CAUSAL CONFERENCE PROGRAM 2016

Causal Conference

DATE: November 10-12

LOCATION: Columbia University

- THURSDAY NOV 10: Uris 326

- FRIDAY NOV 11: Warren 207

- SATURDAY NOV 12: Uris 141

The conference will also be streamed live for those unable to register and attend.

Our first conference on causal inference will focus on ways to estimate a number of different types of treatment effects both when standard assumptions such as ignorability and SUTVA hold, as well as cases where these assumptions fail.

SPEAKERS

1. Edoardo Airoldi: Harvard University
2. Joshua Angrist: Massachusetts Institute of Technology
3. Gary Chan: University of Washington
4. David Choi: Carnegie Mellon University
5. Dean Eckles: Massachusetts Institute of Technology
6. Michael Hudgens: University of North Carolina, Chapel Hill
7. Fan Li: Duke University
8. Jamie Robins: Harvard University
9. Sherri Rose: Harvard University
10. Paul Rosenbaum: University of Pennsylvania
11. Donald Rubin: Harvard University
12. Dylan Small: University of Pennsylvania
13. Mark van der Laan: University of California, Berkeley
14. Stefan Wager: Columbia University
15. Xiaoru Wu: Facebook
16. Cunhui Zhang: Rutgers University

Thursday, November 10, 2016: tutorial

Targeted Learning

This course will introduce targeted learning methods for causal inference. It will emphasize understanding and responding to the challenges posed by observational cohorts and randomized trials including high-dimensional "big data." Examples from the areas of health policy, medicine, and epidemiology will be used as illustrations to translate research questions into statistical estimation problems with accurate interpretation of results. Course content covers material from Chapters 1-6 of "Targeted Learning" by van der Laan & Rose, as well as additional advances.

2:15 3:00 Sherri Rose/Mark vander Laan

3:00 3:45 Sherri Rose/Mark vander Laan

3:45 4:00 (Coffee break)

4:00 4:45 Sherri Rose/Mark vander Laan

4:45 5:30 Sherri Rose/Mark vander Laan

Friday, November 11, 2016: talks

9:00 9:15 Opening remarks

9:15 10:00 Donald Rubin

10:00 10:20 (Coffee break)

10:20 11:05 Stefan Wager

11:05 11:50 Mark vander Laan

11:50 1:20 (Lunch break)

1:20 2:05 Xiaoru Wu

2:05 2:50 Cunhui Zhang

2:50 3:10 (Coffee break)

3:10 3:55 Michael Hudgens

3:55 4:40 David Choi

5:00 6:00 Reception in the Department of Statistics

6:30 (Dinner for speakers)

Saturday, November 12, 2016: talks

9:00	9:45	Fan Li
9:45	10:30	Paul Rosenbaum
10:30	10:50	(Coffee break)
10:50	11:35	Joshua Angrist
11:35	12:20	Dylan Small
12:20	1:50	(Lunch break)
1:50	2:35	Jamie Robins
2:35	3:20	Gary Chan
3:20	3:40	(Coffee break)
3:40	4:25	Edo Airoldi
4:25	5:10	Dean Eckles

Titles and Abstracts

Donald Rubin - TBA

Stefan Wager

Title: Efficient Inference of Average Treatment Effects in High Dimensions via Approximate Residual Balancing

Abstract: There are many settings where researchers are interested in estimating average treatment effects and are willing to rely on the unconfoundedness assumption, which requires that the treatment assignment be as good as random conditional on pre-treatment variables. The unconfoundedness assumption is often more plausible if a large number of pre-treatment variables are included in the analysis, but this can worsen the finite sample properties of standard approaches to treatment effect estimation. There are some recent proposals on how to extend classical methods to the high dimensional setting; however, to our knowledge, all existing method rely on consistent estimability of the propensity score, i.e., the probability of receiving treatment given pre-treatment variables. In this paper, we propose a new method for estimating average treatment effects in high dimensional linear settings that attains dimension-free rates of convergence for estimating average treatment effects under substantially weaker assumptions than existing methods: Instead of requiring the propensity score to be estimable, we only require overlap, i.e., that the propensity score be uniformly bounded away from 0 and 1. Procedurally, out method combines balancing weights with a regularized regression adjustment.

Mark vander Laan

Title: One-step Targeted MLE and the Highly Adaptive Lasso

Abstract: We review targeted maximum likelihood estimation (TMLE), which provides a general template for the construction of asymptotically efficient plug-in estimators of a differentiable target parameter. TMLE involves maximizing a parametric likelihood along a so-called least favorable parametric model through an initial estimator of the data density, and iterating this updating process till convergence. For one-dimensional target parameters, we propose a universal least-favorable submodel that (a) guarantees that the TMLE only takes one step, and thus always exists in closed form, and (b) renders the targeting step of the TMLE maximally effective, resulting in meaningful practical improvements relative to an iterative TMLE. We generalize this to multivariate and infinite-dimensional parameters, and illustrate our proposal in several causal estimation problems.

