Yichen Xu

Email: yichen_xu@berkeley.edu; Tel: 510-697-0346

EDUCATION

08/2023 – Present **University of California, Berkeley (UC Berkeley)**

Ph.D. Student in Biostatistics, Designated Emphasis in Computational and Genomic Biology

M.A. in Biostatistics, May 2025

09/2019 – 06/2023 **Chinese University of HongKong, ShenZhen (CUHK–SZ)**

Bachelor of Science in Mathematics and Applied Mathematics (Stream: Applied Mathematics)

RESEARCH INTERESTS

Reliable and controllable intelligent systems, with an emphasis on developing principled methods to guide, adapt, and evaluate learning under distribution shift, limited supervision, and data or computational constraints. My research spans causal inference, transfer learning, synthetic data, biomedical AI, and large language model adaptation, with the goal of improving how learning systems estimate effects, transfer knowledge, generate data, and extract interpretable structure from complex datasets.

RESEARCH PROJECTS

05/2025 – Present **AI for single-cell and spatial multi-omics**
Researcher, Advisor: Dr. Zhiliang Bai (MIT, Department of Bioengineering)

- Developed a manifold-aware latent network framework for single-cell and spatial omics, embedding genes/features into latent-factor manifolds and constructing latent similarity networks for hub discovery, module detection, pathway analysis, and spatial program mapping.
- Outperformed correlation-network and aggregated differential-expression baselines for latent program discovery in various simulation; validated the framework across mouse/human brain, human PBMC, CAR-T, lung tumor, spatial colon cancer with liver metastasis, and spatial PDAC datasets.
- Recovered interpretable biological programs, including CAR-T hubs with centrality patterns distinct from differential expression, lung tumor hubs with temporal dynamics and *in silico* perturbation validation, and spatial boundary-specific hub programs for cell visualization and classification in colon cancer and PDAC.
- Developed an unpaired latent-optimization method for counterfactual perturbation prediction, recovering substantial perturbational signal without paired perturbation outcomes; on the Adamson benchmark, achieved all-gene delta Pearson 0.5530 and DE20 direction accuracy 0.6724, approaching paired supervised performance under a weaker supervision regime.
- First-author manuscript in preparation.

10/2024 – Present **Adaptive TMLEs**
Researcher, Advisor: Dr. Mark van der Laan (UC Berkeley, Department of Statistics)

- Developed an Adaptive TMLE framework for estimating mean potential outcomes under treatment rules using a HAL-based CATE working model; derived the efficient influence function and second-order remainder for the projected policy-value parameter, and showed that the estimator avoids direct inverse propensity score targeting under practical positivity violations.
- Introduced a regularized TMLE targeting step that replaces the standard inverse-propensity clever covariate with its projection onto the CATE working-model score space; simulations showed lower MSE and improved coverage relative to IPW, AIPW, and standard TMLE under limited overlap, with a real-data Right Heart Catheterization analysis producing more stable estimates and shorter confidence intervals.
- Proposed a Lepski-type adaptive truncation procedure with a brake mechanism and evaluated EIF, plug-in, and targeted-bootstrap variance estimators; in a severe positivity setting, targeted-bootstrap brake achieved

bias 0.027, MSE 0.114, and coverage 0.956, avoiding both weak-truncation instability and over-truncation bias.
- Two first-author and corresponding-author manuscript submitted.

08/2025 – 05/2026 **Generative Synthetic Data for Causal Inference**
Independent Researcher
- Formalized why predictive or distributional fidelity of synthetic tabular data does not guarantee causal validity, showing that ATE preservation requires both accurate covariate generation and treatment-effect contrast preservation.
- Proposed a hybrid synthetic-data framework that generates covariates separately from treatment and outcome mechanisms, enabling randomized synthetic treatment assignment, targeted augmentation under practical positivity problems, and estimator benchmarking before real-data analysis.
- Demonstrated across simulations and ACTG experiments that hybrid LLM-based synthesis improves ATE preservation, reduces finite-sample MSE under limited overlap, and provides useful diagnostics for comparing OR, IPW, AIPW, and TMLE.
- Sole-author paper accepted to STAI-X 2026.

