Bohan Liu

Machine Learning & AI Systems

Charlottesville, VA | bohan@virginia.edu | bohanliu.com
github.com/Liu-524 | linkedin.com/in/bohan-liu | Google Scholar

Computer science PhD student combining multimodal AI research with production market data engineering. Focus: secure and efficient models, faithful retrieval-augmented generation (RAG), distributed training, and agent systems.

Experience

PhD Researcher — University of Virginia, Machine Learning Group
Jan 2025 – Present | Advisor: Prof. Aidong Zhang

- Developed SamplingTAR, a training-free defense against typographic attacks: **+22.8 percentage points** in object accuracy on RTA-100 (CLIP ViT-H/14), with near-zero inference overhead.
- Co-developed RAGLens: lightweight hallucination detection using sparse autoencoders, with interpretable feedback to correct unfaithful outputs; **ICLR 2026**.
- Co-authored Task Operators, encoding in-context demonstrations as attention-weight updates and scaling many-shot learning without longer context windows; **NeurIPS 2026**.

SamplingTAR: first author, ECCV 2026 | Code | 97 citations across 8 listed papers (Scholar, Oct 2026)

Software Developer Intern — DRW Holdings, Chicago
Aug 2022 – Dec 2022

- Built a Java market data ingestor to exchange specifications; diagnosed real-time data errors, developed a React availability monitor, and refactored a Ruby on Rails delivery API.
- Optimized PostgreSQL transactions and built a multithreaded Java traffic repeater for **24/7 Kubernetes operation**, with automated Splunk data fetching and error logging.

Research Assistant — University of Illinois Urbana-Champaign
May 2022 – May 2023 | Shenlong Wang's Research Group

- Contributed to UrbanIR for urban scene inverse rendering (**3DV 2025**); customized NeRF-OSR training on SLURM clusters and integrated Unreal Engine rendering comparisons.

Technical Skills

Deep learning: PyTorch; Hugging Face Transformers, Accelerate, TRL, PEFT; LoRA/QLoRA; supervised fine-tuning (SFT), preference optimization (DPO), reinforcement learning (GRPO).

Training & inference: DeepSpeed, distributed training (DDP/FSDP, ZeRO, NCCL), SLURM; vLLM, continuous batching, KV caching, quantization, FlashAttention, tensor parallelism.

Agents & security: OpenClaw, multi-agent orchestration, tool calling, Model Context Protocol (MCP), sandboxing, least privilege, prompt-injection defenses, evaluation-driven self-improvement.

Systems & software: Python, Java, JavaScript/React, SQL; Linux, Docker, Kubernetes, PostgreSQL; networking (TCP/IP, HTTP/RPC), concurrency, CUDA, CPU/GPU architecture, memory hierarchy.

Selected Projects

Bili-Loop | 2021 – Present — Shipped a browser extension for looping videos on bilibili.com; **200+ users** through the Microsoft Store.

ALFWorld Curriculum Learning | Jan – May 2024 — Designed an embodied-agent curriculum using plan and action length statistics; CS 598, University of Illinois Urbana-Champaign.

Education

PhD, Computer Science — University of Virginia | Expected May 2030

Master of Computer Science; BS, Computer Science — University of Illinois Urbana-Champaign
May 2024 | GPA: **4.00**/**4.00** in both degrees | BS: Highest Honor, Bronze Tablet Award