

310 AI ADVANCES PROTEIN DESIGN WITH AMD

CASE STUDY

310 AI accelerates molecular design and AI-driven discovery cycles on AMD Instinct™ MI300 Series GPUs



310 AI transforms the way therapeutic proteins and peptides are created. The California-based startup builds generalist, generative AI models for protein design, inspired by the principles that drive large language and image models. Its systems can interpret plain-language biological descriptions and generate entirely new functional protein sequences, an approach the company calls text-to-protein. Proteins make up more than half of modern therapeutics, and 310 AI's goal is to automate their design. While therapeutics are the company's first focus, it sees molecule programming as a foundation for materials science, agriculture, energy, and other domains.

“With the large memory on AMD Instinct MI300 Series GPUs, we make fewer trade-offs, iterate faster, and get candidates to the lab sooner.”

Koosh Azimian, CEO and co-founder, 310 AI

Such ambition requires enormous compute power. To reach a useful biological scale, 310 AI trains transformer-based foundation models on billions of biological data points using AMD Instinct™ MI300 Series GPUs and the open-source AMD ROCm™ software stack. The combination of large amounts of memory and accessible, open-standard tools allows the team to move quickly from early research to practical drug design results.

BUILDING A FOUNDATION FOR MOLECULE PROGRAMMING

310 AI's flagship model, MP4 (Molecule Programming version 4), shifts the problem from prediction to creation. Instead of predicting a protein structure from an existing sequence, MP4 generates de novo proteins from functional text prompts. The model was trained on 3.2 billion biological data points across 70 tasks, using about 3,800 AMD Instinct MI250X GPU-days, establishing a new benchmark for generative biology. “We built MP4 on top of AMD ROCm software and PyTorch to take advantage of the large memory and high throughput of AMD Instinct GPUs,” says Koosh Azimian, CEO and co-founder of 310 AI.

Their training results highlight both the scale and pace of 310 AI's work. The models range from 500 million to 15 billion parameters, reflecting the scale required for true generative biology. The team achieved about 200 sequences per second during training and sustained roughly 36 proteins per second in full inference, meaning it can explore massive biological spaces in days rather than weeks.

INDUSTRY

Software & Sciences

CHALLENGES

310 AI needed to train massive foundation models for protein design, requiring high memory, fast throughput, and efficient compute to progress from research to a potential real-world drug discovery.

SOLUTION

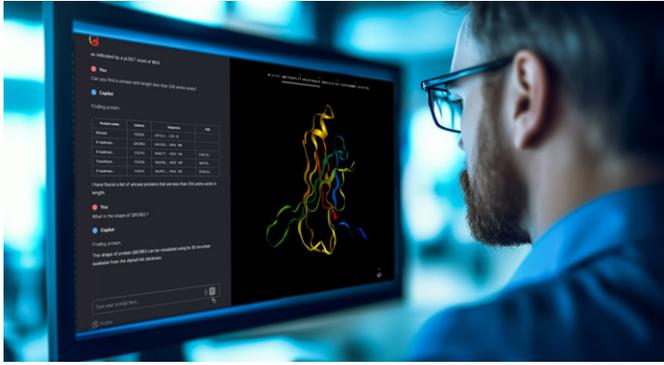
The team built its MP4 model on AMD Instinct™ MI300 Series GPUs and AMD ROCm™ software, leveraging large-memory capacity, strong bandwidth, and open frameworks for rapid model scaling.

RESULTS

310 AI achieved faster convergence, lower costs, and up to an 84% expression rate in lab-tested proteins, cutting 150,000 candidates to 96 finalists in days helping to accelerate discovery cycles.

AMD TECHNOLOGY AT A GLANCE

AMD Instinct™ MI250X GPU
AMD Instinct™ MI300X GPU
AMD Instinct™ MI325X GPU
AMD ROCm™ software



310 AI uses AMD Instinct™ GPUs to help generate new protein structures.

WORKING WITH AMD TO ACCELERATE DISCOVERY

Early generations of 310 AI models ran on competitor’s GPU hardware. In 2024, the team began migrating to AMD to support its MP3 and MP4 models. Engineers expected a difficult port from CUDA to AMD ROCm software. The transition proved effortless. “Our code was already in PyTorch, so translating everything to the AMD stack and ROCm software took one engineer roughly half a day,” says Azimian. The open-source AMD ROCm software also offers 310 AI the confidence that comes from transparency and long-term portability.

AMD engineering teams worked closely with 310 AI on orchestration, the day-to-day R&D runs where dozens of smaller models are tested to refine architecture choices. “AMD supports us on both hardware and software,” says Azimian. “Whenever we have orchestration or networking questions, AMD engineers are quick to jump in.”

“Translating everything to ROCm and the AMD stack took one engineer roughly half a day.”

Koosh Azimian, CEO and co-founder, 310 AI

SCALING PERFORMANCE AND EFFICIENCY WITH AMD INSTINCT GPUS

“AMD Instinct™ MI250X got us started,” says Azimian. “Moving to AMD Instinct MI300X and MI325X GPUs gives us the memory headroom to scale context and batch size without changing our pipeline.”

