



AI Research Bootcamp Brainoid Labs

The **AI Research Bootcamp** by [Brainoid Labs](#) isn't just another course—it's a rare, no-BS gateway into the real frontier of artificial intelligence, where the world's most transformative breakthroughs are actually being forged.

Right now, most people chasing "AI" are stuck on the surface: chaining APIs, fine-tuning someone else's model, launching yet another wrapper app that looks flashy but adds almost zero fundamental value. They feel the hype, chase the LinkedIn posts about "building with GPT," and quietly sense that they're not really in the game—they're just orbiting it. Deep down, many of you already know this truth hurts: the real impact, the Nobel-level discoveries, the trillion-dollar companies of tomorrow, and the jobs that actually move civilization forward are reserved for people who have trained models from scratch, wrestled with massive compute, stared at loss curves for days, and pushed the boundaries of what's possible in pre-training, scaling laws, reinforcement learning from human feedback (RLHF), Group Sequence Policy Optimisation (GSPO) and beyond.

You've probably felt it—that quiet frustration when you realize most online tutorials stop exactly where the hard, meaningful work begins. You've seen brilliant minds publish papers that reshape entire fields while you're still debugging someone else's Hugging Face pipeline. You've wondered: "Will I ever get the chance to do real core AI research? Will I ever contribute something original instead of remixing what's already out there?"

This bootcamp exists for the moment you decide you're done watching from the sidelines.

Brainoid Labs AI Research Bootcamp is built exclusively for the rare individuals who are serious about core AI research—not hype, not wrappers, not quick MVPs, but the raw, gritty, exhilarating work of building foundation models, understanding their internals, scaling them intelligently, and pushing AI to new frontiers.

Who we are?

We are an AI research company focused on advancing real-time 3D reconstruction and building next-generation world models.

Our work sits at the intersection of robotics, perception, and language, aiming to move beyond pattern-based learning in current LLMs toward systems that understand and interact with the real world.

By integrating spatial reasoning, temporal consistency, and multimodal data, we strive to create AI that doesn't just process tokens—but develops a grounded representation of reality.

A hallmark of our approach is extreme efficiency in model design and training. Our latest generation of compact models, with only 290 million parameters, has achieved 67% accuracy on the challenging HellaSwag benchmark—a commonsense reasoning task where conventional autoregressive foundation models typically require 1.2 billion parameters or more to reach comparable levels of performance.

Who This Bootcamp Is For (And Who It Isn't)

It is intensely for

- Ambitious students and working professionals hungry to land true AI research roles at labs like OpenAI, Anthropic, DeepMind, Sarvam AI, or top-tier academic groups—and who want hands-on, portfolio-defining experience actually pre-training large models and designing RL systems.
- Startup founders and entrepreneurs are tired of depending on black-box APIs—who want to train their teams (or themselves) to create custom, proprietary foundation models that become an unfair advantage for the next generation of AI companies.
- Undergraduate and early-career researchers desperate for legitimate, real-world research experience—not toy projects, but deployments and experiments that mirror what top labs do every day.

It is **explicitly not** for:

- Anyone looking to slap a frontend on top of existing APIs and call it "an AI product." (AI Toys)
- People satisfied with surface-level tinkering or resume-padding certificates.
- People who need flashy tutorials and get a paper claiming students are experts.
- People who don't have any knowledge of 12 grade level mathematics (differentiation, partial derivatives, linear algebra).
- People who are quitters (research is not easy).

Curriculum

We have deployed our models in our platform chat.brainoidlabs.com (Rubin v1 2B)

This is one of the early models developed by brainoidlabs from scratch, trained on 800 Billion tokens, did SFT, then GRPO context length of 8k

The Bootcamp is divided into 3 volumes. (vol-1, vol-2, vol-3)

VOLUME-1

Volume-1 is for beginners who have no experience in AI model research at all, We train new models, build docker images and train on H100 GPUs with full support from brainoidlabs research team.

You will build your own custom models, we have prepared and cleaned the datasets (50 Billion tokens) and are available as open source in huggingface [@Brainoidlabs](https://huggingface.co/BrainoidLabs)

Model development is structured as a step-by-step process, where each component is built and explained in depth. The curriculum emphasizes strong mathematical foundations, including detailed derivations presented in LaTeX format, supported by visual illustrations and dedicated doubt-clearing sessions with researchers.

As part of Volume 1, we release three progressively scaled models with 150 million, 500 million, and 2 billion parameters, enabling learners to understand system design across different levels of complexity.

Each volume is released sequentially rather than all at once. This phased approach ensures a guided learning experience while also safeguarding the originality of the course content and preventing unauthorized duplication or redistribution by third parties.

VOLUME-2

Volume 2 includes a custom architecture which has more intelligence density than any open-source model of its kind. (Rubin v2 3B)

It will be a recurrent loop style custom architecture (claude mythos style).

It will solve the issue of attention sink common in most open-source models including deepseek R1.

Also includes custom kernels optimised for blackwell GPUs (TMA + NVFP4 support)

Trained with Distributed training on 16x GB200 GPUs

Build a new VLLM inspired kernel for deployment (we will deploy it live on vast.ai platform)

Rubin v1 3B is a Sparse Mixture of Expert model with 8 experts (2 activated per token)

Support context length of 128K.

VOLUME-3

Volume - 3 is focused on post training techniques including GSPO (Group Sequence Policy Optimisation)

GSPO beats GRPO on many benchmarks including SWBench, Hellaswag, MMLU
SFT (Supervised Fine tuning)

Internship Opportunities

We plan to extend internship opportunities to a select group of learners who are citizens of India.

Shortlisted candidates will be invited via email to participate in an online, MCQ-based assessment conducted after the release of all three volumes.

Internship dates are not confirmed yet and will be offered selectively, based on performance, project timelines, and internal team requirements.

Internship opportunities will only be provided after release of all 3 volumes

We recommend evaluating the program thoroughly prior to enrollment, as access includes comprehensive proprietary materials and therefore all sales are final. Should any technical issues arise with the models, our team is committed to providing timely support and resolution.