

NVIDIA Opens Portals to World of Robotics With New Omniverse Libraries, Cosmos Physical AI Models and AI Computing Infrastructure

- *New NVIDIA Omniverse NuRec 3D Gaussian Splatting Libraries Enable Large-Scale World Reconstruction*
- *New NVIDIA Cosmos Models Enable World Generation and Spatial Reasoning*
- *New NVIDIA RTX PRO Blackwell Servers and NVIDIA DGX Cloud Let Developers Run the Most Demanding Simulations Anywhere*
- *Physical AI Leaders Amazon Devices & Services, Boston Dynamics, Figure AI and Hexagon Embrace Simulation and Synthetic Data Generation*

SIGGRAPH—NVIDIA today announced new [NVIDIA Omniverse™](#) libraries and [NVIDIA Cosmos™](#) world foundation models ([WFMs](#)) that accelerate the development and deployment of robotics solutions.

Powered by new [NVIDIA RTX PRO™ Servers](#) and [NVIDIA DGX™ Cloud](#), the libraries and models let developers anywhere develop physically accurate [digital twins](#), capture and reconstruct the real world in simulation, [generate synthetic data](#) for training [physical AI](#) models and build AI agents that understand the physical world.

“Computer graphics and AI are converging to fundamentally transform robotics,” said Rev Lebedian, vice president of Omniverse and simulation technologies at NVIDIA. “By combining AI reasoning with scalable, physically accurate simulation, we’re enabling developers to build tomorrow’s robots and autonomous vehicles that will transform trillions of dollars in industries.”

New NVIDIA Omniverse Libraries Advance Applications for World Composition

New [NVIDIA Omniverse software development kits \(SDKs\) and libraries](#) are [now available](#) for building and deploying [industrial AI](#) and [robotics simulation](#) applications.

- **New Omniverse SDKs** introduce data interoperability between MuJoCo (MJCF) and [Universal Scene Description \(OpenUSD\)](#), enabling over 250,000 MJCF [robot learning](#) developers to seamlessly simulate robots across platforms.
- **New Omniverse NuRec libraries and AI models** introduce Omniverse RTX ray-traced [3D Gaussian splatting](#), a rendering technique that lets developers capture, reconstruct and simulate the real world in 3D using sensor data.
- **NVIDIA Isaac Sim™ 5.0 and NVIDIA Isaac Lab 2.2** open-source robot simulation and learning frameworks are [now available on GitHub](#). [Isaac Sim](#) now includes NuRec neural rendering and new OpenUSD-based robot and sensor schemas that help robot developers close the simulation-to-reality gap.

[Omniverse NuRec rendering is now integrated in CARLA](#), a leading open-source simulator used by over 150,000 developers. [Autonomous vehicle](#) (AV) toolchain leader [Foretellix](#) is integrating NuRec, [NVIDIA Omniverse Sensor RTX™](#) and Cosmos Transfer to enhance its scalable synthetic data generation with physically accurate scenarios. [Voxel51's](#) data engine for visual and multimodal AI, FiftyOne, supports NuRec to ease data preparation for reconstructions. FiftyOne is used by customers such as Ford and Porsche.

Boston Dynamics, Figure AI, [Hexagon](#), RAI Institute, [Lightwheel](#) and Skild AI are adopting Omniverse libraries, Isaac Sim and [Isaac Lab](#) to accelerate their AI robotics development, while [Amazon Devices & Services](#) is using them to power a new manufacturing solution.

Cosmos Advances World Generation for Robotics

Cosmos WFMs, downloaded over 2 million times, let developers generate diverse data for training robots at scale using text, image and video prompts.

[New models](#) announced at SIGGRAPH deliver major advances in synthetic data generation speed, accuracy, language support and control:

- **Cosmos Transfer-2**, coming soon, simplifies prompting and accelerates photorealistic synthetic data generation from ground-truth 3D simulation scenes or spatial control inputs like depth, segmentation, edges and high-definition maps.
- **A distilled version of Cosmos Transfer** reduces the 70-step distillation process to one so developers can run the model on [NVIDIA RTX PRO Servers](#) at unprecedented speed.

Lightwheel, Moon Surgical and Skild AI are using Cosmos Transfer to accelerate physical AI training by simulating diverse conditions at scale.

