Comparing Synthetic Data Platforms: Synetic AI and NVIDIA Omniverse

This paper evaluates Synetic AI and NVIDIA Omniverse as platforms for generating synthetic data for computer vision, focusing on workflow simplicity, scalability, and dataset utility.

Published by Synetic AI · June 2025

Download as PDF

Table of Contents

Synetic AI vs. NVIDIA Omniverse

	SYNETIC AI	NVIDIA OMNIVERSE
Workflow Approach	Fully managed service	Customizable platform
Ease of Use	No 3D expertise needed	Steep learning curve
Data Outputs	RGB, bounding boxes, depth, segmentation, keypoints	RGB, bounding boxes, depth, segmentation
Price	1 cent per image	Varies
Cost Transparency	Clear and straightforward	Less predictable

> *This white paper compares Synetic AI and NVIDIA Omniverse for synthetic data generation, focusing on deployment-ready computer vision models. Whether you're exploring simulation tools or evaluating dataset creation platforms, this guide outlines key differences and trade-offs.*

Why Compare These Platforms?

NVIDIA Omniverse and Synetic AI represent two distinct approaches to synthetic data generation. This section explores their comparative strengths and limitations, particularly for organizations seeking to develop production-ready vision models as opposed to building custom simulation environments.

Evolution of Simulation Tools in Synthetic Data Generation

Simulation platforms like Omniverse have played a central role in advancing synthetic data generation by enabling physics-based rendering and complex 3D scenarios. These platforms remain powerful tools for teams with specialized needs and simulation expertise. At the same time, the landscape is evolving. Newer platforms such as Synetic AI are emerging with a focus on simplifying synthetic data pipelines, emphasizing ease of use, scalability, and alignment with downstream vision model requirements. This section explores how the field is shifting toward broader accessibility without compromising on output fidelity.

Simulation Team Requirements

Omniverse typically necessitates a team of simulation engineers and 3D artists to support dataset creation workflows. Synetic's methodology is positioned as an intermediary between manual and fully autonomous processes, allowing users to specify requirements and receive high-quality datasets. This approach may reduce or eliminate the need for a dedicated simulation team in certain contexts.

Using Omniverse Cosmos Assets

Omniverse Cosmos provides a curated library of 3D assets primarily intended for simulation and visualization. While these assets can be imported into Synetic, they often require modification to support procedural variation, physics-based interactions, or licensing suitable for synthetic dataset generation. Synetic emphasizes the use of reusable, procedurally generated assets that are specifically designed for vision AI applications.

Workflow Design and Automation

Omniverse offers extensive control over the simulation stack, which can increase complexity in environment setup, asset management, and output configuration. These processes may extend engineering timelines. In contrast, Synetic AI eliminates the need for scripting, manual 3D setup, or simulation configuration. Users simply specify what they want the model to recognize and provide context about the intended use case. From there, Synetic AI handles everything—from asset creation and image generation to

scenario variability, model training, and SDK wrapping—streamlining the entire synthetic data workflow into a single, guided experience.

Balancing Simulation Fidelity with Dataset Utility

Omniverse is capable of generating high-fidelity simulations suitable for robotics demonstrations or immersive visualization. However, high realism does not inherently guarantee the utility of training data, which frequently depends on coverage, variation, and annotation quality. Synetic AI emphasizes dataset utility, including edge case density, physics-based variation, and consistent annotation, to align with the requirements of contemporary vision models.

Iteration Speed and Productivity

Iteration speed is a critical factor in synthetic dataset generation. Omniverse provides considerable flexibility, but this can introduce additional pipeline overhead and increase iteration times. Synetic AI is designed to support rapid dataset iteration through the use of prebuilt assets, procedural environments, and automation, enabling the generation of large volumes of annotated images in shorter timeframes.

Optimizing for Vision AI Tasks

Omniverse serves as a robust simulation platform and is widely used in robotics, visualization, and digital twin applications. While simulation is foundational, further optimization is often necessary for effective model

training. Synetic AI is developed with a focus on vision tasks, incorporating features and automations that target training accuracy, deployment efficiency, and inference robustness. Its datasets are structured to address practical challenges encountered in real-world computer vision scenarios.

