Stratified Permutational Berry--Esseen Bounds and Their Applications to Statistics

Peng Ding
University of California, Berkeley

The stratified linear permutation statistic arises in various statistics problems, including stratified and post-stratified survey sampling, stratified and post-stratified experiments, conditional permutation tests, etc. Although we can derive the Berry--Esseen bounds for the stratified linear permutation statistic based on existing bounds for the non-stratified statistics, those bounds are not sharp, and moreover, this strategy does not work in general settings with heterogeneous strata with varying sizes. We first use Stein's method to obtain a unified stratified permutational Berry--Esseen bound that can accommodate heterogeneous strata. We then apply the bound to various statistics problems, leading to stronger theoretical quantifications and thereby facilitating statistical inference in those problems.

Bio:

Peng Ding is an Associate Professor in the Department of Statistics at UC Berkeley. He obtained my Ph.D. from the Department of Statistics, Harvard University in May 2015 and worked as a postdoctoral researcher in the Department of Epidemiology, Harvard T. H. Chan School of Public Health until December 2015. His main interests are causal inference, experimental design, survey sampling, missing data, and applied statistics.

Optimal Treatment Regimes with Policy-Relevant Guarantees

Mats Stensrud
EPFL

I will study how to define, choose, and estimate effects that inform individualized treatment decisions. I introduce a new class of treatment regimes that can outperform conventional “optimal” regimes, and I analyze the estimation challenges for both. A central difficulty, common to most existing approaches, is that as we include more covariates and aim for finer personalization, the required assumptions grow stronger.

As an alternative, I propose a strategy for detecting and estimating customized group-level effects, interpretable as coarsened versions of conventional optimal regimes. I provide explicit frequentist guarantees that these groups differ in their effects. Finally, I show that, in realistic settings, group-based strategies can substantially outperform state-of-the-art optimal-regime methods, even when those methods are correctly implemented with modern doubly robust machine-learning estimators.

Bio:

Mats Stensrud is an MD and associate professor at the Institute of Mathematics, EPFL, where he holds the Chair of Biostatistics (since 2020) and serves as director of the doctoral program in mathematics (since 2025). He develops causal and statistical methodology with two main goals: make the target of inference scientifically meaningful, and keep assumptions transparent and testable. He is particularly interested in settings where exposures and outcomes evolve over time (longitudinal and time-to-event data). Many of his works are inspired by applications in clinical medicine and epidemiology.

Before he came to EPFL, Stensrud worked with Miguel Hernán and Jamie Robins at Harvard School of Public Health as a Kolokotrones Research Fellow and Fulbright Research Scholar. Before he became a full time academic, he had a short career as resident doctor in internal medicine.

He received his MD, Dr.Philos in Neuroscience and BSc in Mathematics from the University of Oslo. He also holds a Msc in Statistics from the University of Oxford.

Robust generalization and transfer learning with sensitivity analysis

Hyunseung Kang
University of Wisconsin-Madison

In recent years, a popular task in data science has been to learn a feature of a new, target distribution Q, denoted as f(Q), using an existing dataset from a source distribution P. This task, broadly referred to as generalization or transfer learning, relies on a key assumption called conditional exchangeability. Unfortunately, this assumption is often untenable in practice and worse, the assumption cannot be verified with data.

The main theme of the talk is to study the sensitivity of learning about f(Q) when conditional exchangeability is violated. The first part of the talk presents how to conduct nonparametric and efficient inference for f(Q) under a well-known sensitivity model from causal inference. We also propose a method to benchmark or calibrate the sensitivity parameter using an idea from design sensitivity by Rosenbaum (2004). The second part of the talk proposes a novel measure to assess the sensitivity of learning f(Q) under local violations of conditional exchangeability. The measure, which we call SLOPE, is inspired by an idea from Hampel (1974)'s influence curve and sensitivity analysis from causal inference. Practically, SLOPE helps investigators address questions about which datasets to use for generalization and broadly reflects a philosophy that robust generalization should start with a robust study design.

This work is joint work with Xinran Miao (UW-Madison) and Jiwei Zhao (UW-Madison).

