
Deep Learning Methods for the Noniterative Conditional Expectation G-Formula for Causal Inference from Complex Observational Data

Sophia M. Rein*

CAUSALab
Harvard T.H. Chan School of Public Health
Boston, MA 02115
srein@hsph.harvard.edu

Jing Li*

CAUSALab
Harvard T.H. Chan School of Public Health
Boston, MA 02115
jing_li@hsph.harvard.edu

Miguel Hernan†

CAUSALab
Harvard T.H. Chan School of Public Health
Boston, MA 02115
mhernan@hsph.harvard.edu

Andrew Beam†

CAUSALab
Harvard T.H. Chan School of Public Health
Boston, MA 02115
Andrew_Beam@hms.harvard.edu

Abstract

The g-formula can be used to estimate causal effects of sustained treatment strategies using observational data under the identifying assumptions of consistency, positivity, and exchangeability. The non-iterative conditional expectation (NICE) estimator of the g-formula also requires correct estimation of the conditional distribution of the time-varying treatment, confounders, and outcome. Parametric models, which have been traditionally used for this purpose, are subject to model misspecification, which may result in biased causal estimates. Here, we propose a unified deep learning framework for the NICE g-formula estimator that uses multitask recurrent neural networks for estimation of the joint conditional distributions. Using simulated data, we evaluated our model’s bias and compared it with that of the parametric g-formula estimator. We found lower bias in the estimates of the causal effect of sustained treatment strategies on a survival outcome when using the deep learning estimator compared with the parametric NICE estimator in settings with simple and complex temporal dependencies between covariates. These findings suggest that our Deep Learning g-formula estimator may be less sensitive to model misspecification than the classical parametric NICE estimator when estimating the causal effect of sustained treatment strategies from complex observational data.

1 Introduction

The g-formula (or g-computation algorithm formula) identifies the causal effect of sustained treatment strategies under the conditions of consistency, positivity, and exchangeability[9]. Unlike previously proposed methods for causal inference, the g-formula is valid even in the presence of treatment-confounder feedback, i.e., when time-varying confounders are affected by past treatment as commonly occurs in medical settings [2].

A common representation of the g-formula is a non-iterative conditional expectation (NICE) of the outcome weighted by the joint density of covariate history[12]. The NICE g-formula estimator has 2

*Equal contribution.

†Co-senior authors.

stages: (i) estimation of the conditional distributions of time-varying confounders, treatments, and outcome, and (ii) calculation of an integral (or sum) over all possible covariate histories, which is usually approximated via Monte Carlo simulation.

In realistic settings that require adjustment for a high-dimensional set of confounders, nonparametric estimation of the components of the NICE g-formula is usually not possible. Therefore, in real-world applications[6][11][1][4], the NICE g-formula estimator has been based on parametric models such as generalized linear models (GLMs).

Because correct specification of parametric models may be difficult [10][7], an alternative is the use of more flexible deep neural networks (aka "deep learning") that do not require an explicit specification of the functional form of covariates. Specifically, recurrent neural networks, such as long short-term memory (LSTM) models[3], were developed to model longitudinal data, and thus are well suited to model evolving covariate trajectories with potentially long-range dependencies between them. Therefore, LSTMs can be used to jointly predict all covariate trajectories in stage (i) of the NICE g-formula estimator. A previous study [5] found that g-formula estimates based on LSTMs were less biased than those based on GLMs. However, this study implemented GLMs that are too simplistic for most practical applications. Therefore, it is unclear whether LSTMs would similarly outperform the traditional parametric approach when compared with sufficiently complex and richly parameterized GLMs.

Here we describe a deep learning (DL) NICE g-formula estimator that uses LSTMs to estimate the joint distribution of covariates and outcomes. Using simulated data that mimics real-world data, we compare the bias of g-formula estimates between this DL-NICE g-formula estimator and the parametric NICE g-formula estimator when targeting the causal effect of sustained treatment strategies on a survival outcome.

2 Methods

2.1 Simulated data

We simulated a cohort of people with HIV who started antiretroviral treatment for the first time. We used this simulated dataset to mimic a recent study that estimated the effect of integrase inhibitors (INSTI; a class of antiretrovirals) on the risk of cardiovascular events[8] in the HIV-CAUSAL Collaboration, a consortium of observational studies from Europe and North America[6].

The data simulation procedures are described in the supplementary material (see Section A.1). Briefly, we simulated data for each individual at months $k = 0, 1, 2, \dots, K$, where $k = 0$ is baseline and $K = 60$ is the end of the follow-up. For each individual and month, we simulated a time-varying indicator A_k for initiation of an antiretroviral combination regimen containing an INSTI drug during k , a vector of time-varying confounders L_k measured at the start of month k , and a time-varying indicator Y_k of experiencing a cardiovascular event (myocardial infarction, stroke, or invasive cardiovascular procedure) during month k . No individual was lost to follow-up.

We simulated datasets with two different dependencies between variables across time. First, we used a 'simple' time-dependency between covariates: the value of every variable at k does not depend on covariate values before time $k - 1$. Second, we used a 'complex' time-dependency between covariates: the value of every variable at k depends on a function of the entire history of all covariates through k . For each type of dependency, we simulated datasets of 1,000 and 10,000 individuals.

2.2 The g-formula and the plug-in NICE estimator

The risk by K had all individuals received treatment according to an intervention distribution $f^{int}(a_k|Y_k = 0, \bar{l}_k, \bar{a}_{k-1})$ is identified by the g-formula (see equation S21 in supplementary material) under exchangeability, positivity and consistency[9]. For a deterministic treatment strategy (i.e. $f^{int}(a_k|Y_k = 0, \bar{l}_k, \bar{a}_{k-1})$ either equals 0 or 1 for all $(\bar{l}_k, \bar{a}_{k-1})$) the joint density g-formula equals

$E(Y_K^g)$ and is given by:

$$\sum_{\forall \bar{l}_{K-1}} \sum_{k=1}^K P(Y_k = 1 | Y_{k-1} = 0, \bar{L}_{k-1} = \bar{l}_{k-1}, \bar{A}_{k-1} = \bar{a}_{k-1}^g) \times \prod_{s=0}^{k-1} P(Y_s = 0 | Y_{s-1} = 0, \bar{L}_{s-1} = \bar{l}_{s-1}, \bar{A}_{s-1} = \bar{a}_{s-1}^g) f(l_s | Y_s = 0, \bar{l}_{s-1}, \bar{a}_{s-1}^g), \quad (1)$$

where $\bar{X}_k = (X_0, \dots, X_k)$ denotes the history of random variable X , through k . By definition, $Y_0 = 0$ and $\bar{L}_{-1} = \bar{A}_{-1} = \emptyset$.

