fragments an already scarce data pool into many small silos. While some labs have shown promising results across embodiments, it remains unclear whether a single model can truly generalize across bodies for longer-horizon, complex tasks.

How to address counterfactuals: Real World RL, World Models, something else? Imitation learning only shows a model what a demonstrator did, not what would have happened had it acted differently. A "counterfactual" is knowing what would have happened had a robot taken an alternative approach, which for a robotics AI model, allows the robot to recover from mistakes and handle unfamiliar situations. One camp argues for real-world reinforcement learning (RL), where robots learn from actual consequences of their own actions, which is grounded in true physics but slow and potentially unsafe at scale (and may also need many thousands of deployed robots to work reliably). The other camp argues for "world models" that imagine the range of possible outcomes so a policy

can imagine and evaluate alternatives (also called model-based reasoning), although their predictions can drift from reality, especially for contact-rich tasks, and they are only as good as the data they were trained on.

- *For more on World Models, see our primer from this spring: Embodied AI: World Models: AI's Journey from Digital to Physical (22 Mar 2026)*

Exhibit 4: Notable Robot Foundation Model Companies by Primary Training Data Source (Not all Inclusive)



Source: Company Data, Morgan Stanley Research

Exhibit 5: Robotics Data Pyramid for Training Models. *Generalist's approach is that real-world physical data, collected at scale, will matter more than synthetic or internet-scale data, and that over time the pyramid could invert as deployed robots generate far more useful data than any other source.*



Source: NVIDIA

Physical World Complexities

Often overlooked today. We don't think robotics should be compared to LLMs where once the model "works", rapid adoption follows. The physical world is an order of magnitude messier and more complex than the digital world.

Unforgiving nature of physical world. Law of 9's. With LLMs, hallucinations are a frustration, but don't create serious safety or operational issues like with robotics. In manufacturing specifically, modern supply chains can be incredibly tight, a robotic system

that goes down at one step does can put the broader chain at risk.

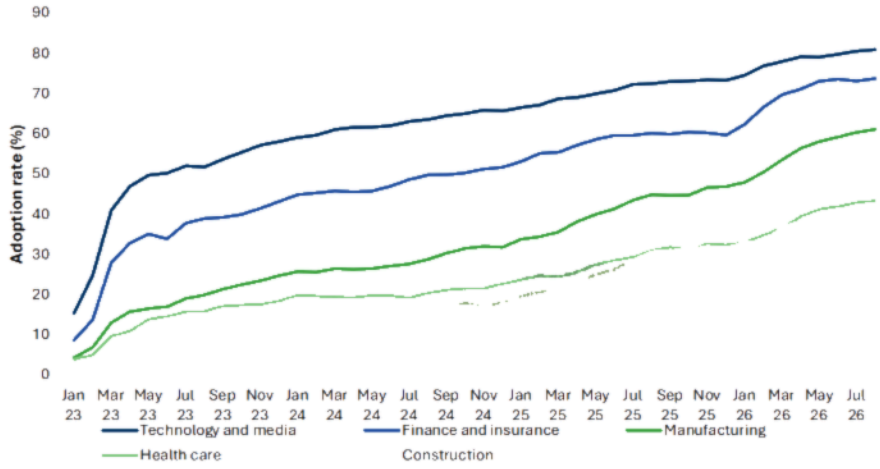
Conservative customer set. Selling robotics is not like an AI lab selling a new LLM tool to the SaaS company across the street. There can be major cultural differences between a robotics lab and a physical world company, and many companies frankly don't understand how to deploy and operate robotic systems.

Lack of scale manufacturing outside of China. From our conversations, the large majority of the robotics BoM is concentrated in Asia, particularly China. Even Tesla, known for its vertical integration, sources many robotics components for Optimus from China today.