Bio:

Hyunseung (pronounced Hun-Sung) is an Associate Professor in the Department of Statistics at the University of Wisconsin–Madison. In 2015, he received his Ph.D. in Statistics from the University of Pennsylvania, where he was advised by Tony Cai and Dylan Small. From 2015 to 2016, he completed his NSF postdoctoral fellowship under Guido Imbens at Stanford University, and he has been at Madison since 2017. His research focuses on developing methods for causal inference, with particular emphasis on (a) instrumental variables and unmeasured confounding, (b) semi/nonparametric inference, and (c) dependence. He is also interested in applications to genetics, education, and political science.

Transfer Learning of CATE with Kernel Ridge Regression

Molei Liu
Peking University

The proliferation of data has sparked significant interest in leveraging findings from one study to

estimate treatment effects in a different target population without direct outcome observations. However, the transfer learning process is frequently hindered by substantial covariate shift and limited overlap between (i) the source and target populations, as well as (ii) the treatment and control groups within the source. We propose a novel method for overlap-adaptive transfer learning of conditional average treatment effect (CATE) using kernel ridge regression (KRR). Our approach involves partitioning the labeled source data into two subsets. The first one is used to train candidate CATE models based on regression adjustment and pseudo-outcomes. An optimal model is then selected using the second subset and unlabeled target data, employing another pseudo-outcome-based strategy. We provide a theoretical justification for our method through sharp non-asymptotic MSE bounds, highlighting its adaptivity to both weak overlaps and the complexity of CATE function. Extensive numerical studies confirm that our method achieves superior finite-sample efficiency and adaptability. We conclude by demonstrating the effectiveness of our approach on two real-world examples.

Bio:

Liu Molei is a researcher and assistant professor at the Peking University Health Science Center, joint with International Center for Mathematical Sciences. He earned his Ph.D. in Biostatistics from Harvard University in 2022 and served as an assistant professor at Columbia University from 2022 to 2024. His primary research focuses on statistical theory and methodology, including high-dimensional complex data inference, data fusion, semi-supervised learning, transfer learning, and distributionally robust optimization. His work has been published in top-tier journals such as the Journal of the Royal Statistical Society: Series B, Journal of the American Statistical Association, Biometrika, and the Journal of Machine Learning Research. Additionally, he engages in extensive collaborative research in applied fields such as biomedical informatics and genomics, with findings published in prestigious journals including Science and npj Digital Medicine. He has received the Sanford Bolton Scholar Award from the Mailman School of Public Health at Columbia University.

New results and methods for the categorical instrumental variable model

Thomas Richardson
University of Washington

In this talk we consider categorical instrumental variable (IV) models with instrument, treatment and outcome taking finitely many values.

We will first review existing results for the simplest case in which the treatment and outcome are both binary.

We then consider the general case and present a simple closed-form characterization of the set of joint distributions of potential outcomes that are compatible with a given observed data distribution in terms of a set of inequalities. These inequalities unify several different IV models defined by versions of the independence and exclusion restriction assumptions and are shown to be non-redundant.

Lastly we consider statistical inference for a set of linear functionals of the joint counterfactual distribution, such as pairwise average treatment effects. Using a tail bound on the Kullback–Leibler divergence we construct a set of confidence intervals that will simultaneously contain both the upper and lower bounds for each estimand with probability at least $(1-\alpha)$. We illustrate our method using data from the Minneapolis Domestic Violence Experiment.

This is joint work with Yilin Song (Columbia University), F. Richard Guo (University of Michigan) and Gary Chan (University of Washington).

Bio:

Thomas S. Richardson is a Professor in the Department of Statistics at the University of Washington, Seattle. He received his PhD in Logic, Computation and Methodology from Carnegie-Mellon University. Richardson is a Fellow of the Center for Advanced Studies in the Behavioral Sciences at Stanford University and a 2014 U.S. Fulbright Fellowship recipient. In 2022 he was a co-recipient of the inaugural Rousseeuw Prize for Statistics.

Richardson's research interests include causal inference, graphical models and multivariate statistics. Together with James M. Robins, he created a new unification of graphs and potential outcome models via Single World Intervention Graphs (SWIGs). He has also developed two new classes of graphical models, called ancestral graphs and nested Markov models, which represent acyclic causal systems in which unmeasured confounding variables may be present.

