

---

# DEEPBLIP: ESTIMATING CONDITIONAL AVERAGE TREATMENT EFFECTS OVER TIME

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Structural nested mean models (SNMMs) are a principled approach to estimate the treatment effects over time. A particular strength of SNMMs is to break the joint effect of treatment sequences over time into localized, time-specific “blip effects”. This decomposition promotes interpretability through the incremental effects and enables the efficient offline evaluation of optimal treatment policies without re-computation. However, neural frameworks for SNMMs are lacking, as their inherently sequential g-estimation scheme prevents end-to-end, gradient-based training. Here, we propose DeepBlip, the first neural framework for SNMMs, which overcomes this limitation with a novel double optimization trick to enable simultaneous learning of all blip functions. Our DeepBlip seamlessly integrates sequential neural networks like LSTMs or transformers to capture complex temporal dependencies. By design, our method correctly adjusts for time-varying confounding to produce unbiased estimates, and its Neyman-orthogonal loss function ensures robustness to nuisance model misspecification. Finally, we evaluate our DeepBlip across various clinical datasets, where it achieves state-of-the-art performance.

## 1 INTRODUCTION

Predicting the effects of treatment sequences is crucial for personalized medicine to choose the best therapeutic strategy for a patient based on their history (Feuerriegel et al., 2024). Methodologically, the **conditional average treatment effect (CATE) over time** captures the combined effect of multiple treatments in the next  $\tau$  time steps (see Fig. 1). In clinical practice, CATEs over time can be frequently estimated from observational data with patient histories, such as the electronic health records (Allam et al., 2021; Bica et al., 2021). Yet, a key challenge is *time-varying confounding*, which arises when past treatments affect future covariates that, in turn, influence both subsequent treatment assignment and the outcome of interest.

Several statistical strategies exist to address time-varying confounding. While these strategies all target the same estimand, they differ in their practical strengths, especially in finite-sample settings. ① One approach is *G-computation* (Robins, 1986; Robins et al., 1999), with neural instantiations such as **G-Net** (Rui Li et al., 2021) and the **G-transformer** (Hess et al., 2025)). This approach sequentially models the conditional outcome given treatment and covariates, and then averages (integrates) over the distribution of time-varying covariates under specific treatment regimes. ② Another approach is *marginal structural models (MSNs)* (Robins, 1994; 2004), with neural instantiations such as **RMSNs** (Lim, 2018). However, MSNs rely on inverse propensity weighting and can be unstable as extreme weight often occurs in long prediction windows. ③ A third approach is *structural nested mean models (SNMMs)* (Robins, 1994; 2004). A particular strength of SNMMs is their ability to break the joint effect of treatment sequences over time into localized, time-specific “*blip effects*”. This yields a simpler estimand that does not require learning full counterfactual outcome trajectories, which can be especially beneficial over long horizons. In practice, this can have several benefits: the incremental effects are interpretable, and, the blip effects can be reused to predict treatment effects for new treatment sequences without re-computation, which allows for identifying optimal treatment sequences through offline evaluation.  $\Rightarrow$  *Our work is located in the third stream and contributes to research on SNMMs.*

However, neural implementations of SNMMs are missing. The closest effort in this direction is **Dynamic DML** by Lewis & Syrkanis (2021); however, this approach has several *limitations*: (i) The original implementation is based on linear models. While, in principle, other machine learning models *could* be used, this is impractical because essentially *a different machine model must be estimated for each time step*, so the complexity grows linearly with the prediction horizon. (ii) Dynamic DML uses a sequential minimization scheme that does not scale to large datasets. Hence, their training follows an *iterative* procedure, while our second stage later is *end-to-end*. (iii) Dynamic DML is unable to process the full, accumulating history of a patient. Its architecture requires a fixed-size input, preventing it from conditioning on patient data (clinical records) that grows over time. Yet, building a neural framework for SNMMs

is non-trivial. One major obstacle is that the SNMM formulation inherently prevents end-to-end learning due to its sequential scheme. For the same reason, Dynamic DML cannot be directly adapted. As a remedy, we propose a new *double optimization trick* to break the sequential dependence and enable gradient-based training. This approach allows us to address (i)–(iii).

Here, we propose **DeepBlip**, the first *neural* framework to estimate CATE over time by leveraging the blip function from SNMMs. DeepBlip decomposes the joint effect of treatment sequences over time into localized, time-specific blip effects, which enables more tractable and stable learning. As a result, our DeepBlip adjusts for time-varying confounding and is thus *unbiased*. To do so, a key novelty is that we propose a new *double optimization trick* to break the sequential dependence in SNMMs and allow for gradient-based training. Our DeepBlip is built on top of sequential neural networks (e.g., LSTMs, transformers) to capture complex temporal dependencies. For this, it employs a two-stage architecture: Stage 1 models the probability of time-varying treatments and mean outcomes conditioned on a patient’s history, while Stage 2 reformulates *g*-estimation (Robins, 1994; 2004) as a risk minimization task to directly learn the blip functions.

Our DeepBlip has several further strengths for medical practice: (i) The estimated blip effects are *interpretable*, so that clinicians can understand the incremental gain from treatment. (ii) The loss function in DeepBlip is *Neyman-orthogonal*, which makes the estimates robust against model misspecification. (iii) The learned blip effects are reusable, which allows clinicians to predict outcomes under novel treatment sequences *without re-computation*. This offers an efficient approach for offline evaluation of different therapeutic strategies and thus to search for optimal treatment sequences. Formally, at inference time, DeepBlip can identify the optimal treatment sequence within just one forward pass. This is unlike other methods, which typically require either re-training (Hess et al., 2025; Rui Li et al., 2021) or multiple forward passes due to exhaustive search (Bica et al., 2020; Lim, 2018; Melnychuk et al., 2022).

Our **contributions** are three-fold:<sup>1</sup> (1) We introduce the first neural framework to predict CATE over time via the SNMM framework. (2) Our framework is carefully tailored to medical practice by benefiting from robustness due to Neyman-orthogonality and efficient offline evaluation. (3) We conduct extensive experiments across multiple medical datasets to demonstrate that our DeepBlip is effective and also robust across long time horizons.

## 2 RELATED WORK<sup>2</sup>

**Estimating CATE in the static setting:** There has been extensive research in estimating CATE with neural networks in the static setting (e.g., Alaa & van der Schaar; Curth & van der Schaar, 2021; Künzel et al., 2019; Shalit et al., 2017; Wager & and, 2018; Yoon et al., 2018; Künzel et al., 2019). However, these methods are aimed at static settings and thus struggle with medical datasets such as electronic health records, where patient histories are recorded *over time* and give rise to *time-varying confounding*.

**Estimating ATE over time:** One line of work has developed methods for the ATE over time (Frauen et al., 2023; Shirakawa et al., 2024). However, the ATE captures only population-level effects and thus overlooks differences in treatment effectiveness across patients. In contrast, we focus on the CATE, which provides a more granular, individualized estimate of treatment outcomes, which is highly relevant for personalized medicine (Feuerriegel et al., 2024).

**Estimating CATE over time:** There are several *neural* methods for this task<sup>3</sup>, which can be broadly categorized into different streams (see Table 1):

<sup>1</sup>Code for review is available at <https://github.com/anonymous> Upon acceptance, we will move our code to a public GitHub repository.

<sup>2</sup>We provide an extended related work in Appendix D

<sup>3</sup>There have been some attempts to use non-parametric models (Schulam & Saria, 2017; Soleimani et al., 2017; Xu et al., 2016), yet these approaches impose strong assumptions on the outcomes, and their scalability is limited. For these reasons, we focus on neural methods, which offer better flexibility and scalability for complex, high-dimensional medical data.

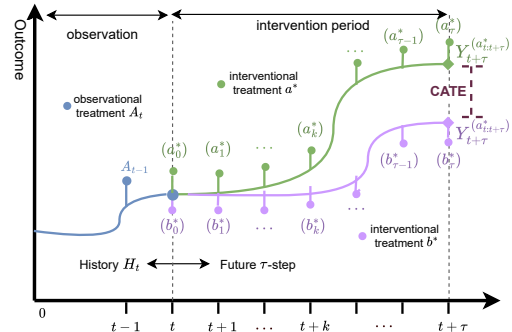


Figure 1: **CATE over time.** Trajectories of potential outcomes under two interventional sequences  $a_{t:t+\tau}^*, b_{t:t+\tau}^*$  given the shared observed history  $H_t$ . The difference between the two curves is the CATE over time.

• *No proper adjustment for confounding and thus bias*: Some methods for CATE over time fail to properly adjust for time-varying confounding, which leads to estimates that are *biased*. Here, prominent examples are counterfactual recurrent network (CRN) (Bica et al., 2020) and the causal transformer (CT) (Melnichuk et al., 2022). These methods attempt to alleviate time-varying confounding via balanced representations. However, balancing was originally designed for reducing finite-sample estimation variance and *not* for mitigating confounding bias (Shalit et al., 2017). Hence, such methods act as heuristics without a theoretical justification. The difficulty of enforcing balanced representations may even introduce further confounding bias (Melnichuk et al., 2024). Unlike these methods, our DeepBlip allows for proper adjustments and is thus unbiased.

• *Non-SNMM-based adjustment strategies*: Other neural methods build on frameworks from statistics such as marginal structural models (MSMs) (Orellana et al., 2010; Robins et al., 2000) and G-computation (Robins, 1986; Robins et al., 1999). Examples are: ① **G-Net** (Rui Li et al., 2021) and G-transformer (**GT**) (Hess et al., 2025), which are both based on G-computation and thus compute nested conditional expectations over time. Hence, these methods require the entire counterfactual outcome trajectories and thus model the full data-generating process of covariates and outcomes, which becomes exponentially more complex as time horizons grow. ② MSMs-based methods like **R-MSNs** (Lim, 2018) and the **DR-learner** (for time-varying settings) (Frauen et al., 2025) use inverse propensity weighting (IPW) to re-weight outcomes as in a randomized control trial. In time-varying settings, the propensity is a multiplication of a sequence of probabilities, which has known to become unstable when the horizon is large. In sum, methods from both streams can often become *unstable* over long time horizons (different from SNMMs that break the CATE into localized effects at each time step).

• *SNMM-based adjustment strategies*: SNMMs (Robins, 1994; 2004) offer a principled framework for estimating CATE over time by directly estimating incremental treatment effects via so-called blip functions. In principle, SNMMs could address both of the above limitations; however, SNMMs are only an abstract theoretical foundation, *not* an off-the-shelf model that can be directly applied. Early implementations (Robins, 1994; 2004) relied on strong parametric assumptions (e.g., linearity) and were limited to short time horizons.

Stream	Method	Methodological advantage		Benefits for medical practice	
		Unbiased	Neural	Orthogonal	Offline efficiency
W/o adjustment	CRN (Bica et al., 2020)	✗	✓	✗	✓
	CT (Melnichuk et al., 2022)	✗	✓	✗	✓
G-computation	G-Net (Rui Li et al., 2021)	✓	✓	✗	✗
	GT (Hess et al., 2025)	✓	✓	✗	✗
MSNs	R-MSNs (Lim, 2018)	✓	✓	✗	✗
	DR-learner (Frauen et al., 2025)	✓	✓	✓	✗
SNMMs	Dynamic DML (Lewis & Syrgkanis, 2021)	✓	✗	✓	✓
	<b>DeepBlip (ours)</b>	✓	✓	✓	✓

Table 1: **Methods for learning CATE over time.**

Despite the theoretical appeal of SNMMs, a neural implementation is missing. Closest to our work is Dynamic DML (Lewis & Syrgkanis, 2021), which builds upon SNMMs but still has key *limitations*. The original implementation is provided for *linear* models, which cannot capture complex, non-linear relationships. In principle, the framework in Lewis & Syrgkanis (2021) is model-agnostic and *could* be used along with other deep learning models. Yet, this is impractical because essentially *a different network must be estimated for each time step*, so the complexity *grows linearly with the prediction horizon*. Because of this, extending Dynamic DML into a scalable neural framework is *non-trivial*: the sequential g-estimation of SNMM (as used by Dynamic DML) is iterative and essentially prevents efficient gradient-based and thus end-to-end optimization. As a remedy, we later propose a new *double optimization trick*.

**Research gap**: To the best of our knowledge, there is no neural implementation of SNMMs. For this, we introduce a new neural architecture together with a new *double optimization trick* to leverage the blip function from SNMMs for CATE estimation over time.

### 3 PROBLEM FORMULATION

**Setup**: We follow the standard setup (Bica et al., 2020; Frauen et al., 2025; Hess et al., 2025; Melnychuk et al., 2022) for estimating CATE over time  $t \in \{1, 2, \dots, T\} \subset \mathbb{N}$  given by: (1) the target outcome  $Y_t \in \mathbb{R}$ ; (2) time-varying covariates  $X_t \in \mathcal{X} \subset \mathbb{R}^{d_x}$ ; and (3) treatment  $A_t \in \mathcal{A} \subset \mathbb{R}^{d_a}$  that can be either discrete or continuous. We assume w.l.o.g. that the static features (e.g., age, sex) are included in the covariate.

**Notation**: To simplify notation, we use overlines to denote the full sequence of a variable (e.g.,  $\overline{X}_t = (X_1, \dots, X_t)$ ). We refer to a sequence of variables that starts at  $t$  and ends at  $t + \tau$  via  $A_{t:t+\tau} = (A_t, A_{t+1}, \dots, A_{t+\tau})$ . We use the

lowercase letter to denote a realization of a random variable (e.g.,  $X_t = x_t$ ). We use an asterisk  $*$  to indicate a constant quantity (e.g., a fixed treatment  $a_t^*$ ). We denote the patient history by  $H_t = (\bar{X}_t, \bar{A}_{t-1}, \bar{Y}_{t-1})$ .

**CATE estimation over time:** We build upon the potential outcome framework (Rubin, 1972) for the time-varying setting (Robins et al., 1999). We aim to estimate the CATE over time between two treatment sequences for a given patient history, i.e.,

$$\mathbb{E} \left[ Y_{t+\tau}^{(a_{t:t+\tau}^*)} - Y_{t+\tau}^{(b_{t:t+\tau}^*)} \mid H_t = h_t \right], \quad 0 \leq t \leq T - \tau, \quad (1)$$

where  $Y_{t+\tau}^{(a_{t:t+\tau}^*)}$  and  $Y_{t+\tau}^{(b_{t:t+\tau}^*)}$  represent the  $\tau$ -step-ahead potential outcomes under interventions  $\text{do}(A_{t:t+\tau} = a_{t:t+\tau}^*)$  and  $\text{do}(A_{t:t+\tau} = b_{t:t+\tau}^*)$ , respectively (see Appendix A for a formal definition of potential outcomes and interventions).

**Identifiability:** We make the following identifiability assumptions (Petersen & van der Laan, 2014; Robins et al., 2000) that are standard in the time-varying setting (Bica et al., 2020; Lim, 2018; Melnychuk et al., 2022; Rui Li et al., 2021): (1) *Consistency:* The potential outcome under the intervention by the observed treatment equals the observed outcome, namely,  $Y_t^{(A_t)} = Y$ . (2) *Overlap:* Given an observed history  $H_t = h_t$ , if  $\Pr(H_t = h_t) > 0$ , then any possible treatment has a positive probability of being received:  $\forall a \in \mathcal{A}, \Pr(A_t = a \mid H_t = h_t) > 0$ . (3) *Sequential ignorability:* The potential outcome under an arbitrary intervention is independent of the treatment assignment conditioned on the history, i.e.,  $Y_t^{(a_t^*)} \perp A_t \mid H_t = h_t$ .

However, estimating the CATE over time is non-trivial due to *time-varying confounding* (Coston et al., 2020; Hess et al., 2025). In the time-varying setting, covariates act as confounders because they are influenced by earlier treatments and then affect later treatments and outcomes. This means they cannot be adjusted for naively as in the static setting without introducing bias (see Appendix A.2). In this paper, we address this issue via SNMMs.

**Blip function:** SNMMs model the incremental effect of treatments (which are called “blips”) at time  $t + k$  on the mean outcome at  $t + \tau$ , given observed patient history  $H_t$  (Vansteelandt & Joffe, 2014). These “blips” accumulate over time into the total treatment effect and thus allow to rigorously adjust for time-varying confounding (Robins, 2004) (see Thm 1 below). Formally, let  $d$  be a treatment policy. The blips are defined via a **blip function** (Robins, 1994; 2004):

$$\gamma_{t,k}(\bar{x}_{t+1:t+k}, \bar{a}_{t:t+k}; h_t) = \mathbb{E} \left[ Y_{t+\tau}^{(a_{t:t+k}, d_{t+k+1:t+\tau})} - Y_{t+\tau}^{(a_{t:t+k-1}, 0, d_{t+k+1:t+\tau})} \mid A_{t:t+k} = a_{t:t+k}, X_{t+1:t+k} = x_{t+1:t+k}, H_t = h_t \right]. \quad (2)$$

Intuitively, the blip function  $\gamma_{t,k}$  isolates the causal effect of each treatment decision *locally*. By isolating the local effect of each treatment, this formulation avoids modeling the full outcome trajectory, making SNMMs often stable over long horizons.

**Theorem 1. Adjustment via blip functions** (Theorem 3.1 from Robins (2004)) *Given a policy  $d$  between time  $t$  and  $t + \tau$ , the following identification holds under the sequential ignorability assumption in 2:*

$$\mathbb{E} \left[ Y^{(d)} \mid H_t = h_t \right] = \mathbb{E} \left[ Y + \sum_{k=0}^{\tau} \rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) \mid H_t = h_t \right] \quad (3)$$

where  $\rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) = \gamma_{t,k}(X_{t+1:t+k}, (A_{t:t+k-1}, d_{t+k}); h_t) - \gamma_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t)$ .

The theorem demonstrates that SNMMs adjust for time-varying confounding by making local adjustments at each step. Furthermore, if we can correctly estimate the blip functions, the CATE over time estimation will converge to the unbiased causal effect given sufficiently large data.

## 4 OUR DEEPBLIP FRAMEWORK

In this section, we present DeepBlip. First, we introduce how we learn the CATE via blip functions using a neural parameterization (Sec. 4.1), then introduce our  $L^2$ -moment loss (Sec. 4.2), our model architecture (Sec. 4.3), and the training and inference procedure (Sec. 4.4).

### 4.1 LEARNING THE CATE VIA BLIP FUNCTIONS

**Overview:** Our DeepBlip leverages Eq. (4) to adjust for time-varying confounding. Our task thus reduces to estimating the blip functions – in particular, so-called *blip coefficients* that parametrize the blip functions. However, we do **not**

attempt to estimate the coefficients via the iterative g-estimation. Instead, we optimize a  $L^2$ -moment loss that directly predicts the blip coefficients and which allows us to estimate Eq. (4) more efficiently.

**Parameterization trick:** We first explain how we estimate the CATE via the blip function. For this, we adopt a similar parametrization for the blip function as in (Lewis & Syrgkanis, 2021), namely,  $\gamma_{t,k}(\tilde{x}_{t+1:t+k}, \tilde{a}_{t:t+k}; h_t) = \psi_{t,k}(h_t)' a_{t+k}$ , but where  $\psi_{t,k}$  is a **neural network**.<sup>4</sup> Under identifiability assumptions and the parametrization for  $\gamma_{t,k}$  defined above, the CATE of  $a^*$  against  $b^*$  for any two treatment sequences  $a^*, b^* \in \mathbb{R}^{(\tau+1) \cdot d_a}$  is (see (Lewis & Syrgkanis, 2021) for a formal derivation):

$$\mathbb{E}[Y_{t+\tau}^{(a^*)} - Y_{t+\tau}^{(b^*)} | H_t = h_t] = \sum_{k=0}^{\tau} \psi_{t,k}(h_t)' (a_{t+k}^* - b_{t+k}^*). \quad (4)$$

We refer to  $\psi_{t,k}(h_t)$  as the conditional *blip coefficients* of the blip function  $\gamma_{t,k}$ .

