

# SHENG LI

[!\[\]\(c8d96c8885d3000a912c2582004aed63\_img.jpg\) shl188@pitt.edu](mailto:shl188@pitt.edu) [!\[\]\(3ad821e3ca7dd4cb7003e9c8d982e254\_img.jpg\) Homepage](#) [!\[\]\(177bde115c7ebbeffa559d05eea9e94b\_img.jpg\) LinkedIn](#) [!\[\]\(cab2e95699b614c49dd80341e1932607\_img.jpg\) Google Scholar](#)

## EDUCATION

---

<b>University of Pittsburgh</b> <i>Pittsburgh, PA, USA</i> Ph.D. in Computer Science (with a focus on Efficient Machine Learning Algorithms and Systems) Advisor: Dr. Xulong Tang	<i>Aug. 2021 ~ Present</i>
<b>Sichuan University</b> <i>Chengdu, Sichuan, China</i> B.Eng. in Software Engineering, Graduated with Honor Advisor: Dr. Ziliang Feng	<i>Sept. 2016 ~ June 2020</i>

## PUBLICATIONS

---

*\*equal contribution*

- [1] **Sheng Li\***, Qitao Tan\*, Yue Dai, Zhenglun Kong, Tianyu Wang, Jun Liu, Ao Li, Ninghao Liu, Yufei Ding, Xulong Tang, and Geng Yuan. “Mutual Effort for Efficiency: A Similarity-based Token Pruning for Vision Transformers in Self-Supervised Learning.” *In Proceedings of 13th International Conference on Learning Representations. (ICLR 2025)*
- [2] Huidong Ji, **Sheng Li**, Yue Cao, Chen Ding, Jiawei Xu, Qitao Tan, Ao Li, Jun Liu, Xulong Tang, Lirong Zheng, Geng Yuan, and Zhuo Zou. “A Computation and Energy Efficient Hardware Architecture for SSL Acceleration.” *In Proceedings of 30th Asia and South Pacific Design Automation Conference. (ASP-DAC 2025)*
- [3] **Sheng Li**, Chao Wu, Ao Li, Yanzhi Wang, Xulong Tang, and Geng Yuan. “Waxing-and-Waning: a Generic Similarity-based Framework for Efficient Self-Supervised Learning.” *In Proceedings of 12th International Conference on Learning Representations. (ICLR 2024)*
- [4] **Sheng Li\***, Geng Yuan\*, Yue Dai\*, Youtao Zhang, Yanzhi Wang, and Xulong Tang. “SmartFRZ: An Efficient Training Framework using Attention-Based Layer Freezing.” *In Proceedings of 11th International Conference on Learning Representations. (ICLR 2023, Honored with Spotlight Award)*
- [5] **Sheng Li**, Geng Yuan, Yawen Wu, Yue Dai, Tianyu Wang, Chao Wu, Alex K. Jones, Jingtong Hu, Yanzhi Wang, and Xulong Tang. “ETuner: A Redundancy-Aware Framework for Efficient Continual Learning Application on Edge Devices.” **(Preprint)**
- [6] Tianyu Wang\*, **Sheng Li\***, Bingyao Li, Yue Dai, Ao Li, Geng Yuan, Yufei Ding, Youtao Zhang, and Xulong Tang. “Improving GPU Multi-Tenancy Through Dynamic Multi-Instance GPU Reconfiguration.” **(Preprint)**
- [7] Geng Yuan\*, Yanyu Li\*, **Sheng Li**, Zhenglun Kong, Sergey Tulyakov, Xulong Tang, Yanzhi Wang, and Jian Ren. “Layer Freezing & Data Sieving: Missing Pieces of a Generic Framework for Sparse Training”. *In Proceedings of the 36th Conference on Neural Information Processing Systems. (NeurIPS 2022)*
- [8] Sébastien Ollivier, **Sheng Li**, Yue Tang, Stephen Cahoon, Ryan Caginalp, Chayanika Chaudhuri, Peipei Zhou, Xulong Tang, Jingtong Hu, and Alex K. Jones. “Sustainable AI Processing at the Edge.” *IEEE Micro, 2022. (IEEE Micro)*
- [9] Bingyao Li\*, Qi Xue\*, Geng Yuan\*, **Sheng Li**, Xiaolong Ma, Yanzhi Wang, and Xulong Tang. “Optimizing Data Layout for Training Deep Neural Networks.” *In Companion Proceedings of the Web Conference 2022. (WWW 2022 workshop)*
- [10] Mingzheng Hou, Ziliang Feng, Haobo Wang, Zhiwei Shen, and **Sheng Li**. “An adaptive regression based single-image super-resolution.” *Multimedia Tools and Applications, 2022. (Multimed. Tools Appl.)*

[11] **Sheng Li**, Zanhan Ding, and Honglv Chen. “A Neural Network-Based Teaching Style Analysis Model.” *In 2019 11th International Conference on Intelligent Human-Machine Systems and Cybernetics, IEEE, 2019. (IHMSC 2019)*

[12] **Sheng Li**, and Hanxin Feng. “EEG signal classification method based on feature priority analysis and CNN.” *In 2019 International Conference on Communications, Information System and Computer Engineering, IEEE, 2019. (CISCE 2019)*

## EXPERIENCE

---

### Adobe Research

2025.05 - 2025.11

*Research Scientist/Engineer Intern*

Mentor: Dr. Yifan Gong

- We develop a content-aware patchification router that adaptively adjusts patch size across different spatiotemporal regions of the input video, reducing redundant tokens and achieving up to 1.8× speedup in video generation without affecting the visual quality.
- Draft and submit a manuscript for this project to CVPR 2026.

### University of Pittsburgh

2020 - Present

*Research Assistant / Teaching Assistant*

Advisor: Dr. Xulong Tang

- We propose a temporal-aware token pruning framework for diffusion-based video generation that preserves cross-frame coherence during pruning and achieves up to 1.5× speedup with minimal quality loss.
- We propose an efficient self-supervised learning technique that prunes less important parts of the input image, cutting training costs by over 40% without sacrificing model performance.
- We propose an efficient layer freezing approach that can reduce computation cost by 40% and memory costs by 50% for ML model training, without affecting the model performance.
- We propose to leverage the tensor core to accelerate the process of 3D Gaussian Splatting and the text-to-3D diffusion application.

## SELECTED HONORS AND AWARDS

---

Spotlight Award, The 11th ICLR Conference

2023

National Scholarship, Ministry of Education of China

2018, 2019

Comprehensive Scholarship, Sichuan University

2017

## TEACHING

---

Teaching Assistant

CS 0447: Computer Organization and Assembly Language

Teaching Assistant

CS 1501: Algorithms and Data Structures

## PROFESSIONAL SERVICE

---

Reviewer of NeurIPS, ICLR, ICML, AAAI, ACML, TPAMI, CVPR, and NPL