

## Transformation

# Physical AI, part 1: The basics

26 February 2026

### Key takeaways

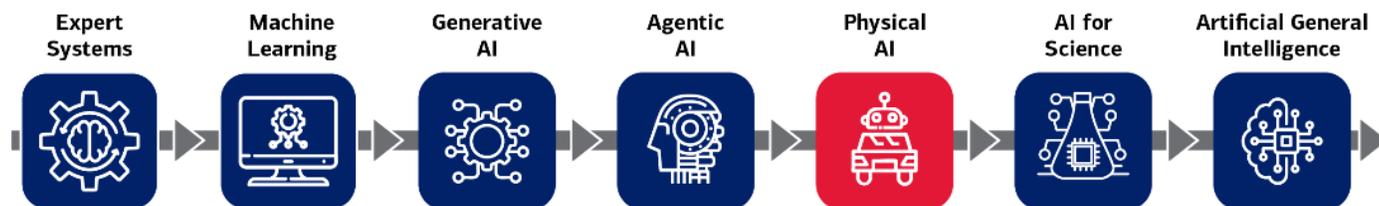
- Physical AI marks the next major phase of AI commercialization, extending intelligence beyond software and into machines that can see, decide and act in the real world - unlocking multi-trillion-dollar opportunities across robotics, autonomous vehicles and drones.
- Progress is accelerating due to converging advances in models, data, compute and simulation, with multimodal foundation models, synthetic data and world models enabling robots to learn, interact and train more safely and efficiently.
- With these advancements, we're seeing robotics shift from rules-based systems to data-driven, end-to-end AI, supported by open-source innovation and improved observability paving the way for adaptable, general-purpose robots and more autonomous industrial systems.
- This is the first publication in a three-part series that aims to explore what happens when AI leaves the chat, and shows up in the real world.

### What is physical AI?

Physical AI, or embodied AI, refers to the shift from systems that exist only in digital environments to machines that can sense, reason, and act in the physical world – such as humanoid robots, autonomous vehicles and drones. This evolution marks the next major phase of AI commercialization, pushing intelligence beyond software and digital output, like chatbots and content generation, and into machines capable of manipulating their surroundings and delivering real-world productivity gains (Exhibit 1). New to these terms? Read [AI Dictionary, part 1: The basics](#) for quick primers on core concepts.

**Exhibit 1: AI capabilities are evolving, with practical applications beginning to emerge in physical domains**

Infographic illustration of AI's evolution



Source: BofA Global Research

BANK OF AMERICA INSTITUTE

This transition, however, didn't happen overnight. As BofA Global Research notes, the rise of physical AI stems from cumulative technological progress, rather than a single breakthrough. Physical AI systems combine sensors and computer vision with AI models that learn through interaction. Rather than relying on explicit programming for every action, these systems learn and improve through trial, feedback and repeated experience. For instance, instead of hard-coding a robot to pick up an object, physical AI enables it to experiment, adjust and refine its approach until it masters the task – resulting in far greater adaptability and autonomy.

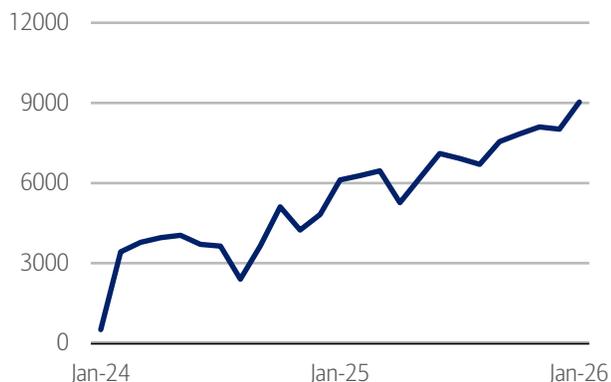
Early applications focus on automating repetitive or hazardous tasks in areas such as warehouse operations, manufacturing and logistics. As self-learning capabilities improve, physical AI is expected to evolve from task-based automation toward greater levels of autonomy – reshaping mobility, fulfilment, and industrial operations in the process.

