

Yaoliang Bian

First-year Ph.D. Student | Westlake University & Shenzhen Loop Area Institute

+86 183-2637-2810 | bianaoliang@westlake.edu.cn | jeffyaoliang.github.io

Research Interests

Embodied AI and multimodal foundation models, with a focus on (1) world action models and physical agents; (2) cross-modal alignment and projector mechanisms in vision-language models; (3) distributed training and fine-tuning of VLA models; and (4) long-horizon memory construction and retrieval for embodied agents. Hands-on experience with multi-GPU distributed training (DDP), LoRA fine-tuning, and inference deployment (vLLM); day to day I drive experiments and cluster jobs with AI agents such as Claude Code.

Education

Westlake University & Shenzhen Loop Area Institute — Joint Ph.D. Program Sep 2026 – Present

Advisors: Prof. [Haibin Ling](#), Prof. [Hengshuang Zhao](#). First-year Ph.D. student; expected graduation: June 2031.

University of Science and Technology of China (USTC) Sep 2022 – Jun 2026

B.Eng. in Electronic Science and Technology.

Publications and Open-Source Contributions

Benchmark Radar: A Living Database and Search Engine for AI Benchmarks and Evaluation

Koutian Wu, Wenjun Zhang, Junjie Zhou, Pengqian Han, Qiming Bao, Ergan Shang, Jiayu Wang, Junkai Wang, Wanghan Xu, Songyuanyi Lu, Yixuan Li, Rudy Mizrahi Celekli, Lazizbek Ravshanov, **Yaoliang Bian**, Xudong Han, Lin Shi.

arXiv:2609.11115, 2026 (v4). **Hugging Face Daily Papers: #1 Paper of the Day.**

- **Co-author and open-source contributor.** Contributed an agent-benchmark digest connector, evidence preservation across paper revisions, recovery of cited benchmark release dates, and handling of unreported adoption counters. Merged contributions: [#632](#), [#840](#), [#834](#), [#831](#).

Honors and Awards

USTC Sunshine Scholarship (1/49) 2025

USTC “Outstanding Student Leader”, “Outstanding Youth League Cadre”, “Outstanding Youth League Member” 2025

McMaster-USTC Summer Research Internship Program — Completion Certificate 2024

USTC Fellowship Undergraduate A-Class Funding (Top 5%) 2023

Robogame Competition at USTC, 2nd Prize (2/39) 2023

USTC Sunshine Scholarship (1/45); Scholarship for Electrical Engineering Program of Excellence; “Star of Diligent Learning” 2023

Research Highlights

MemAff: Retrieving and Grounding Affordance Objects in Long Visual Memory (Memory-Grounded Affordance) May 2026 – Present

- **Task and benchmark.** Defined memory-grounded affordance prediction: given a past stream of visual observations and one intent-style instruction, retrieve the frame containing the target object and localize it at the pixel level, with the two stages scored jointly under a strict end-to-end mIoU. Built **MemAff-Bench: 3,628 tasks** drawn from three real sources (Ego4D, EPIC-KITCHENS, PAP), filtered for object-level uniqueness with human spot-checking, and scaling memory in controlled steps from **120 to 960 frames**.
- **Main result.** Piling on context does not solve it: single-pass VLMs that ingest everything hit a context-window or memory wall first, and even a 32B model reading 960 frames in a 262k window picks the right frame only **6–21%** of the time. The difficulty is choosing the one frame among hundreds of near-duplicates, not fitting them into context.
- **Method and diagnosis.** Proposed training-free **MemAff** (three-way coarse retrieval to 8 candidates → VLM rerank → Grounding-DINO + SAM 2): frame accuracy matches or beats a 4× larger single-pass Qwen3-VL-32B while using **14–116×** fewer tokens per query. Ablations localize the remaining bottleneck in coarse recall (65–84% of the loss at 960 frames comes from the target never entering the top 8); a larger reranker does not help, and a trained selector does not transfer across sources.
- **Output: First-author** paper submitted to **ICLR 2027** (under review). Comparison systems — A4-Agent, Affordance-R1, Qwen3-VL-8B/32B one-shot, and retrieval-style ablations — all run in their best known configuration.

The Attention-Sink Cliff in VLA Finetuning: Diagnostic Probes Reveal the Load-Bearing Mechanism Mar 2026 – Jul 2026

- **Diagnosed the attention-sink cliff.** Built a 4-GPU PyTorch DDP pipeline (OpenVLA-7B LoRA fine-tuning, 200K steps, BF16, full checkpoint resume) and ran a controlled sweep over 14 projector normalization variants: success rate forms a **cliff-plateau** structure in per-token norm T . MLLM calibration recipes that compress vision tokens to the text scale (+Normalize from *Beyond Semantics*, Qi et al. 2025; LayerNorm- γ from *The Unseen Bias*, Li et al. 2025) fall into the cliff (**64.8–71.4%**), while baselines preserving the original magnitudes hold a plateau (**77–82%**) across a $16\times T$ range, a >10 pp gap. Reproduced the official checkpoint within 0.6 pp, ruling out setup artifacts.
- **Localized the mechanism and tested its cross-architecture predictive power.** Attention probes attribute the cliff to BOS attention-sink collapse and vision tokens crowding out the attention budget, locked in early in fine-tuning (resuming from a cliff checkpoint with only γ unfrozen leaves T nearly unchanged and costs a further 3.4 pp,

beyond recovery). A **custom cross-framework probe** (OpenVLA PyTorch $\leftrightarrow$ $\pi_{0.5}$ JAX, via `jax.debug.callback` inside jit) validated the mechanism on Gemma-2B $\pi_{0.5}$, whose norm ordering is reversed: under the same perturbation ($T=1$) OpenVLA's visual attention spikes $4\times$ and collapses, while $\pi_{0.5}$'s decreases and success drops only 4.6 pp — as the mechanism predicts.