*Why do we need a tailored architecture and learning algorithm?* A key component of our framework is that the function  $\psi_{t,k}(h_t)$  is parameterized by a sequential neural network (e.g., LSTM or transformer). This is a crucial difference from traditional SNMMs, which were developed for estimating the ATE over a fixed number of time steps (i.e.  $\mathcal{H}_t = \emptyset \wedge t \equiv 0 \wedge T \equiv \tau$ ) and where, as a result, blip coefficients are constants. These constants are typically estimated through *iteratively* solving a set of moment equations via g-estimation (Robins, 1994; 2004; Vansteelandt & Joffe, 2014) (see Appendix B.1). However, such an approach is not compatible with neural network-based learning. In contrast, DeepBlip introduces a tailored neural architecture (Sec. 4.3) that we can train via gradient-based optimization (Sec. 4.2).

## 4.2 $L^2$ -MOMENT LOSS

We adapt the moment-based iterative minimization from Lewis & Syrgkanis (2021) to our history-conditioned setting. For  $k = \tau, \dots, k = 0$ , at each time step  $k$ , we aim to find the minimizer  $\psi_{t,k}^*(\cdot)$ , which is a function that maps the history  $h_t$  to the blip coefficients, via

$$\psi_{t,k}^* = \arg \min_{\hat{\psi}_{t,k}(\cdot) \in \Phi_{t,k}} \mathbb{E} \left[ \left( \tilde{Y}_{t,k} - \sum_{j=k+1}^{\tau} \psi_{t,j}^*(h_t)' \tilde{A}_{t,j,k} - \hat{\psi}_{t,k}(h_t)' \tilde{A}_{t,k,k} \right)^2 \right], \quad (5)$$

where  $\Phi_{t,k}$  is the function space for the blip coefficient predictors and where

$$\tilde{Y}_{t,k} = Y_{t+\tau} - \mathbb{E}[Y_{t+\tau} | H_{t+k} = h_{t+k}], \quad \tilde{A}_{t,j,k} = A_{t+j} - \mathbb{E}[A_{t+j} | H_{t+k} = h_{t+k}]. \quad (6)$$

are residuals of the conditional means for  $Y$  and  $A$ , respectively. We name the target as the  $L^2$ -**moment loss** for step  $k$  and denote it as  $\mathcal{L}_k$ . Notice that only one  $\hat{\psi}_{t,k}(\cdot)$  is solved by minimizing  $\mathcal{L}_k$ ; hence, in order to solve the entire task (i.e., all the  $\hat{\psi}_{t,k}$ ,  $(0 \leq k \leq \tau)$ ), one would normally need to conduct a series of minimizations **sequentially** from  $k = \tau$  down to  $k = 0$  such as in Dynamic DML from (Lewis & Syrgkanis, 2021).

**Double optimization trick for our  $L^2$ -moment loss:** In order to find  $\psi_{t,k}^*$ , all the previous blip predictors  $\psi_{t,j}^*$ ,  $j \geq k$  are required. To avoid solving  $\psi_{t,k}^*$  sequentially, we propose a *double optimization trick* that allows *simultaneous* training of all the blip predictors: During each iteration, first, the blip predictor  $\hat{\psi}_t$  makes two forward passes to generate two sets of the blip coefficients  $\hat{\psi}_t^1(h_t)$  and  $\hat{\psi}_t^2(h_t)$ . Then,  $\hat{\psi}_t^2(h_t)$  serves as a fixed, "pseudo-ground-truth" target for the future blip effects in Eq. (5) and is detached from the computation graph. For  $k = 0, \dots, \tau$ , the adapted  $L^2$ -moment loss at step  $k$  is then given empirically by

$$\mathcal{L}_{\text{blip}}^k = \frac{1}{T - \tau} \sum_{t=1}^{T-\tau} \sum_{i=1}^n \left( \tilde{Y}_{t+k}^i - \sum_{j=k+1}^{\tau} \hat{\psi}_j^2(H_t^i)' \tilde{A}_{t,j,k}^i - \hat{\psi}_k^1(H_t^i)' \tilde{A}_{t,k,k}^i \right)^2. \quad (7)$$

**Why is our double optimization trick necessary for scalability?** Existing SNMM methods, like Dynamic DML (Lewis & Syrgkanis, 2021), rely on a sequential scheme that solves for blip coefficients one at a time. For neural networks, this would require *training a separate model for each of the  $\tau + 1$  time steps, which is computationally prohibitive*. In contrast, our double optimization trick *bypasses this sequential dependency*, allowing us to instead optimize a single, unified objective:  $\mathcal{L}_{\text{blip}} = \sum_{k=0}^{\tau} \mathcal{L}_{\text{blip}}^k$ . This makes our framework significantly more scalable, as

<sup>4</sup>This formulation is chosen mainly for clarity. In fact, it is readily extended to capture nonlinear treatment interactions by incorporating a high-dimensional, nonlinear feature map  $\phi(\cdot)$ . The blip function can then be parametrized as  $\gamma_{t,k}(h_t) = \psi_{t,k}'(h_t) \phi(a_{t+k})$ .

a neural version of the sequential approach would multiply the total training time by a factor of at least  $\tau + 1$  (see Figure 4).

We show in Appendix C that the double optimization trick is a single-step gradient descent realization of an iterative scheme for the complete blip predictor  $\hat{\psi}$ . Under this scheme, all the intermediate blip predictors  $\{\hat{\psi}_j \mid (j < \tau)\}$  converge to the true value as long as the final predictor  $\hat{\psi}_\tau$  converges to the true value. Further, the convergence of the final predictor  $\hat{\psi}_\tau$  is guaranteed by the  $L^2$ -moment loss being the correct objective, since  $\mathcal{L}_{\text{blip}}^\tau = \mathcal{L}_\tau$ . Therefore, all the blip predictors will converge under the double optimization trick. We further perform ablation studies on the double optimization trick in Appendix E.1 and E.2 to empirically show its effectiveness.

**Theoretical properties:** Below, we state that our loss satisfies *Universal Neyman-orthogonality* (see Foster & Syrgkanis (2019) for a formal definition), which ensures double robustness. This means that the target loss is *robust* against perturbations of the nuisance functions (Chernozhukov et al., 2018; Kennedy, 2023). The remark below follows from the theory in Lewis & Syrgkanis (2021), which can be extended to our double optimization trick (see Appendix B.2 for a formal derivation, respectively).

*Remark 1. Our moment loss is Neyman-orthogonal.*

### 4.3 MODEL ARCHITECTURE

DeepBlip works in two stages (see Fig. 2): **• Stage ① (nuisance network):** models the nuisance functions to estimate the residuals in Eq. (6). **• Stage ② (blip prediction network):** estimates the blip coefficients given the observed history  $h_t$ . The neural networks in both stages have a similar structure: (i) a *sequential encoder* that encodes the observed history  $H_t$ , and (ii) multiple *prediction heads* that take the encoded history as input to predict the targets.

**Why we need a two-stage design:** To construct the  $L^2$ -moment loss defined Eq. (5), we need the variables  $\tilde{Y}_{t,k}, \tilde{A}_{t,j,k}$ , defined in Eq. (6) (see Appendix B.1 for details). Based on the definitions, we must compute the *residuals* between the outcome variable and the regressed means *before* we optimize  $\mathcal{L}_{\text{blip}}^k$ . We thus follow previous literature (Lewis & Syrgkanis, 2021; Frauen et al., 2025; Kennedy, 2023) and treat the conditional expectations  $\mathbb{E}[Y_{t+\tau} \mid H_{t+k} = h_{t+k}]$  and  $\mathbb{E}[A_{t+j} \mid H_{t+k} = h_{t+k}]$  as *nuisance functions*, so that we first estimate the nuisance function (Stage ①) and then train the blip prediction network to minimize the  $L^2$ -moment loss (Stage ②).

**Neural backbone:** Our DeepBlip is flexible and allows for different neural backbones (e.g., LSTM or transformer). These necessary to capture the patient history: We notice that the networks at both Stage ① and ② take the history variable  $H_t = (\bar{X}_t, \bar{A}_{t-1}, \bar{Y}_{t-1}) \in \mathcal{H}_t$  as input. Hence, both stages can be written as a function  $f: \cup_{t=1}^T \mathcal{H}_t \rightarrow \mathbb{R}^c$ , where  $c$  is the number of outputs. However,  $\dim(H_t)$  varies over time, which makes  $H_t$  not suitable as a direct input to a neural network. A standard way to handle this is by using a sequential model to iteratively take the inputs  $(X_t, A_{t-1}, Y_{t-1})$  and then maintain a vector  $Z_t \in \mathbb{R}^{d_z}$  with fixed dimension that encodes all the necessary information (Hess et al., 2025; Lim, 2018; Melnychuk et al., 2022; Rui Li et al., 2021). Here, we thus use LSTMs and transformers (see details of the architectures in Appendix I). Finally, we stress that each stage uses a *separate* encoder:  $\mathcal{E}_{\theta_N}^N$  for Stage ①, and  $\mathcal{E}_{\theta_B}^B$  for Stage ② ( $N$  for Nuisance and  $B$  for Blip) with different model weights  $\theta_N$  and  $\theta_B$ .

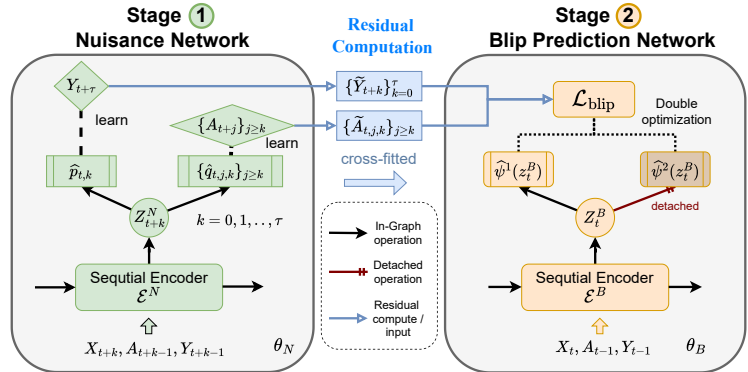


Figure 2: Neural architecture of the two-stage DeepBlip framework.

**Stage ①: Nuisance network.** The nuisance network  $(\mathcal{E}_{\theta_N}^N, \{\text{gp}_{\theta_N}^k\}_{k=0}^{\tau}, \{\text{gq}_{\theta_N}^{j,k}\}_{0 \leq k \leq j \leq \tau})$  consists of a sequential encoder  $\mathcal{E}_{\theta_N}^N$  and a collection of prediction heads  $\{\text{gp}_{\theta_N}^k\}_{k=0}^{\tau}, \{\text{gq}_{\theta_N}^{j,k}\}_{0 \leq k \leq j \leq \tau}$ . The nuisance networks are responsible for computing the following nuisance functions:

$$p_{t,k}(h_{t+k}) := \mathbb{E}[Y_{t+\tau} \mid H_{t+k} = h_{t+k}], \quad 1 \leq t \leq T - \tau, 0 \leq k \leq \tau \quad (8)$$

$$q_{t,j,k}(h_{t+k}) := \mathbb{E}[A_{t,j} \mid H_{t+k} = h_{t+k}], \quad 1 \leq t \leq T - \tau, 0 \leq k \leq j \leq \tau \quad (9)$$

For a patient with history  $H_t$  and subsequent covariates  $X_{t+1:t+k}, A_{t:t+k-1}$ , we proceed as follows: First, the encoder  $\mathcal{E}_{\theta_N}^N$  processes history  $H_{t+k}$  to produce a representation  $Z_{t+k}^N = \mathcal{E}_{\theta_N}^N(H_{t+k})$ . Then, the prediction heads receive  $Z_{t+k}^N$  to compute the regressed outcomes for the nuisance functions via:

$$\mathbf{gp}_{\theta_N}^k(Z_{t+k}^N) = \hat{p}_{t,k}(H_{t+k}), \quad \mathbf{gq}_{\theta_N}^{j,k}(Z_{t+k}^N) = \hat{q}_{t,j,k}(H_{t+k}) \quad \text{for } k = 0, \dots, \tau \text{ and } k \leq j \leq \tau \quad (10)$$

where  $Z_{t+k}^N = \mathcal{E}_{\theta_N}^N(H_{t+k})$ . Third, the residuals are computed via

$$\tilde{Y}_{t,k} \approx Y_{t+\tau} - \mathbf{gp}_{\theta_N}^k(Z_{t+k}^N), \quad \tilde{A}_{t,j,k} \approx A_{t+j} - \mathbf{gq}_{\theta_N}^{j,k}(Z_{t+k}^N) \quad (11)$$

**Stage ②: Blip prediction network.** The blip prediction network ( $\mathcal{E}_{\theta_B}^B, \{\mathbf{gb}_{\theta_B}^k\}_{k=0}^\tau$ ) is responsible for predicting the blip coefficients  $\psi_t(h_t) = (\psi_{t,0}, \dots, \psi_{t,\tau}) \in \mathbb{R}^{r(\tau+1)}$  as described in Eq. (equation 5). Here, we proceed as follows. First, the sequential encoder  $\mathcal{E}_{\theta_B}^B$  (B for Blip) processes the patient's history  $H_t$  into a representation  $Z_t^B = \mathcal{E}_{\theta_B}^B(H_t)$ . Then, for each horizon  $k \in \{0, 1, \dots, \tau\}$ , the prediction head  $\mathbf{gb}_{\theta_B}^k$  maps  $Z_t^B$  onto the corresponding blip coefficient:

$$\hat{\psi}_{t,k}(H_t) = \mathbf{gb}_{\theta_B}^k(Z_t^B) \sim \psi_{t,k}(H_t) \in \mathbb{R}^r \quad \text{where } Z_t^B = \mathcal{E}_{\theta_B}^B(H_t). \quad (12)$$

#### 4.4 TRAINING AND INFERENCE

Taken together, the training procedure of DeepBlip now follows two steps (see Fig. 2): (1) train the nuisance networks and compute the residuals, and (2) train the blip prediction network. In contrast, inference with DeepBlip is highly efficient as it involves *only* the second-stage blip prediction network. Details are below. We provide the pseudocode in Alg. 1 and Alg. 2 in the appendix.

**Step ①: Train nuisance network.** The nuisance network is trained to predict nuisance functions  $p_{t,k}(h_{t+k})$  and  $q_{t,j,k}(h_{t+k})$  simultaneously. Since  $p_{t,k}(h_{t+k})$  is the conditional expectation of real outcome  $Y_{t+\tau} \in \mathbb{R}$ , we use the squared error loss  $\mathcal{L}_p = \frac{1}{(T-\tau)(\tau+1)} \sum_{t=1}^{T-\tau} \sum_{k=0}^\tau (\mathbf{gp}_{\theta_N}^k(Z_{t+k}^N) - Y_{t+\tau})^2$ . For  $q_{t,j,k}(h_{t+k})$ , which denotes the treatment response, we proceed for the  $i$ -th treatment in  $A_{t+j} \in \mathbb{R}^{d_a}$  as follows. If  $(A_{t+j})_i$  is a continuous variable, then we apply the squared loss:  $\mathcal{L}_{q,i} = \frac{2}{(T-\tau)(\tau+1)(\tau+2)} \sum_{t=1}^{T-\tau} \sum_{0 \leq k \leq j \leq \tau} (\mathbf{gq}_{\theta_N}^{k,j}(Z_{t+k}^N)_i - (A_{t+j})_i)^2$ . If  $(A_{t+j})_i$  is a binary variable, then we apply the binary cross entropy loss  $\mathcal{L}_{q,i} = \frac{2}{(T-\tau)(\tau+1)(\tau+2)} \sum_{t=1}^{T-\tau} \sum_{0 \leq k \leq j \leq \tau} \text{BCE}((A_{t+j})_i, \mathbf{gq}_{\theta_N}^{k,j}(Z_{t+k}^N)_i)$ . For categorical variables with more than 2 classes, we preprocess the variable into a one-hot vector of binary variables. Since the network predicts these targets simultaneously, we update the parameter  $\theta_N$  by backpropagating the sum of all the losses discussed above, i.e.,  $\mathcal{L}_N = \mathcal{L}_p + \frac{1}{d_a} \sum_{i=1}^{d_a} \mathcal{L}_{q,i}$

**Step ②: Train blip prediction network.** After having trained the nuisance network, we freeze its parameters and then compute the residuals of each sample as in Eq. (6) for the  $L^2$ -moment loss. To accelerate the training process, we adopt the double optimization trick from above: For  $k = 0, \dots, \tau$ , we perform two forward passes that create two predictions  $\hat{\psi}_{t,k}^1(H_t)$  and  $\hat{\psi}_{t,k}^2(H_t)$ . The latter is then *detached* from the computation graph before feeding into the adapted  $L^2$ -moment loss at step  $k$ . The final loss target is then given by:  $\mathcal{L}_{\text{blip}} = \sum_{k=0}^\tau \mathcal{L}_{\text{blip}}^k$ .

*Remark 2. Under standard assumptions, the output of our DeepBlip has a mean squared error guarantee:*

$$\max_{t \leq T-\tau} \max_{k \in \{0, \dots, \tau\}} \mathbb{E} \left[ \left\| \hat{\psi}_{t,k} - \psi_{t,k} \right\|_{2,2}^2 \right] = O(r^2 \delta_n^2), \quad \delta_n^2 \propto \frac{\log \log(n)}{n} \quad (13)$$

*Details are in Appendix B.3).*

**Inference at runtime:** Once trained, our DeepBlip predicts the CATE over time (i.e.,  $\mathbb{E}[Y_{t+\tau}^{(a^*)} - Y_{t+\tau}^{(b^*)} \mid H_t = h_t]$ ) through only the blip prediction network:

$$\sum_{k=0}^\tau \mathbf{gb}_{\theta_B}^k(z_t^B)'(a_k^* - b_k^*), \quad \text{where } z_t^B = \mathcal{E}_{\theta_B}^B(h_t) \quad (14)$$

**Efficient offline evaluation:** Once we have estimated the blip coefficients, then we can instantly identify the treatment sequence  $a^*$  with the best effect compared to the baseline  $b^*$  (e.g., a treatment sequence with no interventions). The reason is that the blip coefficients do *not* depend on treatments. Hence, our DeepBlip is much more efficient for

evaluating the personalized effects of different treatment sequences compared to existing methods that require re-computation (Bica et al., 2020; Lim, 2018; Melnychuk et al., 2022) or even re-training (Hess et al., 2025; Rui Li et al., 2021). This is highly relevant in personalized medicine where clinicians and patients jointly reason about different treatment strategies (Feuerriegel et al., 2024).

**Implementation details.** We instantiate our DeepBlip with a transformer architecture (see Appendix I). We also provide a variant based on an LSTM, which, despite the simpler architecture, is still highly competitive (see Appendix E.1).

## 5 EXPERIMENTS

Our experiments primarily serve two purposes to demonstrate DeepBlip is an effective neural framework for SNMMs: (i) we aim to show that our DeepBlip performs adequate adjustment for time-varying confounding, which is a key property of SNMMs; and (ii) we aim to show that our DeepBlip is especially robust for longer time horizons, which is a natural strength of SNMMs.

