
Time Series Deconfounder: Estimating Treatment Effects over Time in the Presence of Hidden Confounders

Ioana Bica^{1,2} Ahmed M. Alaa³ Mihaela van der Schaar^{2,3,4}

Abstract

The estimation of treatment effects is a pervasive problem in medicine. Existing methods for estimating treatment effects from longitudinal observational data assume that there are no hidden confounders, an assumption that is not testable in practice and, if it does not hold, leads to biased estimates. In this paper, we develop the Time Series Deconfounder, a method that leverages the assignment of multiple treatments over time to enable the estimation of treatment effects in the presence of multi-cause hidden confounders. The Time Series Deconfounder uses a novel recurrent neural network architecture with multitask output to build a factor model over time and infer latent variables that render the assigned treatments conditionally independent; then, it performs causal inference using these latent variables that act as substitutes for the multi-cause unobserved confounders. We provide a theoretical analysis for obtaining unbiased causal effects of time-varying exposures using the Time Series Deconfounder. Using both simulated and real data we show the effectiveness of our method in deconfounding the estimation of treatment responses over time.

1. Introduction

Forecasting the patient’s response to treatments assigned over time represents a crucial problem in the medical domain. The increasing availability of observational data makes it possible to learn individualized treatment responses from longitudinal disease trajectories containing information about patient covariates and treatment assignments (Robins et al., 2000a; Robins & Hernán, 2008; Schulam

& Saria, 2017; Lim et al., 2018; Bica et al., 2020a). Existing methods for estimating individualized treatment effects over time assume that all confounders—variables affecting the treatment assignments and the potential outcomes—are observed, an assumption which is not testable in practice¹ and probably not true in many situations.

To understand why the presence of hidden confounders introduces bias, consider the problem of estimating treatment effects for patients with cancer. They are often prescribed multiple treatments at the same time, including chemotherapy, radiotherapy and/or immunotherapy based on their tumor characteristics. These treatments are adjusted if the tumor size changes. The treatment strategy is also changed as the patient starts to develop drug resistance (Vlachostergios & Faltas, 2018) or the toxicity levels of the drugs increase (Kroschinsky et al., 2017). Drug resistance and toxicity levels are multi-cause confounders since they affect not only the multiple causes (treatments)², but also the patient outcome (e.g. mortality, risk factors). However, drug resistance and toxicity may not be observed and, even if observed, may not be recorded in the electronic health records. Estimating, for instance, the effect of chemotherapy on the cancer progression in the patient without accounting for the dependence on drug resistance and toxicity levels (hidden confounders) will produce biased results.

Wang & Blei (2019) developed theory for deconfounding—adjusting for the bias introduced by the existence of hidden confounders in observational data—in the *static* causal inference setting and noted that the existence of multiple causes makes this task easier. Wang & Blei (2019) observed that the dependencies in the assignment of multiple causes can be used to infer latent variables that render the causes independent and act as substitutes for the hidden confounders.

In this paper, we propose the Time Series Deconfounder, a method that enables the unbiased estimation of treatment responses *over time* in the presence of hidden confounders, by taking advantage of the dependencies in the sequential assignment of multiple treatments. We draw from the main

¹University of Oxford, Oxford, United Kingdom ²The Alan Turing Institute, London, United Kingdom ³UCLA, Los Angeles, USA ⁴University of Cambridge, Cambridge, United Kingdom. Correspondence to: Ioana Bica <ioana.bica@eng.ox.ac.uk>.

¹Since counterfactuals are never observed, it is not possible to test for the existence of hidden confounders that could affect them.

²Causes and treatments are used interchangeably throughout the paper.

idea in Wang & Blei (2019), but note that the estimation of hidden confounders in the longitudinal setting is significantly more complex than in the static setting, not just because the hidden confounders may vary over time but in particular because the hidden confounders may be affected by previous treatments and covariates. Thus, standard latent variable models are no longer applicable, as they cannot capture these time dependencies.

The Time Series Deconfounder relies on building a factor model *over time* to obtain latent variables which, together with the observed variables render the assigned causes conditionally independent. Through theoretical analysis we show that these latent variables can act as substitutes for the multi-cause unobserved confounders and can be used to satisfy the sequential strong ignorability condition in the potential outcomes framework for time-varying exposures (Robins & Hernán, 2008) and obtain unbiased estimates of individualized treatment responses, using weaker assumptions than standard methods. Following our theory, we propose a novel deep learning architecture, based on a recurrent neural network with multi-task outputs and variational dropout, to build such a factor model and infer the substitutes for the hidden confounders in practice.

The Time Series Deconfounder shifts the need for observing all multi-cause confounders (untestable condition) to constructing a good factor model over time (testable condition). To assess how well the factor model captures the distribution of assigned treatments, we extend the use of predictive checks (Rubin, 1984; Wang & Blei, 2019) to the temporal setting and compute p -values at each timestep. We perform experiments on a simulated dataset where we control the amount of hidden confounding applied and on a real dataset with patients in the ICU (Johnson et al., 2016) to show how the Time Series Deconfounder allows us to deconfound the estimation of treatment responses in longitudinal data. To the best of our knowledge, this represents the first method for learning latent variables that can act as substitutes for the unobserved confounders in the time series setting.

2. Related Work

Previous methods for causal inference mostly focused on the static setting (Hill, 2011; Wager & Athey, 2017; Alaa & van der Schaar, 2017; Shalit et al., 2017; Yoon et al., 2018; Alaa & Schaar, 2018; Zhang et al., 2020; Bica et al., 2020c), and less attention has been given to the time series setting. We discuss methods for estimating treatment effects over time, as well as methods for inferring substitute hidden confounders in the static setting.

Potential outcomes for time-varying treatment assignments. Standard methods for performing counterfactual inference in longitudinal data are found in the epidemiol-

ogy literature and include the g -computation formula, g -estimation of structural nested mean models, and inverse probability of treatment weighting of marginal structural models (Robins, 1994; Robins et al., 2000a; Robins & Hernán, 2008). Alternatively, (Lim et al., 2018) improves on the standard marginal structural models by using recurrent neural networks to estimate the propensity weights and treatment responses, while (Bica et al., 2020a) propose using balancing representations to handle the time-dependent confounding bias when estimating treatment effects over time. Despite the wide applicability of these methods in forecasting treatment responses, they are all based on the assumption that there are no hidden confounders. Our paper proposes a method for deconfounding such outcome models, by inferring substitutes for the hidden confounders which can lead to unbiased estimates of the potential outcomes.

The potential outcomes framework has been extended to the continuous-time setting by (Lok et al., 2008). Several methods have been proposed for estimating treatment responses in continuous time (Soleimani et al., 2017; Schulam & Saria, 2017), again assuming that there are no hidden confounders. Here, we focus on deconfounding the estimation of treatment responses in the discrete-time setting.

