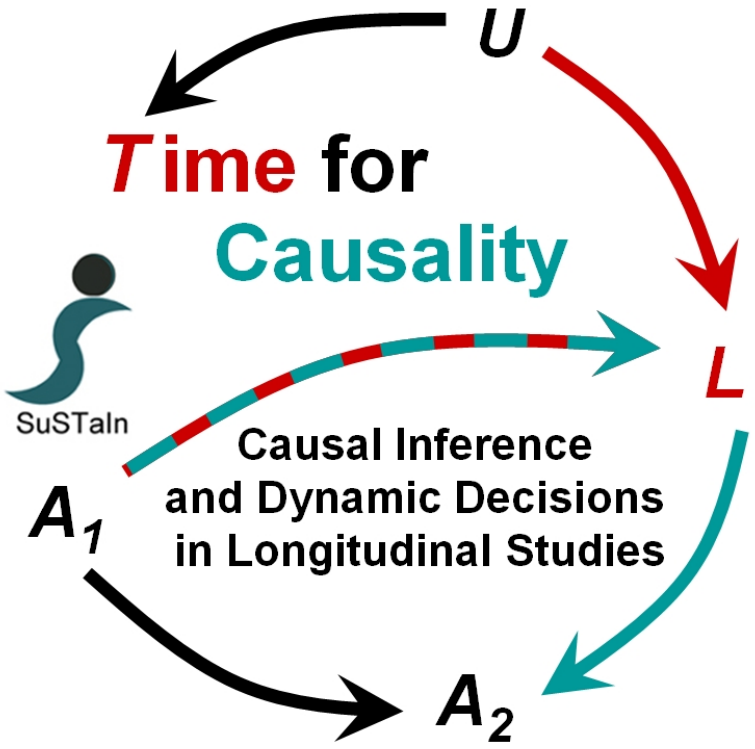

Book of Abstracts

10–13 April 2012

Bristol, UK

Talks

Dynamic causal models

Odd O. Aalen

University of Oslo, Norway

o.o.aalen@medisin.uio.no

Time obviously is an essential aspect of our causal understanding. Causality unfolds in time and cause has to precede effect. This is closely related to the kind of mechanistic understanding that is continuously sought in biology, medicine and other fields. Nevertheless, time dynamics play a very limited role in much causal inference.

One relevant method is Granger causality, the value of which is usually underestimated by statisticians. The closely allied idea of "local independence" due to Tore Schweder (1970) has in recent years been developed much by Vanessa Didelez, becoming an increasingly attractive concept.

I will discuss some aspects of dynamic modelling related to the above approaches.

Schweder T. Composable Markov processes. Journal of Applied Probability 1970; 7:400–410.

Didelez V. Graphical models for composable finite Markov processes. Scandinavian Journal of Statistics 2007; 34:169–185. DOI: 10.1111/j.1467-9469.2006.00528.x

Time to consider time, and time to predict?

Elja Arjas

University of Helsinki and National Institute for Health and Welfare (THL), Finland

elja.arjas@helsinki.fi

Time has a key role in all attempts to arrive even at a most elementary mechanistic understanding of causal phenomena. Where such a goal appears too ambitious for empirical research, it is naturally replaced by an attempt to predict the value of a future observable based on information obtained in the past. Nevertheless, in the current practices of causal modeling and statistical inference, aspects relating to time are mostly left only implicit, or even ignored. In this talk I consider three causal case studies, each supported with real data, where an explicit consideration of time is essential. In each case, existing substantive knowledge from the problem area is used in setting up a suitable stochastic model, and using non-parametric specification of functions if such prior knowledge seems vague. A combination of Bayesian inferential methods and an application of MCMC sampling in the numerical work are then shown to lead to immediately useful and understandable results, formulated in terms of predictive distributions. The approach adopted here differs in several respects from those currently considered 'standard' in causal modeling and inference. The talk ends with some conclusions of a general nature.

References:

E. Arjas & L. Liu: Assessing the losses caused by an industrial intervention: a hierarchical Bayesian approach. Applied Statistics 44 (1995): 357 - 368.

A. Andreev & E. Arjas: Acute middle ear infection in small children: a Bayesian analysis using multiple time scales. Lifetime Data Analysis 4 (1998): 121-137.

E. Arjas & A. Andreev: Predictive inference, causal reasoning, and model assessment in nonparametric Bayesian analysis: a case study. Lifetime Data Analysis 6 (2000):187-205.

