

JOHN WILMES

Staff AI Scientist

📍 Chicago, IL 📞 (224) 475-1565 ✉️ jw@johnwilmes.name 🌐 johnwilmes.name 🔗 [john-wilmes](https://www.linkedin.com/in/john-wilmes) 🐙 [johnwilmes](https://github.com/johnwilmes)

Machine learning scientist and technical leader with a decade of experience spanning theoretical foundations and production systems. Proven track record from first-of-kind training guarantees for neural networks (COLT, NeurIPS) to production AI shipping at Copilot scale. Contributed core research to a \$31M Series A. Deep expertise in LLMs, symbolic reasoning, and algorithm design; experience leading small research teams and translating business objectives into engineering deliverables.

EXPERIENCE

Staff Applied Researcher

[GitHub](#)

📅 Feb 2025 – Present

📍 Chicago, IL (remote)

- Designed and implemented new contextualization framework for Copilot code completions, correcting 73% of previously unsolved evaluation scenarios and reducing latency even while adding 30% more relevant context.
- Authored offline evaluations strategy for Copilot Applied Science and developed automated pipeline for benchmark creation, reducing benchmark creation time from days to minutes. Built optimization pipeline to align offline evaluations with real-world usage from telemetry.

Director of Research Engineering

Senior Machine Learning Researcher

[Symbolica AI](#)

📅 Nov 2023 – Nov 2024

📅 Aug 2023 – Oct 2023

📍 Chicago, IL (remote)

- Developed proof-of-concept models beating industry-leading LLMs on a formal language task using 10000x fewer parameters, leading to \$31m Series A.
- Led research and engineering of SoTA mathematical theorem proving models combining LLMs, RAG, neurosymbolic integration, and reinforcement learning. Grew team of researchers and engineers to eight ICs and wrote technical roadmap to translate business objectives into R&D deliverables.

Founder and Principal Scientist

[Wilmes Consulting](#)

📅 Feb 2018 – Dec 2023

📍 Chicago, IL

- Achieved >200x speedup of NLP training pipeline in C and Rust for Gen AI client
- Integrated 11 data sources using Python and SQL to build XGBoost predictive model of member retention and design behavioral member segmentation for professional society.
- Designed and implemented deep learning algorithms for categorizing human activities from motion sensor data using Python, Tensorflow, and AWS, improving accuracy by 20% for medical AI client.

Assistant Professor of Mathematics

[Brandeis University](#)

📅 Jul 2018 – Jun 2021

📍 Waltham, MA

- Awarded \$175k NSF grant “Guarantees for Training Neural Networks,” producing first training guarantees for convolutional graph neural networks on SBM data, published in ICLR.
- Supervised two Ph.D. theses, producing new state-of-art algorithms using deep learning and Markov chain Monte Carlo techniques.

EDUCATION

Postdoc at Algorithms and Randomness Center

📅 Sep 2016 – Jun 2018

[Georgia Institute of Technology](#)

- Outstanding Post-Doctoral Research Award for advances in provable guarantees for deep learning and Markov chain Monte Carlo algorithms.

Ph.D. and M.S. in Mathematics

📅 Sep 2010 – Aug 2016

[University of Chicago](#)

- Algorithms research as NSF Graduate Research Fellow led to most significant breakthroughs in decades on infamous Graph Isomorphism problem.

B.A. in Mathematics

📅 Sep 2006 – Jun 2010

[Reed College](#)

- Barry M. Goldwater scholarship. Undergraduate thesis result appears in textbook as “Wilmes’ Theorem.”

PUBLICATIONS

(Mathematics and theoretical computer science communities follow a convention of alphabetical author listings.)

1. Santosh Vempala and John Wilmes. Gradient descent for one-hidden-layer neural networks: Polynomial convergence and SQ lower bounds. In *Proceedings of the 32nd Conference on Learning Theory (COLT)*, 2019.
2. Daniel Štefankovič, Eric Vigoda, and John Wilmes. On counting perfect matchings in general graphs. In *Proceedings of the 13th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 873–885, 2018.
3. Le Song, Santosh Vempala, John Wilmes, and Bo Xie. On the complexity of learning neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5514–5522, 2017.
4. László Babai and John Wilmes. Asymptotic Delsarte cliques in distance-regular graphs. *Journal of Algebraic Combinatorics*, 43(4):771–782, 2016.
5. Xiaorui Sun and John Wilmes. Faster canonical forms for primitive coherent configurations. In *Proceedings of the 47th ACM on Symposium on Theory of Computing (STOC)*, pages 693–702, 2015.
6. Madhusudan Manjunath, Frank-Olaf Schreyer, and John Wilmes. Minimal free resolutions of the G -parking function ideal and the toppling ideal. *Transactions of the American Mathematical Society*, 367(4):2853–2874, 2015.
7. László Babai, Xi Chen, Xiaorui Sun, Shang-Hua Teng, and John Wilmes. Faster canonical forms for strongly regular graphs. In *Proceedings of the 54th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 157–166, 2013.
8. Laszlo Babai and John Wilmes. Quasipolynomial-time canonical form for Steiner designs. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 261–270, 2013.
9. David Perkinson, Jacob Perlman, and John Wilmes. Primer for the algebraic geometry of sandpiles. *Tropical and non-Archimedean geometry*, 605:211–256, 2013.