Yaoxu Song

Senior Undergraduate Student

No. 800, Dongchuan Road, Minhang District, Shanghai 200240, China

(+86) 13429635208 | hbeillschmidt@gmail.com | HeydrichBeillschmidt

Education

Shanghai Jiao Tong University - Zhiyuan College — *Shanghai, China*

B.S. in Computer Science (ACM Class) — *Sept. 2019 - Present*

- GPA: 3.72/4.3, TOFEL: 104
- Zhiyuan Honor Fellowship (every year)

Skills

Computer Science	Proficiency in parallelization and acceleration for deep learning models and scientific computation
	Understanding about graph representation learning
Programming	C++, CUDA, Python, Java
Mathematics	Mathematical Analysis, Abstract Algebra, Probability Theory
	Cyclist of Point Set Topology & Differential Geometry
Physics	Statistical mechanics & hydrodynamics
	Cyclist of nonlinear dynamics & soft condensed matter

Publication

µGrapher: High-performance Graph Operator Computation via Unified Abstraction for Graph Neural Networks — *ASPLOS'23*

Coauthor — *Accepted (To Appear)*

- Yangjie Zhou, Jingwen Leng, ***Yaoxu Song***, Shuwen Lu, Mian Wang, Chao Li, Minyi Guo, Wenting Shen, Yong Li, Wei Lin, Xiangwen Liu, Hanqing Wu

Research Experience

SKY Computing — *Berkeley, California, United States*

Undergraduate Research Assistant — *Feb. 2022 - PRESENT*

- Committed in building a high-performance parallel computing library based on JAX for scientific computation
- Supervised by Joseph E. Gonzalez, Associate Professor in SKY Computing at UC Berkeley

John Hopcroft Computer Science Center — *Shanghai, China*

Undergraduate Research Assistant — *Aug. 2021 - Jul. 2022*

- Contributed in building a high-performance graph operator computation framework for Graph Neural Networks
- Supervised by Jingwen Leng, Associate Professor in John Hopcroft Computer Science Center at Shanghai Jiao Tong University

Code Projects

AdaptGear — *Shanghai Jiao Tong University*

Coauthor — *Sept. 2022 - Feb. 2023*

- Proposed AdaptGear, a novel high-performance GNN training system that exploits both the intra- and inter-graph sparsities via adaptive subgraph-level customized kernels
- Submitted to the 20th ACM International Conference on Computing Frontiers (CF '23)

Sci-Alpa — *Berkeley, California, United States*

Author — *Feb. 2022 - PRESENT*

- Extended Alpa, a compiler system for distributed deep learning on GPU clusters, to scientific computation workloads
- Introduced new semantics and optimized performance for several typical scientific computation tasks

µGrapher — *Shanghai Jiao Tong University*

Coauthor — *July. 2021 - April. 2022*

- Built a GNN-specific operator abstraction that incorporates the semantics of graph tensors and graph loops
- Explored various schedule strategies based on the abstraction that can balance the well-established trade-off relationship between parallelism, locality, and efficiency