E. Arjas & O. Saarela: Optimal dynamic regimes: presenting a case for predictive inference. The International Journal of Biostatistics: 6 (2010), Article 10.

On Pairwise Granger-Causality Analysis of Selected Economic Indicators in Nigeria

Olawale Awe

Obafemi Awolowo University, Nigeria

olawaleawe@yahoo.co.uk

The goal of most empirical studies in econometrics and other social sciences is to determine whether a change in one variable causes a change in another. Granger-causality approach is more popular in such non-experimental subjects which involves some dynamic time series methodology. It is premised upon the fact that the cause occurs before the effect and it must contain some unique information that is not available and helps to predict the effect.

In this paper, we consider the pairwise Granger-Causality Tests involving (seven) major economic indicators in the Nigerian Economy by formalizing Granger's approach to causality suitable for multivariate/longitudinal time series as well as the corresponding concept of intervention. The basic concepts of probability theory are used to formalize the sense in which a change in one variable causes a change in the other variable in the pairwise tests.

In all the analyses considered, we find that the Granger-Causality structure of a multivariate time series depends on the components chosen to make up the series.

Doubly robust estimation of optimal dynamic treatment regimes using regret-regression

Jessica Barrett

MRC Biostatistics Unit, Cambridge, UK

Jessica.Barrett@mrc-bsu.cam.ac.uk

Murphy [1] has proposed a method for estimating optimal dynamic treatment regimes through the use of regret functions. The general scenario is assumed to be an alternating sequence of measurements of a state variable and treatment decisions, followed by a final response which we aim to maximise. The regret function at a given timepoint is the expected loss in the final outcome incurred by a given treatment decision compared to an optimal treatment decision at that timepoint, assuming an optimal treatment regime is to be followed at all future times, and conditional on previous state and treatment history. The optimal dynamic treatment regime is therefore by definition the treatment regime for which all regret functions are equal to zero. For estimation of regret function models Murphy proposes a method which relies on correct specification of the probability of treatment allocation, conditional on state and treatment history. More recently Henderson et al. [2] and Almirall et al. [3] have proposed a new method for estimation of optimal dynamic treatment regimes in which models must be specified both for the regret functions and for the distribution of the state variables. A model for the probability of treatment allocation is, however, not required. We consider a doubly robust version of Henderson et al.'s regret-regression model which is robust to misspecification of the state variable distribution when the probability of treatment allocation is correctly specified.

[1] Murphy, S. A. (2003). Optimal dynamic treatment regimes (with discussion). Journal of the Royal Statistical Society, Series B 65, 331-366.

[2] Henderson, R., Ansell, P., and Alshibani, D. (2010). Regret-regression for optimal dynamic treatment regimes. Biometrics 66, 1192-1201.

[3] Almirall, D., Have, T. T. , and Murphy, S. A. (2010). Structural nested mean models for assessing time-varying effect moderation. Biometrics 66, 131-139.

An exploration of the dynamic longitudinal relationship between mental health and alcohol consumption

Steven Bell

Department of Epidemiology and Public Health, University College London, UK

steven.bell.09@ucl.ac.uk

Objective: To examine the longitudinal relationship between the frequency of heavy drinking days and mental health.

Design: Data from waves 1 (1997) to 6 (2007) of the Birmingham Untreated Heavy Drinkers project were used.

Participants: 471 respondents (74% male) aged 25-55 years at baseline who were drinking a minimum of 50/35 UK units of alcohol for men/women on a weekly basis and were not seeking treatment for their alcohol use (nor had sought treatment in the previous 10 years) upon recruitment into the study.

Variables: Heavy drinking days were defined as consuming 10/7+ UK units of alcohol in a single day for men/women. Mental health was assessed using the SF-36 health survey tool.

Methods: Dynamic longitudinal structural equation models were used to test hypotheses of dominant underlying processes (frequency of heavy drinking days leading to changes in mental health scores; and, mental health scores leading to changes in the frequency of heavy drinking days) and a reciprocal relationship (both mental health scores and the frequency of heavy drinking days influencing changes in each other).

Results: A model whereby mental health scores were predictors of change in the frequency of heavy drinking days was of best fit (no significant effects were observed for the model where the frequency of heavy drinking days was a predictor of change in mental health scores) after adjustment for age, sex, socioeconomic status, marital status, educational level, labour market status, physical functioning, BMI, smoking status, illicit drug use, weekly number of units consumed and symptoms of alcohol dependence.