Sensitivity analysis methods that evaluate the potential impact that an unmeasured confounder could have on the estimation of treatment effects have also been developed (Robins et al., 2000b; Roy et al., 2016; Scharfstein et al., 2018). However, these methods assess the suitability of applying existing tools, rather than propose a direct solution for handling the presence of hidden confounders in observational data.

Latent variable models for estimating hidden confounders. The most similar work to ours is the one of Wang & Blei (2019), who proposed the deconfounder, an algorithm that infers latent variables that act as substitutes for the hidden confounders and then performs causal inference in the static multi-cause setting. The deconfounder involves finding a good factor model of the assigned causes which can be used to estimate substitutes for the hidden confounders. Then, the deconfounder fits an outcome model for estimating the causal effects using the inferred latent variables. Our paper extends the theory for the deconfounder to the time-varying treatments setting and shows how the inferred latent variables can lead to sequential strong ignorability. To estimate the substitute confounders, Wang & Blei (2019) used standard factor models (Tipping & Bishop, 1999; Ranganath et al., 2015), which are only applicable in the static setting. To build a factor model over time, we propose an RNN architecture with multitask output and variational dropout.

Several other methods have been proposed for taking advantage of the multiplicity of assigned treatments in the

static setting and capture shared latent confounding (Tran & Blei, 2018; Heckerman, 2018; Ranganath & Perotte, 2018). These works are based on Pearl’s causal framework (Pearl, 2009) and use structural equation models. Alternative methods for dealing with hidden confounders in the static setting use proxy variables as noisy substitutes for the confounders (Lash et al., 2014; Louizos et al., 2017; Lee et al., 2018).

A different line of research involves performing causal discovery in the presence of hidden confounders (Spirtes et al., 2000). In this context, several methods have been proposed to perform causal graphical model structure learning with latent variables (Leray et al., 2008; Jabbari et al., 2017; Raghu et al., 2018). However, in this paper, we are not aiming to discover causal relationships between patient covariates over time. Instead, we improve existing methods for estimating the individualized effects of time-dependent treatments by accounting for multi-cause unobserved confounders.

3. Problem Formulation

Let the random variables $\mathbf{X}_t^{(i)} \in \mathcal{X}_t$ be the time-dependent covariates, $\mathbf{A}_t^{(i)} = [A_{t1}^{(i)} \dots A_{tk}^{(i)}] \in \mathcal{A}_t$ be the possible assignment of k treatments (causes) at timestep t and let $\mathbf{Y}_{t+1}^{(i)} \in \mathcal{Y}_t$ be the observed outcomes for patient (i). Treatments can be either binary and/or continuous. Static features, such as genetic information, do not change our theory, and, for simplicity, we assume they are part of the observed covariates.

The observational data for patient (i), also known as the patient trajectory, consists of realizations of the previously described random variables $\zeta^{(i)} = \{\mathbf{x}_t^{(i)}, \mathbf{a}_t^{(i)}, \mathbf{y}_{t+1}^{(i)}\}_{t=1}^{T^{(i)}}$, with samples collected for $T^{(i)}$ discrete and regular timesteps. Electronic health records consist of data for N independent patients $\mathcal{D} = \{\tau^{(i)}\}_{i=1}^N$. For simplicity, we omit the patient superscript (i) unless it is explicitly needed.

We leverage the potential outcomes framework proposed by Rubin (1978) and Neyman (1923), and extended by Robins & Hernán (2008) to take into account time-varying treatments. Let $\mathbf{Y}(\bar{\mathbf{a}})$ be the potential outcome, either factual or counterfactual, for each possible course of treatment $\bar{\mathbf{a}}$.

Let $\bar{\mathbf{A}}_t = (\mathbf{A}_1, \dots, \mathbf{A}_t) \in \bar{\mathcal{A}}_t$ be the history of treatments and let $\bar{\mathbf{X}}_t = (\mathbf{X}_1, \dots, \mathbf{X}_t) \in \bar{\mathcal{X}}_t$ be the history of covariates until timestep t . For each patient, we want to estimate individualized treatment effects, i.e. potential outcomes conditional on the patient history of covariates and treatments:

$$\mathbb{E}[\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t], \quad (1)$$

for any possible treatment plan $\bar{\mathbf{a}}_{\geq t}$ that starts at timestep t and consists of a sequence of treatments that ends just before the patient outcome \mathbf{Y} is observed. The observational data can be used to fit a regression model to estimate

$\mathbb{E}[\mathbf{Y} \mid \bar{\mathbf{a}}_{\geq t}, \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t]$. Under certain assumptions, these estimates are unbiased so that $\mathbb{E}[\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \mid \bar{\mathbf{X}}_t, \bar{\mathbf{A}}_{t-1}] = \mathbb{E}[\mathbf{Y} \mid \bar{\mathbf{a}}_{\geq t}, \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t]$. These conditions include Assumptions 1 and 2, which are standard among the existing methods for estimating treatment effects over time and can be tested in practice (Robins & Hernán, 2008).

Assumption 1. Consistency. *If $\bar{\mathbf{A}}_{\geq t} = \bar{\mathbf{a}}_{\geq t}$, then the potential outcomes for following the treatment plan $\bar{\mathbf{a}}_{\geq t}$ is the same as the observed (factual) outcome $\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) = \mathbf{Y}$.*

Assumption 2. Positivity (Overlap) (Imai & Van Dyk, 2004): *If $P(\bar{\mathbf{A}}_{t-1} = \bar{\mathbf{a}}_{t-1}, \bar{\mathbf{X}}_t = \bar{\mathbf{x}}_t) \neq 0$ then $P(\mathbf{A}_t = \mathbf{a}_t \mid \bar{\mathbf{A}}_{t-1} = \bar{\mathbf{a}}_{t-1}, \bar{\mathbf{X}}_t = \bar{\mathbf{x}}_t) > 0$ for all \mathbf{a}_t .*

The positivity assumption means that at each timestep t , each treatment has a non-zero probability of being given to the patient. This assumption is testable in practice.

In addition to these two assumptions, existing methods also assume *sequential strong ignorability*:

$$\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \perp\!\!\!\perp \mathbf{A}_t \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \quad (2)$$

for all possible treatment plans $\bar{\mathbf{a}}_{\geq t}$ and for all $t \in \{0, \dots, T\}$. This condition holds if there are no hidden confounders and it cannot be tested in practice. To understand why this is the case, note that the sequential strong ignorability assumption requires the conditional independence of the treatments with all of the potential outcomes, both factual and counterfactual, conditional on the patient history. Since the counterfactuals are never observed, it is not possible to test for this conditional independence.