The plug-in parametric estimator of the NICE g-formula has often been implemented with the following steps:

1. Fit models for each component in Expression 1. Specifically,
 - (a) Fit a pooled (over persons and time) model for the conditional distribution of each confounder in L_k at time k as a function of k , past treatment, and confounder history based on those who are alive at k .
 - (b) Fit a pooled logistic regression model for the probability of experiencing the outcome of interest by time $k + 1$ as a function of k , past treatment, and confounder history based on those who are alive at k .
2. Approximate the integral in expression 1 via Monte Carlo simulation. For each simulation sample $m = 1, \dots, n$, baseline confounders at $k = 0$ are sampled from observed values, and treatment at $k = 0$ is assigned according to the intervention rule. For $k > 0$:
 - (a) Simulate confounders from the fitted models in Step 1(a) using previously simulated confounders and assigned treatment values through time $k - 1$. Assign treatment according to the intervention based on simulated confounders and assigned treatment values through time $k - 1$.
 - (b) Compute the discrete-time hazard of the outcome at time $k + 1$ from the fitted model in Step 1(b) using previously simulated confounders and assigned treatment values through time k .
3. Calculate the average cumulative probability of failure by time K over all generated intervention histories.

2.3 Deep learning NICE algorithm

The DL-NICE g-formula algorithm replaces the parametric models for covariates and outcomes by a single multitask recurrent neural network for joint prediction of all covariates and a recurrent neural network for the outcome. The recurrent neural networks do not require specification of the functional form of the models (i.e., the relationships between covariates), which may be complex and unknown. The DL-NICE algorithm has the following steps depicted in Figure 1:

Step 1: Representation learning by the recurrent neural network

(a) Fit a recurrent neural network (we fit a LSTM in the present study but other recurrent neural networks could be considered) to jointly estimate the conditional distributions of the confounders. The model receives as input all covariate and treatment values up to and including at time $k - 1$ (i.e. the entire history of treatment and confounders). (b) Fit a recurrent neural network (here, LSTM) to estimate the probability of experiencing the outcome of interest at each time $k + 1$

based on those who are alive at k . The model receives as input all covariate and treatment values up to and including at time k ; i.e. the entire history of treatment and confounders).

Step 2: Monte Carlo simulation of covariate trajectories and computation of outcome hazards Monte Carlo simulation is performed for n number of times based on the intervention of interest. The covariate values are sequentially simulated by assuming a normal distribution for the mean μ of the

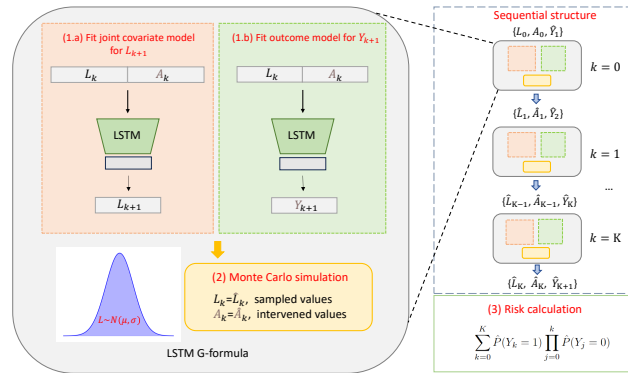


Figure 1: Deep Learning NICE g-formula framework

covariate distributions estimated by the LSTM. The procedure for computation of the hazards is the same as in the classic parametric g-formula algorithm described above.

Step 3: Calculate the cumulative incidence of the outcome Calculate the average cumulative probability of failure by time K over all generated intervention histories as is done in the classical g-formula algorithm.

2.4 Analysis

We investigated the 5-year risk of cardiovascular events under two deterministic static treatment strategies (see section A.3 in the supplementary material for a formal definition of treatment strategies): (1) "initiate ART containing INSTI at time 0 and continue to treat with this combination during the study" (or "always treat with INSTI"), and (2) "initiate ART not containing INSTI at time 0 and continue with this strategy throughout the study" (or "never treat with INSTI"). We first estimated the absolute risk under no intervention on treatment (the "natural course") and under the "always treat" and "never treat" strategies and then estimated the causal contrast (risk ratio and risk difference) comparing the "always treat" and "never treat" strategies using first the classical parametric g-formula estimator and then our DL-NICE estimator. We then evaluated the absolute bias in the estimated risks over time under each strategy and the bias in the estimates of the causal effect (risk ratio and risk difference) under both methods and averaged the bias across all time points. We obtained the 'ground truth' risks from simulated datasets with 1 million observations for the 'simple' and 'complex' time dependency scenario. Let $\hat{R}_k^{\text{method}}$ denote the estimated risk at time k obtained using either the DL-NICE estimator or the classical parametric g-formula estimator, and let R_k^{true} represent the 'ground truth' risk at time k . The bias in the estimated risk at time k is then defined as:

$$\text{Bias}(\hat{R}_k^{\text{method}}) = \hat{R}_k^{\text{method}} - R_k^{\text{true}}.$$

Similarly, for the causal effect estimates, the risk ratio (RR) and risk difference (RD), the biases were calculated as the differences between the estimates obtained from the methods and the corresponding ground truth values. For the risk ratio, the bias is defined as:

$$\text{Bias}(\hat{\text{RR}}_k^{\text{method}}) = \hat{\text{RR}}_k^{\text{method}} - \text{RR}_k^{\text{true}}.$$

For the risk difference, the bias is defined as:

$$\text{Bias}(\hat{\text{RD}}_k^{\text{method}}) = \hat{\text{RD}}_k^{\text{method}} - \text{RD}_k^{\text{true}}.$$

We conducted these analyses in both the dataset with the 'simple' and 'complex' time dependency scenario and for both sample sizes (1,000 and 10,000). To optimize hyperparameter values for each LSTM, we used the 'hyperopt' Python package with 50 trials for each search. The hyperparameter search strategy and settings that we used for the LSTMs are shown in the supplementary material in section A.4.

For the classic parametric g-formula, we fit different parametric models. For the simple time dependency scenario, we fit parametric models for the covariates and outcome that were equal to the data generating process (i.e. each covariate depends on covariate values at time $k-1$; see supplementary material A.1) except for not including the unmeasured covariate U . For the complex time dependency scenario, we fit two parametric models that differed in the complexity of covariate histories included:

1. *Least flexible*: We fit models for L_k and Y_k that include terms for the lagged value of L_k and A_k (i.e. the values of L_{k-1} and A_{k-1}).

2. *Moderately flexible*: We fit models for L_k and Y_k that include terms for the lagged value as well as the lagged cumulative average of the covariate values from $k = 0$ up to $k - 1$.