Conclusion: Mental health is the stronger underlying process in the relationship between mental health and frequency of heavy drinking days.

Stochastic kinetic processes: biochemical mechanisms, information and network design

Clive Bowsher

University of Bristol, UK

C.Bowsher@bristol.ac.uk

Molecular mechanisms are of widespread importance for understanding causality in the biosciences. A stochastic kinetic model (SKM) is a mechanistic description of the dynamics of the molecular networks inside living cells. We discuss SKMs, their graphical representation and conditional independence structure. For SKMs with dynamic input processes, we consider how to control the output by choosing the network design or the input process appropriately. We introduce a new method based on conditional moments and orthogonal noise components, and provide formal connections with mutual information and channel capacity. The optimal design of a synthetic gene network whose output faithfully tracks the input dynamics illustrates the techniques.

Can we attribute increased autumn asthma consultations in children to the new school year?

Mike Campbell and Steven Julious

ScHARR, University of Sheffield, UK

m.j.campbell@sheffield.ac.uk

General practice consultations by children for asthma increase at the start of a new school year. We used data from the General Practice Research DataBase (GPRD) to investigate reasons for this.

We employed contemporary controls of children without asthma, to control for other reasons for consultation. Children from Scotland start school on average two weeks before children from England. We used this natural experiment to assess the extent to which the start of a new school year can be said to cause the increase, compared to other factors present in autumn.

A formal treatment of sequential ignorability

Panayiota Constantinou and Philip Dawid

Statistical Laboratory, University of Cambridge, UK

A.P.Dawid@statslab.cam.ac.uk

We consider sequential decision problems, involving observable variables L_i, unobservable variables U_i, and action variables A_i, that exhibit *extended stability* — which would support G-computation if the U_i were also observed. Under certain specific additional conditions we can infer *simple stability* (or *sequential ignorability*), so supporting G-computation based only on the L_i. One such specific condition is "sequential randomization", where the U_i essentially behave as random noise in generating the actions. Another is "sequential irrelevance", where the U_i do not affect future L_i. However in this latter case, to deduce sequential ignorability in full generality requires additional positivity conditions. We show that these positivity conditions are automatically satisfied when all variables are discrete, but not in the continuous case.

Applications of Dynamic Models in Monitoring and Fault Detection

Roger Dixon

Loughborough University, UK

r.dixon@lboro.ac.uk

The aim of the talk will be to discuss how dynamic models can be used as the basis of algorithms that enable both fault detection and condition/health monitoring of critical components in engineering systems. The concepts of fault detection and monitoring will be introduced, as will the models and modelling approaches. Practical real-world examples will be used to illustrate the concepts: (e.g. based on research applications in areas such as: wind-turbines, electro-mechanical actuators, nuclear fusion, rail-vehicle monitoring, aircraft fuel systems).

Learning for exploration-exploitation in reinforcement learning: The dusk of the small formulas' reign

Damien Ernst

University of Liège, Belgium

dernst@ulg.ac.be

We propose a learning approach to pre-compute K-armed bandit playing policies by exploiting prior information describing the class of problems targeted by the player. Our algorithm first samples a set of K-armed bandit problems from the given prior, and then chooses in a space of candidate policies one that gives the best average performances over these problems. The candidate policies use an index for ranking the arms and pick at each play the arm with the highest index; the index for each arm is computed in the form of a linear combination of features describing the history of plays (e.g., number of draws, average reward, variance of rewards and higher order moments), and an estimation of distribution algorithm is used to determine its optimal parameters in the form of feature weights. We carry out simulations in the case where the prior assumes a fixed number of Bernoulli arms, a fixed horizon, and uniformly distributed parameters of the Bernoulli arms. These simulations show that learned strategies perform very well with respect to several other strategies previously proposed in the literature (UCB1, UCB2, UCB-V, KL-UCB and e_n-GREEDY); they also highlight the robustness of these strategies with respect to wrong prior information.

A sequential Cox approach for estimating treatment effects in the presence of time-dependent confounding

Jon Michael Gran

University of Oslo, Norway

j.m.gran@medisin.uio.no

A typical problem when estimating the effect of treatment for HIV using observational data is time-dependent confounding. Such confounding can be present when a covariate, affected by past exposure, is both a predictor of the future exposure and the outcome. One example is the CD4 cell count, being a marker for disease progression for HIV patients, but also a marker for treatment initiation, and influenced by treatment. Fitting a marginal structural model (MSM) using inverse probability weights has become the standard way to give appropriate adjustment for this type of confounding.

