Joint Point Process Model for Counterfactual Treatment–Outcome Trajectories Under Policy Interventions

Çağlar Hızlı, ST John
Department of Computer Science
Aalto University
{caglar.hizli, ti.john, pekka.marttinen}@aalto.fi

Anne Juuti, Tuure Saarinen, Kirsi Pietiläinen
Helsinki University Hospital and University of Helsinki
{anne.juuti, tuure.saarinen}@hus.fi, kirsi.pietilainen@helsinki.fi

Pekka Marttinen
Department of Computer Science
Aalto University
pekka.marttinen@aalto.fi

Abstract

Policy makers need to predict the progression of an outcome before adopting a new treatment policy, which defines when and how a sequence of treatments affecting the outcome occurs in continuous time. Commonly, algorithms that predict interventional future outcome trajectories take a fixed sequence of future treatments as input. This excludes scenarios where the policy is unknown or a counterfactual analysis is needed. To handle these limitations, we develop a joint model for treatments and outcomes, which allows for the estimation of treatment policies and effects from sequential treatment–outcome data. It can answer interventional and counterfactual queries about interventions on treatment policies, as we show with a realistic semi-synthetic simulation study. This abstract is based on work that is currently under review for AAAI-23.

1 Introduction

What policy should we adopt? In healthcare, for example, we observe patients' physiological markers (*outcomes*) changing over time. We want to affect these outcomes by actions (*treatments*) such as doses of a medicine. Sequences of outcomes and treatments are recorded as a time series. The choice of when to take what action constitutes the *policy*. To improve our policies, we must be able to assess their consequences: What is the effect of a given policy? What will be the effect of a change to a different policy? What would have happened if a patient had followed a different treatment policy? These questions correspond to observational, interventional, and counterfactual queries.

In high-risk domains such as public health and healthcare [Schulam and Saria, 2017, Bica et al., 2020], it is important to quantify risks and expectations accompanying the policy decision, as well as to evaluate the performance of past decisions [Oberst and Sontag, 2019, Tsirtsis and Gomez Rodriguez, 2020, Tsirtsis et al., 2021]. This requires estimating the causal effect of an intervention affecting the treatment policy on the outcome progression using a *causal* model.

2022 NeurIPS Workshop on Gaussian Processes, Spatiotemporal Modeling, and Decision-making Systems.

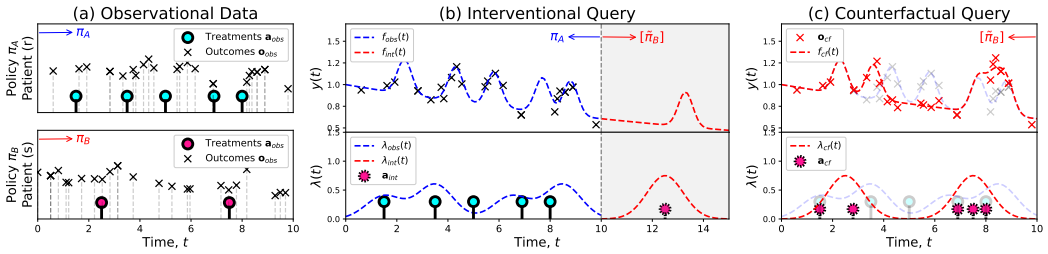

Figure 1: (a) Treatment–outcome data for two patients, following distinct policies π_A and π_B in the observation period $[0, 10]$. (b) The interventional query corresponds to how the outcome trajectory (r) will progress under a different policy π_B after the observation period $[0, 10]$ (shaded area). (c) The counterfactual query corresponds to how the outcome trajectory (r) would have progressed if the policy had been set to $\pi^{(r)} = \pi_B$ instead, in the observation period $[0, 10]$. Notice how the algorithm chooses to keep some of the observed treatments as counterfactual treatments $\mathbf{a_{cf}}$, where the counterfactual intensity is higher than the original observational intensity.

Observed treatment–outcome data are always created by some existing policy; however, the policy is generally not recorded and may be known only implicitly through the observed data. Consequently, this link from past outcomes to future treatments is largely neglected in the sequential treatment–outcome literature, and the causal analysis is generally limited to a fixed sequence of treatment interventions set by hand or generated by a simplistic parametric model [Schulam and Saria, 2017, Lim et al., 2018, Bica et al., 2020, Seedat et al., 2022]. Such models cannot generalize beyond simulations to the analysis of realistic, alternative treatment policies in real-world applications. Also, evaluating treatment policies using counterfactual reasoning is not considered by most of the literature, which focuses on future progression.

