

CHAO CAO

✉ caochao0306@gmail.com  linkedin.com/in/chao-cao-167037246  Google Scholar

SUMMARY

Ph.D. candidate in Computer Science and Engineering at the University of Texas at Arlington, graduating in December 2026. My first-author research applies deep learning to **brain MRI**, from cortical folding patterns on surface meshes to cross-subject landmark matching that combines **T1-weighted and diffusion MRI**. As a co-author, I contributed to continuous **staging of Alzheimer's disease and Lewy body dementia** and to a **multimodal LLM** that pairs 3D brain MRI with clinical text. Outside imaging, I worked on CPT → SFT → RL post-training of a **32B vision-language foundation model** with **PyTorch** and **Megatron-LM**. On the language side, I built a radiology NLI benchmark for LLMs and a multi-agent clinical decision support system. I write research code with AI coding assistants (Claude Code, Cursor). First or co-first author on three IEEE ISBI papers (**two 2026 orals**) and an IEEE Transactions on Big Data article; co-author at NeurIPS, MICCAI and AAAI.

EDUCATION

University of Texas at Arlington <i>Ph.D. in Computer Science and Engineering</i>	Jan 2023 – Dec 2026 (Expected)
Research focus: deep learning for medical imaging (brain MRI) and neurodegenerative disease	
University of Maryland at College Park <i>M.S. in Data Science and Analytics</i>	Jan 2021 – Dec 2022
Nanjing University of Posts and Telecommunications <i>B.S. in Computer Science and Technology</i>	Sep 2017 – Jul 2021

RESEARCH EXPERIENCE

University of Texas at Arlington Medical Imaging and Brain Networks	Jan 2023 – Present
<i>Graduate Research Assistant (Advisor: Dr. Dajiang Zhu)</i>	<i>Arlington, TX</i>
- Developed a cross-subject matching method for 3-hinge gyri that combines T1-weighted MRI and diffusion MRI (DTI) from 1,065 Human Connectome Project subjects: anatomical embeddings propose candidate matches, then DTI-derived structural connectivity optimization resolves one-to-many matches into one-to-one correspondences; first-author IEEE ISBI 2024. - Co-developed Gyral-Sulcal-Net (one of three co-first authors), a network representation that unifies gyral and sulcal folding patterns, evaluated on more than 1,600 brain scans spanning 34 gestational weeks to older adults, including pathological cohorts, with collaborators at Indiana University and Massachusetts General Hospital; IEEE ISBI 2026 Oral. - Co-authored neurodegenerative disease modeling papers: continuous staging of Alzheimer's disease and Lewy body dementia that projects each subject onto continuous tree structures built from hierarchical anatomical and structural connectivity features (MICCAI 2025), and a spatio-temporal Transformer on dynamic functional connectivity from resting-state fMRI for mild cognitive impairment (MCI) classification (ISBI 2025). - Co-authored Brain-Adapter (ISBI 2025), an adapter-tuned multimodal LLM that aligns 3D brain MRI with clinical text through CLIP-style contrastive training, and DCMM-Transformer (AAAI 2026), a degree-corrected mixed-membership attention mechanism for medical imaging.	
University of North Carolina at Chapel Hill BRAIN Lab	May 2024 – Aug 2024
<i>Research Assistant, Department of Radiology and BRIC (Advisor: Dr. Gang Li)</i>	<i>Chapel Hill, NC</i>
- Led Deep-GyralNet, a deep learning replacement for the conventional GyralNet and 3-hinge gyrus extraction pipeline: a Spherical U-Net in PyTorch on 40,962-vertex spherical cortical surface meshes, trained and cross-validated on 1,098 Human Connectome Project subjects; co-first-author IEEE ISBI 2026 Oral. - Paired focal loss for class imbalance with a Connectivity-Aware Path Enforcement (CAPE) loss for topological continuity, followed by geodesic refinement; reached an F1 of 96.4 (recall 98.2) and cut per-subject runtime from 237s to 14s (17×).	

RESEARCH PROJECTS

Large Vision-Language Model Post-Training for Multimodal Math Reasoning	Jan 2026 – May 2026
<i>Qwen2.5-VL-32B PyTorch, Megatron-LM, verl, vLLM</i>	<i>Arlington, TX</i>
- Post-trained Qwen2.5-VL-32B, a 32B-parameter vision-language foundation model, through a CPT → SFT → RL pipeline: continued pretraining with Megatron-LM on 20M image-text pairs, then full-parameter SFT on 770K reasoning samples. - Ran the GSPO-based RL stage in verl, with vLLM accelerating inference, on 80K curated samples with rule-based rewards.	
Radiology NLI Benchmarking with Large Language Models	Nov 2023 – Jun 2024
<i>Clinical NLP, LLM Evaluation</i>	<i>Arlington, TX</i>
Built RadQNLI, a 9,838-pair radiology natural language inference benchmark converted from RadQA, and evaluated ChatGPT/GPT-4 with zero-shot, few-shot and chain-of-thought prompting against open-source LLMs and fine-tuned BERT-family models; co-first-author IEEE Transactions on Big Data article (2025).	
Multi-Agent Clinical Decision Support System	Feb 2026 – May 2026
<i>LangGraph, GraphRAG, FastAPI, PostgreSQL, Docker</i>	<i>Arlington, TX</i>
Built a LangGraph system of 5 agents (StateGraph orchestration, conditional routing, shared state, traceable execution) that combines GraphRAG retrieval, HIPAA-oriented PHI de-identification, ICD-10 coding and drug-drug interaction checks, served through FastAPI in Docker.	