In this paper, we assume that there are hidden confounders. Consequently, using standard methods for computing $\mathbb{E}[\mathbf{Y} \mid \bar{\mathbf{a}}_{\geq t}, \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t]$ from the dataset will result in biased estimates since the hidden confounders introduce a dependence between the treatments at each timestep and the potential outcomes ($\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \not\perp\!\!\!\perp \mathbf{A}_t \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t$) and therefore:

$$\mathbb{E}[\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t] \neq \mathbb{E}[\mathbf{Y} \mid \bar{\mathbf{a}}_{\geq t}, \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t]. \quad (3)$$

By extending the method proposed by Wang & Blei (2019), we take advantage of the multiple treatment assignments at each timestep to infer a sequence of latent variables $\bar{\mathbf{Z}}_t = (\mathbf{Z}_1, \dots, \mathbf{Z}_t) \in \bar{\mathcal{Z}}_t$ that can be used as substitutes for the unobserved confounders. We will then show how $\bar{\mathbf{Z}}_t$ can be used to estimate the treatment effects over time.

4. Time Series Deconfounder

The idea behind the Time Series Deconfounder is that multi-cause confounders introduce dependencies between the treatments. As treatment assignments change over time we infer substitutes for the hidden confounders that take advantage of the patient history to capture these dependencies.

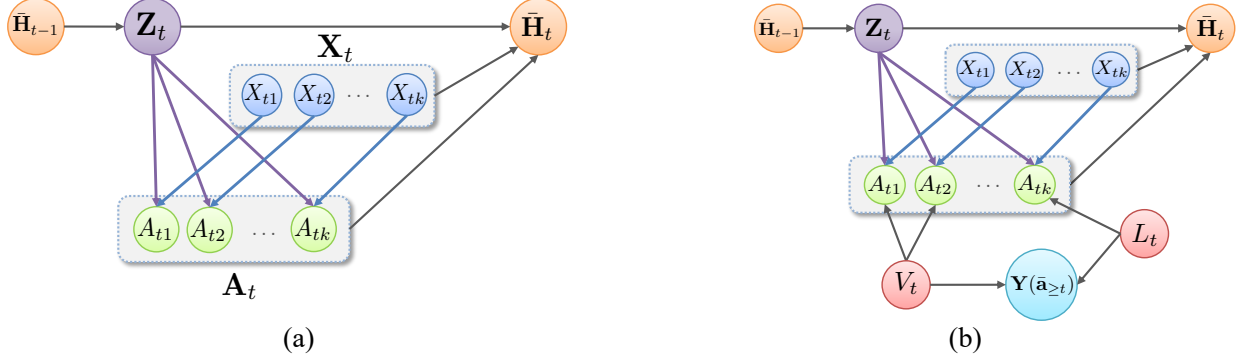


Figure 1. (a) Graphical factor model. Each Z_t is built as a function of the history, such that, with X_t , it renders the assigned causes conditionally independent: $p(a_{t1}, \dots, a_{tk} | z_t, \mathbf{x}_t) = \prod_{j=1}^k p(a_{tj} | z_t, \mathbf{x}_t)$. The variables can be connected to $Y(\bar{a}_{\geq t})$ in any way. (b) Graphical model explanation for why this factor model construction ensures that Z_t captures all of the multi-cause hidden confounders.

4.1. Factor Model

The Time Series Deconfounder builds a factor model to capture the distribution of the causes (treatments) over time. At time t , the factor model constructs the latent variable $\mathbf{z}_t = g(\bar{\mathbf{h}}_{t-1})$, where $\bar{\mathbf{h}}_{t-1} = (\bar{\mathbf{a}}_{t-1}, \bar{\mathbf{x}}_{t-1}, \bar{\mathbf{z}}_{t-1})$ is the realization of history $\bar{\mathbf{H}}_{t-1}$. The latent variable \mathbf{z}_t , together with the observed patient covariates \mathbf{x}_t , render the assigned treatments conditionally independent:

$$p(a_{t1}, \dots, a_{tk} | z_t, \mathbf{x}_t) = \prod_{j=1}^k p(a_{tj} | z_t, \mathbf{x}_t). \quad (4)$$

Figure 1(a) illustrates the corresponding graphical model for timestep t . The factor model of the assigned treatments is built as a latent variable model with joint distribution:

$$p(\theta_{1:k}, \bar{\mathbf{x}}_T, \bar{\mathbf{z}}_T, \bar{\mathbf{a}}_T) = p(\theta_{1:k})p(\bar{\mathbf{x}}_T) \cdot \prod_{t=1}^T (p(\mathbf{z}_t | \bar{\mathbf{h}}_{t-1}) \prod_{j=1}^k p(a_{tj} | z_t, \mathbf{x}_t, \theta_j)), \quad (5)$$

where $\theta_{1:k}$ are parameters. The distribution of the treatments $p(\bar{\mathbf{a}}_T)$ is the corresponding marginal. Notice that we do not assume that, in the observational data, the patient covariates \mathbf{x}_t at timestep t are independent of the patient history. The graphical factor model shows how the latent variables \mathbf{z}_t that can act as substitutes for the multi-cause unobserved confounders are built. As we will see in Section 4.2 these latent variables will be used as part of an outcome model that estimates the potential outcomes $Y(\bar{a}_{\geq t})$.

By taking advantage of the dependencies between the multiple treatment assignments, the factor model allows us to infer the sequence of latent variables $\bar{\mathbf{Z}}_t$ that can be used to render the assigned causes conditionally independent. Through this factor model construction and under correct model specifications, we can rule out the existence of other multi-cause confounders that are not captured by Z_t . To understand why this is the case, consider the graphical model

in Figure 1(b). By contradiction, assume that there exists another multi-cause confounder V_t not captured by Z_t . Then, by d -separation the conditional independence between the assigned causes given Z_t and X_t does not hold anymore. This argument cannot be used for single-cause confounders, such as L_t , which are only affecting one of the causes and the potential outcomes. Thus, we assume sequential single strong ignorability (no hidden single cause confounders).

Assumption 3. Sequential single strong ignorability:

$$Y(\bar{a}_{\geq t}) \perp\!\!\!\perp A_{tj} | X_t, \bar{H}_{t-1}, \quad (6)$$

for all $\bar{a}_{\geq t}$, for all $t \in \{0, \dots, T\}$, and for all $j \in \{1, \dots, k\}$.

Causal inference relies on assumptions. Existing methods for estimating treatment effects over time assume that there are no multi-cause and no single-cause hidden confounders. In this paper, we make the *weaker* assumption that there are no single-cause hidden confounders. While this assumption is also untestable in practice, as the number of treatments increases for each timestep, it becomes increasingly weaker: the more treatments we observe, the more likely it becomes for a hidden confounder to affect multiple of these treatments rather than a single one of them.

