The virtual European Causal Inference Meeting April 22nd and 23rd, 2020

Abstracts

Invited presentation 1

Apples and oranges: visualising and explaining non-collapsibility

Rhian Daniel
School of Medicine, Cardiff University, UK

When marginal and conditional effect measures differ beyond what can be explained by confounding and other such structural phenomena, the measure is labelled "non-collapsible". For example, the risk difference and ratio for a binary outcome are collapsible, but the odds ratio is not. Although by no means at the top of the list of things applied statisticians seeking to make causal inferences should worry about, it is an issue that infects most of what we do in one way or another, eg confounder selection, mediation analysis, meta-analysis... It's also a topic that often confuses people, with some subtleties missed by even those who understand it quite well. For example, is non-collapsibility essentially the same issue in models for binary and time-to-event outcomes? Does it matter if our time-to-event outcome model is in discrete or continuous time? Should "conditional" and "adjusted" be used interchangeably? Is there a paradox when contrasting power and precision in non-linear regression models? By drawing on existing literature (e.g. Neuhaus and Jewell, 1993; Sjölander et al, 2016) I will attempt to answer these questions, in a way that hopefully helps an epidemiologist travelling on a bus...

Invited presentation 2

What makes causal inferences honest?

Stijn Vansteelandt
Department of Applied Mathematics, Computer Science and Statistics, Ghent University, Belgium

The lack of identification and need for extrapolation can make modelling assumptions dangerously influential in a causal analysis. Over the past few decades, we have therefore seen a growing tendency to move away from pure parametric approaches in causal inference. Despite these advances, we often continue to rely on modelling assumptions more than we may realise, and treat them as representing a priori biological knowledge in how we do inference. This complicates the interpretation of the estimands we report, and may invalidate our inferences. In this talk, I will illustrate these problems, and review a roadmap which aims to overcome them, as has been put forward very explicitly by Mark van der Laan (and has been independently adopted by several other scholars). I will argue why, in my opinion, this roadmap is the right pathway towards honest causal inferences, and highlight challenges that must be addressed before it can be rolled out on a wide scale.

Contributed presentation 1

Proportionally-representative interventions: New Causal Estimands for Settings with Limited Resources

Aaron Sarvet, Kerollos Nashat Wanis, Jessica Young, Miguel A. Hernán and Mats J. Stensrud
Harvard School of Public Health

Investigators often evaluate treatment effects by considering settings in which all individuals are assigned a treatment of interest, assuming that an unlimited number of treatment units are available. However, many real-life treatments may be constrained and cannot be provided to all individuals in the population. For example, patients on the liver transplant waiting list cannot be assigned a liver transplant immediately at the time they reach highest priority because a suitable organ is not likely to be immediately available. In these cases, investigators may still be interested in the effects of treatment strategies in which a finite number of organs are available at a given time, that is, treatment regimes that satisfy resource constraints. Here, we describe an estimand that can be used to define causal effects of treatment strategies that satisfy resource constraints. We introduce a novel and simple class of inverse probability weighted estimators, and apply one such estimator to evaluate the effect of restricting or expanding utilization of 'increased risk' liver organs to treat patients with end-stage liver disease. Our method is designed to evaluate plausible and policy-relevant interventions in the setting of finite treatment resources.

Contributed presentation 2

Causal inference for semi-competing risks data: estimands, partial identifiability, sensitivity analysis and models

