

Zihan Liang

✉ zihan.liang@emory.edu ☎ (404)310-6818 🌐 Personal Website in zihan-liang 🌐 zihanliang

Education

Emory University *Honors BS in Applied Mathematics and Statistics; AI Minor* Aug 2023 – May 2026

- **GPA:** 3.98/4.00
- **Awards and Honors:** Phi Beta Kappa; Dean’s List; ASA DataFest Best Visualization Award; AIMO Bronze Medal
- **Coursework:**
 - Mathematics: Multivariable Calculus, Linear Algebra, Discrete Mathematics, Ordinary Differential Equations, Graph Theory (Graduate Level), [Real Analysis \(UIUC NetMath\)](#)
 - Statistics & Computation: Mathematical Statistics, Numerical Analysis, Data Science Computing, Machine Learning, Text Data Science, Experimental Methods, Causal Inference (Doctoral Level)
 - Biostatistics & Bioinformatics: [Systems Biology and Biotechnology Specialization](#)

Publications

Learning Dynamic Representations and Policies from Multimodal Clinical Time-Series with Informative Missingness [ACL 2026 Findings](#)
Mar. 2025 – Jan. 2026

Zihan Liang*, Ziwen Pan*, *Ruoxuan Xiong*

- Proposed an MNAR-aware multimodal framework for clinical time-series that treats observation patterns as signal rather than noise, combining explicit missingness features, adaptive fusion of sparse clinical text, and action-conditioned latent dynamics to learn patient states for both prediction and offline decision-making
- Demonstrated strong results on MIMIC-IV and eICU: mortality prediction reached AUROC 0.876 on MIMIC-IV (+3.9% over GRU-D), while the learned treatment policy improved FQE by 20.3% over clinician behavior and delivered the largest gains for high-severity patients

DART: Mitigating Harm Drift in Difference-Aware LLMs via Distill-Audit-Repair Training [ACL 2026 Findings](#)
Mar. 2025 – Jan. 2026

Ziwen Pan*, Zihan Liang*, *Jad Kabbara, Ali Emami*

- Identified the problem of harm drift in safety-tuned LLMs—where fine-tuning improves decision accuracy but makes rationales more harmful—and introduced DART, a Distill-Audit-Repair pipeline for improving difference-awareness classification while preserving safer explanations
- On eight benchmarks, DART improved Llama-3-8B-Instruct accuracy from 39.0% to 68.8%, boosted equal-treatment accuracy from 11.3% to 72.6%, and reduced harm drift cases by 72.6%, showing that accuracy and safety can be improved together

MambaDATG: Domain-Adaptive Tri-Plane-Gated Pre-training for 3D Abdominal Segmentation [ICASSP 2026 Oral](#)
Mar. 2025 – Jan. 2026

Y. Shen*, D. Jiang*, *Zihan Liang**, Y. Wei*, W. Yang, S. Wang, et al.

- Proposed MambaDATG, a tri-plane-gated selective SSM pre-training that fuses axial, coronal, and sagittal scans via voxel gating—capturing directional anisotropy with linear-time efficiency and zero inference overhead
- Introduced a domain-adaptive MIM stage adapting high-level layers on unlabeled CTs, improving small-organ segmentation (e.g., +1.8 Dice on BTCV) and revealing the tri-plane gate as the key performance driver

Causal Representation Learning from Multimodal Clinical Records under Non-Random Modality Missingness [EMNLP 2025 Main](#)
Feb. 2025 – May 2025

*Zihan Liang**, Ziwen Pan*, *Ruoxuan Xiong*

- Proposed CRL-MMNAR, a causal multimodal framework that treats modality missing-not-at-random as signal—combining missingness-aware fusion (gated by observation patterns) with cross-modal reconstruction + contrastive learning, and a multitask predictor with a cross-fitted rectifier to correct observation-pattern bias
- Demonstrated consistent gains on MIMIC-IV and eICU, including AUC 0.8687→0.9693 for ICU admission (+11.6%) and 0.7989→0.8657 for readmission (+8.4%) on MIMIC-IV, and readmission AUC improving to 0.9294 (+13.8%) on eICU, with lower Brier scores indicating better calibration