The asymptotic efficiency of the TMLE relies on the asymptotic negligibility of a second-order term. This typically requires the initial data density estimator to converge fast enough. We propose a new estimator, the Highly Adaptive LASSO (HAL), of the data density (and its functionals) that converges at a sufficient rate regardless of the dimensionality of the problem, under almost no additional regularity. This allows us to propose a one-step TMLE that is asymptotically efficient in great generality across all models and differentiable target parameters. We demonstrate the practical performance of HAL and its corresponding TMLE for the average causal effect.

Xiaoru Wu - TBA

Cunhui Zhang

Title: Lasso Adjustments of Treatment Effect Estimates in Randomized Experiments

Abstract: We provide a principled way for investigators to analyze randomized experiments when the number of covariates is large. Investigators often use linear multivariate regression to analyze randomized experiments instead of simply reporting the difference of means between treatment and control groups. Their aim is to reduce the variance of the estimated treatment effect by adjusting for covariates. If there are a large number of covariates relative to the number of observations, regression may perform poorly because of overfitting. In such cases, the Lasso may be helpful. We study the resulting Lasso-based treatment effect estimator under the Neyman-Rubin model of randomized experiments. We present theoretical conditions that guarantee that the estimator is more efficient than the simple difference-of-means estimator, and we provide a conservative estimator of the asymptotic variance, which can yield tighter confidence intervals than the difference-of- means estimator. Simulation and data examples show that Lasso-based adjustment can be advantageous even when the number of covariates is less than the number of observations. Specifically, a variant using Lasso for selection and OLS for estimation performs particularly well, and it chooses a smoothing parameter based on combined performance of Lasso and OLS. This talk is based on joint work with Adam Bloniarz, Hanzhong Liu, Jasjeet Sekhon and Bin Yu

Michael Hudgens

Title: Causal Inference in the Presence of Interference

Abstract: A fundamental assumption usually made in causal inference is that of no interference between individuals (or units), i.e., the potential outcomes of one individual are assumed to be unaffected by the treatment assignment of other individuals. However, in many settings, this assumption obviously does not hold. For example, in infectious diseases, whether one person becomes infected depends on who else in the population is vaccinated. In this talk we will discuss recent approaches to assessing treatment effects in the presence of interference. Inference about different direct and indirect (or spillover) effects will be considered in a population where individuals form groups such that interference is possible between individuals within the same group but not between individuals in different groups. An analysis of an individually-randomized, placebo controlled trial of cholera vaccination in 122,000 individuals in Matlab, Bangladesh will be presented which indicates a significant indirect effect of vaccination.

David Choi

Title: Estimation of monotone treatment effects under interference

Abstract: Randomized experiments on social networks pose statistical challenges, due to the possibility of interference between units. We propose new methods for finding confidence intervals on the attributable treatment effect in such settings. The methods do not require partial interference, but instead require an identifying assumption that is similar to requiring "no defiers" or non-negative treatment effects. Network or spatial information can be used to customize the test statistic; in principle, this can increase power without making formal assumptions on the data generating process.

Fan Li

Title: Weighting beyond Horvitz-Thompson in causal inference

Abstract: Covariate balance is crucial for unconfounded descriptive or causal comparisons. However, lack of balance is common in observational studies. This article considers weighting strategies for balancing covariates. We define a general class of weights---the balancing weights---that balance the weighted distributions of the covariates between treatment groups. These weights incorporate the propensity score to weight each group to an analyst-selected target population. This class unifies existing weighting methods, including commonly used weights such as inverse-probability weights as special cases. General large-sample results on nonparametric estimation based on these weights are derived. We further propose a new weighting scheme, the overlap weights, in which each unit's weight is proportional to the probability of that unit being assigned to the opposite group. The overlap weights are bounded, and minimize the asymptotic variance of the weighted average treatment effect among the class of balancing weights. The overlap weights also possess a desirable small-sample exact balance property, based on which we propose a new method that achieves exact balance for means of any selected set of covariates. Two applications illustrate these methods and compare them with other approaches. This is a joint work with Alan Zaslavsky and Kari Lock Morgan.

Paul Rosenbaum

Title: Addressing bias from unmeasured dispositions in observational studies

Abstract: There are two treatments, each of which may be applied or withheld, yielding a 2×2 factorial arrangement with three degrees of freedom between groups. The differential effect of the two treatments is the effect of applying one treatment in lieu of the other. In randomized experiments, the differential effect is of no more or less interest than other treatment contrasts. Differential effects play a special role in certain observational studies in which treatments are not assigned to subjects at random, where differing outcomes may reflect biased assignments rather than effects caused by the treatments. Differential effects are immune to certain types of unobserved bias, called generic biases, which are associated with both treatments in a similar way. This is exemplified using three familiar models, a Rasch model, a symmetric multivariate logit model and a preference tree model. Differential effects are not immune to differential biases, whose possible consequences are examined by sensitivity analysis. Under certain conditions, the differential comparison of two treatments balances other treatments, including unmeasured treatments, that are governed by the same unmeasured disposition. Three scientific examples are presented.