Daniel Nevo and Malka Gorfine
Tel Aviv University

An emerging challenge for time-to-event data is studying semi-competing risks, where two event times are of interest: the non-terminal event (e.g. disease diagnosis) time, and a terminal event (e.g. death) time. The non-terminal event is observed only if it precedes the terminal event, which may occur before or after the non-terminal event, leading to the latter being unobserved or even undefined. Studying treatment or intervention effects on the event times is complicated because for some units, the non-terminal event time may occur only under one treatment value but not the other. Until recently, existing approaches (e.g., the survivor average causal effect) generally disregarded the time-to-event nature of both outcomes. More recent research focused on principal strata effects within time-varying populations under either fully-Bayesian or a non-parametric Bayesian with an additional copula assumption approaches. In this talk, I will present and discuss alternative related non time-varying estimands, based on stratification of the population, which do not change with time. These estimands correspond to the scientific questions of interest, as I will exemplify using a real-data example of the effect of the APOE gene on Alzheimer's disease and death. I will then present a novel assumption utilizing the time-to-event nature of the data, that is generally weaker than the often-invoked monotonicity assumption. Our new assumption enables partial identifiability of causal effects of interest, namely bounds. Because the bounds might be too wide for practical use, I will also present and discuss a sensitivity analysis approach, and semi-parametric models.

Contributed presentation 3

Decomposing Causal Mechanisms in Duration Models with Unobserved Heterogeneity

Stephen Kastoryano[1] *and Jad Beyhum*[2]
Unaffiliated[1] *and University of Toulouse*[2]

This paper presents an econometric framework for optimal policy design which identifies the causal effects of a treatment policy regime and of the actual implementation of treatment when the outcome of interest is a duration variable. We consider a situation in which agents are randomized to a policy regime upon entering an initial state that prescribes a stochastic propensity to future treatment. Thereafter, at different moments in time and depending on their policy regime, surviving agents are randomized to actually receive treatment. Our dynamic potential outcomes framework provides non-parametric identification of: the ex-ante effect of the policy regime on the duration to exit, which may include placebo-type information effects, the ex-post baseline effect of actually receiving treatment on the duration to exit within a given policy regime, and the additional ex-post interaction effect of the policy regime and actually receiving treatment. We also extend our model to allow for dynamic information accumulation due to the past history of unobserved intermediate variables in a hazard model, which weakens several proportionality assumptions usually imposed. We further present an estimation procedure and produce simulation results. Lastly we illustrate the framework using data from the experimental design of the US National Job Corps Study which includes duration measures for employment and crime outcomes. In this setting we decompose the different causal effects of the Job Corps program on the probability of arrest operating via employment.

Contributed presentation 4

An efficient and robust approach to Mendelian randomization with measured pleiotropic effects in a high-dimensional setting

Andrew Grant and Stephen Burgess
University of Cambridge

Valid estimation of a causal effect using instrumental variables requires that all of the instruments are independent of the outcome conditional on the risk factor of interest and any confounders. In Mendelian randomization studies with large numbers of genetic variants used as instruments, it is unlikely that this condition will be met. Any given genetic variant could be associated with a large number of traits, all of which represent potential pathways to the outcome which bypass the risk factor of interest. Such pleiotropy can be accounted for using standard multivariable Mendelian randomization with all possible pleiotropic traits included as covariates. However, the estimator obtained in this way will be inefficient if some of the covariates do not truly sit on pleiotropic pathways to the outcome. We present a method which uses regularization to identify which out of a set of potential covariates need to be accounted for in a Mendelian randomization analysis in order to produce an efficient and robust estimator of a causal effect. The method can be used in the case where individual-level data are not available and the analysis must rely on summary- level data only. It can also be used in the case where there are more covariates under consideration than instruments, which is not possible using standard multivariable Mendelian randomization. We illustrate the method in an applied example which looks at the causal effect of urate plasma concentration on coronary heart disease.

