Licong Lin

📱+1 6282495717 | ✉ liconglin@berkeley.edu | 🏠 licong-lin.github.io | mailing address: 1606 MLK Jr Way, Berkeley, CA

Education

University of California at Berkeley *Berkeley, CA*

PhD student in Statistics *Aug 2021 - Current*

- Advised by Song Mei and Peter Bartlett.
- Worked as a GSI (teaching assistant) for STAT 153, STAT 210B, STAT 135; as a GSR (research assistant).
- **Graduate Courses:** STAT 205B (A), STAT 210A (A+), STAT 210B (A+), STAT 215A (A+), STAT 241 (Reinforcement learning, A+), STAT 256 (Causal inference, A+), EE 227C (Convex optimization, A).

Peking University *Beijing, China*

B.S. in Statistics *Sep 2017 - July 2021*

- Overall GPA: **3.82**/4 (rank: **2**/45)

Research Interests

Deep learning theory, LLM alignment, high-dimensional statistics, statistical inference

Publication and preprints

Transformers as Decision Makers: Provable In-Context Reinforcement Learning via Supervised Pretraining
  Licong Lin, Yu Bai, Song Mei
  International Conference on Learning Representations (ICLR) (2024). 2024

Statistical Limits of Adaptive Linear Models: Low-Dimensional Estimation and Inference
  Licong Lin, Mufang Ying, Suvrojit Ghosh, Koulik Khamaru, Cun-Hui Zhang
  Advances in Neural Information Processing Systems 36 (2024). 2024

Mean-field variational inference with the TAP free energy: Geometric and statistical properties in linear models
  Michael Celentano, Zhou Fan, Licong Lin, Song Mei
  arXiv preprint arXiv:2311.08442 (2023). 2023

Semi-parametric inference based on adaptively collected data
  Licong Lin, Koulik Khamaru, Martin J Wainwright
  arXiv preprint arXiv:2303.02534 (2023). 2023

Plug-in Performative Optimization
  Licong Lin, Tijana Zrnic
  arXiv preprint arXiv:2305.18728 (2023). 2023

Near-optimal multiple testing in Bayesian linear models with finite-sample FDR control
  Taejoo Ahn, Licong Lin, Song Mei
  arXiv preprint arXiv:2211.02778 (2022). 2022

What causes the test error? going beyond bias-variance via anova
  Licong Lin, Edgar Dobriban
  The Journal of Machine Learning Research 22.1 (2021) pp. 6925–7006. JMLRORG, 2021

Selected Research Projects

Transformers as Decision Makers: Provable In-Context Reinforcement Learning *Berkeley, CA*

Advised by Prof. Song Mei (UCB) *May 2023 - Oct 2023*

- Proposed a theoretical framework for in-context reinforcement learning via supervised pretraining using transformers.
- Theoretically showed that transformers have the ability to in-context approximate near-optimal RL algorithms, e.g., LinUCB, Thompson sampling for stochastic linear bandits, and UCB-VI for tabular MDPs.
- Performed preliminary experiments on in-context RL using GPT-2 to verify our theoretical findings.
- The paper is accepted as a poster to ICLR 2024.

Inference of one coordinate in GLM given adaptively collected data via online debiasing. *Berkeley, CA*

Joint work with Koulik Khamaru (former Statistics PhD at UCB) *Dec 2022 - August 2023*

- Follow-up work of the paper on semi-parametric inference (see below). Aim to provide $O(1/\sqrt{n})$-consistent asymptotically normal estimator of a single coordinate of the unknown parameter vector, when the data points are sequentially collected.
- Start with constructing an estimator of a single coordinate that is $O(1/\sqrt{n})$-consistent; then perform the online-debiasing trick to modify the estimator so that it is asymptotically normal. The paper is accepted as a poster to Neurips 2023.

Semi-parametric inference with adaptively collected data

Berkeley, CA

Advised by Prof. Martin Wainwright (UCB & MIT)

Dec 2021 - Feb 2023

- Statistical inference of the target parameter in a generalized linear model at the presence of potentially high-dimensional nuisance parameter.
- Assumed the data points are sequentially collected instead of i.i.d. Obtained asymptotic normal estimators via solving a weighted estimating equation.
- Our proposed algorithm works for offline data collected from bandit algorithms or sequential experiments.

Deep learning theory on overparametrization

Philadelphia, PA

Advised by Prof. Edgar Dobriban (UPenn)

May 2020 - Mar 2021

- Aimed to understand the non-monotonicity (e.g. double descent phenomenon) of the test error as a function of model complexity or number of training samples using asymptotic random matrices theory.
- Published in *Journal of Machine Learning Research, 2021*.

Design a new MCMC algorithm

Beijing, China

Advised by Prof. Cheng Zhang (PKU)

Dec 2019 - Dec 2020

- Combined neural networks with HMC algorithm to design a sampling algorithm suitable for high-dimensional problems.
- Introduced the normalizing flow model into the design of the kinetic energy function of HMC. With the strong representability of the neural networks, our algorithm finds a kinetic energy function that significantly improves the sampling performance of HMC.

Ongoing Projects

Jackknife debiasing of high-dimensional Z-estimators

Berkeley, CA

Joint work with Fangzhou Su, advised by Peng Ding, Martin Wainwright (UCB & MIT)

Dec 2022 - Current

- Analyze the performance of jackknife debiased Z-estimators of one-dimensional functionals. Prove asymptotic normality under the high dimensional regime $p^{3/2} << n$, where p is the problem dimension and n is the number of samples.
- The proposed method is an automatic debiasing procedure. It has wide applications in econometrics, causal inference, generalized linear model, etc.
- Apply concentration results of random matrices and U-statistics in the analysis. Derive non-asymptotic bounds for the estimation error of Z-estimators.

Theoretical understanding of scaling law in LLMs

Berkeley, CA

Joint work with Jingfeng Wu, advised by Peter Bartlett, Jason Lee (UCB & Princeton)

Jan 2024 - Current

- Aim to provide a theoretical interpretation of the scaling laws of LLMs.
- Analyze the simple linear model with SGD training. We have theoretically established that the risk of the model satisfies a power law formula, which is consistent with the empirical observations in LLMs.

LLM unlearning

Berkeley, CA

Joint work with Ruiqi Zhang, advised by Song Mei (UCB)

Jan 2024 - Current

- Aim to design a new algorithm for LLM unlearning; develop a statistical framework for LLM unlearning.
- This is an empirical project on LLM finetuning. We finetune Llama 2 and evaluate algorithms' performance under various metrics.

Honors & Awards

2021	**Huaixin Bachelor**, Peking University	*China*
2021	**Honor graduate of Applied Mathematics and Statistics Program**, Peking University	*China*
2018,19,20	**Academic Excellence Award**, Peking University	*China*
2019,20	**Peking University Scholarship**, Peking University	*China*
2020	**Gold Medal in Probability & Statistics**, S.-T. Yau College Student Mathematics Contest, placed 1st nationally	*China*
2018	**1st Prize**, Beijing College Student Mathematics Competition	*China*
2016	**2nd Prize**, China National Mathematical Olympiad	*China*

Skills

Programming and software   Python (proficient), R (proficient), Matlab, Git, LaTeX.

References available upon request.