And people are taking notice. Coverage, interest, and public awareness of physical AI have been growing fast across mainstream business and financial media, social platforms, and investors and corporates as they focus on robotics, autonomy, and AI monetization beyond software. BofA Global Research notes that news mentions and market commentary around physical AI

continue to rise, with January 2026 generating the highest number of citations in the past two years and almost nine times as many citations as January 2024 (Exhibit 2). Meanwhile, social engagement on X is also trending higher (Exhibit 3).

**Exhibit 2: Physical AI is increasingly cited in the news**

Number of physical AI citations across trade publications and news sources



Source: AlphaSense, BofA Global Research

BANK OF AMERICA INSTITUTE

**Exhibit 3: Physical AI social media (X) posts are trending up**

Number of physical AI social media posts per day on X



Source: ListenFirst, BofA Global Research

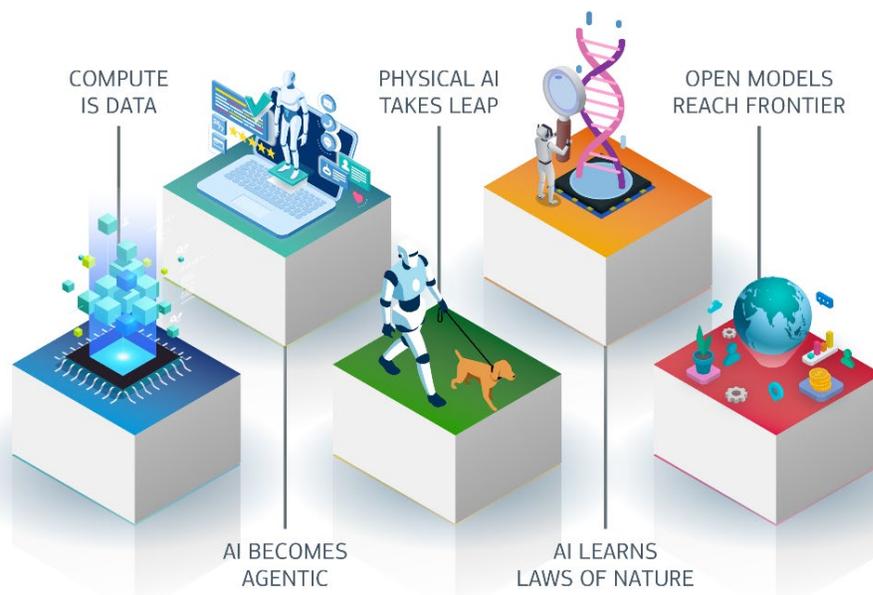
BANK OF AMERICA INSTITUTE

**Technology enablement: Scaling beyond LLMs**

Rapid cost reductions and performance improvements across AI models, chips, sensors, batteries, and manufacturing capacity are enabling the development of machines that can see, decide, and act in real time. Together, these advances are unlocking what could be multi-trillion-dollar opportunities across robotics, autonomous vehicles, drones and the broader AI infrastructure.

**Exhibit 4: A combination of better models, data and compute is enabling more sophisticated AI applications**

Infographic illustrating how AI is scaling beyond large language models (LLMs)



Source: BofA Global Research, based on NVIDIA

BANK OF AMERICA INSTITUTE

According to BofA Global Research, advancing physical AI will require fundamentally new model architectures, training methods, and compute strategies – approaches that go beyond the techniques that previously scaled LLMs. However, several converging trends are now accelerating the types of tasks AI can perform, helping explain why progress is spilling over from digital cognition into the physical AI world (Exhibit 4). Trends enabling the rise of physical AI include:<sup>1</sup>

<sup>1</sup> NVIDIA. (2026). *NVIDIA Live at CES 2026 Deck*.

- **Compute and data:** Advances in model architecture demand massive amounts of compute and data. Despite limitations, scaling laws continue to hold: models improve as data and compute increase. To supplement limited real-world interaction data, synthetic data generation is expanding rapidly, providing richer training signals for physical tasks.
- **Agentic AI:** New AI systems are becoming more autonomous, able to gather information, conduct research, plan and simulate outcomes on their own. These capabilities allow models to perform “inference-time compute” as it is required, often at the edge, enabling more adaptive and context-aware decision-making. For deeper reading, we covered the latest advances in [The new wave: Agentic AI](#).
- **Physical AI models:** Progress in AI systems capable of understanding and operating in the real world is accelerating, fueled by models that go beyond traditional LLMs. Vision-language-action (VLA) models, for example, integrate perception, reasoning, and motor control, creating the foundation for machines that understand their environment.
- **AI physics understanding:** For AI to operate safely and effectively in real environments, it must grasp the laws of physics and nature. Emerging “world models” are beginning to provide this understanding, allowing systems to predict how objects behave, anticipate outcomes, and interact reliably with the physical world.
- **Open-source models:** Open-source models are publicly and freely accessible, meaning that anyone can access and/or modify them, which allows for more transparency, collaboration and democratized access. This broad participation accelerates experimentation, lowers barriers to entry and speeds up the development of new techniques that feed directly into physical AI progress.

