Name:	Yanbo Xu
Affiliation:	Johns Hopkins University
Email:	yxu68@jhu.edu
Presentation type	Poster

A Non-parametric Bayesian Approach for Estimating Treatment-Response Curves from Sparse Time Series

Yanbo Xu, MS[1], Yanxun Xu, PhD[1], Suchi Saria, PhD[1]
[1]Johns Hopkins University, Baltimore, MD

Abstract

We study the problem of estimating the continuous response over time of actions from observational time series—a retrospective dataset where the policy by which the data are generated are unknown to the learner. We are motivated by applications where response varies by individuals and therefore, estimating responses at the individual-level are valuable for personalizing decision-making. We refer to this as the problem of estimating individualized treatment response (ITR) curves. In statistics, G-computation[1] has been commonly used for estimating treatment responses from observational data containing sequential treatment assignments. However, past studies have focused predominantly on obtaining point in time estimates at the population level. We leverage G-computation and develop a novel method based on Bayesian nonparametrics (BNP) that can flexibly model functional data and provide posterior inference over the treatment response curves both at the individual and population level. On a challenging dataset containing time series from patients admitted to a hospital, we estimate treatment responses for 2 different treatments used in managing kidney function and show that the resulting fits are more accurate than alternative approaches. Accurate methods for obtaining ITRs from observational data can dramatically accelerate the pace at which personalized treatment plans become possible.

Introduction

Accurate models of actions and their effects on the state of the agent are critical for decision-making. Learning of action-effect models is most straightforward from data where the learner can control the choice of actions and observe their responses. But, such data are not always possible to acquire. Alternatively, retrospective data may be available that contain time series generated from observing other agents act. Estimating action-effect models from observational data—data where the learner cannot control the actions that are prescribed, and the actions may be prescribed by a mechanism that is not known to the learner—are more challenging. We study an instance of this problem: specifically, we consider the problem of estimating the continuous response over time to an action. We are particularly motivated by applications in medicine where accurate action-effect models for estimating treatment effects can be used for personalizing therapy. We propose a new Bayesian nonparametric method for estimating, at the individual-level, the posterior density of the continuous response over time to treatments (or actions more broadly) from observational time-series data. A key practical advantage of using a nonparametric approach is that they often provide better fits to challenging data than can be obtained using parametric model based methods. This is particularly important in our application of estimating treatment response curves for physiologic time series.

Method

As a running example, we will use our application of estimating response curves of two treatments called intermittent hemodialysis (IHD) and continuous renal replacement therapy (CRRT) on creatinine, a measurement of kidney function. Towards this, we use data obtained from the electronic health record from a patient admitted to a

hospital. Our goal is to obtain posterior inference for the treatment response curves at the individual and population levels, and for the outcomes given any sequence of treatments conditioned upon historical data about the individual and the population. We model the outcome using a generalized mixed-effects model combining two parts: 1) the baseline progression over time in the response variable if no treatments were prescribed; 2) the change in response due to all active treatments. Specifically, for the baseline progression, we model a fixed-effect component using linear regression, a random-effect component using a Gaussian process regression and a Gaussian distributed noise term. For the treatment response, we focus on the scenario where treatment choices are discrete and assume that the treatment effects are additive. Further, we assume that the effect of each treatment type lasts at most within window W, and define a parametric function to characterize the individual treatment response (ITR) curves. The deviations from the parametric form are captured by a zero-mean Gaussian process. Finally, to borrow strength across the individual-specific estimates, we generalize the Gaussian Process component for both the baseline progression and the treatment response by using a Dirichlet process mixture of Gaussian processes[2,3]. We use Markov Chain Monte Carlo (MCMC) approach to approximate the posterior inference.

Numerical Analysis

We fit our models on electronic health record data from patients admitted to the Beth Israel Deaconess Medical Center in Boston. The data are publicly available in the MIMIC-II Clinical Database[4]. The creatinine data contains time series from 123 individuals with average duration of 20.75 days and a total of 6,992 observations. We model 4 treatments including IHD and CRRT prescribed at three different levels: < 500 ml/hr, = 500 ml/hr and > 500 ml/hr. We compare ITR's performance to two baselines: 1) the model parameters are drawn from a broad prior distribution but each individual samples it's own set of parameters, which we refer to as pop; 2) the model parameters are drawn from a DP instead of a DPM that allows treatment responses to vary by subgroups but no explicit differences across individuals within a subgroup, which we refer to as sub-pop. We report held-out prediction root mean squared errors (RMSE) averaged within a day for 10 days following the time of prediction. From the left graph in Figure 1, we see that the proposed ITR model improves significantly over the baselines. We show trajectory fits for two different randomly chosen patients in the right two graphs, where the darker shaded background denotes data in the training period and the lighter background is the test period. In Figure 2, we show the distribution over the individual-specific response curves for IHD and CRRT at the three different dose levels. As is clear, there is significant treatment heterogeneity across all treatments.

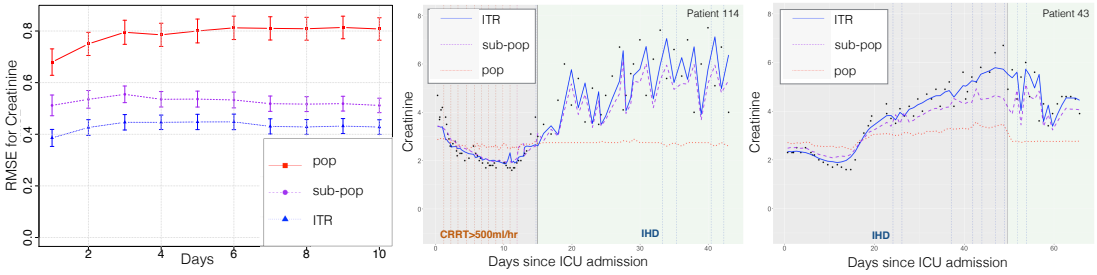

Figure 1. Comparison of model prediction

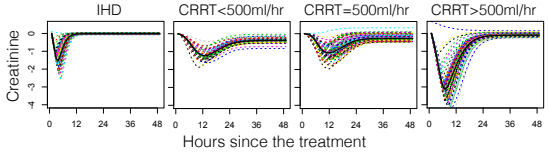

Figure 2. Treatment-response curves estimated by ITR

References

1. James Robins. A new approach to causal inference in mortality studies with a sustained exposure period - application to control of the healthy worker survivor effect. Mathematical Modeling. 1986; 7(9): 1393–1512,
2. Hemant Ishwaran and Lancelot F James. Gibbs sampling methods for stick-breaking priors. Journal of the American Statistical Association. 2001; 96(453).
3. Fernando A Quintana, WO Johnson, E Waetjen, and E Gold. Bayesian nonparametric longitudinal data analysis. Journal of the American Statistical Association. 2015; doi: 10.1080/ 01621459.2015.1076725.
4. Mohammed Saeed, C Lieu, G Raber, and RG Mark. MIMIC II: a massive temporal ICU patient database to support research in intelligent patient monitoring. In Computers in Cardiology, 2002; 641–644.