

Avik Pal

PH.D. CANDIDATE · MASSACHUSETTS INSTITUTE OF TECHNOLOGY · AI COMPILERS · AI FOR SCIENCE

✉ avikpal@mit.edu | 🏠 <https://avik-pal.github.io> | 🌐 [avik-pal](https://avik-pal.github.io) | 💬 avikpal1410 | 🐦 [@avikpal1410](https://avikpal1410) | 🎓 Avik Pal

Education

Massachusetts Institute of Technology

Cambridge, MA

S.M. & PH.D. IN COMPUTER SCIENCE (TRACK: AI & DECISION MAKING), GPA: **4.9 / 5.0**

Sep '21 – Feb '26 (est.)

- Advisors: [Dr. Alan Edelman](#), & [Dr. Chris Rackauckas](#)
- PhD Thesis: *Performance Portability through Composable Compiler Abstractions*
- Masters' Thesis: *On Efficient Training and Inference of Neural Differential Equations*

Indian Institute of Technology Kanpur

Kanpur, India

BACHELOR OF TECHNOLOGY (B.TECH.) IN COMPUTER SCIENCE AND ENGINEERING, GPA: **9.9 / 10.0**

Jul '17 – May '21

Work Experience

Google Cloud, XLA TPU Performance Optimization | Student Researcher

New York City, NY

MANAGER: [DR. MEHRDAD KHANI](#)

Jun - Aug '25

- Designed learned cost models as an alternative to expert-designed heuristic cost models typically used in AI compilers.
- Built a Mixture of Experts (MoE) model achieving under **3% runtime prediction error** for fusions. Optimized MoE inference for batched predictions, **reducing runtime to 1s** on mid-tier CPUs.
- Improved **end-to-end TPU model performance by 1%**, realizing **over 70% of available headroom**. Integrated and deployed the learned cost model in the XLA:TPU compiler tool-chain.

MIT CSAIL, Julia Lab | Graduate Research Assistant

Cambridge, MA

PI(s): [DR. ALAN EDELMAN](#) & [DR. CHRIS RACKAUCKAS](#)

Sep '21 - Present

- Developed an MLIR-based compiler enabling automatic differentiation and tensor program optimizations for scientific computing and deep learning.
- Deployed the compiler to scale ocean simulation codes from single-node execution to **1700 TPUs** and **6144 A100 CUDA GPUs** with automated parallelization and communication optimizations.
- Introduced a framework for **enforcing physical constraints in neural surrogates of hypersonic flows** (*AIAA SciTech*), made scalable through tensor optimizations. This technology is being transitioned for use inside **US Department of Defense (DoD)**.
- Created a nonlinear root-finding framework (*ACM TOMS*) (with over **10_000 unique monthly downloads**) that **outperforms PETSc and Sundials**, and applied it to training **constrained neural networks** achieving faster runtimes and improved numerical convergence (tolerances up to 10^{-10} vs. 10^{-8} in prior work).

Intel Labs, Parallel Computing Lab | Graduate Research Intern

Santa Clara, CA

MANAGER: [DR. DIPANKAR DAS](#)

May - Aug '24

- Explored Kolmogorov–Arnold Networks (KANs) as a **parameter-efficient alternative to MLPs** in scientific ML and vision tasks.
- Developed a systematic evaluation of KANs across Neural ODEs, computer vision, and operator learning benchmarks. Introduced Hessian eigenvalue spectrum analysis to study convergence dynamics.
- Quantified computational trade-offs providing insights into the scalability and practical applicability of KANs.
- Publication: Convergence Dynamics and Eigenvalue Analysis of B-Spline KANs (*NeurIPS 2025 Workshop on Science for Deep Learning*).

Google AI | Student Researcher

Mountain View, CA

MANAGER(S): [DR. ANDREY ZHMOGINOV](#), & [DR. LILY HU](#)

May - Aug '22

- Proposed a deep learning method to **augment coarse-grained simulations for approximating fine-grained wildfire dynamics**.
- Showed that black-box neural simulators fail to capture stable long-term dynamics in limited-data regimes.
- Developed probabilistic neural simulators that model the full range of possible wildfire outcomes.
- Released open-source **wildfire simulator based on percolation models in JAX**.

University of Toronto & Vector Institute | Research Intern

Toronto, CAN

PI: [DR. SANJA FIDLER](#)

Jan - Dec '20

- Developed multi-agent reinforcement learning (MARL) environments where **driving rules emerge naturally** from optimizing traffic flow.
- Analyzed how POMDP design factors, such as perception noise and agent density, shape cooperative driving behaviors.
- Demonstrated emergent behaviors including lane following, right of way, fast lanes, communication, and safe distances.
- Released an **open-source suite of driving environments** to support MARL research in self-driving.
- Publication: Emergent Road Rules in Multi-Agent Driving Environments (*ICLR 2021*).

Selected Publications

550+ citations across all publications and pre-prints. For a complete list of publications check [my Google Scholar profile](#).