Contributed presentation 5

An instrumental variable approach to estimating precision in linked data

James Doidge
Intensive Care National Audit and Research Centre

Linked data are increasingly being employed in observational and experimental research across many disciplines. Data linkage typically involves some degree of error or uncertainty. Estimates of precision— the proportion of a set of candidate links that are true—are useful when designing linkage algorithms and when accounting for linkage error or uncertainty in analysis of linked data. Without supplementary information or training data, the main option for estimating precision requires implausible assumptions of independence. While linkage algorithms may be robust to violation of these assumptions, the resulting estimates of precision can be strongly biased and are generally unsuitable for bias analysis or causal inference. There are few alternatives and most can only be applied in limited contexts. A novel approach for estimating precision in linked data, akin to instrumental variable analysis, is proposed. The approach involves identifying one single matching variable for which the standard independence assumption holds, and exchanging its discriminatory power for estimation of precision that is robust to dependencies between other matching variables and to many mechanisms of missing data. Consistency of the technique is demonstrated using simulated data and performance and generalisability is then explored using real-world data and compared to conventional techniques. Options for establishing the validity of potential instruments are discussed. While some questions remain unanswered, this approach represents a new variety of statistical tool that provides an angle on an otherwise difficult-to-estimate bias parameter and may have applications to unsupervised classification beyond the realm of data linkage.

Contributed presentation 6

Mediation analysis when outcome and mediator are semi-competing events with application in health disparities research

Linda Valeri[1], *Helene Jacqmin-Gadda*[2] *and Cecile Proust-Lima*[2]
Columbia University[1] *and Universite de Bordeaux*[2]

We propose novel methodology for mediation analysis to explain how much of the effect of the exposure on a terminal time-to-event outcome is attributed to the, non-terminal, potentially intermediate, time-to-event. Addressing this question is important in health disparities research when we seek to quantify inequities in timely delivery of treatment and its impact on patients' survival time. We formalize a type of direct and indirect effects using the potential outcome framework in the presence of semi-competing risks. Mediation is studied in a multistate model in continuous time. Simulation based as well as closed form estimators of the causal contrasts are developed. We show via simulations that mediation analysis ignoring censoring in mediator and outcome time-to event-processes and/or ignoring competing risks may give misleading results. Rigorous definition of the direct and indirect effects and estimation of the joint outcome and mediator distributions in the presence of semi-competing risks is crucial for valid investigation of mechanisms in continuous time. We employ this novel methodology to investigate the role of delaying treatment uptake in explaining racial disparities in cancer survival in a multicenter cohort study of colorectal cancer patients.

Contributed presentation 7

Introducing the causaloptim R package for symbolically deriving causal bounds

Michael Sachs, Arvid Sjölander and Erin E Gabriel
Karolinska Institutet

Given a graph that encodes structural assumptions and a causal query, one can determine whether that query is identified from observable data, and if so, provide an estimand. In situations where a causal query is not identified, one way to proceed is to attempt to find bounds for it. The approach to symbolically deriving bounds that was used in Balke and Pearl (1994) has been applied in other settings, but has not previously been broadly accessible to the research community. We have generalized and extended Balke and Pearl's approach for translating a graph plus a causal query into a constrained optimization problem. We have identified a class of graphs plus causal queries that are guaranteed to lead to linear programming problems and derived an algorithm for translating them into the optimization problem. We have implemented this algorithm in an R package with a graphical user interface to allow users to draw graphs interactively and specify causal queries that are in this class. The bounds are derived symbolically using a vertex enumeration algorithm and returned to the user as R functions, text formulas, and LaTeX equations. In this talk I will describe the approach and implementation, demonstrate a use of the package in a novel setting, and attempt to demystify the underlying optimization algorithm.

Alexander Balke and Judea Pearl. Counterfactual probabilities: Computational methods, bounds and applications. In Proceedings of the Tenth international conference on Uncertainty in artificial intelligence, pages 46–54, 1994.

