




Brought to you by
THE **ROBOTREPORT**

THE RACE TO MAKE AI PHYSICAL

Physical intelligence, embodied AI, and the engineering stack that moves models from prediction into action

Executive Summary

AI has become remarkably good at understanding language, images, and code. The next challenge is harder, understanding reality.

Physical AI moves intelligence out of digital environments and into machines that must sense, decide, and act in the world. Unlike software, those systems face imperfect measurements, timing constraints, changing conditions, and the consequences of action.

Success will depend on more than model performance. It will require advances across sensing, simulation, computing, connectivity, and control. In many cases, the technologies

that determine whether a system succeeds sit beneath the application layer itself.

A machine that shuts down when vibration crosses a preset limit is following a rule. A machine that reads vibration, temperature, current, and motor behavior, recognizes the signature of bearing wear, and adjusts before failure occurs is doing something closer to judgment.

That distinction helps explain why Physical AI has become one of the most closely watched areas in technology. These systems do more than analyze information. They interact with the world around them, adapting to changing conditions and responding to situations that cannot be fully predicted in advance.

The next phase of AI is not simply about producing better answers. It is about producing useful actions in the presence of imperfect information.

Physical Intelligence versus Physical AI

The terms Physical Intelligence and Physical AI are often used interchangeably. The distinction matters because they describe different parts of the same system.

Physical Intelligence


Physical Intelligence is the capacity of a system where models interpret contextual data whilst interacting with the physical world in real time. A camera frame, accelerometer trace, current waveform, or tactile signal is not intelligence by itself. Intelligence emerges when those signals can be synchronized, interpreted, and trusted well enough to support decisions.

Physical AI

Physical AI is the broad system built on top of that context. It is where systems collect signals, plan, adapt, and interact with the physical world. Physical AI is a system of algorithms, neural networks, sensors, and hardware that all bridge from the digital world to the physical world.

Physical Intelligence is the decision and action layer built on that understanding. A robot can have excellent perception and poor judgment. A model can have sophisticated reasoning and still fail because its understanding of the environment was incomplete or incorrect.

Physical Intelligence determines how well a system understands the world and adapts. Physical AI is the physical embodiment of that system.



The Physical AI Technology Stack

The visible product may be a robot, vehicle, instrument, or wearable. Underneath, many systems share the same stack.

The stack is only as strong as the measurement layer

The current wave of AI has trained people to look for brakthroughs in models. Physical systems push attention back toward measurement. A robot cannot reason about

slip if it cannot detect slip. A vehicle cannot respond to road noise if the audio network introduces unpredictable delay. A wearable cannot infer context if the sensors and power budget force the system to discard the relevant signal.

This is where the old engineering disciplines re-enter the AI conversation: signal integrity, deterministic timing, power management, thermal design, mechanical tolerances, and control stability. Physical AI does not replace those disciplines. It depends on them.

Layer	Role in a physical AI system
Sensing	Captures light, sound, motion, force, position, current, pressure, temperature, vibration, and other physical signals. The constraint is not only accuracy; timing, drift, cost, power, packaging, and calibration all matter.
Signal chain and conversion	Turns raw physical measurements into usable data. This layer often determines whether downstream inference is useful or misleading.
Perception and world modeling	Fuses sensor streams into a working picture of objects, surfaces, constraints, intent, and likely future motion.
Reasoning and policy	Chooses what to do. Current approaches range from end-to-end policies to libraries of smaller skills that can be composed into larger tasks.
Control and actuation	Turns decisions into motion, torque, grip, sound, heat, or other physical effects. Fast local loops often matter more than high-level intelligence here.
Compute placement	Splits work across sensors, embedded processors, edge accelerators, central compute, and cloud infrastructure.
Validation	Tests the system in simulation, controlled lab conditions, and real deployments. The hard question is whether success transfers.



Sim-to-Real Gap, Simulation, and the Digital Twin

The most important difference between conventional AI and Physical AI is not embodiment itself. It is consequence. A language model can be wrong and produce a bad sentence. A physical system can be wrong and bend a connector, damage equipment, waste a production run, or put a person at risk.

The defining challenge is not generating the next token. It is predicting what happens after an action is taken in a dynamic world governed by physics.

This helps explain why advances in language models do not automatically translate into advances in robotics. Language operates in a symbolic environment where mistakes are often inexpensive and reversible. Physical systems operate in environments governed by timing, force, uncertainty, and consequences.

That is why hybrid approaches matter. Many teams are combining learned models with physics-informed methods, model-based control, digital twins, and explicit constraints.

A learned policy may decide the task strategy while a controller enforces torque limits. A world model may learn from data while simulation validates behavior against physical reality.

Physical AI does not replace traditional engineering. It builds on it.

Simulation has become one of the most important tools in Physical AI development. It allows teams to generate training data, explore failure modes, and evaluate systems long before hardware enters production.

The challenge is that simulation is ultimately an approximation of reality. Friction, cable dynamics, contact forces, wear, acoustic reflections, lighting conditions, and thousands of other small variables are difficult to model perfectly.

Progress increasingly depends on calibrated feedback loops between simulation and deployment. Teams measure real-world performance, update simulation models, retrain systems, and validate against physical hardware. Today's goal is not to be a perfect digital twin. It is knowing when a model is accurate enough to support a decision.

Current State of Industry Adoption

The field is past the pure research stage, but far from mature. Adoption is practical where the task is valuable, bounded, and measurable.

