Qianyi Chen (陈谦益)

Tsinghua University | cqy22@mails.tsinghua.edu.cn | (+86) 188-7388-2370

EDUCATION

Engineering Physics, Tsinghua University — **Beijing, China**
Mathematics and Physics Major Category — **Aug. 2018 - Apr. 2019**

Industrial Engineering, Tsinghua University — **Beijing, China**
Major in Industrial Engineering (graduated) — **May. 2019 - Jun. 2022**

School of Software, Tsinghua University — **Beijing, China**
Minor in Data Science and Technology (graduated) — **May. 2019 – Jun. 2022**

School of Economics and Management, Tsinghua University — **Beijing, China**
Major in Management Science and Engineering, Ph.D. Candidate — **Aug. 2022 – Present**

- **Supervisor:** Bo Li
- **Research Focus:** Causal inference and conformal prediction in complex systems.
  I aim to develop effective, computationally scalable, and statistically principled methodologies for addressing problems in complex systems, including social networks, bipartite marketplaces, predictive systems, and large language models.
- **Additional Expertise:**
  - Deep learning, reinforcement learning, and statistical learning
  - Post-training of large language models.

PUBLICATIONS & WORKING PAPERS

***Branch 1, Causal Inference under Interference*: Experimental Design, Estimation, and Policy Learning.**

Just Ramp-up: Debiasing Regression-based Estimator for A/B Tests under Network Interference

- **Qianyi Chen**, Bo Li
- *The 40th Annual Conference on Neural Information Processing Systems (NeurIPS 2026)*
- **TL;DR:** We develop a data-centric engineering approach that selectively merges experimental data naturally generated during the ramp-up process (gradually increase the experimental traffic in the industry) to **debias regression-based extrapolation**. We provide both a systematic theoretical explanation and an intuitive illustration, showing that the improvement stems from **increased exposure variation**.
- **Technical details:** Motivated by the common industry practice of ramp-up A/B testing, we perform an exact, non-asymptotic bias and variance analysis of regression-based estimators under general linear network interference. We demonstrate that **substantial bias reduction** can be achieved by training the regression model on the **merged experimental data aggregated across the incremental ramp-up stages**. The generality of these theoretical findings is subsequently validated through a series of novel and challenging simulation studies.

Journey to the Centre of Cluster: Harnessing Interior Nodes for A/B Testing under Network Interference

- **Qianyi Chen**, Anpeng Wu, Bo Li, Lu Deng, Yong Wang
- *The 14th International Conference on Learning Representations (ICLR 2026)*
- **TL;DR:** We construct a **semi-supervision perspective for causal estimation under network interference** and effectively leverages all experimental units for reducing bias and variance in the meanwhile, based on the cluster structure in the network.
- **Technical details:** Based on the wide application of cluster-level randomization for tackling network interference in A/B testing, we propose a Mean-in-Interior (MII) estimator to remove the unnecessary weights imposed on outcomes, given the clustering being weak structure. We then demonstrate that most existing estimators heavily relies on the interior nodes of clusters, while there exists potential distributional discrepancy between interior and boundary of clusters. Hence, we propose an augmented MII estimator, which utilizes a regression function to adjust for the selection bias of cluster interior, i.e., "**debiasing using predictions**".

Optimized Covariance Design for AB Test on Social Network under Interference

- **Qianyi Chen**, Bo Li, Lu Deng, Yong Wang
- *The 37th Annual Conference on Neural Information Processing Systems (NeurIPS 2023)*
- **TL;DR:** We propose a new experimental design paradigm, covariance design. Specifically, we optimize the covariance matrix of the treatment vector to minimize a tractable upper bound on the mean-squared error, while preserving samplability and differentiability.
- **Technical details:** We propose an elaborate characterization of the bias and variance of Horvitz-Thompson estimator under network interference, which first delineate the source of bias with the between-clusters edges and covariance matrix of treatment allocation vector. We then propose an innovative upper bound on the mean-squared error, which organically incorporate network topology instead of high-level statistics in the existing literature. This upper bound enable **gradient-based optimization to induce the experimental design on covariance matrix of treatment vector**. Technically, we further leverage the Grothendieck's identity and Cholesky-based parameterization to guarantees the validity of optimized covariance matrix of treatment allocation.

Practical Performative Policy Learning with Strategic Agents

- **Qianyi Chen**, Ying Chen, Bo Li
- *Presented at INFORMS MSOM 2025, London*
- **TL;DR**: Scalable policy learning when the data distribution echoes the deployed policy.
- **Technical details:** We address the challenge of **performative** (i.e., the data distribution echoes the deployed policy) **policy learning** by developing a scalable, gradient-based policy optimization framework. The core of this method is an evaluation vector, an innovative **mediator** that models agents' decision outcomes directly, thus avoiding fragile parametric assumptions about their strategic response process. Furthermore, we establish a novel theoretical path for proving convergence based on realizability assumptions in RKHS. Our approach relies on weaker technical assumptions and circumvents the **scalability** issues inherent in classic finite difference-based guarantees, and exhibits outstanding performance for various high-dimensional settings.