- *We recently published on how the US could reshore domestic actuator manufacturing:* Actuators: Robot Muscle Needs American Hustle (13 Sep 2026)

In our opinion, the ability for US robotics model companies to eventually build a sustainable revenue model will ultimately depend on: 1) The US's ability to make robots; 2) The willingness of the old-world economy to adopt drastically new technologies. We imagine many researchers would argue that once the models work to a certain extent, it would incentivize additional manufacturing capacity and willingness to adopt. We are sympathetic to that argument, but also believe it's more of a circular dynamic vs. sequential where AI informs the hardware, hardware informs the manufacturing, and data collected in the real-world leads to ongoing improvements in the AI. *We are not saying that Generalist or other model companies can't be very successful businesses, we're just making the case that the physical world is messy, and building a model that can do useful work does not immediately translate to steep revenue/ARR growth in the same way as LLMs.*

Exhibit 6: AI Adoption in Finance & Technology vs. Major Physical Economy Sectors (% of Businesses Spending on AI; Ramp AI Index)



Source: Ramp, Morgan Stanley Research

About Generalist AI

Founding Story: Generalist AI was founded in 2024 by Pete Florence and Andy Zeng, two former Google DeepMind robotics scientists, along with Andrew Barry, who used to work on robots at Boston Dynamics. Florence earned his PhD at MIT under robotics researcher Russ Tedrake, then joined DeepMind, where he helped create RT-2, one of the first "vision-language-action" (VLA) models, and PaLM-E, an embodied multimodal model. Zeng spent years at Google as well, contributing to SayCan, which paired large language models with physical robot skills, and other early work on grounding language in robot action. Together as a team, Florence brought research in robot manipulation and multimodal models, Zeng brought expertise in robot data and grounding language in action, and Barry brought hands-on experience with real machines and hardware product lifecycles. Barry claims the idea for the company began with a phone call from Florence, who felt it was time to scale up the robot-learning work he had been doing. In March 2024, Boldstart Ventures wrote the company its first check, and the company stayed quiet until Florence appeared at Nvidia's GTC conference in March 2025 as co-founder and CEO. The founders saw physical data as the biggest bottleneck, so they built a way to collect it with cheap handheld grippers and focused on tasks that companies already pay people to do. Their mission is to build robots that can generalize (hence the name "Generalist") across any task, with Florence describing the goal as bringing general intelligence into the physical world.

Exhibit 7: Generalist AI Founders Left to Right: Pete Florence, Andy Zeng, Andrew Barry



Source: Generalist AI

Data/Model Approach

Generalist's approach leans on real-world data, but at unprecedented scale: Teleoperation, where a human remotely controls a robot to demonstrate a task, has been the traditional way to train robot foundation models, and competitors such as Physical Intelligence, Google DeepMind, Dyna, and Skild AI all use it to some degree. However, every teleoperated hour requires a working robot and usually a trained operator. Generalist's claimed differentiator is that its collection runs with no robot in the loop, which lets it gather data from many people in many places at once. Per the company, it has produced a dataset of roughly 270,000 hours at GEN-0 (November 2025), growing by about 10,000 hours a week, and more than 500,000 hours by GEN-1 (April 2026). The data also comes from people performing real tasks in homes, warehouses, and workplaces, which gives it a breadth of objects, environments, and edge cases that is hard to reproduce in a lab. Generalist reports that it has developed a model that may obey traditional LLM scaling laws based on its research, meaning its performance improves predictably as pretraining data grows. Many demos are shown on a single DoF gripper, but the company claims it has tested the same architecture on robots with 6, 7, and 16+ degrees of freedom.

