# Bohan Liu

PhD Student in Computer Science | University of Virginia

Charlottesville, Virginia | Email: bohan@virginia.edu | bohanliu.com
GitHub: github.com/Liu-524 | LinkedIn: linkedin.com/in/bohan-liu

Computer science PhD student focused on multimodal models, robustness, and interpretability. Research in faithful retrieval-augmented generation; industry experience in market data systems.

## Technical Skills

**Machine learning:** PyTorch, Transformers, LLaVA, Qwen-VL, model fine-tuning and evaluation, sparse autoencoders, retrieval-augmented generation (RAG), vision-language reasoning.

**Languages and frameworks:** Python, Java, JavaScript, React.js, C#, Unity, Ruby on Rails.

**Systems:** PostgreSQL, Kubernetes, Linux, SLURM, high-performance computing, Splunk.

## Education

**PhD in Computer Science, University of Virginia**
Expected May 2030 | Advisor: Prof. Aidong Zhang

**Master of Computer Science, University of Illinois Urbana-Champaign**
May 2024 | GPA: 4.00/4.00

**Bachelor of Science in Computer Science, University of Illinois Urbana-Champaign**
May 2024 | Highest Honor | GPA: 4.00/4.00

## Research and Industry Experience

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

- Study spurious biases and robustness in multimodal large language models through MM-SpuBench (KDD 2026), SAGE (AAAI 2026), and training-free defenses against typographic attacks.
- Investigate sparse autoencoders for faithful RAG (ICLR 2026), latent computation, and mechanistic interpretability of language-model reasoning (EMNLP 2026).

**Graduate Research Assistant — University of Virginia, Collaborative Robotics Lab**
Aug 2024 – Dec 2024

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

- Researched large-scale urban scene inverse rendering from a single video (3DV 2025); investigated NeRF-OSR relighting models and customized training on SLURM clusters.
- Integrated luma-plugin with Unreal Engine for rendering comparisons.

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

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

**Research Assistant — Illinois Space Grant Consortium**
May 2022 – Jul 2022 | Undergraduate Research Program

- Built Unity structural-monitoring software for Oculus Quest 2 and Vive Pro Eye; extended ARENA-xr for web-based multi-agent virtual reality with client-side computation.

**Research Member — University of Illinois Urbana-Champaign**
Oct 2021 – May 2022 | Mobility and Fall Prevention Research Lab

- Automated EEG processing and statistical analysis of anxiety-related channels; evaluated state-anxiety classifiers in virtual reality (IEEE EMBC 2022).

## Selected Publications

1. **Towards Robustness against Typographic Attack with Training-free Concept Localization**
   **Bohan Liu**, Wenqian Ye, Guangzhi Xiong, Zhenghao He, Sanchit Sinha, Aidong Zhang. **ECCV 2026**, pp. 598–616.
2. **Triggering Chain-of-Thought via Latent Feature Interventions in Large Language Models**
   Zhenghao He, Guangzhi Xiong, **Bohan Liu**, Sanchit Sinha, Aidong Zhang. **EMNLP 2026** (Main Conference, accepted).
3. **MM-SpuBench: Towards Better Understanding of Spurious Biases in Multimodal LLMs**
   Wenqian Ye, **Bohan Liu**, Guangtao Zheng, Di Wang, Yunsheng Ma, Xu Cao, Bolin Lai, James M. Rehg, Aidong Zhang. **KDD 2026** (Dataset and Benchmark Track).
4. **Toward Faithful Retrieval-Augmented Generation with Sparse Autoencoders**
   Guangzhi Xiong, Zhenghao He, **Bohan Liu**, Sanchit Sinha, Aidong Zhang. **ICLR 2026.**
5. **SAGE: Spuriousness-Aware Guided Prompt Exploration for Mitigating Multimodal Bias**
   Wenqian Ye, Di Wang, Guangtao Zheng, **Bohan Liu**, Aidong Zhang. **AAAI 2026.**

Research impact: 97 citations across 8 Scholar-listed papers (snapshot supplied Oct 2026).

Full publication list: bohanliu.com/pubs.html | Google Scholar

## Projects

**Bili-Loop — Browser Extension**
2021 – Present

- Built a Chrome extension for looping videos on bilibili.com; distributed on the Microsoft Store with 200+ users.

**Curriculum Learning on ALFWorld — University of Illinois Urbana-Champaign**
Jan 2024 – May 2024 | CS 598 Embodied NLP | Equal Contributor

- Designed a curriculum using plan and action length statistics, with potential for automatic curriculum generation using large language models.

**Few-shot Document Classification — University of Illinois Urbana-Champaign**
Aug 2023 – Jan 2024 | CS 512 Data Mining | Team Lead

- Estimated event role affinity with RoBERTa prompting; used graph neighborhoods for zero-shot role merging and language-model embeddings with graph convolutional networks for document clustering.

## Teaching and Mentoring

**Graduate Teaching Assistant — University of Illinois Urbana-Champaign**

- CS 415 Game Development | Jan 2024 – May 2024
- CS 410 Text Information System | Aug 2023 – Dec 2023

**WYSE Mentor — Worldwide Youth in Science and Engineering, UIUC**
Jun 2022 – Aug 2022

- Mentored two high school students in introductory computer science and 3D application development.

**Course Assistant — University of Illinois Urbana-Champaign**
Jan 2020 – May 2020 | CS 125 Introduction to Computer Science

- Guided a 20-student lab in coding and debugging; held office hours on algorithms, Java, and Android Studio.

## Professional Service

- **Organizer:** 5th Workshop on Large Language and Vision Models for Autonomous Driving (LLVM-AD), WACV 2026.
- **Reviewer:** AAAI 2027; NeurIPS 2026.