Perturbed Double Machine Learning: Nonstandard Inference Beyond the Parametric Length

Zijian Guo
Zhejiang University

We study inference on a low dimensional functional β in the presence of possibly infinite dimensional nuisance parameters. Classical inferential methods are typically based on the Wald interval, whose large sample validity rests on the asymptotic negligibility of the nuisance error; for example, estimators based on the influence curve of the parameter (Double/Debiased Machine Learning DML estimators) are asymptotically Gaussian when the nuisance estimators converge at rates faster than $n^{-1/4}$. Although, under suitable conditions, such negligibility can hold even in nonparametric classes, it can be restrictive. To relax this requirement, we propose Perturbed Double Machine Learning (Perturbed DML) to ensure valid inference even when nuisance estimators converge at rates slower than $n^{-1/4}$. Our proposal is to 1) inject randomness into the nuisance estimation step to generate a collection of perturbed nuisance models, each yielding an estimate of β and a corresponding Wald interval, and 2) filter out perturbations whose deviations from the original DML estimate exceed a threshold. For Lasso nuisance learners, we show that, with high probability, at least one perturbation produces nuisance estimates sufficiently close to the truth, so that the associated estimator of β is close to an oracle estimator with knowledge of the true nuisances. Taking the union of the retained intervals delivers valid coverage even when the DML estimator converges more slowly than $n^{-1/2}$. The framework extends to general machine learning nuisance learners, and simulations show that Perturbed DML can have coverage when state of the art methods fail.

Bio:

Zijian Guo is a Qiushi Chair Professor at Zhejiang University. He earned his B.S. in Mathematics from The Chinese University of Hong Kong in 2012 and his Ph.D. in Statistics from the University of Pennsylvania in 2017, advised by Prof. Tony Cai (COPSS Award laureate). From 2017 to 2025, he served on the faculty of the Department of Statistics at Rutgers University, advancing from Assistant Professor to tenured Associate Professor, and in 2025 he returned to China to join the Data Science Center at Zhejiang University. His research interests include causal inference, high-dimensional statistics, multi-source learning and distributionally robust optimization, non-regular statistical inference, and methods at the interface of optimization and statistics, with an emphasis on robust learning and inference under heterogeneity and distribution shift and applications in health and genetics. His work has appeared in the Annals of Statistics, Journal of the Royal Statistical Society: Series B, Cell Genomics, Journal of the American Statistical Association, Biometrika, Journal of Machine Learning Research, and the Journal of Econometrics. He serves on the editorial boards of the Journal of the American Statistical Association (Theory and Methods) and TEST.

Identifying Average Causal Effects in Regression Discontinuity Design with Auxiliary Data

Xiaohua Zhou
Peking University

Regression discontinuity design (RDD) is widely employed in both medical and social science research, where treatment assignment is determined by whether a running variable exceeds a predefined threshold. While most studies focus on estimating local causal effects at the threshold, the challenge of identifying global treatment effects remains largely unexplored. In this talk, we introduce a novel and practically feasible framework for identifying and estimating the global average causal effect in RDDs, leveraging the existence of an auxiliary variable and an additional dataset.

This is a joint with Xinqin Feng, a PhD. candidate at Peking University

Bio:

Xiao-Hua Zhou is PKU Endowed Chair Professor and Chair of the Department of Biostatistics at Peking University. He was selected into the National High-level Oversea Talent Program of China. He is a Fellow of the American Association for the Advancement of Science, the American Statistical Association, and the Institute of Mathematical Statistics.

For over three decades, Dr. Zhou has dedicated his effort to develop new biostatistical methods and their applications to medicine and public health, striving to extract statistical structures from complex data and develop mathematically interpretable statistical methods. His research has primarily centered on developing new statistical methods in three closely interrelated medicine areas, diagnostic test evaluation, precision medicine, and medical cost analysis. Dr. Zhou has published over 300 statistical methodology and medical papers and is either the corresponding author or senior author on most of them; many of them have been published in top statistical journals, such as JJRSSB, JASA, Biometrics, Biometrika, Annals of Statistics. Dr. Zhou is a main contributor for correcting the verification and imperfect gold standard bias in evaluation of the accuracy of diagnostic test. Dr. Zhou has also made important methodologic contributions in causal inference, particularly in statistical methods for handling broken randomized experiments and complex data scenarios. He has led multiple U.S. national-level statistical methodology projects, as well as key initiatives such as the Key Program of National Natural Science Foundation of China, the National Key Research and Development Program, and the National Development and Reform Commission of China.