**Baselines:** We demonstrate the performance of our DeepBlip against key baselines from the literature (see Table 1)<sup>5</sup> for the task of estimating CATE (or conditional average potential outcomes) on medical datasets. Descriptions of the baseline methods are available in Appendix G. We select the HA-PI-learner from Frauen et al. (2025) instantiated by transformer (named **HA-TRM**) as a naïve baseline. We provide additional implementation details – including architecture choices, training procedures, and hyperparameter tuning – in Appendix I. To ensure a fair comparison, all methods – including baselines – use the **same** neural backbone architecture, so any performance differences can be attributed to the respective learning frameworks. All results are averaged over 5 runs.

**Ablations:** We include ablation studies of both the architecture and the double optimization trick in Appendix E.1. We also validate our component-wise blip coefficient estimates in Appendix E.2.

### 5.1 TUMOR GROWTH DATASET

**Setting:** We use the pharmacokinetic-pharmacodynamic tumor growth dataset (Geng et al., 2017), which is commonly used for benchmarking CATE methods over time (Bica et al., 2020; Hess et al., 2025; Lim, 2018; Melnychuk et al., 2022; Rui Li et al., 2021). The dataset describes the time-varying effects of chemotherapies and radiotherapies on cancer volumes. The treatment assignments depend on previous outcomes, subject to time-varying confounding. The amount of confounding is controlled by the simulation parameter  $\gamma_{\text{conf}}$ . Details are in Appendix F.1.

**Results:** Figure 3 shows the average RMSE of CATE against increasing confounding  $\gamma_{\text{conf}}$  and under  $\tau = 2$ . Our **DeepBlip** outperforms all baselines for  $\gamma_{\text{conf}} \geq 2$ . This matches the purpose of our method to deal with time-varying confounding. More importantly, our DeepBlip achieves large performance gains under strong confounding levels ( $\gamma_{\text{conf}} > 6$ ). This highlights that our DeepBlip is robust against time-varying confounding by providing adequate adjustment for time-varying confounding.

We could further make the following observations: ① The MSM-based **R-MSN** performs poorly across all confounding levels and even has a higher RMSE than HA-LSTM for  $\gamma_{\text{conf}} \leq 6$ . This aligns with the inherent instability of inverse propensity weighting in MSMs, which was the motivation for our method. ② Baselines like **CT** and **CRN** that use balanced representation (in orange) are ineffective. This is expected as balanced representations were originally developed for reducing finite-sample estimation variance and *not* for proper adjustment (see the original work (Shalit et al., 2017) on balanced representations for a discussion), because of which these baselines are known to be biased. ③ G-computation-based methods like **G-Net** and **GT** show slightly lower RMSE for  $\gamma_{\text{conf}} \leq 1$  but still perform

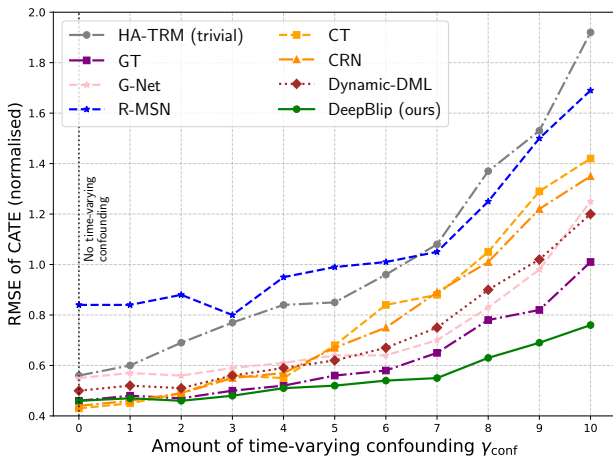


Figure 3: **Results for tumor growth dataset.** Normalized RMSE (averaged over 5 runs) of CATE predictions against ground-truth over growing confounding. Here:  $\tau = 2$

<sup>5</sup>We also compare against Dynamic DML (Lewis & Syrgkanis, 2021); In principle, it does not qualify as a standard baseline because it cannot handle patient histories  $H_t$  of varying lengths. Nonetheless, to provide a thorough comparison, we adapt it by training a separate regression model for each individual time step  $t$ .

significantly worse than our DeepBlip for  $\gamma_{\text{conf}} \geq 6$ . We attribute this to the fact that the learning is unstable, which we empirically verify in the following by varying the prediction horizon  $\tau$ . ④ Dynamic DML is competitive, thereby showing that SNMMs are a suitable modeling framework for this task, while our DeepBlip is better by a clear margin.

## 5.2 MIMIC-III DATASET

**Setting:** Next, we evaluate the performance for longer prediction horizons  $\tau$ . We build upon MIMIC-III (Johnson et al., 2016), a widely used benchmark for evaluating CATE over time (Bica et al., 2020; Hess et al., 2025; Melnychuk et al., 2022). Following previous literature (Hess et al., 2025; Melnychuk et al., 2022; Schulam & Saria, 2017), we extract patient vitals from MIMIC-III and then simulate the patient outcome over time with the mixed dynamics of exogenous dependency, endogenous dependency, and treatment effects combined. Treatments are assigned based on previous outcomes and patient vitals, which again, introduces time-varying confounding (see Appendix F.2).

**Results:** Table 2 shows the average RMSE (with std. dev.) over five different runs with  $\gamma_{\text{conf}} = 1$  and varying prediction horizon  $\tau$ . First, our **DeepBlip** consistently achieves lower RMSE compared to other baselines for  $\tau \geq 3$ . Of note, the performance gain from DeepBlip becomes larger as the prediction horizon  $\tau$  increases. For  $\tau = 10$ , DeepBlip achieves  $\sim 34\%$  performance gain compared to the second-best model (here: GT). This highlights that our DeepBlip is temporally stable over longer horizons.

We further make the observations that all baselines either struggle with high-dimensional covariates or become unstable as  $\tau$  increases. ① The MSM-based method (**R-MSNs**) exhibits the highest variance across all  $\tau$ , indicating that it struggles with high-dimensional propensity modeling and becomes unstable over time with increasing standard deviation. The reason is that inverse propensity weighting produces unstable weights. ② Methods like **CRN** and **CT** perform better than baselines with high variance like R-MSNs due to the way they handle high-dimensional covariates. However, both **CRN** and **CT** are known to be biased and thus inferior to GT and our DeepBlip. ③ G-computation-based methods (i.e., **G-Net** and **GT**) achieve a lower RMSE than the other baselines due to proper adjustments, but still are not as stable as our method. This is because G-computation accumulates error over time due to modeling nested expectations. ④ The performance gain of DeepBlip over Dynamic DML is large, which can be attributed to the complexity of the dataset, especially in light of the long time windows.

**Scalability of DeepBlip:** To demonstrate the computational advantage of our framework, we tested the second-stage training time of both **Dynamic DML** and our **DeepBlip** over increasing  $\tau$ . The results are in Figure 4: we observed a large difference in scalability. The training time for Dynamic DML grows linearly with  $\tau$ , due to its sequential scheme, which requires training a separate model at *each* time step. In contrast, the runtime for DeepBlip remains *nearly constant*, which confirms the efficiency of our double optimization trick and thus demonstrates that our DeepBlip overcomes the scalability bottleneck of SNMMs.

	$\tau = 3$	$\tau = 4$	$\tau = 5$	$\tau = 6$	$\tau = 7$	$\tau = 8$	$\tau = 9$	$\tau = 10$
HA-LSTM (naïve)	0.89 ± 0.03	0.97 ± 0.04	1.02 ± 0.10	1.42 ± 0.20	1.92 ± 0.40	2.57 ± 0.44	2.58 ± 0.56	3.11 ± 0.72
R-MSNs	0.98 ± 0.17	1.12 ± 0.21	1.25 ± 0.28	1.65 ± 0.57	2.25 ± 1.02	2.85 ± 1.18	3.20 ± 1.42	3.55 ± 1.50
CRN	0.66 ± 0.11	0.82 ± 0.12	1.05 ± 0.22	1.22 ± 0.35	1.43 ± 0.33	1.62 ± 0.42	1.83 ± 0.43	2.04 ± 0.54
CT	0.64 ± 0.12	0.79 ± 0.11	1.01 ± 0.18	1.18 ± 0.33	1.77 ± 0.52	1.85 ± 0.49	1.99 ± 0.63	1.98 ± 0.60
G-Net	0.58 ± 0.08	0.73 ± 0.12	1.05 ± 0.25	1.38 ± 0.40	1.75 ± 0.60	2.15 ± 0.80	2.55 ± 0.90	3.12 ± 1.05
GT	0.52 ± 0.02	0.63 ± 0.08	0.75 ± 0.17	0.85 ± 0.13	0.95 ± 0.26	1.10 ± 0.34	1.25 ± 0.37	1.50 ± 0.45
Dynamic DML	0.72 ± 0.15	0.88 ± 0.23	1.12 ± 0.27	1.34 ± 0.68	1.67 ± 0.43	1.98 ± 0.55	2.33 ± 0.57	2.46 ± 0.72
DeepBlip (ours)	<b>0.50 ± 0.12</b>	<b>0.58 ± 0.16</b>	<b>0.64 ± 0.19</b>	<b>0.75 ± 0.21</b>	<b>0.86 ± 0.24</b>	<b>0.89 ± 0.27</b>	<b>0.95 ± 0.28</b>	<b>0.98 ± 0.32</b>
Improvement	3.8%	8.9%	14.7%	11.7%	9.5%	19.1%	24.0%	34.6%

Table 2: **MIMIC-III with longer time horizons  $\tau$** . Normalized RMSE (mean  $\pm$  std. dev. over 5 runs) for  $\tau$ -step-ahead CATE estimation on the MIMIC-III dataset. We highlight the relative improvement over the best-performing baseline.




Figure 4: **Scalability of Dynamic DML vs. DeepBlip**. Shown is the training time.

**Ablations on the transformer architecture:** To further verify our framework’s robustness, we replace the transformer architecture by an LSTM (Hochreiter & Schmidhuber, 1997) for our DeepBlip and other baselines in the ablation study (see Appendix E.1). Even with a simpler architecture, our model is highly competitive and outperforms some of the transformer-based baselines (see Appendix E.1), underscoring the effectiveness of our framework.

## 6 CONCLUSION

We introduced DeepBlip, the first neural framework for SNMMs that overcomes the scalability issues of prior works via a double optimization trick. By using blip functions to decompose total effects into localized increments, our method provides unbiased estimates that are stable over long horizons. Experiments across clinical datasets confirm that DeepBlip achieves state-of-the-art performance, offering a powerful and scalable tool for personalized medicine.

---

## ETHICS STATEMENT

Our research aims to advance personalized medicine by developing reliable methods for treatment effect estimation. Our work relies on the MIMIC-III dataset, which is a publicly available, de-identified critical care database. All research was conducted in compliance with the data use agreement for this dataset, which was created with the approval of the Massachusetts Institute of Technology (MIT) and Beth Israel Deaconess Medical Center (BIDMC).

We acknowledge several ethical considerations in this work. First, any model trained on historical clinical data might exhibit discrimination present in the data. Hence, we stress that any deployment of such a model would require rigorous auditing for fairness and performance equity across different demographic and social subgroups. Second, we emphasize that our DeepBlip is intended to function as a decision-support tool to aid clinicians, not to replace them. Finally, before any consideration for clinical use, the safety and reliability of this method must be tested through clinical trials. We thus emphasize the importance of following a responsible development of machine learning tools for healthcare and advocate for a cautious, evidence-based approach to their translation into practice.

## REFERENCES

- Ahmed M. Alaa and Mihaela van der Schaar. Bayesian inference of individualized treatment effects using multi-task Gaussian processes. URL <http://arxiv.org/pdf/1704.02801>.
- Ahmed Allam, Stefan Feuerriegel, Michael Rebhan, and Michael Krauthammer. Analyzing patient trajectories with artificial intelligence. *Journal of Medical Internet Research*, pp. e29812, 2021. doi: 10.2196/29812. URL <https://www.jmir.org/2021/12/e29812>.
- Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar. Estimating Counterfactual Treatment Outcomes over Time Through Adversarially Balanced Representations. In *ICLR, 2020*.
- Ioana Bica, Ahmed M. Alaa, Christoph Lambert, and Mihaela van der Schaar. From real-world patient data to individualized treatment effects using machine learning: Current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics*, 2021. doi: 10.1002/cpt.1907. URL <https://doi.org/10.1002/cpt.1907>.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/Debiased Machine Learning for Treatment and Causal Parameters. *The Econometrics Journal*, 2018.
- Amanda Coston, Edward Kennedy, and Alexandra Chouldechova. Counterfactual predictions under runtime confounding. In *Advances in Neural Information Processing Systems*, 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/2b64c2f19d868305aa8bbc2d72902cc5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/2b64c2f19d868305aa8bbc2d72902cc5-Paper.pdf).
- Alicia Curth and Mihaela van der Schaar. Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*, 13–15 Apr 2021. URL <https://proceedings.mlr.press/v130/curth21a.html>.
- William Falcon, Jirka Borovec, Adrian Wälchli, Katharina Eggensperger, Roman Schiele, Massimo Patacchiola, Ryutaro Tanno, Alexander J. Landgraf, Chad Baskin, Jeff Jordan, Raghuvvar Sinha, et al. PyTorch Lightning. In *NeurIPS 2019 Reproducibility Challenge*, 2019. URL <https://github.com/Lightning-AI/lightning>.
- Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Alicia Curth, Stefan Bauer, Niki Kilbertus, Isaac S. Kohane, and Mihaela van der Schaar. Causal machine learning for predicting treatment outcomes. *Nature Medicine*, 04 2024. ISSN 1546-170X. doi: 10.1038/s41591-024-02902-1. URL <https://doi.org/10.1038/s41591-024-02902-1>.
- Dylan J. Foster and Vasilis Syrgkanis. Orthogonal statistical learning. In *Proceedings of the 32nd Conference on Learning Theory (COLT)*, pp. 1362–1385, 2019. URL <https://arxiv.org/abs/1901.09036>. arXiv preprint arXiv:1901.09036.
- Dennis Frauen, Tobias Hatt, Valentyn Melnychuk, and Stefan Feuerriegel. Estimating average causal effects from patient trajectories. 37:7586–7594, Jun. 2023. doi: 10.1609/aaai.v37i6.25921. URL <https://ojs.aaai.org/index.php/AAAI/article/view/25921>.

- 540 Dennis Frauen, Konstantin Hess, and Stefan Feuerriegel. Model-agnostic meta-learners for estimating heterogeneous  
541 treatment effects over time. In *International Conference on Learning Representations (ICLR)*, 2025. URL <https://arxiv.org/abs/2407.05287>.  
542
- 543 Changran Geng, Harald Paganetti, and Clemens Grassberger. Prediction of Treatment Response for Combined Chemo-  
544 and Radiation Therapy for Non-Small Cell Lung Cancer Patients Using a Bio-Mathematical Model. *Scientific Reports*,  
545 2017. doi: 10.1038/s41598-017-13646-z. Author correction published in *Scientific Reports*, 2018, 8(1):12631,  
546 doi:10.1038/s41598-018-30761-7.  
547
- 548 Miguel A. Hernán, Babette Brumback, and James M. Robins. Marginal structural models to estimate the joint causal  
549 effect of nonrandomized treatments. *Journal of the American Statistical Association*, 2001. ISSN 01621459,  
550 1537274X. URL <http://www.jstor.org/stable/2670282>.
- 551 Konstantin Hess, Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Bayesian neural controlled differential  
552 equations for treatment effect estimation, 2024. URL <https://arxiv.org/abs/2310.17463>.  
553
- 554 Konstantin Hess, Dennis Frauen, Valentyn Melnychuk, and Stefan Feuerriegel. G-transformer for conditional average  
555 potential outcome estimation over time. In *NeurIPS*, 2025. URL <https://arxiv.org/abs/2405.21012>.  
556 arXiv preprint arXiv:2405.21012.
- 557 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 1997.  
558
- 559 Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li wei H. Lehman, Mengling Feng, Mohammad Mahdi Ghassemi,  
560 Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. MIMIC-III, a freely accessible critical care  
561 database. *Scientific Data*, 2016. URL <https://api.semanticscholar.org/CorpusID:33285731>.
- 562 Edward H. Kennedy. Towards Optimal Doubly Robust Estimation of Heterogeneous Causal Effects. *Electronic Journal*  
563 *of Statistics*, 2023.  
564
- 565 Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Meta-learners for estimating heterogeneous treatment  
566 effects using machine learning. *Proceedings of the National Academy of Sciences*, 116:4156–4165, 2019. ISSN  
567 1091-6490. doi: 10.1073/pnas.1804597116. URL <http://dx.doi.org/10.1073/pnas.1804597116>.
- 568 Greg Lewis and Vasilis Syrgkanis. Double/Debiased Machine Learning for Dynamic Treatment Effects via g-Estimation.  
569 In *Advances in Neural Information Processing Systems*, Virtual, Dec 2021. URL [https://proceedings.](https://proceedings.neurips.cc/paper/2021/file/db72d405246268496e0580e426ca7f93-Paper.pdf)  
570 [neurips.cc/paper/2021/file/db72d405246268496e0580e426ca7f93-Paper.pdf](https://proceedings.neurips.cc/paper/2021/file/db72d405246268496e0580e426ca7f93-Paper.pdf).  
571
- 572 Bryan Lim. Forecasting treatment responses over time using recurrent marginal structural networks. In *Advances*  
573 *in Neural Information Processing Systems*, volume 31, 2018. URL [https://proceedings.neurips.cc/](https://proceedings.neurips.cc/paper_files/paper/2018/file/56e6a93212e4482d99c84a639d254b67-Paper.pdf)  
574 [paper\\_files/paper/2018/file/56e6a93212e4482d99c84a639d254b67-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/56e6a93212e4482d99c84a639d254b67-Paper.pdf).
- 575 Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Causal Transformer for Estimating Counterfactual  
576 Outcomes. In *ICML*, 2022.
- 577 Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Bounds on representation-induced confounding bias for  
578 treatment effect estimation. In *ICLR*, 2024. URL <https://openreview.net/forum?id=1n4rPUq1lJ>.  
579 arXiv:2311.11321.  
580
- 581 Chuizheng Meng, Yihe Dong, Sercan Ö Arık, Yan Liu, and Tomas Pfister. COSTAR: Improved Temporal Counterfactual  
582 Estimation with Self-Supervised Learning, February 2024. URL <http://arxiv.org/abs/2311.00886>.
- 583 Liliana Orellana, Andrea Rotnitzky, and James M. Robins. Dynamic regime marginal structural mean models for  
584 estimation of optimal dynamic treatment regimes, part i: main content. *The International Journal of Biostatistics*,  
585 2010. URL <https://pubmed.ncbi.nlm.nih.gov/21969994/>.  
586
- 587 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor  
588 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Ed-  
589 ward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,  
590 Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance  
591 deep learning library. In *NeurIPS*, 2019. URL [http://papers.neurips.cc/paper/](http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf)  
592 [9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf](http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf).  
593
- 594 Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning*  
595 *Algorithms*. MIT Press, 2017. ISBN 0262037319.

594 Maya L. Petersen and Mark J. van der Laan. Causal models and learning from data: integrating causal modeling  
595 and statistical estimation. *Epidemiology*, 2014. ISSN 1044-3983. doi: 10.1097/EDE.000000000000078. URL  
596 <https://pmc.ncbi.nlm.nih.gov/articles/PMC4077670/>.  
597

598 J. M. Robins, M. A. Hernan, and B. Brumback. Marginal structural models and causal inference in epidemiology.  
599 *Epidemiology*, 2000. ISSN 1044-3983. doi: 10.1097/00001648-200009000-00011. URL [https://pubmed.  
600 ncbi.nlm.nih.gov/10955408/](https://pubmed.ncbi.nlm.nih.gov/10955408/).

