Peter Chen

lc3826@columbia.edu
peterlaulukchen.github.io

Research Statement

I pursue empirical and application-driven work supported by theoretical foundations. My research lies at the intersection of optimization, machine learning, and geometric data analysis, with applications to LLMs and vision tasks. My areas of expertise include:

1. **LLM Post Training**: LLM Preference Optimization/Alignment (RLHF/DPO); LLM Reasoning via Reinforcement Learning (GRPO/PPO); Agent Training.
2. **Optimization Theory**: Optimal Transport; Non-convex Optimization; Minimax Optimization & Games; Zeroth-order Optimization.
3. **Geometric Data Analysis**: Geometry Processing; 3D Segmentation/Reconstruction.

Education

Columbia College, Columbia University — New York, US
B.A in Mathematics, Computer Science — May 2026
Advisor: Prof. Andrew J. Blumberg (Math), Prof. Tianyi Lin (IEOR)

Experience

Research Intern | AI Lab, Princeton University
Hosted by Prof. Mengdi Wang — Feb 2025 – Ongoing
Topic: LLM RL Reasoning; LLM Agent Training

Research Intern | HKU Musketeers Foundation Institute of Data Science
Hosted by Prof. Yue Xie, Prof. Qingpeng Zhang — Apr 2024 – Jan 2025
Topic: Neural Optimal Transport, Convex Networks

Teaching Assistant | Department of Mathematics, Columbia University
TA for MATH 2500 Analysis & Optimization over SP24, FA24, SP25, FA25, SP26

Publications

ComPO: Preference Alignment via Comparison Oracles
***Peter Chen**, Xi Chen, Wotao Yin, Tianyi Lin*
Advances in Neural Information Processing Systems 38 (NeurIPS 2025)

Spectral Policy Optimization: Coloring your Incorrect Reasoning in GRPO
***Peter Chen**, Xiaopeng Li, Ziniu Li, Xi Chen, Tianyi Lin*
arxiv-2505.11595

Exploration v.s. Exploitation: Rethinking RLVR through Clipping, Entropy, and Spurious Reward
***Peter Chen**, Xiaopeng Li, Ziniu Li, Xi Chen, Wotao Yin, Tianyi Lin*
under review

Displacement-Sparse Neural Optimal Transport
***Peter Chen**, Yue Xie, Qingpeng Zhang*
arxiv-2502.01889

Geometric Framework for Cell Oversegmentation
***Peter Chen**, Bryan Chang, Olivia Creasey, Julie Sneddon, Zev Gartner, Yining Liu*
arxiv-2502.01890

SICNN: Sparsity-induced Input Convex Neural Network for Optimal Transport
***Peter Chen**, Yue Xie, Qingpeng Zhang*
NeurIPS 2024 Optimization for Machine Learning

Talks

2025 INFORMS Annual Meeting, Atlanta — Oct 2025
Invited Speaker; *LLM Post Training: Turning "Trash" Samples into Value*

Services

Reviewers for: *NeurIPS, AAAI, ICLR, ICML, Optimization for Machine Learning*