

# 邹畅

上海交通大学 | 人工智能学院 | 计算机科学与技术



## 基本信息

男 21岁 电话/微信: 13608211202 邮箱: shenyizou@outlook.com  
上海交通大学 人工智能学院 博士生 Ph.D. Student. Advised by Prof. Linfeng Zhang 2026-2031(预估)  
研究方向: 精准高效的图像/视频生成 Precise and Efficient AIGC

电子科技大学 英才实验学院 本科生 Undergraduate 2022-2026

- 均分 93.35, GPA 3.96/4.00, 专业排名 3/26, (从全校范围选拔后的荣誉专业, 全员优先推免)
- CET-4: 556, CET-6: 520; 连续多年获得优秀学生奖学金

## 实习经历

- 腾讯-青云计划 混元大模型方向 多模态模型部 基础模型中心 实习生 2025.05 至今  
作为青云人才计划实习生开展长期实习, 研究下一代图像与视频生成大模型的算法设计, 并在超大规模集群上开展训练。参与HunyuanImage3.0, HunyuanVideo1.5 等基础图像与视频生成模型相关研究, 并以核心贡献成员身份 (Core Contributors) 参与 HunyuanVideo 1.5 技术报告, 主要负责基础模型的特征缓存加速与模型蒸馏技术等基础研究。实习期间主导并开展以可学习的特征缓存为主的相关研究, 提出了 MeanFlow 蒸馏的改进方案, 并结合可学习缓存技术, 实现了在蒸馏后模型上的进一步加速, 成功达到行业领先水准, 相关部分算法改进已被成功应用于多款基础模型, 并于本人实习期间以第一作者身份发表相关论文于计算机视觉领域顶会 CVPR 2026。



Technical Report (Core Contributor): [arXiv:2511.18870]

"HunyuanVideo-1.5 Technical Report" - Tencent Hunyuan Foundation Model Team

## 主要学术成果

\* Equal contribution

本人已在多个高水平学术会议/期刊上以第一作者发表论文4篇, 一作在投论文2篇, 非一作论文发表10篇。

1. [ICLR 2025] "Accelerating Diffusion Transformers with Token-wise Feature Caching."  
- Chang Zou\*, Xuyang Liu\*, Ting Liu, Siteng Huang and Linfeng Zhang.
2. [ICCV 2025] "Features of Diffusion Models are Predictable: Accelerating Diffusion Transformers with Taylor Seers." - Jiacheng Liu\*, Chang Zou\*, Yuanhuiyi Lyu, Junjie Chen and Linfeng Zhang.
3. [ACM MM 2025] "SpeCa: Accelerating Diffusion Models with Speculative Sampling."  
Jiacheng Liu\*, - Chang Zou\*, Yuanhuiyi Lyu, Fei Ren, Shaobo Wang, Kaixin Li, and Linfeng Zhang.
4. [CVPR 2026] "DisCa: Accelerating Video Diffusion Transformers with Distillation-Compatible Learnable Feature Caching." - Chang Zou, Changlin Li, Yang Li, Patrol Li, Jianbing Wu, Xiao He, Songtao Liu, Zhao Zhong, Kailin Huang and Linfeng Zhang.
5. [ICCV 2025] "Efficient Image Editing via Spatial Locality Caching."  
Zexuan Yan, Yue Ma, Chang Zou, Wenteng Chen and Linfeng Zhang.
6. [ACM MM 2025] "Compute only 16 tokens in one timestep: Accelerating Diffusion Transformers with Cluster-Driven Feature Caching." - Zhixin Zheng\*, Xinyu Wang\*, Chang Zou, Shaobo Wang, Linfeng Zhang.
7. [NeurIPS 2025] "EfficientVLA: Training-Free Acceleration and Compression for Vision-Language-Action Models" - Yantai Yang, Yuhao Wang, Zichen Wen, Luo Zhongwei, Chang Zou, Zhipeng Zhang, Chuan Wen and Linfeng Zhang.
8. [AAAI 2026] "Forecast then Calibrate: Feature Caching as ODE for Efficient Diffusion Transformers."  
- Shikang Zheng, Liang Feng, Xinyu Wang, Qinming Zhou, Peiliang Cai, Chang Zou, Jiacheng Liu, Yuqi Lin, Junjie Chen, Yue Ma and Linfeng Zhang.
9. [ICLR 2026] "HiCache: A Plug-in Scaled-Hermite Upgrade for Taylor-Style Cache-then-Forecast Diffusion Acceleration." - Liang Feng\*, Shikang Zheng\*, Jiacheng Liu, Yuqi Lin, Qinming Zhou, Peiliang Cai, Xinyu Wang, Junjie Chen, Chang Zou, Yue Ma and Linfeng Zhang.