3 Results

3.1 Simple time-dependency scenario

Figure 2 shows the ground truth and predicted (using both DL-NICE and parametric NICE) risk of the outcome over time under no intervention (the natural course) and suggests that the risk predicted using the DL-NICE estimator is closer to the ground truth risk. Tables 1 and 2 show the absolute bias averaged across all 60 time points in the estimates of the risk under natural course, the 'always treat' and 'never treat' interventions and the bias in the risk ratio and risk difference comparing the always and never treat interventions (see the supplementary material, section A.5.1, for the risks under intervention, effect estimates and bias over time). In both sample sizes, the bias in the risk estimates was lower when using the DL-NICE method compared to the parametric method. However,

when using the DL-NICE method, the bias in the effect estimates (risk difference and risk ratio) was somewhat higher in the smaller dataset with sample size of 1,000, while it was lower compared to the parametric NICE method in the larger dataset.

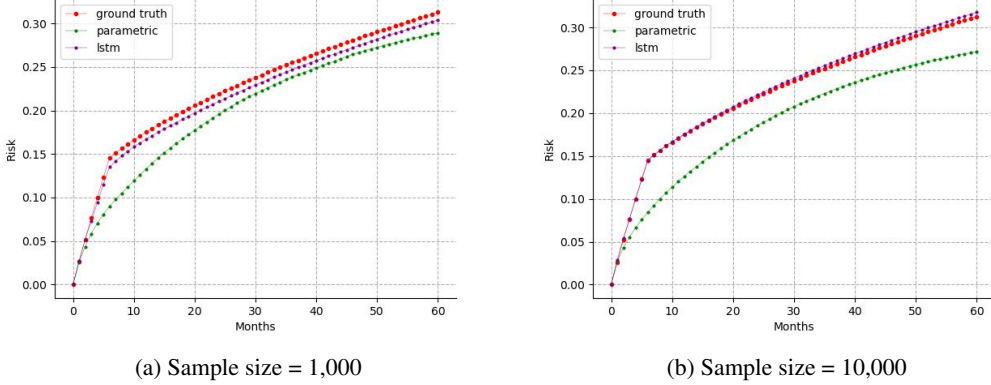


Figure 2: Ground truth and estimated risk under the 'natural course' using the parametric and DL-NICE methods on 'simple' time-dependency data

Method	Bias in risk estimates (\downarrow is better)			Bias in causal effect estimates (\downarrow is better)	
	Natural course	Always treat	Never treat	Risk difference	Risk ratio
Parametric	0.024	0.036	0.014	0.022	0.096
DL-NICE (Ours)	0.008	0.029	0.004	0.032	0.156


Table 1: Bias in risk and causal effect estimates (comparing always vs. never treat; averaged over all 60 time points) using the parametric and DL-NICE methods on 'simple' time-dependency data with 1,000 individuals.

Method	Bias in risk estimates (\downarrow is better)			Bias in causal effect estimates (\downarrow is better)	
	Natural course	Always treat	Never treat	Risk difference	Risk ratio
Parametric	0.035	0.042	0.029	0.013	0.047
DL-NICE (Ours)	0.003	0.003	0.004	0.006	0.031


Table 2: Bias in risk and causal effect estimates (comparing always vs. never treat; averaged over all 60 time points) using the parametric and DL-NICE methods on 'simple' time-dependency data with 10,000 individuals.

3.2 Complex time-dependency scenario

Figure 3 shows the ground truth and predicted outcome risk over time under the natural course on the complex time-dependency data, again suggesting that the risk predicted using the DL-NICE estimator is closer to the ground truth risk than that predicted using parametric NICE. Tables 3 and 4 show the absolute bias averaged over all 60 time points in the risk and effect estimates (risk ratio and risk difference) in the datasets with sample sizes of 1,000 and 10,000, respectively (see section A.5.2 in the supplementary material for the estimated risks under the interventions, effect estimates and bias over time). In both sample sizes, the bias in the risk estimates was highest when using the parametric g-formula with the least flexible models for the covariates and the outcome (covariate history includes lagged values only), followed by the parametric g-formula with moderately flexible models (covariate history includes lagged values and the lag of the cumulative average of values over time) and was lowest when using the DL-NICE estimator. However, this difference in bias comparing the parametric and DL-NICE estimator was more pronounced in the larger dataset. The bias in the effect estimates was also lower when using the DL-NICE method for both sample sizes, except for the risk ratio bias in the larger sample size.



(a) Sample size = 1,000



(b) Sample size = 10,000

Figure 3: Ground truth and estimated risk under the 'natural course' using the parametric and DL-NICE methods on 'complex' time-dependency data

Method	Bias in risk estimates (\downarrow is better)			Bias in causal effect estimates (\downarrow is better)	
	Natural course	Always treat	Never treat	Risk difference	Risk ratio
Parametric 1*	0.066	0.079	0.053	0.026	0.175
Parametric 2**	0.037	0.093	0.036	0.131	0.661
DL-NICE (Ours)	0.008	0.005	0.014	0.019	0.112

Table 3: Bias in risk and causal effect estimates (comparing always vs. never treat; averaged over all 60 time points) using the parametric and DL-NICE methods on 'complex' time-dependency data with 1,000 individuals. * Parametric 1: Least flexible, ** Parametric 2: Moderately flexible.

Method	Bias in risk estimates (\downarrow is better)			Bias in causal effect estimates (\downarrow is better)	
	Natural course	Always treat	Never treat	Risk difference	Risk ratio
Parametric 1*	0.074	0.081	0.065	0.016	0.079
Parametric 2**	0.051	0.102	0.023	0.124	0.672
DL-NICE (Ours)	0.006	0.012	0.003	0.014	0.089

Table 4: Bias in risk and causal effect estimates (comparing always vs. never treat; averaged over all 60 time points) using the parametric and DL-NICE methods on 'complex' time-dependency data with 10,000 individuals. * Parametric 1: Least flexible, ** Parametric 2: Moderately flexible.

4 Conclusion

We introduced a deep learning NICE g-formula estimator for causal inference from complex observational data that uses LSTMs to estimate the joint distribution of covariates and outcomes over time. This approach eliminates the need to fit separate parametric models for the conditional distribution of each covariate and specification of the functional form of covariate relationships. Our results showed that the DL-NICE estimator generally achieved lower bias in risk and causal effect estimates compared to parametric NICE, particularly in settings with complex temporal dependencies. Even in simpler settings, the DL-NICE estimator generally showed less bias.

However, the reduction in bias under the DL-NICE estimator was not uniform across scenarios, with the bias in the effect estimates being marginally higher in the 'simple' data with smaller sample size when using DL-NICE compared to parametric NICE. This finding suggests that the relative performance of the DL-NICE estimator may be sensitive to sample size and the complexity of the underlying data structure. Future work should explore statistical inference under the DL-NICE method, the conditions under which it may provide the most benefit over parametric NICE and model selection strategies.