INDUSTRY EXPERIENCE

GyriQAI, Inc.

May 2025 – Aug 2026

Full-Stack Software Development Engineer Intern | Python, asyncio, Playwright, AWS Lambda

Athens, GA

- Engineered 20+ fault-tolerant crawling and **data pipelines** (asyncio + Playwright, checkpoint/resume, rotating proxy pools) that produced a national dataset of 40K+ schools and **1M+ structured Q&A samples**.
- Built multi-agent LLM orchestrators across 5 platforms on a dual GPT-4o + Gemini client with **160+ pytest tests**; red-teamed the production chatbot with a 232-cell adversarial matrix and confirmed 3 vulnerabilities.

SELECTED PUBLICATIONS

1,000+ citations, h-index 11 | *equal contribution | [Google Scholar](#)

- Cao, C.***, Cheng, J.*, Chen, M., Zhang, J., Liu, T., Zhang, L., Zhu, D., Li, G. (2026). Deep-GyralNet: Enabling machine learning in gyral folding pattern extraction on cortical surface. *IEEE ISBI 2026* (**Oral**).
- Cao, C.***, Chen, T.*, Zhao, N.*, Chen, M., Qu, M., Zhang, Z., Shi, X., Li, X., Liu, T., Zhang, L. (2026). Gyral-Sulcal-Net: An integrated network representation of brain folding patterns. *IEEE ISBI 2026* (**Oral**).
- Cao, C.**, Yu, X., Zhang, L., Chen, T., Lyu, Y., Liu, T., Zhu, D. (2024). Enhancing group-wise consistency in 3-hinge gyrus matching via anatomical embedding and structural connectivity optimization. *IEEE ISBI 2024*.
- Wu, Z.*, Zhang, L.*, **Cao, C.***, Yu, X., et al. (2025). Exploring the trade-offs: Unified large language models vs local fine-tuned models for highly-specific radiology NLI task. *IEEE Transactions on Big Data*, 11(3).
- Chen, T., Chen, M., Zhang, J., Zhuang, Y., **Cao, C.**, et al. (2025). A unified continuous staging framework for Alzheimer's disease and Lewy body dementia via hierarchical anatomical features. *MICCAI 2025*.
- Zhang, J., Yu, X., Lyu, Y., Zhang, L., Chen, T., **Cao, C.**, et al. (2025). Brain-Adapter: Enhancing neurological disorder analysis with adapter-tuning multimodal large language models. *IEEE ISBI 2025*.
- Zhang, J., Lyu, Y., Yu, X., Zhang, L., **Cao, C.**, et al. (2025). Classification of mild cognitive impairment based on dynamic functional connectivity using spatio-temporal transformer. *IEEE ISBI 2025*.
- Yu, X., Zhang, J., Chen, T., Zhuang, Y., Chen, M., **Cao, C.**, et al. (2025). Domain-adaptive diagnosis of Lewy body disease with transferability aware transformer. *MICCAI 2025*.
- Cheng, H., Yu, X., Wu, S., Fang, L., **Cao, C.**, et al. (2026). DCMM-Transformer: Degree-corrected mixed-membership attention for medical imaging. *AAAI 2026*.
- Yu, X., Zhang, L., **Cao, C.**, Chen, T., Lyu, Y., et al. (2024). Gyri vs. sulci: Core-periphery organization in functional brain networks. *MICCAI 2024*.
- Li, X., Zhao, L., Zhang, L., Wu, Z., Liu, Z., Jiang, H., **Cao, C.**, et al. (2025). Artificial general intelligence for medical imaging analysis. *IEEE Reviews in Biomedical Engineering*, 18.
- Li, Y., Wu, Z., Zhao, H., Zhou, Y., **Cao, C.**, et al. (2026). A world model of radiologist reading for medical image representation learning. *NeurIPS 2026*.
- Liu, Z., Huang, Y., Yu, X., Zhang, L., Wu, Z., **Cao, C.**, et al. (2023). DeID-GPT: Zero-shot medical text de-identification by GPT-4. *arXiv:2303.11032*.

TECHNICAL SKILLS

ML & LLM: PyTorch, Megatron-LM, verl, vLLM, LangGraph, GraphRAG, OpenAI and Gemini APIs

Medical Imaging: T1-weighted and diffusion MRI (DTI), cortical surface meshes, brain network analysis, Human Connectome Project data, radiology report NLP

Languages: Python, MATLAB, SQL (PostgreSQL, SQLite), JavaScript

AI Coding Assistants: Claude Code, Cursor (daily use for research code)

Engineering: Linux, Git, Docker, FastAPI, Flask, pytest, AWS Lambda

TEACHING & SERVICE

Teaching Assistant, UT Arlington: CSE 3313 Introduction to Signal Processing (Spring 2024 – Fall 2025); CSE 1320 Intermediate Programming (Summer 2025); CSE 5325 Software Engineering II (Fall 2023)

Reviewer: MICCAI 2024; Pattern Recognition; IEEE Transactions on Neural Networks and Learning Systems; Engineering Applications of Artificial Intelligence