Contributed presentation 8

Panel Based Experiments and Dynamic Causal Effects: A Finite Sample Perspective

Iavor Bojinov[1], *Ashesh Rambachan*[2] *and Neil Shephard*[2]
Harvard Business School[1] *and Harvard University*[2]

Abstract: In panel (or longitudinal) experiments, we randomly expose multiple units to different treatments and measure their subsequent outcomes, sequentially repeating the procedure numerous times. Using the extended potential outcomes framework, we define finite sample dynamic causal effects that capture the relative effectiveness of the alternative treatments in panel experiments. For the p-lag dynamic causal effects, an important subclass, we provide unbiased estimators over the randomization distribution (i.e., the distribution obtained by conditioning on the potential outcomes and assuming the treatment assignment is the only source of randomness). We further derive the finite sample limiting distribution of our estimators as either the sample size or the duration of the experiment increases, or both. Our approach provides a new framework for deriving finite sample central limit theorems in causal inference that exploits the underlying Martingale property of unbiased estimators. Using the limiting distribution, we derive a conservative test for the null hypothesis of zero average causal effect, a generalization of the usual Neyman style null hypothesis in classical causal inference. We also provide an exact randomization-based test for the sharp null of no treatment effect at any point in time, which is often called the Fisher null hypothesis. Our analysis is illustrated with simulations and through the study of a panel experiment that compares the relative effectiveness of humans versus algorithms at executing large financial trades.

Contributed presentation 9

Bayesian principal stratification with longitudinal data and truncation by death

Giulio Grossi[1], *Marco Mariani*[2], *Alessandra Mattei*[1] *and Fabrizia Mealli*[1]
University of Florence[1] *and Tuscan Regional Institute for Economic Planning*[2]

In many causal studies, outcomes are 'censored by death,' in the sense that they are neither observed nor defined for units who die. In such studies, focus is usually on the stratum of 'always survivors' up to a single fixed time s. Building on a recent strand of the literature, we propose an extended framework for the analysis of longitudinal studies, where units can die at different time points, and the main endpoints are observed and well-defined only up to the death time. We develop a Bayesian longitudinal principal stratification framework, where units are cross-classified according to the longitudinal death status. Under this framework, focus is on causal effects for the principal strata of units that would be alive up to a time point s irrespective of their treatment assignment, where these strata may vary as a function of s. We can get precious insights on the effects of treatment by inspecting the distribution of baseline characteristics within each longitudinal principal stratum, and by investigating the time trend of both principal stratum membership and survivor-average causal effects. We illustrate our approach for the analysis of a longitudinal observational study aimed to assess, under the assumption of strong ignorability of treatment assignment, the causal effects of a policy promoting start-ups on firms' survival and hiring policy, where firms' hiring status is censored by death.

Contributed presentation 10

Conditional separable effects: New estimands for causal inference conditional on post-treatment variables

Mats J. Stensrud[1,2], *Jessica G. Young*[1], *Aaron Sarvet*[1] *and Miguel A. Hernan*[1]
Harvard T. H. Chan School of Public Health[1] *and University of Oslo*[2]

Many researchers aim to study treatment effects on outcomes in individuals characterized by status on a particular post-treatment variable. For example, we may be interested in the effect of cancer therapies on quality of life, and quality of life is only well-defined in the those individuals who are alive. Similarly, we may be interested in the effect of vaccines on post-infections outcomes, which are only of interest in those individuals who become infected. In these settings, a naive contrast of outcomes conditional on the post-treatment variable does not have a causal interpretation, even in a randomized experiment. Therefore the effect in the principal stratum of those who would have the same value of the post-treatment variable regardless of treatment, such as the survivor average causal effect, is often advocated for causal inference. Whereas this principal stratum effect is a well defined causal contrast, it cannot be identified without strong untestable assumptions, and its practical relevance is ambiguous because it is restricted to an unknown subpopulation of unknown size. Here we formulate alternative estimands, which allow us to define the conditional separable effects. We describe the causal interpretation of the conditional separable effects, e.g. in settings with truncation by death, and introduce three different estimators. As an illustration, we use data from a randomized clinical trial to estimate a conditional separable effect of chemotherapies on quality of life in patients with prostate cancer.