Data Ownership and Pipeline Transparency

Omniverse's complexity can present challenges in understanding internal data processing and pipeline operations, which may be a concern for organizations with regulatory or transparency requirements. Synetic AI provides data pipelines intended to be transparent and auditable, offering users greater control and insight into dataset generation and output alignment with workflow standards.

Pricing Models and Licensing Considerations

Licensing and cost transparency are critical when selecting a synthetic data platform. Omniverse can involve complex pricing structures, including hardware requirements, asset licensing, and cloud resource costs—often making it difficult to estimate total spend.

Synetic AI offers a simpler, usage-based pricing model with clearly defined licensing terms. There are no hidden fees or dependencies on proprietary hardware. Everything from asset generation to dataset export and model training is included in transparent tiers, enabling teams to predict costs and scale confidently.

Estimated Cost Comparison

To understand how these models compare in real-world budgets, consider a typical vision project requiring 100,000 annotated images:

Component	Synetic AI	Omniverse-Based Workflow
Image Cost	~$0.01 per image (fully annotated)	$0.10–$5.00+ per image (including labor and tools)
3D Asset Licensing	Included or procedurally generated	Typically licensed per asset
Engineering Overhead	Low — no simulation team required	High — 3D artists, simulation engineers, pipeline devs
Training Infrastructure	Included in platform	Must be built and managed separately
Setup Requirements	Browser-based	GPU-equipped workstation with Omniverse installed
Total Project Cost (100K images)	**~$1,000**	**$50,000–$250,000+**

In many real-world scenarios, Synetic AI is 50x to 250x more cost-effective than traditional simulation-based approaches like Omniverse—especially when speed, automation, and minimal overhead are important.

Estimates are based on publicly available reports and customer case studies from companies using NVIDIA Omniverse for simulation workflows,

compared with internal usage data from Synetic AI platform deployments.

Platform Ecosystem and Community Support

Omniverse benefits from a large ecosystem, which may introduce additional complexity such as enterprise licensing, structured development roadmaps, and broad tooling that may not always align with specific user needs. Synetic AI is positioned as a more specialized platform, with development cycles responsive to user feedback and new use cases. The platform offers direct user support based on submitted feedback, rather than relying solely on traditional ticketing systems.

Approach to Dataset Delivery

Synetic AI emphasizes turnkey dataset delivery, allowing teams to define desired model learning objectives while the platform manages asset generation and annotation. This approach can reduce internal resource allocation for infrastructure and may accelerate the transition from concept to trained model.

Support for Continuous Feedback and Refinement

Rapid iteration and continuous improvement are increasingly important in dynamic industries. Synetic AI is designed to facilitate short feedback loops, supporting frequent updates and refinements to datasets and models. This

enables users to address edge cases or explore new scenarios in alignment with evolving project requirements.

Scalability and Operational Overhead

While Omniverse enables extensive customization, scaling operations may require dedicated personnel and detailed configuration management. Synetic AI incorporates automated scaling features such as procedural generation, template-driven variation, and API integration, supporting dataset growth with potentially reduced engineering overhead.

Output Format Flexibility and Model Compatibility

Omniverse is capable of generating diverse simulation outputs, though aligning these with specific training architectures may necessitate additional customization. Synetic AI provides outputs that are pre-aligned with popular vision model architectures, including YOLO, DINO, and transformer-based pipelines. Annotations and metadata are formatted for direct model ingestion, minimizing the need for post-processing.

Sensor Simulation and Depth Buffer Support

Synetic AI provides automated depth buffer generation and supports multi-sensor outputs, including RGB, segmentation maps, and keypoints, all synchronized for multi-modal vision model training. In Omniverse, achieving comparable depth and sensor alignment generally involves custom scripting

and manual configuration, which may increase development time and the potential for configuration errors.

Scene Variation Capabilities

Synetic AI supports real-time procedural variation across lighting, weather, asset arrangement, and camera parameters. Scene variation can be introduced without full reconstruction. In contrast, Omniverse often requires timeline reauthoring or scripting of environment parameters to achieve similar variation, which may extend iteration cycles.

Vendor Lock-In and Portability

Synetic AI generates datasets in open, widely adopted formats, facilitating integration with diverse training pipelines across cloud and on-premises environments. Omniverse, by contrast, frequently relies on NVIDIA-specific hardware, connectors, and drivers, which may constrain portability and infrastructure choices.