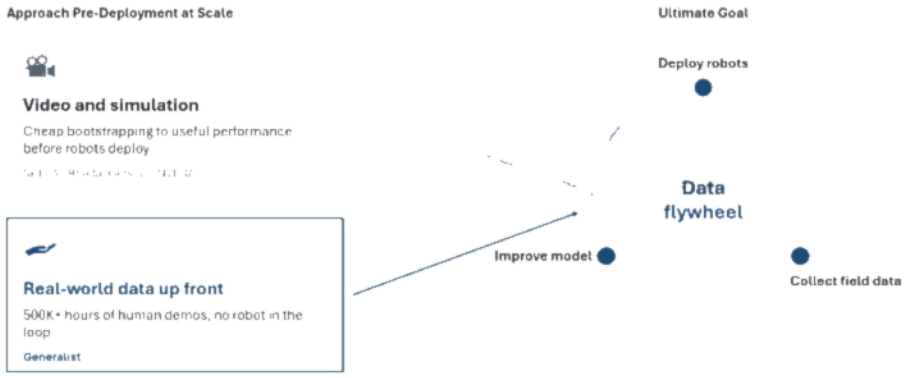
Generalist's "Data Hands": Generalist currently collects its training data with handheld devices it calls "data hands", which are handheld tools that turn a person's hands into pincer-like robot hands that capture visual and sensory data. From an interview with Forbes, co-founder Andy Zeng first thought of the idea years earlier while watching someone use a simple trash-grabber tool on a walk in Newport Beach. The devices are handheld grippers, a design that comes from academic research on the "Universal Manipulation Interface" or UMI, a framework from researchers at Stanford, Columbia, and the Toyota Research Institute built around a handheld 3D-printed gripper with a camera attached which lets people record robot demonstrations anywhere with no robot present. The company has reportedly deployed thousands of collection devices across multiple regions to perform real tasks in homes, warehouses, and factories.

Generalist's foundation model can support a broad range of end-effectors. In July 2026, Generalist reported that its GEN-1 policy now works with end-effectors ranging from five-finger anthropomorphic hands to specialized tools such as screwdrivers, tongs, and spatulas. The company says the demonstration data totaled about 80 hours at most, according to founding mechanical engineer Samantha Castellanos, which suggests the same model could run on a customer's existing grippers and tools, though the results are company-reported lab demonstrations.

Which data approach is correct? That is an open question today. However, there are real scenarios where Generalist's strategy could gain momentum as robotics scales over the coming years. It remains an open question for researchers today. We do not necessarily believe the answer has to be mutually exclusive. Generalist's strategy could gain momentum if its claimed scaling behavior continues to hold and human-hand data transfers well across robot bodies. Video and simulation-heavy approaches (Rhoda, Skild, Genesis, NVIDIA) could be seen as a form of bootstrapping: a cheap way to reach useful performance before deployed robots generate real-world data at scale. While others could gain similar data flywheels once they begin deploying at scale, Generalist's counter is that a larger, more diverse dataset up front produces a better model sooner and wins

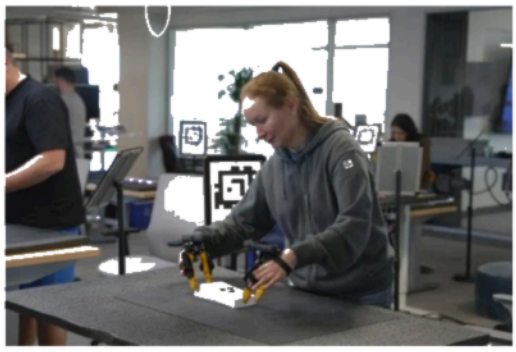
the deployments that start the flywheel.

Exhibit 8: Generalist AI is taking an approach that relies heavily on real-world collected data.



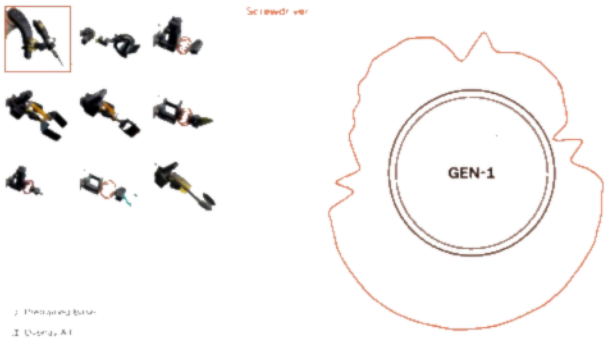
Source: Morgan Stanley Research

Exhibit 9: Generalist "Data Hands" Used by a Researcher to Collect Robotics Training Data



Source: Generalist AI

Exhibit 10: Despite being trained on its pincer-like "Data Hands", Generalist's model can adapt to a range of end-effectors with moderate changes to the model weights.