06/2025 – 06/2026 **Layer Placement in Parameter-Efficient Fine-Tuning & Source–Target Transfer Prediction of LLMs**
Independent Researcher
- Developed a projected-residual theory for PEFT layer placement, showing that adapter effectiveness is governed by residual signal, activation conditioning, inter-layer coupling, and depth-dependent cost rather than raw gradient magnitude alone.
- Introduced Layer Cards, reusable metadata for selecting LoRA layer budgets and placement regimes under accuracy–efficiency trade-offs.
- Validated the method across GPT-2, LLaMA2-7B, and Qwen3-8B on NLG, reasoning, and commonsense benchmarks; selective 10-layer LoRA often matched or exceeded all-layer LoRA, and Layer Card transfer improved 5-layer accuracy over raw-gradient placement by 9.14% on average.
- Formulated transfer prediction as target-conditional source ranking, introducing target-normalized rank percentiles to remove target-difficulty and metric-scale confounding when comparing source tasks.
- Developed late-layer voting for gradient-based transfer prediction, aggregating per-layer source rankings rather than raw gradient scores; improved within-target Spearman correlation over global gradient aggregation from 0.351 to 0.519 on T5 and 0.117 to 0.439 on LLaMA.
- Demonstrated strong directional transfer asymmetry across task pairs, with mean absolute asymmetry of 0.246 on T5 and 0.240 on LLaMA; showed that smaller-model transfer measurements can serve as low-cost pilots for predicting larger-model source-target transfer behavior.
- Two first-author and corresponding-author manuscripts submitted.

08/2023 – 05/2024 **Residual feature integration & LLM synthetic data**
Researcher, Advisor: Dr. Lexin Li (UC Berkeley, School of Public Health) and Dr. Linjun Zhang (Rutgers University, Department of Statistics)

- Developed the first provably safe mechanism residual feature integration (ReFine) for deep pretrained models. ReFine solves long standing negative transfer problem in deep transfer learning theory by providing rigorous generalization bound.
- Demonstrated that ReFine significantly outperforms linear probing, adapters, knowledge distillation, LoRA, and DANN-GATE under large domain gaps across modalities and architectures. Provided theoretical and empirical insight into why existing transfer methods fail and showed that ReFine supports parallel multi-source scaling and adapt-time multimodal extension, enabling pretrained single-cell foundation models to incorporate spatial signals for lymph-node anatomical classification despite being trained without them.

Yichen Xu

Email: yichen_xu@berkeley.edu; Tel: 510-697-0346

- Developed an LLM-based synthetic oversampling pipeline that outperforms numerous baselines such as SMOTE and duplication for imbalanced classification and mitigation of spurious correlation, with empirical validation of scaling laws for high-quality synthetic data.
- First-author paper accepted to *ICLR 2026*; first/co-first-author manuscripts under revision at JASA.

03/2022 – 05/2023 **Efficient and effective algorithms for densest subgraph discovery and maintenance**
Research Assistant, Advisor: Dr. Chenhao Ma and Dr. Yixiang Fang (CUHK–SZ, School of Data Science)

- Proposed the c-core framework and extended to dynamic graphs. Developed two algorithms for rapid insertions and deletions of edges and nodes in dynamic graphs, specifically designed to maintain the c-core structure for efficient DS retrieval. Demonstrated that these algorithms transition updates from a full re-computation to localized updates.
- Empirically demonstrated that c-core provided over 1000 times speedup and the algorithms for insertions and deletions can achieve a speed increase of over 300 times.
- Two first-author papers published in *The VLDB Journal* and *SIGMOD 2023*.

PUBLICATIONS (* indicates equal contribution, † indicates corresponding author)

