

Donald Ye

New York City | 917-885-8083 | donaldye827@gmail.com | [LinkedIn](#) | <https://donaldye.com>

PROFESSIONAL EXPERIENCE

Election Data Engineer | The Associated Press

Nov 2024 – Present

- Ensured data integrity across 1,000+ county sources, monitoring pipeline outputs for accuracy and consistency.
- Engineered fault-tolerant Python pipelines delivering real-time election results with 99.9% accuracy.
- Deployed scalable data pipelines used by NYT, Politico, and Google News during 2024 presidential election coverage.

AI Interpretability Research | Algorverse Research Program

Aug 2025 – Present

Mentor: Jonas Rohweder, LMU Munich | Full scholarship recipient

- Authored two first-author preprints on LLM interpretability, accepted and under review at ICLR 2026 and ACL venues.
- Developed a novel causal faithfulness metric demonstrating reasoning collapse at 70–85% of chain length across LLM architectures.
- Designed and executed large-scale ablation experiments on A100 GPUs identifying gradient-causal dissociation.

AI Software Engineering Intern | Headstarter AI

Jul 2024 – Sep 2024

- Developed full-stack customer support chatbot using RAG (Retrieval-Augmented Generation) and LangChain.
- Evaluated chatbot response quality and iterated on retrieval pipeline to improve accuracy and consistency of outputs.
- Improved output consistency through context-aware response generation from structured support documentation.

PUBLICATIONS & PREPRINTS

Ye, D., et al. (2026). *Mechanistic Evidence for Faithfulness Decay in Chain-of-Thought Reasoning*

- Introduced Normalized Logit Difference Decay (NLDD), a metric revealing a reasoning horizon where causal influence collapses.
- Demonstrated faithfulness collapse emerges consistently at 70–85% of chain length across model architectures and task types.
- Accepted ICLR workshop 2026 and under review at EMNLP 2026. [[arXiv:2602.11201](https://arxiv.org/abs/2602.11201)]

Ye, D. (2026). *The Gradient-Causal Gap: Why Gradient Importance Fails on Complex Tasks*

- Identified that gradients negatively correlate with causal importance on tasks, challenging standard pruning assumptions.
- Demonstrated pruning high-gradient 'Bloats' improves accuracy +3.5% while removing low-gradient 'Heroes' devastates it -32%.
- Under review at ACL workshop 2026. [[arXiv:2602.01442](https://arxiv.org/abs/2602.01442)]

TECHNICAL PROJECTS

Political Bias & Misinformation Detection | Technical Project

Jan 2024 – May 2024

- Analyzed 34,000 news articles to determine political patterns and cluster misinformation.
- Authored 11-page [report](#) on multi-class political bias classification and misinformation detection.

Marathon Performance Analysis & Prediction | Technical Project

Sept 2025 – Dec 2025

- Scraped and cleaned 56,000+ NYC Marathon runner records and built a predictive model for finish time estimation.
- Authored 13-page technical [report](#); dataset and code published on GitHub.

EDUCATION

Fordham University, New York City

B.S. in Mathematics & Computer Science

Technical Skills: Python, C++, Javascript, SQL, R, Bash, Pytorch, Anaconda

Leadership: Fordham Mathematics Society (VP), Fordham South Asian Club (Event Manager)