Utkarsh

Ph.D. Candidate at MIT CSAIL

Post-Training • Reinforcement Learning • AI for Science • ML Systems

utkarsh5@mit.edu • utkarsh530.github.io • github.com/utkarsh530 • Google Scholar • linkedin.com/in/utkarsh530

Research Focus: LLM post-training, reinforcement learning with verifiable feedback, generative models for scientific and physical systems, and high-performance distributed training and inference.

EDUCATION

Massachusetts Institute of Technology — **Cambridge, Massachusetts**
Ph.D., Computational Science and Engineering, CSAIL — *September 2024 – Present*

Expected completion: Spring 2027 **GPA: 5.0/5.0**
Advisors: **Prof. Alan Edelman** and **Dr. Chris Rackauckas**
Research: **LLM Post-Training, Reinforcement Learning, Scientific Machine Learning, High-Performance Computing**

Massachusetts Institute of Technology — **Cambridge, Massachusetts**
Master of Science in Computational Science and Engineering, CSAIL — *September 2022 – May 2024*

Thesis: **Automating Heterogeneous Parallelism in Numerical Differential Equations**, **GPA: 5.0/5.0**

Indian Institute of Technology Kanpur — **Kanpur, India**
*Bachelor of Technology, Double Major: **Electrical Engineering**, Chemical Engineering* — *July 2017 – April 2022*

Minor: Computer Science & Engineering with focus on **Machine Learning and Systems**; Graduation with Distinction, **GPA: 9.1/10.0 (EE Major)**

PROFESSIONAL EXPERIENCE

Developer Technology (DevTech), NVIDIA
Ph.D. Research Intern, Mentors: Dominique LaSalle, Daniel Galvez & Aurelien Chartier — *May'26 – Aug'26*

- Built high-performance RL post-training infrastructure for **Qwen3-30B-A3B, a 128-expert MoE**, in **Megatron-LM/NeMo-RL** on **8×B200**, achieving bitwise-identical rollout/training likelihoods while retaining fused MoE kernels and CUDA graphs at only **2.5% generation-throughput overhead**.
- Diagnosed and eliminated cross-engine numerical mismatch across GEMMs, MoE routing/reductions, batch-dependent kernels and CUDA-graph execution; the work resulted in **upstream Megatron-LM/NeMo-RL contributions and a follow-up research paper**.

Developer Technology (DevTech), NVIDIA – Earth-2
Research Intern, Mentors: Akshay Subramaniam & Noah Brenowitz — *May'25 – Aug'25*

- Led first-author work on **multi-resolution diffusion for global 5 km climate emulation**, unifying coarse generation and super-resolution through a Laplacian-pyramid diffusion process with a shared probability-flow ODE and exact spherical cross-resolution operators.
- Built a single model that generates coarse and fine atmospheric structure jointly, reproducing spectra, field statistics and extreme precipitation across **12 atmospheric channels**; the internship led to a **first-author research manuscript** on kilometer-scale climate emulation.

AWS AI Labs, Amazon Science
Applied Scientist Intern, Mentors: Danielle Robinson, Bernie Wang & Michael W. Mahoney — *June'24 – Aug'24*

- Developed a **Differentiable Probabilistic Projection Layer** that maps learned distributions exactly onto hard constraint manifolds; built the custom batched optimization solver and KKT-based implicit differentiation end-to-end. *ICLR 2026.*
- Extended the framework to nonlinear constraints for PDE learning and probabilistic forecasting, achieving up to **15× lower mean-forecast MSE** and **2.5× better uncertainty estimates** than baseline methods, with a sampling-free formulation that reduced training overhead by **3–5×**.

Julia Lab, MIT CSAIL
Research Assistant, PI(s): Dr. Alan Edelman, Dr. Chris Rackauckas — *Sept'22 – Present*

- Led **GEOPROBE**, **up to 200× faster than traditional CFD**, an accurate geometry-aware neural surrogate for million-point 3-D high-speed CFD that transfers across vehicle geometries rather than requiring a new model per shape. Built the architecture, uncertainty model, and differentiable formulation that enforces **compressible mass conservation exactly by construction**.
- Built award-winning **DiffEqGPU.jl**, turning high-level differential-equation models into specialized GPU solver kernels with stiff integration and automatic differentiation across NVIDIA, AMD, Intel and Apple accelerators; achieves **20–100× speedups over JAX** and **100–200× over PyTorch** in fixed-step ensemble ODE benchmarks.
- Developed **Physics-Constrained Flow Matching (PCFM)**, a zero-shot sampling method that imposes hard nonlinear physical constraints on pretrained generative models **without retraining**; demonstrated on PDEs with conservation laws, shocks and discontinuities. *NeurIPS 2025.*
- **Post-trained LLMs to write executable PDE solver code using physics as the reward.** A single **Qwen2.5-Coder-7B** policy learned numerical strategies across multiple PDE families and transferred to held-out equations, showing out-of-distribution generalization beyond equation-specific training. *Under review, NeurIPS 2026.*

- Developed **parallel explicit/implicit extrapolation methods for stiff systems** and **GPU-accelerated inverse-problem optimization**.