Here we present another approach for estimating similar treatment effects, selecting subsets of the observational data to mimic randomized trials. Each mimicked trial is constructed based on patients starting treatment in a certain time interval, comparing them with the patients not yet on treatment.

An overall effect estimate, given all possible times of treatment start, can be found using composite likelihood inference. The method avoids the use of inverse probability of treatment weights, which can be unstable in certain situations. The parameter being estimated in the overall analysis is not identical to the one of a MSM, as it is conditioned on covariate values at the start of each mimicked trial. This opens for some additional possibilities. The analysis can be performed using a stratified Cox analysis on the joint dataset of all the mimicked trials, where each trial is treated as one stratum. We have applied the method to data from the Swiss HIV cohort study.

Causal network identification in multivariate dynamic systems using genetic programming

Manoj Kandpal

National University of Singapore

manoj@nus.edu.sg

Current challenges in determining the cause-effect relationships in various systems are the major source of inspiration for this work on causality network construction. In this work we are focusing on biological systems, however, the proposed methodology can be applied to many fields for predicting the causal outcome of the particular process.

With the involvement of improved biological technologies, a huge amount of experimental data is available in different fields such as basic sciences, pharmaceuticals, health etc. Various statistical tools like PCA, PLS etc have already being been used to find out the relationship among such different biological entities and in the last few years, the focus has been shifting towards representing these biological genes/metabolites/proteins associations in the form of networks. Many chem-bio systems comprise of a cascade of various interrelated reactions/metabolites such that the occurrence of one event depends on the occurrence of a prior event or a set of events of the same or different nature, e.g. several genes can be expressed together and can cause the expression of other genes which in turn can induce/repress the expression of other sets of genes/proteins. These causal relationships can often be represented in form of multivariate ARX models. However, this causality structure expressed in the form of nodes and branches of a network can become more and more complex with the involvement of more biological entities. While the analysis of static data from biological systems has received much attention, relatively less work has been reported on the analysis of multivariate dynamic biological data. This work takes care of finding out the causalities in such multivariate dynamic data sets using a genetic programming (GP) based variable interaction methodology (GPBVIM).

The methodology developed in this work includes fitting data of different inputs and outputs using a Genetic Modeling System (GeMS). This method can identify non-linear causal relationships as well. In the methodology, possible causal relationship models are generated using GP, which are subsequently validated (for parameter confidence level) using parameter estimation tools. We have tested the proposed methodology on various simulated and experimental datasets. These datasets have been obtained by either simulating the causality based models available in literature or have been kindly provided by the authors on request. We have also tested GPBVIM on real systems; for analysis of

HeLA cell cycle gene expression data, for modeling the spread of an epidemic in 343 US county and for determining causal interactions among various protein dynamics signals. A comparison of the results obtained has indicated that the GPBVIM outperforms other methodologies like the Granger Causality and Bayesian Networks, for it can even give the quantitative measure of the interactions. Qualitatively, these results are in accordance with the earlier published studies on explaining the relationship between the predictors and the response variable. In addition, the developed network also provides a quantitative measure of how a specific predictor variable influences one another and in turn, the response variable.

Using Marginal Structural Models for Nested Counterfactuals to Assess Mediation Through a Longitudinal Measured Mediator

Theis Lange

University of Copenhagen, Denmark

thlan@sund.ku.dk

A recent paper by Lange, Vansteelandt, and Bekaert (1) established a simple unified approach for quantifying mediation through the use of Marginal Structural Models (MSMs) for nested counterfactuals. In this paper we show how the technique can be applied when the mediator has been measured repeatedly/longitudinal through follow-up. The approach is used to quantify how much of the effect of poor sleep on mortality is mediated through obesity based on data from the French GAZEL cohort study. This part of the paper augments the analysis presented in Rod et al. (2).

My talk based on this work would naturally include both a short presentation of the novel results in Lange et al. (2011) and how to use the techniques when the mediator has been measured repeatedly.

References:

1. T. Lange, S. Vansteelandt, and M. Bekaert, A simple unified procedure for assessing mediation by marginal structural models?, forthcoming in American Journal of Epidemiology, (2011).

2. N.H. Rod, J. Vahtera, H. Westerlund, M. Kivimaki, M. Zins, M. Goldberg, and T. Lange, Sleep Disturbances and Cause-Specific Mortality: Results From the GAZEL Cohort Study, American Journal of Epidemiology, Vol. 173, p. 300-309, (2011).