Optimization for Real-World Deployment

Synetic AI is designed with real-world deployment considerations, allowing datasets to be generated with attention to hardware constraints, camera placement, and operational conditions in diverse environments such as agriculture, warehousing, and edge computing. Omniverse is highly capable

in simulation but may require additional configuration to address downstream inference requirements.

Inference-Time Constraints in Dataset Design

Many platforms focus primarily on dataset creation, leaving training and inference considerations to downstream processes. However, inference-time constraints—such as sensor configuration, latency, and edge deployment—can significantly influence model effectiveness. Synetic AI incorporates these factors into dataset design, considering camera placement, hardware limitations, and intended use cases alongside training requirements.

Access Requirements and Deployment Environment

Omniverse typically requires installation on GPU-equipped systems, along with associated drivers and dedicated workstations. In contrast, Synetic AI operates entirely within a web browser, enabling dataset setup and download without installation or configuration. This allows users to generate datasets and initiate model training from a variety of locations and devices, without command-line interaction.

Platform Differentiators at a Glance

Capability	Synetic AI	NVIDIA Omniverse
Deployment	Browser-based, zero install	GPU workstation with local install
Workflow Setup	Prompt-based, no scripting required	Manual asset management and scripting
Asset Source	Procedurally generated or user-provided	User-provided or Omniverse Cosmos
Training & SDK Wrapping	Built-in and automated	Requires external tooling
Depth & Sensor Outputs	Standardized RGB, segmentation, keypoints, depth	Manually configured per scene
Real-World Optimization	Includes field constraints, inference-time awareness	Must be manually integrated
Estimated Cost (100K images)	~$1,000	$50,000–$250,000+

Use Case Snapshots

AgTech: Detect crop emergence and plant health without needing to fly drones over every field.

Retail: Monitor shelf stock, track foot traffic, and test store layouts, all from synthetic footage.

Robotics: Train object recognition and spatial navigation systems across hundreds of randomized layouts before a single physical test.

Security: Simulate edge-case scenarios for threat detection without compromising real-world privacy or safety.

Customer Perspectives

ML Engineer: “We used to spend weeks wiring up simulation scripts. With Synetic AI, we were training on new datasets the same day.”

Computer Vision PM: “It’s the only platform where I can say what I want and get training data back without babysitting a pipeline.”

Data Scientist: “We finally had consistent annotations and full variation coverage, without labeling a single image by hand.”

Next Steps

If you're ready to build a production-ready dataset, you can place an order directly with Synetic AI—no demo required, and no installation necessary. For a preview of output quality, download a free sample dataset.

If you're interested in exploring simulation-based workflows in more depth, you can learn more about NVIDIA Omniverse here.

Synetic AI Implementation Considerations

Can I use my own 3D assets? It depends. Synetic AI's 3D assets are custom-built to be dynamic, realistic, and optimized for high-quality vision model training. While you can provide your own assets, they may not deliver ideal results without modification. Asset creation is included at no extra cost—every dataset we generate comes with fully tailored, simulation-ready components.

Which models are supported? We support popular architectures like YOLOv8, RT-DETR, DINOv2, and custom pipelines.

What's required to get started? Nothing to install. The entire platform runs in your browser.

Can I try it before committing? Yes. We offer sample datasets and free trials for qualified teams.

Who owns the 3D models after dataset creation? All custom 3D assets created by Synetic AI remain the property of Synetic AI. This is because clients are not charged separately for asset development—it is included as part of the platform service. Clients retain full rights to the generated datasets and model outputs, but not to the underlying 3D models used to create them.

Conclusion

Omniverse is well-suited for interactive 3D simulation development. For organizations focused on training vision models for deployment in real-world conditions, Synetic AI represents an alternative approach that emphasizes workflow efficiency, dataset utility, and reduced operational complexity. The choice between platforms depends on specific project requirements and resource considerations.

It's also worth noting that some organizations may value the deep configurability and hands-on control that platforms like Omniverse provide. For teams that prefer building highly customized simulations from the ground up, or have existing infrastructure and personnel tailored for those workflows, Omniverse can offer unmatched flexibility. The decision ultimately depends on whether speed, automation, and accessibility—or fine-grained simulation control—are higher priorities for the team.

← Back to Explainer