For a background on many of these terms and more, read [AI Dictionary, part 2: The next generation](#).

## Open-source innovation fuels rise of physical AI applications – and aspirations

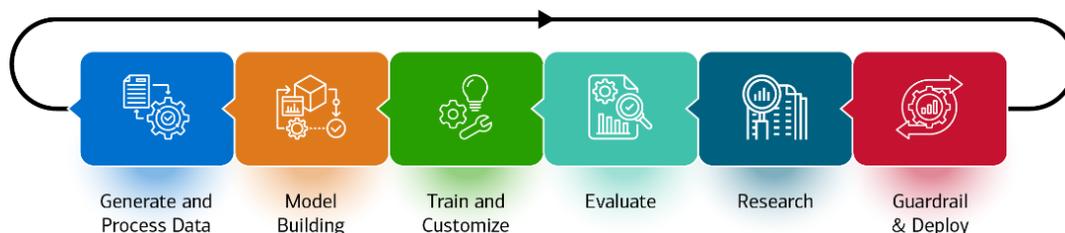
With physical AI, open-source model innovation is allowing for rapid advancements to be made, as robotics developers can access frontier-level reasoning, perception, and world-modeling capabilities without proprietary barriers and prohibitive development costs. By releasing weights, data, and training frameworks openly, the ecosystem can iterate faster, fine-tune for niche hardware, and leverage community-driven improvements – dramatically shortening development cycles for humanoids, autonomous systems, and industrial robots.

### Why open source?

In his keynote at the Consumer Electronics Show earlier this year, NVIDIA CEO Jensen Huang explained that open models accelerate every stage of the physical AI pipeline – from data generation and model building to custom training, evaluation, research, and deployment.<sup>2</sup> By making models, data, and evaluation tools open, developers can skip costly pre-training, customize faster, benchmark safely, and build transparent, auditable guardrails for real-world robotic systems (Exhibit 5).

#### Exhibit 5: Open-model ecosystem accelerates innovation – allowing for faster, safer, and customized physical AI

Infographic illustrating the impact of open-source model innovation on the physical AI pipeline



Source: BofA Global Research, based on NVIDIA

BANK OF AMERICA INSTITUTE

## Accelerating model capabilities power “robot brains”

Advances in AI models have pushed robot intelligence from fragile lab demos to fast, robust, and adaptable skills. Developments are emerging across vision-language models (VLMs), vision-language-action models (VLAs), and world models – together forming the foundational stack of physical AI.

- **Vision-language models (VLMs):** Provide robots with perception and scene understanding by combining visual inputs with natural-language grounding. They act as the base layer that later models build on.

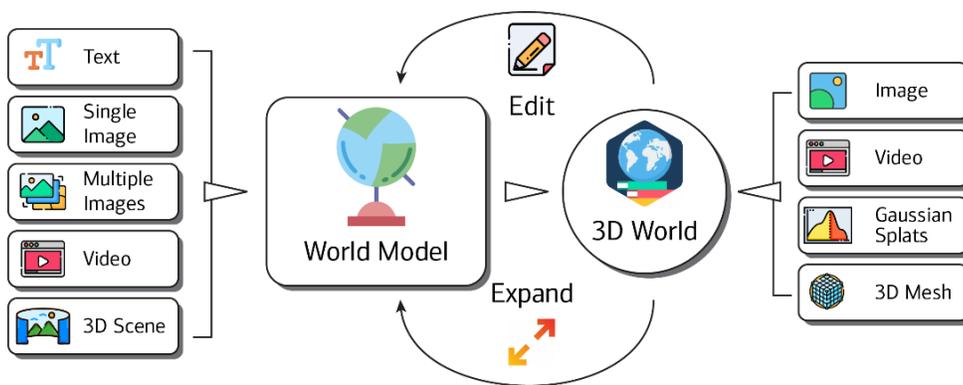
<sup>2</sup> NVIDIA. (2026). *NVIDIA Live at CES 2026 Deck*.