Q-learning for estimating optimal dynamic treatment rules from observational data

Erica E. M. Moodie

McGill University, Canada

erica.moodie@mcgill.ca

Q-learning, a popular reinforcement learning approach, has recently been applied to estimate dynamic treatment regimes. While, in principle, Q-learning can be used for both randomized and observational data, literature thus far has focused exclusively on the randomized treatment setting. We extend an existing implementation of Q-learning to incorporate confounding covariates. We provide results of an extensive simulation study to compare different adjustment methods and show how some forms of model mis-specification may be diagnosed. Methods are illustrated using the PROBIT data to study the effect of breastfeeding on IQ.

This is joint work with Bibhas Chakraborty (Columbia University).

Confidence Intervals, Q-Learning and Dynamic Treatment Regimes

Susan A. Murphy

University of Michigan, USA

samurphy@umich.edu

Dynamic treatment regimes (or treatment policies) are used to operationalize multi-stage decision making in the medical field. Common approaches to constructing the dynamic treatment regimes from data, such as Q-Learning, employ non-smooth functionals of the data. The non-smoothness leads to non-regular asymptotics under certain generative models. Methods that ignore the non-regularity have poor performance in small samples. In this talk, we propose a bootstrap based method for constructing asymptotically valid confidence sets. This method is adaptive in the sense that it provides exact coverage when the true underlying generative model leads to regular asymptotics and is conservative otherwise. Empirical studies show that the amount of conservatism is small.

Estimation and Extrapolation of Optimal Treatment

James Robins[1], Liliana Orellana[2], Andrea Rotnitzky[1]

[1]*Harvard School of Public Health, USA*
[2]*University of Buenos Aires, Argentina*

robins@hsph.harvard.edu andrea@hsph.harvard.edu

We will discuss methods for using the data obtained from an observational database in one health care system to determine the optimal treatment regime for biologically similiar subjects in a second health care system when, for cultural, logistical, and financial reasons, the two health care systems differ (and will continue to differ) in the frequency of, and reasons for, both laboratory tests and physician visits. We also describe methods for estimating the optimal timing of expensive and/or painful diagnostic or prognostic tests. Diagnostic or prognostic tests are only useful in so far as they help a physician to determine the optimal dosing strategy, by providing information on both the current health state and the prognosis of a patient because, in contrast to drug therapies, these tests have no direct causal effect on disease progression. Our methods explicitly incorporate this no direct effect restriction.

Determination of optimal dynamic treatment regimes from incomplete data

Susanne Rosthøj[1] and Robin Henderson[2]

[1]*University of Copenhagen, Denmark*
[2]*Newcastle University, UK*

S.Rosthoej@biostat.ku.dk

Murphy (2003) introduced a new approach to estimation of optimal dynamic treatment regimes from experimental or observational data with a time varying treatment. An optimal dynamic strategy is a strategy maximising a response measured at the end of the study. Robins (2004) demonstrated that Murphy's model is a special case of his Structural Nested Mean Models (SNMMs). Several applications and discussion of these methods have now been seen in the statistical literature (e.g. Rosthøj et al. (2006), Moodie et al. (2009), Henderson et al. (2010), Almirall et al. (2010)). In general, these methods require that all patients are seen at the same time points during the treatment (e.g. weekly or monthly visits to a clinic). However in practice, it might be difficult to obtain completely balanced data, since some patients might skip some of the visits. We demonstrate that the optimal dynamic treatment regime as defined by Murphy and Robins is not optimal when the visit pattern is incomplete and discuss how to define and estimate optimal treatment strategies in this setting. It turns out that these models are not identifiable, i.e. that we cannot distinguish between the optimal dynamic strategy for complete visits and the optimal dynamic strategy for incomplete visits.

References:

Murphy (2003). Optimal dynamic treatment regimes. JRSS B, 65, 331-355.

Robins (2004). Optimal structural nested models for optimal sequential decisions. In Proceedings of the Second Seattle Symposium on Biostatistics, 189-326.

Moodie, Platt, Kramer (2009). Estimating response-maximized decision rules with applications to breastfeeding. JASA 104,155-165.

Rosthøj, Fullwood, Henderson, Stewart (2006). Estimation of optimal dynamic anticoagulation regimes from observational data: a regret-based approach. Stats in Med 25, 4197-4215.