Theorem 1. *If the distribution of the assigned causes $p(\bar{\mathbf{a}}_T)$ can be written as the factor model $p(\theta_{1:k}, \bar{\mathbf{x}}_T, \bar{\mathbf{z}}_T, \bar{\mathbf{a}}_T)$, we obtain sequential ignorable treatment assignment:*

$$Y(\bar{a}_{\geq t}) \perp\!\!\!\perp (A_{t1}, \dots, A_{tk}) | \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \bar{\mathbf{Z}}_t, \quad (7)$$

for all $\bar{a}_{\geq t}$ and for all $t \in \{0, \dots, T\}$.

Theorem 1 is proved by leveraging Assumption 3, the fact that the latent variables Z_t are inferred without knowledge of the potential outcomes $Y(\bar{a}_{\geq t})$ and the fact that the causes (A_{t1}, \dots, A_{tk}) are jointly independent given Z_t and X_t . The result means that, at each timestep, the variables

$\bar{\mathbf{X}}_t, \bar{\mathbf{Z}}_t, \bar{\mathbf{A}}_{t-1}$ contain all of the dependencies between the potential outcomes $\mathbf{Y}(\bar{\mathbf{a}}_{\geq t})$ and the assigned causes \mathbf{A}_t . See Appendix A for the full proof.

As discussed in Wang & Blei (2019), the substitute confounders \mathbf{Z}_t also need to satisfy positivity (Assumption 2), i.e. if $P(\bar{\mathbf{A}}_{t-1} = \bar{\mathbf{a}}_{t-1}, \bar{\mathbf{Z}}_t = \bar{\mathbf{z}}_t, \bar{\mathbf{X}}_t = \bar{\mathbf{x}}_t) \neq 0$ then $P(\mathbf{A}_t = \mathbf{a}_t \mid \bar{\mathbf{A}}_{t-1} = \bar{\mathbf{a}}_{t-1}, \bar{\mathbf{Z}}_t = \bar{\mathbf{z}}_t, \bar{\mathbf{X}}_t = \bar{\mathbf{x}}_t) > 0$ for all \mathbf{a}_t . After fitting the factor model, this can be tested (Robins & Hernán, 2008). When positivity is limited, the outcome model estimates of treatment responses will also have high variance. In practice, positivity can be enforced by setting the dimensionality of \mathbf{Z}_t to be smaller than the number of treatments (Wang & Blei, 2019).

Predictive Checks over Time: The theory holds if the fitted factor model captures well the distribution of assigned treatments. This condition can be assessed by extending predictive model checking (Rubin, 1984) to the time-series setting. We compute p -values over time to evaluate how similar the distribution of the treatments learned by the factor model is with the distribution of the treatments in a validation set of patients. At each timestep t , for the patients in the validation set, we obtain M replicas of their treatment assignments $\{\mathbf{a}_{t,\text{rep}}^{(i)}\}_{i=1}^M$ by sampling from the factor model. The replicated treatment assignments are compared with the actual treatment assignments, $\mathbf{a}_{t,\text{val}}$, using the test statistic $T(\mathbf{a}_t)$:

$$T(\mathbf{a}_t) = \mathbb{E}_Z[\log p(\mathbf{a}_t \mid Z_t, X_t)], \quad (8)$$

which is related to the marginal log likelihood (Wang & Blei, 2019). The predictive p -value for timestep t is computed as follows:

$$\frac{1}{M} \sum_{i=1}^M \mathbf{1}\left(T\left(\mathbf{a}_{t,\text{rep}}^{(i)}\right) < T\left(\mathbf{a}_{t,\text{val}}\right)\right), \quad (9)$$

where $\mathbf{1}(\cdot)$ represents the indicator function.

If the model captures well the distribution of the assigned causes, then the test statistics for the treatment replicas are similar to the test statistics for the treatments in the validation set, which makes 0.5 the ideal p -value in this case.

4.2. Outcome Model

If the factor model passes the predictive checks, the Time Series Deconfounder fits an outcome model (Robins et al., 2000a; Lim et al., 2018) to estimate individualized treatment effects over time. After sampling the sequence of latent variables $\hat{\mathbf{Z}}_t = [\hat{\mathbf{Z}}_1 \dots \hat{\mathbf{Z}}_t]$ from the factor model, the outcome model can be used to estimate $\mathbb{E}[\mathbf{Y} \mid \bar{\mathbf{a}}_{\geq t}, \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \hat{\mathbf{Z}}_t] = \mathbb{E}[\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \hat{\mathbf{Z}}_t]$.

To compute uncertainty estimates of the potential outcomes, we can sample $\hat{\mathbf{Z}}_t$ repeatedly and then fit an outcome

model for each sample to obtain multiple point estimates of $\mathbf{Y}(\bar{\mathbf{a}}_{\geq t})$. The variance of these point estimates will represent the uncertainty of the Time Series Deconfounder.

D’Amour (2019) raised some concerns about identifiability of the mean potential outcomes using the deconfounder framework in Wang & Blei (2019) in the static setting and illustrated some pathological examples where identifiability might not hold.³ In practical settings, the outcome estimates from the Time Series Deconfounder are identifiable, as supported by the experimental results in Sections 6 and 7. Nevertheless, when identifiability represents an issue, the uncertainty in the potential outcomes can be used to assess the reliability of the Time Series Deconfounder. In particular, the variance in the potential outcomes indicates how the finite observational data inform the estimation of substitutes for the hidden confounders and subsequently the treatment outcomes of interest. When the treatment effects are non-identifiable, the estimates of the Time Series Deconfounder will have high variance.

By using this framework to estimate substitutes for the hidden confounders we are trading off confounding bias for estimation variance (Wang & Blei, 2019). The treatment effects computed without accounting for the hidden confounders will inevitably be biased. Alternatively, using the latent variables from the factor model will result in unbiased, but higher variance estimates of treatment effects.

5. Factor Model over Time in Practice

Since we are dealing with time-varying treatments, we cannot use standard factor models, such as PCA (Tipping & Bishop, 1999) or Deep Exponential Families (Ranganath et al., 2015), as they can only be applied in the static setting. Using the theory developed for the factor model over time we introduce a practical implementation based on a recurrent neural network (RNN) with multitask output and variational dropout as illustrated in Figure 2.

The recurrent part of the model infers the latent variables \mathbf{Z}_t such that they depend on the patient history:

$$\mathbf{Z}_1 = \text{RNN}(\mathbf{L}), \quad (10)$$

$$\mathbf{Z}_t = \text{RNN}(\bar{\mathbf{Z}}_{t-1}, \bar{\mathbf{X}}_{t-1}, \bar{\mathbf{A}}_{t-1}, \mathbf{L}), \quad (11)$$

where \mathbf{L} consists of randomly initialized parameters that are trained with the rest of the parameters in the RNN.

The size of the RNN output is D_Z and this specifies the size of the latent variables that are inferred as substitutes for the hidden confounders. In our experiments, we use an LSTM unit (Hochreiter & Schmidhuber, 1997) as part of the RNN. Moreover, to infer the assigned treatments at

³See Wang & Blei (2019) for a longer discussion addressing the concerns in (D’Amour, 2019).

