Garrett Bingham

gjb.ai
garrett@gjb.ai
linkedin.com/in/garrettbingham
x.com/gjb_ai
github.com/garrettbingham
Last Updated: November 4, 2026

EDUCATION

The University of Texas at Austin — Austin, TX · 2019 - 2023
Ph.D., Department of Computer Science — Advisor: Risto Miikkulainen

Yale University — New Haven, CT · 2015 - 2019
B.S., Computer Science & Mathematics, *cum laude*

RESEARCH AND WORK EXPERIENCE

Google DeepMind — Mountain View, CA · Oct 2023 - Present
Staff Research Scientist

- Implemented the Self-Verification pattern in Antigravity Teamwork, enabling research-level math at Flash model scale.
- Co-led engineering contributions behind *Aletheia*, a math research agent which solved 13 open Erdős problems, contributed intermediate propositions on two research papers, collaborated with a human author on a third, and produced a fourth standalone research paper.
- Proposed and co-led Deep Think, a method which enables Gemini to consider multiple hypotheses in parallel before responding.
- Co-led development of systems that achieved gold and silver medals at the 2025 and 2024 International Mathematical Olympiads.
- Contributed to IMO-Bench dataset for advancing mathematical reasoning and HaloQuest dataset to improve multimodal reasoning.

Cognizant AI Labs — San Francisco, CA · May 2020 - Oct 2023
Machine Learning Researcher

- Developed AQuaSurF, a technique for efficient activation function optimization through surrogate modeling.
- Fisher eigenvalues and function outputs are used to accurately predict performance, accelerating search for better activation functions.
- AQuaSurF is orders of magnitude more efficient than comparable methods and scales up to large datasets like ImageNet.

- Developed AutoInit, an approach for analytic signal-preserving weight initialization for neural networks.
- AutoInit improves performance of convolutional, residual, and transformer networks across a range of activation function, dropout, weight decay, learning rate, and normalizer settings.
- In neural architecture search and activation function discovery, AutoInit automatically calculates specialized weight initialization strategies for thousands of unique architectures and hundreds of unique activation functions.
- AutoInit improves performance in vision, language, tabular, multi-task, and transfer learning scenarios.

Neural Networks Research Group — Austin, TX · Aug 2019 - May 2023
Graduate Research Assistant

- Developed PANGAEA, a process combining evolutionary and gradient-based optimization to design deep network activation functions.
- Activation functions discovered by PANGAEA give statistically significant increases in performance over 31 baseline functions.

- Demonstrated that evolutionary algorithms can design novel deep learning activation functions that outperform ReLU.
- Replacing ReLU with an evolved activation function in Wide ResNet gives a statistically significant increase in accuracy on CIFAR-100.
- Showed that an evolutionary algorithm discovers better activation functions more quickly than random search.

Amazon Robotics AI — Seattle, WA · May - Aug 2019
Software Development Engineer Intern

- Migrated production large-scale video classification workflow to internal computation framework, optimizing for specific use cases.
- Parallelized 10TB dataset construction and preprocessing; used cached results from prior datasets to incrementally build new ones.
- Created model versioning scheme to allow framework to substitute precomputed per-epoch results instead of requiring two weeks of retraining for every experiment.
- Assisted visiting scientists with onboarding and added GPU support to the framework.

Language, Information, and Learning at Yale (LILY) Lab New Haven, CT · Aug 2018 - May 2019
Natural Language Processing Researcher

- Developed BiDARTS, a novel approach to automatically design bidirectional recurrent neural network architectures.
- BiDARTS approached state-of-the-art accuracy on word token part of speech tagging across 60 treebanks (91.32% vs. 91.83%).
- Empirically demonstrated that random search outperforms current gradient-based approaches to neural network architecture design.
- Improved English-Swahili and English-Tagalog cross-lingual document retrieval relevance scores by 5.2% by training support vector regressor to predict query-specific document cutoffs. Research done for IARPA MATERIAL Program.

Reservoir Labs New York, NY · Jun - Aug 2018
Machine Learning Engineer Intern

- Utilized polyhedral compiler to speed up neural network inference time by 12% without sacrificing accuracy.
- Implemented methods to translate TensorFlow operations to equivalent C code, increasing percentage of optimizable operations from 84% to 99% for Inception V3/V4 and from 58% to 95% for ResNet V2 50.
- Designed algorithm to optimize disjoint subgraphs of the neural network, decreasing optimization time to less than one minute.
- Improved robustness by developing algorithms that guarantee the optimized neural network remains acyclic.

University of North Carolina Wilmington Wilmington, NC · May - Jul 2017
NSF Research Experience for Undergraduates Intern

- Developed novel face recognition algorithm which classifies by weighting predictions made in multiple random feature subspaces.
- Improved accuracy from 60.9% to 78.8% on MORPH-II and from 90.3% to 94.8% on ORL datasets vs. related algorithms.
- Corrected 2,700 errors in the 55,000 entry MORPH-II face image dataset that previously published research missed.

