

Yifei Liu | Curriculum Vitae

✉ yifei23liu@gmail.com • 🌐 kaikai23.github.io

Education

Beihang University (BUAA) and Shanghai AI Laboratory

Ph.D. Candidate

Co-supervised by Prof. Yu Qiao and Prof. Xiao Sun.

Research focus: sparse computer vision and vlm spatial understanding.

Beijing/Shanghai, China

2024.9–present

University of Zurich

M.Sc. in Informatics, focus: artificial intelligence

GPA: 5.9/6.0 (rank 1/180) with distinction

Supervisor: Prof. Davide Scaramuzza.

Zurich, Switzerland

2020.9–2023.6

University of Science and Technology of China (USTC)

B.Sc. in Mathematics and Applied Mathematics

Hefei, China

2016.9–2020.6

Experience

Shanghai AI Laboratory and BUAA

Ph.D. Researcher

Research on 3D spatial intelligence, Gaussian splatting, sparse visual representations, and visual language models.

China

2024.9–present

Shanghai AI Laboratory

Research Assistant / Computer Vision Engineer, mentor: Xiao Sun

- Developed universal sports field calibration system, which laid the foundation for the team's 3D sports analysis.
- Core contributor on vision models for Olympic 2024 badminton and tennis tactical analytical broadcast, and 2024 China Open (tennis) tactical analysis; the latter was used in CCTV 5 live broadcast.

Shanghai, China

2024.1–2024.9

ETH Zurich

Research Assistant, mentor: Prof. Otmar Hilliges

- Worked on using depth prior to improve human-object reconstruction from a single view under self- and mutual-occlusion.
- Published a best paper in ECCV 2024 T-CAP workshop.

Zurich, Switzerland

2023.6–2023.11

Robotics and Perception Group (RPG)

Master Thesis, supervisor: Dr. Mathias Gehrig & Prof. Davide Scaramuzza

Selective ViT for object detection and instance segmentation.

- First proposed to prune tokens in dense tasks, which preserve and reactivate pruned tokens for improved performance.
- One WACV paper.

Zurich, Switzerland

2023.1–2023.6

Research Interests

Computer vision; 3D spatial intelligence; novel view synthesis; 3D Gaussian splatting; visual language models; sparse attention and token pruning; SLAM and robotic perception; agentic reinforcement learning for spatial tasks.

Selected Publications

- Yuanyuan Gao*, Hao Li*, **Yifei Liu***, Xinhao Ji*, Yuning Gong*, Yuanjun Liao, Fangfu Liu, Manyuan Zhang, Yuchen Yang, Dan Xu, et al., "Holi-Spatial: Evolving Video Streams into Holistic 3D Spatial Intelligence", *ICML*, 2026. **(Oral)**
Summary: We build an automatic pipeline that turns raw videos into accurate and semantically aligned 3DGS and constructs instance- and camera-level spatial QA data; finetuning on it can improve open-source VLMs across multiple public spatial benchmarks.
- Yifei Liu***, Yuanyuan Gao*, Yifan Zhan, Sheng Xu, Yuning Gong, Jiaqi Chen, Yuchen Yang, Dingwen Zhang, Yinqiang Zheng, Xiao Sun, Zhihang Zhong, "MaskGaussian++: Scalable and Neural Gaussian Splatting with Probabilistic Masks", *TPAMI*, under review. (Major Revision)
Summary: We revisit and adapt masked rasterization to neural GS, 4DGS, and large-scale GS, showing that masked Gaussians are a generally efficient representation. We also identify a log-linear scaling law between the number of primitives and rendering quality, and show that using masks can lead to 2x more efficient scaling.

- Yuanyuan Gao, Yuning Gong, **Yifei Liu**, Jingfeng Li, Dingwen Zhang, Yanci Zhang, Dan Xu, Xiao Sun, Zhihang Zhong, “Proxy-GS: Unified Occlusion Priors for Training and Inference in Structured 3D Gaussian Splatting”, *CVPR*, 2026. **(Oral, Best Paper Award Candidate)**
 Summary: We render a depth map from a proxy mesh to cull occluded Gaussians, eliminating long-overlooked occlusion redundancy in 3DGS. This hybrid rendering strategy leads to large speedups for large-scale rendering under heavy occlusions.
- Yuning Gong, **Yifei Liu**, Yifan Zhan, Muyao Niu, Xueying Li, et al., “Visionary: The World Model Carrier Built on WebGPU-Powered Gaussian Splatting Platform”, *arXiv*, 2025.
 Summary: A web-native WebGPU platform for real-time Gaussian splatting, mesh rendering, and per-frame ONNX inference.
- Yifei Liu**, Zhihang Zhong, Yifan Zhan, Sheng Xu, Xiao Sun, “MaskGaussian: Adaptive 3D Gaussian Representation from Probabilistic Masks”, *CVPR*, 2025.
 Summary: We propose a learnable mask for selecting 3D Gaussians during training by incorporating it into the strict rendering equation of 3DGS, and write a custom backward CUDA kernel to avoid vanishing gradients for zeroed Gaussians. This leads to 2x–3x faster rendering with almost no quality degradation.
- Yifei Liu**, Mathias Gehrig, Nico Messikommer, Marco Cannici, Davide Scaramuzza, “Revisiting Token Pruning for Object Detection and Instance Segmentation”, *WACV*, 2024.
 Summary: We first adapt token pruning to dense prediction tasks, where directly discarding tokens hurts dense feature maps. By dynamically reactivating unselected tokens, we achieve high pruning rates with little performance loss.
- Alexey Gavryushin, **Yifei Liu**, Daoji Huang, Yen-Ling Kuo, Julien Valentin, Luc Van Gool, Otmar Hilliges, Xi Wang, “ROMEO: Revisiting Optimization Methods for Reconstructing 3D Human-Object Interaction Models From Images”, *T-CAP Workshop at ECCV*, 2024. **(Best Paper Award)**
 Summary: Using depth prior to reconstruct human-object interactions from a single view under self- or mutual occlusion.