601 James Robins. A new approach to causal inference in mortality studies with a sustained exposure period. *Mathematical  
602 Modelling*, January 1986. ISSN 0270-0255. doi: 10.1016/0270-0255(86)90088-6.  
603

604 James M. Robins. Correcting for non-compliance in randomized trials using structural nested mean models. *Communi-  
605 cations in Statistics - Theory and Methods*, 1994. ISSN 0361-0926. doi: 10.1080/03610929408831393.  
606

607 James M. Robins. Optimal structural nested models for optimal sequential decisions. In *Proceedings  
608 of the Second Seattle Symposium in Biostatistics*. 2004. ISBN 978-1-4419-9076-1. doi: 10.1007/  
609 978-1-4419-9076-1\_{\text{underscore}}11. URL [https://link.springer.com/chapter/10.1007/  
610 978-1-4419-9076-1\\_11](https://link.springer.com/chapter/10.1007/978-1-4419-9076-1_11).

611 James M. Robins, Sander Greenland, and Fu-Chang Hu. Estimation of the causal effect of a time-varying exposure  
612 on the marginal mean of a repeated binary outcome. *Journal of the American Statistical Association*, 1999. ISSN  
613 01621459. doi: 10.2307/2669978. URL <http://www.jstor.org/stable/2669978>.

614 Donald Rubin. Estimating causal effects of treatments in experimental and observational studies. *ETS Research Bulletin  
615 Series*, 1972. ISSN 2333-8504. doi: 10.1002/j.2333-8504.1972.tb00631.x.  
616

617 Rui Li, Stephanie Hu, Mingyu Lu, Yuria Utsumi, Prithwish Chakraborty, Daby M. Sow, Piyush Madan, Jun Li,  
618 Mohamed Ghalwash, Zach Shahn, and Li-wei Lehman. G-Net: a recurrent network approach to G-Computation for  
619 counterfactual prediction under a dynamic treatment regime. *Machine Learning for Health*, 2021. ISSN 2640-3498.  
620 URL <https://proceedings.mlr.press/v158/li21a.html>.

621 Peter Schulam and Suchi Saria. Reliable Decision Support using Counterfactual Models. In *Proceedings of the  
622 31st International Conference on Neural Information Processing Systems*, pp. 6076–6086, Long Beach, CA,  
623 USA, 04–09 Dec 2017. Curran Associates, Inc. URL [https://papers.nips.cc/paper/2017/file/  
624 5b69b9cb83065d403869739ae7f3325e-Paper.pdf](https://papers.nips.cc/paper/2017/file/5b69b9cb83065d403869739ae7f3325e-Paper.pdf).

625

626 Nabeel Seedat, Fergus Imrie, Alexis Bellot, Zhaozhi Qian, and Mihaela van der Schaar. Continuous-Time Modeling of  
627 Counterfactual Outcomes Using Neural Controlled Differential Equations. In *ICML*, June 2022.  
628

629 Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating Individual Treatment Effect: Generalization Bounds  
630 and Algorithms. In Doina Precup and Yee Whye Teh (eds.), *ICML*, volume 70 of *Proceedings of Machine Learn-  
631 ing Research*, pp. 3076–3085. PMLR, 06–11 Aug 2017. URL [https://proceedings.mlr.press/v70/  
632 shalit17a.html](https://proceedings.mlr.press/v70/shalit17a.html).

633 Toru Shirakawa, Yi Li, Yulun Wu, Sky Qiu, Yuxuan Li, Mingduo Zhao, Hiroyasu Iso, and Mark van der Laan.  
634 Longitudinal targeted minimum loss-based estimation with temporal-difference heterogeneous transformer. *arXiv  
635 preprint arXiv:2404.04399*, apr 2024. doi: 10.48550/arXiv.2404.04399. URL [https://arxiv.org/abs/  
636 2404.04399](https://arxiv.org/abs/2404.04399).

637 Hossein Soleimani, Adarsh Subbaswamy, and Suchi Saria. Treatment-response models for counterfactual reasoning  
638 with continuous-time, continuous-valued interventions. In *UAI*, Sydney, Australia, 2017. AUAI Press.  
639

640 Stijn Vansteelandt and Marshall Joffe. Structural nested models and g-estimation: The partially realized promise.  
641 *Statistical Science*, 29, November 2014. ISSN 0883-4237. doi: 10.1214/14-sts493. URL [http://dx.doi.org/  
642 10.1214/14-sts493](http://dx.doi.org/10.1214/14-sts493).

643 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia  
644 Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.  
645

646 Stefan Wager and Susan Athey and. Estimation and inference of heterogeneous treatment effects using random  
647 forests. *Journal of the American Statistical Association*, 2018. doi: 10.1080/01621459.2017.1319839. URL  
<https://doi.org/10.1080/01621459.2017.1319839>.

---

648 Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48 of *Cambridge Series in*  
649 *Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, UK, 2019. ISBN 9781108498029.  
650 doi: 10.1017/9781108627771.

651 Haotian Wang, Haoxuan Li, Hao Zou, Haoang Chi, Long Lan, Wanrong Huang, and Wenjing Yang. Effective and  
652 Efficient Time-Varying Counterfactual Prediction with State-Space Models. In *ICLR*, 2025a.

653 Shirly Wang, Matthew B. A. McDermott, Geeticka Chauhan, Marzyeh Ghassemi, Michael C. Hughes, and Tristan  
654 Naumann. MIMIC-Extract: a data extraction, preprocessing, and representation pipeline for MIMIC-III. In  
655 *Proceedings of the ACM Conference on Health, Inference, and Learning*, ACM CHIL '20, April 2020. doi:  
656 10.1145/3368555.3384469. URL <http://dx.doi.org/10.1145/3368555.3384469>.

657 Xin Wang, Shengfei Lyu, Lishan Yang, Yibing Zhan, and Huanhuan Chen. A Dual-module Framework for Counterfac-  
658 tual Estimation over Time. In *ICML*, July 2024.

659 Xin Wang, Shengfei Lyu, Chi Luo, Xiren Zhou, and Huanhuan Chen. Variational Counterfactual Intervention Planning  
660 to Achieve Target Outcomes. In *ICML*, 2025b.

661 Yanbo Xu, Yanxun Xu, and Suchi Saria. A non-parametric bayesian approach for estimating treatment-response curves  
662 from sparse time series. In *MLAH*, Proceedings of Machine Learning Research, Northeastern University, Boston,  
663 MA, USA, aug 2016. URL <https://proceedings.mlr.press/v56/Xu16.html>.

664 Jinsung Yoon, James Jordon, and Mihaela van der Schaar. GANITE: Estimation of individualized treatment effects using  
665 generative adversarial nets. In *ICLR*, 2018. URL <https://openreview.net/forum?id=ByKWUeWA->.

666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701

## 702 A PRELIMINARIES

703 In this section we briefly introduce the preliminaries of causal inference.

### 704 A.1 CAUSAL MODELING AND TREATMENT EFFECT

705 **Structural causal model** A structural causal model is a tuple  $\mathcal{M} = (\mathcal{U}, \mathcal{V}, \mathcal{F}, P)$  where:

- 706 1.  $\mathcal{U}$  is a set of exogenous variables, which are not caused by any other variables, namely the unknown background factors or simply noises. Each  $U_i \in \mathcal{U}$  has a value domain denoted  $\text{dom}(U)$ .  $\mathcal{U}$  is defined on a probability space  $(\Omega, \mathcal{B}, P)$ .
- 707 2.  $\mathcal{V}$  is a set of endogenous variables that are determined within the model from other variables. The domain of  $V \in \mathcal{V}$  is  $\text{dom}(V)$ .
- 708 3.  $\mathcal{F}$  is a set of structural functions:  $\mathcal{F} = \{f_i : \text{dom}(\text{Pa}_i) \times \text{dom}(U_i) \rightarrow \text{dom}(V_i) \mid V_i \in \mathcal{V}\}$ , where  $\text{Pa}_i$  is the set of endogenous variable that directly influences  $V_i$ , and  $U_i \in \mathcal{U}$  is the exogenous variable that directly influences  $V_i$ . Here,  $f_i$  is a deterministic function that describes the structural causal mechanism for generating  $V_i$ :  $V_i = f_i(\text{Pa}_i, U_i)$ .

709 Furthermore, the causal relationships specified by  $\mathcal{F}$ 's  $\{\text{Pa}_i \mid V_i \in \mathcal{V}\}$  induce a directed graph  $\mathcal{G} = (\mathcal{V}, E)$ , for which nodes correspond to the endogenous variables and  $(V_i, V_j) \in E \iff V_i \in \text{Pa}_j$ . It is important to note that the causal graph  $\mathcal{G}$  must be a **directed acyclic graph (DAG)**, meaning that there should not be a causal cycle, which is contradictory to both real world and mathematics.

710 For SCM, the joint probability distribution  $P$  induces the probability on the whole variable sets. Once the set  $\mathcal{U}$ 's variables takes a fixed value, the whole endogenous set's values are determined through  $\mathcal{F}$ . Let  $P_M$  be the joint probability distribution over  $\mathcal{V}$ , then for any value  $v \in \prod_{V_i \in \mathcal{V}} \text{dom}(V_i)$ :

$$711 P_M(V = v) = \Pr(\{u \in \prod_{U_i \in \mathcal{U}} \text{dom}(U_i) \mid \mathcal{F}(u) = v\}) \quad (15)$$

712 **Intervention:** Intervention could be defined under the context of SCM as forcing a subset of variables (here we assume endogenous)  $X \subset \mathcal{V}$  to take a specific value  $x \in \prod_{V_i \in X} \text{dom}(V_i)$ <sup>6</sup>, independent of the original induced probability. This action simulates the process of controlled experiment or a hypothetical scenario.

713 An intervention, denoted as  $\text{do}(X = x)$ , creates a **new** SCM  $\mathcal{M}_x = (\mathcal{U}, \mathcal{V}, \mathcal{F}_x, P)$ , which is a mutation from the original SCM  $\mathcal{M}$ , where:

- 714 1. The set of endogenous  $\mathcal{V}$  and exogenous  $\mathcal{U}$  remain unchanged with the same probability space  $(\Omega, \mathcal{B}, P)$ .
- 715 2. For each  $V_i \in X$ , the old structural equation  $V_i = f_i(\text{Pa}_i, U_i)$  is replaced by a degenerated constant equation  $V_i \equiv x_i$ .
- 716 3. For all the other endogenous variables, the mapping  $f_i$  stays unchanged, but their distributions may adjust.

717 On the whole, the intervention  $\text{do}(X = x)$  induces a new probability distribution  $P_{M_x}$ . The interventional distribution of a variable, say  $Y \in \mathcal{V}$ , is denoted as  $\Pr(Y \mid \text{do}(X = x))$  or  $\Pr(Y^{\text{do}(X=x)})$ . Under certain context for simplicity, we could directly use  $\Pr(Y^{(x)})$  without ambiguity.

718 **Potential Outcome and Treatment effect:** From an intuition perspective, intervention is like cutting the causal links to  $X$  from the original diagram, setting the value to  $x$  and then observe how the system responds. We emphasize that  $\Pr(Y \mid \text{do}(X = x)) \neq \Pr(Y \mid X = x)$ , which is just a conditional distribution under the original SCM.

719 A **potential outcome (PO)** of a variable  $Y \in \mathcal{V}$  is the value under a hypothetical scenario, where we make an intervention  $\text{do}(X = x)$ , given a specific realization of  $U = u$ . We denote this as  $Y_x(u)$ . When  $U \sim P$ . This random variable  $Y_x(U)$  is just  $Y^{\text{do}(X=x)}$ , which follows the distribution  $\Pr(Y \mid \text{do}(X = x))$ . The average potential outcome (APO) is the expectation of PO:  $\mathbb{E}[Y^{\text{do}(X=x)}]$ . A **conditional average potential outcome (CAPO)** is nothing but the expectation of the random variable  $Y_x(U)$  conditioned on some subgroups  $W = w$ , where  $W \subset \mathcal{V}$ . This conditioning on  $W = w$  could be traced back to a posterior distribution on  $U$ , namely  $\Pr(U \mid W = w) := P_w$ , then the CAPO is just  $\mathbb{E}_{U_w \sim P_w}[Y_x(U_w)]$  or equivalently  $\mathbb{E}[Y^{(x)} \mid W = w]$ .

720 <sup>6</sup>General intervention could be defined as setting some variables to follow a specific **distribution**(Peters et al., 2017), however we follow a common practice to assume the intervention means forcing a fixed value

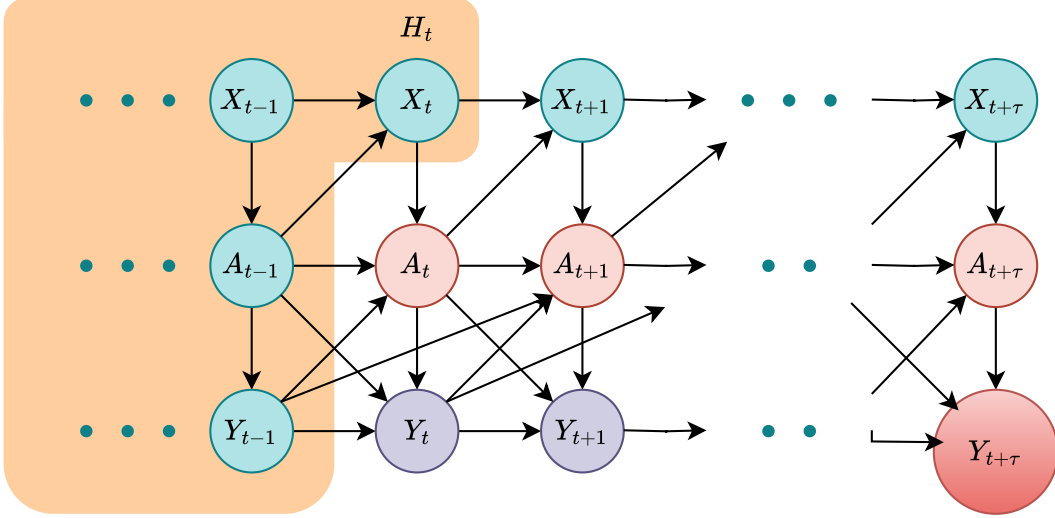


Figure 5: An example of causal diagram for the SCM in the setting of estimating CATE over time. The variable sequences are truncated at  $t + \tau$  (maximal length is  $T$ ). Exogenous variables  $\mathcal{U}$  are hidden.

An **average treatment effect (ATE)** is the expectation difference between two POs under intervention  $x_1$  and  $x_2$  respectively:  $\mathbb{E}[Y^{(x_1)} - Y^{(x_2)}]$ . A **conditional average treatment effect (CATE)** is then the difference between two CAPOs:  $\mathbb{E}[Y^{(x_1)} - Y^{(x_2)} | W = w]$ .

So far we have strictly defined these foundational concepts. Overall, the potential outcomes are counterfactual variables in intervened SCM while treatment effects are the contrast between these potential outcomes. Within this rigorous framework, we are able to clearly define the goal for our work on estimating CATE over time conditioned on patient history  $H_t$  as in Sec. 3.

## A.2 SCM FOR CATE ESTIMATION OVER TIME

We formalize the setting for estimating CATE over time with the following structural causal model  $\mathcal{M}$ :

1. Exogenous variables:  $\mathcal{U} = \{U_{X_t}, U_{A_t}, U_{Y_t}\}_{t=1}^T$ , where  $U_{X_t}, U_{A_t}, U_{Y_t}$  are mutually independent noise.
2. Endogenous variables are  $\mathcal{V} = \{X_t, A_t, Y_t\}_{t=1}^T$ . Denote  $H_t = \{X_s\}_{s=1}^t \cup \{A_j\}_{j=1}^{t-1} \cup \{Y_i\}_{i=1}^{t-1}$  as the history up until time  $t$ .
3. Structural functions  $\mathcal{F}$ : for  $t \in \{1, \dots, T\}$

$$X_t = f_{X_t}(H_{t-1} \cup \{A_{t-1}\}, U_{X_t})$$

$$A_t = f_{A_t}(H_t, U_{A_t})$$

$$Y_t = f_{Y_t}(H_t, A_t, U_{Y_t})$$

These equations generate the full observed patient trajectory in the following temporal order:  $(X_1 \rightarrow A_1 \rightarrow Y_1 \rightarrow \dots \rightarrow X_T \rightarrow A_T \rightarrow Y_T)$ .

A possible instantiation of temporal graph of this given structural causal model is presented in Figure.5 When talking about SCMs, it is important to point out that SCM represents perfect knowledge of the data generating process. In real-world settings, there are possibly omitted variables (e.g. unmeasured confounders).

**Time-varying treatment and potential outcomes:** A **time-varying intervention** is intervention on a continuous sequence of treatment variables  $\text{do}(A_{t:t+\tau} = a_{t:t+\tau})$ , which modifies the structural equations for  $A_t, \dots, A_{t+\tau}$ :

$$\forall k \in \{t, \dots, t + \tau\} : A_k \equiv a_k.$$

This induces a new SCM  $\mathcal{M}^{(a_{t:t+\tau})}$ , generating the potential outcome:

$$Y_{t+\tau}^{(a_{t:t+\tau})} = f_{Y_{t+\tau}} \left( H_{t+\tau}^{(a_{t:t+\tau})}, X_{t+\tau}^{(a_{t:t+\tau})}, a_{t+\tau}, U_{Y_{t+\tau}} \right),$$

where  $H_{t+\tau}^{(a_{t:t+\tau})}$ ,  $X_{t+\tau}^{(a_{t:t+\tau})}$  are recursively computed under the intervention using structural equations defined in 3.

**Identification assumptions:** The structural functions, although might be unspecified, already contains rich information about the data generating process. It is easy to verify that under the specified data generating process and the mutual independence between the exogenous variables, the following two conditions are satisfied:

1. **Consistency:** If  $A_{t:t+\tau} = a_{t:t+\tau}$  is the observed treatment sequence for a given patient  $u$ , then  $Y_{t+\tau}^{(a_{t:t+\tau})}(u) = Y(u)$ . Or, equivalently  $Y_{t+\tau}^{(A_{t:t+\tau})} = Y_{t+\tau}$ . This means when we intervene with the observed treatment, the potential outcome equals the observed outcome.
2. **Sequential ignorability:** The potential outcome under a fixed intervention is conditionally independent with the treatment assignment:

$$Y_t^{(a_t)} \perp A_t \mid H_t, \quad \forall a_t, 1 \leq t \leq T$$

This ignorability assumption implies that there are no hidden confounders that affect both outcome and the treatment, which is already decided with the SCM.

3. **Sequential overlap:** If  $\Pr(H_t = h_t) > 0$ , then:

$$\Pr(A_t = a_t \mid H_t = h_t) > 0 \quad \forall a_t \in \mathcal{A}_t, 1 \leq t \leq T.$$

These three assumptions together allow the use of observational data to compute causal quantities, such as potential outcome over time and treatment effect over time (Robins et al., 2000). Without these assumptions, exact identification would be impossible. While these assumptions themselves are untestable, they could be justified under appropriate study design and domain knowledge.