10. [ICLR 2026] **“Z-Cache: Accelerating Diffusion Transformers via Self-Reflection.”**  
- Zegang Cheng, Zhikai Wang, Jiacheng Liu, **Chang Zou**, Junjie Chen, Yue Ma, Zhiyuan Ma and Linfeng Zhang
11. [CVPR 2026] **“LESA: Learnable Stage-Aware Predictors for Diffusion Model Acceleration.”**  
- Peiliang Cai, Jiacheng Liu, Haowen Xu, Xinyu Wang, **Chang Zou** and Linfeng Zhang.
12. [CVPR 2026] **“From Sketch to Fresco: Efficient Diffusion Transformer with Progressive Resolution”**  
- Shikang Zheng, Guantao Chen, Lixuan He, Jiacheng Liu, Yuqi Lin, **Chang Zou** and Linfeng Zhang.
13. [CVPR 2026] **“Beyond Fixed Formulas: Data-Driven Linear Predictor for Efficient Diffusion Models”**  
- Zhirong Shen\*, Rui Huang\*, Jiacheng Liu, **Chang Zou**, Peiliang Cai, Shikang Zheng, zhengyi shi, Liang Feng and Linfeng Zhang.
14. [(Under Review) TIP (Major Revision)] **“Rethinking Token-wise Feature Caching: Accelerating Diffusion Transformers with Dual Feature Caching”** - **Chang Zou\***, Shikang Zheng\*, Evelyn Zhang, Runlin Guo, Haohang Xu, Conghui He, Xuming Hu and Linfeng Zhang.
15. [(Under Review) TIP] **“Accelerate Diffusion Transformers with Feature Momentum.”**  
- Jiaxin Fang\*, **Chang Zou\***, Jiacheng Liu, Yuanhuiyi Lyu, Xuming Hu and Linfeng Zhang.
16. [(Under Review) ICML 2026] **“dLLM-Cache: Accelerating Diffusion Large Language Models with Adaptive Caching.”** - Zhiyuan Liu, Yicun Yang, Yaojie Zhang, Junjie Chen, **Chang Zou**, Qingyan Wei, Shaobo Wang, Yichen Zhu, Linfeng Zhang
17. [(Under Review) ICML 2026] **“Let Features Decide Their Own Solvers: Hybrid Feature Caching for Diffusion Transformers.”** - Shikang Zheng, Guantao Chen, Qinming Zhou, Yuqi Lin, Lixuan He, **Chang Zou**, Peiliang Cai, Jiacheng Liu and Linfeng Zhang.
18. [(Under Review) ICML 2026] **“FreqCa: Accelerating Image generation and editing via Frequency-Aware Caching”** - Jiacheng Liu\*, Peiliang Cai\*, Qinming Zhou, Yuqi Lin, Deyang Kong, Benhao Huang, Yupei Pan, Haowen Xu, **Chang Zou**, Junshu Tang, Shikang Zheng and Linfeng Zhang.
19. [(Under Review) ECCV 2026] **“Accelerating Diffusion Models with Gaussian Process Rectified Cache”** - Zhirong Shen\*, Rui Huang\*, **Chang Zou**, Shikang Zheng, Jiacheng Liu, Peiliang Cai, zhengyi shi, Zheng Lu, Liang Feng, tuxiaobing, Jinkui Ren, Xiantao Zhang and Linfeng Zhang.
20. [(Under Review) ECCV 2026] **“AViTS: Adaptive Spatiotemporal Token Selection for Efficient Dynamic-Resolution Generation”** - Haoran Qin, Zhengan Yan, Shikang Zheng, Xiaobing Tu, Jiacheng Liu, Yuqi Lin, **Chang Zou**, JinShan Liu, Peiliang Cai, Xiantao Zhang, Jinkui Ren and Linfeng Zhang.