To address these limitations, we propose a joint treatment–outcome model. Our model can be learned from observational sequential treatment–outcome data (Figure 1(a)) and can estimate future and counterfactual progression. We show that an intervention on a treatment policy is equivalent to a stochastic intervention on sequential treatments, which we can model with our joint model, and use it to answer interventional (Figure 1(b)) and counterfactual (Figure 1(c)) queries.

2 Problem Definition

Consider an observational data set $\mathcal{D}$,

$$\mathcal{D} = \Big\{ \underbrace{\pi_{[0,T]}}_{\text{policy label}}, \underbrace{\{(t_i, m_i)\}_{i=1}^{N_a}}_{\text{treatments } \mathbf{a}}, \underbrace{\{(t_j, y_j)\}_{j=1}^{N_o}}_{\text{outcomes } \mathbf{o}} \Big\},$$

observed in the period $\mathcal{T} = [0, T]$. For notational simplicity, the data set is defined for a single individual. Our model can be trivially generalized to multiple individuals.

A policy label $\pi_{[0,T]} \in \Pi$ specifies a treatment intensity function $\lambda^*_\pi(t, m)$ in the interval $[0, T]$ that defines when and how a sequence of continuous-time treatments occur. We assume the data set $\mathcal{D}$ contains the policy label $\pi_{[0,T]}$, but its corresponding intensity function $\lambda^*_\pi(t, m)$ is unobserved. A treatment tuple $a_i = (t_i, m_i)$ consists of an arrival time t_i and a treatment mark m_i. An outcome tuple $o_j = (t_j, y_j)$ consists of a measurement time t_j and an outcome value y_j. The history $\mathcal{H}_{\leq t} = \{\pi_{\leq t}, \mathbf{a}_{\leq t}, \mathbf{o}_{\leq t}\}$ contains the information about the past policy $\pi_{\leq t}$, past treatments $\mathbf{a}_{\leq t} = \{(t_i, m_i) : t_i \leq t\}$ and past outcomes $\mathbf{o}_{\leq t} = \{(t_j, y_j) : t_j \leq t\}$.

We observe a continuous-time process $\mathbf{Y}_{\leq T} = \{y(\tau) : \tau \leq \mathcal{T}\}$ as outcome tuples $\mathbf{o}$ measured at times $\mathbf{t_o} = \{t_j\}_{j=1}^{N_o}$. To answer causal queries, we model the potential outcome trajectory $\mathbf{Y}_{>\tilde{\tau}}[\tilde{\pi}_{>\tilde{\tau}}]$, under an intervened policy specified by $\tilde{\pi}_{>\tilde{\tau}}$. When the intervention time $\tilde{\tau}$ is set to the end of the observation period $\tilde{\tau} = T$, we call the estimation task a *policy intervention*, as its computation requires access to the interventional distribution (Figure 1(b)):

$$P(\mathbf{Y}_{>T}[\tilde{\pi}_{>T}] \mid \mathcal{H}_{\leq T}). \tag{1}$$

Also, we can set the intervention time $\tilde{\tau}$ to the start of the observation period $\tilde{\tau} = 0$ and consider a hypothetical scenario under an alternative treatment policy specified by $\tilde{\pi}_{\leq T}$. We call this estimation task a *policy counterfactual*, as its computation requires access to the counterfactual distribution (Figure 1(c)):

$$P(\mathbf{Y}_{\leq T}[\tilde{\pi}_{\leq T}] \mid \mathcal{H}_{\leq T}). \tag{2}$$

Under a set of causal assumptions defined in Appendix C, we show in Appendix D that (i) the potential outcome trajectory $\mathbf{Y}_{>T}[\tilde{\pi}_{>T}]$ is equivalent to the potential outcome trajectory under a sequence of stochastic interventions on treatments and (ii) it is identified using two statistical quantities [Pearl, 2009, Hernán and Robins, 2010]. To estimate these statistical quantities from the observational data, we propose a joint treatment–outcome model.

3 Treatment–Outcome Model

Our joint model is a combination of two conditional intensity functions: (i) treatment intensity: $\lambda^*_\pi(t, m) = \lambda^*_\pi(t)p^*(m \mid t)$ and (ii) outcome intensity: $\lambda^*_o(t, y) = \lambda^*_o(t)p^*(y \mid t)$. We assume the measurement times $\mathbf{t_o}$ of the outcomes are given, which is valid for example when the data are collected through automated patient monitoring in healthcare. Then, the joint distribution for the data set $\mathcal{D}$ can be written in terms of two intensity functions as follows [Daley and Vere-Jones, 2003, Rasmussen, 2011]:

$$p(\mathcal{D}) = \prod_{i=1}^{I} \underbrace{\lambda^*_\pi(t_i)p^*(m_i \mid t_i)}_{\text{Treat. Intensity}} \prod_{t_j \in \mathbf{t_o}} \underbrace{p^*(y_j \mid t_j)}_{\text{Out. Model}} \times \exp(-\Lambda), \tag{3}$$

with the integral term $\Lambda = \int_{\mathcal{T}} \lambda^*_\pi(\tau)d\tau$.