PUBLICATIONS

Journal Articles

1. G. Bingham and R. Miikkulainen. Discovering Parametric Activation Functions. *Neural Networks*, 2022

Refereed Conference Papers

2. M.-T. Luong, D. Hwang, H. H. Nguyen, G. Ghiasi, Y. Chervonyi, I. Seo, J. Kim, G. Bingham, J. Lee, S. Mishra, et al. Towards robust mathematical reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 35406–35430, 2025

3. Z. Wang*, G. Bingham*, A. W. Yu, Q. V. Le, T. Luong, and G. Ghiasi. HaloQuest: A Visual Hallucination Dataset for Advancing Multimodal Reasoning. In *European Conference on Computer Vision (ECCV)*, 2024

4. G. Bingham and R. Miikkulainen. Efficient activation function optimization through surrogate modeling. *Advances in Neural Information Processing Systems*, 36:6634–6661, 2023

5. G. Bingham and R. Miikkulainen. AutoInit: Analytic Signal-Preserving Weight Initialization for Neural Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2023

6. G. Bingham*, W. Macke*, and R. Miikkulainen. Evolutionary Optimization of Deep Learning Activation Functions. In *Genetic and Evolutionary Computation Conference (GECCO '20), July 8–12, 2020, Cancún, Mexico*, 2020

7. R. Zhang, C. Westerfield, S. Shim, G. Bingham, A. R. Fabbri, W. Hu, N. Verma, and D. Radev. Improving Low-Resource Cross-Lingual Document Retrieval by Reranking with Deep Bilingual Representations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3173–3179, 2019

8. B. Yip, G. Bingham, K. Kempfert, J. Fabish, T. Kling, C. Chen, and Y. Wang. Preliminary Studies on a Large Face Database. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 2572–2579. IEEE, 2018

Technical Reports

9. T. Feng, J. Jung, S. hyun Kim, C. Pagano, S. Gukov, C.-C. Tsai, D. Woodruff, A. Javanmard, A. Mokhtari, D. Hwang, Y. Chervonyi, J. N. Lee, G. Bingham, T. H. Trinh, V. Mirrokni, Q. V. Le, and T. Luong. Aletheia tackles firstproof autonomously. *arXiv:2602.21201*, 2026
10. T. Feng, T. H. Trinh, G. Bingham, D. Hwang, Y. Chervonyi, J. Jung, J. Lee, C. Pagano, S.-h. Kim, F. Pasqualotto, et al. Towards autonomous mathematics research. *arXiv:2602.10177*, 2026
11. T. Feng, T. Trinh, G. Bingham, J. Kang, S. Zhang, S.-h. Kim, K. Barreto, C. Schildkraut, J. Jung, J. Seo, et al. Semi-autonomous mathematics discovery with gemini: A case study on the Erdős problems. *arXiv:2601.22401*, 2026
12. Gemini Team, Google. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. 2025
13. Gemini Team, Google. Gemini 1.5: Unlocking Multimodal Understanding across Millions of Tokens of Context. *arXiv:2403.05530*, 2024
14. G. Bingham. Optimizing Neural Networks through Activation Function Discovery and Automatic Weight Initialization. *Ph.D. Dissertation*, 2023
15. G. Bingham. Part of Speech Tagging with Neural Architecture Search. *Yale University Senior Project*, 2019
16. G. Bingham. Random Subspace Two-Dimensional LDA for Face Recognition. *arXiv:1711.00575*, 2017
17. G. Bingham, B. Yip, M. Ferguson, C. Nansalo, C. Chen, Y. Wang, and T. Kling. MORPH-II: Inconsistencies and Cleaning. *University of North Carolina Wilmington NSF REU*, 2017

PRESENTATIONS

Oral Presentations

18. G. Bingham. Optimizing Neural Networks through Activation Function Discovery and Automatic Weight Initialization. Ph.D. Dissertation Defense. The University of Texas at Austin. Austin, TX. April 3, 2023
19. G. Bingham. AutoInit: Analytic Signal-Preserving Weight Initialization for Neural Networks. AAAI 2023. Washington, DC. February 7-14, 2023
20. G. Bingham. Automated Activation Function Design with Evolutionary Computation. Ph.D. Proposal. The University of Texas at Austin. Austin, TX. December 1, 2020
21. G. Bingham. Evolutionary Optimization of Deep Learning Activation Functions. BEACON Congress. August 12-14, 2020
22. G. Bingham. Evolutionary Optimization of Deep Learning Activation Functions. GECCO 2020. Cancún, Mexico. July 8-12, 2020
23. G. Bingham. Evolutionary Optimization of Deep Learning Activation Functions. Ph.D. Research Preparation Exam. The University of Texas at Austin. Austin, TX. April 21, 2020
24. G. Bingham. Random Subspace Two-Dimensional LDA for Face Recognition. NES Mathematical Association of America Meeting. Sacred Heart University. Fairfield, CT. November 17-18, 2017