**Time-varying confounding:** Estimating treatment effect over time poses a significant challenge due the existence of time-varying confounding (Coston et al., 2020), which differentiate the task from the static setting ( $\tau = 0$ ). The underlying reason is that, the future time-varying confounders, such as  $X_{t+1:t+tau}$  or  $Y_{t:t+\tau-1}$ , are *unobserved* during the inference phase. As a result, the naïve adjustment with the history variable  $H_t$  is not sufficient for the backdoor criteria (Peters et al., 2017) and, thus, leads to biased estimation:

$$\text{When } \tau > 0 : \quad \mathbb{E} \left[ Y_{t+\tau}^{(a_{t:t+\tau})} \mid H_t = h_t \right] \neq \mathbb{E} [Y_{t+\tau} \mid A_{t:t+\tau} = a_{t:t+\tau}, H_t = h_t]. \quad (16)$$

## B THEORY OF STRUCTURAL NESTED MEAN MODELS

In this section, we show by strict causal reasoning how to use the structural nested mean models to estimate CATE over time. The theory consists of the foundations of SNMM (Robins, 2004) and a specific realization (Lewis & Syrgkanis, 2021), which we adopt for our method. We show the mathematical derivations for completeness. Next, in sec. B.2 we show that our loss is Neyman orthogonal. Finally, we briefly explain the mean squared error rate for the estimated blip coefficients which is proved by previous works in SNMM (Lewis & Syrgkanis, 2021).

To simplify the notation, we denote  $Y = Y_{t+\tau}$ . And we use  $\underline{W}_{t+k}$  to denote  $\underline{W}_{t+k:t+\tau}$  for any variable  $W \in \mathcal{V}$ .

### B.1 SNMM

**Identification:** The blip functions can be accumulated to identify the potential outcomes, as summarized by the theorem (Robins, 2004) below. We show the proof for completeness.

**Theorem 2. Adjustment via blip functions** (Theorem 3.1 from Robins (2004)) *Given a policy  $d$  between time  $t$  and  $t + \tau$ , the following identification holds under the sequential ignorability assumption in 2:*

$$\mathbb{E}\left[Y^{(d)} \mid H_t = h_t\right] = \mathbb{E}\left[Y + \sum_{k=0}^{\tau} \rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) \mid H_t = h_t\right] \quad (17)$$

where  $\rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) = \gamma_{t,k}(X_{t+1:t+k}, (A_{t:t+k-1}, d_{t+k}); h_t) - \gamma_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t)$ .

*Proof.* By the definition of  $\gamma_{t,k}$ , we could write  $\gamma_{t,k}$  into the following form:

$$\begin{aligned} & \gamma_{t,k}(x_{t+1:t+k}, (a_{t:t+k-1}, d_{t+k}); h_t) \\ \stackrel{\text{Def.}}{=} & \mathbb{E}\left[Y^{(a_{t:t+k-1}, d_{t+k})} - Y^{(a_{t:t+k-1}, 0, d_{t+k+1})} \mid X_{t+1:t+k} = x_{t+1:t+k}, A_{t:t+k-1} = a_{t:t+k-1}, A_{t+k} = d_{t+k}, H_t = h_t\right] \\ \stackrel{\text{Ignor.}}{=} & \mathbb{E}\left[Y^{(a_{t:t+k-1}, d_{t+k})} - Y^{(a_{t:t+k-1}, 0, d_{t+k+1})} \mid X_{t+1:t+k} = x_{t+1:t+k}, A_{t:t+k} = a_{t:t+k}, H_t = h_t\right] \end{aligned} \quad (18)$$

Similarly:

$$\begin{aligned} & \gamma_{t,k}(x_{t+1:t+k}, a_{t:t+k}; h_t) \\ \stackrel{\text{Def.}}{=} & \mathbb{E}\left[Y^{(a_{t:t+k}, d_{t+k+1})} - Y^{(a_{t:t+k-1}, 0, d_{t+k+1})} \mid X_{t+1:t+k} = x_{t+1:t+k}, A_{t:t+k} = a_{t:t+k}, H_t = h_t\right] \end{aligned} \quad (19)$$

This implies that:

$$\begin{aligned} & \rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) \\ & = \gamma_{t,k}(X_{t+1:t+k}, (A_{t:t+k-1}, d_{t+k}); h_t) - \gamma_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) \\ & \stackrel{\text{By 18,19}}{=} \mathbb{E}\left[Y^{(A_{t:t+k-1}, d_{t+k})} - Y^{(A_{t:t+k}, d_{t+k+1})} \mid X_{t+1:t+k}, A_{t:t+k}, H_t = h_t\right] \end{aligned} \quad (20)$$

By consistency 1,  $Y^{(A_{t:t+\tau})} = Y$ , and that by towering law of the conditional expectation:

$$\begin{aligned} & \mathbb{E}\left[\rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) \mid H_t = h_t\right] \\ & = \mathbb{E}\left[\mathbb{E}\left[Y^{(A_{t:t+k-1}, d_{t+k})} - Y^{(A_{t:t+k}, d_{t+k+1})} \mid X_{t+1:t+k}, A_{t:t+k}, H_t = h_t\right] \mid H_t = h_t\right] \\ & = \mathbb{E}\left[Y^{(A_{t:t+k-1}, d_{t+k})} - Y^{(A_{t:t+k}, d_{t+k+1})} \mid H_t = h_t\right] \end{aligned} \quad (21)$$

Summing up (21) for  $k = 0, 1, \dots, \tau$ , we observe that all the middle terms cancel off:

$$\begin{aligned} & \mathbb{E}\left[\sum_{k=0}^{\tau} \rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) \mid H_t = h_t\right] \\ & = \mathbb{E}\left[\sum_{k=0}^{\tau} Y^{(A_{t:t+k-1}, d_{t+k})} - Y^{(A_{t:t+k}, d_{t+k+1})} \mid H_t = h_t\right] \\ & = \mathbb{E}\left[Y^{(d)} - Y \mid H_t = h_t\right] \end{aligned} \quad (22)$$

This is equivalent with the Eq. (3).  $\square$

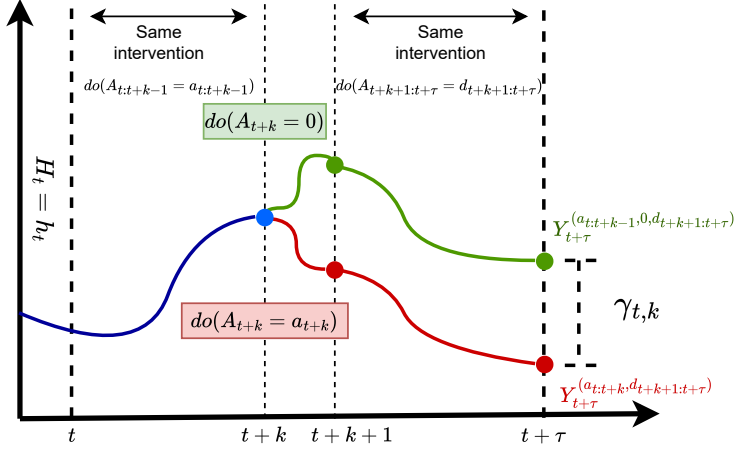


Figure 6: Visualization of the blip function  $\gamma_{t,k}$ . The blip function quantifies the localized treatment effect at each time step.

Lemma 2 shows that, in order to estimate the potential outcomes, it suffices to estimate the blip functions. Now the identification problem is reduced into correctly estimating these blip functions. Suppose  $\gamma_t^\dagger(\cdot, \cdot; h_t) = (\gamma_{t,0}^\dagger, \gamma_{t,1}^\dagger, \dots, \gamma_{t,\tau}^\dagger)$  is an **arbitrary** candidate function that takes  $x_{t+1:t+k}$ ,  $a_{t:t+k}$  as input and is subject to the following conditions:

$$\forall k \in [0, \tau], \quad \gamma_{t,k}^\dagger(x_{t+1:t+k}, (a_{t:t+k-1}, 0); h_t) = 0. \quad (23)$$

Let:

$$H_{t,k}(\gamma_t^\dagger) = Y + \sum_{j=k}^T \left( \gamma_{t,k}^\dagger(X_{t+1:t+j}, (A_{t:t+j-1}, d_{t+j})) - \gamma_{t,k}^\dagger(X_{t+1:t+j}, A_{t:t+j}) \right) \quad (24)$$

Then, with Lemma 2 and the sequential conditional ignorability, theorem 3.2 from Robins (2004) states that the moment restriction below is satisfied. Here we present it as a crucial lemma and show the proof for completeness.

**Lemma 1. Moment restriction for blip functions** When (23) is satisfied, then for arbitrary function  $S_k$ , the following two conditions are equivalent:

1.  $\gamma_{t,k}^\dagger(\cdot, \cdot; h_t)$  is the true  $\gamma_{t,k}(\cdot, \cdot; h_t)$ .
2. The following moment restriction is satisfied:

$$\mathbb{E} \left[ H_{t,k}(\gamma_t^\dagger) (S_k(X_{t+1:t+k}, A_{t:t+k}) - \mathbb{E} [S_k(X_{t+1:t+k}, A_{t:t+k}) \mid X_{t+1:t+k}, A_{t:t+k-1}]) \mid H_t = h_t \right] = 0 \quad (25)$$

*Proof.* We only prove the direction of  $1 \rightarrow 2$ , which already conveys the core thought of the moment restriction.

We denote:

$$R_k(X_{t+1:t+k}, A_{t:t+k}) = S_k(X_{t+1:t+k}, A_{t:t+k}) - \mathbb{E} [S_k(X_{t+1:t+k}, A_{t:t+k}) \mid X_{t+1:t+k}, A_{t:t+k-1}, H_t = h_t]$$

Observing the LHS of (25):

$$\begin{aligned}
\text{L.H.S.} &= \mathbb{E} \left[ H_{t,k}(\gamma_t^\dagger) R_k(X_{t+1:t+k}, A_{t:t+k}) \mid H_t = h_t \right] \\
&\stackrel{\text{Cond.}}{=} \mathbb{E} \left[ \mathbb{E} [H_{t,k}(\gamma_t) R_k(X_{t+1:t+k}, A_{t:t+k}) \mid X_{t+1:t+k}, A_{t:t+k}] \mid H_t = h_t \right] \\
&= \mathbb{E} \left[ \mathbb{E} [H_{t,k}(\gamma_t) \mid X_{t+1:t+k}, A_{t:t+k}] \cdot R_k(X_{t+1:t+k}, A_{t:t+k}) \mid H_t = h_t \right] \\
&\stackrel{\text{Lem. 2}}{=} \mathbb{E} \left[ \mathbb{E} \left[ Y^{(A_{t:t+k-1}, \underline{d}_{t+k})} \mid X_{t+1:t+k}, A_{t:t+k} \right] \cdot R_k(X_{t+1:t+k}, A_{t:t+k}) \mid H_t = h_t \right] \\
&\stackrel{\text{Ignor.}}{=} \mathbb{E} \left[ \mathbb{E} \left[ Y^{(A_{t:t+k-1}, \underline{d}_{t+k})} \mid X_{t+1:t+k}, A_{t:t+k-1} \right] \cdot R_k(X_{t+1:t+k}, A_{t:t+k}) \mid H_t = h_t \right] \\
&\stackrel{\text{Tower}}{=} \mathbb{E} \left[ \mathbb{E} \left[ Y^{(A_{t:t+k-1}, \underline{d}_{t+k})} \mid X_{t+1:t+k}, A_{t:t+k-1} \right] \cdot \mathbb{E} [R_k(X_{t+1:t+k}, A_{t:t+k}) \mid X_{t+1:t+k}, A_{t:t+k-1}] \mid H_t = h_t \right] \\
&\stackrel{\text{Def.}}{=} \mathbb{E} \left[ \mathbb{E} \left[ Y^{(A_{t:t+k-1}, \underline{d}_{t+k})} \mid X_{t+1:t+k}, A_{t:t+k-1} \right] \cdot 0 \mid H_t = h_t \right] = 0 \tag{26}
\end{aligned}$$

□

To further reduce the complexity of the moment restriction to achieve locally efficient doubly robust estimator (see Section 3.3 from Robins (2004)), we could subtract  $H_{t,k}(\gamma^*)$  (25) by its conditional expectation  $\mathbb{E}[H_{t,k}(\gamma_t^\dagger) \mid X_{t+1:t+k}, A_{t:t+k-1}]$ :

$$\mathbb{E} \left[ (H_{t,k}(\gamma_t^\dagger) - \mathbb{E}[H_{t,k}(\gamma_t^\dagger) \mid X_{t+1:t+k}, A_{t:t+k-1}]) R_k(X_{t+1:t+k}, A_{t:t+k}) \mid H_t = h_t \right] = 0 \tag{27}$$

**Parametrization of the blip function:** To achieve parametric rates of the estimation, we use the common semi-parametrization of the blip function<sup>7</sup> (Lewis & Syrgkanis, 2021; Robins, 2004):

$$\gamma_{t,k}(x_{t+1:t+k}, a_{t:t+k}; h_t) = \psi_{t,k}(h_t)' a_{t+k} \tag{28}$$

Here  $\psi_{t,k}$  is a function that maps a given  $H_t = h_t$  to  $\psi_{t,k}(h_t)$ , which is the parameter of the blip function  $\gamma_{t,k}(\cdot, \cdot; h_t)$ . Intuitively, these  $\psi_{t,k}(h_t)$  works as the coefficients in a linear parametrization. Hence, we call them in our method *blip coefficients* and the mapping  $\psi_{t,k}(\cdot)$  *blip coefficient predictor* or *blip coefficient estimator*.

**Deriving the moment restriction of SNMM:** We abuse the notation of  $H_{t,k}(\gamma^*)$  with using the parameter  $\psi_t$  to represent the  $H_{t,k}(\gamma_t^{\psi_t})$  as  $H_{t,k}(\psi_t)$ . For convenience, we introduce several notation for residuals:

$$\tilde{A}_{t,j,k} = A_{t+j} - \mathbb{E}[A_{t+j} \mid X_{t+1:t+k}, A_{t:t+k}, H_t = h_t], \quad 0 \leq k \leq j \leq \tau \tag{29}$$

$$\tilde{Y}_{t,k} = Y - \mathbb{E}[Y \mid X_{t+1:t+k}, A_{t:t+k}, H_t = h_t], \quad 0 \leq k \leq \tau \tag{30}$$

Then we could derive the following (with definition of  $H_{t,k}(\cdot)$  in (24)), which is an adaptation from Lewis & Syrgkanis (2021):

$$\mathbb{E} \left[ H_{t,k}(\psi_t) - \mathbb{E}[H_{t,k}(\psi_t) \mid X_{t+1:t+k}, A_{t:t+k-1}] \mid H_t = h_t \right] = \tilde{Y}_{t,k} - \sum_{j=k}^{\tau} \psi'_{t,j} \tilde{A}_{t,j,k} \tag{31}$$

Moreover, we specify  $S_k(x_{t+1:t+k}, a_{t:t+k-1}) = a_{t+k}$  from Lemma 1, then:

$$\begin{aligned}
&R_k(X_{t+1:t+k}, A_{t:t+k-1}) \\
&= S_k(X_{t+1:t+k}, A_{t:t+k}) - \mathbb{E} [S_k(X_{t+1:t+k}, A_{t:t+k}) \mid X_{t+1:t+k}, A_{t:t+k-1}, H_t = h_t] = \tilde{A}_{t,j,k} \tag{32}
\end{aligned}$$

Therefore, the moment restriction in (27) could be simplified as:

$$\forall k \in \{0, \dots, \tau\} : \quad \mathbb{E} \left[ \left( \tilde{Y}_{t,k} - \sum_{j=k+1}^{\tau} \psi_{t,j}(h_t)' \tilde{A}_{t,j,k} - \psi_{t,k}(h_t)' \tilde{A}_{t,k,k} \right) \tilde{A}_{t,k,k} \mid H_t = h_t \right] = 0 \tag{33}$$

which is the primary moment restriction for our setting, adapted from Lewis & Syrgkanis (2021). To estimate  $\psi_t(h_t)$  we need to solve the set of equations defined in (33) in an iterative way (See Algorithm 3 from Lewis & Syrgkanis (2021)):

<sup>7</sup>In our paper, we choose the simplest parametrization for clarity. In fact, the blip functions could have more sophisticated parametrizations, such as  $\gamma_{t,k}(x_{t+1:t+k}, a_{t:t+k}; h_t) = \psi_{t,k}(h_t)' \phi_{t,k}(x_{t+1:t+k}, a_{t:t+k}; h_t)$ , where  $\phi$  is a nonlinear, high-dimensional feature map (Lewis & Syrgkanis, 2021; Robins, 2004).

1. Step 0: for  $k = \tau$ , solve  $\widehat{\psi}_{t,\tau}(h_t) \in \mathbb{R}^r$
2. Step  $\tau - k$ , ( $k \geq 1$ ), With the given  $\widehat{\psi}_{t,j}(h_t), j > k$ , solve  $\widehat{\psi}_{t,k}(h_t)$ .

Each step we need to solve a linear equation, which has a unique solution under standard assumptions (see Theorem 8 from Lewis & Syrgkanis (2021)). And since the solution is unique, the solution equals the true blip coefficients. However, we argue that this classical method has limitations due to its sequential nature which prevents us from batch training. Noticing that solving (33) is equivalent to solving this minimization problem, we successfully establish the estimation of  $\psi_t$  under risk minimization paradigm, suitable for training with neural networks:

$$\psi_{t,k}^* = \underset{\widehat{\psi}_{t,k}(\cdot) \in \Phi_{t,k}}{\arg \min} \mathbb{E} \left[ \left( \widetilde{Y}_{t,k} - \sum_{j=k+1}^{\tau} \psi_{t,j}^*(h_t)' \widetilde{A}_{t,j,k} - \widehat{\psi}_{t,k}(h_t)' \widetilde{A}_{t,k,k} \right)^2 \right] \quad (34)$$

**Estimating CATE over time:** Once the parameters  $\psi_t(h_t) = (\psi_{t,0}(h_t), \dots, \psi_{t,\tau}(h_t))$  are estimated, we could proceed to estimate the potential outcome from Eq. (3):

$$\begin{aligned} & \mathbb{E} \left[ Y^{(d)} \mid H_t = h_t \right] \\ &= \mathbb{E} \left[ Y + \sum_{k=0}^{\tau} \rho_{t,k}(X_{t+1:t+k}, A_{t:t+k}; h_t) \mid H_t = h_t \right] \\ &= \mathbb{E} \left[ Y + \sum_{k=0}^{\tau} \widehat{\psi}_{t,k}(h_t)' (d_{t+k} - A_{t+k}) \mid H_t = h_t \right]. \end{aligned} \quad (35)$$

Then the CATE over time under treatment sequences  $a^*$  and  $b^*$  is computed by taking difference between the potential outcomes respectively Lewis & Syrgkanis (2021); Robins (1994; 2004):

$$\begin{aligned} & \mathbb{E} \left[ Y^{(a^*)} - Y^{(b^*)} \mid H_t = h_t \right] \\ &= \mathbb{E} \left[ \sum_{k=0}^{\tau} \psi_{t,k}(h_t)' [(a_{t+k}^* - A_{t+k}) - (b_{t+k}^* - A_{t+k})] \mid H_t = h_t \right] \\ &= \mathbb{E} \left[ \sum_{k=0}^{\tau} \psi_{t,k}(h_t)' (a_{t+k}^* - b_{t+k}^*) \mid H_t = h_t \right] \\ &= \sum_{k=0}^{\tau} \psi_{t,k}(h_t)' (a_{t+k}^* - b_{t+k}^*) \end{aligned} \quad (36)$$

The expectation term vanishes since all the terms left are constant in the conditional expectation, enabling us to directly estimate CATE over time.