Dr. Zhou has received many honors, including Mitchell Prize from International Society for Bayesian Analysis, a prestigious Research Career Scientist Award from U.S. Federal Government Department of Veterans Affairs, Advisory Committee Service Award from U.S. FDA.

On the least favorable path

Lin Liu
Shanghai Jiaotong University

In this talk, we will discuss some recent work on the application and extension of semi-parametric theory, which has been the core mathematical framework for causal inference in the past decades. We will discuss the role of the least favorable paths in such theory and how to construct some non-standard least favorable paths in several examples. One example is somewhat classical, related to IV with continuous instruments, whereas the other examples motivated us to extend the current semi parametric theory framework from linear spaces to Riemannian manifolds. We will discuss some open-ended problems at the end.

Bio:

Lin Liu is an assistant professor at Shanghai Jiao Tong University, with an interest in causal inference, semi parametric theory and related theoretical and applied problems.

Doubly Robust Alignment for Large Language Models

Chengchun Shi
London School of Economics

This talk presents an application of classical causal inference methods to a modern AI problem -- reinforcement learning from human feedback (RLHF) for aligning large language models with human preferences. While RLHF has demonstrated promising results, many algorithms are highly sensitive to misspecifications in the underlying preference model (e.g., the Bradley-Terry model), the reference policy, or the reward function, resulting in undesirable fine-tuning. To address model misspecification, we draw on doubly robust methods from the causal inference literature and propose a doubly robust preference optimization algorithm that remains consistent when either the preference model or the reference policy is correctly specified (without requiring both). Our proposal demonstrates superior and more robust performance than state-of-the-art algorithms, both in theory and in practice. The code is available at https://github.com/DRPO4LLM/DRPO4LLM.

Statistical Inference with Mixed-Effect Model for Covariate-Adaptive Randomized Experiments

Yang Liu
Renmin University of China

Recent clinical studies increasingly involve a long factor with many levels, e.g., investigation sites, resulting in a large number of strata that must be accounted for either through design or subsequent analysis. This complication has raised concerns by the U.S. Food and Drug Administration (FDA, 2023) regarding the adequacy of standard statistical methods, whose performance may deteriorate, and their properties become unclear when the number of strata is relatively large. In this work, we offer a first-time rigorous solution by employing mixed-effect models in covariate-adaptive randomized experiments. We show that the mixed-effect estimate achieves lower variance in treatment effect estimation than its fixed-effect counterpart in the presence of the long factor with many levels. This variance reduction is attributable to marginal imbalances induced by the randomization procedure, suggesting that designs promoting finer covariate balance lead to more efficient inference and increased statistical power. Furthermore, we demonstrate that, as the sample size grows, the mixed- and fixed-effect estimators become asymptotically equivalent. Our theoretical findings are validated through simulations and a clinical trial case study and provide new insights on the design and analysis of clinical trials in the presence of a large number of strata.

Bio:

Dr. Liu received his Ph.D. in Statistics from The George Washington University in 2022. He is currently an Assistant Professor at the Institute of Statistics and Big Data, Renmin University of China. His research focuses on the development of innovative methods for adaptive design and statistical inference, with applications in clinical trials and industrial A/B testing.

An Empirical Tale for the Predictive Role of Covariate Shift in Treatment Effect Generalization

Ying Jin
University of Pennsylvania

Many existing approaches to generalizing statistical inference amidst distribution shift operate under the covariate shift assumption, which posits that the conditional distribution of unobserved variables given observable ones is invariant across populations. However, recent empirical investigations have demonstrated that adjusting for shift in observed variables (covariate shift) is often insufficient for generalization. In other words, covariate shift does not typically "explain away" the distribution shift between settings. Here, we present a series of empirical evidence from large-scale multi-site replication studies to support a new role of covariate shift in "predicting" the strength of the unknown conditional shift. Analyzing 680 studies across 65 sites, we find that even though the conditional shift is non-negligible, its strength can often be bounded by that of the observable covariate shift. However, this pattern only emerges when the two sources of shifts are quantified by our proposed standardized, "pivotal" measures. We then interpret this phenomenon by connecting it to similar patterns that can be theoretically derived from a random distribution shift model. Finally, we demonstrate that exploiting the predictive role of covariate shift leads to reliable and efficient uncertainty quantification for target estimates in generalization tasks with partially observed data.

