Anxin (Bob) Guo

anxinbguo@gmail.com | anxin-bob-guo.org

Education

Northwestern University, Evanston, Illinois — September 2023 – Present

Ph.D. student in Computer Science.

Northwestern University, Evanston, Illinois — September 2019 – June 2023

B.A. in Mathematics and M.S. in Computer Science. GPA: 3.92/4.00.

Research Interests

Theoretical computer science and the theoretical foundations of machine learning, including learning theory, neural network expressivity, and information-theoretic interpretations of language models.

Research Projects and Papers

AXIOM: Algorithmic Lemma Discovery via Formal Library Compression

With Ruize Xu, Yi Gu, Tianchong Jiang, Chenxiao Yang, and Zhiyuan Li.

Manuscript.

Copying Under Repetition: From Prefix Matching to Position Hashing

With Tianlong Huang, Chenxiao Yang, and Zhiyuan Li.

Manuscript.

- Formalized finite-horizon uniformity, an expressivity regime between non-uniform circuits and fully uniform transformers, where one model must handle all inputs up to length N while resources may scale with N.
- Proved impossibility results for a class of transformers spanning NoPE and a practical implementation of ALiBi: they cannot correctly terminate the copy task even with depth $N^{0.01}$, near-logarithmic precision, and arbitrary width.
- Designed a modular attention head, a plug-in augmentation using value rotation with integer period p to encode query-key distance through the value stream.
- Showed that the augmented architecture can copy arbitrary strings using logarithmic width and $\log \log N$ precision.
- Built diagnostic experiments showing that frontier 400B-parameter language models struggle to copy strings with extreme repetition.

Hallucination is a Consequence of Space-Optimality: A Rate-Distortion Theorem for Membership Testing

With Jingwei Li.

International Conference on Machine Learning (ICML), 2026, spotlight. arXiv.

- Defined membership testing as a unifying abstraction for approximate set membership, Bloom-filter-style error metrics, and factual errors in language models.
- Established the optimal space-error tradeoff in this framework, with an asymptotic lower bound governed by relative entropy, analogous to rate-distortion theory.
- Showed that optimally trained language models under cross-entropy loss with bounded memory may assign high confidence to unseen non-facts as a structural consequence of limited capacity.
- Extended known bounds for Bloom-type filters into a tight form, including regimes where false negatives are allowed.

Agnostic Learning of Arbitrary ReLU Activation under Gaussian Marginals

With Aravindan Vijayaraghavan.

Conference on Learning Theory (COLT), 2025. arXiv; conference version; recorded talk.

- Gave the first polynomial-time constant-factor approximation agnostic learner for arbitrarily biased ReLU neurons under Gaussian marginals.
- Established a formal separation between Statistical Query (SQ) and Correlational Statistical Query (CSQ) models for this problem, proving limitations of gradient-based algorithms in this setting.
- Designed a method combining Thresholded PCA and Reweighted Gradient Descent to overcome limitations of standard gradient-based approaches.

To Store or Not to Store: A Graph-Theoretic Approach for Dataset Versioning

With Jingwei Li, Pattara Sukprasert, Samir Khuller, Amol Deshpande, and Koyel Mukherjee.

IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2024. arXiv; conference version.

- Formulated and analyzed approximation algorithms and heuristics for optimizing the storage-retrieval cost tradeoff in large-scale dataset versioning systems.
- Designed provably near-optimal algorithms for tree-like graphs by exploiting low-treewidth structure.
- Implemented and benchmarked proposed heuristics on experiments based on real-world GitHub repository data, improving performance by up to 1000x over prior methods.

Awards and Honors

Ph.D. Student Research Award, Computer Science Department, Northwestern University 2024 – 2025
Barris Award for Outstanding TA, CS 335: Introduction to the Theory of Computing Fall 2024
Junior Career Award in Mathematics, Mathematics Department, Northwestern University 2021 – 2022

Research and Teaching

IDEAL-funded summer research exchange, University of Chicago Summer 2026

Hosted by Chao Gao.

Research Intern / IDEAL-funded summer research exchange, TTIC Summer 2025

Toyota Technological Institute at Chicago. Hosted by Zhiyuan Li.

- Initiated the transformer copying project during the internship.
- Built experiments validating the failure mode of frontier language models and the proposed architecture.
- Presented at a reading group.

Teaching Assistant, CS 336: Design and Analysis of Algorithms Spring 2026
Teaching Assistant, CS 335: Introduction to the Theory of Computing Fall 2024

- Gave a guest lecture on busy beaver numbers in addition to regular TA duties.
- Awarded the Barris Award for Outstanding TA.

Service and Programs

Program Committee / Reviewer, Reliable ML from Unreliable Data, NeurIPS Workshop 2025
Volunteer, 65th IEEE Symposium on Foundations of Computer Science (FOCS) 2024
Ross Mathematics Program Asia, student, junior counselor, counselor 2018 – 2020
Directed Reading Program, Linear Representations of Finite Groups Spring 2021

Read with Wenyuan Li.

Skills

Technology: C++, Python (NumPy, PyTorch, pandas, matplotlib), Mathematica, and Gurobi Optimizer.

Writing

I have written two Chinese-language articles on theoretical computer science for my Zhihu column.