Poster Presentations

25. G. Bingham. Efficient Activation Function Optimization through Surrogate Modeling. NeurIPS 2023. New Orleans, LA. December 10-16, 2023
26. G. Bingham. AutoInit: Analytic Signal-Preserving Weight Initialization for Neural Networks. AAAI 2023. Washington, DC. February 7-14, 2023
27. G. Bingham. Random Subspace Two-Dimensional LDA for Face Recognition. Council on Undergraduate Research REU Symposium. The Westin Alexandria. Alexandria, VA. October 22-23, 2017
28. G. Bingham. Random Subspace Two-Dimensional LDA for Face Recognition. REU Research Showcase. The University of North Carolina Wilmington. Wilmington, NC. July 24, 2017

PATENTS

29. G. Bingham and R. Miikkulainen. Method and System for Optimizing Activation Functions in Neural Networks, Dec. 13 2023. US Patent App. 18/537,924
30. G. Bingham and R. Miikkulainen. System and Method for Evaluating Weight Initialization for Neural Network Models, Mar. 23 2023. US Patent App. 17/855,955
31. G. Bingham and R. Miikkulainen. System and Method For Generating Parametric Activation Functions, Feb. 17 2022. US Patent App. 17/399,470
32. One patent pending

OPEN-SOURCE REPOSITORIES

33. M.-T. Luong, D. Hwang, H. H. Nguyen, G. Ghiasi, Y. Chervonyi, I. Seo, J. Kim, G. Bingham, J. Lee, S. Mishra, et al. IMO-Bench. https://github.com/google-deepmind/superhuman, 2025
34. Z. Wang, G. Bingham, et al. HaloQuest. https://github.com/google/haloquest, 2024
35. G. Bingham. AQuaSurF. https://github.com/cognizant-ai-labs/aquasurf, 2023
36. G. Bingham. Act-Bench. https://github.com/cognizant-ai-labs/act-bench, 2023
37. G. Bingham et al. AutoInit. https://github.com/cognizant-ai-labs/autoinit, 2021

AWARDS

38. Neural Networks Best Paper Award. https://www.sciencedirect.com/journal/neural-networks/about/awards/announcement-of-the-neural-networks-best-paper-award, 2024.
39. SIGEVO Best Dissertation Award Honorable Mention. https://sig.sigevo.org/index.html/tiki-index.php?page=SIGEVO+Dissertation+Award#SIGEVO_Dissertation_Award, 2024.
40. University Graduate Continuing Fellowship, The University of Texas at Austin, 2021 - 2022, $44,000. (declined)

NEWS COVERAGE

41. The Antigravity Team. Teamwork: When AI Becomes a Research Partner. https://antigravity.google/blog/teamwork-when-ai-becomes-a-research-partner, 2026
42. T. Luong and V. Mirrokni. Accelerating Mathematical and Scientific Discovery with Gemini Deep Think. https://deepmind.google/blog/accelerating-mathematical-and-scientific-discovery-with-gemini-deep-think/, 2026

43. H. M. Lin and H.-T. Cheng. Gemini achieves gold-medal level at the International Collegiate Programming Contest World Finals. https://deepmind.google/discover/blog/gemini-achieves-gold-level-performance-at-the-international-collegiate-programming-contest-world-finals/, 2025

44. The Deep Think team. Try Deep Think in the Gemini app. https://blog.google/products/gemini/gemini-2-5-deep-think/, 2025

45. T. Luong and E. Lockhart. Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the International Mathematical Olympiad. https://deepmind.google/discover/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-international-mathematical-olympiad/, 2025

46. T. Doshi. Gemini 2.5: Our most intelligent models are getting even better. https://blog.google/technology/google-deepmind/google-gemini-updates-io-2025/#deep-think, 2025

47. A. Tseng. Risto Miikkulainen's Leadership in AI Research and Industry Recognized Globally. https://www.cs.utexas.edu/news/2024/risto-miikkulainens-leadership-ai-research-and-industry-recognized-globally, 2024

48. AlphaProof and AlphaGeometry teams. AI achieves silver-medal standard solving International Mathematical Olympiad problems. https://deepmind.google/discover/blog/ai-solves-imo-problems-at-silver-medal-level/, 2024

49. T. Erales. A More Efficient Future For Neural Network Systems. https://www.cs.utexas.edu/news/2023/more-efficient-future-neural-network-systems, 2023

50. A. Schiff. TXCS Researchers Design Evolutionary Algorithms for Neural Networks. https://www.cs.utexas.edu/news/2020/txcs-researchers-design-evolutionary-algorithms-neural-networks, 2020

REVIEWING

Journals

- ACM Journal on Emerging Technologies in Computing Systems
- IEEE Transactions on Evolutionary Computation
- Nature Machine Intelligence

Conferences

- AAAI (Association for the Advancement of Artificial Intelligence)
- GECCO (Genetic and Evolutionary Computation Conference)
- ICML (International Conference on Machine Learning)
- NeurIPS (Neural Information Processing Systems)