- **Vision-language-action models (VLAs):** Connect what robots see (vision), what they are told (language), and what they should do next (action). They translate multimodal inputs directly into behaviors, but often lack deeper predictive understanding of how the world works.
- **World models:** Give robots an internal simulation of reality – physics, object dynamics, cause-and-effect – allowing them to anticipate consequences, plan multi-step tasks, and recover from errors. They also help validate robotic systems in simulation, making training safer, cheaper and faster. These systems can unify perception, reasoning, and action in a single model to enable the intelligence and multi-purpose tasks that are starting to be demonstrated from robot developers (Exhibit 6).

While these model architectures emerged as separate research pathways, they are now converging as complementary components of physical AI – VLMs for perception, VLAs for instruction-following and action generation, and world models for reasoning, planning and robust performance under uncertainty. This combined intelligence could be key to enabling general-purpose robotics, bringing physical AI closer to scalable commercial reality.

**Exhibit 6: Similar to multimodal LLMs that combine multiple data sources, world models do so to simulate behaviors and outcomes in the real world, with 3D outputs**

Infographic depicting how world models combine data sources to simulate and predict outcomes



Source: BofA Global Research

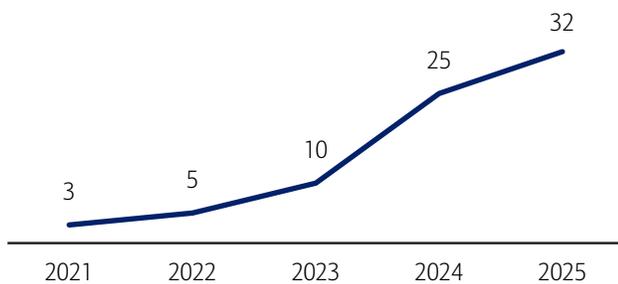
BANK OF AMERICA INSTITUTE

**Growing list of companies targeting world models for robotics**

World models require rich multimodal training data from real environments, with several companies targeting this vertical to provide the data and models necessary for robotics, including tech firms and privates alike. Investment in companies looking to develop such models is also increasing, with 32 equity deals with robot foundation model developers in 2025 compared to just three in 2021 (Exhibit 7).

**Exhibit 7: Funding deals to robot model developers grew 10x since 2021**

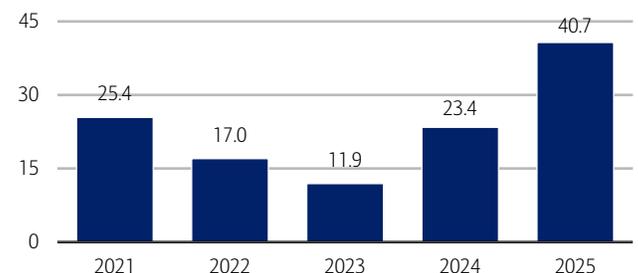
Equity deals for robot foundation model developers (yearly)



Source: CB Insights, Tech Trends 2026: 14 emerging trends to watch closely this year (https://www.cbinsights.com/research/report/top-tech-trends-2026/); BofA Global Research  
BANK OF AMERICA INSTITUTE

**Exhibit 8: Funding for robotics companies reached \$41 billion in 2025**

Annual funding to companies developing robotics hardware and software (yearly, \$ billions)



Source: CB Insights, State of Venture 2025 (https://www.cbinsights.com/research/report/venture-trends-2025/); BofA Global Research  
BANK OF AMERICA INSTITUTE

