Zack Hassman

zhassman [at] uchicago [dot] edu | zackhassman.com
(last updated August 27, 2025)

I'm an undergraduate at UChicago studying mathematics, statistics, and computer science. My current research focuses on quantum information science and deep learning. Previously, I published work on hardware-inspired techniques to accelerate training and inference for machine learning models.

Education

The University of Chicago | Chicago, IL   **September 2022 - Present**
- Pursuing distinct undergraduate degrees in Mathematics, Statistics, and Computer Science
- GPA: 3.759/4
- Relevant graduate coursework: Information and Coding Theory; Statistical Theory and Methods; Complexity Theory; and Geometric Methods in Computer Science
- Favorite undergraduate coursework: Markov Chains, Martingales, and Brownian Motion; and Mathematical Probability

Research and Work Experience

EPiQC Group, Department of Computer Science, UChicago – *Research Assistant*   **August 2024-Present**
- Leading a collaboration between Professor Fred Chong's group, IBM, and the Center for Theoretical Physics PAS
- Developed a high-performance quantum circuit simulator in Rust & Python compatible with a class of quantum circuits used by IBM for chemistry simulations
- Used simulator to improve ground state energy estimates for electronic configuration data obtained on the IBM Heron Quantum Processor

Infleqtion — *Quantum Software Engineer*   **June 2024-August 2024**
- Investigated physics-inspired machine learning techniques to be considered for Infleqtion's submission to the $50 million Wellcome Leap Q4Bio challenge program
- Developed the first practical open-source Quantum Boltzmann Machine with complete training and inference demonstrations on classical hardware
- Used Quantum Boltzmann Machines to classify types of cancer using DNA features representing genetic mutations

YLAB Group, Department of Computer Science, UChicago — *Research Assistant*   **March 2023-May 2024**
- Produced two publications advised by Professor Yanjing Li
  - SySMOL – a novel hardware-efficient training algorithm for convolutional neural networks and transformers
  - YFlows – an analysis of dataflow techniques that demonstrate how to accelerate convolutional neural networks up to 4.8x (published at the International Conference on Compiler Construction 2024)
- Leveraged the UChicago AI Cluster and the Slurm Workload Manager to train hundreds of large machine learning models on GPUs
- Optimized the transformer architecture in Meta's Fairseq repository, reducing its size by 80% while maintaining original performance
- Implemented a variety of iconic convolutional neural network architectures (VGG, ResNet, DenseNet, MobileNet, SqueezeNet, etc.) in PyTorch
- Benchmarked and tuned neural networks using the Apache Tensor Virtual Machine (TVM) compiler

Google's Computer Science Summer Institute — *Participant*   **July 2022 - August 2022**
- Participated in an intensive computer science summer program led by Google software engineers
- Integrated an API to build an automatically-updating hackathon database for computer science students

Miscellaneous

- **Awards and Grants:** Quad Undergraduate Research Scholar (2x), Jeff Metcalf Grant (2x)
- **Technical:** Python, C, Rust, Java, JavaScript, PyTorch, Qiskit, LaTeX, Git, Linux, Vim
- **Other skills:** Proficient in Spanish. Previously a classical cellist.