3.1 Treatment Intensity

We model the treatment time intensity $\lambda^*_\pi(\tau)$ using a constant baseline β_0 and three functions with GP priors, $g_b, g^*_a, g^*_o \sim \mathcal{GP}$. The latent-state function g_b models the baseline intensity. The regressive components g^*_a and g^*_o model the dependence on past treatments and outcomes respectively [Liu and Hauskrecht, 2019]. The treatment intensity $\lambda^*_\pi(\tau)$ is defined as follows:

$$\lambda^*_\pi(\tau) = \big(\underbrace{\beta_0}_{\substack{\text{PP}\\ \text{Baseline}}} + \underbrace{g_b(\tau)}_{\substack{\text{NHPP}\\ \text{Baseline}}} + \underbrace{g^*_a(\tau)}_{\substack{\text{Treat.}\\ \text{Effect}}} + \underbrace{g^*_o(\tau)}_{\substack{\text{Out.}\\ \text{Effect}}} \big)^2.$$

The model and kernel definitions are detailed in Appendix F.1-2.

3.2 Outcome Model

We model the outcome trajectory $\mathbf{Y} = \{y(\tau) : \tau \in \mathbb{R}_{\geq 0}\}$ over time τ, combining three independent components: (i) a baseline progression, (ii) treatment effects and (iii) a noise variable [Schulam and Saria, 2017, Xu et al., 2016, Zhang et al., 2020]:

$$y(\tau) = \underbrace{f_b(\tau)}_{\text{Baseline}} + \underbrace{f_a(\tau; \mathbf{a})}_{\text{Treatment Effect}} + \underbrace{\epsilon(\tau)}_{\text{Noise}}. \tag{4}$$

The baseline progression and the treatment effect functions are modeled by GP priors, with an independent Gaussian noise $\epsilon(\tau) \sim \mathcal{N}(0, \sigma^2_\epsilon)$. The model and kernel definitions are detailed in Appendix G.

4 Experiments

We evaluate our model on two causal inference tasks: (i) the policy intervention (Equation 1) and (ii) the policy counterfactual (Equation 2), by setting up a realistic semi-synthetic simulation scenario.

We fit our joint model to a challenging real-world data set on meal–blood glucose dynamics [Zhang et al., 2020, Wyatt et al., 2021] to obtain the ground-truth data generators. The ground-truth models

Table 1: DACC results for two policy interventions $\{[\tilde{\pi}_{>T} = \pi_A], [\tilde{\pi}_{>T} = \pi_B]\}$ over 10 runs. The observed policy is $\pi_{[0,T]} = \pi_A$. DACC closer to 50% is better, as it suggests estimated trajectories are inseparable from ground-truth trajectories.

	$[\tilde{\pi} = \pi_A]$	$[\tilde{\pi} = \pi_B]$
JOINT MOD.	DACC ↓	DACC ↓
OBS-EST	$51.8 \pm 2.8\%$	$86.4 \pm 2.5\%$
INT-EST	$51.8 \pm 2.8\%$	$51.8 \pm 2.8\%$
INT-ORACLE	$50.3 \pm 2.3\%$	$50.3 \pm 2.3\%$

Table 2: DACC results for the policy counterfactual in the observed period $[0, T]$ with the policy intervention $[\tilde{\pi}_{[0,T]} = \pi_B]$ over 10 runs. The observed policy is $\pi_{[0,T]} = \pi_A$. DACC closer to 50% is better, as it suggests estimated trajectories are inseparable from ground-truth trajectories.

	$[\tilde{\pi} = \pi_B]$
JOINT MOD.	DACC ↓
INT-EST	$90.1 \pm 4.1\%$
CF-EST	$60.8 \pm 2.2\%$
CF-ORACLE	$51.8 \pm 2.7\%$

are used to simulate samples from observational, interventional and counterfactual distributions. Simulated patients are divided into two policy groups $\{\pi_A, \pi_B\}$, representing different treatment policies of different hospitals, countries, etc. The details of the simulation study are presented in Appendix J.