- **Extension: independently reproduced and challenged *The Unseen Bias: How Norm Discrepancy in Pre-Norm MLLMs Leads to Visual Information Loss* (Li et al., ICLR 2026).** Its headline text-only MMLU +6.41 gain is largely recovery from catastrophic forgetting (native Llama-3.2-3B MMLU 60.6 vs. their baseline 45.2); on a non-forgetting Qwen backbone the gain collapses to +0.72.

Industry Experience

Known AI — Algorithm Intern

Dec 2025 – Apr 2026

Advisor: Prof. [Yingcong Chen](#) (Information Hub, [HKUST\(GZ\)](#)).

- **M3-Agent reproduction and deployment** (ICLR 2026, ByteDance Seed). The system builds a persistent memory graph from streaming video by encoding visual, auditory, and spatial signals into a unified graph, and answers long-horizon questions through multi-turn retrieval. Reached **30.9% accuracy** on M3-Bench-robot (1,276 QA, 100 indoor scenes), matching the paper's 30.7%, while **cutting per-query latency from 30s to 4s ($7\times$ speedup)** via 50+ controlled experiments on retrieval parameters, embedding models, and batched inference.
- **Speaker-identity pipeline diagnosis.** Found a three-stage cascading failure: (1) SpeakerLab clustering fragmented 3 speakers into 47 IDs; (2) VLM equivalence judgement produced false positives; (3) Union-Find transitive merging collapsed all IDs into one identity. Replaced the speech model with **CAM++ (3D-Speaker)** embeddings and adaptive- k spectral clustering, **reducing per-scene speaker fragments from 25.4 to 5.4**.
- **Closed-loop fault-recovery framework with episodic memory.** Designed a Detect $\rightarrow$ Reflect $\rightarrow$ Correct $\rightarrow$ Memorize $\rightarrow$ Pre-check loop synthesizing REFLECT (CoRL 2024), Inner Monologue (CoRL 2022), and Reflexion (NeurIPS 2023); the memory and pre-check components extend prior work.
- **Long-term memory retrieval for LLM agents.** Proposed an embedding-free retrieval scheme based on LLM query rewriting + BM25; matches the embedding baseline overall on LoCoMo and **outperforms it by 18.4% on temporal-reasoning subsets**. Provided a theoretical analysis of why embeddings struggle with multi-hop temporal reasoning.
- **Output:** Multiple manuscripts in preparation.

Fiveages Ltd — VLA Algorithm Intern

Sep 2025 – Dec 2025

Advisor: Prof. [Yan Huang](#) (NLPR, Institute of Automation, Chinese Academy of Sciences).

- Worked on VLA training optimization and vision-backbone replacement. On top of [BridgeVLA](#), processed and integrated the VastTrack dataset for pre-training, then ran distributed re-training to improve cross-scene generalization.
- Replaced the fine-tuning-stage PaLiGemma vision backbone with the [SANSa](#) vision network to strengthen object segmentation and visual grounding. Surveyed VLA + RL literature and explored RL-VLM combinations.

Selected Projects

USTC Robogame 2023

May – Oct 2023

Built an autonomous robot that recognizes and grasps minerals of different colors and delivers them to designated zones on schedule. Used AprilTag for real-time localization, and applied bilateral filtering, edge detection, and HSV-space color segmentation for ore classification. Second Prize among 39 teams.

AI Final-Exam Review Agent — Nankesong Hackathon

Jun 28, 2025

Developed a Gemini 2.5 Flash-based agent that lets students upload course material and past exams to build a personal knowledge base, then generates predicted exam questions linked back to the relevant knowledge points. Built both backend and frontend UI. Won the top “Star of Nanke” prize and 1st place at WaytoAGI Nanjing.

Custom Hardware for Deep Reinforcement Learning

Jul – Sep 2024

Advisor: Prof. [Ameer Abdelhadi](#) (ECE, [McMaster University](#)).

Surveyed GPU and FPGA implementations of DQN, PPO, A3C, and SAC; compared parallelization strategies and bottlenecks across hardware architectures. Authored a comprehensive survey of RL hardware acceleration and proposed near-memory computation + advanced layer fusion as future directions.

International Exchange

[National University of Singapore](#) — Dexterous Manipulation Research

Jul – Sep 2025

[University of Texas at Austin](#) — Software Engineering Winter School

Winter 2024

University of Oxford, Oriel College — Quantum Computing Summer School

Summer 2023

University of Warsaw — Information Technology Winter School

Winter 2023

Skills

Programming: Python, PyTorch, C, C++, Java, Verilog

Training: PyTorch DDP, JAX/Flax (FSDP), LoRA / PEFT, BF16 mixed precision, gradient checkpointing, checkpoint resume, vLLM inference deployment

Systems: Git, Docker, SLURM cluster scheduling, GCP ($4\times$ A100), Linux

Hardware: FPGA, STM32, Raspberry Pi, nRF52840 SoC

English: CET-6 568 · CET-4 597 · TOEFL 85