Yuxiang Wei

Ph.D. Student | Multimodal Foundation Models · LLM Post-Training · Evaluation & Benchmarking

weiyuxiang@gatech.edu (404) 203-8538 yuxiangwei0808.github.io Scholar GitHub LinkedIn

Research Profile

I build **multimodal foundation models for fMRI** that align brain activity with language and vision for neural decoding, stimulus reconstruction, and clinical neuroimaging. Beyond neuroscience, I study **LLM post-training**, **long-term memory**, and **evaluation & benchmarking** for coding and language agents. My work includes **11 first-author papers** and two additional equal-contribution papers.

Education

Georgia Institute of Technology, Ph.D. in Electrical and Computer Engineering – Atlanta, GA — Expected May 2028

- Advisor: Vince D. Calhoun

Georgia Institute of Technology, M.S. in Electrical and Computer Engineering – Atlanta, GA — May 2025

UESTC, B.S. in Electronic Information Engineering – Chengdu, China — June 2023

Research & Industry Experience

Research Intern, Handshake AI – San Francisco, CA — Summer 2026

- Led research on **interactive benchmarks for LLM coding agents**, designed a grounded **user simulator** to control specification, clarification, and execution feedback across **3 benchmarks, 3 models, 5 agent harnesses, and 11.5K runs** revealed how **information access and interaction design** shape measured agent capability.
- Developed an item-level **LLM prompt-robustness evaluation** across **6 models and 9 benchmarks**, revealing how **background brittleness** and aggregate metrics can obscure perturbation effects and distort model comparisons.
- Contributed to **LLM post-training** pipelines for realistic **RLEs**, integrating **GSPO**, rubric-based rewards, and held-out evaluation to improve agents' execution of **multi-step, tool-using workflows**.

Graduate Research Assistant, TReNDS Center, Georgia Institute of Technology – Atlanta, GA — Aug 2023 – present

- Designed a **stateful world model of fMRI dynamics** that combines observed brain history with multimodal sensory inputs to simulate future neural trajectories, enabling **closed-loop forecasting**, alternative-stimulus simulation, and **in-silico experimental planning**.
- Built **fMRI-LM**, a **multimodal foundation model** that **post-trains** an LLM via **RLVR** and preference tuning so that it can understand fMRI tokens and reason over the data for diverse tasks (prelimary work accpeted to **CVPR 2026**).
- Developed **MoRE-Brain**, an **interpretable hierarchical MoE** whose dynamic space-time routing reflects brain functional specialization and guides a **diffusion model** for visual reconstruction, enabling cross-subject adaptation with **~90% fewer trainable parameters** (**NeurIPS 2025**).

Research Collaborator, Quanta Alpha — Dec 2025 – present

- Designed a modular framework for **long-term dialog memory** spanning representation, organization, retrieval, and maintenance; controlled experiments characterized when graph complexity helps and when it adds diminishing returns. **ACL 2026**
- Introduced **CloneMem**, a longitudinal AI-clone benchmark grounded in non-conversational digital traces, and identified persistent failures in **temporal coherence** (**ACL 2026**).

Applied Scientist Intern, Amazon – Austin, TX — May 2025 – Aug 2025

- Diagnosed **GNN over-smoothing** and developed a **graph-aware Transformer** with graph-derived positional encodings, paired with a **spatiotemporal diffusion model** for probabilistic **time-series forecasting** across a **42K+ node network**.
- Built a counterfactual framework for out-of-distribution conditions and optimized large-scale training and inference, reducing **WAPE by >30%** versus production models.

Applied Scientist Intern, Tencent – Shenzhen, China — May 2024 – Aug 2024

- Developed **causal inference** methods for **multi-treatment, multi-task uplift modeling** and bias-aware estimation on nonrandom observational data, contributing to a **>20% increase in monthly active users** (paper accepted to **Neurocomputing**).

Research Areas, Skills & Honors

Foundation & Multimodal AI: multimodal foundation models, multimodal representation learning, vision-language alignment, brain-AI alignment, neuroimaging, biomedical AI

LLMs & Agents: LLM post-training, agent evaluation, long-term memory, user simulation, benchmark design

Modeling: Transformers, mixture-of-experts, diffusion models, state-space models, graph neural networks, generative modeling, knowledge distillation, probabilistic modeling, time-series forecasting, causal inference, uplift modeling

Programming & Systems: Python, C++, MATLAB, PyTorch, scikit-learn, CUDA, Ray, AWS, Docker, Kubernetes, Git, Linux

Honors: CVPR Travel Award (2026); Georgia Tech SGA Conference Fund (2025, 2026); RECOMB Student Travel Grant (2025)

Publications

⋆ Selected publication; † equal contribution.