Henderson, Ansell, Alshibani (2010). Regret-Regression for Optimal Dynamic Treatment Regimes. Biometrics 66, 1192-1201.

Almirall, Ten Have, Murphy (2010). Structural Nested Mean Models for Assessing Time-Varying Effect Moderation. Biometrics 66, 131-139.

Martingale measures in causal inference

Kjetil Røysland

Univeristy of Oslo and Cancer Registry of Norway

kjetil.roysland@medisin.uio.no

We show that one can do causal inference for several continuous time scenarios using some powerful tools from stochastic calculus. A probability distribution that governs the frequency of observations in the counterfactual scenario can be characterized in terms of a so called martingale problem. The counterfactual and observational probability distributions may be related through a likelihood ratio given by a stochastic differential equation. If this is possible, one can do inference for counterfactual scenarios based on the original observations, re-weighted according to the likelihood ratio. The crucial assumption of positivity in marginal structural modelling translates, in this sense, into absolute continuity of a counterfactual probability distribution with respect to the observational distribution. The stability of the weights can analogously be characterized as distance between these distributions.

The role of exchangeability in causal inference

Olli Saarela

McGill University, Canada

olli.saarela@mcgill.ca

The notion of exchangeability has been recognized in the causal inference literature in various interpretations (e.g. Greenland and Robins, 2009), but only rarely in the original Bayesian meaning as a symmetry property between individual units in statistical inference. A connection between exchangeability and causal reasoning was first suggested by Lindley and Novick (1981); here we relate the Bayesian notion of exchangeability to alternative confounding related conditions, commonly stated in terms of potential outcome variables, with the aim to provide an update of this classic account with the hindsight of the vast developments that have taken place in causal inference theory and methodology since. Although the usefulness of the concept of exchangeability was disputed by Pearl (2009, p. 177–180), in Bayesian inference the model building and estimation tasks are inseparable, and thus we will argue that between unit exchangeability is a relevant concept in Bayesian causal inference, not least because through de Finetti's representation theorem it links assumptions about the underlying causal mechanism directly into quantities (parameters) relevant to statistical estimation of causal contrasts of interest, which in turn can be defined in terms of limits of posterior predictive expectations for further exchangeable units. Even though in longitudinal settings the well known g-computation formula can be derived through the assumption of sequential randomization or even weaker conditions (e.g. Dawid and Didelez, 2010), in practice, due to the curse of dimensionality, stronger conditions are required for the actual estimation task of the probability distributions involved, in order to avoid the phenomenon known as the null paradox. We interpret the null paradox in terms of exchangeability assumption made on a too coarse scale. The message is that in applying the g-computation formula or its Bayesian counterparts in longitudinal settings it may be more feasible to try to avoid the paradox by latent variable modeling of individual level frailty-type factors instead of pursuing completely nonparametric model specifications.

References:

Dawid, A. P. and Didelez, V. (2010). Identifying the consequences of dynamic treatment strategies: a decision-theoretic overview. Statistical Surveys, 4:184-231.

Greenland, S. and Robins, J. M. (2009). Identifiability, exchangeability, and epidemiological confounding revisited. Epidemiologic Perspectives & Innovations, 6. doi:10.1186/1742-5573-6-4.

Lindley, D. V. and Novick, M. R. (1981). The role of exchangeability in inference. The Annals of Statistics, 9:45-58.

Pearl, J. (2009). Causality, Second Edition. Cambridge University Press, Cambridge.

Targeted Maximum Likelihood Methods for Longitudinal Data

Mireille Schnitzer

McGill University, Canada

mireille.schnitzer@mail.mcgill.ca

In longitudinal problems where time-dependent confounding is present, several causal inference methods have been developed that correctly estimate treatment-specific means and marginal structural model summaries of time-dependent treatment effects. Targeted Maximum Likelihood Estimation (van der Laan and Rubin, 2006) is a class of methods in causal inference that produces efficient, doubly robust estimators. These methods also benefit from the stability and respect for global bounds arising from plug-in estimation. In this talk, we demonstrate the construction of a recently proposed flexible longitudinal targeted method (van der Laan and Gruber, 2011) and its natural use for censored data. We apply this method and several alternatives to the PROmotion of Breastfeeding Intervention Trial (PROBIT) where we analyze the effect of breastfeeding on gastrointestinal infections in infants.
(Joint work with Mark van der Laan, Erica Moodie and Robert Platt)

