Vimarsh Sathia

Ph.D. Student in Computer Science, University of Illinois Urbana-Champaign
✉ vimarsh.sathia@gmail.com · vimarsh6739.github.io · GitHub · Google Scholar · LinkedIn

Research Interests

Compilers and programming languages for scientific computing and machine learning; automatic differentiation; MLIR-based tensor compilers; numerical program optimization; performance engineering for CPUs, GPUs, and accelerators.

Education

University of Illinois Urbana-Champaign Aug 2022 – Present
Ph.D. in Computer Science
Advisor: William S. Moses

Indian Institute of Technology Madras 2017 – 2021
B.Tech. in Computer Science and Engineering

Honors and Awards

2026 Distinguished Paper Award, CGO 2026
for "Thinking Fast and Correct: Automated Rewriting of Numerical Code through Compiler Augmentation"

2021 C. S. Krishnamoorthy Endowment Prize, IIT Madras
for the best B.Tech./M.Tech. project in Genetic Algorithms and Evolutionary Computation

Research Appointments

University of Illinois Urbana-Champaign Aug 2022 – Present
Graduate Research Assistant, Urbana, IL

Modal Labs Jun 2026 – Present
Member of Technical Staff (full-time Jun – Aug 2026; part-time since Aug 2026)

AMD Research and Advanced Development May 2025 – Aug 2025
Research Associate, Austin, TX
- Verified GPU data-movement instructions and transformations using SMT.
- Built a translation-validation pipeline in MLIR using the `smt` dialect.

Qualcomm, Compiler Labs May 2024 – Aug 2024
Research Intern, New York, NY
- Prototyped a Hexagon backend for IREE with the MLIR team.
- Designed and implemented compiler passes for automatic tiling of HVX code.

Microsoft, Azure Jun 2021 – Jul 2022
Software Engineer, Bangalore, India

Research Affiliations

NSF–Simons AI Institute for the Sky (SkAI) Sep 2024 – Present
Graduate Researcher, project "Automating Bayesian Inference of Millimeter Source Association"
PIs: William S. Moses, Joaquin Vieira

MIT Center for the Exascale Simulation of Coupled High Enthalpy Fluid–Solid Interactions (CHEFSI) Sep 2025 – Present
Graduate Student Researcher; DOE NNSA PSAAP IV Predictive Science Center

Publications

Conference Papers

[CGO'27] **Floating up the Peak: Automatic Numerical Rewriting Enables Fast and Tensorized Accelerator Code**
Siyuan Brant Qian, **Vimarsh Sathia**, Jan Hückelheim, Jens Domke, Paul Hovland, William S. Moses
25th ACM/IEEE International Symposium on Code Generation and Optimization (CGO), Salt Lake City, UT, 2027. To appear.

[CGO'26a] **Thinking Fast and Correct: Automated Rewriting of Numerical Code through Compiler Augmentation**
Siyuan Brant Qian, **Vimarsh Sathia**, Ivan R. Ivanov, Jan Hückelheim, Paul Hovland, William S. Moses
24th ACM/IEEE International Symposium on Code Generation and Optimization (CGO), Sydney, Australia, 2026.
Distinguished Paper Award

[CGO'26b] **GRANII: Selection and Ordering of Primitives in GRAph Neural Networks using Input Inspection**
Damitha Lenadora, **Vimarsh Sathia**, Gerasimos Gerogiannis, Serif Yesil, Josep Torrellas, Charith Mendis
24th ACM/IEEE International Symposium on Code Generation and Optimization (CGO), Sydney, Australia, 2026, pp. 14–27. *doi:10.1109/CGO68049.2026.11395219*

Workshop Papers

[SC-W'26] **SkyMatch.jl: Fast, Differentiable Modeling for Millimeter-Sky Inference**
Vimarsh Sathia, Kedar Phadke, Siyuan Brant Qian, Joaquin Vieira, William S. Moses
Scale to See: Scalable AI for Scientific Imaging and Spatiotemporal Data, Workshop at SC'26, Chicago, IL, 2026.

[SC-W'25] **Data Race Detection through Vibe Translation**
Jan Hückelheim, **Vimarsh Sathia**, Siyuan Brant Qian
9th International Workshop on Software Correctness for HPC Applications (Correctness), Workshop at SC'25, St. Louis, MO, 2025.
doi:10.1145/3731599.3767363
Best Paper Presentation Award

[DiffProg'25] **Compiler Optimizations for Higher-Order Automatic Differentiation**
Vimarsh Sathia, Siyuan Brant Qian, Jan Hückelheim, Paul Hovland, William S. Moses
2nd Workshop on Differentiable Parallel Programming (DiffProg), co-located with PPoPP'25, Las Vegas, NV, 2025.

Under Review

[Submitted] **Reactant: Abstraction Raising Automates Kernel Portability over Specialized Accelerators**
William S. Moses, Paul Berg, **Vimarsh Sathia**, Avik Pal, Ivan R. Ivanov, Jules Merckx, Arpit Jaiswal, Lorenzo Chelini, Albert Cohen, Oleksandr Zinenko
Under review, 2026.

[Submitted] **Impulse: Momentously Fast, General, and Portable Probabilistic Programming via Compiler Augmentation**
Siyuan Brant Qian, **Vimarsh Sathia**, Jesse Michel, William S. Moses
Under review, 2026.

Preprints

[arXiv'21] **Accelerating Genetic Programming using GPUs**
Vimarsh Sathia, Venkataramana Ganesh, Shankara Rao Thejaswi Nanditale
arXiv:2110.11226, 2021.

Talks

SkyMatch.jl: Fast, Differentiable Modeling for Millimeter-Sky Inference
- Scale to See Workshop at SC'26, 11/20/2026, Chicago, IL

Optimizing Across the Differentiation Boundary with EnzymeMLIR
- 2026 US LLVM Developers' Meeting, 10/2026, Santa Clara, CA

Raising, Differentiating, and Optimizing Tensor Programs in EnzymeMLIR
- Modal Labs, 07/08/2026, New York, NY
- Seminar talk, University of Illinois Urbana-Champaign, 05/04/2026

AD-aware Compiler Optimizations in Enzyme MLIR
- MFEM Community Workshop 2025, 09/10/2025

Compiler Optimizations for Higher-Order Automatic Differentiation
- DiffProg Workshop at PPoPP'25, 03/01/2025, Las Vegas, NV

Teaching

Teaching Assistant, University of Illinois Urbana-Champaign

Fall 2026	CS 598 APE Advanced Performance Engineering (William S. Moses)
Spring 2026	CS 421 Programming Languages (Mattox Beckman)
Spring 2025	CS 598 APE Advanced Performance Engineering (William S. Moses)
Fall 2024	CS 426 Compiler Construction (Charith Mendis)

Lectures

Fall 2026	CS 598 APE Advanced Performance Engineering, UIUC (1 lecture)

Mentoring

- **Hashmita Nittala** (UIUC undergraduate): Bayesian forward modeling in SkyMatch.jl.
- **Yuansui Xu** (UIUC): Portable performance for GPU code across NVIDIA and AMD accelerators.
- **Aditya Prerepa** and **Ayush Raman** (UIUC undergraduates): Performance engineering with LLMs (pesquared).
- **Siddarth Aananth** (UIUC): Pruning and semi-supervised learning.