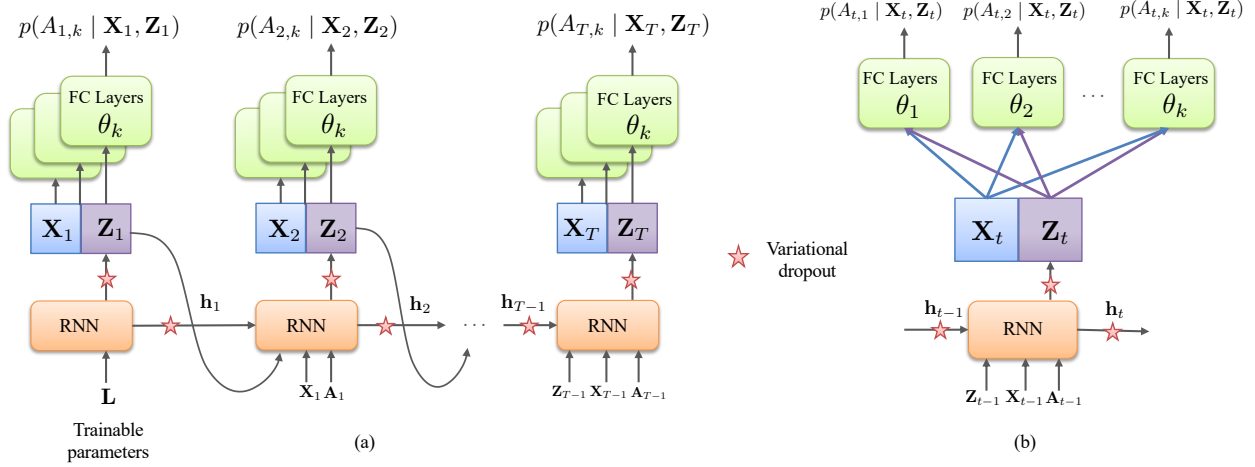


Figure 2. (a) Proposed factor model implementation. \mathbf{Z}_t is generated by RNN as a function of the history $\bar{\mathbf{H}}_{t-1}$, given by the hidden state \mathbf{h}_t , and current input. Multitask output is used to construct the treatments such that they are independent given \mathbf{Z}_t and \mathbf{X}_t . (b) Closer look at a single timestep.

timestep t , $\mathbf{A}_t = [A_{t1}, \dots, A_{tk}]$ such that they are conditionally independent given the latent variables \mathbf{Z}_t and the observed covariates \mathbf{X}_t , we propose using multitask multilayer perceptrons (MLPs) consisting of fully connected (FC) layers:

$$A_{tj} = \text{FC}(\mathbf{X}_t, \mathbf{Z}_t; \theta_j), \quad (12)$$

for all $j = 1, \dots, k$ and for all $t = 1, \dots, T$, where θ_j are the parameters in the FC layers used to obtain A_{tj} . We use a single FC hidden layer before the output layer. For binary treatments, the sigmoid activation is used in the output layer. For continuous treatments, MC dropout (Gal & Ghahramani, 2016a) can instead be applied in the FC layers to obtain $p(A_{tj} | \mathbf{X}_t, \mathbf{Z}_t)$.

To model the probabilistic nature of factor models we incorporate *variational dropout* (Gal & Ghahramani, 2016b) in the RNN as illustrated in Figure 2. Using dropout enables us to obtain samples from \mathbf{Z}_t and treatment assignments A_{tj} . These samples allow us to obtain treatment replicas and to compute predictive checks over time, but also to estimate the *uncertainty* in \mathbf{Z}_t and potential outcomes.

Using the treatment assignments from the observational dataset, the factor model can be trained using gradient descent based methods. The proposed factor model architecture follows from the theory developed in Section 4 where at each timestep the latent variable \mathbf{Z}_t is built as a function of the history (parametrized by an RNN). The multitask output is essential for modeling the conditional independence between the assigned treatments given the latent variables generated by the RNN and the observed covariates. The factor model can be extended to allow for irregularly sampled data by using a PhasedLSTM (Neil et al., 2016).

Note that our theory does not put restrictions on the factor model that can be used. Alternative factor models over time

are generalized dynamic-factor model (Forni et al., 2000; 2005) or factor-augmented vector autoregressive models (Bermanke et al., 2005). These come from the econometrics literature and explicitly model the dynamics in the data. The use of RNNs in the factor model enables us to learn complex relationships between $\bar{\mathbf{X}}_t$, $\bar{\mathbf{Z}}_t$, and $\bar{\mathbf{A}}_t$ from the data, which is needed in medical applications involving complex diseases. Nevertheless, predictive checks should be used to assess any selected factor model.

6. Experiments on Synthetic Data

To validate the theory developed in this paper, we perform experiments on synthetic data where we vary the effect of hidden confounding. It is not possible to validate the method on real datasets since the true extent of hidden confounding is never known (Wang & Blei, 2019; Louizos et al., 2017).

6.1. Simulated Dataset

To keep the simulation process general, we propose building a dataset using p -order autoregressive processes. At each timestep t , we simulate k time-varying covariates $X_{t,k}$ representing single cause confounders and a multi-cause hidden confounder Z_t as follows:

$$X_{t,j} = \frac{1}{p} \sum_{i=1}^p (\alpha_{i,j} X_{t-i,j} + \omega_{i,j} A_{t-i,j}) + \eta_t \quad (13)$$

$$Z_t = \frac{1}{p} \sum_{i=1}^p (\beta_i Z_{t-i} + \sum_{j=1}^k \lambda_{i,j} A_{t-i,j}) + \epsilon_t, \quad (14)$$

for $j = 1, \dots, k$, $\alpha_{i,k}, \lambda_{i,j} \sim \mathcal{N}(0, 0.5^2)$, $\omega_{i,k}, \beta_i \sim \mathcal{N}(1 - (i/p), (1/p)^2)$, and $\eta_t, \epsilon_t \sim \mathcal{N}(0, 0.01^2)$. The value of Z_t changes over time and is affected by the treatment

assignments.