Parametric g-estimation of treatment effects in randomized controlled trials

Tibor Schuster

Technical University of Munich, Germany

Tibor.Schuster@tum.de

In contradiction to an often presumed individual-level therapy impact, estimation of treatment effects in randomized controlled trials is usually based on associational rather than on causal inference approaches. We introduce a novel inference concept based on a g-estimation procedure linking a subject's observed survival time to an assumed latent counterfactual survival time which would have occurred if innovative treatment had always been withheld. By conditioning on the rank-preserving assumption, the suggested procedure allows for estimation of a treatment effect parameter (or parameter vector) which is connected by a real-valued function to the parameters of the assumed and empirically verifiable distribution of survival times in the reference treatment arm. In this term, causal inference for treatment effects on survival is feasible based on the likelihood function and its corresponding test statistics.

Causal inference for dynamic treatment regimens: how analyses of observational data changed international guidelines on when to start antiretroviral therapy

Jonathan Sterne

University of Bristol, UK

Jonathan.Sterne@bristol.ac.uk

At the beginning of 2010 the US Department of Health and Human Services and the WHO issued revised guidelines recommending that HIV-positive people start antiretroviral therapy (ART) at higher CD4 counts. These changes were made after two large HIV cohort collaborations published analyses using novel statistical methods to deal with time-dependent confounding and lead time bias. The ART Cohort Collaboration (Lancet 2009; 373: 1352-1363) concluded that rates of AIDS and death are reduced when ART is initiated at above 350 cells/mm^3, but found little evidence of benefit at higher CD4 thresholds. In contrast, NA-ACCORD (N Engl J Med 2009; 360: 1815-26) found strong evidence of reduced mortality when ART is initiated at above 500 cells/mm^3. The use of dynamic marginal structural models in that paper was subsequently criticised, and a recent publication using the same approach (HIV-CAUSAL collaboration, Ann Intern Med. 2011; 154: 509-515) did not replicate these findings. I will describe the methods used in these papers, and discuss possible reasons for differences in their results.

Control System Design: Feedback and Sequential Decision Making

James Taylor

Lancaster University, UK

c.taylor@lancaster.ac.uk

The presentation will start with an overview of some standard control engineering terms, such as optimal control and feedback, using a nuclear decommissioning robot as a worked example. The aim is to find concepts and modelling tools that could be analogous/useful in the context of the workshop. With this goal in mind, the speaker will also show how control engineering techniques have fed into other areas of science, such as recent research into climate change modelling and the setting of optimal CO2 emission targets for policymakers, a form of sequential decision making.

Adjusting for time-varying confounders in survival analysis using structural nested cumulative failure time models

Stijn Vansteelandt

Ghent University, Belgium

stijn.vansteelandt@ugent.be

Time-varying confounding forms a pervasive problem in observational studies that attempt to assess the total effect of a time-varying exposure on a survival outcome. It is tempting to see prosperity in the use of routine survival models with time-varying covariates, but unfortunately the resulting estimates of the exposure effect are prone to bias when, as often, time-varying confounders are themselves affected by earlier exposures. This is because standard regression adjustment for such confounders eliminates indirect effects of early exposures that are mediated via those covariates, and in addition, may induce a so-called collider-stratification bias. Martinussen et al. (2011) demonstrated how a valid adjustment for time-varying confounding is attainable when effects are parameterized on the additive hazard scale. Because they focused on the special case of 2 exposures, the first of which is dichotomous and randomly assigned, we here extend their results to general time-varying exposures. The extension explicates a close link with inference under structural nested cumulative failure time models (Young et al., 2010; Picciotto et al., 2011), but yields a different class of estimators that are deemed more efficient. Relative to G-estimation for structural nested accelerated failure time models, the attraction of the proposed approach is that it naturally accommodates non-informative censoring without requiring an artificial recensoring procedure to maintain unbiased estimating equations. Relative to inverse probability weighting for marginal structural survival models, the advantages are that it tends to yield more stable inferences by not involving inverse weights, and that it can incorporate effect modification by time-varying covariates.

This is based on joint work with Torben Martinussen and Eric Tchetgen Tchetgen.