A Supplementary Material

A.1 Data simulation design

In our motivating example, baseline ($k = 0$) was defined as meeting the following eligibility criteria for the first time: at least 18 years old, diagnosis of HIV-1 infection, starting ART for the first

time, an HIV-RNA measurement while treatment-naïve that had to be detectable (>50 copies/ml) and no history of a cardiovascular event (myocardial infarction, stroke, or invasive cardiovascular procedure) or cancer. At baseline, L_0 included sex (binary, 1=male, 0=female), age (continuous) and smoking status (categorical, 0=never smoked, 1=currently smoking, 2=ex-smoker). The time-varying confounders in L_k included a person's HIV-RNA viral load, CD4 cell count, and an indicator of being overweight or obese at time k based on a person's body mass index (BMI >25). The binary treatment A_k is an indicator for whether a person is using an ART regimen with an INSTI ($A_k = 1$) or using an ART regimen without an INSTI ($A_k = 0$). The outcome Y_k is an indicator of whether the person experienced a cardiovascular event ($Y_k = 1$). For each individual, data were simulated for up to 60 observation months. If and when an individual experienced an event ($Y_{k+1} = 1$), follow up ends and there will be no subsequent observation months for that individual. The data generating processes in both the 'simple' and 'complex' time-dependency datasets are detailed below.

A.1.1 Data generating processes for the "simple time dependency" dataset

We first drew an unmeasured variable U and simulated the time-fixed covariates sex (L_{10}), age (L_{20}) and smoking (L_{30}). We then generated the time-varying covariates in L_k and A_k and the survival outcome Y_k at each time point. To generate increasing and decreasing trends in the time-varying covariates CD4 count and HIV-RNA, respectively, that would be observed over the first six months after ART is initiated, we used different data generating equations for the first six observation months and for the months after for these variables. For all time-varying covariates we also generated values at time $k = -1$ (pre-baseline), so that these could be used to generate the covariate values at time $k = 0$ (baseline). The data generating functions for each variable are shown below.

Baseline Covariates

- $Age \sim \text{TruncatedNormal}(\mu = 50, \sigma = 12, a = 18, b = 80)$
- $Sex \sim \text{Ber}(p = 0.8)$
- $Smoking \sim \text{Multinomial}(1, [0.45, 0.4, 0.15])$
- $U \sim U(0.0, 1.0)$

Time-Varying Covariates

CD4 Count:

Pre-baseline value at time $k = -1$:

$$CD4_{k=-1} \sim \text{TruncatedNormal}(\mu = 450, \sigma = 100, a = 350, b = 800) \quad (S1)$$

Time $0 \leq k \leq 5$:

$$\begin{aligned} CD4_k &\sim \text{TruncatedNormal}(\mu, \sigma = 100, a = 350, b = 800) \\ \mu &\sim -0.8 \cdot U + 0.7 \cdot Sex - 0.8 \cdot Age - 0.05 \cdot Age^2 + 0.6 \cdot Smoking \\ &\quad + 2 \cdot CD4_{k-1} + 0.005 \cdot CD4_{k-1}^2 - 1 \cdot RNA_{k-1} - 0.1 \cdot HighBMI_{k-1} \quad (S2) \\ &\quad + 0.05 \cdot Insti_{k-1} + 0.1 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\ &\quad + 0.08 \cdot HighBMI_{k-1} \cdot Sex - 0.1 \cdot CD4_{k-1} \cdot Age \end{aligned}$$

At time $k = 0$, the coefficient for $Insti_{k-1}$ was set to 0 as there is no treatment before baseline.

From time $k \geq 6$:

$$\begin{aligned} CD4_k &\sim \text{TruncatedNormal}(\mu, \sigma = 80, a = 400, b = 800) \\ \mu &\sim -0.8 \cdot U + 0.7 \cdot Sex - 0.8 \cdot Age - 0.05 \cdot Age^2 + 0.6 \cdot Smoking \\ &\quad + 1.8 \cdot CD4_{k-1} + 0.008 \cdot CD4_{k-1}^2 - 1 \cdot RNA_{k-1} - 0.1 \cdot HighBMI_{k-1} \quad (S3) \\ &\quad + 0.05 \cdot Insti_{k-1} + 0.1 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\ &\quad + 0.08 \cdot HighBMI_{k-1} \cdot Sex - 0.1 \cdot CD4_{k-1} \cdot Age \end{aligned}$$

HIV RNA:

Pre-baseline value at time $k = -1$:

$$RNA_{k=-1} \sim \text{TruncatedNormal}(\mu = 60, \sigma = 30, a = 40, b = 90) \quad (S4)$$

Time $0 \leq k \leq 5$:

$$\begin{aligned}
RNA_k &\sim \text{TruncatedNormal}(\mu, \sigma = 30, a = 40, b = 80) \\
\mu &\sim 0.5 \cdot U + 0.5 \cdot Sex + 0.3 \cdot Age + 0.006 \cdot Age^2 + 0.8 \cdot Smoking \\
&\quad - 0.5 \cdot CD4_{k-1} - 0.0001 \cdot CD4_{k-1}^2 + 2 \cdot RNA_{k-1} + 0.3 \cdot HighBMI_{k-1} \\
&\quad + 0.05 \cdot INSTI_{k-1} + 0.08 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\
&\quad + 0.1 \cdot HighBMI_{k-1} \cdot Sex - 0.01 \cdot CD4_{k-1} \cdot Age
\end{aligned} \tag{S5}$$

From time $k \geq 6$:

$$\begin{aligned}
RNA_k &\sim \text{TruncatedNormal}(\mu, \sigma = 20, a = 20, b = 70) \\
\mu &\sim 0.5 \cdot U + 0.5 \cdot Sex + 0.3 \cdot Age + 0.006 \cdot Age^2 + 0.8 \cdot Smoking \\
&\quad - 0.5 \cdot CD4_{k-1} - 0.0001 \cdot CD4_{k-1}^2 + 1.8 \cdot RNA_{k-1} + 0.3 \cdot HighBMI_{k-1} \\
&\quad + 0.05 \cdot INSTI_{k-1} + 0.08 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\
&\quad + 0.1 \cdot HighBMI_{k-1} \cdot Sex - 0.01 \cdot CD4_{k-1} \cdot Age
\end{aligned} \tag{S6}$$