Each treatment assignment $A_{t,j}$ depends on the single-cause confounder $X_{t,j}$ and multi-cause hidden confounder Z_t :

$$\pi_{tj} = \gamma_A \hat{Z}_t + (1 - \gamma_A) \hat{X}_{tj} \quad (15)$$

$$A_{tj} | \pi_{tj} \sim \text{Bernoulli}(\sigma(\lambda \pi_{tj})), \quad (16)$$

where \hat{X}_{tj} and \hat{Z}_t are the sum of the covariates and confounders respectively over the last p timesteps, $\lambda = 15$, $\sigma(\cdot)$ is the sigmoid function and γ_A controls the amount of hidden confounding applied to the treatment assignments. The outcome is also obtained as a function of the covariates and the hidden confounder:

$$\mathbf{Y}_{t+1} = \gamma_Y Z_{t+1} + (1 - \gamma_Y) \left(\frac{1}{k} \sum_{j=1}^k X_{t+1,j} \right), \quad (17)$$

where γ_Y controls the amount of hidden confounding applied to the outcome. We simulate datasets consisting of 5000 patients, with trajectories between 20 and 30 timesteps, and $k = 3$ covariates and treatments. To induce time dependencies we set $p = 5$. Each dataset undergoes a 80/10/10 split for training, validation and testing respectively. Hyperparameter optimization is performed for each trained factor model as explained in Appendix B. Using the training observational dataset, we fit the Time Series Deconfounder to perform one-step ahead estimation of treatment responses.

6.2. Evaluating Factor Model using Predictive Checks

Our theory for using the inferred latent variables as substitutes for the hidden confounders and obtain unbiased treatment responses relies on the fact that the factor model captures well the distribution of the assigned causes.




Figure 3. Predictive checks over time. We show the mean p -values at each timestep and the std error.

To assess the suitability of our proposed factor model architecture, we compare it with the following baselines: RNN without multitask output (predicting the k treatment assignments by passing \mathbf{X}_t and \mathbf{Z}_t through a FC layer and output

layer with k neurons) and multilayer perceptron (MLP) used instead of the RNN at each timestep to generate \mathbf{Z}_t . The MLP factor model does not use the entire history for generating \mathbf{Z}_t . See Appendix C for details.

Figure 3 shows the p -values over time computed for the test set in 30 simulated datasets with $\gamma_A = \gamma_Y = 0.5$. The p -values for the MLP factor model decrease over time, which means that there is a consistent distribution mismatch between the treatment assignments learned by this model and the ones in the test set. Conversely, the predictive checks for our proposed factor model are closer to the ideal p -value of 0.5. This illustrates that having an architecture capable of capturing time-dependencies and accumulating past information for inferring the latent confounders is crucial. Moreover, the performance for the RNN without multitask is similar to our model, which indicates that the factor model constraint does not affect the performance in capturing the distribution of the causes.

6.3. Deconfounding the Estimation of Treatment Responses over Time

We evaluate how well the Time Series Deconfounder⁴ can remove hidden confounding bias when used in conjunction with the following outcome models:

Standard Marginal Structural Models (MSMs). MSMs (Robins et al., 2000a; Hernán et al., 2001) have been widely used in epidemiology to estimate treatment effects over time. MSMs use inverse probability of treatment weighting (IPTW) to adjust for the time-dependent confounding bias present in observational datasets (Mansournia et al., 2017; Bica et al., 2020b). MSMs compute the propensity weights by using logistic regression; through IPTW, these models construct a pseudo-population from the observational data where the treatment assignment probability no longer depends on the time-varying confounders. The treatment responses over time are computed using linear regression. For full implementation details in Appendix D.1.

Recurrent Marginal Structural Networks (R-MSNs). R-MSNs (Lim et al., 2018) also use IPTW to remove the bias from time-dependent confounders when estimating treatment effects over time. However, R-MSNs estimate the propensity scores using RNNs instead. The use of RNNs is more robust to changes in the treatment assignment policy. To estimate the treatment responses over time R-MSNs also use a model based on RNNs. For implementation details, see Appendix D.2.

⁴The implementation of the Time Series Deconfounder can be found at https://bitbucket.org/mvdschaar/mlforhealthlabpub/src/master/alg/time_series_deconfounder/ and at <https://github.com/ioanabica/Time-Series-Deconfounder>.




Figure 4. Results for deconfounding the one-step ahead estimation of treatment responses in two outcome models: (a) Marginal Structural Models (MSM) and (b) Recurrent Marginal Structural Networks (R-MSN). The average RMSE and the standard error in the results are computed for 30 dataset simulations for each different degree of confounding, as measured by γ .

Notice that these outcome models were also chosen because they are capable of handling multiple treatments (that may be assigned simultaneously) at the same timestep. In the simulated dataset, parameters γ_A and γ_Y control the amount of hidden confounding applied to the treatments and outcomes respectively. We vary this amount through $\gamma_A = \gamma_Y = \gamma$.

The Time Series Deconfounder is used to obtain unbiased estimates of one-step ahead treatment responses. For a comparative evaluation, the outcome models are trained to estimate these treatment responses in the following scenarios: without information about $\bar{\mathbf{Z}}_t$ in which case they estimate $\mathbb{E}[\mathbf{Y}_{t+1}(\mathbf{a}_t) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t]$ (Confounded), with information about the simulated (oracle) $\bar{\mathbf{Z}}_t$ which leads to the following regression model $\mathbb{E}[\mathbf{Y}_{t+1}(\mathbf{a}_t) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \bar{\mathbf{Z}}_t]$ (Oracle), as well as after applying the Time Series Deconfounder with different model specifications and in this case $\mathbb{E}[\mathbf{Y}_{t+1}(\mathbf{a}_t) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \hat{\mathbf{Z}}_t]$ is estimated (Deconfounded). To highlight the importance of Assumption 3, we also apply the Time Series Deconfounder after removing the single-cause confounder X_1 , thus violating the assumption.

Figure 4 shows the root mean squared error (RMSE) for the one-step ahead estimation of treatment responses for patients in the test set. We notice that the Time Series Deconfounder gives unbiased estimates of treatment responses, i.e. close to the estimates obtained using the simulated (oracle) confounders. The method is robust to model misspecification, performing similarly when $D_Z = 1$ (simulated size of hidden confounders) and when $D_Z = 5$ (misspecified size of inferred confounders). When there are no hidden confounders ($\gamma = 0$), the extra information from $\hat{\mathbf{Z}}_t$ does not harm the estimations (although they have higher variance).

When the sequential single strong ignorability assumption (Assumption 3) is invalidated, namely when the single cause confounder X_1 is removed from the observational dataset,

we obtain biased estimates of the treatment responses. The performance in this case, however, is comparable to the performance when there is no control for the unobserved confounders.

In Appendix E, we consider an experimental set-up with a different simulated size for the hidden confounders (true $D_Z = 3$) and show results when the size of the hidden confounders is underestimated in the Time Series Deconfounder. We also include additional results on a simulated setting with static hidden confounders.

Source of gain: To understand the source of gain in the Time Series Deconfounder, consider why the outcome models fail in the scenarios when there are hidden confounders. MSMs and R-MSNs make the implicit assumption that the treatment assignments depend only on the observed history. The existence of any multi-cause confounders not captured by the history results in biased estimates of both the propensity weights and of the outcomes. On the other hand, the construction in our factor model rules out the existence of any multi-cause confounders which are not captured by \mathbf{Z}_t . By augmenting the data available to the outcome models with the substitute confounders, we eliminate these biases.