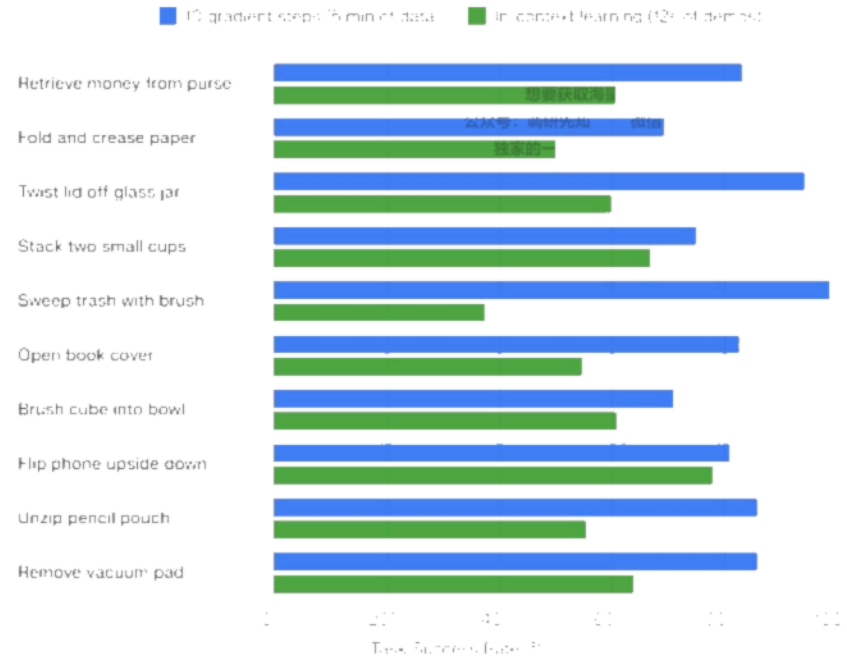
Source: Generalist AI

Model Releases

- **GEN-0 (Nov 2025):** GEN-0 was pretrained on more than 270,000 hours of real-world manipulation data, with the dataset growing by about 10,000 hours a week. Generalist presented it as the first demonstration of scaling laws in robotics, with the same architecture tested on 6 degrees-of-freedom (DoF), 7 DoF, and 16+ DoF semi-humanoid robots.
- **GEN-1 (Apr 2026):** GEN-1 was trained from scratch on a dataset of half a million hours, nearly double the data behind GEN-0. Average task success rose from 64% for previous models to 99%, tasks ran roughly three times faster, and in some tests it matched GEN-0's performance with ten times less task-specific data and fine-tuning steps.
- **GEN-1.5 (Aug 2026):** Where GEN-1 reached its 99% success rates after post-training on individual tasks, GEN-1.5 can perform a new task from a single 3 to 12 second demonstration with no training steps (in context learning), averaging 59% success across 10 tasks, a result Generalist attributes to over eight months of continuous pretraining. Ten gradient steps on five minutes of data lift can reportedly lift that figure to 83%.

GEN-1.5 Claims to Learn a New Task in Seconds, One-Shot: Generalist describes the capability as **"in-context learning"**, the same mechanism behind GPT-3's few-shot prompting, in which a sensorimotor demonstration recorded by a person with handheld grippers or by the robot itself is placed in the model's 30-second context window and the robot performs the task immediately. The company says it engineered none of this, with no architectural changes, meta-learning loop, or special training objective, and that the ability emerged from continued pretraining at scale. It also reports that two prompts can be chained into one longer behavior, that a demonstration recorded in simulation can prompt a real robot even though pretraining contained no simulation data, and that the model improvises with unfamiliar tools. However, the company also notes the tasks are simple and short-horizon and in-context skills are more brittle than fine-tuned ones.

Exhibit 11: 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.