SELECTED RESEARCH

230+ citations across publications and preprints. Full list: Google Scholar.

- **Achieving High Performance in True On-Policy Reinforcement Learning for Mixture-of-Experts LLMs**
  Utkarsh, et al. *Under review, NeurIPS 2026; NVIDIA*
- **HyperHadamard: Layout-Aware Hadamard Transforms on Tensor Cores**
  Utkarsh *Sole author; under review, ICLR 2027*
- **Reinforcement Learning with Verifiable Physics (RLVP): Post-training LLMs with Continuous Rewards**
  Pengfei Cai*, **Utkarsh***, Alan Edelman, Christopher V. Rackauckas, Rafael Gómez-Bombarelli *Under review, NeurIPS 2026*
- **Multi-Resolution Laplacian Diffusion Models for Kilometer-Scale Climate Emulation**
  Utkarsh, Akshay Subramaniam, Noah Brenowitz *Under review, ICLR 2027; NVIDIA Earth-2*
- **Physics-Constrained Flow Matching: Sampling Generative Models with Hard Constraints**
  Utkarsh*, Pengfei Cai*, Alan Edelman, Rafael Gómez-Bombarelli, Christopher Rackauckas *NeurIPS 2025*
- **SNAP-FM: Sparse Nonlinear Accelerated Projection for Physics-Constrained Generative Modeling**
  A. Kolli*, T. Xenakis*, **Utkarsh***, ..., A. Edelman, C. Rackauckas *IEEE HPEC 2026 – Best Paper Award*
- **GEOPROBE: Geometry-Aware Probabilistic Neural Surrogates for High-Speed Aerodynamic Flows**
  Utkarsh, Steven Spreizer, Matthew C. Jones, Chris Rackauckas, Alan Edelman *To appear, AIAA SciTech 2027*
- **End-to-End Probabilistic Framework for Learning with Hard Constraints**
  Utkarsh, Danielle C. Maddix, Ruijun Ma, Michael W. Mahoney, Yuyang Wang *ICLR 2026*
- **Automated Translation and Accelerated Solving of Differential Equations on Multiple GPU Platforms**
  Utkarsh, et al. *Computer Methods in Applied Mechanics and Engineering, 2024*

OPEN-SOURCE & ML SYSTEMS

- **Megatron-LM / NeMo-RL:** Upstream implementations for **batch-invariant MoE inference, exact RL rollout/training execution, CUDA-graph generation, and training–generation consistency**.
- **HyperHadamard:** Sole-author tensor-core FWHT design reaching **94–97% of measured device-copy bandwidth** and up to **1.59**× the throughput of prior tensor-core implementations; integrated with vLLM.
- **DiffEqGPU.jl:** Lead developer of the **300+ star** SciML GPU solver package; **20–100**× **faster than JAX** and **100–200**× **over PyTorch** in published ensemble benchmarks, portable across NVIDIA, AMD, Intel and Apple GPUs, and used in pharmaceutical, aerospace and physics workflows.

TECHNICAL SKILLS

- **LLM Post-Training:** RL / RLVR, GRPO, SFT, Verifiable Rewards, Rollout Generation, Model Evaluation
- **Generative Models:** Diffusion, Flow Matching, Test-Time Optimization, Constrained Generation
- **ML Systems:** Megatron-LM, NeMo-RL, vLLM, PyTorch, MoE, Distributed Training/Inference, CUDA Graphs
- **Performance / Programming:** CUDA, Triton, GPU Kernels, Tensor Cores, Python, C/C++, Julia, Parallel Computing

SELECTED PREPRINTS / WORKSHOP PUBLICATIONS

- **On the Parisi PDE and the Overlap Gap Property in Ising Spin Glasses**
  Utkarsh, Theodoros Xenakis
- **Enforcing Constraints in Molecular and Crystalline Generative Models via Physics-Constrained Flow Matching**
  Pengfei Cai, **Utkarsh**, ..., Chris Rackauckas, Rafael Gómez-Bombarelli *ICLR 2026 Workshop on AI4Mat*
- **Developing Surrogates for Epidemic Agent-Based Models via Scientific Machine Learning**
  Sharv Murgai, **Utkarsh**, ..., Chris Rackauckas *NeurIPS 2025 Workshop on ML4PS*
- **Efficient GPU-Accelerated Global Optimization for Inverse Problems**
  Utkarsh, Vaibhav Kumar Dixit, Julian Samaroo, Avik Pal, Alan Edelman, Chris Rackauckas *ICLR 2024 Workshop on AI4DifferentialEquations in Science*

PROFESSIONAL SERVICE & AWARDS

- IEEE HPEC Outstanding Paper Award (2026)
- Reviewer, NeurIPS (2026)
- Reviewer, The International Conference on Learning Representations (ICLR) (2025, 2026)
- Oral Presentation, Amazon Machine Learning Conference (2025)
- SIAM BGCE Best Student Paper Award Finalist (2025)
- Talk: Efficient Numerical Methods for SciML, MIT Beaverworks (2024)
- Mentor, WIDS Cambridge Datathon (2024, 2025)