## B.2 NEYMAN ORTHOGONALITY OF THE $L^2$ -MOMENT LOSS

**Fréchet Derivative:** The Fréchet derivative of any functional  $\mathcal{L}(f)$  is defined as:

$$\forall v, \quad D_f \mathcal{L}(f)[v] = \left. \frac{\partial}{\partial t} \mathcal{L}(f + tv) \right|_{t=0}. \quad (37)$$

where  $h = (p, q)$  is the nuisance parameter,  $\theta$  is the trainable weight of the blip prediction model from Stage 2. Our objective is to show that:

$$D_h(\partial_{\theta} \mathcal{L}_{D,k}(\theta; h)) = 0 \quad (38)$$

where  $D_h$  is the Fréchet derivative w.r.t. the nuisance function  $h$ . This property is called the **universal orthogonality** (Foster & Syrgkanis, 2019), which mandates that the gradient of the loss  $\partial_{\theta} \mathcal{L}_{D,k}(\theta; h)$  has a vanishing derivative over the nuisance parameters  $h$ . A functional loss that satisfies this property is more robust to the perturbation of the nuisance parameter, since the first-order error of  $h$  would only have a second-order effect on the gradient. Secondly, with the universal Neyman orthogonality, it is possible to achieve lower prediction error (Foster & Syrgkanis, 2019), as stated in the theorem in Section B.3.

Recall that our double-optimization-adapted  $L^2$ -moment loss at step  $k$  is:

$$\mathcal{L}_{\text{blip}}^k := \mathcal{L}_{D,k}(\theta; h) = \mathbb{E} \left[ \left( Y - p_{t,k}(H_{t+k}) - \sum_{j=k+1}^{\tau} \psi_{t,j}^{(\text{ps})}(H_t)' (A_{t+j} - q_{t,j,k}(H_{t+k})) - \psi_{t,k}(H_t; \theta)' (A_{t+k} - q_{t,k,k}(H_{t+k})) \right)^2 \right], \quad (39)$$

where  $h = (\mathbf{p}, \mathbf{q})$  is the nuisance parameter, where  $\mathbf{p} = \{p_{t,k} \mid 0 \leq k \leq \tau\}$  and  $\mathbf{q} = \{q_{t,j,k} \mid 0 \leq k \leq j \leq \tau\}$  and  $\theta$  is the trainable weight of the blip prediction model from Stage 2.  $\psi_{t,j}^{(\text{ps})}(H_t)$  is the pseudo target resulting from gradient detachment in the adapted loss.

First, we denote:

$$\phi_k(\theta, h) := Y - p_{t,k}(H_{t+k}) - \sum_{j=k+1}^{\tau} \psi_{t,j}^{(\text{ps})}(H_t)' (A_{t+j} - q_{t,j,k}(H_{t+k})) - \psi_{t,k}(H_t; \theta)' (A_{t+k} - q_{t,k,k}(H_{t+k})), \quad (40)$$

then, we derive the gradient of  $\mathcal{L}_{D,k}$  over  $\theta$ . Note that the nuisance functions  $p_{t,k}, q_{t,j,k}$  does not depend on the model weight  $\theta$ , the gradient is then:

$$\partial_{\theta} \mathcal{L}_{D,k}(\theta, h) = \mathbb{E} \left[ 2 \cdot \phi_k(\theta, h) \cdot -\partial_{\theta} \psi_{t,k}(H_t; \theta)' (A_{t+k} - q_{t,k,k}(H_{t+k})) \right]. \quad (41)$$

Let

$$\delta_k(\theta, h) = -\partial_{\theta} \psi_{t,k}(H_t; \theta)' (A_{t+k} - q_{t,k,k}(H_{t+k})), \quad (42)$$

Then we have

$$\partial_{\theta} \mathcal{L}_{D,k}(\theta, h) = 2 \cdot \mathbb{E} \left[ \phi_k(\theta, h) \cdot \delta_k(\theta, h) \right]. \quad (43)$$

Then we proceed to show that the Fréchet derivative of  $\partial_{\theta} \mathcal{L}_{D,k}(\theta, h)$  over  $p_{t,k}, q_{t,j,k}$  (where  $j > k$ , and  $q_{t,k,k}$  are all zero.

**(1) Orthogonality over  $p_{t,k}$ :** Note that  $D_{p_{t,k}} \delta_k(\theta, h) = 0$  and  $D_{p_{t,k}} \phi_k(\theta, h)[v] = v(H_{t+k})$ . then

$$\begin{aligned} D_{p_{t,k}} \partial_{\theta} \mathcal{L}_{D,k}(\theta, h)[v] &= 2 \cdot \mathbb{E} \left[ v(H_{t+k}) \cdot \delta_k(\theta, h) \right] \\ &= 2 \cdot \mathbb{E} \left[ \mathbb{E}[v(H_{t+k}) \cdot \delta_k(\theta, h) \mid H_{t+k}] \right] \\ &= 2 \cdot \mathbb{E} \left[ v(H_{t+k}) \cdot \mathbb{E}[\delta_k(\theta, h) \mid H_{t+k}] \right] \\ &= 2 \cdot \mathbb{E} \left[ v(H_{t+k}) \cdot 0 \right] = 0. \end{aligned} \quad (44)$$

**(2) Orthogonality over  $q_{t,j,k}$ :** Note that  $D_{q_{t,j,k}} \phi_k(\theta, h)[u] = \psi_{t,j}^{(\text{ps})}(H_t)' u(H_{t+k})$  and  $D_{q_{t,j,k}} \delta_k(\theta, h) = 0$ , then

$$\begin{aligned} D_{q_{t,j,k}} \partial_{\theta} \mathcal{L}_{D,k}(\theta, h)[u] &= 2 \cdot \mathbb{E} \left[ \psi_{t,j}^{(\text{ps})}(H_t)' u(H_{t+k}) \cdot \delta_k(\theta, h) \right] \\ &= 2 \cdot \mathbb{E} \left[ \mathbb{E}[\psi_{t,j}^{(\text{ps})}(H_t)' u(H_{t+k}) \cdot \delta_k(\theta, h) \mid H_{t+k}] \right] \\ &= 0. \text{ (similar to the proof of } p_{t,k} \text{)} \end{aligned} \quad (45)$$

**(3)Orthogonality over  $q_{t,j,k}$ :** Note that  $D_{q_{t,k,k}}\phi_k(\theta, h)[u] = \psi_{t,k}(H_t, \theta)'u(H_{t+k})$  and  $D_{q_{t,k,k}}\delta_k(\theta, h) = \partial_\theta\psi_{t,k}(H_t; \theta)'u(H_{t+k})$ , then

$$D_{q_{t,j,k}}\partial_\theta\mathcal{L}_{D,k}(\theta, h)[u] = 2 \cdot \mathbb{E} \left[ \psi_{t,k}(H_t, \theta)'u(H_{t+k}) \cdot \delta_k(\theta, h) + \right. \quad (46)$$

$$\left. \phi_k(\theta, h) \cdot \partial_\theta\psi_{t,k}(H_t; \theta)'u(H_{t+k}) \right] \\ = 2 \cdot \mathbb{E} \left[ \mathbb{E} \left[ \psi_{t,k}(H_t, \theta)'u(H_{t+k}) \cdot \delta_k(\theta, h) + \right. \quad (47)$$

$$\left. \phi_k(\theta, h) \cdot \partial_\theta\psi_{t,k}(H_t; \theta)'u(H_{t+k}) \mid H_{t+k} \right] \\ = 2 \cdot \mathbb{E} \left[ \psi_{t,k}(H_t, \theta)'u(H_{t+k}) \cdot \mathbb{E}[\delta_k(\theta, h) \mid H_{t+k}] + \quad (48)$$

$$\mathbb{E}[\phi_k(\theta, h) \mid H_{t+k}] \cdot \partial_\theta\psi_{t,k}(H_t; \theta)'u(H_{t+k}) \right] \\ = 2 \cdot \mathbb{E} \left[ \psi_{t,k}(H_t, \theta)'u(H_{t+k}) \cdot 0 + 0 \cdot \partial_\theta\psi_{t,k}(H_t; \theta)'u(H_{t+k}) \right] \\ = 0. \quad (49)$$

Putting (1), (2), (3) together, we prove that the overall Fréchet derivative of  $\partial_\theta\mathcal{L}_{D,k}(\theta, h)$  over nuisance parameter  $h$  is zero. And hence, the adapted  $L^2$ -moment loss satisfies universal Neyman orthogonality.

### B.3 THE MSE RATE THEOREM

A mean squared error (MSE) guarantee under finite samples for the estimator  $\hat{\psi}_t$  in a primary heterogeneous setting, where the conditioning variable is the static feature  $X_0$  (see theorem 10 from Lewis & Syrgkanis (2021)). We point out that the conclusion of the theorem could be extended to our setting, where conditioning variable is the history  $H_t$ . Before we show the theorem, we clarify the prerequisite notations and regular assumptions:

#### Mathematical notations:

1. **Norm of a function in function space** Define the norm for a vector-valued function  $\psi : \mathcal{X} \rightarrow \mathbb{R}^d$ , then the  $u, v$ -norm is defined as:

$$\|\psi\|_{u,v} = \mathbb{E}_{X \sim P_X} \left[ \left( \sum_{i=1}^d \psi_i(X)^u \right)^{\frac{v}{u}} \right]^{\frac{1}{v}}. \quad (50)$$

We point out that when  $u = v$ , the  $u, v$ -norm degenerate into the normal norm:  $\|\cdot\|_{u,u} = \|\cdot\|_u$ . Further, we could define the product norm between two functions  $f, g$  as:

$$\|f \circ g\|_{u,v} = \mathbb{E} \left[ \|f(\hat{X})\|_u^v \cdot \|g(\hat{Y})\|_v^v \right]^{1/v} = \mathbb{E} \left[ \left( \sum_{i=1}^r f_i(X)^u \right)^{\frac{v}{u}} \cdot \sum_{i=1}^r g_i(X)^u \right]^{\frac{1}{v}} \quad (51)$$

2. **Localized Rademacher complexity:** The localized Rademacher complexity is an extended concept of the standard Radmacher complexity. It measures the the fitting capacity of a function space  $\mathcal{F}$  where functions have bounded second moment. Formally it is defined as:

$$R_{P_X, n}(\mathcal{F}; \delta) = \mathbb{E}_{\epsilon_{1:n}, X_{1:n} \sim P_X} \left[ \sup_{f \in \mathcal{F}: \|f\|_2 \leq \delta} \frac{1}{n} \sum_{i=1}^n \epsilon_i f(X_i) \right] \quad (52)$$

**Assumptions:** We adapt the assumptions from the original heterogeneous setting (Lewis & Syrgkanis, 2021) to conditioning on the history variable  $H_t$ :

1. All the variables and functions in our problem setting are bounded.
2. A regularity assumption regarding the treatment variables: The following quantities are also bounded:

$$c_{k,j} := \sup_{t \leq T-\tau} \sup_{h_t \in \mathcal{H}_t} \mathbb{E} \left[ \|C_{t,k,j}\|_{u,u}^v \mid H_t = h_t \right]^{2/v} < +\infty \quad (53)$$

where  $C_{t,k,j} := \text{Cov}(A_{t+k}, A_{t+j} \mid H_{t+k})$ , and:

$$\forall h_t \in \mathcal{H}_t, \quad \mathbb{E}[C_{t,k,k} \mid H_t = h_t] \succeq \lambda I \quad (54)$$

3. The  $u, \infty$ -norms of  $\psi_{t,k}$  are bounded by:

$$M := \max_{t \leq T-\tau} \max_{0 \leq k \leq \tau, \psi_{t,k} \in \Psi_{t,k}} \|\psi_{t,k}\|_{u,\infty} < +\infty \quad (55)$$

**Theorem 1 (Adapted from Heterogeneous R-learner(Lewis & Syrgkanis, 2021)):** Suppose the assumptions above are satisfied. Let  $\delta_n$  be an upper bound on the critical radius of the star hull of all the function spaces

$$\{\Psi_{t,k,i} - \psi_{t,k,i}\}_{t \in [SL-\tau], k \in [\tau], i \in [r]}$$

and that  $\delta_n = \Omega\left(\sqrt{\frac{r \log \log(n)}{n}}\right)$ . Suppose:

$$\forall 1 \leq t \leq T - \tau, \quad \max_{0 \leq k \leq j \leq \tau} \mathbb{E}_{\hat{q}_{t,k,k}, \hat{q}_{t,j,k}} \left[ \left\| (\hat{q}_{t,k,k} - q_{t,k,k}) \circ (\hat{q}_{t,j,k} - q_{t,j,k}) \right\|_{2,2}^2 \right] = O(r^2 \delta_n^2) \quad (56)$$

$$\forall 1 \leq t \leq T - \tau, \quad \max_{0 \leq k \leq \tau} \mathbb{E}_{\hat{q}_{t,k,k}, \hat{q}_{t,k}} \left[ \left\| (\hat{q}_{t,k,k} - q_{t,k,k}) \circ (\hat{p}_{t,k} - p_{t,k}) \right\|_{2,2}^2 \right] = O(r^2 \delta_n^2), \quad (57)$$

then the blip coefficient predictor  $\hat{\psi}_{t,k}$  trained under the  $L^2$ -moment loss satisfies:

$$\max_{t \leq T-\tau} \max_{k \in \{0, \dots, \tau\}} \mathbb{E} \left[ \left\| \hat{\psi}_{t,k} - \psi_{t,k} \right\|_{2,2}^2 \right] = O(r^2 \delta_n^2), \quad \delta_n^2 \propto \frac{\log \log(n)}{n}. \quad (58)$$

The theorem uses the notion of critical radius which is often applied to describe the statistical complexity of functional spaces(Wainwright, 2019). Intuitively, the theorem states that when:

- our sequential neural architecture is correctly specified and trained with empirical risk minimization,
- the product norm between nuisance networks has a error small enough,
- the boundness assumption stated above are satisfied,

then our trained neural network for predicting the sequential blip effect will only have an  $\Omega\left(\frac{\log \log(n)}{n}\right)$  root mean squared error rate.

## C THEORETICAL JUSTIFICATION FOR THE DOUBLE OPTIMIZATION TRICK

Our double optimization trick implements the blip estimation stage with a simultaneous gradient update. This section provides the theoretical rationale for this design. We show that our approach is a practical, one-step approximation of an iterative process that is guaranteed to converge to the correct solution.

1. We define an idealized, iterative solution operator and prove that it converges to the true blip coefficients.
2. We derive an error propagation bound for a single application of this operator.
3. We discuss the practical implications of this theory for our model’s training and error propagation.

### C.1 ITERATIVE OPERATOR FOR LEARNING BLIP ESTIMATORS

The  $L^2$ -moment loss from Eq. (5) in the main chapter defines a system of interdependent optimizations. The solution for the blip coefficients at step  $k$ , i.e.  $\psi_{t,k}$ , depends on all the future steps  $j > k$ . This sequential solution scheme is, however, extremely slow. Hence, we replace it with a fixed point iteration scheme, which leads to the same exact solution.

Let  $\Psi$  be the space of all possible blip coefficient functions  $\psi = (\psi_0, \dots, \psi_\tau)$ . We define an operator  $T : \Psi \rightarrow \Psi$  that maps an input set of functions  $\psi^{(\text{old})}$  to an output set  $\psi^{(\text{new})}$ . The  $k$ -th component of the output is defined as the exact minimizer of the loss at step  $k$ , using the input functions as targets:

$$\psi_{t,k}^{(\text{new})}(h_t) = \arg \min_{\hat{\psi}_{t,k}(\cdot) \in \Psi} \mathbb{E} \left[ \left( \tilde{Y}_{t,k} - \sum_{j=k+1}^{\tau} \psi_{t,j}^{(\text{old})}(h_t)' \tilde{A}_{t,j,k} - \hat{\psi}_{t,k}(h_t)' \tilde{A}_{t,k,k} \right)^2 \right] \quad (59)$$

**Our double optimization trick is a practical, gradient-based implementation of a single iteration of this operator.**

The input  $\psi^{(\text{old})}$ , in practice, is just the output of the detached network  $\hat{\psi}^2$  in Eq. (7). The loss function  $\mathcal{L}_{\text{blip}} = \sum_k \mathcal{L}_{\text{blip}}^k$  is constructed to optimize the summed targets from Eq.59. The goal of the update step of the double optimization is to make the trained blip prediction network one step closer to the ideal target output  $\psi_{t,k}^{(\text{new})}$ . Actually, this kind of approach is not a new heuristic. It resembles the policy update in actor-critic methods in reinforcement learning: A single gradient step of the policy is taken to approximate the arg max operation on the Q-value function.

**Convergence of the operator:** A fundamental property of this operator is that it looks backward. The output at step  $k$ ,  $\psi_{t,k}^{(\text{new})}$ , depends only on the input functions from future steps,  $j > k$ . At the final time step  $k = \tau$ , the term  $\sum_{j=\tau+1}^{\tau} (\dots)$  is an empty sum that equals zero. The definition of  $\psi_{t,\tau}^{(\text{new})}$  is therefore independent of the input  $\psi^{(\text{old})}$ . This means for any two different input blip estimators  $\psi$  and  $\phi$ , the operator produces the exact same output for the final step:  $T(\psi)_\tau = T(\phi)_\tau$ .

This provides a stable anchor at the end of the horizon. When the operator  $T$  is applied again, the output for the second-to-last step,  $\psi_{t,\tau-1}$ , will now also be uniquely determined, as its only dependency is on the now-fixed final step. This logic propagates backward through all time steps. After  $\tau + 1$  applications of the operator, any initial difference between two blip predictors will vanish completely. This means that, starting with an arbitrary  $\psi_t^{\text{initial}}$ , the output of the operator converges to the true blip coefficient after  $\tau + 1$  iterations.

### C.2 ERROR BOUND OF THE OUTPUT OF THE OPERATOR

In previous section, we have shown the convergence property of the transformation  $T$ . In practice, our double optimization trick is only an approximation of  $T$ . Therefore, in this section, we analyze how estimation errors propagate through a single application of the operator  $T$ .

#### C.2.1 FORMAL ASSUMPTIONS

We invoke two standard regularity assumptions:

1. **Strong Overlap:** The smallest eigenvalue of the conditional covariance matrix  $C_k(h_t) = \mathbb{E}[\tilde{A}_{t,k,k} \tilde{A}'_{t,k,k} | H_t = h_t]$  is larger than some  $\sigma_{\min}^2 > 0$ .
2. **Bounded Moments:** The conditional cross-moments of treatment residuals are bounded. There exists a constant  $M < \infty$  such that  $\mathbb{E} \left[ \|\tilde{A}_{t,k,k} \tilde{A}'_{t,j,k}\|_2^2 | H_t = h_t \right] \leq M^2$ .

### C.2.2 DERIVATION OF THE ERROR BOUND

Let  $\Delta^{(\text{in})} = \psi^{(\text{in})} - \phi^{(\text{in})}$  be the difference between two input estimators, and  $\Delta^{(\text{out})} = T(\psi^{(\text{in})}) - T(\phi^{(\text{in})})$  be the difference in their outputs. The closed-form solution to the least-squares problem in Eq. equation 59 gives the difference at step  $k$ :

$$\Delta_k^{(\text{out})}(h_t) = \mathbb{E} \left[ -C_k(h_t)^{-1} \sum_{j=k+1}^{\tau} \tilde{A}_{t,k,k} \tilde{A}'_{t,j,k} \Delta_j^{(\text{in})}(h_t) \middle| H_t = h_t \right]$$

By taking the squared norm, applying Jensen's inequality, Cauchy-Schwarz, and our assumptions, we can bound the expected error (see derivation below):

$$\mathbb{E}[\|\Delta_k^{(\text{out})}\|^2] \leq C \sum_{j=k+1}^{\tau} \mathbb{E}[\|\Delta_j^{(\text{in})}\|^2] \quad (60)$$

where the constant  $C = \frac{\tau M^2}{\sigma_{\min}^4}$  depends on the horizon and the data properties (overlap and moments).