We define three joint estimation models: OBS-EST, INT-EST and CF-EST. OBS-EST is trained on the observational data of each individual to generate predictions. INT-EST adjusts predictions of OBS-EST by accounting for the fact that the treatments are generated by the estimated policy for another individual, as a consequence of a policy intervention. CF-EST additionally conditions predictions of INT-EST with the posterior of the individual's noise terms. We denote ground-truth versions of these models as OBS-ORACLE, INT-ORACLE and CF-ORACLE, which represent the performance of the estimated models if infinite training data were available.

To measure how similar predicted trajectories are to samples from the ground-truth distribution, we train discriminators. Ideally, for samples of the same distribution, predicted trajectories should be inseparable from ground-truth trajectories, leading to a 50% discriminator accuracy (DACC).

For the policy intervention task (Table 1), we see that the INT-EST model is able to sample observational and interventional trajectories close to ground-truth distributions when the intervention policy is (i) same as the observed policy $[\tilde{\pi} = \pi_A]$ and (ii) different from the observed policy $[\tilde{\pi} = \pi_B]$, while the OBS-EST model fails in the latter case. For the policy counterfactual task (Table 2), we see the INT-EST model fails to sample counterfactual trajectories close to the ground-truth counterfactual distribution $[\tilde{\pi} = \pi_B]$, as it does not take the individual's noise posterior into account. On the other hand, the CF-EST model is able to sample counterfactual trajectories close to the ground-truth counterfactual distribution.

5 Conclusion

To study what happens if the (possibly implicit) treatment policy of one individual (hospital, country, ...) is or had been adopted by another individual, we proposed a model that jointly considers sequences of treatments and outcomes of each individual. Theoretically, we showed that an intervention on a treatment policy is equivalent to a sequence of stochastic interventions on treatments, whose potential outcomes can be estimated from observational data with the joint model. In a semi-synthetic experiment, we demonstrated that the joint model can answer causal queries about the interventional and counterfactual distributions of the outcome after an intervention on the treatment policy.

References

Ioana Bica, Ahmed M Alaa, James Jordon, and Mihaela van der Schaar. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. *arXiv preprint arXiv:2002.04083*, 2020.

Daryl J Daley and David Vere-Jones. *An introduction to the theory of point processes: volume I: elementary theory and methods*. Springer, 2003.

Miguel A Hernán and James M Robins. Causal inference, 2010.

Bryan Lim, Ahmed M Alaa, and Mihaela van der Schaar. Forecasting treatment responses over time using recurrent marginal structural networks. *advances in neural information processing systems*, 31, 2018.

Siqi Liu and Milos Hauskrecht. Nonparametric regressive point processes based on conditional gaussian processes. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 1064–1074, 2019.

Michael Oberst and David Sontag. Counterfactual off-policy evaluation with gumbel-max structural causal models. In *International Conference on Machine Learning*, pages 4881–4890. PMLR, 2019.

Judea Pearl. *Causality*. Cambridge university press, 2009.

Jakob Gulddahl Rasmussen. Temporal point processes: the conditional intensity function. *Lecture Notes, Jan*, 2011.

Peter Schulam and Suchi Saria. Reliable decision support using counterfactual models. *Advances in neural information processing systems*, 30, 2017.

Nabeel Seedat, Fergus Imrie, Alexis Bellot, Zhaozhi Qian, and Mihaela van der Schaar. Continuous-time modeling of counterfactual outcomes using neural controlled differential equations. *arXiv preprint arXiv:2206.08311*, 2022.

Stratis Tsirtsis and Manuel Gomez Rodriguez. Decisions, counterfactual explanations and strategic behavior. *Advances in Neural Information Processing Systems*, 33:16749–16760, 2020.

Stratis Tsirtsis, Abir De, and Manuel Rodriguez. Counterfactual explanations in sequential decision making under uncertainty. *Advances in Neural Information Processing Systems*, 34:30127–30139, 2021.

Patrick Wyatt, Sarah E Berry, Graham Finlayson, Ruairi O'Driscoll, George Hadjigeorgiou, David A Drew, Haya Al Khatib, Long H Nguyen, Inbar Linenberg, Andrew T Chan, et al. Postprandial glycaemic dips predict appetite and energy intake in healthy individuals. *Nature metabolism*, 3(4): 523–529, 2021.

Yanbo Xu, Yanxun Xu, and Suchi Saria. A bayesian nonparametric approach for estimating individualized treatment-response curves. In *Machine learning for healthcare conference*, pages 282–300. PMLR, 2016.

Guangyi Zhang, Reza A Ashrafi, Anne Juuti, Kirsi Pietiläinen, and Pekka Marttinen. Errors-in-variables modeling of personalized treatment-response trajectories. *IEEE Journal of Biomedical and Health Informatics*, 25(1):201–208, 2020.