1. **Yichen Xu**†, Mark van der Laan. Adaptive Targeted Maximum Likelihood Estimation of the Mean Potential Outcome under a Treatment Rule. 2026, under review, ***Journal of Causal Inference***.
2. **Yichen Xu.** Generative Synthetic Data for Causal Inference: Pitfalls, Remedies, and Opportunities. Full paper accepted at the First Conference on Statistics and Trustworthy AI for Cross (X)-Domain Acceleration (STAI-X 2026). arXiv: https://arxiv.org/abs/2604.23904
3. **Yichen Xu**†, Susan Gruber, Mark van der Laan. Investigating Targeting Strategies and Truncation in TMLE for the Average Treatment Effect under Practical Positivity Violations. 2026, under review, ***International Journal of Biostatistics***.
4. **Yichen Xu***†, Yuyang Liang*, Shan Dai, Tianyang Hu, Tsz Nam Chan, Chenhao Ma. Understanding and Guiding Layer Placement in Parameter-Efficient Fine-Tuning of Large Language Models. submitted, ***NeurIPS***. Preprint: *https://arxiv.org/abs/2602.04019*
5. **Yichen Xu***, Ryumei Nakada*, Linjun Zhang, Lexin Li. Residual Feature Integration is Sufficient to Prevent Negative Transfer. ***Proceedings of the International Conference on Learning Representations (ICLR),*** 2026, pp. 77738–77763. arXiv: *https://arxiv.org/abs/2505.11771*
6. Ryumei Nakada*, **Yichen Xu***, Lexin Li, Linjun Zhang. Synthetic Oversampling: Theory and A Practical Approach Using LLMs to Address Data Imbalance. ***Journal of the American Statistical Association***, in revision. Preprint: *https://arxiv.org/abs/2406.03628v2*
7. **Yichen Xu**, Chenhao Ma, Yixiang Fang, Zhifeng Bao. Efficient and Effective Algorithms for Densest Subgraph Discovery and Maintenance. ***The VLDB Journal***, 33, 1427–1452 (2024). https://doi.org/10.1007/s00778-024-00855-y
8. **Yichen Xu**, Chenhao Ma, Yixiang Fang, Zhifeng Bao. Efficient and effective algorithms for generalized densest subgraph discovery. ***Proceedings of the ACM on Management of Data (SIGMOD)***, 1(2), Article 169, 1–27 (2023). https://doi.org/10.1145/3589314

HONORS & AWARDS

2024	ACM SIGMOD ARI Reproducibility Badges Awarded
2023	B.S. degree with Honors, First Class
2019 – 2023	Dean's List

SCIENTIFIC PRESENTATIONS

Yichen Xu

Email: yichen_xu@berkeley.edu; Tel: 510-697-0346

- Residual Feature Integration is Sufficient to Prevent Negative Transfer. ICLR 2026, Rio de Janeiro, Brazil.
- Efficient and Effective Algorithms for Generalized Densest Subgraph Discovery. ACM SIGMOD 2023, Seattle, USA

SERVICES

Reviewer for ICML, NeurIPS, STAI-X

SOFTWARES

- Libretto: Open-source platform for deterministic symbolic-music generation and LLM-based structural assessment, built with a custom grammar, corpus-calibrated evaluation, and iterative refinement. Offers free public access to downloadable practice scores and audio examples across keys, rhythms, melodic patterns, and clefs, supporting beginning music learners and music-education use. Website: https://libretto.site/. GitHub: https://github.com/Xyc-arch/Libretto. arXiv: https://arxiv.org/abs/2606.22708.

TEACHING EXPERIENCES (Graduate Student Instructor, UC Berkeley)

07/2025 – 08/2025	Department of Statistics, Stat135: Concept of Statistics
08/2024 – 12/2024	Department of Statistics, Stat2: Introduction to Statistics
07/2024 – 08/2024	Department of Public Health, PH142: Intro Probability and Statistics
01/2024 – 05/2024	Department of Public Health, PH244: Big Data: Public Health Perspective
08/2023 – 12/2023	Department of Mathematics, MATH54: Linear algebra and Differential Equations

INTERNSHIP

06/2026 – 08/2026 KEIJI AI, AI Research Engineer Intern

- Developed self-improving coding agents for CRF-to-SDTM transformation, using execution-grounded gates to prevent fabricated records and missingness in clinical data prepared for regulatory review.
- Built a constraint-guided generate–validate–revise loop with residual-error traceability and resumable remediation, reducing Pinnacle 21 findings and warnings without rerunning the full workflow.

SKILLS

- Programming Skills: Python (Expert); C++ (Expert); R (Expert); Linux (Expert); LaTex (Expert); Java (Medium); MATLAB (Medium); MySQL (Medium)
- Hobbies: Flute and guitar (Member, Chinese Orchestra, CUHK–SZ, 2019 – 2022); Calligraphy; Tennis