Bio:

Ying Jin is an Assistant Professor in Statistics and Data Science at the Wharton School, University of Pennsylvania. Prior to that, she was a Wojcicki-Troper Postdoctoral Fellow at Harvard Data Science Initiative from 2024 to 2025. She obtained her PhD in Statistics from Stanford University in 2024, advised by Professors Emmanuel Candès and Dominik Rothenhäusler. Her research centers around statistical uncertainty quantification for black-box AI models, generalizability, distributional robustness, causal inference, and their applications in biomedical discovery and human decisions.

Robust Estimation of Gene Regulatory Networks in Perturb-seq Experiments under Unmeasured Confounding Using Proxy and Instrumental Variables

Hongzhe Li
University of Pennsylvania

Emerging single-cell technologies that integrate CRISPR-based genetic perturbations with single-cell RNA sequencing, such as Perturb-seq, have substantially advanced our understanding of gene regulation. While Perturb-seq data provide valuable causal insights into gene–gene interactions, statistical concerns remain regarding unobserved confounders that may bias inference. These latent factors may arise not only from technical sources such as batch effects but also from intrinsic molecular features of regulatory elements captured in Perturb-seq experiments. Although methods for analyzing large-scale Perturb-seq data are rapidly maturing, approaches that explicitly account for such unobserved confounders are still lacking. Here, we propose a novel method to recover gene regulatory networks from Perturb-seq experiments with robustness to arbitrarily omitted confounders. Our framework leverages proxy and instrumental variable strategies to exploit the rich information embedded in perturbations, enabling unbiased estimation of the underlying Directed Acyclic Graph (DAG) of gene expression. Simulation studies and analyses of data-driven experiments demonstrate that our method outperforms baseline approaches that ignore unmeasured confounding, yielding more accurate recovery of the true gene regulatory DAG.

Bio:

Dr. Hongzhe Li is Perelman Professor of Biostatistics, Epidemiology and Informatics at the Perelman School of Medicine at the University of Pennsylvania. He is Vice Chair for Research Integration, Director of Center of Statistics in Big Data and former Chair of the Graduate Program in Biostatistic at Penn. He is also a Professor of Statistics and Data Science at the Wharton School and a Professor of Applied Mathematics and Computational Science at Penn. Dr. Li's research focuses on developing statistical and modern machine learning and AI methods for analysis of large-scale genetic, genomics and metagenomics data and theory on high dimensional statistics.

XMR: A cross-population Mendelian randomization method for causal inference using genome-wide summary statistics

Can Yang
The Hong Kong University of Science and Technology

The rapid advancement of genome-wide association studies (GWASs) has led to significant discoveries, yet the persistent lack of representation of non-European populations limits the generalizability and equity of these findings. In Mendelian randomization (MR), small sample sizes in underrepresented populations often result in insufficient instrumental variables (IVs) and unreliable causal estimates. To address these challenges, we propose XMR, a novel cross-population MR method that leverages data from large-sample populations to improve causal inference in smaller, data-limited groups. By accounting for the shared genetic architecture across populations and population-specific genetic effects, XMR effectively increases the number of IVs and carefully controls for confounding factors. This approach improves statistical power while maintaining rigorous control over false positive rates.

Through extensive simulations and real-data analyses, XMR demonstrates superior performance, including well-calibrated results in negative-control experiments, higher statistical power, and the ability to identify novel causal relationships in East Asian, Central/South Asian, and African populations that conventional methods fail to detect. These findings reveal potential heterogeneity in causal patterns across populations, highlighting the importance of population-specific research. We anticipate that XMR will serve a valuable tool for advancing global equity in genetic studies.

Short Bio:

Prof. Yang Can is now a professor at the Department of Mathematics, HKUST. He obtained BS. and M.Phil. Degrees at Zhejiang University 2003, and 2006, and Ph.D. Degree at HKUST in 2011. He was a postdoc (2011-2012) and associate scientist (2012-2014) at Yale. He serves as an associate editor for Annals of Applied Statistics, associate editor for Genetics, and academic editor for PLOS Computational Biology. His research focuses on the development of statistical and computational methods with their applications in large-scale data analysis. His research papers have appeared in high impact journals, such as Nature, Nature Machine Intelligence, Nature Computational Science, Nature Communications, Proceedings of the National Academy of Sciences (PNAS), IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), Annals of Statistics, and The American Journal of Human Genetics. Prof. Yang has also established industrial collaborations supported by the Innovation and Technology Fund of the Hong Kong Government.

