Avik **Pal**

Ph.D. Student · Massachusetts Institute of Technology

avikpal@mit.edu | https://avik-pal.github.io | avik-pal | avikpal1410 | @avikpal1410 | Avik Pal

Education

Massachusetts Institute of Technology — *Cambridge, MA, USA*

S.M. & Ph.D. in Electrical Engineering and Computer Science, GPA: **4.9 / 5.0** — *Sept. '21 – Dec '25 (Expected)*

- Advisors: Dr. Alan Edelman, & Dr. Christopher Rackauckas
- Masters' Thesis: *On Efficient Training and Inference of Neural Differential Equations*

Indian Institute of Technology Kanpur — *Kanpur, UP, India*

Bachelor of Technology (B.Tech.) in Computer Science and Engineering, GPA: **9.9 / 10.0** — *July '17 – May '21*

Publications

357 citations across all publications and preprints.

Conference Proceedings

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) (Oral Presentation) **(Best Student Paper Award)**

Opening the Blackbox: Accelerating Neural Differential Equations by Regularizing Internal Solver Heuristics

Avik Pal, Yingbo Ma, Viral B. Shah, & Chris Rackauckas — *2021*

International Conference on Machine Learning (ICML)

Emergent Road Rules In Multi-Agent Driving Environments

Avik Pal, Jonah Philion, Andrew Liao, & Sanja Fidler — *2021*

International Conference on Learning Representations (ICLR)

Stably Accelerating Stiff Quantitative Systems Pharmacology Models: Continuous-Time Echo State Networks as Implicit Machine Learning

Ranjan Anantharaman, Anas Abdelrehim, Anand Jain, Avik Pal, Danny Sharp, Utkarsh, & Chris Rackauckas — *2021*

International Federation of Automatic Control (IFAC)

Composing Modeling and Simulation with Machine Learning in Julia

Chris Rackauckas, Ranjan Anantharaman, Alan Edelman, Shashi Gowda, Maja Gwozdz, Anand Jain, Chris Laughman, Yingbo Ma, Francesco Martinuzzi, Avik Pal, Utkarsh, Elliot Saba, & Viral B. Shah — *2021*

International Modelica Conference

TorchGAN: A Flexible Framework for GAN Training and Evaluation

Avik Pal, & Aniket Das — *2021*

Journal of Open Source Software (JOSS)

RayTracer.jl: A Differentiable Renderer that supports Parameter Optimization for Scene Reconstruction

Avik Pal — *2019*

Proceedings of the JuliaCon Conferences

Workshop Papers

Efficient GPU-Accelerated Global Optimization for Inverse Problems

Utkarsh, Vaibhav K. Dixit, Julian Samaroo, Avik Pal, Alan Edelman, & Chris Rackauckas — *2024*

ICLR Workshop on AI4DifferentialEquations in Science

Humor@IITK at SemEval-2021 Task 7: Language Models for Quantifying Humor And Offensiveness

Avik Pal, Aishwarya Gupta, Bholeshwar Khurana, Lakshay Tyagi, & Ashutosh Modi — *2021*

Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval)

Fashionable Modelling with Flux

Mike Innes, Elliot Saba, Keno Fischer, Dhairya Gandhi, M.C. Rudilosso, Neethu Mariya Joy, Tejan Karmali, Avik Pal, & Viral B. Shah — *2019*

NeurIPS Workshop on Systems for Machine Learning

Preprints

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

Avik Pal, Flemming Holtorf, Axel Larsson, Torkel Loman, Utkarsh, Frank Schaefer, Qingyu Qu, Alan Edelman, & Chris Rackauckas — *2024*

In Submission at SIAM Journal on Scientific Computing (SISC)

Work Experience

Research Experience

Julia Lab, MIT, Graduate Research Assistant — *Cambridge, MA, USA*

PI(s): Dr. Alan Edelman & Dr. Chris Rackauckas — *Sept. '21 - Present*

- Developing a Generalized Framework for Embedding Equality and Inequality Constraints in Neural Networks
- Working on Sparse Automatic Differentiation for fast Boundary Value Problems and Complementarity Problems
- Generalized the framework of infinitely deep discrete neural networks to the continuous domain (Continuous DEQs)
- Improved Training and Prediction Time of Deep Equilibrium Models by over 2x using Predictor-Corrector Methods

Google AI, Student Researcher — *Mountain View, CA, USA*

Manager(s): Dr. Andrey Zhmoginov, & Dr. Lily Hu — *May - Aug. '22*

- Proposed a novel deep learning method to augment fast coarse-grained simulations to approximate slow fine-grained simulations
- Demonstrated that Black-Box Neural Simulators cannot capture stable long-term dynamics
- Developed Probabilistic Neural Simulators to capture all possible wildfire dynamics
- Code: Wildfire Simulator in JAX

