

Shikang Zheng

✉ zhengshikang0201@gmail.com

EDUCATION

South China University of Technology

Sep 2023 – Jun 2027 (expected)

- B.Eng. in Artificial Intelligence (GPA: 3.77 / 4.00)

University of California, Berkeley

Dec 2024 – May 2025

- Visiting Student of Berkeley Global Access

PUBLICATIONS

PUBLISHED PAPERS:

- [1] **Shikang Zheng**, Guantao Chen, *et al.*, Linfeng Zhang[†]. “*Let Features Decide Their Own Solvers: Hybrid Feature Caching for Diffusion Transformers*”. **ICLR 2026 (Oral Presentation)**. arXiv:2510.04188
- [2] **Shikang Zheng***, Guantao Chen*, *et al.*, Linfeng Zhang[†]. “*From Sketch to Fresco: Efficient Diffusion Transformer with Progressive Resolution*”. **CVPR 2026 (Poster)**. arXiv:2601.07462
- [3] Guantao Chen*, **Shikang Zheng***, Yuqi Lin, Linfeng Zhang[†]. “*Forecast the Principal, Stabilize the Residual: Subspace-Aware Feature Caching for Diffusion Transformers*”. **CVPR 2026 (Poster)**. arXiv:2601.07396
- [4] **Shikang Zheng**, Liang Feng, *et al.*, Linfeng Zhang[†]. “*Forecast then Calibrate: Feature Caching as ODE for Efficient Diffusion Transformers*”. **AAAI 2026 (Poster)**. arXiv:2508.16211
- [5] **Shikang Zheng**, Suraj Sahu, Valery M. Nakariakov, Ding Yuan[†], Song Feng. “*Stellar White-light Emission Reveals Two Types of Flaring Light Curve Patterns*”. **Proceedings of the Royal Society A (Journal)**.
- [6] Liang Feng, **Shikang Zheng**, *et al.*, Linfeng Zhang[†]. “*HiCache: Training-free Acceleration of Diffusion Models via Hermite Polynomial-based Feature Caching*”. **ICLR 2026 (Poster)**. arXiv:2508.16984
- [7] Zhirong Shen, Rui Huang, *et al.*, **Shikang Zheng**, Zhengyi Shi, Liang Feng, Linfeng Zhang[†]. “*Beyond Fixed Formulas: Data-Driven Linear Predictor for Efficient Diffusion Models*”. **CVPR 2026 (Poster)**.

UNDER REVIEW PAPERS:

- [1] Chang Zou*, **Shikang Zheng***, *et al.*, Conghui He, Xuming Hu, Linfeng Zhang[†]. “*Accelerating Diffusion Transformers with Dual Feature Caching*”. Submitted to **IEEE Transactions on Image Processing**. arXiv:2412.18911
- [2] Lixuan He, **Shikang Zheng**. “*A Flat Vocabulary or a Rich Hierarchy? Re-introducing Intrinsic Structure Transforms the Autoregressive Image Generation*”. Submitted to **ICML 2026**. arXiv:2510.04220
- [3] Lixuan He, **Shikang Zheng**, Guantao Chen. *Restoring Generalization in Fine-tuned Multimodal LLMs via Geometric Alignment*. Submitted to **ICML 2026**.
- [4] Jiacheng Liu, *et al.*, **Shikang Zheng**, Linfeng Zhang[†]. “*FreqCa: Accelerating Diffusion Models via Frequency-Aware Caching*”. Submitted to **ICML 2026**. arXiv:2510.08669

RESEARCH EXPERIENCES

Research Assistant

Mar 2026 – Now

MMLab, The University of Hong Kong

- Conducted research on efficient world model, focusing on accelerating World Action Model to real-time generation.
- Experimented with distillation and KV-cache compression for autoregressive video diffusions.
- Explored distillation method for large World Action Model.

Industry Research Collaboration

Nov 2025 – Feb 2026

Alibaba Cloud, Alibaba

- Collaborated with Alibaba on joint research projects, resulting in several co-authored papers on accelerating few-step generative models (e.g., Z-Image / Wuying DiT).

Research Assistant

Sep 2024 – Nov 2025

EPIC-Lab, Shanghai Jiao Tong University

- Conducted research on efficient machine learning, focusing on inference acceleration of Diffusion Transformers.
- Experimented with feature caching and model compression strategies for large-scale generative models.
- Explored KV-cache compression and codebook modeling to accelerate inference and training of autoregressive models.

Research Assistant

Dec 2024 – Mar 2025

Shanghai Artificial Intelligence Laboratory

- Worked on post-training quantization for diffusion models, with a focus on enhancing robustness and security in generative model deployment under constrained environments.

SKILLS

Professional Services: Reviewer of CVPR 2026, ICML 2026, ECCV 2026

Models Familiar with: DiT, FLUX, Qwen-Image, Qwen-Image-Edit, Hunyuan Video, Inf-DiT

Programming Languages: Python, C++, CUDA, Bash, L^AT_EX

Frameworks & Libraries: PyTorch, HuggingFace Diffusers, Transformers, FlashAttention

Systems & Tools: Linux, Windows, Git, Docker, Slurm

COMPETITION EXPERIENCE & SCHOLARSHIP

- First Runner-up, Young Professionals Exhibition & Competition, 2024, Hong Kong.
- First Runner-up, RAICOM AI Robot Competition, 2024 National Finals.
- Second Runner-up, Overwatch University League (OUL), 2024 National Finals.
- Second Runner-up, Overwatch University League, 2025 National Finals.
- Future Technology Taihu Scholarship.
- Hongping Changqing Scholarship.