Identification and estimation of the oracle decision rule for achieving fairness with biased data

Wang Miao
Peking University

The fairness of statistical and machine learning models has become a prominent concern. When data contain bias towards certain demographic groups, models trained with such data may inherit and amplify the bias, leading to unfair or unreasonable decisions. In this paper, we propose a novel framework for characterizing and addressing fairness issues by introducing the oracle decision, a latent variable representing the ideal and fair decision that should have been made. The oracle decision differs from the observed decision for individuals who are treated unfairly. In contrast to existing approaches that aim to mitigate the observed decision under certain fairness constraints, we propose to assess the degree of unfairness in the data and advocate for making decisions by identifying the oracle decision rule with biased data. We establish the identification of oracle decision and develop estimation methods, including a sieve maximum likelihood estimator of the oracle decision rule and an estimator of the degree of unfairness.

Bio:

Wang Miao is currently tenured Associate Professor at the Department of Probability and Statistics and Center for Statistical Science, Peking University. He obtained BS and PhD degrees in 2012 and 2017 at Peking University, and did postdoctoral research at the Department of Biostatistics at Harvard University during 2017-2018. His research interest spans in causal inference, missing data analysis, data fusion, semiparametrics, and their application.

Nuisance parameter tuning for inference in observational studies

Rajarshi Mukherjee
Harvard University

The purpose is to discuss the issue of nuisance parameter tuning for estimating quantities in observational studies, such as the average treatment effect and measures of conditional dependence. Typical methods of estimating such quantities of interest rely on estimating nuisance functions often through the lens of nonparametric and/or high-dimensional machine learning methods. Whereas many popular ideas pertain to tuning these nuisance function estimation from a prediction perspective and subsequently perform downstream bias correction for valid inference of low dimensional summaries of interest in the observational studies of interest, cases are explored to show that there exists a delicate interplay between nuisance function estimation strategies, type of estimators that uses these nuisance functions in its pipeline of estimation of the final object of interest, and sample splitting strategies that are now popular to allow flexible methods of nuisance function estimation without jeopardizing the standard errors of estimators of the downstream objects of interest. The above is explored through the lens of specific functionals that arise in the context of causal inference, and both are studied in nonparametric and high-dimensional regimes.

Bio:

I am an Associate Professor in the Department of Biostatistics at Harvard T.H. Chan School of Public Health. Previously, I was an Assistant Professor in the Division of Biostatistics at UC Berkeley following my time as a Stein Fellow in the Department of Statistics at Stanford University. I obtained my PhD in Biostatistics from Harvard University, advised by Prof. Xihong Lin. I am generally interested in understanding broad aspects of causal inference in observational studies in modern data settings, with a focus on learning about fundamental challenges in the statistical analysis of environmental mixtures and their effects on the cognitive development of children and cognitive decline in aging populations. My research is also motivated by learning through applications in large-scale genetic association studies, developing statistical methods to quantify the effects of climate change on human health, and understanding the effects of homelessness on human health.

Classification Learning under Selective Labels

Xiaojie Mao
Tsinghua University

We study the problem of classification with selectively labeled data, whose distribution may differ from the full population due to historical decision-making. We exploit the fact that in many applications historical decisions were made by multiple decision-makers, each with different decision rules. We analyze this setup under a principled instrumental variable (IV) framework and rigorously study the identification of classification risk. We establish conditions for the exact identification of classification risk and derive tight partial identification bounds when exact identification fails. We further propose a unified cost-sensitive learning (UCL) approach to learn classifiers robust to selection bias in both identification settings. We further theoretically and numerically validate the efficacy of our proposed method.

Bio:

Xiaojie Mao is an associate professor in Management Science and Engineering at Tsinghua University. He did his undergraduate in Mathematical Economics at Wuhan University and Ph.D. in Statistics and Data Science at Cornell University. His research interest is in causal inference and data-driven decision-making. His research has appeared in top journals and conferences, such as Operations Research, Management Science, Information Systems Research, Journal of Machine Learning Research, Journal of the Royal Statistical Society Series B, NeurIPS, ICML, AISTATS, COLT, etc.