7. Experiments on MIMIC III

Using the Medical Information Mart for Intensive Care (MIMIC III) (Johnson et al., 2016) database consisting of electronic health records from patients in the ICU, we show how the Time Series Deconfounder can be applied on a real dataset. From MIMIC III we extracted a dataset with 6256 patients for which there are three treatment options at each timestep: antibiotics, vasopressors, and mechanical ventilator (all of which can be applied simultaneously). These treatments are common in the ICU and are often used to treat patients with sepsis (Schmidt et al., 2016; Scheeren et al.,

Table 1. Average RMSE $\times 10^2$ and standard error in the results for predicting the effect of antibiotics, vassopressors and mechanical ventilator on three patient covariates. The results are for 10 runs.

Outcome model	White blood cell count		Blood pressure		Oxygen saturation	
	MSM	R-MSN	MSM	R-MSN	MSM	R-MSN
Confounded	3.90 ± 0.00	2.91 ± 0.05	12.04 ± 0.00	10.29 ± 0.05	2.92 ± 0.00	1.74 ± 0.03
Deconfounded ($D_Z = 1$)	3.55 ± 0.05	2.62 ± 0.07	11.69 ± 0.14	9.35 ± 0.11	2.42 ± 0.02	1.24 ± 0.05
Deconfounded ($D_Z = 5$)	3.56 ± 0.04	2.41 ± 0.04	11.63 ± 0.10	9.45 ± 0.10	2.43 ± 0.02	1.21 ± 0.07
Deconfounded ($D_Z = 10$)	3.58 ± 0.03	2.48 ± 0.06	11.66 ± 0.14	9.20 ± 0.12	2.42 ± 0.01	1.17 ± 0.06
Deconfounded ($D_Z = 20$)	3.54 ± 0.04	2.55 ± 0.05	11.57 ± 0.12	9.63 ± 0.14	2.40 ± 0.01	1.28 ± 0.08

2019). For each patient, we extracted 25 patient covariates consisting of lab tests and vital signs measured over time that affect the assignment of treatments. We used daily aggregates of the patient covariates and treatments and patient trajectories of up to 50 timesteps. We estimate the effects of antibiotics, vasopressors, and mechanical ventilator on the following patient covariates: white blood cell count, blood pressure, and oxygen saturation.

Hidden confounding is present in the dataset as patient comorbidities and several lab tests were not included. However, since this is a real dataset, it is not possible to evaluate the extent of hidden confounding or to estimate the true (Oracle) treatment responses.

Table 1 illustrates the RMSE when estimating one-step ahead treatment responses by using the MSM and R-MSN outcome models directly on the extracted dataset (Confounded) and after applying the Time Series Deconfounder and augmenting the dataset with the substitutes for the hidden confounders of different dimensionality D_Z (Deconfounded). We notice that in all cases, the Time Series Deconfounder enables us to obtain a lower error when estimating the effect of antibiotics, vasopressors, and mechanical ventilator on the patients’ white blood cell count, blood pressure, and oxygen saturation. By modeling the dependencies in the assigned treatments for each patient, the factor model part of the Time Series Deconfounder was able to infer latent variables that account for the unobserved information about the patient states. Using these substitutes for the hidden confounders in the outcome models resulted in better estimates of the treatment responses. While these results on real data require further validation from doctors (which is outside the scope of this paper), they indicate the potential of the method to be applied in real medical scenarios.

In Appendix E, we include results for an additional experimental set-up where we remove several patient covariates from the dataset and we show how the Time Series Deconfounder can be used to account for this bias. Moreover, in Appendix F, we provide further discussion and directions for future work.

8. Conclusion

The availability of observational data consisting of longitudinal information about patients prompted the development of methods for modeling the effects of treatments on the disease progression in patients. All existing methods for estimating the individualized effects of time-dependent treatment from observational data make the untestable assumption that there are no hidden confounders. In the longitudinal setting, this assumption is even more problematic than in the static setting. As the state of the patient changes over time and the complexity of the treatment assignments and responses increases, it becomes much easier to miss important confounding information.

In this paper, we proposed the Time Series Deconfounder, a method that takes advantage of the patterns in the multiple treatment assignments over time to infer latent variables that can be used as substitutes for the hidden confounders. Moreover, we developed a deep learning architecture based on an RNN with multitask output and variational dropout for building a factor model over time and computing the latent variables in practice. Through experimental results on both synthetic and real datasets, we show the effectiveness of the Time Series Deconfounder in removing the bias from the estimation of treatment responses over time in the presence of multi-cause hidden confounders.

Acknowledgements

The authors would like to thank the reviewers for their helpful feedback. The research presented in this paper was supported by The Alan Turing Institute, under the EPSRC grant EP/N510129/1 and by the US Office of Naval Research (ONR).

References

Alaa, A. and Schaar, M. Limits of estimating heterogeneous treatment effects: Guidelines for practical algorithm design. In *International Conference on Machine Learning*, pp. 129–138, 2018.