Joshua Angrist

Title: Research Design Meets Market Design: Using Centralized Assignment for Impact Evaluation

A growing number of school districts use centralized assignment mechanisms to allocate school seats in a manner that reflects student preferences and school priorities. Many of these assignment schemes use lotteries to ration seats when schools are oversubscribed. The resulting random assignment opens the door to credible quasi-experimental research designs for the evaluation of school effectiveness. Yet the question of how best to separate the lottery-generated variation integral to such designs from non-random preferences and priorities remains open. This paper develops easily-implemented empirical strategies that fully exploit the random assignment embedded in the widely-used deferred acceptance mechanism and its variants. We use these methods to evaluate charter schools in Denver, one of a growing number of districts that integrate charter and traditional public schools in a unified assignment system. The resulting estimates show large achievement gains from charter school attendance. Our approach expands the scope for impact evaluation by maximizing the number of students and schools that can be studied using random assignment. We also show how to use DA to identify causal effects in models with multiple school sector.

Dylan Small

Title: Estimating the Malaria Attributable Fever Fraction Accounting for Parasites Being Killed by Fever and Measurement Error

Abstract: Malaria is a parasitic disease that is a major health problem in many tropical regions. The most characteristic symptom of malaria is fever. The fraction of fevers that are attributable to malaria, the malaria attributable fever fraction (MAFF), is an important public health measure for assessing the effect of malaria control programs and other purposes. Estimating the MAFF is not straightforward because there is no gold standard diagnosis of a malaria attributable fever; an individual can have malaria parasites in her blood and a fever, but the individual may have developed partial immunity that allows her to tolerate the parasites and the fever is being caused by another infection. We define the MAFF using the potential outcome framework for causal inference and show what assumptions underlie current estimation methods.

Current estimation methods rely on an assumption that the parasite density is correctly measured. However, this assumption does not generally hold because (i) fever kills some parasites and (ii) the measurement of parasite density has measurement error. In the presence of these problems, we show current estimation methods do not perform well. We propose a novel maximum likelihood estimation method based on exponential family g-modeling. Under the assumption that the measurement error mechanism and the magnitude of the fever killing effect are known, we show that our proposed method provides approximately unbiased estimates of the MAFF in simulation studies. A sensitivity analysis can be used to assess the impact of different magnitudes of fever killing and different measurement error mechanisms. We apply our proposed method to estimate the MAFF in Kilombero, Tanzania. This is joint work with Kwonsang Lee.

Jamie Robins

An Interventionist Redefinition of Pure Direct and Indirect Effects and Path Specific Effects

Pure direct and indirect effects and path specific effects are defined in terms of cross-world counterfactuals. Hence they do not have an interventionist interpretation. It follows that no set of randomized experiments suffices for their non-parametric identification without further assumptions. In this talk I provide an interventionist redefinition of pure direct and indirect effects and path specific effects that makes no reference to cross-world counterfactuals. I describe identification conditions for these redefined effects and show how they can in principle be tested in future randomized experiments. I show these identification conditions imply those required to identify the cross world pure direct and indirect effects and path specific effects, but that the converse is false.

Gary Chan

Title: Empirical balancing scores and balancing weights

Abstract: Propensity scores have been central to causal inference and are often used as balancing scores or balancing weights. Estimated propensity scores, however, may exhibit undesirable finite-sample performance. We take a step back to understand what properties of balancing scores and weights are desirable. For balancing scores, the dimension reduction aspect is important; whereas for balancing weights, a conditional moment balancing property is crucial. Based on these considerations, a joint sufficient dimension reduction framework is proposed for balancing scores, and a covariate functional balancing framework is proposed for balancing weights.

This presentation includes joint works with Ming-Yueh Huang, Raymond Wong, Phillip Yam and Zheng Zhang.

Edo Airoldi - TBA

Dean Eckles

Title: Massive meta-analysis using regularized instrumental variables, with an application to peer effects

Abstract: The widespread adoption of randomized experiments (i.e. A/B tests) in the Internet industry means that there are often numerous well-powered experiments on a given product. Individual experiments are often simple "bake-off" evaluations of a new intervention: They allow us to estimate effects of that particular intervention on outcomes of interest, but they are often not informative about the mechanisms for these effects or what other inventions might do. We consider what else we can learn from a large set of experiments. In particular, we use many experiments to learn about the effects of the various endogenous variables (or mechanisms) via which the experiments affect outcomes. This involves treating the experiments as instrumental variables, and so this setting is similar to, but somewhat different from, "many instrument" settings in econometrics and biostatistics. Motivated by the distribution of experiment first-stage effects, we present and evaluate regularization methods for improving on standard IV estimators.

Joint work with Alex Peysakhovich (Facebook AI Research).