BMI > 25:

$$\begin{aligned}
HighBMI_k &\sim \text{Ber}(\text{expit}(\mu)) \\
\mu &\sim -8 - 2 \cdot U + 0.03 \cdot Sex + 0.01 \cdot Age + 0.0001 \cdot Age^2 + 0.04 \cdot Smoking \\
&\quad - 0.0001 \cdot CD4_{k-1} + 0.001 \cdot RNA_{k-1} + 10 \cdot HighBMI_{k-1} \\
&\quad + 5 \cdot INSTI_{k-1} + 0.001 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\
&\quad + 0.004 \cdot HighBMI_{k-1} \cdot Sex + 0.00001 \cdot CD4_{k-1} \cdot Age
\end{aligned} \tag{S7}$$

Treatment

$$\begin{aligned}
INSTI_k &\sim \text{Ber}(\text{expit}(\mu)) \\
\mu &\sim -4.5 + 0.5 \cdot Sex + 0.01 \cdot Age + 0.0001 \cdot Age^2 + 0.1 \cdot Smoking \\
&\quad + 0.001 \cdot CD4_k + 0.01 \cdot RNA_k - 7 \cdot HighBMI_k \\
&\quad + 10 \cdot INSTI_{k-1} + 0.0001 \cdot HighBMI_k \cdot RNA_k \\
&\quad + 0.001 \cdot HighBMI_k \cdot Sex + 0.00001 \cdot CD4_k \cdot Age
\end{aligned} \tag{S8}$$

Outcome

$$\begin{aligned}
Y_k &\sim \text{Ber}(\text{expit}(\mu)) \\
\mu &= -0.08 \cdot U + 0.005 \cdot Sex + 0.015 \cdot Age + 0.00000005 \cdot Age^2 + 0.025 \cdot Smoking \\
&\quad - 0.015 \cdot CD4_k + 0.03 \cdot RNA_k + 0.0000004 \cdot RNA_k^2 + 0.1 \cdot HighBMI_k \\
&\quad + 0.09 \cdot INSTI_k
\end{aligned} \tag{S9}$$

A.1.2 Data generating processes for the "complex time dependency" dataset

The data generation functions for the unmeasured confounder U and the baseline covariates in L_0 were the same as in the simple time dependency scenario. The time-varying covariates in L_k , treatment A_k and the outcome Y_k were generated to depend on a function of the entire history of covariates and treatment, defined as the cumulative average of values over the previous 6, 7-24 and over >24 time points from the current time point k for L_k and the cumulative average over all previous time points after baseline for A_k . As in the 'simple' time-dependency dataset, we used different data generating equations for the first six observation months and for the months after for the time-varying covariates CD4 count and HIV-RNA. For all time-varying covariates we also generated values for 30 pre-baseline observation months ($k = -30$ to $k = -1$) to be used to generate the covariate values from time $k = 0$ (baseline) onwards. The data generating functions for each covariate L_k , A_k and the outcome Y_k are shown below.

Time-Varying Covariates

CD4 Count:

Initial value at time $k = -30$:

$$CD4_{k=-30} \sim \text{TruncatedNormal}(\mu = 450, \sigma = 100, a = 350, b = 800) \tag{S10}$$

Value evolution from $k = -30$ to $k = -1$:

$$CD4_k = CD4_{k-1} \cdot 0.995 \quad \text{for } k = -29 \quad \text{to } k = -1 \quad (\text{S11})$$

Time $0 \leq k \leq 5$:

$CD4_k \sim \text{TruncatedNormal}(\mu, \sigma = 100, a = 350, b = 800)$

$$\begin{aligned} \mu = & -2 \cdot U + 0.1 \cdot Sex - 1 \cdot Age - 0.05 \cdot Age^2 + 0.3 \cdot Smoking \\ & + 1 \cdot CD4_{k-1} + 0.005 \cdot CD4_{k-1}^2 + 0.8 \cdot \text{mean}(CD4_{k-6:k-1}) \\ & + 0.6 \cdot \text{mean}(CD4_{k-24:k-7}) + 0.5 \cdot \text{mean}(CD4_{k-25:k=-30}) \\ & - 0.8 \cdot RNA_{k-1} - 0.5 \cdot \text{mean}(RNA_{k-6:k-1}) - 0.3 \cdot \text{mean}(RNA_{k-24:k-7}) \\ & - 0.1 \cdot \text{mean}(RNA_{k-25:k=-30}) - 0.1 \cdot HighBMI_{k-1} - 0.08 \cdot \text{mean}(HighBMI_{k-6:k-1}) \\ & - 0.04 \cdot \text{mean}(HighBMI_{k-24:k-7}) - 0.02 \cdot \text{mean}(HighBMI_{k-25:k=-30}) \\ & + 0.05 \cdot INSTI_{k-1} + 0.02 \cdot \text{mean}(INSTI_{k-1:k=0}) + 0.1 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\ & + 0.08 \cdot HighBMI_{k-1} \cdot Sex - 0.1 \cdot CD4_{k-1} \cdot Age \end{aligned} \quad (\text{S12})$$

From time $k \geq 6$:

$CD4_k \sim \text{TruncatedNormal}(\mu, \sigma = 80, a = 400, b = 800)$

$$\begin{aligned} \mu = & -2 \cdot U + 0.1 \cdot Sex - 1 \cdot Age - 0.05 \cdot Age^2 + 0.3 \cdot Smoking \\ & + 1.8 \cdot CD4_{k-1} + 0.008 \cdot CD4_{k-1}^2 + 1 \cdot \text{mean}(CD4_{k-6:k-1}) \\ & + 0.5 \cdot \text{mean}(CD4_{k-24:k-7}) + 0.2 \cdot \text{mean}(CD4_{k-25:k=-30}) \\ & - 0.8 \cdot RNA_{k-1} - 0.5 \cdot \text{mean}(RNA_{k-6:k-1}) - 0.3 \cdot \text{mean}(RNA_{k-24:k-7}) \\ & - 0.1 \cdot \text{mean}(RNA_{k-25:k=-30}) - 0.1 \cdot HighBMI_{k-1} - 0.08 \cdot \text{mean}(HighBMI_{k-6:k-1}) \\ & - 0.04 \cdot \text{mean}(HighBMI_{k-24:k-7}) - 0.02 \cdot \text{mean}(HighBMI_{k-25:k=-30}) \\ & + 0.05 \cdot INSTI_{k-1} + 0.02 \cdot \text{mean}(INSTI_{k-1:k=0}) + 0.1 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\ & + 0.08 \cdot HighBMI_{k-1} \cdot Sex - 0.1 \cdot CD4_{k-1} \cdot Age \end{aligned} \quad (\text{S13})$$

HIV RNA:

Initial value at time $k = -30$:

$$RNA_{k=-30} \sim \text{TruncatedNormal}(\mu = 60, \sigma = 30, a = 40, b = 90) \quad (\text{S14})$$

Value evolution from $k = -30$ to $k = -1$:

$$RNA_k = RNA_{k-1} \cdot 1.01 \quad \text{for } k = -29 \quad \text{to } k = -1 \quad (\text{S15})$$