- **Warehousing and logistics remain early leaders.** The environment is structured enough to constrain the problem, and the economic value of throughput is clear. Amazon has publicly described reaching one million deployed robots across its operations, a sign that fleet-level autonomy is already operating at industrial scale. [Amazon operations announcement](#)
- **Fixed vision systems are spreading quietly.** Many deployments do not look like robots. Cameras and models monitor motion, detect anomalies, improve safety, and help optimize paths through facilities.
- **Healthcare and inspection are moving carefully.** Autonomous imaging, drones for grid inspection, and remote field robotics are attractive because the work can be dangerous, repetitive, or difficult to staff.
- **Autonomous driving remains a proving ground.** It exposes the full difficulty of the problem: multimodal sensing, fast decisions, edge computing, simulation, regulation, and public safety in one system.

Challenges on the Path to Scale

Challenge	Why it matters
Measurement trust	Bad physical data can make a very capable model confidently wrong. Calibration, drift detection, synchronization, and field robustness are not optional.
Safety certification	A system that learns and adapts is harder to certify than a fixed machine. Regulators and customers will ask for evidence, not just performance claims.
Generalization	A policy trained on one fixture, route, object, or lighting condition may fail when the environment shifts.
Human trust and collaboration	Many deployments will be mixed human-machine environments, where predictability and explainability affect adoption as much as raw capability.
Cost and maintainability	The system has to survive dirty sensors, worn parts, software updates, component substitutions, and technicians who need to service it quickly.

The likely adoption path is therefore uneven. Specialized systems will scale before general-purpose robots. Narrow tasks can be measured, bounded, and certified. General machines may arrive, but they will be built from many narrow competencies that first prove themselves in less glamorous environments.

Industry Leaders Shaping Physical AI

Physical AI is no longer confined to research labs. Advances in computing, simulation, and robotics have pushed autonomous systems from controlled demonstrations into real-world deployment.

Much of the infrastructure behind this progress has been enabled by advances in accelerated computing and simulation. Companies such as NVIDIA have helped establish many of the tools used to train,

test, and validate Physical AI systems, including simulation environments, digital twins, robotics frameworks, and accelerated computing platforms.

At the system level, companies including Boston Dynamics, Tesla, Figure AI, and Agility Robotics are pursuing increasingly capable mobility and manipulation systems. Their approaches differ, but they are ultimately confronting many of the same challenges: autonomy, dexterity, safety, and operation in dynamic environments.

Large-scale deployments are also beginning to emerge. Amazon's robotics fleet demonstrates how Physical AI can create measurable value when applied to well-defined operational problems.

Together, these developments illustrate a broader shift. Physical AI is moving from experimentation toward deployment.



Physical AI Is a Full-Stack Challenge: The Role of Semiconductor Companies in Physical Intelligence

Much of the public conversation around Physical AI centers on foundation models, humanoid platforms, and GPU infrastructure. Far less attention is paid to the technologies that connect those systems to the physical world.

Every Physical AI system depends on sensing, signal conversion, timing, power delivery, connectivity, embedded processing, and control. While robotics companies build machines and AI companies build models, semiconductor companies provide many of the technologies that allow those systems to perceive, understand, and interact with the physical world.

Analog Devices provides a useful example of how the semiconductor industry is adapting to the demands of Physical AI. The company has increasingly positioned itself around what it describes as Physical Intelligence, systems that perceive, reason, and act locally using signals drawn from motion, sound, vibration, magnetics, timing, and other physical domains. ADI has been positioning its portfolio and research efforts around these foundational layers of the stack.


Robotic dexterity and tactile intelligence. At NVIDIA GTC 2026, ADI highlighted a software-defined humanoid hand platform that combines sensing, motor control, real-time connectivity, AI, and high-fidelity simulation. Demonstrations focused on multimodal

tactile sensing, cable manipulation, industrial dexterity benchmarks, and real-to-sim validation for high-precision manipulation tasks. [ADI at NVIDIA GTC 2026](#)

Audio as a reasoning modality. ADI has described a future in which audio evolves beyond a user interface and becomes a contextual sensing modality. Combined with spatial audio, sensor fusion, and on-device AI, systems could use sound to infer environmental context, intent, and situational awareness while enabling more natural interactions between humans and intelligent machines. [ADI 2026 predictions](#)

Deterministic connectivity for software-defined systems. In 2026, ADI introduced A2B 2.0, expanding its audio and connectivity platform with four times the bandwidth of the previous generation while maintaining deterministic latency as low as 62 microseconds. The platform also introduces Ethernet tunneling capabilities and compatibility with existing A2B infrastructure, reflecting the need for predictable, real-time communication inside increasingly software-defined systems. [A2B 2.0 announcement](#)

Edge-native intelligence. ADI and other semiconductor companies has proposed a class of compact, domain-specific models capable of operating entirely at the edge. These systems sit between traditional embedded software and large cloud-hosted foundation models, offering deeper reasoning within narrowly defined tasks while maintaining low latency, local autonomy, and reduced dependence on cloud connectivity.



Analog Devices is one example of how semiconductor companies are responding to the demands of Physical Intelligence. Similar efforts are underway across industrial, automotive, communications, and edge-computing markets. These efforts share a recognition that AI alone is insufficient. Physical AI depends equally on the infrastructure that connects algorithms to the physical world.

Conclusion

Physical AI is sometimes presented as the next robotics boom. That is too narrow. The deeper transition is that AI is being pulled into systems where timing, energy, force, noise, heat, and safety shape the answer.

The winners will not simply be the teams with the largest models. They will be the teams that make the whole loop work: clean measurement, grounded reasoning, dependable control, realistic simulation, and architecture that puts the right intelligence in the right place. Some of that work will look like advanced AI research. Much of it will look like careful engineering.

The future of AI will not be built only in data centers. It will be built at the boundary between the digital and physical worlds, where sensing, computation, and control converge into action.

Source note: company examples and statistics are included to illustrate industry direction, not to endorse any vendor. Links are embedded inline for source review.