

Aayush Mishra

📍 Baltimore, USA ✉ amishr24@jhu.edu ☎ +1 667-379-9669 🌐 aamixsh.github.io in aamixsh

Towards efficient and reliable adaptation in AI models.

Education

Johns Hopkins University	Jan 2022 -
Ph.D. in Computer Science	- ongoing
M.S.E. in Computer Science – GPA: 4.0/4.0 (transcript 🔗)	- Dec 2023
Indian Institute of Technology Mandi	Aug 2015 - Jun 2019
B.Tech. in Computer Science and Engineering – GPA: 8.88/10.0 (transcript 🔗) with Minor in Management	

Selected Publications

◦ Steered LLM Activations are Non-Surjective	2026
Aayush Mishra, Daniel Khashabi, Anqi Liu [pre-print 🔗 , (Sci4DL 🔗 , Re-Align 🔗)@ICLR]	
◦ Genomic Next-Token Predictors are In-Context Learners	2026
Nathan Breslow, Aayush Mishra, et al. [TMLR 🔗 , (Sci4DL 🔗 , MLGenX (Oral) 🔗)@ICLR]	
◦ IA2 : Alignment with ICL Activations Improves Supervised Fine-Tuning	2025
Aayush Mishra, Daniel Khashabi, Anqi Liu [ICLR 🔗]	
◦ ICL CIPHERS: Quantifying “Learning” in ICL via Substitution Ciphers	2025
Zhouxiang Fang, Aayush Mishra, et al. [EMNLP (Main) 🔗]	
◦ ODD: Overlap-aware Estimation of Model Performance under Distribution Shift	2025
Aayush Mishra, Anqi Liu [UAI 🔗]	
◦ EigenLoRA: Recycle trained Adapters for Resource Efficient Adaptation and Inference	2024
Aayush Mishra*, Prakhar Kaushik*, et al. [pre-print 🔗]	
◦ Do pretrained Transformers Learn In-Context by Gradient Descent?	2024
Aayush Mishra*, Lingfeng Shen*, Daniel Khashabi [ICML (Oral) 🔗]	
◦ Source-Free and Image-Only Unsupervised Domain Adaptation for Category Level OPE	2024
Prakhar Kaushik, Aayush Mishra, et al. [ICLR 🔗]	
◦ VStegNET: Video Steganography Network using Spatio-Temporal features and Micro-Bottleneck	2019
Aayush Mishra*, Suraj Kumar*, et al. [BMVC 🔗]	

Industry Experience

Qualcomm [AI Research Intern] <i>Amsterdam, Netherlands</i>	May 2026 -
◦ Developing memory+compute efficient early exit algorithms for Looped Language Models.	
Adobe [AI Research Intern] <i>College Park, USA</i>	May 2024 - Aug 2024
◦ Developed EigenLoRA , to recycle pre-trained LoRAs for efficient inference and training of new LoRAs.	
Microsoft [Data & Applied Scientist] <i>Hyderabad, India</i>	Aug 2021 - Jan 2022
◦ Developed 1) a query → product class (~18k classes) model; 2) a relevance metric to improve offer ranking.	
Siemens [Research Engineer] <i>Bengaluru, India</i>	Jul 2019 - Aug 2021
◦ Developed a CNN explainability and compression algorithm aimed at edge deployment.	
◦ Used GAN-Dissection to procedurally generate traffic images for stress-testing traffic monitors (YOLO).	
◦ Trained an RL agent to find edge-case safety violations of self-driving agents in a simulator.	
◦ Developed a semantic search tool for codebases using Latent Semantic Analysis, and a Siemens social network post classification tool using statistical NLP. Work done during internships during 2017–2018.	

Skills

Python ●●●●● C++ ●●●●○ Data Engineering ●●●●● ML toolkits (torch/TRL/etc.) ●●●●●