

Sourav Pal

✉ spal9@wisc.edu ☎ +1-608-772-4907 🏠 sourav-roni.github.io 🌐 [sourav-roni](#) 🎓 [Scholar](#) [in](#) [souravpalkgp](#)

Research Interests

My research interfaces with scientific computing, representation learning and privacy in Machine Learning. My work includes long sequence modeling, guidance for diffusion models, adaptive sampling and training scalable foundational models for healthcare data. It enables efficient deep learning through interpretable neural networks, with applications in large-scale models like LLMs and VLMs.

Education

University of Wisconsin-Madison

MS-PhD in Computer Science

Minor in Mathematics

Aug 2019 - Nov 2025 (expected)

Advisor: Prof. Vikas Singh

Indian Institute of Technology (IIT) Kharagpur

B.Tech. (Hons.) in Computer Science and Engineering

Jul 2014 - May 2018

CGPA: 9.33/10.00

Work Experience

Adobe, Noida

Jul'18 - Jul'19

Member of Technical Staff

- Developed voice-first interface for Acrobat Reader, enabling interaction via Amazon Alexa, and prototyped ML-driven features for extracting insights from documents-awarded 1st place at Adobe Document Cloud HackWeek.
- Led migration of Acrobat Reader app from ndk-build to CMake, and drove app-size optimization initiatives reducing APK size by 40%, resulting in measurable increase in monthly active users (MAU) in pilot deployment.

Internships

Microsoft Research, Redmond

Jun'22 - Aug'22

Research Intern

- Developed bilevel optimization framework for synthetic data generation, learning cut-and-paste strategies by optimizing object placement and composition over diverse backgrounds using gradient-based supervision.
- Achieved 10–15% mAP improvement on low-resource classes compared to heuristic baselines, enhancing generalization under domain shifts and challenging compositions.

Adobe, San Jose

May'21 - Aug'21

Data Science Intern

- Designed and implemented an end-to-end controlled text summarization system with fine-grained control over content, style, and tone, leveraging transformer-based encoder-decoder architectures (BART, T5).
- Improved summary fidelity and control by 15% over baselines on standard metrics (ROUGE, BLEU), enabling dynamic generation for marketing and product content applications.

Adobe Research, Bengaluru

May'17 - Jul'17

Research Intern

- Developed deep learning models for predicting visual saliency on mobile UIs using eye-gaze tracking data, integrating computer vision and attention modeling to improve interface usability.
- Deployed a large-scale crowdsourcing pipeline on Amazon Mechanical Turk to collect gaze data, designed novel evaluation metrics for saliency prediction on mobile-specific layouts, enabling scalable benchmarking.

Publications

- **Sourav Pal**, Kamyar Azizzadenesheli, Vikas Singh; “*PINNs with learnable quadrature.*” [In Review]
- Xizheng Yu, Justin Torok, Sneha Pandya, **Sourav Pal**, Vikas Singh, Ashish Raj; “*Brain-wide interpolation and conditioning of gene expression in the human brain using Implicit Neural Representations.*” [In Review]
- **Sourav Pal**, Harshavardhan Adepur, Clinton Wang, Polina Golland, Vikas Singh; “*Implicit Representations via Operator Learning.*” [ICML 24]
- Jurijs Nazarovs, Zhichun Huang, Xingjian Zhen, **Sourav Pal**, Rudrasis Chakraborty, Vikas Singh; “*Variational Sampling of Temporal Trajectories.*” [arXiv 24]
- **Sourav Pal**, Zhanpeng Zeng, Sathya N. Ravi, Vikas Singh; “*Controlled Differential Equations on Long Sequences via Non-standard Wavelets.*” [ICML 23]
- Zhanpeng Zeng, **Sourav Pal**, Jeffery Kline, Glenn Fung, Vikas Singh; “*Multi Resolution Analysis (MRA) for Approximate Self-Attention.*” [ICML 22]
- Ronak Mehta*, **Sourav Pal***, Vikas Singh, Sathya N. Ravi; “*Deep Unlearning via Randomized Conditionally Independent Hessians.*” [CVPR 22]
- Gupta Prakhar, **Sourav Pal***, Shubh Gupta*, Ajaykrishnan Jayagopal*, and Ritwik Sinha; “*Saliency Prediction for Mobile User Interfaces.*” [WACV 18]
- **Sourav Pal***, Tharun Mohandoss*, Pabitra Mitra; “*Visual Attention for Behavioral Cloning in Autonomous Driving.*” [ICMV 18]
- Ankan Mullick, **Sourav Pal**, Projjal Chanda, Arijit Panigrahy, Anurag Bharadwaj and Siddhant Singh; “*D-FJ: Deep Neural Network Based Factuality Judgment.*” [SIGKDD-W 19]

* Authors contributed equally.

Patents

- Prakhar Gupta, Sourav Pal, Shubh Gupta, Ritwik Sinha and Ajaykrishnan Jayagopal; “*Saliency Prediction for a Mobile User Interface.*” [US Patent 10,664,999]
- Prakhar Gupta, Shubh Gupta, Ritwik Sinha, Sourav Pal and Ajaykrishnan Jayagopal; “*Saliency Prediction for Informational Documents.*” [US Patent 11,263,470]

Ongoing Research

- **Efficient Foundation Model for High-Dimensional Brain Imaging:** Developed an architecture for scalable foundation model for fMRI data, incorporating tensor-product, wavelets and operator learning to enable efficient training on long-sequence, high-dimensional neuro-imaging signals. When fine-tuned, it achieves state-of-the-art performance on multiple downstream clinical and cognitive tasks, with applicability to real-world healthcare AI.
- **Physics-Guidance for Multi-Modal Diffusion Models:** Designed a theoretically grounded guidance strategy for text-to-image diffusion models inspired by PDEs from soft-matter physics. Current implementation using Stable Diffusion XL improves multi-modal generative quality and controllability, achieving competitive FID, CLIP-score, and LPIPS performance on benchmark text-guided image synthesis tasks.

Technical Skills

Programming Languages: Python, Matlab, C++, HTML, Bash, Julia

Tools & Frameworks: PyTorch, Lightning, Jax, Keras, Numpy, Scikit-learn, OpenCV, Hugging Face, LangChain, L^AT_EX, Git

Relevant Coursework

Mathematical Foundation of Machine Learning	Methods of Applied Mathematics	Real & Complex Analysis
Learning Based Computer Vision	Natural Language Processing	Non-Linear Optimization
Methods of Computational Mathematics	Partial Differential Equations	Measure Theory

Reviewing

NeurIPS 25, ICML 25, ICLR 25, NeurIPS 24, ICML 24, ICLR 24, NeurIPS 23, ICML 22, CVPR 22, WACV 22

Talks

- ML+X SEMINAR SERIES: Modeling long-sequence dynamics using differential equations and wavelets.
- COMPUTER VISION ROUND TABLE, UW–Madison: Efficient machine unlearning in deep neural networks.
- ADOBE TECH SUMMIT 2019: Deep learning for visual saliency prediction in mobile user interfaces.