Posters

Conditional Independence in the Decision-Theoretic framework

Panayiota Constantinou

Statistical Laboratory, University of Cambridge, UK

pc393@cam.ac.uk

The calculus of conditional independence arises to be the basic element which enables us to express causal concepts in the Decision-Theoretic framework of Statistical Causality. In this framework, we differentiate between observational and interventional regimes using a non-stochastic variable (to index the regimes) and use the notion of conditional independence to state conditions under which we can relate them. Here, we rigorously extend the language and calculus of conditional independence to incorporate stochastic and non-stochastic variables and use the properties that accrue to identify causal quantities.

Simulation of Time-Dependent Confounding in Marginal Models

William Havercroft and Vanessa Didelez

University of Bristol, UK

w.havercroft@bristol.ac.uk

In longitudinal studies where the causal effect of treatment on patient outcome is focal, the problem of time-dependent confounding, in particular where past treatment has an effect on present covariates, is an important issue. In cases where the key problem is ascertaining the best static treatment regime, methods to account for such confounding have been established for several years. The issue of finding and evaluating an optimal dynamic treatment regime based on the patient's observed history is more recent, with new methods being proposed that still require investigation into their practical usefulness.

In both cases it is edifying to test such methods on simulated data with known values of interest, in order to compare properties such as the accuracy or efficiency of different methods against one another. To this end, we examine the problem of designing a sampling algorithm which produces the confounded data structure of interest, but is based on the model of a 'true' (marginal) causal treatment effect. In particular, we build upon algorithms proposed in a few recent papers, to create a new, flexible and more general algorithm for simulating time-dependent confounding of longitudinal data, in which survival time is the outcome of interest.

Monitoring the illegal ivory trade and assessing the impact of CITES policy on the trade

Fiona Underwood

University of Reading, UK

f.m.underwood@reading.ac.uk

The illegal ivory trade is monitored by the Elephant Trade Information System (ETIS). ETIS was mandated by the UN body CITES (Convention for International Trade in Endangered Species) which makes international policy decisions regarding the trade in ivory. Specifically commercial trade in ivory is prohibited from all Asian and all but four African countries. Only three conditional one-off sales of raw ivory to Japan and China have been allowed by CITES in the last two decades. These sales are controversial due to differing opinions about whether these sales drive the illegal trade.

Two of the aims of ETIS were that it should (1) establish trends in the illegal ivory trade through time and (2) determine whether or not such trends are related to CITES decisions. ETIS collects reports of illegal ivory seizures and currently holds partial information on over 16,000 seizures from the last 20 years from over 80 countries. Simple data summaries and traditional statistical methods cannot provide a true picture of the trade: the data are opportunistically collected and are an intervention in the process being measured; countries differ in the proportion of shipments they seize and report - these are unknown and vary through time, and; countries differ in the flow of ivory shipments and their position in the trade chain, which may also change over time. Furthermore the impact of CITES decisions must be viewed against the backdrop of all other potential drivers and interventions in the trade many of which may be more dominant than the impact of CITES decisions.

Statistical modelling to describe relative trends in the illegal ivory trade have been developed; but there has been less progress on examining the causal effect of CITES decisions. One issue is that In many cases proxy variables based on global datasets on governance and development and knowledge of local ivory markets are required to capture the potential drivers of this dynamic process.

In this presentation I describe the problem in detail and first steps to addressing the issue of causality. The aim of the presentation is to provide a starting point for discussion and advice on this topic.

Quadratic Programming for Weight Adjustment in Longitudinal Studies of Treatment Effects

Jose Zubizarreta

University of Pennsylvania, USA

josezubi@wharton.upenn.edu

Inverse probability weighting is a widely used method in longitudinal studies of treatment effects. Typically the researcher estimates the probability of treatment using a logistic model and then inverts the estimated probabilities. However, the resulting weights tend to have high variance and thus to yield estimates with high standard errors. In view of this problem, common practice suggests trimming the weights, but this is likely to introduce bias. Also, where to trim is an arbitrary decision, and little attention is paid to diagnosing the resulting weights. Furthermore, with this approach there is no systematic way of integrating subject matter knowledge of the problem at hand, and the efforts are placed in estimating the "true" probability model instead of making the ignorability assumptions plausible. In this study we propose a new method for weight adjustment that overcomes these issues. This method consists of solving a sequence of quadratic programming problems to obtain weights (i) with minimum variance, that (ii) finely adjust the distributions of the observed covariates. The quadratic programming problems optimize the bias-variance tradeoff, while making easy to integrate subject matter knowledge of the problem at hand and make the ignorability assumptions more plausible. We implement this method computationally. In a number of examples we show that the proposed method yields weights that adjust better and that are less variable than those obtained with common methods.