

HAOMING WANG

PH.D. CANDIDATE

+1 412-909-7571

haw200@pitt.edu

Pittsburgh, PA, USA

[Homepage](#) [Scholar](#) [LinkedIn](#) [GitHub](#)

SUMMARY

- **Full-Stack ML Execution:** Bridging the gap between theoretical AI and real-world application. Consistently architect and validate end-to-end VLM/LLM pipelines on physical hardware.
- **Multimodal Foundation Models:** Deep domain expertise in visual spatial reasoning. Systematically diagnosed model bottlenecks via latent-state analysis and benchmarking, and developed novel training-free inference pipelines and latent regularization methods for performance improvement.
- **Recognized Academic Excellence:** First-author publications across top-tier computer vision, AI, and mobile systems venues (CVPR'26 Oral, ECCV'26, MobiCom'26, AAAI'25, MobiSys'25, MobiCom'25).

EDUCATION

Sept 2022 – May 2027 (Anticipated)	Ph.D. in Electrical and Computer Engineering University of Pittsburgh, Pittsburgh, PA, USA Advisor: Prof. Wei Gao, Intelligent System Lab	
Sept 2018 – May 2022	B.Eng. in Automation Zhejiang University, Hangzhou, China Chu Kochen Honors College	GPA: 3.8/4.0

EXPERIENCE

Graduate Student Researcher

Intelligent System Lab, University of Pittsburgh, Pittsburgh, PA, USA

- **Spatial Reasoning Performance Improvement (Sep 2025 – Present).** Improved VLM spatial reasoning across training and inference: Dirichlet-energy regularizer shaping a latent subspace encoding 3D scene topology [2]; training-free multi-view fusion deployed on Jetson Orin, AR Glass, and smartphone [3]; co-developed VLM-based autonomous mobile systems for small-object search and infrastructure-assisted navigation on real quadruped and wheeled robots.
- **Spatial Reasoning Benchmarking and Analysis (May 2025 – Jan 2026).** Established benchmarks and diagnostic methodologies for VLM spatial reasoning: customizable framework generating photo-realistic 3D test scenes with user-controllable complexity axes [1]; cognitively grounded multi-view benchmark with reasoning-path / latent-state analysis [4].
- **Efficient Model Training on Mobile Devices (Sep 2022 – May 2025).** Designed efficient training systems for heterogeneous mobile / IoT fleets: explainable warm-start for on-device LLM personalization [5]; gradient-corrected federated AIoT training under data and device heterogeneity [6]; staleness-tolerant federated learning via gradient inversion [7].

SELECTED PAPERS (FIRST-AUTHOR)

1. **[CVPR'26 Oral] InfiniBench: Infinite Benchmarking for Visual Spatial Reasoning with Customizable Scene Complexity**
Haoming Wang, Qiyao Xue, Wei Gao.
IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2026 (Acceptance 25.4%, Oral).
[\[arXiv\]](#) [\[Code\]](#) [\[Data\]](#)
2. **Uncovering and Shaping the Latent Representation of 3D Scene Topology in Vision-Language Models**
Haoming Wang, Wei Gao.
Preprint, 2026. [\[Paper\]](#) [\[Code\]](#)

3. **[MobiCom'26] SpatialMind: Spatially Aware On-Device Embodied AI via Viewpoint Integration**
Haoming Wang, Qiyao Xue, Weichen Liu, Wei Gao.
ACM International Conference on Mobile Computing and Networking (MobiCom), 2026. [\[Paper\]](#) [\[arXiv\]](#)
4. **[ECCV'26] Reasoning Path and Latent State Analysis for Multi-view Visual Spatial Reasoning: A Cognitive Science Perspective**
 Qiyao Xue, **Haoming Wang (co-first)**, Weichen Liu, Shiqi Wang, Yuyang Wu, Wei Gao.
European Conference on Computer Vision (ECCV), 2026. [\[arXiv\]](#)
5. **[MobiSys'25] Never Start from Scratch: Expediting On-Device LLM Personalization via Explainable Model Selection**
Haoming Wang, Boyuan Yang, Xiangyu Yin, Wei Gao.
Proceedings of the 23rd ACM International Conference on Mobile Systems, Applications and Services (MobiSys), 2025 (Acceptance 18.0%). [\[Paper\]](#)
6. **[MobiCom'25] When Device Delays Meet Data Heterogeneity in Federated AIoT Applications**
Haoming Wang, Wei Gao.
Proceedings of the 31st ACM International Conference on Mobile Computing and Networking (MobiCom), 2025 (Acceptance 17.1%). [\[Paper\]](#)
7. **[AAAI'25] Tackling Intertwined Data and Device Heterogeneities in Federated Learning with Unlimited Staleness**
Haoming Wang, Wei Gao.
Proceedings of the 39th AAAI Conference on Artificial Intelligence, 2025 (Acceptance 23.4%). [\[arXiv\]](#)
8. **[ECCV'26 Workshop] FreezeAsGuard: Mitigating Illegal Adaptation of Diffusion Models via Selective Tensor Freezing**
 Kai Huang, **Haoming Wang (co-first)**, Wei Gao.
LifeGenIP Workshop at the European Conference on Computer Vision (ECCV), 2026. [\[Paper\]](#) [\[Code\]](#)

Full publication list (including co-authored papers on VLM-based navigation) available on [My homepage](#).

HONORS & AWARDS

April 2026 **ECE Research Assistant of the Year**
 Department of Electrical and Computer Engineering, University of Pittsburgh

TECHNICAL SKILLS

Programming & ML Frameworks: Python, PyTorch, C/C++, Hugging Face Transformers, vLLM, DeepSpeed.
On-Device Inference & Optimization: Model quantization and acceleration using TensorRT, ONNX Runtime, and NNAPI; deployed end-to-end pipelines on Jetson AGX Orin and Android smartphones.
Model Training & Alignment: SFT, GRPO / DPO, LoRA / PEFT;
Embodied AI & Simulation: Isaac Sim, Habitat, Blender.