Source: Generalist AI

Go-To-Market Strategy

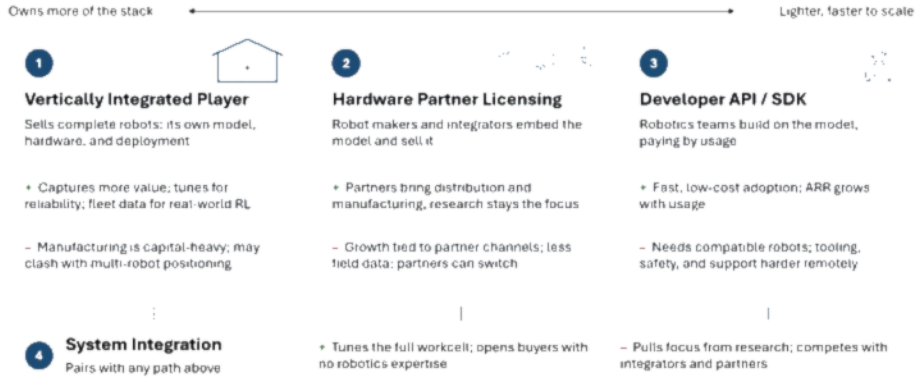
Generalist has described its go-to-market approach as still taking shape. Its current routes to end customers include hardware partners such as Universal Robots, which names Generalist a "UR preferred model partner", and an early access program in which a handful of customers help tailor the model for specific use cases. Delivery options so far range from an all-in-one package of model + hardware to running the model on arms a customer already owns. *Note that the company has not disclosed named customers, deployment counts, or revenue.* Generalist AI appears to be in a similar position to other robotics foundation model companies, such as Physical Intelligence, who have described the current focus as research-first and productization & go-to-market later since focusing on a specific go-to-market path too early could restrict the company's research efforts (over-focusing on a narrow customer set) as it competes against other highly advanced robotics labs to develop general purpose robotics models. ***We describe 4 potential GTM approaches we could envision the company taking going forward, none of which are mutually exclusive:***

- **Becoming a Vertically Integrated Robotics Player:** Generalist would build and sell complete robotic systems, combining its own models with robot hardware and handling deployment end to end. It would move from supplying intelligence to owning the full product that customers buy. Competitors such as Dyna and Genesis AI have recently pivoted to taking this approach, unveiling their own robots.
  - *Pros:* Allows tuning for reliability and captures more of the value created, plus direct access to data captured from its own fleet to improve future model iterations (real world RL).
  - *Cons:* Manufacturing is hard. Building robots demands far more capital and operational complexity, and it may conflict with the company's positioning across many robot types.
- **Licensing through Hardware Partners:** Generalist would license its models to robot makers and integrators, who embed them in their own products and sell to end customers.
  - *Pros:* Partners provide distribution and manufacturing, so Generalist can stay focused on research while scaling without a large services team.
  - *Cons:* Generalist's revenue growth would be tied to its partners' channels. Generalist would also have more limited visibility into field use and data capture, and partners could always switch to competing models.
- **Developer Access Through an API or SDK:** Generalist would give robotics teams and integrators direct access to its models, likely with usage-based pricing, so they can build their own applications on compatible robots.
  - *Pros:* Self-serve access could spread adoption quickly and at low cost, with ARR that grows with usage.
  - *Cons:* Customers still need compatible robots, and reliable tooling, safety guardrails, and support are harder to provide at a distance.
- **System Integration (Could be Paired with Any of Above):** Generalist could help design and install complete robotic workcells for customers, handling the arms, tooling, cameras, safety systems, and installation around its model. Most robotics deployments today rely on some system integration partner to build and optimize

their robotic systems. With a potential proliferation of various robotic form factors with AI, we could see it becoming increasingly challenging for physical world companies to most efficiently integrate these robots into their operations, creating a real opening for companies that best understand how these systems operate - such as the robot foundation model players.

- *Pros:* Generalist knows how its model behaves better than any third party, so it can tune the whole workcell for performance and reliability. Turnkey deployments also remove a major hurdle for customers lacking robotics expertise- expanding the penetrable TAM.
- *Cons:* Could pull focus from model research and put Generalist in competition with established integrators and, potentially, its own hardware partners.

