Arthur Gretton

Title: Causal modelling with kernels: treatment effects, counterfactuals, mediation, and proxies

Abstract: A fundamental causal modelling task is to predict the effect of an intervention (or treatment) D=d on outcome Y in the presence of observed covariates X. We can obtain an average treatment effect (ATE: the expected outcome of intervening with d), by marginalising our estimate \gamma(X,D) of the conditional mean E(Y|X,D) over P(X). More complex causal questions require more delicate treatment. For instance, the average treatment on the treated (ATT) addresses a counterfactual: what is the outcome of treatment d' on a subpopulation that received treatment d? In this case, we must marginalise \gamma over the conditional distribution P(X\d), which becomes challenging for continuous multivariate d. Many additional causal questions require us to marginalise over conditional distributions, including Conditional ATE, mediation analysis, dynamic treatment effects, and correction for unobserved confounders using proxy variables.

We address these questions in the nonparametric setting using kernel methods, which may be applied for very general treatments D and covariates X (continuous multivariate, strings, groups, ...). We learn \gamma by kernel ridge regression, and perform marginalization over conditional distributions using the kernel conditional mean embedding, which may be thought of as a generalization of two-stage least-squares. We provide strong statistical guarantees under general smoothness assumption, and a straightforward and robust implementation (a few lines of code). The method is demonstrated by addressing causal modelling questions arising from the US Job Corps program for Disadvantaged Youth.

Bio: Arthur Gretton is a Professor with the Gatsby Computational Neuroscience Unit, and director of the Centre for Computational Statistics and Machine Learning (CSML) at UCL. He received degrees in Physics and Systems Engineering from the Australian National University, and a PhD with Microsoft Research and the Signal Processing and Communications Laboratory at the University of Cambridge. He previously worked at the MPI for Biological Cybernetics, and at the Machine Learning Department, Carnegie Mellon University.
Arthur's recent research interests in machine learning include the design and training of generative models, both implicit (e.g. GANs) and explicit (exponential family and energy-based models), causal modeling, and nonparametric hypothesis testing.

He has been an associate editor at IEEE Transactions on Pattern Analysis and Machine Intelligence from 2009 to 2013, an Action Editor for JMLR since April 2013, a Senior Area Chair for NeurIPS in 2018 and 2021, and a member of Royal Statistical Society Research Section Committee since January 2020. Arthur was a program chair for AISTATS in 2016, a tutorials chair for ICML 2018, a workshops chair for ICML 2019, a

program chair for the Dali workshop in 2019, and an organsier of the Machine Learning Summer School 2019 in London.

Photo: