EUROPEAN CAUSAL INFERENCE MEETING (EuroCIM - 2021)

Wednesday 26th May 2021

Theme: survival / time-to-event / competing risks / longitudinal / trials & estimands

Hosted by the MRC Integrative Epidemiology Unit and Population Health Sciences, Bristol Medical School, University of Bristol, Bristol, UK

Free online event on Zoom; registration details at https://www.eurocim.org/

Time (British Summer Time)	Session / talk information
	Intro session: chair Tom Palmer, helper Chin Yang Shapland
9:30 – 9:40	Welcome, Introductions, plan for the day
	Session 1: chair Eleanor Sanderson, helper Apostolos Gkatzionis Trials and estimands
9:40 – 10:00	Hege Michiels. Ghent University. A novel estimand to adjust for rescue treatment in randomised clinical trials
10:00 – 10:20	Camila Olarte Parra. University of Bath. Hypothetical estimands in clinical trials: a unification of causal inference and missing data methods
10:20 – 10:40	Matias Janvin. École Polytechnique Fédérale de Lausanne. Causal estimands in settings with recurrent events
10:40 – 11:00	Break
	Session 2: chair Apostolos Gkatzionis, helper Rosie Cornish Competing risks and time-to-event analyses
11:00 – 11:20	Pawel Morzywolek. Ghent University. On Estimation and Cross-validation of Dynamic Treatment Regimes with Competing Risks
11:20 – 11:40	Awa Diop. Laval University. Marginal Structural Models with Latent Class Growth Modeling of Time-varying Treatment
11:40 – 12:00	Daniel Nevo. Tel Aviv University. A matching framework for truncation by death problems
12:00 – 13:00	Lunch with virtual chat. Tom Palmer, Chin Yang Shapland Platform/invite link TBA
	Session 3: chair Rosie Cornish, helper Chin Yang Shapland Longitudinal analyses
13:00 – 13:20	Ruth Keogh. London School of Hygiene & Tropical Medicine. Estimating treatment effects on survival in an entirely treated cohort
13:20 – 13:40	Emily Granger. London School of Hygiene and Tropical Medicine. Estimating the effects of multi-level exposures on outcomes using longitudinal data: a simulation study
13:40 – 14:00	Matea Skypala. Delft University of Technology. Extending multistate models with g-computation to evaluate the effect of treatment timing

Time (British Summer Time)	Session / talk information
14:00 - 14:30	Break
	Session 4: chair Tom Palmer, helper Eleanor Sanderson Time-to-event and longitudinal analyses
14:30 – 14:50	Judith Lok. Boston University. Truncation by death and the survival-incorporated median: What are we measuring? And why?
14:50 – 15:10	Andrew Ying. University of Pennsylvania. Proximal Causal Learning for Complex Longitudinal Studies
15:10 – 15:50	Invited talk: Jessica Young. Harvard Medical School. Choosing a definition of causal effect in the face of competing events: considerations for interpretation and identification

Contact: Please direct any questions to tom.palmer@bristol.ac.uk

A novel estimand to adjust for rescue treatment in randomised clinical trials

Hege Michiels; Cristina Sotto; An Vandebosch; Stijn Vansteelandt

Ghent University

The interpretation of randomised clinical trial results is often complicated by intercurrent events. For instance, rescue medication is sometimes given to patients in response to worsening of their disease, either in addition to the randomised treatment or in its place. The use of such medication complicates the interpretation of the intention-to-treat analysis. In view of this, we propose a novel estimand defined as the intention-to-treat effect that would have been observed, had patients on the active arm been switched to rescue medication if and only if they would have been switched when randomised to control. This enables us to disentangle the treatment effect from the effect of rescue medication on a patient's outcome, while tempering the strong extrapolations that are typically needed when inferring what the intention-to-treat effect would have been in the absence of rescue medication. We propose a novel inverse probability weighting method for estimating this effect in settings where the decision to initiate rescue medication is made at one pre-specified time point. This estimator relies on specific untestable assumptions, in view of which we propose a sensitivity analysis. We use the method for the analysis of a clinical trial conducted by Janssen Pharmaceuticals, in which patients with type 2 diabetes mellitus can switch to rescue medication for ethical reasons. Monte Carlo simulations confirm that the proposed estimator is unbiased in moderate sample sizes.

Hypothetical estimands in clinical trials: a unification of causal inference and missing data methods

Camila Olarte Parra; Rhian Daniel; Jonathan Bartlett

University of Bath

The analysis and interpretation of randomised trials is often complicated by the occurrence of certain events that affect the interpretation of the treatment effect or preclude the observation of the outcome of interest. Treatment discontinuation, addition of rescue medication, or death prior to measurement of the outcome of interest are examples of such events. The ICH E9 addendum on estimands labels them as intercurrent events. The addendum outlines different target estimands in the presence of intercurrent events but does not suggest statistical methods for their estimation. In this talk, we focus on the hypothetical estimand, where the treatment effect is estimated under the hypothetical scenario in which we (somehow) intervene to prevent the intercurrent event from occurring. To estimate a hypothetical estimand, we consider methods from causal inference (G-formula and inverse probability of treatment weighting) and missing data (multiple imputation and mixed models). We establish that certain causal inference estimators are identical to certain missing data estimators. These links may help those familiar with one set of methods but not the other. Moreover, they allow us to transparently show using potential outcome language the assumptions missing data methods are relying on to estimate hypothetical estimands. We also present Monte Carlo simulations that provide evidence of the performance of the methods in different settings including intercurrent events happening at different time points during follow-up, the presence of the intercurrent event affecting the outcome, a potential deterministic nature of the intercurrent event and failing to account for time-varying covariates.

Causal estimands in settings with recurrent events

Matias Janvin; Jessica G. Young; Pål C. Ryalen; Mats J. Stensrud

École Polytechnique Fédérale de Lausanne

Many research questions concern treatment effects on outcomes that can reoccur several times in the same individual. For example, medical researchers are interested in treatment effects on hospitalizations in patients with heart failure, seizures in individuals with epilepsy, and sport injuries in athletes. However, causal inference can be challenging in settings with recurrent events for several reasons, including the presence of competing events like death. Many statistical estimands have been studied in recurrent event settings. However, the causal interpretations of these estimands, and the conditions that are required to identify these estimands from observed data, are often ambiguously described. Here we use a counterfactual framework for causal inference to formulate causal estimands in recurrent event settings, with and without competing events. We clarify that common statistical estimands can be formally defined as (controlled) direct effects and total effects, which are frequently employed in mediation analysis. Furthermore, we show that recent results on interventionsist's mediation estimands (separable effects) allow us to define new causal estimands in recurrent event settings. Using theory on separable effects, we clarify when treatment effects on the recurrent event of interest can be disentangled from treatment effects on the competing event. Furthermore, using counting process theory, we show how our estimands, which are articulated in discrete time, converge to classical continuous time recurrent events estimands in the limit of fine discretizations of time. Finally, we propose several estimators.

On Estimation and Cross-validation of Dynamic Treatment Regimes with Competing Risks

Pawel Morzywolek; Johan Steen; Wim Van Biesen; Johan Decruyenaere; Stijn Vansteelandt

Ghent University

An open question in intensive care medicine is when to best initiate renal replacement therapy in patients with acute kidney injury. Starting too soon may put patients at increased risk of adverse events and may imply an enormous cost to the health care system. Starting too late may likewise be detrimental to patients. In view of this, we use routinely collected observational data from the Ghent University Hospital to investigate different pre-specified treatment strategies for optimizing the initiation time of renal replacement therapy based on serum potassium, pH and fluid

balance in critically ill patients with acute kidney injury with the aim to minimize 30-day ICU mortality. For this purpose, we apply statistical techniques for evaluating the impact of specific dynamic treatment regimes in the presence of intensive care unit discharge as a competing event. We discuss two approaches, a non-parametric one - using an inverse probability weighted Aalen-Johansen estimator and a semiparametric one - using dynamic-regime marginal structural models. Furthermore, we suggest an easy to implement cross-validation technique that can be used as a benchmark to compare the performance of the optimal dynamic treatment regime against. Our work shows the potential of data-driven medical decision support based on routinely collected observational data.

Marginal Structural Models with Latent Class Growth Modeling of Time-varying Treatment.