Exhibit 12: Potential Go-to-Market Approaches for Generalist AI



Source: Morgan Stanley Research

Key Risks

Generalist's concentrated play on one data approach could prove to be the wrong technical path. We see risk in its reliance on human-worn-device data, which may not transfer cleanly across robot bodies, and its scaling and 99% success claims are self-reported on simple, short-horizon tasks with little independent benchmarking. Competitors pursuing teleoperation, simulation, video pretraining or hybrids could erode its data advantage, and pivoting would be costly given its proprietary collection operation.

Robotics foundation model companies do not yet have a proven business model, and the economics and go-to-market path remain uncertain. We think it's not fully clear whether value will accrue to the model layer, hardware makers or integrators, and hardware partners may prefer to build their own models, forcing the foundation model companies to either serve a set of relatively weak hardware-only customers or pivot to their own in-house hardware. Customers in manufacturing and logistics bring long sales cycles and high reliability demands, while data collection and training are capital-intensive and pricing is undefined.

Standalone robotics companies without in-house use cases could struggle to find customers who can bring their models to the scale a competitive data flywheel requires. Physical world companies, especially in manufacturing, are often inflexible, slow-to-adopt, and operate on thin margins with already high capital intensity. Like AVs, we believe robotics adoption is a law of nines. In a manufacturing operation, a robot that breaks can shut down a line and freeze an entire supply chain. Convincing many old-world companies to risk their operations for an unproven technology may be more challenging than many companies or investors appreciate. Humans, at the end of the day, are really good biological machines. Even Elon Musk, known for his highly automated manufacturing, is known have his workers to 'rip' misbehaving robots off a factory line when humans could perform a task more effectively.

Known Hardware Partners & Collaborations

Universal Robots (Owned by Teradyne - Covered by Shane Brett): UR's March 2026 press release calls Generalist a "UR preferred model partner". Two UR7e robots at UR's GTC booth ran a Generalist model on a smartphone packaging task, which the release describes as the first public demonstration of Generalist's models.

Flexiv: At the Automate trade show in June 2026, Generalist used Flexiv arms to repair robot vacuums in a live demonstration.

Notable Robot Foundation Model Competitors

- **Physical Intelligence (Private):** Pi builds general-purpose robot models that run on other companies' hardware. Founded in 2024 in San Francisco by former Google Brain and DeepMind researchers, and is led by CEO Karol Hausman. Its last announced round was a $600 million Series B (November 2025) at a $5.6 billion valuation, and PitchBook shows a Series C of about $1.0 billion (May 2026) at > $11 billion.
  - *Data approach:* Not taking a hard play with data strategy, unlike many others. Teleoperated demonstrations across many robot types, plus corrections and reinforcement learning from the robots' own attempts. Its newest model (April 2026) also uses human video, and researchers have also indicated the company is working with world models.
- **Genesis AI (Private):** Startup based in Paris and Silicon Valley, founded in December 2024 by Zhou Xian and Théophile Gervet. It recently unveiled Eno, its first general-purpose robot, after showing its GENE-26.5 model and a human-scale robotic hand in May 2026. Its last closed round was a $105 million seed (July 2025, valuation undisclosed), and Bloomberg reported in July 2026 that it is in talks to raise about $500 million at a $3 billion pre-money valuation.
  - *Data approach:* Synthetic data from its own physics simulator, which it says runs up to 430,000 times faster than real time, combined with real robot data.
- **Field AI (Private):** Sells hardware-agnostic autonomy software for robots in unstructured sites like construction and energy. Founded in January 2023 in Irvine, California, by NASA JPL veterans who led DARPA autonomy programs, its team also includes former DeepMind researchers. Its last announced round was $405 million across two rounds (August 2025) at a $2 billion valuation, and it is reportedly seeking $700 million at a $10 billion valuation per Business Insider.
  - *Data approach:* Multi-sensor data from paying customers' deployments, supplemented with synthetic data from NVIDIA's Isaac Lab.
- **Skild AI (Private):** Builds the "Skild Brain", a single model meant to control many robot types. Founded in 2023 in Pittsburgh by former Carnegie Mellon professors Deepak Pathak and Abhinav Gupta, who are also former Meta FAIR researchers. Its last round was a roughly $1.4 billion Series C (January 2026, led by SoftBank) at a valuation above $14 billion per PitchBook and the company.
  - *Data approach:* Internet human videos and physics simulation, plus teleoperation and deployment data. Its S1 model (August 2026) learns tasks from a single human video.