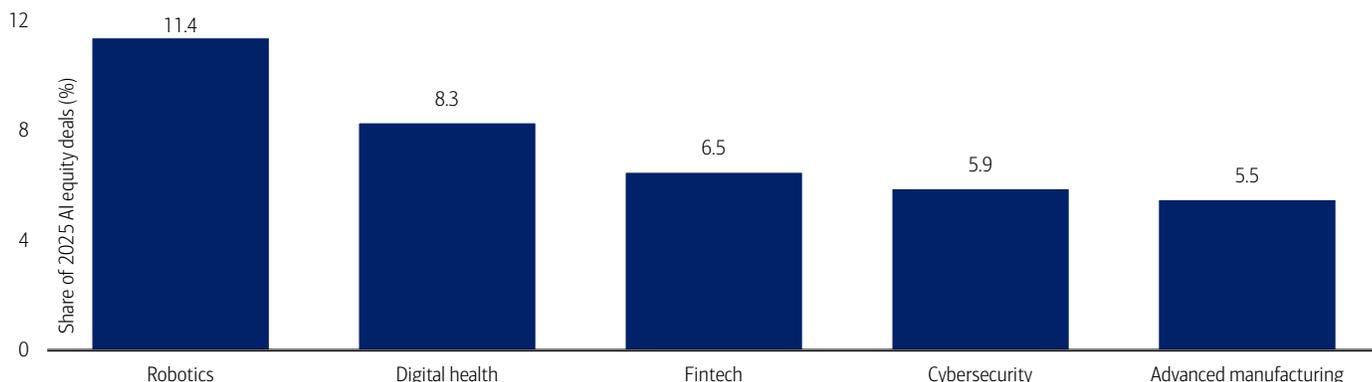
**\$41 billion investment in 2025 driven by physical AI**

With physical AI driving growth, funding and deals for AI robotics startups hit record highs in 2025 with the total reaching \$41 billion, accounting for 9% of all venture capital for humanoid robots, autonomous vehicles, drones, robot foundation models and

supporting hardware (Exhibit 8). Robotics hardware and software companies attracted the largest share of AI equity funding deals in 2025 (11.4%), overtaking digital health in 2023 and sustaining its share as overall AI deal activity expanded (Exhibit 9).

### Exhibit 9: Robotics received the highest share of 2025 AI equity trade deals

Share of 2025 AI equity deals (%)



Source: CB Insights, *State of AI 2025 Report* (<https://www.cbinsights.com/research/report/ai-trends-2025/>); BofA Global Research

BANK OF AMERICA INSTITUTE

### Foundation models: The final robot intelligence package

Foundation models are the end-product that companies deploy or sell in robotics applications. While VLMs, VLAs and world models are the technical ingredients, foundation models are fully trained robot brains built from this data and model architecture ready for real-world use, often applying additional layers for control, safety and memory. Per CB Insights, this enables developers to license and adapt this intelligence rather than build from scratch, across either robots, autonomous vehicles or robot co-ordination.<sup>3</sup> Many companies use multiple types of foundation models, given that training data from one domain can improve performance in others.

### Synthetic data supercharges simulation

Physical AI is shifting robotics from rules-based systems to data-driven learning. As a result, the race is now on to develop foundation models that work across tasks. But unlike text-based AI, which benefits from abundant digital data, physical AI systems depend on scarce real-world interaction data. As BofA Global Research notes, training robots requires vast amounts of real-world interaction data, yet collecting it is slow, costly and often risky.

Simulation is essential in bridging that gap. It lets robots fail safely, rehearse rare edge cases, and scale into new environments without real-world exposure first. Physical AI data stacks increasingly blend synthetic data, real-world demonstrations and advanced simulators.

Building on this need, recent progress in model simulation has emerged from foundation models that are becoming increasingly multimodal, able to interpret or generate text, images, video, and sensor streams. When combined with specialized models, software, and high-performance chips, they enable simulators that are far more realistic and useful for developing, training and optimizing robots and autonomous vehicles. In physical AI, simulation is particularly valuable for:

- **Humanoid robots:** Modeling forces, balance, contact and full-body dynamics so they can practice walking, lifting, grasping and interacting with unpredictable environments – tasks too risky, slow or costly to learn on real robots.
- **Autonomous vehicles:** Recreating complex driving scenarios and long-tail “edge cases” by fusing simulated worlds with ground-truth data (e.g., speed, position, object detection) and vehicle sensor streams to predict system behavior before it’s tested on the road.

