Manish Nagaraj
http://linkedin.com/in/m-nagaraj
Email : mnagara@purdue.edu
Mobile : +1-765-701-7970

Objective: *A PhD candidate with extensive expertise in data efficiency and optimization techniques, focusing on creating compact, accurate, and scalable machine learning systems. Passionate about reducing computational overhead while enhancing system performance for real-world applications.*

Education

- **Purdue University** — West Lafayette, IN
  PhD in Electrical & Computer Engineering; GPA: 3.66 — *May 2019 - Present*
  - **Relevant Areas**: Data Efficiency, Privacy-Preserving Machine Learning, Optimization Techniques
  - **Research Advisor**: Prof. Kaushik Roy
- **Purdue University** — West Lafayette, IN
  MS in Electrical & Computer Engineering; GPA: 3.72 — *Aug 2017 - May 2019*
  - **Thesis**: Energy Efficient Byzantine Agreement Protocols for Cyber-Physical Resilience
- **PES Institute of Technology and Science** — Bangalore, India
  Bachelor of Engineering in Electronics and Communications; GPA: 9.77/10.0 — *Aug 2013 - May 2017*

Work Experience

- **Research Intern, Integrated Systems Team** — Skillman, NJ
  Latent AI — *May 2023 - August 2023*
  - Built data annotation tools to streamline evaluation processes for anomaly detection frameworks.
  - Designed scalable systems to handle noisy datasets efficiently and ensured robust model performance.
  - Improved energy and latency metrics of internal tools for optimized machine learning workflows.

Research Projects

- **Exploring Data Efficiency in Deep Learning Systems** — West Lafayette, IN
  Ph.D. Dissertation, Purdue University — *May 2019 - Present*
  - **Efficient Fine-tuning of LLMs Using Impact-Driven Data Selection**:
    - Designing a lightweight, correlation-based data selection pipeline for fine-tuning LLMs under tight memory and compute constraints.
    - Eliminates the need for computationally expensive gradient/Hessian-based methods, enabling on-device personalization.
    - Ensures robustness to hyperparameter changes and transferability across model scales (e.g., LLaMA-3 8B $\rightarrow$ 70B) using proxy-based scoring.
    - Prioritizes high-impact samples while filtering redundant or harmful data, optimizing fine-tuning efficiency for edge deployment.
    - *Work in progress*
  - **Finding the Muses: Identifying Coresets Through Loss Trajectories**:
    - Proposing *Loss Trajectory Correlation (LTC)*, a scalable and efficient metric to identify influential training samples for coreset construction.
    - Achieves state-of-the-art performance on CIFAR-100 and ImageNet-1k across various subset sizes and architectures (ResNet, VGG, Swin Transformer), with $< 2\%$ performance drop in cross-architecture transfer.
    - Offers significant reduction in computational and memory overhead compared to gradient-based methods.
    - *Manuscript under peer review.*

- **DOTIE: Energy-Efficient Object Detection Using Event Cameras**
  - Proposed a novel lightweight detection framework leveraging event-driven camera data to minimize energy consumption.
  - Surpassed conventional detection techniques in accuracy, efficiency, and latency metrics.
  - Demonstrated at **ICRA 2023** and **CVPR 2023 Workshops**.
- **TOFU: Federated Learning with Data and Communication Efficiency**
  - Developed an innovative federated learning framework to reduce communication overhead by 10×.
  - Ensured privacy-preserving data sharing via encoded weight updates, validated for robustness against inversion attacks.
  - Published at **IEEE Access 2024**.

Selected Publications

- **Nagaraj, Manish**, Deepak Ravikumar, Efstathia Soufleri, and Kaushik Roy. Finding the muses: Identifying coresets through loss trajectories. *arXiv preprint arXiv:2503.09721*, 2025
- **Manish Nagaraj**, Chamika Mihiranga Liyanagedera, and Kaushik Roy. DOTIE - detecting objects through temporal isolation of events using a spiking architecture. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4858–4864, 2023
- Arjun Roy, **Manish Nagaraj**, Chamika Mihiranga Liyanagedera, and Kaushik Roy. Live demonstration: Real-time event-based speed detection using spiking neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4080–4081, 2023
- **Nagaraj, Manish**, Isha Garg, and Kaushik Roy. Tofu: Towards obfuscated federated updates by encoding weight updates into gradients from proxy data. *IEEE Access*, 2024
- Chamika Mihiranga Liyanagedera, **Manish Nagaraj**, Wachirawit Ponghiran, and Kaushik Roy. Low-power real-time sequential processing with spiking neural networks. In *2023 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2023
- **Manish Nagaraj**. *Energy Efficient Byzantine Agreement Protocols for Cyber Physical Resilience*. PhD thesis, Purdue University Graduate School, 2019
- Amogh Joshi, Wachirawit Ponghiran, Adarsh Kosta, **Nagaraj, Manish**, and Kaushik Roy. Fedora: A flying event dataset for reactive behavior. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5859–5866. IEEE, 2024
- Timur Ibrayev, **Nagaraj, Manish**, Amitangshu Mukherjee, and Kaushik Roy. Exploring foveation and saccade for improved weakly-supervised localization. In *Gaze Meets Machine Learning Workshop*, pages 61–89. PMLR, 2024

Relevant Coursework

• Artificial Intelligence • Statistical Machine Learning • Random Processes and Probability • Linear Algebra
• Computational Models and Algorithms(DSA) • Distributed Computer Systems • Computer Networks

Skills

- **Programming Languages and OS**: Python, C, Ubuntu
- **Software Development Tools**: Docker, GitHub
- **DL Frameworks and Libraries**: PyTorch, HuggingFace, OpenCV, numpy, scipy

References

PhD Advisor
Prof. Kaushik Roy
Purdue University,
West Lafayette, USA
kaushik@purdue.edu