Contributed presentation 11

Emulation of target trials to investigate causal effects of organ transplantation on survival

Ruth Keogh[1], *Daniela Schlueter*[2], *Nan van Geloven*[3], *Elisa Allen*[4], *Mo Al-Aloul*[5], *Freddy Frost*[6], *Dilip Nazareth*[6], *Oliver Rayner*[7], *Nicholas Simmonds*[8] *and Thom Daniels*[9]
London School of Hygiene and Tropical Medicine[1], *University of Liverpool*[2], *Leiden University*[3], *NHS*[4], *Wythenshawe Hospital Manchester*[5], *Liverpool Heart and Chest Hospital*[6], *Cystic Fibrosis Trust*[7], *Royal Brompton Hospital London*[8] *and University Hospital Southampton*[9]

Organ transplantation is a treatment option for severe disease. Understanding its impact on survival relies on observational data such as national transplant registers, as randomized trials are infeasible. Investigations have typically quantified the effect of transplant on survival using a hazard ratio, and the causal questions to be addressed have not been clearly articulated. In this work we specify 'target trials' to articulate different causal questions and illustrate how to answer them using observational longitudinal and survival data. The motivation is a study of the impact of lung transplantation on survival in cystic fibrosis using UK registry data. We focus on three causal questions: (i) the effect of joining the transplant wait-list in the 'real world', which may or may not result in a transplant; (ii) the effect of joining the wait-list in an 'ideal world' where patients receive a transplant within a short time, and (iii) the impact of transplant in those who actually received one. Analyses are based on forming a sequence of emulated trials within the longitudinal data, and time-dependent confounding is handled using adjustment and weighting. We show how the effects of transplant can be quantified in terms of functions of survival curves and discuss both population-average and 'personalised' effects. I will discuss the challenges of accounting for the fact that organs are a finite resource, and some possible solutions. Results will be given from the motivating application, where the answers to the three questions differ, but all point to lung transplant being beneficial for survival.

Contributed presentation 12

DAGbase: a database of human-drawn causal diagrams

Johannes Textor, Ankur Ankan, Franka Buytenhuijs, Jeroen Creemers, Laurens Stuurman, Shabaz Sultan and Inge Wortel
Radboud University Medical Center, Nijmegen, The Netherlands

Directed acyclic graphs (DAGs) are increasingly used in causal inference from observational data in fields ranging from Epidemiology to Econometrics. While the theoretical foundations of DAGs are rapidly expanding, there is still considerable debate about their usefulness in practical applications. There is currently only limited information about how DAGs are being used, but a recent review suggests that most researchers draw DAGs specifically for one study and do not synthesize or build on each other's work. One reason for this might be the lack of a platform that allows researchers to share DAGs and build on existing ones. Here we present the DAGbase, a platform for discovering and sharing causal diagrams. DAGbase already contains human-drawn DAGs from two data-sources: (1) a curated set of 6000 diagrams that were made publicly available on the dagitty.net platform; (2) a set of 300 diagrams that we extracted from scientific publications. DAGbase therefore provides a wealth of information of how causal diagrams are used in teaching and research. We illustrate the potential of DAGbase by evaluating the identifiability of desired causal effects through different covariate adjustment criteria in human-constructed DAGs. Further, we cluster DAGs on structural and semantic properties to identify areas of research that might benefit from merging the information provided in multiple DAGs. On the occasion of EUROCIM 2020, we would like to make the DAGbase publicly available as a resource for all causal inference researchers.

Contributed presentation 13

Applying the separable effects causal mediation framework to etiological questions in epidemiological research: Lessons learned and recommendations for the future

Ryan M. Andrews and Vanessa Didelez
Leibniz Institute for Prevention Research and Epidemiology