*Proof.*

$$\begin{aligned} \|\Delta_k^{(\text{out})}(h_t)\|^2 &= \left\| \mathbb{E} \left[ -C_k(h_t)^{-1} \sum_{j=k+1}^{\tau} \tilde{A}_{t,k,k} \tilde{A}'_{t,j,k} \Delta_j^{(\text{in})}(h_t) \middle| H_t = h_t \right] \right\|^2 \\ &\leq \mathbb{E} \left[ \left\| C_k(h_t)^{-1} \sum_{j=k+1}^{\tau} \tilde{A}_{t,k,k} \tilde{A}'_{t,j,k} \Delta_j^{(\text{in})}(h_t) \right\|^2 \middle| H_t = h_t \right] \quad (\text{by Jensen's Inequality}) \\ &\leq \mathbb{E} \left[ \|C_k(h_t)^{-1}\|_2^2 \left\| \sum_{j=k+1}^{\tau} \tilde{A}_{t,k,k} \tilde{A}'_{t,j,k} \Delta_j^{(\text{in})}(h_t) \right\|^2 \middle| H_t = h_t \right] \\ &\leq \frac{1}{\sigma_{\min}^4} \mathbb{E} \left[ (\tau - k) \sum_{j=k+1}^{\tau} \left\| \tilde{A}_{t,k,k} \tilde{A}'_{t,j,k} \Delta_j^{(\text{in})}(h_t) \right\|^2 \middle| H_t = h_t \right] \quad (\text{Assumption 1 \& Cauchy-Schwarz}) \\ &\leq \frac{\tau - k}{\sigma_{\min}^4} \sum_{j=k+1}^{\tau} \mathbb{E} \left[ \|\tilde{A}_{t,k,k} \tilde{A}'_{t,j,k}\|_2^2 \cdot \|\Delta_j^{(\text{in})}(h_t)\|^2 \middle| H_t = h_t \right] \\ &\leq \frac{(\tau - k)M^2}{\sigma_{\min}^4} \sum_{j=k+1}^{\tau} \|\Delta_j^{(\text{in})}(h_t)\|^2 \quad (\text{Assumption 2}) \end{aligned}$$

Taking the expectation over  $H_t$  on both sides yields the final bound in Eq. equation 60.  $\square$

### C.3 PRACTICAL IMPLICATIONS

This theory provides two crucial insights for our practical implementation.

First, the error bound confirms the importance of the stable anchor at  $k = \tau$ . Because the error at any step  $k$  depends only on the errors from future steps, if we can learn an accurate estimator for the final step,  $\hat{\psi}_{t,\tau}$ , its correctness will propagate backward to stabilize the learning of all preceding estimators. The loss for the final step,  $\mathcal{L}_{\text{flip}}^{\tau}$ , is exactly the original  $L^2$ -moment loss  $\mathcal{L}_{\tau}$  in Eq. (5), guaranteeing that  $\hat{\psi}_{t,\tau}$  converges to its true target.

Second, the bound quantifies how errors can accumulate. In practice, our one-step update will not be perfect. Let  $\epsilon_k = \mathbb{E}[\|\hat{\psi}_{t,k} - \psi_{t,k}^*\|^2]$  be the total MSE at step  $k$ , and assume a minimal unavoidable error of  $\delta_k$  at each step (due to finite data or model misspecification). The error recurrence is:

$$\epsilon_k \approx \delta_k + C \sum_{j=k+1}^{\tau} \epsilon_j$$

1350  
1351  
1352  
1353  
1354  
1355  
1356  
1357  
1358  
1359  
1360  
1361  
1362  
1363  
1364  
1365  
1366  
1367  
1368  
1369  
1370  
1371  
1372  
1373  
1374  
1375  
1376  
1377  
1378  
1379  
1380  
1381  
1382  
1383  
1384  
1385  
1386  
1387  
1388  
1389  
1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403

---

From this, we can derive a recursive relationship between adjacent steps:  $\epsilon_{k-1} \approx (1 + C)\epsilon_k + (\delta_{k-1} - \delta_k)$ . Assuming the base error  $\delta_k$  is roughly constant ( $\delta_{k-1} \approx \delta_k$ ), this simplifies to  $\epsilon_{k-1} \approx (1 + C)\epsilon_k$ . Unrolling this recursion from the final step backwards gives:

$$\epsilon_k \approx (1 + C)^{\tau-k} \epsilon_\tau \approx (1 + C)^{\tau-k} \delta_\tau \tag{61}$$

This bound shows that errors can grow exponentially with the distance from the final step,  $\tau - k$ . This underscores the importance of having strong data overlap (to keep  $C$  small) and an accurate model (to keep the base error  $\delta_\tau$  small), especially for applications involving long horizons.

---

## 1404 D EXTENDED RELATED WORKS

1405

### 1406 D.1 NON-PARAMETRIC MODELS

1407

1408 There have been some attempts to use non-parametric models (Schulam & Saria, 2017; Soleimani et al., 2017; Xu et al.,  
1409 2016) to estimate CATE or CAPO over time, yet these approaches impose strong assumptions on the outcomes, and  
1410 their scalability is limited. For these reasons, we focus on neural methods, which offer better flexibility and scalability  
1411 for complex, high-dimensional medical data.

1412

### 1413 D.2 TREATMENT EFFECT ESTIMATION OVER CONTINUOUS TIME

1414

1415 A streamline in the literature, orthogonal to ours, is estimating HTEs or potential outcomes over *continuous* time.  
1416 Prominent examples like Hess et al. (2024); Seedat et al. (2022) leverage Neural Controlled Differential Equations  
1417 (NCDEs), a method well-suited for modeling dynamic systems, to address the problem. However, these continuous-time  
1418 models introduce significant complexity and computational overhead, limiting their practical applicability. Hence, we  
1419 only focus on discrete-time treatment effect estimation.

1420

### 1421 D.3 BALANCED REPRESENTATION BASED APPROACHES

1422

1423 Learning the balanced representation to predict treatment effect or potential outcome is very popular. Similar to the  
1424 previously introduced Causal Transformer (Melnichuk et al., 2022), Wang et al. (2024) adopts an adversarial generative  
1425 approach to learn the balanced representations, which reduces the confounding bias. Wang et al. (2025a) introduces a  
1426 de-correlation strategy that removes cross-covariance between current treatments and representation of past trajectory,  
1427 based on a data selective state-space model architecture called Mamba. However, balanced representation-based methods  
1428 are considered *biased* because they are only heuristics that lack formal theoretical guarantee for adjusting confounding.  
1429 In fact, these methods were originally designed to reduce finite-sample variance, not to mitigate confounding bias.  
1430 Moreover, the process of enforcing such balance would even introduce more bias (Melnichuk et al., 2024).

1431

### 1432 D.4 OTHER MODEL-BASED APPROACHES

1433

1434 Wang et al. (2025b) proposed Variational Counterfactual Intervention Planning (VCIP), a framework that uses variational  
1435 inference to directly model the likelihood of the potential outcome over sequential interventions. The method relies  
1436 on the assumption that the ELBO derived from observational data truly approximates the interventional likelihood.  
1437 However, this is hard to verify. Furthermore, the approach uses simple distributions such as Gaussian (state transition)  
1438 and Beta (outcome distribution) distribution, which may fail to capture the complex relationship, leading to biased  
1439 estimation.

1440 Meng et al. (2024) introduces the COSTAR framework that first pre-trains a transformer-based encoder using a self-  
1441 supervised loss and then fine-tunes the encoder to predict the factual outcome. The work lacks a causal adjustment  
1442 mechanism, hoping that the pre-trained transformer can incidentally transfer to the setting of estimating counterfactual  
1443 outcomes. Hence, the approach is a heuristic and does not properly adjust for time-varying confounding.

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

---

## E ADDITIONAL RESULTS

### E.1 ABLATION STUDIES

We report the ablation studies on (1) the neural backbone of the sequential encoder and (2) the  $L^2$ -moment loss with double optimization trick. To this, we further introduce the following baselines:

1. **DeepBlip-LSTM:** We replace the transformer architecture in the sequential encoders  $\mathcal{E}_{\theta_N}^N, \mathcal{E}_{\theta_B}^B$  of nuisance network and blip prediction network by the LSTM network. From the experimental results below, we find that DeepBlip-LSTM is still highly effective: it outperforms other baselines, which proves the effectiveness of our propose DeepBlip framework regardless of the instantiations. From the examples of DeepBlip instantiated by transformer and LSTM, we demonstrate that our DeepBlip framework can seamlessly integrate popular sequential networks.
2. **DeepBlip-WDO:** We remove the double optimization trick (WDO stands for **Without Double Optimization**) and only make a single forward pass on  $\hat{\psi}_t$  to construct the  $L^2$ -moment loss. In this scenario, we discard the mandate on the order of solving the blip coefficients estimators, which is required for a correct estimation in SNMMs. This way, we are able to identify the contribution of applying the double optimization trick. Our results show that DeepBlip-WDO suffers from high estimation error and thus, demonstrate the necessity of the double optimization trick.

We report the RMSE on both synthetic dataset and semi-synthetic dataset as in the experiment section. We adopt the similar test setting in experiment section, where we test the performance against growing time-varying confounding  $\gamma_{\text{conf}}$  in the synthetic data experiment and test the performance against increasing prediction horizons  $\tau$  in the semi-synthetic dataset.

### E.2 GRANULAR ANALYSIS OF THE BLIP COEFFICIENTS PREDICTION

So far, we only demonstrate the overall RMSE on the CATE over time estimation. Our DeepBlip works by predicting the blip coefficients  $\hat{\psi}_t(H_t) = (\hat{\psi}_{t,0}(H_t), \dots, \hat{\psi}_{t,\tau}(H_t)) \in \mathbb{R}^{(\tau+1)d_a}$ . In this subsection, we report a granular analysis on the prediction of the blip coefficients, bringing more transparency to our DeepBlip.

In fact, we could directly derive the ground truth personalized blip coefficients  $\psi_t = (\psi_{t,0}(H_t), \dots, \psi_{t,\tau}(H_t))$  for both the synthetic dataset in F.1 and semi-synthetic dataset in F.2. This enables us to accurately evaluate the performances of all the blip coefficient predictors. Specifically, we visualize the distributions of all the components of the difference between the true blip coefficients and the prediction  $\psi_t - \hat{\psi}_t \in \mathbb{R}^{(\tau+1)d_a}$ . For synthetic dataset, we set  $\gamma = 5$  and  $\tau = 2$ . For semi-synthetic dataset, we set  $\gamma = 1$  and  $\tau = 4$ . In both settings, we use DeepBlip models instantiated by transformer that achieve the lowest RMSE and DeepBlip models trained without the double optimization trick to further demonstrate its necessity. The results from Figure 9 and Figure 10 show that our DeepBlip model is indeed capable of predicting each blip coefficient with high accuracy while with low variance. This supports our claim that **DeepBlip achieves robust estimation of CATE by decomposing the total effect into incremental blip effects with controlled error**. Further, by comparing the blip prediction of DeepBlip with and without double optimization, we see that the double optimization trick is essential to generating unbiased blip coefficient estimates.

1512  
 1513  
 1514  
 1515  
 1516  
 1517  
 1518  
 1519  
 1520  
 1521  
 1522  
 1523  
 1524  
 1525  
 1526  
 1527  
 1528  
 1529  
 1530  
 1531  
 1532  
 1533  
 1534  
 1535  
 1536  
 1537  
 1538  
 1539  
 1540  
 1541  
 1542  
 1543  
 1544  
 1545  
 1546  
 1547  
 1548  
 1549  
 1550  
 1551  
 1552  
 1553  
 1554  
 1555  
 1556  
 1557  
 1558  
 1559  
 1560  
 1561  
 1562  
 1563  
 1564  
 1565

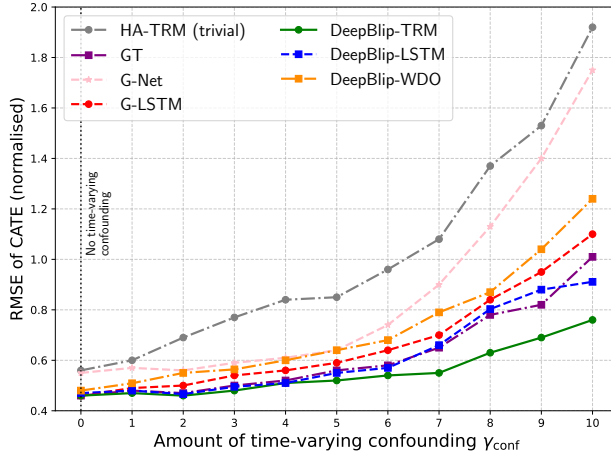


Figure 7: Ablation study on the tumor growth synthetic dataset. We explore two ablations: (1) We replace the transformer sequential encoder of DeepBlip and GT by LSTMs to create **DeepBlip-LSTM** and **G-LSTM**. (2) We remove the double optimization trick to create **DeepBlip-WDO**. As in previous experiment, we report the average RMSE against increasing confounding  $\gamma_{\text{conf}}$ . We note that DeepBlip-LSTM is still more competitive than other baselines, which indicates the *effectiveness of the model-agnosti framework of DeepBlip*. DeepBlip-TRM, as our originally proposed model, remain the strongest.

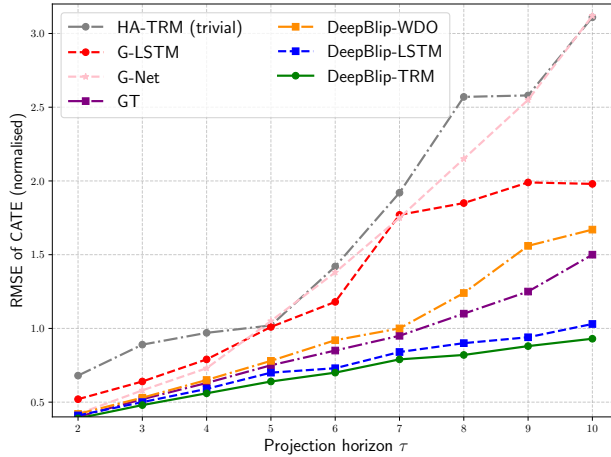
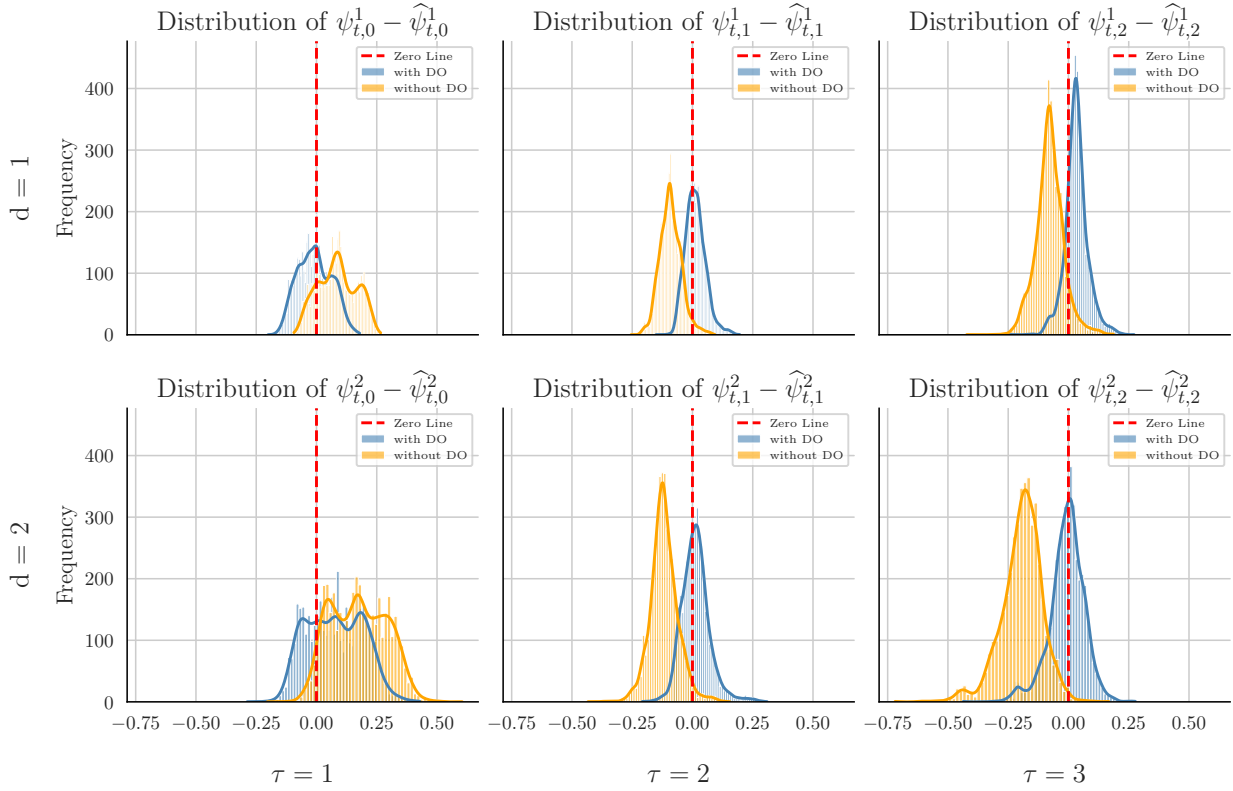


Figure 8: Ablation study on the MIMIC semi-synthetic dataset. Same ablations are investigated using the average RMSE against increasing projection horizon  $\tau$ . Even though DeepBlip-LSTM slightly underperforms the original model DeepBlip-TRM, it is still superior over other baselines. Hence, our proposed DeepBlip framework is *robust* over longer prediction horizon regardless of the backbone. We also observe that DeepBlip-WDO fail to predict CATE as accurately as the proposed DeepBlip-TRM. This shows that the adoption of a stable  $L^2$ -moment loss, together with the double optimization trick, are crucial for learning the blip coefficients unbiasedly.