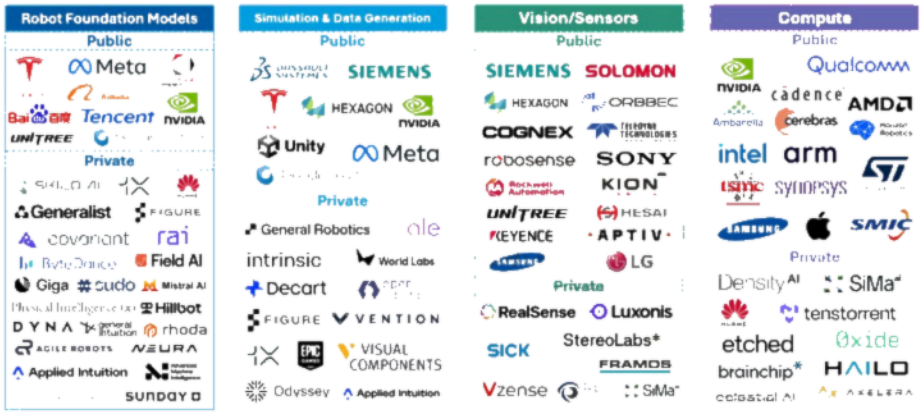
- **Rhoda AI (Private):** Rhoda targets industrial work with video-trained robot models. Founded in 2024 in Palo Alto by Jagdeep Singh, the founder of QuantumScape, together with Stanford researchers, it was incubated within Khosla Ventures. Its last round was a $450 million Series A (March 2026, led by Premji Invest) at a reported $1.7 billion post-money valuation.
  - *Data approach:* Pretrains on hundreds of millions of internet videos, then post-trains on small amounts of robot data, sometimes as little as ten hours of teleoperation.
- **Dyna (Private):** Redwood City startup with commercial deployments in hotels, restaurants, and laundromats. Founded in 2024 by Lindon Gao and York Yang, who sold their previous company Caper AI to Instacart for $350 million, together with former DeepMind researcher Jason Ma. Its last round, per PitchBook, was a $120 million Series A (September 2025) at a reported valuation above $600 million.
  - *Data approach:* Customer deployments feed a continuous reinforcement learning loop, alongside human demonstrations. Its DYNA-2 model is pretrained on over one million hours of human video.
- **Google DeepMind (GOOG - Covered by Brian Nowak):** Released Gemini Robotics 2 in July 2026 as an intelligence layer for other companies' robots, including Apptronik's humanoids. Also major world model developer through 'Genie' and 'Veo'.
  - *Data approach:* A mix of human teleoperation, video, and simulation. Claims it adapts to new robot bodies with typically fewer than 200 examples.
- **NVIDIA (NVDA - Covered by Joe Moore):** Offers open Isaac GR00T humanoid models and the simulation tools (Isaac Lab, etc.) which several competitors also use. It is also an investor in several competitors through NVentures, including Skild AI, Field AI, and Dyna (as well as Generalist AI).
  - *Data approach:* A "data pyramid" of web and human video, synthetic data from Omniverse and Cosmos, and real teleoperation data. But major focus on simulated data leveraging in-house Omniverse platform. NVIDIA reports a 40% gain from adding synthetic data to real data.