Recently, Aalen et al. (2019) proposed a "separable effects" approach to causal mediation analysis with a survival outcome and longitudinal mediator, based on the prior work of Didelez (2018) and Strohmaier et al. (2015). In this approach, the treatment is assumed to be separable into distinct components that (1) affect the survival outcome directly and (2) affect the outcome only through the mediator process. Under various assumptions, direct and indirect effects can be defined in terms of relative survival and cumulative hazards by combining the g-formula with an additive hazards model for the outcome and a sequential linear model for the mediator process. In this talk, we will discuss this method in the context of a novel application that explored whether and how the causal effect of physical activity on time-to-dementia among older adults is mediated by pulse pressure, which is a proxy of vascular health. We found that the separable effects analysis offered several benefits over other available causal mediation survival analysis methods, particularly with respect to the interpretability of effects. On the other hand, we also found the method could not readily handle common actual data features, like a non-linear mediator process, which led us to make compensatory assumptions that were questionable. As one of the first practical applications of the Aalen et al. method, we believe our study demonstrates the value of the approach and highlights several open questions that should be addressed for it to be usable in a wider range of realistic data settings.

Aalen, O. O., Stensrud, M. J., Didelez, V., Daniel, R., Røysland, K., and Strohmaier, S. (2019). Time-dependent mediators in survival analysis: Modeling direct and indirect effects with the additive hazards model. Biometrical Journal, 1-18.

Didelez, V. (2018). Defining causal mediation with a longitudinal mediator and a survival outcome. Lifetime data analysis, 25(4), 593-610.

Strohmaier, S., Røysland, K., Hoff, R., Borgan, Ø., Pedersen, T. R., and Aalen, O. O. (2015). Dynamic path analysis – a useful tool to investigate mediation processes in clinical survival trials. Statistics in medicine, 34(29), 3866-3887.

Contributed presentation 14

Ensuring valid inference for the effect of a treatment on a time-to-event endpoint in Cox models with variable selection

Kelly Van Lancker, Oliver Dukes and Stijn Vansteelandt
Ghent University

In assessing the effect of a treatment on an outcome in an observational study, a common challenge is choosing which covariates should be adjusted for. Routine practice is often based on stepwise methods or the Lasso; unfortunately, selection can lead to biased treatment effect estimators and inferential procedures that are not guaranteed to perform well at any finite sample size. In recent papers, Belloni, Chernozhukov and colleagues have developed a 'double selection' procedure for obtaining valid inferences in generalized linear models. This uses a second selection step which picks up confounders that are weakly predictive of the outcome but strongly associated with the exposure; these are the variables that may be mistakenly ignored by single selection procedures. However, so far there has been limited focus in the post-selection inference literature on survival outcomes. In this work, we develop a framework for obtaining valid tests, estimators and confidence intervals for a treatment effect parameter in a proportional hazards model after variable selection. The major complication that we face with survival data is that the key confounding variables may not be those that explain the censoring mechanism. We propose a simple procedure to overcome this, which can be implemented using existing software for penalized Cox regression. Simulation results show that the proposed methods yield valid inferences even when covariates are high-dimensional. The approaches proposed are illustrated on a real dataset.

Contributed presentation 15

But will it work for me? Personalizing randomized trial evidence with effect heterogeneity structure learned from observational data

Leah Comment[1] *and Giovanni Parmigiani*[2,3]
Foundation Medicine[1]*, Harvard University*[2] *and Dana-Farber Cancer Institute*[3]

In medicine, treatment decisions in the clinic are made based on whether a treatment is likely to be effective for the particular patient in question, making conditional causal effects most relevant for clinical decision making. When there is strong effect modification, traditional meta-analyses of average treatment effects can substantially over- or understate evidence for an individual profile. We propose a method to contextualize individual-level randomized controlled trial (RCT) data with effect heterogeneity information from observational real world data (RWD) using Bayesian additive regression trees (BART). Tree splitting criteria are extracted from BART models fit to the RWD, while the causal contrasts for the patient profile of interest are computed exclusively using RCT patients. When multiple trials are used as input data, the resulting posterior predictive distributions can be used in a personalized meta-analysis that summarizes the current level of evidence about treatment efficacy for patients with a particular covariate profile. With minimal modification, this procedure can be used to construct informative priors for leaf node values in BART models. We evaluate the method using simulations and discuss strategies for handling different data sets' availability of covariate information.