- individualized treatment effects using multi-task gaussian processes. In *Advances in Neural Information Processing Systems*, pp. 3424–3432, 2017.
- Bernanke, B. S., Boivin, J., and Elias, P. Measuring the effects of monetary policy: a factor-augmented vector autoregressive (favar) approach. *The Quarterly journal of economics*, 120(1):387–422, 2005.
- Bica, I., Alaa, A. M., Jordon, J., and van der Schaar, M. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. *International Conference on Learning Representations*, 2020a.
- Bica, I., Alaa, A. M., Lambert, C., and van der Schaar, M. From real-world patient data to individualized treatment effects using machine learning: Current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics*, 2020b.
- Bica, I., Jordon, J., and van der Schaar, M. Estimating the effects of continuous-valued interventions using generative adversarial networks. *arXiv preprint arXiv:2002.12326*, 2020c.
- D’Amour, A. On multi-cause approaches to causal inference with unobserved confounding: Two cautionary failure cases and a promising alternative. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 3478–3486, 2019.
- Forni, M., Hallin, M., Lippi, M., and Reichlin, L. The generalized dynamic-factor model: Identification and estimation. *Review of Economics and statistics*, 82(4):540–554, 2000.
- Forni, M., Hallin, M., Lippi, M., and Reichlin, L. The generalized dynamic factor model: one-sided estimation and forecasting. *Journal of the American Statistical Association*, 100(471):830–840, 2005.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059, 2016a.
- Gal, Y. and Ghahramani, Z. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*, pp. 1019–1027, 2016b.
- Heckerman, D. Accounting for hidden common causes when inferring cause and effect from observational data. *arXiv preprint arXiv:1801.00727*, 2018.
- Hernán, M. A., Brumback, B., and Robins, J. M. Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *Journal of the American Statistical Association*, 96(454):440–448, 2001.
- Hill, J. L. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Imai, K. and Van Dyk, D. A. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467):854–866, 2004.
- Jabbari, F., Ramsey, J., Spirtes, P., and Cooper, G. Discovery of causal models that contain latent variables through bayesian scoring of independence constraints. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 142–157. Springer, 2017.
- Johnson, A. E., Pollard, T. J., Shen, L., Li-wei, H. L., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., and Mark, R. G. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3:160035, 2016.
- Kroschinsky, F., Stölzel, F., von Bonin, S., Beutel, G., Kochanek, M., Kiehl, M., and Schellongowski, P. New drugs, new toxicities: severe side effects of modern targeted and immunotherapy of cancer and their management. *Critical Care*, 21(1):89, 2017.
- Lash, T. L., Fox, M. P., MacLehose, R. F., Maldonado, G., McCandless, L. C., and Greenland, S. Good practices for quantitative bias analysis. *International journal of epidemiology*, 43(6):1969–1985, 2014.
- Lee, C., Mastrorade, N., and van der Schaar, M. Estimation of individual treatment effect in latent confounder models via adversarial learning. *arXiv preprint arXiv:1811.08943*, 2018.
- Leray, P., Megane, S., Maes, S., and Manderick, B. Causal graphical models with latent variables: learning and inference. In *Innovations in Bayesian Networks*, pp. 219–249. Springer, 2008.
- Lim, B., Alaa, A., and van der Schaar, M. Forecasting treatment responses over time using recurrent marginal structural networks. In *Advances in Neural Information Processing Systems*, pp. 7493–7503, 2018.
- Lok, J. J. et al. Statistical modeling of causal effects in continuous time. *The Annals of Statistics*, 36(3):1464–1507, 2008.
- Louizos, C., Shalit, U., Mooij, J. M., Sontag, D., Zemel, R., and Welling, M. Causal effect inference with deep latent-variable models. In *Advances in Neural Information Processing Systems*, pp. 6446–6456, 2017.

- Mansournia, M. A., Etmnan, M., Danaei, G., Kaufman, J. S., and Collins, G. Handling time varying confounding in observational research. *bmj*, 359:j4587, 2017.
- Neil, D., Pfeiffer, M., and Liu, S.-C. Phased lstm: Accelerating recurrent network training for long or event-based sequences. In *Advances in neural information processing systems*, pp. 3882–3890, 2016.
- Neyman, J. Sur les applications de la théorie des probabilités aux expériences agricoles: Essai des principes. *Roczniki Nauk Rolniczych*, 10:1–51, 1923.
- Pearl, J. *Causality*. Cambridge university press, 2009.
- Raghu, V. K., Ramsey, J. D., Morris, A., Manatakis, D. V., Sprites, P., Chrysanthis, P. K., Glymour, C., and Benos, P. V. Comparison of strategies for scalable causal discovery of latent variable models from mixed data. *International journal of data science and analytics*, 6(1):33–45, 2018.
- Ranganath, R. and Perotte, A. Multiple causal inference with latent confounding. *arXiv preprint arXiv:1805.08273*, 2018.
- Ranganath, R., Tang, L., Charlin, L., and Blei, D. Deep exponential families. In *Artificial Intelligence and Statistics*, pp. 762–771, 2015.
- Robins, J. M. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics-Theory and methods*, 23(8):2379–2412, 1994.
- Robins, J. M. and Hernán, M. A. Estimation of the causal effects of time-varying exposures. In *Longitudinal data analysis*, pp. 547–593. Chapman and Hall/CRC, 2008.
- Robins, J. M., Hernan, M. A., and Brumback, B. Marginal structural models and causal inference in epidemiology, 2000a.
- Robins, J. M., Rotnitzky, A., and Scharfstein, D. O. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In *Statistical models in epidemiology, the environment, and clinical trials*, pp. 1–94. Springer, 2000b.
- Roy, J., Lum, K. J., and Daniels, M. J. A bayesian non-parametric approach to marginal structural models for point treatments and a continuous or survival outcome. *Biostatistics*, 18(1):32–47, 2016.
- Rubin, D. B. Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*, pp. 34–58, 1978.
- Rubin, D. B. Bayesianly justifiable and relevant frequency calculations for the applies statistician. *The Annals of Statistics*, pp. 1151–1172, 1984.
- Scharfstein, D., McDermott, A., Díaz, I., Carone, M., Lunardon, N., and Turkoz, I. Global sensitivity analysis for repeated measures studies with informative drop-out: A semi-parametric approach. *Biometrics*, 74(1):207–219, 2018.
- Scheeren, T. W., Bakker, J., De Backer, D., Annane, D., Asfar, P., Boerma, E. C., Cecconi, M., Dubin, A., Dünser, M. W., Duranteau, J., et al. Current use of vasopressors in septic shock. *Annals of intensive care*, 9(1):20, 2019.
- Schmidt, G. A., Mandel, J., Sexton, D. J., and Hockberger, R. S. Evaluation and management of suspected sepsis and septic shock in adults. *UpToDate*. Available online: <https://www.uptodate.com/contents/evaluation-and-management-of-suspected-sepsisand-septic-shock-in-adults> (accessed on 29 September 2017), 2016.
- Schulam, P. and Saria, S. Reliable decision support using counterfactual models. In *Advances in Neural Information Processing Systems*, pp. 1697–1708, 2017.
- Shalit, U., Johansson, F. D., and Sontag, D. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pp. 3076–3085, 2017.
- Soleimani, H., Subbaswamy, A., and Saria, S. Treatment-response models for counterfactual reasoning with continuous-time, continuous-valued interventions. *arXiv preprint arXiv:1704.02038*, 2017.
- Spirtes, P., Glymour, C. N., Scheines, R., and Heckerman, D. *Causation, prediction, and search*. MIT press, 2000.
- Tipping, M. E. and Bishop, C. M. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- Tran, D. and Blei, D. M. Implicit causal models for genome-wide association studies. *International Conference on Learning Representations*, 2018.
- Vlachostergios, P. J. and Faltas, B. M. Treatment resistance in urothelial carcinoma: an evolutionary perspective. *Nature Reviews Clinical Oncology*, pp. 1, 2018.
- Wager, S. and Athey, S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 2017.
- Wang, Y. and Blei, D. M. The blessings of multiple causes. *Journal of the American Statistical Association*, (just-accepted):1–71, 2019.

Yoon, J., Jordon, J., and van der Schaar, M. Ganite: Estimation of individualized treatment effects using generative adversarial nets. *International Conference on Learning Representations (ICLR)*, 2018.

Zhang, Y., Bellot, A., and van der Schaar, M. Learning overlapping representations for the estimation of individualized treatment effects. *International Conference on Artificial Intelligence and Statistics*, 2020.