Y. Wei, Y. Zhang, X. Xiao, C. Qian, T. Wang, V. D. Calhoun.

1. ⋆ **fMRI-LM: Towards a Universal Foundation Model for Language-Aligned fMRI Understanding**

(*CVPR, 2026*).

C. Lane, M. Tripathy, L. K. Murali, R. S. Grandhi, S. S. Z. Yang, S. Gijsen, D. Das, M. Ram, U. K. Singh, C. K. T. Villanueva, **Y. Wei**, W. Beddow, G. Cortés, S. Cho, D. Z. Kaplan, B. Warner, T. M. Abraham, P. S. Scotti.

2. **Scaling Vision Transformers for Functional MRI with Flat Maps**

(*ICML, 2026*).

Y. Wei, Y. Zhang, X. Xiao, T. Wang, X. Wang, V. D. Calhoun.

3. ⋆ **MoRE-Brain: Routed Mixture of Experts for Interpretable and Generalizable Cross-Subject fMRI Visual Decoding**

(*NeurIPS, 2025*).

S. Hu†, **Y. Wei**†, J. Ran†, Z. Yao, X. Han, H. Wang, R. Chen, L. Zou.

4. ⋆ **Does Memory Need Graphs? A Unified Framework and Empirical Analysis for Long-Term Dialog Memory**

(*ACL (Long Papers), 2026*).

S. Hu†, Z. Zhang†, **Y. Wei**†, X. Han, Z. Tang, H. Wang, R. Chen.

5. ⋆ **CloneMem: Benchmarking Long-Term Memory for AI Clones**

(*ACL (Long Papers), 2026*).

Y. Wei, Y. Zhang, X. Xiao, T. Wang, X. Wang, V. D. Calhoun.

6. **4D Multimodal Co-attention Fusion Network with Latent Contrastive Alignment for Alzheimer's Diagnosis**

(*WACV, 2026*).

X. Xiao†, Y. Zhang†, J. Wang, L. Zhao, **Y. Wei**, H. Li, Y. Li, X. Wang, S. K. Roy, H. Xu, T. Wang.

7. **RoadBench: A Vision-Language Foundation Model and Benchmark for Road Damage Understanding**

(*WACV, 2026*).

X. Xiao†, Y. Zhang†, X. Li, T. Wang, X. Wang, **Y. Wei**, J. Hamm, M. Xu.

8. **Visual Instance-aware Prompt Tuning**

(*ACM Multimedia, 2025*).

Y. Wei, A. Abrol, D. Qiu, J. J. Lah, A. I. Levey, V. D. Calhoun.

9. **Spatiotemporal State Space Modeling of Dynamic Brain Connectivity in Cognitively Normal Individuals at Risk for Alzheimer's Disease**

(*ICASSP (Oral), 2026*).

Y. Wei, A. Abrol, J. J. Lah, A. I. Levey, V. D. Calhoun.

10. **From Symptomatic to Pre-Symptomatic: Adaptive Knowledge Distillation for Early Alzheimer's Detection Using Functional MRI**

(*IEEE Transactions on Biomedical Engineering 73(3), 2026*).

Y. Wei, A. Abrol, V. D. Calhoun.

11. **Hierarchical Spatio-Temporal State-Space Modeling for fMRI Analysis**

(*RECOMB, 2025*).

Y. Wei, Z. Qiu, Y. Li, Y. Sun, X. Li.

12. ⋆ **Multi-Treatment Multi-Task Uplift Modeling for Enhancing User Growth**

(*Neurocomputing, accepted*).

Y. Wei, H. Wang.

13. **Wavelet Integrated Attention Network with Multi-resolution Frequency Learning for Mixed-Type Wafer Defect Recognition**

(*Engineering Applications of Artificial Intelligence 121, 2023*).

Y. Wei, Y. Chen, T. Xue, L. Zekelman, N. Makris, Y. Rathi, W. Cai, F. Zhang, L. J. O'Donnell.

14. **A Deep Network for Explainable Prediction of Non-imaging Phenotypes Using Anatomical Multi-view Data**

(*CDMRI at MICCAI, 2023*).

Y. Wei, H. Wang.

15. **Mixed-Type Wafer Defect Recognition With Multi-Scale Information Fusion Transformer**

(*IEEE Transactions on Semiconductor Manufacturing 35(2), 2022*).

Y. Wei, H. Wang.

16. **Mixed-Type Wafer Defect Pattern Recognition Framework Based on Multifaceted Dynamic Convolution**

(*IEEE Transactions on Instrumentation and Measurement 71, 2022*).