Contributed presentation 16

Prediction meets causal inference: the role of treatment in clinical prediction models

Saskia le Cessie[1], *N van Geloven*[1], *SA Swanson*[2], *CL Ramspek*[3], *K Luijken*[3], *M van Diepen*[3], *TP Morris*[4], *RHH Groenwold*[1,3], *HC van Houwelingen*[1] *and H Putter*[1]
Leiden University[1], *Erasmus MC*[2], *Harvard University*[3] *and UCL London*[4]

Causal inference research is typically strictly distinguished from prediction research. However, we argue that when developing a clinical prediction model in the presence of treatment initiation, methods from both domains are needed. Currently in the development of prediction models ad hoc approaches are used for patients who start a treatment. Such patients may for instance be excluded or censored. Consequences of such analysis choices on the interpretation of the calculated risks are typically neglected which may lead to wrongly interpreted risks. We propose that the way treatment is dealt with in clinical prediction models should be guided by a clearly defined prediction estimand. Analogous to the estimand framework recently proposed by the European Medicines Agency for clinical trials, we propose a 'predictimand' framework of different questions that may be of interest when predicting risk in relation to treatment. We provide a formal definition of the estimands matching these questions, give examples of settings in which each is useful and discuss appropriate estimators including their assumptions. Besides 'ignoring treatment', and 'including treatment in a composite outcome', we focus attention on the 'untreated risk' under a hypothetical intervention that would eliminate treatment and on the 'risk preceding treatment', i.e., the cumulative incidence curve. We illustrate the impact of the predictimand choice in a dataset of patients with end-stage kidney disease where part of the patients are receive a kidney transplantation during follow up. Clearly defining an estimand is equally important in prediction research as in causal inference.

Contributed presentation 17

Discovering toxic exposure mixtures and ranking variable set importance via cross-validated causal inference and ensemble machine learning

Chris J. Kennedy, Catherine Metayer, Todd P. Whitehead, Alan E. Hubbard and Mark J. van der Laan
UC Berkeley

There is increasing interest in evaluating the combined impact of groups of correlated exposures, termed "mixtures", in observational data. Compared to single-exposure analyses, mixture-based analyses may better detect interactive relationships between exposures, such as antagonistic or synergistic effects. We sought to create a method for mixture estimation and risk evaluation that avoided limitations of existing approaches. We also aimed to support mixture estimation on subsets of exposures, providing a new form of variable importance for sets of variables. Our method estimates mixtures using backfit semiparametric ensemble regression, which posits a nonparametric functional form for the mixture and also adjusts for confounding nonparametrically. The mixture and the confounding adjustment functions are estimated using ensemble machine learning, and a correction term is added to provide double robustness. We estimate the marginal outcome mean at quantiles of the mixture through an outer loop of cross-validated targeted maximum likelihood estimation (CV-TMLE). When ex- posures can be grouped into subsets (e.g. chemical classes), we estimate separate mixtures for each subset and rank the importance of exposure groups by the statistical significance of a counterfactual intervention setting all observations to the high vs. low mixture quantile. We applied our method to chemical exposures extracted from household dust to estimate their influence on childhood leukemia risk, and to rank the importance of chemical classes in terms of toxicity. We also evaluated the method's performance on standard benchmarks and new data simulations. Our method is provided as an open source R software package tlmixture.

