Leo Xia

Mathematical and Algorithmic Foundations · Model Replication · Research Training

Email: coderxl@stu.xjtu.edu.cn Blog: blog.leosrealms.top

Education

Xi'an Jiaotong University · School of Artificial Intelligence

Major: Artificial Intelligence · 2025-2029 (expected graduation)

Research Experience

Soccer Agent

Participating in research training in Luo's research group, working on the Soccer Agent project in video understanding.

- Preparing a first-author manuscript on this project.

Projects

MiniMind2 Reproduction

Reproduced MiniMind2 and documented the process in a detailed report.

MiniMind2-Small: An End-to-End Reproduction Report (in Chinese)

DocRes Reproduction

Reproduced DocRes and conducted further experiments based on it.

DocRes

Team Experience

RoboMaster Vision Team

Used OpenCV, ROS2, and Ceres; contributed to armor plate detection, pose estimation, camera calibration, and aiming algorithms.

Mathematical Foundations

Linear Algebra: 100

Mathematical Analysis for Engineering I / II: 92 / 98

Probability Theory: 96

Discrete Mathematics: 96

Technical Skills

Programming: Proficient in Python and C++

Deep Learning and Scientific Computing: Proficient in PyTorch, NumPy, Pandas, and Matplotlib

Computer Vision and Optimization: Familiar with OpenCV, ROS2, and Ceres

Competition Awards

2023 NOI Spring Test

Provincial First Prize

Lanqiao Cup, Spring 2026

Provincial Second Prize

Languages

College English Test Band 6 (CET-6)

CET-6: 629; Spoken English: Good

Personal Interests and Research Approach

I have long been interested in computing and artificial intelligence, particularly autonomous driving and robotics. In my second year of high school, I experimented with basic reinforcement learning to train a cart to balance in a simulator. The work remained entirely in simulation, but it was a way to put my extracurricular interests into practice. My interest in embodied and spatial intelligence continues that same curiosity: how can a system learn to understand its environment, make decisions, and act?

The areas I most want to explore are spatial intelligence and machine intelligence, especially perception, learning, and action in the physical world. My interests in model architectures and audio perception are closely connected to these goals. I want to understand what representations and learning mechanisms would allow an intelligent system to use the rich, continuous information in its environment and gradually develop the capacity for continual learning and planning.

I am particularly interested in the computational efficiency of bringing these capabilities to edge devices. For applications such as autonomous driving and robotics, I care about both efficient inference under limited compute and resources, and adaptation to new environments and experiences. Efficient on-device training, fine-tuning, and model updates are therefore questions I hope to explore. Energy Transformer, attractor networks, and Forward-Forward are among the directions I currently follow. Through reading, implementation, and experiments, I hope to understand how these architectures and learning methods work and test whether they offer useful ideas for efficient intelligence on edge devices.

In audio, I am interested in perception and understanding beyond speech transcription. Intelligible speech is only one part of sound; music, background sounds, noise, and sounds from the natural environment also carry rich information. I want to explore how models can retain and understand more of that information, and how it can contribute to spatial understanding, interpreting the environment, and autonomous action. This connects to a more intuitive aspiration: enabling models to perceive the structure and changes in sound itself, beyond extracting content that can be transcribed into text.

My preferred way into a research problem is through code and experiments: implementing methods, reproducing results, and testing my understanding against concrete questions. I enjoy turning an idea into something I can run, observe, and analyze. When a problem calls for mathematical foundations, I am also willing to stop and work through the underlying principles carefully. Implementation and theory should help me understand the problem more precisely.

I hope to pursue patient, sustained exploration of model architectures and foundational learning methods: understanding why a method works, where its limitations lie, and where substantial improvements may be possible. I value research whose insights can build over time, so that a deeper understanding of a problem and a methodological contribution remain useful as applications and foundation models change. I want to develop my own judgment, formulate concrete ideas, and test and advance them through careful implementation and experiments.