Awa Diop; Caroline Sirois; Denis Talbot

Laval University

Latent class growth models (LCGM) are increasingly proposed as a solution to summarize time-varying treatment in a few distinct groups. When combined with standard approaches like Cox proportional hazards models, LCGM can fail to control time-dependent confounding bias. Instead, we proposed to combine LCGM with a working marginal structural model (MSM). The parameter of interest is nonparametrically defined as the projection of the true MSM onto the chosen working model. We showed that the data-driven estimation of the trajectory groups can be ignored. As such, parameters can be estimated using inverse probability of treatment weighting and conservative inferences can be obtained using a standard robust variance estimator. Simulation studies were performed to illustrate our approach and compare it with current practice. For all explored scenarios, we find that our proposed approach yield estimators with little or no bias. As expected, when using stabilized IPTW, all confidence interval coverages are closed to 95% (between 90% and 98%). Opposingly, regardless of the number of follow-up times and number of trajectory classes, alternative LCGM analyses are highly biased with low coverage of their confidence interval (between 2% and 67%). We will apply our LCGM-MSM approach to a database composed by 572 822 Quebecers aged 66 or more and who are statins initiators to estimate the effect of statin-usage trajectories on a first CVD event. Our proposal is relatively simple to implement and we expect it to yield results that are clinically meaningful, easy to interpret and statistically valid.

A matching framework for truncation by death problems

Daniel Nevo; Tamir Zehavi

Tel Aviv University

Even in a carefully designed randomized trial, outcomes for some study participants can be missing, or more precisely, ill-defined, because participants had died prior to date of outcome collection. This problem, known as truncation by death, means that a simple comparison between the treated and untreated among the survivors does not correspond to a causal effect. The treated and untreated are no longer balanced with respect to covariates determining survival. To overcome this problem, researchers often utilize principal stratification and focus on the Survivor Average Causal Effect (SACE). The SACE is the average causal effect among the subpopulation that will survive regardless of treatment status. In this talk, we present a new approach based on matching for SACE identification and estimation. We provide an identification result for the SACE that motivates the use of matching to restore the balance among the survivors. We discuss various practical issues, including the choice of distance measures, possibility of matching with replacement, post-matching SACE estimators, and non-parametric tests. Our simulation results demonstrate the flexibility and advantages of our approach. Because the cross-world assumptions needed for SACE identification can be too strong and are unfalsifiable, we also present sensitivity analysis techniques and illustrate their use in real data analysis. Finally, a recent alternative for SACE that does not demand cross-world unfalsifiable assumptions targets the conditional separable effects. We show how our approach can also be utilized to estimate these causal effects.

Estimating treatment effects on survival in an entirely treated cohort

Ruth Keogh

London School of Hygiene & Tropical Medicine

Treatments are sometimes introduced for all patients in a particular cohort. In this situation, estimating the treatment effect in the treated is not straightforward because there are no directly comparable untreated patients. The motivating application relates to a disease-modifying treatment in cystic fibrosis (CF), ivacaftor, which has been available for everyone in the UK with a specific genetic mutation since 2012. The impact of ivacaftor on several health outcomes has been demonstrated in randomized controlled trials, but it is of interest to understand its effect on survival, which has not previously been investigated. Longitudinal data from the UK Cystic Fibrosis Registry provides a resource for estimating the effects of this treatment on survival.

A strong assumption made in most applications of causal inference methods to observational data is that of lack of positivity, but when the entire cohort of interest receives the treatment this assumption is violated. I will discuss how negative control outcomes combined with use of difference-in-differences applied to survival outcomes can be used to assess bias in treatment effect estimates and obtain unbiased estimates under certain assumptions. Causal diagrams and the potential outcomes framework are used to explain the methods and assumptions. The use of different underlying analysis models, including Cox regression and Aalen's additive hazards model, will be discussed.

In the motivating example, the methods are applied to estimate the causal effect of ivacaftor on survival, making use of outcomes observed in patients before the treatment was introduced or who have an ineligible genotype.

Estimating the effects of multi-level exposures on outcomes using longitudinal data: a simulation study

Emily Granger; Ruth H. Keogh

London School of Hygiene and Tropical Medicine

Longitudinal data are often used in medical research to investigate the causal effect of a time-varying exposure on an outcome. In such settings, time-varying confounders are often themselves affected by exposure, in which case standard analysis methods obtain biased effect estimates. Recent years have seen considerable development and increased interest in generalised methods (g-methods) which can obtain estimate effects of longitudinal treatment regimes in complex longitudinal settings. To date, the majority of research on g-methods has focused on settings with a single binary exposure of interest and there is a lack of research investigating the performance of g-methods in settings with multi-level exposures, for example representing multiple treatment combinations. Our research aims to compare g-methods in situations with a multi-level exposure of interest.

First, we describe the use of three g-methods, namely: inverse probability of treatment weighted estimation of marginal structural models, the g-computation formula and g-estimation of structural nested models, for estimating the causal effects of multi-level exposures using longitudinal data. We then compare the three approaches in a series of simulated data scenarios with multi-level exposures, in terms of bias and relative efficiency.