COSTA: Covariance-Optimized Design and Causal Inference under Network-Temporal Interference

- **Qianyi Chen**, Bo Li
- *Under review at Journal of the American Statistical Association*
- **Technical details:** We study randomized experiments where treatment effects propagate through both **network spillovers and temporal carryover**, and optimize the dependence structure of treatment assignments to reduce estimation error. We derive an exact assignment-cut representation of Horvitz–Thompson bias and combine it with a variance bound to directly optimize treatment covariance. COSTA uses a **scalable thresholded-Gaussian Kronecker parameterization** to jointly design network and temporal dependence, together with **robust inference procedures** that account for both assignment dependence and interference-induced outcome dependence.

Branch 2, Conformal Prediction: Uncertainty-Aware Predictions

Colorful Pinball: Density-Weighted Quantile Regression for Conditional Guarantee of Conformal Prediction

- **Qianyi Chen**, Bo Li
- *The 43rd International Conference on Machine Learning (ICML 2026)*
- **TL;DR:** We directly target the mean-squared error of conditional coverage probabilities and achieve state-of-the-art conditional coverage for conformal prediction in regression task, indicating more uniform and robust coverage guarantee across individual input regions.
- **Technical details:** We refine the quantile regression components that underpin many conformal methods. Leveraging a Taylor expansion, we derive a **sharp surrogate** objective for quantile regression: a density-weighted pinball loss, where the weights are given by the conditional density of the conformity score evaluated at the true quantile. We exploit a **reciprocal structure** and estimate the weights bypassing any density estimation procedure. We then propose a three-headed quantile network that estimates these weights via finite differences using auxiliary quantile levels at $1-\alpha\pm\delta$, subsequently fine-tuning the central quantile by optimizing the weighted loss. We provide a theoretical analysis with exact non-asymptotic guarantees characterizing the resulting excess risk.

How Synthetic Label Improve Conformal Prediction: A Perspective on Conditional Coverage

- **Qianyi Chen**, Bo Li
- Under review
- **Technical details:** We study conformal prediction when trusted labels are scarce but abundant **synthetic labels** can be obtained from domain models or general-purpose LLMs. We introduce **prediction-powered quantile learning**, which combines a large synthetic-labeled pool with paired trusted–synthetic samples to reduce quantile-estimation error while correcting synthetic-label bias, and reserves an independent trusted split for final conformalization to preserve marginal validity. Our theory shows that conformalization removes global threshold shifts and retains only **density-weighted shape errors**, leading to a **three-resource error decomposition** and an explicit covariance–variance rule for when synthetic supervision is beneficial. Experiments across eight regression benchmarks and HelpSteer2 show improved conditional coverage with comparable or smaller prediction sets.

Nonparametric Empirical Bayes for Conformal Inference in LLM Question Answering

- **Qianyi Chen**, Bo Li
- Under review
- **Technical details**: We study open-ended QA where repeated LLM responses form **prompt-specific semantic clusters** and the true answer may remain outside the observed alphabet. We introduce BAYES-STATE, which combines a Pitman–Yor process (PYP) cold-start reference for observed and unseen sampling mass with **cross-prompt truth tilting** that learns how sampling behavior and prompt features map to truth probability, without aligning answer identities across prompts. Split conformal calibration then provides finite-sample marginal validity. Our theory separates sampling-law estimation from truth estimation and shows that posterior accuracy controls excess conditional-coverage error. Across eight dataset–generator–budget settings, **truth tilting drives the main fixed-budget gains**, while PYP is most useful in **scarce-history cold-start regimes** by propagating residual-mass uncertainty.

INTERNSHIP

Post-training algorithm engineer [Ant Star], Ant Group — **Beijing, China**
Worked on *agentic post-training for **data agent*** — **Jun. 2026 – Sep. 2026**

Research intern, Tencent — **Beijing, China**
Worked on *causal inference under network interference* — **Aug. 2022 – Aug. 2023**

Machine learning algorithm engineer, Bytedance — **Beijing, China**
Worked on training *visual and multimodal models (CLIP)* — **Jul. 2021 – Dec. 2021**

SERVICE

- Reviewer at ICML (2025, 2026 (**Gold Reviewer**)), NeurIPS (2024, 2025, 2026), ICLR (2025, 2026, 2027), AISTATS (2024)
- TA at *Prob & Stats* (Bo Li and Xiaojie Mao) Tsinghua University, 2022-2025

AWARDS

Tinker Research Grant (from Thinking Machines Lab)
Greenberg High-Level Academic Employment Potential Scholarship (from Tsinghua University)
The Scholarship for Comprehensive Performance in Tsinghua University (1st level)
The Scholarship for Comprehensive Performance in Tsinghua University (2nd level)