University of Toronto & Vector Institute, Research Intern — *Toronto, ON, CAN*

PI: Dr. Sanja Fidler — *Jan. - Dec. '20*

- Worked on Multi Agent Reinforcement Learning for Emergence of Social Driving Rules
- Developed a suite of driving environments with the intention of stimulating interest within the MARL community to solve fundamental self-driving problems
- Analyzed what choices in the definition of the Partially Observed Markov Decision Process leads to the emergence of social driving rules
- Proposed metrics for empirical analysis of emergent driving behavior
- Publication: **Emergent Road Rules In Multi-Agent Driving Environments**

Engineering Experience

Julia Computing, Engineering Simulation Intern — *Remote*

Manager: Dr. Christopher Rackauckas — *Jan. '21 - July. '21*

- Researching Machine Learning Models for Surrogatizing Complex Physical Models
- Developed Reservoir Computing Models to deal with Stiff Physical Systems
- Publications(s): **Stably Accelerating Stiff Quantitative Systems Pharmacology Models: Continuous-Time Echo State Networks as Implicit Machine Learning**, & **Composing Modeling and Simulation with Machine Learning in Julia**

JuliaLang, Google Summer of Code Participant — *Remote*

Mentor: Mike Innes — *Summer '18, '19*

- Worked on the development of a Deep Learning Framework Flux.jl in Julia
- Experimented with Computer Vision Models and developed a clean way to use them in Flux through model-zoo and Metalhead.jl
- Developed a library, RayTracer.jl for differentiable rendering. We use raytracing as the underlying algorithm for the renderer and use source-to-source automatic differentiation for computing the derivatives
- Assisted in the creation of a differentiable simulator, Duckietown.jl, for training and testing the self-driving car
- Developed an algorithm to train agents in the Duckietown environment by differentiating through the 3D renderer
- Integrated the backend with CuDNN which provided a 17x performance boost on GPUs

Open Source Software

7561 github stars across significant open-source projects

2018	**FluxML/Flux.jl**, Relax! Flux is the ML library that doesn't make you tensor.	*4376* ★	</>
2018	**torchgan/torchgan**, Research Framework for easy and efficient training of GANs based on Pytorch.	*1377* ★	</>
2022	**LuxDL/Lux.jl**, Explicitly Parameterized Neural Networks in Julia.	*409* ★	</>
2021	**SciML/SciMLSensitivity.jl**, A component of the DiffEq ecosystem for enabling sensitivity analysis for scientific machine learning (SciML). Optimize-then-discretize, discretize-then-optimize, adjoint methods, and more for ODEs, SDEs, DDEs, DAEs, etc.	*305* ★	</>
2018	**FluxML/NNlib.jl**, Neural Network primitives with multiple backends.	*189* ★	</>
2023	**SciML/NonlinearSolve.jl**, High-performance and differentiation-enabled nonlinear solvers (Newton methods), bracketed rootfinding (bisection, Falsi), with sparsity and Newton-Krylov support.	*183* ★	</>
2019	**avik-pal/RayTracer.jl**, Differentiable RayTracing in Julia.	*146* ★	</>
2020	**fidler-lab/social-driving**, Design multi-agent environments and simple reward functions such that social driving behavior emerges.	*134* ★	</>
2020	**FluxML/Flux3D.jl**, 3D computer vision library in Julia.	*100* ★	</>
2021	**avik-pal/Wandb.jl**, Unofficial Julia bindings for logging experiments to wandb.ai.	*74* ★	</>
2023	**SciML/SimpleNonlinearSolve.jl**, Fast and simple nonlinear solvers for the SciML common interface. Newton, Broyden, Bisection, Falsi, and more rootfinders on a standard interface.	*59* ★	</>
2021	**avik-pal/FluxMPI.jl**, Distributed Data Parallel Training of Deep Neural Networks.	*55* ★	</>

Teaching

18337J: Parallel Computing and Scientific Machine Learning — *MIT*

Teaching Assistant — *Spring 2023*

High-Performance Computing and Scientific Machine Learning course primarily geared toward graduate students

CS633: Parallel Computing — *IIT Kanpur*

Teaching Assistant — *Spring 2021*

Parallel Computing Course for advanced undergraduates and graduate students

CS771: Introduction to Machine Learning — *IIT Kanpur*

Teaching Assistant — *Fall 2020*

Introductory Machine Learning Course designed primarily for 3^{rd}-year undergraduates and 1^{st} year graduate students

Research Talks

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

2023	**Neural Information Processing Systems (NeurIPS)**, Reviewer for 6 papers
2022	**International Conference on Machine Learning (ICML)**, Reviewer for 3 papers

Honors and Awards

2023	**Best Student Paper Award**, IEEE High Performance Extreme Computing Conference	*USA*
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	*India*
2017	**National Standard Examination in Astronomy**, Qualified for Indian National Astronomy Olympiad	*India*