Cosmos Reason Breaks Through World Understanding

Since the introduction of OpenAI’s CLIP model, [vision language models](#) (VLMs) have transformed computer-vision tasks like

object and pattern recognition. However, they have not yet been able to solve multistep tasks nor handle ambiguity or novel experiences.

NVIDIA Cosmos Reason — a new open, customizable, 7-billion-parameter reasoning VLM for physical AI and robotics — lets robots and vision AI agents reason like humans, using prior knowledge, physics understanding and common sense to understand and act in the real world. Cosmos Reason topped the [Physical Reasoning Leaderboard](#) on Hugging Face.

Cosmos Reason can be used for robotics and physical AI applications including:

- **Data curation and annotation**, which enables developers to automate high-quality curation and annotation of massive, diverse training datasets.
- **Robot planning and reasoning**, acting as the brain for deliberate, methodical decision-making in a [robot vision language action \(VLA\) model](#). Cosmos Reason lets robots interpret environments and, given complex commands, break them down into tasks and execute them using common sense, even in unfamiliar environments.
- **Video analytics AI agents** built on the [NVIDIA Blueprint for video search and summarization](#) that can extract valuable insights and perform root-cause analysis on massive volumes of video data.

NVIDIA's robotics and **NVIDIA DRIVE™** teams are using Cosmos Reason for data curation and filtering, annotation and VLA post-training. Uber is using it to annotate and caption AV training data.

[Magna](#) is developing with Cosmos Reason as part of its City Delivery platform — a fully autonomous, low-cost solution for instant delivery — to help vehicles adapt more quickly to new cities. Cosmos Reason adds world understanding to the vehicles' long-term trajectory planner. [VAST Data](#), Milestone Systems and [Linker Vision](#) are adopting Cosmos Reason to automate traffic monitoring, improve safety and enhance visual inspection in cities and industrial settings.

New NVIDIA AI Infrastructure Powers Robotics Workloads Anywhere

To enable developers to take full advantage of these advanced technologies and software libraries, NVIDIA announced AI infrastructure designed for the most demanding workloads.

- **NVIDIA RTX PRO Blackwell Servers** offer a single architecture for every robot development workload across training, [synthetic data generation](#), [robot learning](#) and [simulation](#).
- **NVIDIA DGX Cloud**, available on [Microsoft Azure Marketplace](#), now offers Omniverse developers a fully managed platform to simplify streaming OpenUSD- and NVIDIA RTX™-based applications at scale from the cloud, minimizing infrastructure orchestration and management. Accenture and Hexagon are among the first industry leaders to adopt the platform.

Accelerating the Developer Ecosystem

To help robotics and physical AI developers advance 3D and simulation technology adoption, NVIDIA also announced:

- **OpenUSD Curriculum and Certification**, which addresses demand for USD expertise, with support from [AOUSD](#) members Adobe, Amazon Robotics, Ansys — part of Synopsys, Autodesk, Pixar, PTC, Rockwell Automation, SideFX, Siemens, TCS and Trimble, as well as industry leaders such as Hexagon.
- **Open-source collaboration with Lightwheel** to integrate robot policy training and evaluation frameworks into NVIDIA Isaac Lab, featuring parallel reinforcement learning training capabilities, benchmarks and simulation-ready assets for robot manipulation and locomotion.

Watch the [NVIDIA Research special address at SIGGRAPH](#).

About NVIDIA

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

Certain statements in this press release including, but not limited to, statements as to: computer graphics and AI converging to fundamentally transform robotics; by combining AI reasoning with scalable, physically accurate simulation, NVIDIA enabling developers to build tomorrow's robots and autonomous vehicles that will transform trillions of dollars in industries; the benefits, impact, performance, and availability of NVIDIA's products, services, and technologies; expectations with respect to NVIDIA's third party arrangements, including with its collaborators and partners; expectations with respect to technology developments; 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 product 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; 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.

© 2025 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA Cosmos, NVIDIA DGX, NVIDIA DRIVE, NVIDIA Isaac Sim, NVIDIA Omniverse, NVIDIA Omniverse Cloud Sensor RTX, NVIDIA RTX and NVIDIA RTX PRO 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. Features, pricing, availability, and specifications are subject to change without notice.

Quentin Nolibois
+1 415-741-8356
qnolibois@nvidia.com