Contributed presentation 18

Robust double machine learning for conditional exposure effects

Oliver Dukes and Stijn Vansteelandt
Ghent University

Due to concerns about parametric model misspecification, there is interest in using machine learn- ing to adjust for confounding when evaluating the causal effect of an exposure on an outcome. Unfortunately, exposure effect estimators that rely on machine learning predictions are generally subject to so-called plug-in bias, which can render naive p-values and confidence intervals invalid. Progress has been made via proposals like Targeted Maximum Likelihood Estimation and more recently Double Machine Learning, which rely on learning the conditional mean of both the outcome and exposure. Valid inference can then be obtained so long as both algorithms converge (sufficiently fast) to the truth. We will show that by updating the machine learning predictions in a specific way, we can develop exposure effect estimators that have good properties even when one of the first-stage algorithms does not converge to the truth, along with honest tests and confidence intervals. Our proposal leads to reduced bias and improved confidence interval coverage in moderate-samples, as we observe in simulations studies, and moreover enables the development of double machine learning procedures for effect measures for which no double robust estimator is known. We illustrate the proposal in a case study looking at the effect of obesity on the probability of survival within patients in the Ghent University Hospital Intensive Care Unit.

Contributed presentation 19

Challenges in emulating target trials

Bianca L De Stavola[1], *Daniel Tompsett*[1] *and John Robson*[2]
University College London[1] *and Queen Mary University London*[2]

The framework of target trial emulation (TTE) is increasingly adopted in comparative effectiveness research based on observational data. TTE has multiple advantages, starting from the clarity of explicitly specifying the hypothetical target experimental trial for the questions of interest. However, the challenges of analysing observational data to address causal questions, and in particular of dealing with time-varying con- founding and informative censoring, require careful selection of the estimand(s) of interest and, consequently of the causal model from which the estimand can be identified, as well as of the estimation methods. Intention-to-treat and per-protocol effects may both be of interest when emulating pragmatic trials. In each case inverse probability weighting is often used to deal with confounding of assig Moment and non- adherence. This may be quite inefficient however. In this work we will examine the advantages (and challenges) of adopting instead g-estimation, following its implementation as suggested by Vansteelandt and Sjolander (2016) and Duke and Vansteelandt (2018). These points will be illustrated via an investigation of the cardiovascular risk of type 2 diabetes patients undergoing 2nd line therapy, based on administrative data from general practitioner consultations of 147 East London GPs from 2012 to 2017.

Vansteelandt, S. and Sjolander, A. (2016). Revisiting g-estimation of the effect of a time-varying exposure subject to time-varying confounding. Epidemiologic Methods, 5, 37-56.

Dukes, O. and Vansteelandt, S. (2018). A Note on G-Estimation of Causal Risk Ratios. American Journal of Epidemiology, 187, 1079-1084.

Contributed presentation 20

Exploring heterogeneous treatment effects to inform the targeting of national health insurance programmes

Karla Diaz Ordaz[1] *and Noemi Kreif*[2]
London School of Hygiene and Tropical Medicine[1] *and University of York*[2]

Researchers evaluating the effects of social policies are often interested in identifying individuals who would benefit most. This could inform the design of expansion programmes, common in settings where resources are limited. This requires a better understanding of the drivers of treatment-effect heterogeneity. We explore treatment-effect heterogeneity and estimate conditional average treatment effects (CATE) in a study evaluating the effects of being enrolled in social health insurance scheme on having an assisted birth in a retrospective cohort of Indonesian mothers. CATEs are estimated for pre-specified (receipt of cash transfer) and data-driven effect-modifiers. Assuming no unobserved confounding and positivity, we obtain the CATEs for the pre-specified subgroups using the two-stage meta-learner (T-learner) and X-learner approaches. These approaches use estimates of the individual-level predictions of both potential outcomes, to calculate estimated individual-level treatment effects. The conditional expectation of the potential outcomes are estimated using alternatively parametric models and random forests (RF) with sample splitting. Next, we apply Causal Forests, which train the RFs directly on predicted individual-level treatment effects estimated in a first stage. We obtain a list of data-driven effect modifiers and estimate the corresponding CATEs. For the pre-specified subgroup, all methods resulted in similar inferences, namely there is no evidence that receipt of cash transfers is an effect modifier. Causal Forest test of heterogeneity is significant for the contributory insurance. For the data-driven effect modifiers, we find monotonously decreasing relationships between treatment effects and education levels and household wealth, suggesting that the Indonesian health insurance programme would benefit disadvantaged subpopulations.