Exhibit 13: Notable Robot Foundation Model Companies by Primary Training Data Source (Not all Inclusive)



Source: Company Data, Morgan Stanley Research

Exhibit 14: Notable Robotics AI Enablers: Public & Private



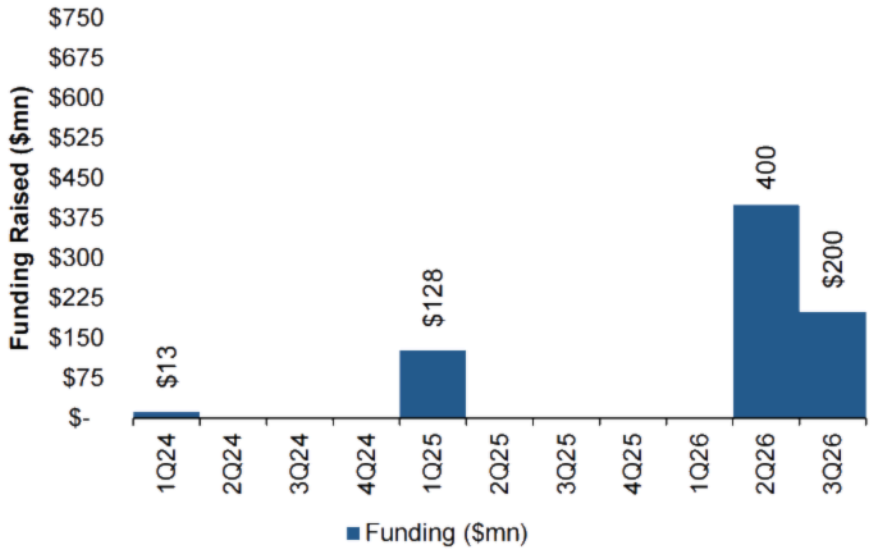
Note: May not capture 100% of players. Does not include autonomous driving-specific companies (see separate section).

Source: Morgan Stanley Research. Company classifications as of September 2025.

Investors & Funding

According to PitchBook, Generalist AI has raised >$740mn to-date. Its most recent round was a $600mn Series B & Series B Extension round at a $2.0 and $3.0bn post-money valuation led by Radical Ventures and 8VC with participation from NVIDIA, Spark Capital, Boldstart Ventures, NFDG, Union Square Ventures, Hanabi Capital Management, Scale Venture Partners, HOF Capital, and others. Notable individual investors include Fei-Fei Li (Notable AI Researcher & Founder World labs), Eric Yuan (Founder Zoom), Lin Bin (Co-Founder Xiaomi), and Jeff Bezos via Bezos Expeditions (Founder Amazon & Blue Origin).

Exhibit 15: Generalist AI Historical Funding Rounds (Sourced from PitchBook)



Source: PitchBook, Morgan Stanley Research

The valuation information provided here is included for illustrative purposes only. The information is based on publicly available information (funding rounds, company disclosure, third-party vendors, etc.) and was not produced or endorsed by Morgan Stanley Research.

Timeline of Major Developments

- **2024:** Founded by ex-DeepMind researchers Pete Florence and Andy Zeng with ex-Boston Dynamics engineer Andrew Barry, and kept in stealth.
- **January 2025:** Raised $128mn Series A at a $400mn post-money valuation led by NVentures and Spark Capital per PitchBook.
- **March 2025:** Florence appears on an Nvidia NVentures panel at GTC, the first hint of the company's goals and Nvidia backing.
- **November 2025:** GEN-0 launches, trained on 270,000+ hours of real-world data, and claims to show scaling laws in robotics.
- **April 2026:** GEN-1 release, trained on 500,000+ hours, claiming 99% success on simple tasks, about 3x faster than prior systems.
- **June 2026:** Raises a $400M Series B led by Radical Ventures at a $2B valuation.
- **August 2026:** GEN-1.5 release, learning new tasks in context from 3-12 second demos, with no gradient updates.
- **August 2026:** 8VC leads a roughly $200M Series B extension at a $3B valuation, bringing the round to about $600M.

Exhibit 16: Timeline of Major Developments



Source: Pitchbook, Company Announcements, Morgan Stanley Research