

PROFILE

Passionate second-year PhD student in Computer Science and Engineering looking for intern software engineering positions with experiences in sparse tensor algebra accelerations, GPU code development, and knowledge in AI/ML.

EDUCATION

- Ph.D. Computer Science, University of California - Riverside

Riverside, California

Sep 2024 – Current
- B.S. Computer Engineering, University of California - San Diego

San Diego, California

Sep 2019 – Sep 2023

PROJECTS & CONTRIBUTIONS

- RTSpMSpM: Harnessing Ray Tracing for Efficient Sparse Matrix Computations

Lab: ESCAL, Riverside Advisor: Hung-Wei, Tseng

September 2023 – Present

 - Designed the algorithm that maps the sparse matrix multiplications (SpMSpM) into ray tracing problems to enable the use of ray tracing accelerators in solving SpMSpM.
 - Developed a library that realized the proposed algorithm using OptiX and CUDA to achieve 1.85× speedup over cuSPARSE on NVIDIA’s RT Cores.
 - Integrated the developed library into a graph neural network training framework using CUDA and Python to achieve more than 3× speedup on NVIDIA RTX 4090 GPUs.
- Machine Learning in Music

June 2023 – August 2023

 - Developed a genre-conditioned music generation system using PyTorch, LSTM, and autoencoders to synthesize short audio clips (~3 minutes) matching target styles.
 - Trained a convolutional neural network (CNN) using mel spectrogram inputs to classify music tracks by genre, achieving over 85% accuracy on a curated dataset.
 - Applied autoencoder-based compression and denoising to enhance audio quality and reduce input dimensionality for downstream learning tasks.

TECHNICAL SKILLS

- Programming and Scripting Languages**
Python, C++, C, Bash
- Libraries and Frameworks**
NVIDIA OptiX, CUDA, cuSPARSE, NumPy, Docker
- Tools and Debuggers**
GitHub, gdb, jdb, vim
- Operating Systems and Simulation Software**
Linux (Ubuntu), Windows, MATLAB, LTSpice, Multisim, Proteus

PUBLICATIONS

Hongrui Zhang, Yunan Zhang, and Hung-Wei Tseng. 2025. *RTSpMSpM: Harnessing Ray Tracing for Efficient Sparse Matrix Computations*. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture (ISCA '25)*. Association for Computing Machinery, New York, NY, USA, 359–373. <https://doi.org/10.1145/3695053.3731072>