Time $0 \leq k \leq 5$:

$RNA_k \sim \text{TruncatedNormal}(\mu, \sigma = 30, a = 40, b = 80)$

$$\begin{aligned} \mu = & 2 \cdot U + 1 \cdot Sex + 1.5 \cdot Age + 0.006 \cdot Age^2 + 0.5 \cdot Smoking \\ & - 0.5 \cdot CD4_{k-1} - 0.0001 \cdot CD4_{k-1}^2 - 0.2 \cdot \text{mean}(CD4_{k-6:k-1}) \\ & - 0.05 \cdot \text{mean}(CD4_{k-24:k-7}) - 0.01 \cdot \text{mean}(CD4_{k-25:k=-30}) \\ & + 6 \cdot RNA_{k-1} + 5 \cdot \text{mean}(RNA_{k-6:k-1}) + 3 \cdot \text{mean}(RNA_{k-24:k-7}) \\ & + 2 \cdot \text{mean}(RNA_{k-25:k=-30}) + 0.3 \cdot HighBMI_{k-1} + 0.1 \cdot \text{mean}(HighBMI_{k-6:k-1}) \\ & + 0.06 \cdot \text{mean}(HighBMI_{k-24:k-7}) + 0.03 \cdot \text{mean}(HighBMI_{k-25:k=-30}) \\ & + 0.05 \cdot INSTI_{k-1} + 0.01 \cdot \text{mean}(INSTI_{k-1:k=0}) + 0.08 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\ & + 0.1 \cdot HighBMI_{k-1} \cdot Sex - 0.01 \cdot CD4_{k-1} \cdot Age \end{aligned} \quad (\text{S16})$$

From time $k \geq 6$:

$$RNA_k \sim \text{TruncatedNormal}(\mu, \sigma = 20, a = 20, b = 70)$$

$$\begin{aligned} \mu = & 2 \cdot U + 1 \cdot Sex + 1.5 \cdot Age + 0.006 \cdot Age^2 + 0.5 \cdot Smoking \\ & - 0.5 \cdot CD4_{k-1} - 0.0001 \cdot CD4_{k-1}^2 - 0.2 \cdot \text{mean}(CD4_{k-6:k-1}) \\ & - 0.05 \cdot \text{mean}(CD4_{k-24:k-7}) - 0.01 \cdot \text{mean}(CD4_{k-25:k=-30}) \\ & + 6 \cdot RNA_{k-1} + 5 \cdot \text{mean}(RNA_{k-6:k-1}) + 3 \cdot \text{mean}(RNA_{k-24:k-7}) \\ & + 2 \cdot \text{mean}(RNA_{k-25:k=-30}) + 0.3 \cdot HighBMI_{k-1} + 0.1 \cdot \text{mean}(HighBMI_{k-6:k-1}) \\ & + 0.06 \cdot \text{mean}(HighBMI_{k-24:k-7}) + 0.03 \cdot \text{mean}(HighBMI_{k-25:k=-30}) \\ & + 0.05 \cdot INSTI_{k-1} + 0.01 \cdot \text{mean}(INSTI_{k-1:k=0}) + 0.08 \cdot HighBMI_{k-1} \cdot RNA_{k-1} \\ & + 0.1 \cdot HighBMI_{k-1} \cdot Sex - 0.01 \cdot CD4_{k-1} \cdot Age \end{aligned} \quad (S17)$$

BMI > 25:

$$HighBMI_k \sim \text{Ber}(\text{expit}(\mu))$$

$$\begin{aligned} \mu = & -6.5 - 1 \cdot U - 0.6 \cdot Sex + 0.01 \cdot Age + 0.0001 \cdot Age^2 - 0.04 \cdot Smoking \\ & - 0.0001 \cdot CD4_{k-1} + 0.000001 \cdot CD4_{k-1}^2 - 0.001 \cdot \text{mean}(CD4_{k-6:k-1}) \\ & - 0.0001 \cdot \text{mean}(CD4_{k-24:k-7}) - 0.001 \cdot \text{mean}(CD4_{k-25:k=-30}) \\ & + 0.01 \cdot RNA_{k-1} + 0.01 \cdot \text{mean}(RNA_{k-6:k-1}) + 0.007 \cdot \text{mean}(RNA_{k-24:k-7}) \\ & + 0.006 \cdot \text{mean}(RNA_{k-25:k=-30}) + 4.5 \cdot HighBMI_{k-1} + 3 \cdot \text{mean}(HighBMI_{k-6:k-1}) \\ & + 1.6 \cdot \text{mean}(HighBMI_{k-24:k-7}) + 1 \cdot \text{mean}(HighBMI_{k-25:k=-30}) \\ & + 2 \cdot INSTI_{k-1} + 1 \cdot \text{mean}(INSTI_{k-1:k=0}) \end{aligned} \quad (S18)$$

Treatment

$$INSTI_k \sim \text{Ber}(\text{expit}(\mu))$$

$$\begin{aligned} \mu = & -4 + 0.5 \cdot Sex + 0.05 \cdot Age + 0.00005 \cdot Age^2 + 0.2 \cdot Smoking \\ & - 0.001 \cdot CD4_k + 0.0000001 \cdot CD4_k^2 - 0.0001 \cdot \text{mean}(CD4_{k-6:k-1}) \\ & + 0.001 \cdot RNA_k + 0.0003 \cdot \text{mean}(RNA_{k-6:k-1}) \\ & - 3 \cdot HighBMI_k - 2 \cdot \text{mean}(HighBMI_{k-6:k-1}) \\ & - 1.3 \cdot \text{mean}(HighBMI_{k-24:k-7}) - 0.8 \cdot \text{mean}(HighBMI_{k-25:k=-30}) \\ & + 6 \cdot INSTI_{k-1} + 4 \cdot \text{mean}(INSTI_{k-1:k=0}) \end{aligned} \quad (S19)$$

Outcome

$$Y_k \sim \text{Ber}(\text{expit}(\mu))$$

$$\begin{aligned} \mu = & -0.05 \cdot U + 0.007 \cdot Sex + 0.02 \cdot Age + 0.00000005 \cdot Age^2 + 0.03 \cdot Smoking \\ & - 0.009 \cdot CD4_k - 0.008 \cdot \text{mean}(CD4_{k-6:k-1}) - 0.006 \cdot \text{mean}(CD4_{k-24:k-7}) \\ & - 0.004 \cdot \text{mean}(CD4_{k-25:k=-30}) + 0.045 \cdot RNA_k + 0.0000004 \cdot RNA_k^2 \\ & + 0.03 \cdot \text{mean}(RNA_{k-6:k-1}) + 0.025 \cdot \text{mean}(RNA_{k-24:k-7}) + 0.02 \cdot \text{mean}(RNA_{k-25:k=-30}) \\ & + 0.14 \cdot HighBMI_k + 0.11 \cdot \text{mean}(HighBMI_{k-6:k-1}) + 0.08 \cdot \text{mean}(HighBMI_{k-24:k-7}) \\ & + 0.06 \cdot \text{mean}(HighBMI_{k-25:k=-30}) + 0.13 \cdot INSTI_k + 0.11 \cdot \text{mean}(INSTI_{k-1:k=0}) \end{aligned} \quad (S20)$$

