

Junwei ZHOU

(765)714-9304 | E-mail: zjw330501@gmail.com

EDUCATION

University of Michigan, Ann Arbor

Sept. 2021 – Apr. 2023

- M.S. in Computer Science and Engineering, GPA: 4.0/4.0
- Online Computer Graphics Courses: Introduction to Computer Graphics (GAMES101), Advanced Physics Engine Implementation (GAMES201), Real-Time High Quality Rendering (GAMES202)

University of Science and Technology of China (Top 5 university in China)

Sept. 2017 – Jun. 2021

- B.S. in Geophysics, GPA: 4.02/4.3, Ranking: 1/66 | B.E. in Computer Science (Dual), GPA: 3.94/4.3
- Awards: Guo Moruo Scholarship (the highest honor of USTC undergraduates), National Scholarship in China

PUBLICATIONS

- **A High Performance Computing Method for Noise Cross-correlation Functions of Seismic Data.** J. Zhou, Q. Wei, C. Wu, G. Sun. *2021 IEEE International Symposium on Parallel and Distributed Processing with Applications*
- **Parallel Ambient Noise Data Processing Algorithm based on CPU-GPU Platform.** C. Wu, Q. Wei, J. Zhou, H. Li, G. Sun. *In Submission*

RESEARCH EXPERIENCES

Position Based Fluids Simulation [Python, Taichi]

Jan. 2023 – Feb. 2023

- Developed an interactive real-time application to simulate fluids in 2D and 3D spaces
- Used particles to simulate fluids. Implemented **position based fluids** algorithm to compute particles' positions
- Implemented **marching cubes** algorithm to construct meshes from particles and rendered the fluid surface

Path Tracing Renderer [C++]

Dec. 2022 – Jan. 2023

- Implemented a path tracing renderer with **Monte Carlo integration** to render the Cornell box scene
- Constructed bounding volume hierarchy (**BVH**) to accelerate the computation of ray-object intersection
- Improved the computation efficiency by sampling the light instead of sampling randomly in all directions

High Performance Computing on Cross-correlation Functions [C/C++, CUDA]

Jul. 2020 – Jun. 2021

- Produced **CUDA** parallel programs to compute cross-correlation functions of large-scale seismic data on **GPU**
- Preloaded the series of files to memory having the parallel programs to process the data based on the content in cache, which reduced reading cost by more than 200 times
- Examined different data segmentation plans and efficiently distributed computing load to threads
- **Achieved 1861 times speedup** for the computation process compared to the serial program

Efficient Graph Attention Networks (GAT) [Python, PyTorch, Graph Neural Networks]

Oct. 2020 – Mar. 2021

- Proposed a new algorithm to accelerate the training of **GAT** by applying clustering and **linear attention**
- Divided nodes into clusters and aggregated features with linear attention, which resulted in linear time complexity
- Improved classification accuracy by 10% and reduced training time by two times on three graph datasets

PROFESSIONAL EXPERIENCES

Software Engineer Intern | Amazon.com Inc [Python, PySpark, TypeScript, AWS]

May 2022 – Aug. 2022

- Built an **automated training** pipeline of substitutable products group model with **natural language processing** and **hierarchical clustering** models for 20 TB product query data

SKILLS

- **Programming Languages:** C, C++, Java, Python, HTML/CSS/JavaScript, TypeScript, C#, Perl, Shell, Rust
- **Computer Graphics:** Taichi, OpenGL, CUDA, Unity
- **Frameworks:** PyTorch, Tensorflow, PySpark, Node.js, React.js, Bootstrap
- **Others:** AWS, Linux, MongoDB, MySQL