CareLab at #SMM4H-HeaRD 2025: Insomnia Detection and Food Safety Event Extraction with Domain-Aware Transformers

[AAAI ICWSM 2025](#)
Feb. 2025 – Apr. 2025

Zihan Liang*, Ziwen Pan*, Sumon Kanti Dey, Azra Ismail

- Ranked 1st on SMM4H-HeaRD 2025 Task 5, developing a RoBERTa + GPT-4-augmented system with class-weighted training and ensembling, achieving $F1 = 0.958$ and revealing the limits of rule-based span extraction in clinical and food-safety texts

Dynamic Policy Design for Autonomous Taxi Adoption: A Hierarchical Game-Theoretic Framework

[TRB Annual Meeting](#)
Jun 2024 – Dec. 2025

Zihan Liang, Ziwen Pan

- Formulated a three-tier Stackelberg game for government–platform–worker interactions in autonomous taxi adoption, deriving equilibrium policies balancing innovation and social welfare

Manuscripts

Adaptive Spatiotemporal Graph Neural Networks with Trend-Aware Prediction and Validation-Gated Calibration for Traffic Flow Forecasting

[IEEE T-ITS](#) Under Review
May 2025 – Present

Zihan Liang*, Ziwen Pan*, Shuyang Yu

- Proposed an adaptive spatiotemporal GNN that fuses distance- and correlation-based graphs with dual temporal–spatial encoders and a trend-aware head, achieving state-of-the-art accuracy on PeMS benchmarks
- Designed a validation-gated calibration pipeline with hour-conditional quantile mapping and AR(1) residual correction, applied only when improving MAPE/MAE to ensure reliable and interpretable post-processing

Collaborative Research Projects

Research Assistant and Software Developer

Atlanta, GA

[Collective Action & Research for Equity \(CARE\) Lab](#), Emory University

Sept 2024 – Present

- Analyzed 1,000 user data points and built an interactive dashboard for MakerGhat using JavaScript, Python, and HTML, improving accessibility and visualization of organizational data
- Developed an offline-first Android app for classroom session recording and metadata management, integrating Google Sheets validation, offline caching, and secure Drive synchronization
- Maintained a WhatsApp EduBot (Python, Twilio) supporting 15,000+ teachers and 600,000+ students

Artificial Intelligence Research Intern

Remote

Google China, Inc.

Jun 2025 – Aug 2025

- Optimized traffic sign detection models (YOLOv8, YOLOX) on TT100K through hparam tuning, data augmentation, and knowledge distillation frameworks, improving mAP by +4% overall and boosting small-object detection by > 15%, retaining real-time efficiency
- Executed cross-domain transfer learning (TT100K → GTSRB, LISA, CTSD, BelgiumTSC), achieving up to +44% mAP in few-shot settings and validating robust model generalization across diverse datasets

Research Assistant - ODE Modeling for C. difficile Transmission

Remote

Polymath Jr REU Summer Program

Jun 2025 – Aug 2025

- Developed and analyzed ODE-based models of C. difficile community transmission and simulated intervention strategies in MATLAB/R to evaluate outbreak control, with results submitted to the [JMM 2026](#)

Leadership Experience

Founder and President | [Emory Artificial Intelligence and Data Association](#)

Dec 2024 – Present

- Founded and led EAIDA, growing it into a 200+ member AI community and launching the “AI for ALL” initiative to promote AI literacy through workshops, discussions, and mentorship

Additional Information

Technologies: Python (Proficient; PyTorch, TensorFlow, NumPy, Pandas, Scikit-learn), R (Proficient; statistical modeling, data visualization), \LaTeX (Proficient), MATLAB (Intermediate; numerical analysis, signal processing), Java (Intermediate), JavaScript & HTML (Intermediate; data visualization & prototyping)