A.2 Identifying assumptions and the g-formula

Let Y_k^g and L_k^g denote the counterfactual outcome and vector of confounders, respectively, at time k had an individual followed a deterministic strategy g ($\forall k = 1, \dots, K$). For simplicity, we will again ignore censoring. Consider the following assumptions for a particular strategy g :

1. *Exchangeability*:

$$(Y_{k+1}^g, \dots, Y_k^g) \perp (A_k) | \bar{L}_k = \bar{l}_k, \bar{A}_{k-1} = \bar{a}_{k-1}^g, Y_k = 0$$

2. *Positivity*:

$$f_{\bar{L}_k, \bar{A}_{k-1}, Y_k}(\bar{l}_k, \bar{a}_{k-1}^g, 0, 0) > 0 \Rightarrow f_{A_k | \bar{L}_k, \bar{A}_{k-1}, Y_k}(a_k^g | \bar{l}_k, \bar{a}_{k-1}^g, 0, 0) > 0$$

3. *Consistency*: If $\bar{A}_k = \bar{A}_k^g$ then \bar{Y}_{k+1}^g and $\bar{L}_k = \bar{L}_k^g$.

Let G be the set of all deterministic interventions g . Provided that these assumptions hold for a subset of $g \in G$ that are observable under $f^{int}(a_k|Y_k = 0, \bar{l}_k, \bar{a}_{k-1})(\forall k = 0, \dots, K-1)$ then the risk by K had all subjects been assigned treatment according to $f^{int}(a_k|Y_k = 0, \bar{l}_k, \bar{a}_{k-1})$ can be estimated from the observational data and written as the g-formula, an expectation weighted by the joint density of covariates:

$$\sum_{\forall \bar{a}_{K-1} \forall \bar{l}_{K-1}} \sum_{k=1}^K P(Y_k = 1|Y_{k-1} = 0, \bar{L}_{k-1} = \bar{l}_{k-1}, \bar{A}_{k-1} = \bar{a}_{k-1}) \times \prod_{s=0}^{k-1} P(Y_s = 0|Y_{s-1} = 0, \bar{L}_{s-1} = \bar{l}_{s-1}, \bar{A}_{s-1} = \bar{a}_{s-1}) f(l_s|Y_s = 0, \bar{l}_{s-1}, \bar{a}_{s-1}) f^{int}(a_s|Y_s = 0, \bar{l}_s, \bar{a}_{s-1}) \quad (\text{S21})$$

where $P(Y_k = 1|Y_{k-1} = 0, \bar{L}_{k-1} = \bar{l}_{k-1}, \bar{A}_{k-1} = \bar{a}_{k-1})$ and $f(l_k|Y_{k-1} = 0, \bar{l}_{k-1}, \bar{a}_{k-1})$ are the observed discrete-time hazards of the outcome and the joint density of the confounders at time k , respectively, conditional on past treatment, confounders and survival through time $k-1$ (as we are ignoring censoring in this study, we are not additionally conditioning on being uncensored through time k but this is typically included in the g-formula for survival outcomes). For notational simplicity, the above expression assumes that all covariates are discrete; otherwise, the sum would be replaced by an integral.

A.3 Treatment strategies and causal contrast

A treatment strategy is a rule that assigns treatment at each time k as an independent draw from an intervention distribution $f^{int}(a_k|Y_k = 0, \bar{l}_k, \bar{a}_{k-1})$ that may, at most, depend on $(\bar{l}_k, \bar{a}_{k-1})$, which is a realization of $(\bar{L}_k, \bar{A}_{k-1})$. Treatment strategies can be either deterministic or random. A treatment strategy is deterministic if at each time point $k(k = 0, \dots, K)$, $f^{int}(a_k|Y_k = 0, \bar{l}_k, \bar{a}_{k-1})$ either equals 0 or 1 for all $(\bar{l}_k, \bar{a}_{k-1})$. Otherwise, the treatment strategy is random. We denote g as a deterministic treatment strategy, and $a_k^g = g_k(\bar{a}_{k-1}^g, \bar{l}_k)$ the value of treatment assigned at time k under $g(\forall k = 0, \dots, K-1)$. In the present study, we consider a deterministic static treatment strategy. A treatment strategy is static if the rule for assigning treatment at each time point does not depend on past covariates. The treatment strategy in our study is “initiate ART containing INSTI at time 0 and continue to treat with this combination during the study” (or “always treat”), which corresponds to $a_k^g = 1$, for all k and for any $(\bar{a}_{k-1}, \bar{l}_k)$. We estimate the causal effect on the risk (cumulative incidence) by time K of the “always treat” treatment strategy compared to a “never treat” strategy ($a_k^g = 0$, for all k and for any $(\bar{a}_{k-1}, \bar{l}_k)$), which in our motivating example is defined as “initiate ART *not* containing INSTI at time 0 and continue with this strategy during the study”. That is, we estimate a contrast in the risk had all individuals been assigned treatment (i.e. initiated ART containing INSTI) vs. had no one been assigned this treatment strategy (had no one taken ART containing INSTI).

A.4 Hyperparameter search and settings

Hyperparameter	Feature dimension	Hidden size	Layers num	Dropout rate	Learning rate	Batch size	Weight decay	Max epochs	Patience epoch
Covariate LSTM	{128, 512}	{64, 128, 256}	{2, 3, 4}	Uniform(0, 0.5)	log-uniform(1e-5, 1e-3)	{512, 1024}	log-uniform(1e-6, 1e-3)	5000	50
Outcome LSTM	{128, 512}	{64, 128, 256}	{2, 3, 4}	Uniform(0, 0.5)	log-uniform(1e-5, 1e-3)	{512, 1024}	log-uniform(1e-6, 1e-3)	1000	50

Table S1: Hyperparameter search space for the LSTMs in the DL-NICE estimator.

