

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 -
<i>Ph.D. in Computer Science</i>	- ongoing
<i>M.S.E. in Computer Science</i> – GPA: 4.0/4.0 (transcript 📄)	- Dec 2023
Indian Institute of Technology Mandi	Aug 2015 - Jun 2019
<i>B.Tech. in Computer Science and Engineering</i> – GPA: 8.88/10.0 (transcript 📄)	
<i>with Minor in Management</i>	

Selected Publications

◦ <i>Genomic Next-Token Predictors are In-Context Learners</i>	2025
Nathan Breslow, Aayush Mishra , et al. [TMLR 📄]	
◦ <i>IA2: Alignment with ICL Activations Improves Supervised Fine-Tuning</i>	2025
Aayush Mishra , Daniel Khashabi, Anqi Liu [ICLR 📄]	
◦ <i>ICL CIPHERS: Quantifying “Learning” in ICL via Substitution Ciphers</i>	2025
Zhouxiang Fang, Aayush Mishra , et al. [EMNLP (Main) 📄]	
◦ <i>ODD: Overlap-aware Estimation of Model Performance under Distribution Shift</i>	2025
Aayush Mishra , Anqi Liu [UAI 📄]	
◦ <i>EigenLoRA: Recycle trained Adapters for Resource Efficient Adaptation and Inference</i>	2024
Aayush Mishra* , Prakhar Kaushik*, et al. [pre-print 📄]	
◦ <i>Do pretrained Transformers Learn In-Context by Gradient Descent?</i>	2024
Aayush Mishra* , Lingfeng Shen*, Daniel Khashabi [ICML (Oral) 📄]	
◦ <i>Source-Free and Image-Only Unsupervised Domain Adaptation for Category Level Object Pose Estimation</i>	2024
Prakhar Kaushik, Aayush Mishra , et al. [ICLR 📄]	

Experience

Adobe	May 2024 - Aug 2024
<i>Research Intern</i> [Document Intelligence Lab]	
◦ Developed EigenLoRA, to recycle pre-trained LoRAs for efficient training and inference of new LoRAs.	
Microsoft	Aug 2021 - Jan 2022
<i>Data & Applied Scientist</i> [Bing Shopping]	
◦ Developed 1) a query → product class (~18k classes) model; 2) a relevance metric to improve offer ranking.	
Siemens	Jul 2019 - Aug 2021
<i>Research Engineer</i> [Automation]	
◦ 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, transformers, etc.) ●●●●● Software Engineering ●●●●○