Our work is motivated by an applied example in cystic fibrosis research. Many people with cystic fibrosis are prescribed at least one treatment to help improve lung function. While existing research has studied the effects of individual treatments, joint treatment effects are unknown. By defining a multi-level exposure that represents different treatment combinations, we can study the individual and joint treatment effects.

Extending multistate models with g-computation to evaluate the effect of treatment timing

Matea Skypala, Gabriela F. Nane; Saskia le Cessie; Nan van Geloven

Delft University of Technology

To avoid unnecessary treatments, a wait-and-see approach is commonly advised in medicine where a natural recovery of the patient is awaited before starting treatment. The impact of the exact duration of such a delay period is often unknown. In this work we develop causal methods that -under certain assumptions- allow evaluating treatment timing using observational data sources.

We combine multistate modelling with g-computation to target the counterfactual cumulative proportion of recovered patients for different delay periods. We first fit a multistate model describing the speed at which patients transition between the disease, treatment and recovery states. It uses Cox proportional hazards models for each of the transitions, employing the clock-reset time scale and including the wait time at which patients entered the treatment state as one of the covariates for the transition to recovery. We then use the multistate model to evaluate the effect of different delay periods by predicting the cumulative percentage of recovered patients had all unrecovered patients transitioned to treatment at the same wait time using g-computation. Uncertainty is quantified by bootstrapping.

We apply the developed methodology to estimate the expected cumulative proportion of pregnancies 1.5 years after diagnosis when delaying the start of intrauterine insemination treatment by 0, 3, 6 or 9 months in a cohort of 1896 couples with unexplained subfertility. Our method allows contrasting the expected number of additional pregnancies with the expected number of extra treatments, thus supporting the evaluation of these treatment delay strategies.

Truncation by death and the survival-incorporated median: What are we measuring? And why?

Judith J. Lok, Qingyan Xiang, Ronald J. Bosch

Boston University

One could argue that if a person dies, their subsequent health outcomes are missing. On the other hand, one could argue that if a person dies, their health outcomes are completely obvious. This talk considers the second point of view, and advocates to not always see death as a mechanism through which health outcomes are missing, but rather as part of the outcome measure. This is especially useful when some people's lives may be saved by a treatment we wish to study. We will show that both the median health score in those alive and the median health score in the always-survivors can lead one to believe that there is a trade-off between survival and good health scores, even in cases where in clinical practice both the probability of survival and the probability of a good health score are better for one treatment arm. To overcome this issue, we propose the survival-incorporated median as an alternative summary measure of health outcomes in the presence of death. It is the outcome value such that 50% of the population is alive with an outcome above that value. Survival-incorporated quartiles can be defined similarly. The survival-incorporated median can be interpreted as what happens to the average person. The survival-incorporated median is particularly relevant in settings with non-negligible mortality. We will illustrate our approach by estimating the effect of statins on neurocognitive function.

Proximal Causal Learning for Complex Longitudinal Studies

Andrew Ying, Shi Xu; Wang Miao; Eric J. Tchetgen Tchetgen

University of Pennsylvania

A standard assumption for causal inference with longitudinal data is that at each follow-up time, one has measured a sufficiently rich set of covariates to ensure that within covariate strata, subjects are exchangeable across observed treatment values, also known as sequential randomization assumption (SRA). Skepticism about SRA is often warranted because it hinges on investigators' ability to accurately measure covariates overtime capturing all potential sources of time-varying confounding. Realistically, confounding mechanisms can rarely be learned with certainty from measured covariates. One can therefore only ever hope that covariate measurements are at best proxies of true underlying confounding mechanisms, thus invalidating causal claims made on basis of SRA. We extend the proximal causal inference framework of Tchetgen Tchetgen et al. (2020), and Cui et al. (2020) to the longitudinal setting under a semiparametric marginal structural mean model (MSMM). The approach we propose offers an opportunity to learn about joint causal effects when SRA on the basis of measured time-varying covariates fails, by formally accounting for the covariate measurements as imperfect proxies of underlying confounding mechanisms. We establish sufficient conditions for nonparametric identification with time-varying proxies when SRA fails to hold. We provide a characterization of all regular and asymptotically linear estimators of the parameter indexing the MSMM, including a rich class of doubly robust estimators, and establish the corresponding semiparametric efficiency bound. Our approach is illustrated via extensive simulation studies and a data application on potential protective effects of the anti-rheumatic therapy Methotrexate (MTX) among patients with rheumatoid arthritis.