



# NVIDIA Launches Alpamayo 2 Super Open Reasoning Model for Robotaxis

*Editor's note: This post, originally published May 31, 2026, has been updated to reflect NVIDIA renamed NVIDIA OmniDreams to NVIDIA Cosmos-Dreams. As of September 2026, the post has also been updated to reflect NVIDIA DRIVE Hyperion has been rebranded to NVIDIA Hyperion.*

## News Summary:

- NVIDIA's most powerful open reasoning model to date, NVIDIA Alpamayo 2 Super is an open 34-billion-parameter reasoning VLA model that reasons, plans and acts across the full driving stack for safer, scalable level 4 development.
- NVIDIA AlphaGym is a new high-throughput, closed-loop reinforcement learning framework that trains AV models on the consequences of their driving decisions in simulation before road deployment.
- NVIDIA Cosmos-Dreams is a new generative world model for photorealistic closed-loop AV scenario generation, enabling developers to simulate rare and long-tail driving scenarios at scale.
- NVIDIA physical AI agent skills for AV development include Neural Reconstruction powered by NVIDIA Omniverse NuRec, enabling developers to reconstruct real-world fleet data into photorealistic 3D scenes and adapt them across vehicle sensor configurations.

**NVIDIA GTC Taipei**--NVIDIA today introduced NVIDIA Alpamayo 2 Super, a 34-billion-parameter reasoning-based vision language action (VLA) model that extends the [NVIDIA Alpamayo](#) family of open AI models, simulation frameworks and physical AI datasets for safe, level 4 [robotaxi](#) development.

Alongside the model, the company announced new tools, models and agent skills that complete the pipeline from real-world data capture to closed-loop training and in-vehicle deployment, including NVIDIA AlphaGym, [NVIDIA Cosmos-Dreams](#) and new [NVIDIA Omniverse NuRec models](#).

Alpamayo 2 Super helps accelerate autonomous vehicle (AV) development by eliminating the need to build key autonomy infrastructure from scratch. It enables humanlike perception, reasoning and action, and provides the interpretability needed for safety validation and regulatory collaboration.

To better train models for on-road deployment, the AlphaGym framework provides a platform for closed-loop reinforcement learning (RL). The NVIDIA Cosmos-Dreams generative world model for photorealistic closed-loop AV scenario generation enables developers to simulate rare and long-tail driving scenarios at scale.

To amplify developer productivity, NVIDIA is providing [physical AI agent skills](#) for all of its AV development tools. For example, the [Neural Reconstruction skill](#) powered by NVIDIA Omniverse NuRec uses real-world fleet driving scenarios for simulation and generates synthetic training data at scale.

"Alpamayo is the moment cars begin to safely reason, not just drive," said Jensen Huang, founder and CEO of NVIDIA. "Only NVIDIA makes available open models, simulation, real-world data and agent skills so the entire global robotaxi ecosystem can develop level 4 capabilities that understand edge cases, explain decisions, earn trust and scale safely to millions of vehicles."

## Alpamayo 2 Super, Now Available for Reasoning-Based AVs

The NVIDIA Alpamayo family now scales from 10 billion to 34 billion parameters with Alpamayo 2 Super -- going beyond trajectory generation to reason, plan and act across the full driving stack. With multitask capabilities spanning reasoning, auto-labeling, scene understanding, model critiquing and distilling knowledge into smaller models, it provides the building blocks for scalable L4 AV development and deployment.

Key features include:

- **3x parameter scale:** Built on [NVIDIA Cosmos](#)™ world foundation models, Alpamayo 2 Super scales to 34 billion parameters compared with previous 10-billion-parameter generations, improving reasoning, 3D spatial understanding and trajectory prediction in long-tail scenarios.
- **Full-surround perception:** Expands from front-focused cameras to 360-degree situational awareness across front, side and rear views, giving the model complete context for safer lane changes, merges and intersection crossing.
- **Meta-Actions:** Adds Meta-Action output -- including macro actions such as yield, lane change and stop -- so the model predicts high-level driving decisions for downstream planning in addition to trajectories and chain-of-causation (CoC) traces.
- **Reasoning auto-labeling and 2D grounding:** Introduces reasoning auto-labeling with 2D grounding so the 34-billion-parameter foundation model can provide high-quality reasoning labels, compressing annotation cycles from months to

days and reshaping AV data pipeline economics.

- **Improved CoC and trajectory quality:** Improved CoC traces and trajectories, especially in rare, complex, long-tail scenarios where traditional imitation-learning AV stacks struggle.

These advancements make Alpamayo 2 Super NVIDIA's most powerful open driving foundation model to date. Designed as a teacher model, Alpamayo 2 Super can be distilled into compact models that run on the accelerated compute of the [NVIDIA Hyperion™](#) platform -- [NVIDIA DRIVE AGX Thor™](#), which runs inside the vehicle.

As the teacher model scales from 10-billion-parameter models like NVIDIA Alpamayo 1 Nano and NVIDIA Alpamayo 1.5 Nano to 34 billion parameters with Alpamayo 2 Super, a downstream AV stack built on Alpamayo inherits higher-quality reasoning and perception from a single open release, without requiring each manufacturer to rebuild from scratch.

Alpamayo was recently recognized by the [COMPUTEX Best Choice Awards](#), winning in the Vehicle Technology and Smart Cockpit category.

Since launch, Alpamayo has been downloaded close to 400,000 times. The Alpamayo open platform also includes post-training scripts that allow researchers and developers to adapt the models to their own datasets, scenarios and driving policies.

Alpamayo 2 Super is expected to be available this summer on GitHub for inference code and Hugging Face for model weights.

### **AlpaGym Enables Closed-Loop Training and Deployment Cycles**

NVIDIA also introduced NVIDIA AlpaGym, an open source, high-throughput, closed-loop RL framework.

Where open-loop training evaluates models against recorded data and generates a single round of actions, AlpaGym runs models through continuous decision and observation cycles in NVIDIA AlpaSim, with every braking, steering and navigation choice affecting the environment.

As a result, AlpaGym exposes the compounding errors and edge-case failures that static datasets miss, allowing models to learn from experience.

Built on the [AlpaSim](#) microservice simulation stack and NVIDIA Omniverse NuRec, AlpaGym enables efficient, scalable, closed-loop RL to push the frontier of driving performance. In combination with the [Physical AI AV Dataset](#), Alpamayo provides a continuous path from [open-loop pretraining to closed-loop refinement](#).

NVIDIA is also releasing the [CoC Auto-Labeling Pipeline](#) as open source on GitHub. The pipeline automatically generates decision-grounded and causally linked CoC labels from raw driving clips with no human annotation required, providing the causal training data foundation needed to train embodied reasoning models at scale.

### **New Physical AI Agent Skills for AV Powered by NVIDIA**

To support reasoning-based AV development, NVIDIA is launching new physical AI agent skills, under NVIDIA Agent Toolkit, to guide developers and their coding agents through the simulation, data generation and closed-loop training workflows needed to build and validate autonomous driving systems at scale.

This includes Neural Reconstruction skills powered by NVIDIA Omniverse NuRec libraries, NVIDIA Cosmos-Dreams skills for photorealistic scenario generation and AlpaGym skills for closed-loop RL.

Watch Huang's [keynote](#) and learn more at [NVIDIA GTC Taipei](#).

### **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 DRIVE AGX Thor and NVIDIA Hyperion 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.

Marie Labrie  
Corporate Communications  
NVIDIA Corporation  
[press@nvidia.com](mailto:press@nvidia.com)