Peer-Reviewed Conference Proceedings / Journal Papers

Geometry & Mesh Invariant Neural Surrogates for Hypersonic Flows

[AVIK PAL](#), [ALAN EDELMAN](#), [CHRIS RACKAUCKAS](#), MATHEW C. JONES, STEVEN SPRIEZER, & TYLER E. KORENYI-BOTH

2025

AIAA SciTech Forum (Accepted, To Appear)

NonlinearSolve.jl: High-Performance and Robust Solvers for Systems of Nonlinear Equations

AVIK PAL, FLEMMING HOLTORF, AXEL LARSSON, TORKEL LOMAN, ..., ALAN EDELMAN, & CHRIS RACKAUCKAS



2025

ACM Transactions on Mathematical Software (TOMS)

Locally Regularized Neural Differential Equations: Some Black Boxes Were Meant to Remain Closed!

AVIK PAL, ALAN EDELMAN, & CHRIS RACKAUCKAS



2023

International Conference on Machine Learning (ICML)

Continuous Deep Equilibrium Models: Training Neural ODEs Faster by Integrating Them to Infinity

AVIK PAL, ALAN EDELMAN, & CHRIS RACKAUCKAS



2023

IEEE High Performance Extreme Computing (HPEC) (Best Student Paper Award)

Opening the Blackbox: Accelerating Neural DEs by Regularizing Internal Solver Heuristics

AVIK PAL, YINGBO MA, VIRAL B. SHAH, & CHRIS RACKAUCKAS



2021

International Conference on Machine Learning (ICML)

Pre-prints / Peer-Reviewed Workshop Publications

Making Waves in the Cloud: A Paradigm-Shift for Scientific Computing and Ocean Modeling through Compiler Technology

WILLIAM S. MOSES, MOSÈ GIORDANO, AVIK PAL, ..., ALBERT COHEN, & OLEKSANDR ZINENKO



2025

Under Review

Semi-Explicit Neural DAEs: Learning Long-Horizon Dynamical Systems with Algebraic Constraints

AVIK PAL, ALAN EDELMAN, & CHRIS RACKAUCKAS



2025

Under Review

Understanding the Limitations of KANs: Convergence Dynamics and Computational Efficiency

AVIK PAL, & DIPANKAR DAS



2024

NeurIPS Workshop on Science for Deep Learning

Differentiable Programming for Differential Equations: A Review

FACUNDO SAPIENZA, JORDI BOLIBAR, ..., AVIK PAL, ..., PER-OLOF PERSSON, & CHRIS RACKAUCKAS



2024

Under Review at SIAM Review

Open Source Software

For a complete list of open source software I have worked on, check [my GitHub profile](#).

	Lux.jl, Elegant and performant deep learning in Julia powered by the XLA compiler.	Julia
	Enzyme-JAX, MLIR-based compiler enabling automatic differentiation, sharded-communication optimizations, and various platform-agnostic high-level optimizations.	C++, MLIR, Jax, XLA
	Reactant.jl, Optimize Julia functions with MLIR and XLA for execution on heterogeneous hardware.	Julia, C++
	NonlinearSolve.jl, High-performance and differentiation-enabled nonlinear solvers with sparsity support.	Julia
	NeuralPDE.jl & NeuralOperators.jl, PINN solvers of PDEs for accelerated simulation.	Julia
	torchgan, Research framework for easy and efficient training of GANs based on Pytorch.	Pytorch, Python
	social-driving, Multi-agent environments and reward functions for social driving behavior emergence.	Pytorch, Python

Selected Presentations

2025 Accelerating Machine Learning in Julia using Lux & Reactant, JuliaCon / CSAIL Alliances Workshop



2025 Semi-Explicit Neural DAEs: Learning Long Horizon Constrained Dynamical Systems, SIAM CSE



2025 The Tricks Required for Scientific Machine Learning to Work on Real Data, SIAM CSE



2025 Accelerating Physics Informed Machine Learning in Julia using Reactant and Lux, AAAI



2022 Lux.jl: Explicit Parameterization of Neural Networks in Julia, JuliaCon



2022 Mixing Implicit and Explicit Deep Learning with Skip DEQs, SciMLCon



2019 Differentiable Rendering and its Applications in Deep Learning, JuliaCon



Professional Activities

'24, '25 International Conference on Learning Representations (ICLR), Reviewer for 6 + 3 papers

'23 - '25 Neural Information Processing Systems (NeurIPS), Reviewer for 6 + 6 + 5 papers

'22 International Conference on Machine Learning (ICML), Reviewer for 3 papers

Honors, Awards & Recognitions

2023 Best Student Paper Award, IEEE High Performance Extreme Computing Conference



2017-20 Academic Excellence, Top 10% students in 3 Consecutive Academic Year

IIT Kanpur

2017 Inspire Scholarship for Higher Education, Top 1% students in 10+2 board results

