Zhen Lu

Quincy, MA | cnzhenlu@gmail.com | LinkedIn | GitHub | Google Scholar

Computer Science Ph.D. student with a first-author IEEE BigData paper and hands-on experience with LLM tool-calling agents, Transformers, GNNs, and time-series pretraining. Seeking an applied AI/ML internship.

EDUCATION

University of Massachusetts Boston | Ph.D. in Computer Science | 2024–2029 (expected)

Grand Valley State University | M.S. in Applied Computer Science | April 2024 | GPA: 3.9/4.0

Grand Valley State University | B.S.E. in Computer Engineering | August 2019

PUBLICATIONS

1. **Zhen Lu**, Chengjie Zheng, Iris Shen, John Quackenbush, Viola Fanfani, Wei Ding, and Ping Chen. "FACT: Gated Fusion-Augmented Causal Mask Transformer for Pseudotime Analysis." *IEEE International Conference on Big Data (BigData)*, pp. 1113–1122, 2025. DOI.
2. Yong Zhuang, Tianyu Kang, **Zhen Lu**, Wei Ding, and Ping Chen. "Improving Generalization in Deep Neural Networks by Mitigating Memorization." *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, pp. 93–104, 2025. DOI.

SELECTED AI PROJECT

Roameo — Multi-Agent LLM Travel Planner | Agentic AI Hackathon, Team Project | Oct 2026 | GitHub | Live demo

- Built a Dockerized FastAPI/React planner with 7 agents (intake, research, composer, critic, feedback, guide, orchestrator); the LLM handles language and tool selection, while budget math stays deterministic.
- Implemented a bounded tool-calling loop (≤8 steps) over flight, hotel, event, attraction, and weather tools, a critic-driven repair loop (≤5 repairs), and human approval checkpoints before calendar export.
- Built a scenario-based evaluation harness covering end-to-end, human-in-the-loop, and failure cases; passed 23/23 core scenarios in LLM mode (Claude Haiku via OpenRouter, 3 runs each).

RESEARCH EXPERIENCE

Doctoral Research | University of Massachusetts Boston | 2024–Present

- Implemented FACT, a Set Transformer with causal masking and gated fusion of language-model embeddings that infers latent orderings (pseudotime) from tabular data, and ran experiments for the first-author BigData 2025 paper.
- Contributed to experimental evaluation of neural network memorization and generalization for the PAKDD 2025 paper.
- **Remaining Useful Life Prediction (first author):** Designed a self-supervised pretraining framework for predictive maintenance on NASA C-MAPSS that needs no failure labels during pretraining, combining a Bradley–Terry–Luce soft ranking loss over within-engine temporal order with condition-aware prototype alignment; benchmarked against 13 baselines (8 supervised, 3 self-supervised, 2 foundation-model) on all 4 subsets with 3 seeds.
- **MIRAGE:** Ran experiments for a time-series pretraining method robust to unreliable timestamps, on clinical, industrial-sensor, and imaging data.
- **Node2Net:** Ran extensive experiments (100 runs per setting) adding node-specific MLPs to GCN, GAT, GraphSAGE, and GraphGPS on benchmarks including Cora and PubMed.

Master's Thesis | Grand Valley State University | 2024

- Implemented and compared 3D CNN, LSTM, CNN–RNN, and Transformer models for crime hotspot forecasting on Portland Police Bureau call-for-service data.

TECHNICAL SKILLS

Programming: Python, Java, C | **ML & Data:** PyTorch, TensorFlow, scikit-learn, NumPy, pandas, Matplotlib

LLM & Agents: LLM tool calling, multi-agent orchestration, structured output, LLM evaluation, Claude via OpenRouter

Modeling: Transformers, GNNs, self-supervised learning, time-series forecasting

Engineering: FastAPI, React, REST APIs, pytest, Git, Docker

ADDITIONAL EXPERIENCE

Field Test Engineer | LG Electronics | Bolingbrook, IL | Nov 2019–Nov 2020 | Mobile-device field testing, diagnostic log analysis, and documentation of findings and corrective actions.

Application Engineer Intern | Kaidi Electrical | Holland, MI | Jan–Dec 2018 | Supported actuator R&D, PLC testing, performance benchmarking, and software debugging.