Model type	Hidden size of LSTM	Layers num	Feature dimension	Dropout	Epoch (best)	Patience parameter	Batch size	Learning rate	Weight decay
Sample size = 10,000									
Covariate LSTM	64	3	128	0.19	5000 (588)	50	1024	4.1e-4	3.3e-6
Outcome LSTM	64	4	128	0.15	1000 (1000)	50	512	4.7e-5	7.4e-4
Sample size = 1,000									
Covariate LSTM	64	2	512	0.45	5000 (716)	50	512	3.8e-4	4.1e-5
Outcome LSTM	128	2	128	0.23	1000 (455)	50	512	6.7e-4	3.6e-5


Table S2: Hyperparameter settings for LSTMs for 'simple' time-dependency dataset with sample sizes 10,000 and 1,000.

Model type	Hidden size of LSTM	Layers num	Feature dimension	Dropout	Epoch (best)	Patience parameter	Batch size	Learning rate	Weight decay
Sample size = 10,000									
Covariate LSTM	64	2	128	0.36	5000 (558)	50	512	4.4e-4	4.7e-5
Outcome LSTM	256	2	512	0.30	1000 (712)	50	1024	2.4e-5	4.4e-4
Sample size = 1,000									
Covariate LSTM	64	2	512	0.36	5000 (2239)	50	1024	1.8e-4	3.1e-4
Outcome LSTM	128	2	512	0.19	1000 (442)	50	1024	5.5e-4	1.4e-6


Table S3: Hyperparameter settings for LSTMs for 'complex' time-dependency dataset with sample sizes 10,000 and 1,000.

A.5 Estimated risks and bias in risk and effect estimates over time

A.5.1 'Simple' time-dependency scenario




(a) Risk under always treat




(b) Risk under never treat


Figure S1: Ground truth and estimated risks under interventions in dataset with 1,000 individuals



(a) Bias in risk under natural course




(b) Bias in risk under always treat




(c) Bias in risk under never treat

Figure S2: Bias in risk in dataset with 1,000 individuals




(a) Bias in risk ratio




(b) Bias in risk difference

Figure S3: Bias in effect estimates (always vs. never treat interventions) in 1,000 sample dataset




(a) Risk under always treat




(b) Risk under never treat


Figure S4: Ground truth and estimated risks under interventions in dataset with 10,000 individuals



(a) Bias in risk under natural course




(b) Bias in risk under always treat




(c) Bias in risk under never treat

Figure S5: Bias in risk in dataset with 10,000 individuals




(a) Bias in risk ratio




(b) Bias in risk difference

Figure S6: Bias in effect estimates (always vs. never treat interventions) in 10,000 sample dataset

A.5.2 'Complex' time-dependency scenario




(a) Risk under always




(b) Risk under never treat


Figure S7: Ground truth and estimated risks under interventions in dataset with 1,000 individuals



(a) Bias in risk under natural course




(b) Bias in risk under always treat




(c) Bias in risk under never treat

Figure S8: Bias in risk in dataset with 1,000 individuals




(a) Bias in risk ratio




(b) Bias in risk difference

Figure S9: Bias in effect estimates (always vs. never treat interventions) in 1,000 sample dataset




(a) Risk under always treat




(b) Risk under never treat


Figure S10: Ground truth and estimated risks under interventions in dataset with 10,000 individuals



(a) Bias in risk under natural course




(b) Bias in risk under always treat




(c) Bias in risk under never treat

Figure S11: Bias in risk in dataset with 10,000 individuals



(a) Bias in risk ratio



(b) Bias in risk difference

Figure S12: Bias in effect estimates (always vs. never treat interventions) in 10,000 sample dataset

References

- [1] B. A. Dickerman, E. Giovannucci, C. H. Parnar, L. A. Mucci, and M. A. Hernán. Guideline-based physical activity and survival among us men with nonmetastatic prostate cancer. *Am J Epidemiol*, 188(3):579–586, 2019.
- [2] M. A. Hernán and J. M. Robins. *Causal Inference: What If*. Chapman Hall/ CRC, 2022.
- [3] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [4] P. Jain, G. Danaei, J. E. Manson, J. M. Robins, and M. A. Hernán. Weight gain after smoking cessation and lifestyle strategies to reduce it. *Epidemiology*, 31(1):7–14, 2020.
- [5] Rui Li, Stephanie Hu, Mingyu Lu, Yuria Utsumi, Prithwish Chakraborty, Daby M. Sow, Piyush Madan, Jun Li, Mohamed Ghalwash, Zach Shahn, and Li-wei Lehman. G-net: a recurrent network approach to g-computation for counterfactual prediction under a dynamic treatment regime. *Proceedings of Machine Learning Research*, 158:282–299, 2021.
- [6] S. Lodi, A. Phillips, R. Logan, A. Olson, D. Costagliola, S. Abgrall, A. van Sighem, P. Reiss, J. M. Miró, E. Ferrer, A. Justice, N. Gandhi, H. C. Bucher, H. Furrer, S. Moreno, S. Monge, G. Touloumi, N. Pantazis, J. Sterne, J. G. Young, L. Meyer, R. Seng, F. Dabis, M. A. Vandehende, S. Pérez-Hoyos, I. Jarrín, S. Jose, C. Sabin, and M. A. Hernán. Comparative effectiveness of immediate antiretroviral therapy versus cd4-based initiation in hiv-positive individuals in high-income countries: observational cohort study. *Lancet HIV*, 2(8):e335–43, 2015.
- [7] S. McGrath, J. G. Young, and M. A. Hernán. Revisiting the g-null paradox. *Epidemiology*, 33(1):114–120, 2022.
- [8] S. M. Rein, S. Lodi, R. Logan, G. Touloumi, A. Antoniadou, L. Wittkop, F. Bonnet, A. van Sighem, M. van der Valk, P. Reiss, M. B. Klein, J. Young, I. Jarrin, A. d’Arminio Monforte, A. Tavelli, L. Meyer, L. Tran, M. J. Gill, R. Lang, B. Surial, A. D. Haas, A. C. Justice, C. T. Rentsch, A. Phillips, C. Sabin, J. M. Miró, A. Trickey, S. M. Ingle, J. A. C. Sterne, and M. A. Hernán. Integrase strand-transfer inhibitor use and cardiovascular events in adults with hiv: An emulation of target trials in the hiv-causal and art-cc collaborations. *Lancet HIV [in press]*, 2023.
- [9] 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 Modelling*, 7(9):1393–1512, 1986.
- [10] James M. Robins and Larry Wasserman. Estimation of effects of sequential treatments by reparameterizing directed acyclic graphs. *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*, page 409–420, 1997.
- [11] S. L. Taubman, J. M. Robins, M. A. Mittleman, and M. A. Hernán. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *Int J Epidemiol*, 38(6):1599–611, 2009.
- [12] L. Wen, J. G. Young, J. M. Robins, and M. A. Hernán. Parametric g-formula implementations for causal survival analyses. *Biometrics*, 77(2):740–753, 2021.