



NVIDIA Launches Cosmos 3, the Open Frontier Foundation Model for Physical AI

News Summary:

- NVIDIA Cosmos 3 is a new leaderboard-topping open physical AI foundation model, built on a breakthrough mixture-of-transformers architecture for physical AI reasoning, world simulation and action generation.
- Cosmos 3 is the world's first fully open omnimodel with native vision reasoning and multimodal generation across text, image, video, ambient sound and action for state-of-the-art synthetic data generation and physical AI policy model development.
- NVIDIA launches the NVIDIA Cosmos Coalition with leading AI labs and robotics leaders -- including Agile Robots, Black Forest Labs, Generalist, LTX, Runway and Skild AI -- to advance the next generation of open world models.

NVIDIA GTC Taipei--NVIDIA today launched [NVIDIA Cosmos™ 3](#), an open world foundation model for physical AI built on a breakthrough [mixture-of-transformers](#) architecture that combines vision reasoning, world generation and action prediction in a single system.

Cosmos 3 is the world's first fully open [omnimodel](#) that can natively understand and generate text, images, video, ambient sound and actions with leading physics accuracy, reducing physical AI training and evaluation cycles from months to days.

NVIDIA also launched the NVIDIA Cosmos Coalition, a global collaboration between [world model](#) builders and AI developers -- including [Agile Robots](#), Black Forest Labs, Generalist, LTX, Runway and Skild AI -- working together to advance next-generation world models.

"The big bang of physical AI is just around the corner thanks to breakthroughs in multimodal reasoning language, vision and world models," said Jensen Huang, founder and CEO of NVIDIA. "The Cosmos 3 family of open, frontier omnimodels gives developers a generational leap in ability to build robots, autonomous vehicles and vision AI that perceive, reason, plan and act in the physical world."

A New Architecture for Physical AI

Cosmos 3 tackles a fundamental challenge in physical AI: enabling robots, autonomous vehicles (AVs) or vision agents to generalize in the real world with limited training data and fragmented simulation stacks.

The model's mixture-of-transformers architecture pairs a reasoning transformer with an expert generation transformer, enabling Cosmos 3 to understand object interactions, motion and spatial-temporal relationships before generating video and action trajectories.

Trained on one of the largest multimodal physical AI datasets -- including billions of samples across text, image, video, sound and action trajectories -- the model gives developers a powerful pretrained foundation for building physical AI systems with less data and lower training costs.

Developers can use Cosmos 3 as:

- A [vision language model](#) that understands and reasons across modalities.
- A **world model or video foundation model** that simulates physical environments and predicts future world states for training and evaluation.
- **The backbone for [world action models](#)** that help train robots to perform specific tasks.

Cosmos 3 models deliver leading results on physical AI benchmarks. Among open models, it ranks first across [Artificial Analysis](#), [Physics-IQ](#), [PAI-Bench](#) and [R-Bench](#) for world generation accuracy, [RoboLab](#) and [RoboArena](#) for action policy, and the [VANTAGE-Bench](#) and [TAR](#) leaderboards for vision understanding.

The Cosmos 3 lineup gives developers options for different stages of physical AI development:

- **Cosmos 3 Super** for post-training robotics and AV models that need the highest physics accuracy and generation quality.
- **Cosmos 3 Nano** for high-quality video and action reasoning in fractions of a second.
- **Cosmos 3 Edge**, coming soon, for real-time inference at the edge.

Cosmos Coalition Accelerates Open World Model Development

The Cosmos Coalition is a global collaboration between world model builders, AI developers and physical AI leaders to advance open world models across industries, enabling members to contribute models, research and evaluation techniques while using Cosmos 3 technologies, training tools and [NVIDIA DGX™ Cloud](#) infrastructure for large-scale training.

Founding coalition members include Agile Robots, Black Forest Labs, Generalist, LTX, Runway and Skild AI. By building in the open and contributing across a shared ecosystem, the coalition aims to enable faster innovation, broader interoperability and more rapid advances in physical AI.

Developers Build on Cosmos

The Cosmos platform powers NVIDIA's physical AI stack to accelerate training and evaluation workflows across industries. The platform now includes new datasets for robotics, physics, human motion, autonomous driving, warehouse safety and spatial reasoning, as well as new [physical AI agent skills](#) for neural scene reconstruction, defect-image generation and video augmentation.

Physical AI developers are building on the Cosmos platform across industries -- Agile Robots, Doosan Robotics, LG Electronics, Samsung Electronics and Skild AI for robotics, Li Auto for AVs, and [Centific](#), [Fogsphere](#), [Linker Vision](#), [Milestone Systems](#) and [Yuan](#) for [vision AI agents](#) to power industrial AI and smart spaces applications.

Availability

Cosmos 3 Super and Cosmos 3 Nano are available now, with Cosmos 3 Edge coming soon for real-time inference. Developers can try Cosmos 3 on [build.nvidia.com](#), download open models from [Hugging Face](#), customize models and generate synthetic data with Hugging Face Diffusers and resources on [GitHub](#), and deploy the models as [NVIDIA NIM™](#) microservices.

Model builders and software providers can accelerate access, customization and deployment of Cosmos for key reasoning and synthetic data generation workloads using physical AI agent skills on GitHub through inference services and cloud infrastructure partners including [Baseten](#), [CoreWeave](#), Microsoft Azure, [Nebius](#), Deep Infra and Classmethod.

Watch the [keynote](#) from Huang, learn more at [NVIDIA GTC Taipei](#) and explore these [physical AI sessions](#).

About NVIDIA

[NVIDIA](#) (NASDAQ: NVDA) is the world leader in AI and accelerated computing.

Certain statements in this press release including, but not limited to, statements as to: expectations with respect to growth, performance, availability, and benefits of NVIDIA's products, services and technologies, and related trends and drivers; expectations with respect to NVIDIA's third party arrangements, including with its collaborators and partners; expectations with respect to technology developments, and related trends and drivers; projected market growth and trends; expectations with respect to AI and related industries; and other statements that are not historical facts are forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, which are subject to the "safe harbor" created by those sections based on management's beliefs and assumptions and on information currently available to management and are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic and political conditions; NVIDIA's reliance on third parties to manufacture, assemble, package and test NVIDIA's products; the impact of technological development and competition; development of new products and technologies or enhancements to NVIDIA's existing products and technologies; market acceptance of NVIDIA's products or NVIDIA's partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of NVIDIA's products or technologies when integrated into systems; NVIDIA's ability to realize the potential benefits of business investments or acquisitions; and changes in applicable laws and regulations, as well as other factors detailed from time to time in the most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its Annual Report on Form 10-K and Quarterly Reports on Form 10-Q. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

Many of the products and features described herein remain in various stages and will be offered on a when-and-if-available basis. The statements above are not intended to be, and should not be interpreted as a commitment, promise, or legal obligation, and the development, release, and timing of any features or functionalities described for our products is subject to change and remains at the sole discretion of NVIDIA. NVIDIA will have no liability for failure to deliver or delay in the delivery of any of the products, features or functions set forth herein.

© 2026 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA Cosmos, NVIDIA DGX and NVIDIA NIM are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and/or other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Quentin Nolibois
Corporate Communications
NVIDIA Corporation
press@nvidia.com