Dynamic treatment effects: high-dimensional doubly robust inference under model misspecification

Yuqian Zhang
Renmin University of China

Estimating dynamic treatment effects is crucial across various disciplines, as it provides insights into the time-dependent causal impact of interventions. However, this task is challenging due to time-varying confounding, which can produce biased estimates. Moreover, accurately specifying the growing number of outcome and treatment assignment models becomes increasingly difficult. Double robustness, which permits model misspecification, is particularly valuable in addressing these challenges. This paper introduces a novel "sequential model doubly robust" estimator. We show that root-N inference can be achieved as long as at least one nuisance model is correctly specified at each exposure time, even in the presence of high-dimensional covariates. To ensure valid inference under minimal model correctness conditions, we demonstrate that the optimal strategy is to estimate treatment propensity score functions sequentially forward in time, followed by outcome regression functions sequentially backward in time, with each step building on previous estimates. Although the nuisance estimates themselves do not achieve root-N rates, the combination of carefully designed loss functions and the sequential estimation strategy ensures final root-N inference for the causal parameter of interest.

Bio:

Yuqian Zhang is an Assistant Professor at the Institute of Statistics and Big Data, Renmin University of China. His current research interests include causal inference, semi-supervised learning, high-dimensional statistics, machine learning, and precision medicine. He received the Junior Researcher Award from ICSA in 2025 and the Best Student Paper Award from the Nonparametric Statistics Section of ASA in 2021.

Causal Underpinnings of Data Synthesis and Batch Effects

Bijan Mazaheri
Dartmouth College

Modern information exists in diverse forms: traditional datasets, large-scale AI models, and accumulated human knowledge. As incremental gains from single data sources plateau, integrating and synthesizing information across these varied forms has emerged as a critical frontier in AI research. However, without careful attention to how we expand and combine data and knowledge,

this synthesis risks leading us away from, rather than closer to, the truths we seek. In this talk, we examine data synthesis both broadly and through the narrow scope of “batch effect correction” algorithms applied to biological data. We will see that the task fundamentally relies on causal dependencies, highlighting causality as a foundational principle underpinning information synthesis in the era of AI.

Bio:

Bijan Mazaheri is an assistant professor at Dartmouth Engineering. He studies issues that arise when synthesizing information from multiple datasets, modalities, and batches, with a focus on applications in information security, biology, and human health. He uses tools from theoretical computer science and statistics—including sample complexity, mixture models, and causal inference. Bijan received his undergraduate degrees in computer science and physics from Williams College, studied mathematics as a Herchel Smith Fellow at the University of Cambridge, and earned his PhD in Computing and Mathematical Sciences from Caltech. Bijan completed postdoctoral work at the Eric and Wendy Schmidt Center at the Broad Institute of MIT and Harvard and remains an affiliated researcher at the center.

Minimax Optimal Design Incorporating Spillover and Carryover Effects

Hanzhong Liu
Tsinghua University

In causal inference, it is often assumed that interference does not occur between experimental units. In practice, however, a unit's potential outcomes may depend not only on its own treatment but also on treatments received by other units (spillover effects) and on its past treatments (carryover effects). Such dependencies violate the stable unit treatment value assumption and introduce significant challenges for experimental design. To address these complexities, we develop a minimax optimal design that explicitly accounts for both spillover and carryover effects, thereby improving the accuracy of causal effect estimation. The proposed design is particularly advantageous in multi-unit experimental settings, as it reduces the required sample size and implementation costs. We further investigate the asymptotic properties of Horvitz–Thompson estimators for direct and spillover effects, establishing their consistency and asymptotic normality under the proposed design. To ensure valid inference, we introduce conservative variance estimators and examine the consequences of potential misspecification in the order of carryover effects. Finally, simulations and application-based analyses demonstrate that the proposed design consistently outperforms existing alternatives across a wide range of scenarios.

Bio:

Hanzhong Liu is a tenured Associate Professor at the Department of Statistics and Data Science, Tsinghua University. He received his Ph.D. in Statistics from Peking University and conducted postdoctoral research at the University of California, Berkeley. His primary research interests include high-dimensional statistical inference and causal inference. His work has been published in leading journals such as PNAS, JASA, and Biometrika.

