Xiang Li

Nanjing, 210023, China lllllimited306@gmail.com +86-138-6772-3723

Education

Nanjing University *Sept 2022 – June 2026 (expected)*
Bachelor of Science in Kuang Yaming Honors School

- **Major**: Computer Science and Artificial Intelligence.
- GPA: 4.61/5.00 (average: 92.2/100); Rank: 1/109.
- **Selected Coursework:** Machine Learning (95/100), Data Structures and Algorithms (93/100), Numerical Analysis (98/100), Linear Algebra (98/100), Computer Systems (96/100), Programming (94/100), Real Analysis and Functional Analysis (93/100), Data Mining (94/100).

University of California, Berkeley *Aug 2024 – Dec 2024*
Berkeley International Student Program (BISP)

- GPA: 4.00/4.00.
- **Relevant Coursework:** Theoretical Statistics (A), Efficient Algorithms and Intractable Problems (A), Linear Modeling: Theory and Applications (A).

Research Interests

Machine Learning, Reinforcement Learning, Optimization, Online Learning

Publications

Softmax for Continuous Actions: Optimality, MCMC Sampling, and Actor-Free Control *Sept 2025*

Pai Liu, ***Xiang Li***, Nan Jiang

ICLR 2026 Conference Submission

Research Experience

Offline Reinforcement Learning Theory with Continuous Action Space and Standalone Policy Class *Champaign, IL, June 2025 – Now*
University of Illinois, Urbana-Champaign

Summer intern supervised by Associate Prof. Nan Jiang.

- Applied pessimism in policy gradient methods, extending theoretical guarantees of existing offline RL works to continuous action spaces and standalone policy classes. And incorporated online learning and optimization techniques into offline RL frameworks to interpret and improve optimization dynamics.
- Jointly proposed Deep Decoupled Softmax Q-learning (DDSQ): a critic-only training framework that directly aligns with energy-based action distributions via Langevin dynamics.
- Gave the convergence analysis under continuous action space of the softmax Q-iteration, and non-asymptotic bound of self-normalized importance sampling (SNIS).

Online Reinforcement Learning and Bandits with Multinomial Logit Function Approximation *Nanjing, China, Feb 2025 – Now*
Nanjing University

Research project supervised by Assistant Prof. Peng Zhao.

- Analyzed sample complexity and regret bounds for MDPs and Markov games under multinomial logit (MNL) function approximation via novel usage of online mirror descent (OMD). And extended theoretical frameworks to multi-player general-sum Markov games and reward-free exploration settings.
- Established regret bounds of $\sqrt{T/\kappa^*}$ for contextual MNL bandits, leveraging structural properties of non-linear feedback models.
- Investigated decentralized Markov games with function approximation.

Neuromorphic Architecture Design, Efficient Spiking Neural Networks and Sparse Mixture-of-Experts *Nanjing, China & Santa Cruz, CA*
Nanjing University & University of California, Santa Cruz *Sept 2023 – Nov 2024*

Research project supervised by Associate Prof. Jie Guo and Assistant Prof. Jason Eshraghian.

- Developed real-time textural recognition systems using Spiking Neural Networks (SNNs) with optimized sparse connectivity, achieving efficient inference.
- Replaced the top-k mechanism in Mixture-of-Experts (MoE) with biologically-inspired sparsity techniques (e.g., LIF neurons), and integrated advanced recurrent architectures such as HGRN2.
- Innovated CNN-to-SNN conversion methods and improved sparsity implementation, reducing inference time and maintaining competitive accuracy on the DTD dataset.
- Explored large-scale configurations ("Mixture of a Million Experts") to mitigate overfitting and enhance model scalability.

Selected Awards

- Mathematical Contest in Modeling (MCM/ICM) - Honorable Mention, 2025
- Special Scholarship for Basic Sciences - First Prize of Year 2024 (top 5% in top-notch personnel)
- Merit Student in Jiangsu Province of Year 2024
- Special Scholarship for Basic Sciences - First Prize of Year 2023 (top 5% in top-notch personnel)
- People's Scholarship - First Prize of Year 2023 (top 5%)

Miscellaneous

Language: Mandarin (native), English (fluent)

Programming: Python, C, C++, Java, Latex, Matlab, Markdown

TOEFL score: 104/120 (listening 25, speaking 24, reading 29, writing 26)

GRE score: 323/340 (verbal 153, quantitative 170, writing 4/6)

Personal Interests: Swimming, Singing, Video Games, Basketball.

Channel: I also maintain Bilibili Channel, a Chinese video-sharing platform similar to YouTube. I have uploaded more than 100 videos covering tutorials on machine learning, deep learning, and glimpses of my personal life. The channel has garnered over 8,000 subscribers to date and over 200,000 times of watches.