That massive amount of memory, along with up to 5.3 TB/s of memory bandwidth, allows 310 AI to train larger batches and process more samples per iteration, improving throughput and overall learning performance. “Life-science models are big and getting bigger,” says Azimian.

“AMD Instinct MI300 Series GPUs provide the memory capacity we need and deliver comparable or better speed across many workloads than what we have experienced before. Convergence happens faster without us having to grow the cluster or redesign our architecture. Our query-per-second rate improved, and the training phase stabilized sooner.”

“Moving to AMD Instinct MI300X and MI325X GPUs gives us the memory headroom to scale context and batch size without changing our pipeline.”

Koosh Azimian, CEO and co-founder, 310 AI

As its workload shifts from training to screening and discovery, AMD Instinct GPUs continue to deliver for 310 AI. “In drug design, discovery is one of the most compute-intensive stages, and AMD handles it efficiently,” Azimian says. “The large amount of memory lets us run many samples in parallel, which is critical for high-throughput inference. AMD Instinct GPUs are the more economical choice.”

310 AI pairs its AMD Instinct accelerators with AMD EPYC™ CPUs to keep utilization high and node counts low. “That combination helps us stay within a tight node budget and avoid cross-GPU bottlenecks,” Azimian says. “It’s a straightforward, efficient setup that performs reliably for our biology workloads.” The mix of generous memory and efficient single-node performance makes the platform a good fit for inference and large-scale screening, where models may need to run hundreds of millions of times to surface viable candidates.

TURNING AMD COMPUTE POWER INTO REAL-WORLD BIOLOGY

310 AI’s laboratory successes reflect the payoff from this computational efficiency. Its AI pipeline on AMD Instinct GPUs cut 150,000 candidates to 96 finalists in days.



AMD Instinct™ GPUs help 310 AI achieve up to an 84% expression rate for AI-designed proteins.

The MP4 model's designs achieved a remarkable up to 84% expression rate¹, far above the 20–30% rates typical of de novo design, confirming that most AI-generated proteins can be physically produced and remain soluble.

“Whenever we have orchestration or networking questions, AMD engineers are quick to jump in.”

Koosh Azimian, CEO and co-founder, 310 AI

All the designs were created entirely in silico without manual edits or template structures. In their GLP-1 peptide study, 310 AI scientists tested 93 AI-designed candidates. Sixty-two of 93 (67%) matched or exceeded native GLP-1 activation, and the best hit achieved up to a 19.4% improvement over native GLP-1. In contrast, traditional structure-based methods typically yield success rates of just 0.1%. “Bigger batches and larger context windows significantly improved learning speed and consistency,” Azimian says. “With the large memory on AMD Instinct™ MI300 Series GPUs, we make fewer trade-offs, iterate faster, and get candidates to the lab sooner.”

BUILDING THE NEXT GENERATION OF AI-DRIVEN SCIENCE

310 AI is already training its next model, MP5, with deeper multimodal capabilities to link chemistry, physics, and biology. The company anticipates a future in which GPUs generate synthetic data for science, just as they currently create text and images.

Azimian refers to this emerging frontier as the “gigawatt molecule-design machine,” a computational platform potentially larger and more influential than today’s language-model clusters.

Azimian is encouraged that AMD is looking to build standards and ecosystem support for physics, chemistry, and biology workloads, and advises other founders to evaluate AMD for their own AI stacks. “We recommend mixing and matching GPUs for different workloads,” Azimian said. “We find AMD to be especially economical for high-throughput discovery.”

“We find AMD to be especially economical for high-throughput discovery.”

Koosh Azimian, CEO and co-founder, 310 AI

Azimian believes 310 AI’s experience demonstrates how open software, high-memory accelerators, and responsive engineering partnerships can turn scientific ambition into working biology. With AMD Instinct MI300 Series GPUs, AMD EPYC™ server CPUs, and AMD ROCm™ software, the company is converting its AI training runs into viable drug candidates and bringing the promise of molecule programming within reach of every lab.



EXPLORE AMD INSTINCT GPUS

Sign up to receive [AMD data center communications](#).

ABOUT 310 AI

310 AI is a California-based team of scientists and technologists building the generative AI engine for programmable biology. Founded by experts from leading research institutions, biotech firms, and technology companies, 310 AI develops foundation models that translate biological language into functional proteins. Their molecular programming platform compresses billions of years of natural evolution into a computational framework that accelerates drug discovery and expands the reach of biology into domains such as materials science, agriculture, and energy. For more information visit [310.ai](#).

ABOUT AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses, and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) [website](#), [blog](#), [LinkedIn](#), and [X](#) pages.

ENDNOTES

1. 310.ai, “[AI-Generated Proteins Achieve 84% Expression Success in the Lab](#),” Nov. 18, 2024

DISCLAIMERS

All performance and cost savings claims are provided by 310 AI and have not been independently verified by AMD. Performance and cost benefits are impacted by a variety of variables. Results herein are specific to 310 AI and may not be typical. GD-181

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD products are as set forth in a signed agreement between the parties or in AMD’s Standard Terms and Conditions of Sale. GD-18u.

COPYRIGHT NOTICE

© 2026 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, CDNA, EPYC, Infinity Fabric, Instinct, ROCm, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names contained herein are for identification purposes only and may be trademarks of their respective owners. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure.