

H100 GPU: A New Era for Deep Learning in Scientific Discovery

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: H100

Cat. No.: B15585327

[Get Quote](#)

A comparative guide for researchers, scientists, and drug development professionals on the performance of the NVIDIA **H100** Tensor Core GPU for deep learning model training in science.

The NVIDIA **H100** Tensor Core GPU, built on the Hopper architecture, represents a significant leap forward in computational power for scientific research. This guide provides an objective comparison of the **H100**'s performance against its predecessors, the A100 and V100, in the context of deep learning model training for scientific applications. We present supporting experimental data from industry-standard benchmarks and detailed methodologies to enable researchers to make informed decisions for their computational needs.

Key Architectural Advancements

The **H100** introduces several key architectural improvements over the Ampere (A100) and Volta (V100) architectures, leading to substantial performance gains. These include fourth-generation Tensor Cores, a new Transformer Engine, and support for the FP8 data format, which collectively accelerate AI training and inference.^{[1][2]} The **H100** also boasts a significantly higher memory bandwidth with HBM3, enabling faster data access for large-scale models.^{[3][4]}

Performance Benchmarks: H100 vs. A100 vs. V100

The following tables summarize the key specifications and performance benchmarks of the NVIDIA **H100**, A100, and V100 GPUs. The data is compiled from publicly available information

and results from the MLPerf HPC benchmark suite, a standardized test for high-performance computing applications.

Table 1: GPU Specifications[3][4][5][6]

Feature	NVIDIA H100 (SXM5)	NVIDIA A100 (SXM4 80GB)	NVIDIA V100 (SXM2 32GB)
Architecture	Hopper	Ampere	Volta
CUDA Cores	16,896	6,912	5,120
Tensor Cores	528 (4th Gen)	432 (3rd Gen)	640 (1st Gen)
Memory	80 GB HBM3	80 GB HBM2e	32 GB HBM2
Memory Bandwidth	3.35 TB/s	2.0 TB/s	900 GB/s
FP64 Performance	60 TFLOPS	19.5 TFLOPS	7.8 TFLOPS
FP32 Performance	60 TFLOPS	19.5 TFLOPS	15.7 TFLOPS
TF32 Tensor Core	989 TFLOPS	312 TFLOPS	N/A
FP16/BF16 Tensor Core	1,979 TFLOPS	624 TFLOPS	125 TFLOPS
FP8 Tensor Core	3,958 TFLOPS	N/A	N/A
Interconnect	NVLink 4.0 (900 GB/s)	NVLink 3.0 (600 GB/s)	NVLink 2.0 (300 GB/s)
Max Power Consumption	700W	400W	300W

Table 2: MLPerf HPC v2.0 Training Benchmarks (Time-to-Train in Minutes)[1][7][8]

Model	NVIDIA H100 (80GB)	NVIDIA A100 (80GB)
CosmoFlow	1.83	4.58
DeepCAM	2.95	8.23
OpenCatalyst	10.88	27.22

Lower is better. Results are based on submissions to the MLPerf HPC v2.0 benchmark and represent the time to train the model to a predefined quality target.

Experimental Protocols

To ensure the reproducibility of the presented benchmarks, this section details the methodologies for the key experiments cited.

CosmoFlow

The CosmoFlow benchmark trains a 3D convolutional neural network to predict cosmological parameters from N-body simulation data.[\[9\]](#)

- Dataset: The dataset consists of 3D dark matter simulation outputs, with each sample being a 128x128x128 voxel cube with 4 channels. The total dataset size is approximately 5.1 TB. [\[1\]](#)[\[10\]](#)
- Model: A 3D convolutional neural network implemented in TensorFlow with Keras.[\[5\]](#)
- Training Parameters:
 - Framework: TensorFlow
 - Optimizer: Adam
 - Batch Size: Varies depending on the submission, typically scaled with the number of GPUs.
 - Learning Rate: A custom learning rate schedule is often employed.
 - Distributed Training: Horovod is used for multi-GPU and multi-node training.[\[1\]](#)
- Target: The model is trained to a validation mean absolute error of less than 0.124.[\[5\]](#)

DeepCAM

The DeepCAM benchmark trains a deep learning segmentation model to identify extreme weather phenomena in climate simulation data.[\[11\]](#)

- Dataset: The CAM5 dataset, which contains 16-channel images of size 768x1152 pixels. The total dataset size is 8.8 TB.[\[9\]](#)[\[11\]](#)
- Model: A U-Net like convolutional encoder-decoder architecture for semantic segmentation, implemented in PyTorch.[\[11\]](#)
- Training Parameters:
 - Framework: PyTorch
 - Optimizer: Adam
 - Batch Size: Scaled with the number of GPUs.
 - Learning Rate: A polynomial decay learning rate schedule is commonly used.
 - Distributed Training: PyTorch's native distributed data-parallel is used.[\[11\]](#)
- Target: The model is trained to a validation Intersection over Union (IoU) of greater than 0.82.[\[11\]](#)

Genomics: NVIDIA Parabricks Germline Workflow

NVIDIA Parabricks is a suite of GPU-accelerated tools for genomic analysis. The germline pipeline performs variant calling from FASTQ files.