A quantum experiment with joint exogeneity violation

Yuhao Wang
Tsinghua University

In randomized experiments, the assumption of potential outcomes is usually accompanied by the joint exogeneity assumption. Although joint exogeneity has faced criticism as a counterfactual assumption since its proposal, no evidence has yet demonstrated its violation in randomized experiments. In this paper, we reveal such a violation in a quantum experiment, thereby falsifying this assumption, at least in regimes where classical physics cannot provide a complete description. We further discuss its implications for potential outcome modelling, from both practial and philosophical perspectives.

Bio:

Yuhao Wang is an assistant professor in the Institute for Interdisciplinary Information Sciences (IIIS), Tsinghua University. Before joining Tsinghua, Yuhao was a postdoctoral research associate at the University of Cambridge. Yuhao received his Ph.D. from Massachusetts Institute of Technology and his Bachelor from Tsinghua University. Yuhao's main research interests include causal inference and distribution-free test. Yuhao received the Forbes China 30 under 30 award in 2021 and is currently serving as an Associate Editor for the Electronic Journal of Statistics.

Bayes methods automatically debias semiparametric estimators

Wenlong Mou
University of Toronto

A central problem in causal estimation is to estimate a non-linear functional of the underlying high-dimensional parameters. When the problem dimension is high, plug-in approaches suffer from large bias due to non-linearity. To overcome this problem, recent years have seen the development of debiasing techniques with improved guarantees.

In this talk, we reveal a surprising phenomenon for this classical problem: the posterior mean estimator automatically reduces the bias of plug-in estimators for a large class of non-linear functionals. In particular, we show that the $1/\sqrt{n}$ non-asymptotic convergence rate for the posterior mean in the challenging high-dimensional regime, without the need for explicit bias correction. This new finding shows the advantage of Bayesian methods for causal estimation, and opens up new avenues for research in this area.

Bio:

Wenlong Mou is an Assistant Professor in the Department of Statistical Sciences at University of Toronto. In 2023, he received his Ph.D. degree in Electrical Engineering and Computer Sciences (EECS) from UC Berkeley. Prior to Berkeley, he received his B.Sc. degree in Computer Science and B.A. degree in Economics, both from Peking University. Wenlong's research interests include machine learning theory, mathematical statistics, optimization, and applied probability. He is particularly interested in data-driven decision-making in modern AI paradigms. His works have been published in many leading journals in statistical machine learning. His research has been recognized by the INFORMS Applied Probability Society as a Best Student Paper finalist.

MR2G: A novel framework for causal network inference using GWAS summary data

Haoran Xue
City University of Hong Kong

Inferring a causal network among multiple traits is essential for unraveling complex biological relationships and informing interventions. Mendelian randomization (MR) has emerged as a powerful tool for causal inference, utilizing genetic variants as instrumental variables (IVs) to estimate causal effects. However, when the directions of causal relationships among traits are unknown, reconstructing the underlying causal network becomes challenging. In particular, the presence of cycles or feedback loops, which are common in biological systems, poses additional challenges for causal network inference, and remains largely under-studied with standard MR approaches and existing IV-based network inference methods. To address these issues, we introduce MR2G, a new statistical framework that enables robust inference of causal networks, including those with cycles, directly from GWAS summary statistics. MR2G is built on a formally defined recursive causal graph model that rigorously links direct causal effects to MR estimands. It recovers a biologically interpretable causal network from pairwise MR effect estimates, while incorporating a network-informed IV screening strategy to reduce pleiotropic bias and improve robustness. Through realistic simulations, MR2G demonstrates superior accuracy and robustness in recovering complex causal structures, including those involving feedback loops. We apply MR2G to GWAS summary statistics for six complex diseases and nine cardiometabolic risk factors. MR2G not only recovers well-established causal pathways but also uncovers multiple feedback relationships, highlighting its utility in disentangling complex and biologically plausible causal networks from large-scale genetic data.

Bio:

Haoran Xue is currently an Assistant Professor in the Department of Biostatistics, City University of Hong Kong. Before this, he was a Postdoctoral Associate (2021-2023) in the Division of Biostatistics, University of Minnesota. He obtained his Ph.D. (2021) in statistics at the School of Statistics, University of Minnesota, where he was fortunately advised by Professor Xiaotong Shen and Professor Wei Pan. Prior to that, he got his B.S. (2016) in statistics from the School of the Gifted Young, University of Science and Technology of China. Haoran's research mainly focuses on development of novel methods in causal inference and statistical genetics, as well as their applications in biomedical studies.