### Synthetic and real-world data combine to accelerate progress

Many companies increasingly combine synthetic data with real-world robot data, as alone, neither is sufficient for reliable training. Companies with proprietary datasets or simulation platforms hold strategic advantages, according to BofA Global Research, while those without must license data or simulation infrastructure – risking slower progress in critical modalities like touch, pressure and physical interaction.

<sup>3</sup> CB Insights. (2026, January 22). *The physical AI models market map: Behind the arms race to control robot intelligence*. <https://www.cbinsights.com/research/the-physical-ai-models-market-map/>

### **Observability – robots need compliance too**

As physical AI applications launch commercially, specialized monitoring, debugging and optimization platforms are being developed. Observability tools monitor how robots behave in real-world deployment, helping engineers diagnose failures, replay incidents and push fixes quickly. These platforms turn unexpected edge cases into new training data, feeding insights back into simulation and model development. Strong observability accelerates model improvement and is often what separates production-ready robots from lab-only prototypes. For more on training a responsible AI, read [AI Dictionary, part 3: Resources and risks](#).

### **Autonomy’s transition: From rules-based to end-to-end**

A major shift is underway in autonomy, from hierarchical, rules-based systems towards end-to-end, data-driven AI models. Traditionally, both robots and autonomous vehicles relied on layered control stacks, where perception, planning and action were engineered and optimized separately using hand-coded rules. However, based on advances in large-scale AI – particularly generative and multimodal models – a single neural network can now learn how to turn raw sensor directly into actions. This “end-to-end” paradigm improves generalization, adaptability and scalability, allowing systems to operate across new environments, tasks and even physical form factors with far less manual re-engineering.

### **End-to-end could unlock a shared platform across robots and autonomous vehicles**

This shift underpins the rise of physical AI across both humanoid robots and autonomous vehicles, with significant commercial implications. The same end-to-end principles can be applied to robots manipulating objects and vehicles navigating roads, enabling shared AI infrastructure across platforms. If successful, this approach promises faster deployment, lower costs and more flexible scaling than traditional methods. But it comes with trade-offs, including far higher data and compute requirements, and greater challenges around safety validation, given the difficulty of tracing errors, leading companies to add models or software to address.

While a few robotics companies are developing solely end-to-end AI tech stacks, several are adopting hybrid strategies in the near term, combining end-to-end learning with selective rules and simulation, while progressing towards more unified AI systems as data, compute and regulatory confidence mature.

## **Contributors**

### **Vanessa Cook**

Content Strategist, Bank of America Institute

### **Lynelle Huskey**

Analyst, Bank of America Institute

## **Sources**

### **Martyn Briggs**

Equity Strategist, BofA Global Research

### **Haim Israel**

Equity Strategist, BofA Global Research

### **Felix Tran**

Equity Strategist, BofA Global Research

### **Lauren-Nicole Kung**

Equity Strategist, BofA Global Research

# Disclosures

These materials have been prepared by Bank of America Institute and are provided to you for general information purposes only. To the extent these materials reference Bank of America data, such materials are not intended to be reflective or indicative of, and should not be relied upon as, the results of operations, financial conditions or performance of Bank of America. Bank of America Institute is a think tank dedicated to uncovering powerful insights that move business and society forward. Drawing on data and resources from across the bank and the world, the Institute delivers important, original perspectives on the economy, sustainability and global transformation. Unless otherwise specifically stated, any views or opinions expressed herein are solely those of Bank of America Institute and any individual authors listed, and are not the product of the BofA Global Research department or any other department of Bank of America Corporation or its affiliates and/or subsidiaries (collectively Bank of America). The views in these materials may differ from the views and opinions expressed by the BofA Global Research department or other departments or divisions of Bank of America. Information has been obtained from sources believed to be reliable, but Bank of America does not warrant its completeness or accuracy. These materials do not make any claim regarding the sustainability of any product or service. Any discussion of sustainability is limited as set out herein. Views and estimates constitute our judgment as of the date of these materials and are subject to change without notice. The views expressed herein should not be construed as individual investment advice for any particular person and are not intended as recommendations of particular securities, financial instruments, strategies or banking services for a particular person. This material does not constitute an offer or an invitation by or on behalf of Bank of America to any person to buy or sell any security or financial instrument or engage in any banking service. Nothing in these materials constitutes investment, legal, accounting or tax advice. Copyright 2026 Bank of America Corporation. All rights reserved.