- Workflow: The workflow typically includes BWA-MEM for alignment, sorting, marking duplicates, and BQSR, followed by a variant caller like HaplotypeCaller or DeepVariant.[\[12\]](#)[\[13\]](#)
- Command-line Example (for HaplotypeCaller):
- Key Parameters: The choice of variant caller (HaplotypeCaller vs. DeepVariant) and specific parameters for each tool can be configured.[\[14\]](#)

Protein Folding: OpenFold

OpenFold is a trainable, memory-efficient, and open-source reproduction of AlphaFold2, a deep learning model for protein structure prediction.[\[15\]](#)

- Dataset: A preprocessed dataset including PDB mmCIF files and sequence alignments. The original OpenFold dataset is available on the Registry of Open Data on AWS (RODA).[\[16\]](#)
- Model: A complex neural network architecture that processes multiple sequence alignments and templates to predict 3D protein structures.
- Training Parameters:
 - Framework: PyTorch
 - Configuration Preset: `initial_training` is a common starting point.[\[16\]](#)
 - Distributed Training: DeepSpeed can be used for training on multiple GPUs.[\[17\]](#)
 - Key flags for training include `--config_preset`, `--num_nodes`, `--gpus`, and paths to the dataset and alignment directories.[\[16\]](#)

Drug Discovery: Generative Models with NVIDIA BioNeMo


NVIDIA BioNeMo is a framework for training and deploying large language models for biology and chemistry, enabling applications like de novo molecular design.[\[4\]](#)[\[18\]](#)

- Models: BioNeMo includes pretrained models like MegaMolBART for generative chemistry.[\[19\]](#)
- Workflow:
 - Data Preparation: Curate and preprocess a dataset of chemical structures (e.g., in SMILES format).
 - Model Fine-tuning: Fine-tune a pretrained generative model like MegaMolBART on the custom dataset to learn the chemical space of interest.
 - Inference: Use the fine-tuned model to generate novel molecules with desired properties.

- Tools: The BioNeMo framework provides tools and tutorials for each step of this workflow.[\[3\]](#)
[\[19\]](#)


Visualizing Workflows and Relationships

The following diagrams, generated using Graphviz, illustrate a typical experimental workflow for genomics analysis and the logical relationship between the key components of the NVIDIA GPU architectures discussed.



[Click to download full resolution via product page](#)

Genomics Analysis Workflow with NVIDIA Parabricks.



[Click to download full resolution via product page](#)

Evolution of NVIDIA GPU Architectures.

Conclusion

The NVIDIA **H100** GPU delivers substantial performance improvements for deep learning model training in scientific applications compared to its predecessors. Its architectural enhancements, particularly the fourth-generation Tensor Cores and the Transformer Engine with FP8 support, translate to significantly faster training times for a wide range of scientific models. For research and development in fields like genomics, drug discovery, and climate science, the **H100** offers a powerful tool to accelerate discovery and innovation. While the A100 remains a capable platform, the **H100** is the clear choice for computationally intensive, large-scale deep learning tasks at the forefront of scientific research.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. GitHub - sparticlesteve/cosmoflow-benchmark: Benchmark implementation of CosmoFlow in TensorFlow Keras [github.com]
- 2. gitlab.com [gitlab.com]
- 3. youtube.com [youtube.com]
- 4. medium.com [medium.com]
- 5. proxyapps.exascaleproject.org [proxyapps.exascaleproject.org]
- 6. germline - NVIDIA Docs [docs.nvidia.com]
- 7. GitHub - azrael417/mlperf-deepcam: This is the public repo for the MLPerf DeepCAM climate data segmentation proposal. [github.com]
- 8. OpenFold — NVIDIA BioNeMo Framework [docs.nvidia.com]
- 9. mlcommons.org [mlcommons.org]
- 10. arxiv.org [arxiv.org]
- 11. proxyapps.exascaleproject.org [proxyapps.exascaleproject.org]
- 12. germline (GATK Germline Pipeline) - NVIDIA Docs [docs.nvidia.com]
- 13. NVIDIA Parabricks Germline DNA Workflow | Element Biosciences Software Documentation [docs.elembio.io]
- 14. germline - NVIDIA Docs [docs.nvidia.com]
- 15. Optimizing OpenFold Training for Drug Discovery | NVIDIA Technical Blog [developer.nvidia.com]
- 16. Training OpenFold - OpenFold documentation [openfold.readthedocs.io]
- 17. openfold/docs/source/original_readme.md at main · aqlaboratory/openfold · GitHub [github.com]
- 18. GitHub - NVIDIA/digital-biology-examples: NVIDIA Digital Biology examples for optimized inference and training at scale [github.com]
- 19. docs.nvidia.com [docs.nvidia.com]
- To cite this document: BenchChem. [H100 GPU: A New Era for Deep Learning in Scientific Discovery]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15585327#benchmarking-h100-for-deep-learning-model-training-in-science]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com