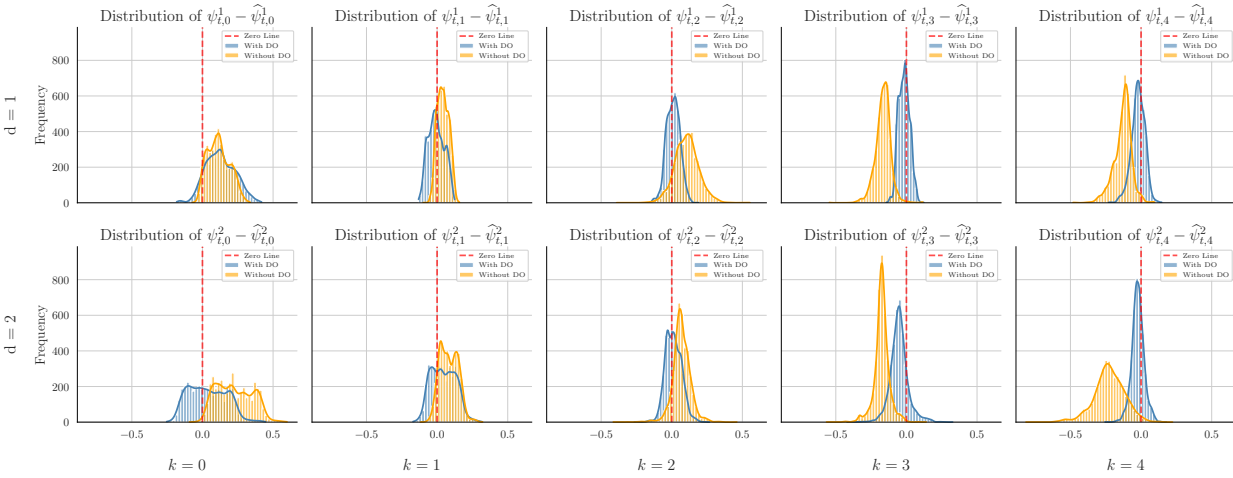
1566  
 1567  
 1568  
 1569  
 1570  
 1571  
 1572  
 1573  
 1574  
 1575  
 1576  
 1577  
 1578  
 1579  
 1580  
 1581  
 1582  
 1583  
 1584  
 1585  
 1586  
 1587  
 1588  
 1589  
 1590  
 1591  
 1592



1593  
 1594  
 1595  
 1596

Figure 9: Tumor synthetic dataset ( $\gamma_{\text{conf}} = 5$ ,  $\tau = 2$ ): We visualize the distribution of the difference between the ground truth blip coefficients and the predictions made by our blip prediction network. Observing the histogram marked in blue represented by our DeepBlip, we find that the blip prediction network unbiasedly predict the blip coefficients, which offers proper adjustment for the confounding.

1597  
 1598  
 1599  
 1600  
 1601  
 1602  
 1603  
 1604  
 1605  
 1606  
 1607  
 1608  
 1609  
 1610  
 1611  
 1612  
 1613



1615  
 1616  
 1617  
 1618  
 1619

Figure 10: MIMIC semi-synthetic dataset ( $\gamma_{\text{conf}} = 1$ ,  $\tau = 4$ ): We conduct the same visualization. When  $k$  grows, the prediction gets more accurate. We expect this to happen, since  $\hat{\psi}_{t,\tau}$  is the first to optimize by the double optimization trick and then  $\hat{\psi}_{t,\tau-1}$  and so on. In contrast, the blip prediction network trained without double optimization exhibits significant bias and generally higher variance. This implies that the double optimization trick is crucial for the training DeepBlip.

## F DATASET

In this section we will give a more detailed elaboration of the data generating process as well as the key simulation of the counterfactual outcome (namely the potential outcome). The treatment effect is then achieved by subtracting two counterfactual outcomes.

### F.1 TUMOR GROWTH DATASET

We specify the dynamic of the tumor growth model to be:

$$Y_{t+1} = Y_t + \underbrace{\left( \rho \log\left(\frac{K}{m(\bar{Y}_{t-l})}\right) + \epsilon_{t+1} \right)}_{\text{Tumor Growth}} m(\bar{Y}_{t-l}) - \underbrace{m(\bar{Y}_{t-l})\alpha_c c_{t+1}}_{\text{Chemotherapy}} + \underbrace{m(\bar{Y}_{t-l})(\alpha_r d_{t+1} + \beta_r d_{t+1}^2)}_{\text{Radiotherapy}} \quad (62)$$

$\alpha_c, \alpha_r, \beta_r$  are the coefficients on the treatment effect of chemotherapy and radiotherapy,  $\rho, K$  together control the growth model.  $m(\bar{Y}_{t-l})$  represents the mean of tumor volumes  $Y_1, \dots, Y_{t-l}$ , where  $l$  is a lag parameter. The averaging of the volumes avoids abrupt changes for the tumor volume. The chemotherapy drug dosage  $c_t$  and the radiation dosage  $d_t$  are applied with probability:

$$A_t^c, A_t^d \sim \text{Ber}\left(\sigma \cdot \left(\frac{\gamma_{conf}}{D_{max}} (\bar{D}_{15}(\bar{Y}_t) - \frac{D_{max}}{2})\right)\right) \quad (63)$$

Since the probability distribution is determined on the mean of previous 15 steps' tumor volumes  $\bar{D}_{15}(\bar{Y}_t)$ , time-varying confounding exists through future outcomes with its influence controlled by the strength parameter  $\gamma_{conf}$ .

To generate the observational dataset, we simulate the trajectories of  $N = 1000$  patients with maximal sequence length  $T = 30$  under confounding strength  $\gamma_{conf} \in \{0, 1, \dots, 10\}$ , resulting in 11 datasets  $\mathcal{D}_{N,T,\gamma_{conf}}$ . Then we split the simulated dataset into training, validation and test set by a split of (70%, 15%, 15%).

**Simulating counterfactuals:** For the test dataset, we compute the ground truth CATE by simulating the counterfactual outcomes  $Y_{t+\tau}^{a^*}$  and  $Y_{t+\tau}^{b^*}$  via a sliding window treatment along the sequence. To simulate the counterfactual outcome  $Y_{t+\tau}^{(a^*)}$  under an intervention  $\mathbf{a}^* \in \mathbb{R}^{(\tau+1) \times 2}$  during  $[t, t + \tau]$  for the tumor growth model, we iteratively compute the counterfactual trajectory starting from the observed history  $H_t$ . The DGP of outcome is consisted by the lagged dependency on tumor volume averages and the cumulative effects of treatments. Here we provide a strict formulation of the counterfactual outcome generation process, including the mathematical formulation and the steps for simulation.

Given the observed history  $H_t = (\bar{X}_t, \bar{A}_{t-1}, \bar{Y}_{t-1})$ , the counterfactual trajectory  $\{Y_s^{(a^*)}\}_{s=t}^{t+\tau}$  under intervention  $\mathbf{a}^* = [a_t^*, a_{t+1}^*, \dots, a_{t+\tau}^*]$  (where  $a_s^* = (c_s^*, d_s^*)$ ) is generated with the following steps:

1. **Initialization:** Get observed history:  $Y_s$  for  $s < t$ . Note  $Y_{t-1}^{(a^*)} = Y_{t-1}$ .
2. **Iterative simulation:** For each time  $s \in [t, t + \tau - 1]$ , we simulate  $Y_{s+1}^{(a^*)}$  based on previous simulated outcomes. Notice that since  $l > \tau$ ,  $s - l < t$ , then  $m(\bar{Y}_{s-l})$  is unaffected by the interventions and therefore could be treated as constant when conditioning on  $H_t$ . Therefore we have:

$$Y_{s+1}^{(a^*)} = Y_s^{(a^*)} + \underbrace{\rho \log\left(\frac{K}{m(\bar{Y}_{s-l})}\right)}_{\text{Growth Term}} m(\bar{Y}_{s-l}) + \epsilon_{s+1} - \underbrace{m(\bar{Y}_{s-l})\alpha_c c_{s+1}^*}_{\text{Chemotherapy Effect}} - \underbrace{m(\bar{Y}_{s-l})(\alpha_r d_{s+1}^* + \beta_r (d_{s+1}^*)^2)}_{\text{Radiotherapy Effect}}, \quad (64)$$

where  $\epsilon_{s+1}$  is the **original** noise term during the generation of the observational data.  $c_{s+1}^*, d_{s+1}^*$  are the intervened chemotherapy and radiotherapy dosages at  $s + 1$ , drawn from  $\mathbf{a}^*$

Under this iterative scheme, we could compute the final counterfactual outcome  $Y_{t+\tau}^{(a^*)}$ .

**Treatment effect:** We then proceed to compute the individual conditioned treatment effect for a patient under treatment  $\mathbf{a}^*$  and the null treatment  $\mathbf{0}$ . Observing the process in (64), we find that the growth term **cancel off** each other during the subtraction, leaving only the terms for chemotherapy effect and radiotherapy effect:

$$Y_{t+\tau}^{(a^*)} - Y_{t+\tau}^{(0)} = \sum_{s=t}^{t+\tau} \left\{ m(\bar{Y}_{s-l}) \alpha_c c_{s+1}^* + m(\bar{Y}_{s-l}) (\alpha_r d_{s+1}^* + \beta_r (d_{s+1}^*)^2) \right\} \quad (65)$$

The individual treatment effect is calculated for all the units on the test dataset  $D_{test}$  in a sliding treatment window rolling over  $t \in \{1, \dots, T - \tau\}$ , producing  $N_{test} \times (T - \tau)$  targets of individual treatment effect. We use these individual treatment effects as realizations of the CATE target  $\mathbb{E}[Y_{t+\tau}^{(\mathbf{a}^*)} - Y_{t+\tau}^{(0)} \mid H_t]$ .

## F.2 MIMIC-III SEMI-SYNTHETIC DATASET

**MIMIC Data and Semi-Synthetic Benchmark** We use the **MIMIC-Extract** pipeline (Wang et al., 2020) to preprocess the **MIMIC-III** ICU dataset (Johnson et al., 2016), aggregating hourly clinical measurements. Missing values are addressed via forward and backward filling, and continuous time-varying covariates are normalised. Our semi-synthetic dataset extends Schulam & Saria (2017) to achieve ground-truth treatment effects.

**Cohort & Features:** A cohort of 1,000 patients is selected, restricted to ICU stays of 50–100 hours ( $T^{(i)} \in [50, 100]$ ). The feature space includes 18 time-varying vital signs and 3 static features (gender, ethnicity, age), where categorical variables are one-hot encoded. These variables, dynamic and static, are concatenated together to form a covariate vector  $X \in \mathcal{X}$  with  $d_x = \dim(\mathcal{X}) = 18 + 2 + 3 + 1 = 25$ .

**Untreated Outcomes:** For each patient  $i$ , the untreated trajectory  $Z_t^{(i)}$  are simulated as:

$$Z_t^{(i)} = \underbrace{\alpha_s \mathbf{B}\text{-spline}(t)}_{\text{Endogenous}} + \underbrace{\alpha_g g^{(i)}(t)}_{\text{Exogenous}} + \underbrace{\alpha_f f_Z(X_t^{(i)})}_{\text{Noise}} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, 0.005^2), \quad (66)$$

where  $\alpha_s^j, \alpha_g^j, \alpha_f^j \in [0, 1]$  control component weights. The **endogenous** term combines: (1) a global trend  $\mathbf{B}\text{-spline}(t)$  sampled from three cubic splines; (2) patient-specific local variations  $g^{j,(i)}(t)$  from a Gaussian process with Matérn kernel (length-scale  $\ell = 2.0$ ). The **exogenous** term is modeled by the effect  $f_Z(X_t^{(i)})$  via random Fourier features (RFF), which approximates Gaussian process. By combining the three terms we are able to simulate a complex untreated outcome that has an intrinsic trend modeled by (1) endogenous dependencies on a different scale (B-spline for the global trend and RFF for the local perturbations), and (2) an exogenous dependency on the time-varying covariates  $X_t$ . The untreated outcome model presents a challenging task which requires complex deep neural networks.

**Treatment Assignment:** 2 binary treatments  $\mathbf{A}_t^l, l \in \{1, 2\}$  are assigned sequentially using:

$$p_t^l = \sigma(\gamma_A^l m(Y_{t-w:t-1}) + \gamma_X^l f_X^l(X_t) + b^l), \quad A_t^l \sim \text{Bernoulli}(p_t^l), \quad (67)$$

where ( $\gamma_A^1 = 2.5, \gamma_X^1 = 1.0, \gamma_A^2 = 2.0, \gamma_X^2 = 0.25$ ) control confounding strength,  $b = (-3.5, -1.75)$  biases treatment probability, and  $f_X^l(\cdot)$  uses RFF approximations of Gaussian process similar to  $f_Y$ . The term  $m(Y_{t-w:t-1})$  computes the average of prior outcomes over a window  $[t - w, t - 1]$ .

**Treatment Effects:** Treatments affect outcomes via cumulative effects that decay over time:

$$E(t) = \sum_{i=t-w}^t \sum_{l=1}^2 \frac{\beta_l A_i \kappa^l(X_i)}{\sqrt{t-i+1}}, \quad \text{where } \kappa^l(X_i) = \tanh((\omega^l)' X_i) + 1 \quad (68)$$

where  $w = 5$  defines the effect window, and the inverse square root decay ensures maximal effect  $\beta_l$  at application time  $t$ . Effects are summed across all treatments to model the local treatment effect at time  $t$ . Sparse dependencies are enforced: Outcomes are affected by at most 3 covariates. The function  $\kappa^l$  further provides heterogeneity to the treatment effect by scaling the full effect  $\beta$ , parametrized by the coefficients  $\omega^l \in \mathbb{R}^{d_x}$  (also sparse). These arrangements **increase** the challenge to estimate treatment effect compared to previous benchmarks (Hess et al., 2025; Melnychuk et al., 2022).

**Observed outcomes:** Finally, we combine the untreated outcome and the treatment effects together to form the observed outcome trajectories:

$$Y_t^i = Z_t^i + E(t)^i \quad (69)$$

**Simulating counterfactuals:** To generate the counterfactual outcome under an intervention  $\mathbf{a}^* \in \mathbb{R}^{(\tau+1) \times 2}$  during  $[t, t + \tau]$ , we intervene on the treatment effect  $E(t)$  by replacing the observed treatments  $A_t$  by the intervened treatments  $\mathbf{a}^*$  while keeping the untreated trajectory  $Z_t$  and covariates  $X_t$  unchanged (note that the covariates  $X_t$  are predefined vital signs which remain unaffected during the DGP). The counterfactual outcome  $Y_{t+\tau}^{\text{CF}}$  at time  $s \in [t, t + \tau]$  is:

$$Y_{t+\tau}^{(\mathbf{a}^*)} \mid H_t = Z_{t+\tau} + \underbrace{\sum_{i=t+\tau-w}^s \sum_{l=1}^2 \frac{\beta_l \cdot a_i^{*,l} \cdot \kappa^l(X_i)}{\sqrt{t+\tau-i+1}}}_{E^*(t+\tau)},$$

---

1728 where:  $Z_{t+\tau}$  is the untreated outcome (identical to the observed scenario). The treatments are divided by intervention  
1729 period and the observed period (history)  
1730

$$1731 A_i^{*,l} = \begin{cases} a_i^{*,l} & \text{if } t \leq i \leq t + \tau \text{ (intervention period),} \\ A_i^l & \text{otherwise (observed treatment).} \end{cases}$$

1733  
1734 And  $\kappa^l(X_i) = \tanh((\omega^l)^\top X_i) + 1$  encodes effect heterogeneity via covariates. The simulation is performed for the  
1735 units only on the **test** set in a sliding window treatment pattern. This means we start by conditioning on  $H_1$  to simulate  
1736 counterfactual outcome at time  $\tau + 1$  and end at  $t = T - \tau$  to simulate counterfactual outcome at time  $T$ . This results  
1737 in  $N_{test} \times (T - \tau)$  counterfactual outcomes under intervention  $\mathbf{a}^*$ .

1738 The full implementation is available in our GitHub repository, including parameter configurations for reproducibility.  
1739  
1740  
1741  
1742  
1743  
1744  
1745  
1746  
1747  
1748  
1749  
1750  
1751  
1752  
1753  
1754  
1755  
1756  
1757  
1758  
1759  
1760  
1761  
1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781

## 1782 G BASELINE METHODS

1783 We select six baselines to demonstrate the performance of our DeepBlip for CATE estimation over time. They are (1)  
 1784 History adjusted plug-in learner with LSTM instantiation (**HA-TRM**) (Frauen et al., 2025), (2) recurrent marginal  
 1785 structural networks (**R-MSNs**) (Lim, 2018), (3) counterfactual recurrent network (**CRN**), (Bica et al., 2020) (4) **G-Net**  
 1786 (Rui Li et al., 2021), (5) G-transformer (**GT**) (Hess et al., 2025), and (6) causal transformer (**CT**) (Melnychuk et al.,  
 1787 2022). We will provide details for each baseline and briefly introduce the theory framework of MSM and G-computation.  
 1788 The hyperparameter tuning is detailed in Appendix I.5 and Appendix I.6.

### 1791 G.1 HISTORY ADJUSTED WITH TRANSFORMER (HA-TRM)

1792 A naïve approach for estimating CATE over time is to create a regressor with the treatment as the condition:

$$1793 \hat{\delta}_\theta(a_{t:t+\tau}, h_t) \approx \mathbb{E}[Y_{t+\tau} \mid A_{t:t+\tau} = a_{t:t+\tau}, H_t = h_t]$$

1794 where  $\hat{\delta}_\theta$  is a non-parametric model (like neural networks). Then we estimate the CATE of  $\mathbf{a}^*$  and  $\mathbf{b}^*$  as:

$$1795 \hat{\delta}_\theta(\mathbf{a}^*, h_t) - \hat{\delta}_\theta(\mathbf{b}^*, h_t) \approx \mathbb{E}[Y_{t+\tau} \mid A_{t:t+\tau} = \mathbf{a}^*, H_t = h_t] - \mathbb{E}[Y_{t+\tau} \mid A_{t:t+\tau} = \mathbf{b}^*, H_t = h_t]$$

1796 This method is given the name **PI-HA learner** (plugged-in history adjusted meta-learner) from the work of Frauen et al.  
 1800 (2025). We instantiate this method with the transformer architecture to encode the history into a latent representation  
 1801 (see Appendix I), which is then fed into a fully connected linear network to predict the outcome.

1802 Since this approach does not adjust for time-varying confounding, it is biased (Frauen et al., 2025). While an advantage  
 1803 of the HA-TRM is low-variance, it is subject to the level of confounding within the dataset as we observed in the  
 1804 experiment.

### 1806 G.2 MARGINAL STRUCTURAL MODELS (MSM) AND ITS NEURAL EXTENSION

1807 **MSM Framework:** Marginal Structural Models (MSMs) (Hernán et al., 2001; Robins et al., 2000) address time-varying  
 1808 confounding via inverse probability of treatment weighting (**IPTW**). IPTW re-weights the targets to create a pseudo  
 1809 population that approximates the randomized controlled trial. With projection horizon  $\tau$ , the stabilized weight (**SW**) at  
 1810 time  $t$  for each sample is defined as:

$$1811 SW(t, \tau) = \prod_{n=t}^{t+\tau} \frac{f(\mathbf{A}_n | \bar{\mathbf{A}}_{n-1})}{f(\mathbf{A}_n | \bar{\mathbf{H}}_n)}, \quad (70)$$

1812 where  $f(\mathbf{A}_n | \bar{\mathbf{A}}_{n-1})$  and  $f(\mathbf{A}_n | \bar{\mathbf{H}}_n)$  represent treatment probabilities conditioned on past treatments or full history.  
 1813 These treatments are assumed discrete, often binary. The probabilities are estimated via logistic regressions. In practice  
 1814 the weights are truncated at 1st and 99th percentiles to avoid numeric overflow and then renormalised (Lim, 2018).  
 1815 However, this division of probabilities still creates significant instability due to the sequential multiplication of the  
 1816 propensities in Eq. (70).

1817 **Recurrent Marginal Structural networks (R-MSNs):** R-MSNs (Lim, 2018) extend MSMs via LSTM networks.  
 1818 Four components constitute the whole R-MSNs: propensity treatment network, propensity history network, encoder,  
 1819 decoder. Each sub-network uses LSTM as the sequential modeling architecture and a fully connected network to  
 1820 produce the desired prediction. Here is a detailed description of each model:

- 1821 • *Propensity Network (conditioned on past treatment)* estimates numerator of  $SW(t, \tau)$ .
- 1822 • *Propensity Network (conditioned on history)* estimates denominator of  $SW(t, \tau)$ .
- 1823 • *Encoder* maps history  $\bar{\mathbf{H}}_t$  to latent representation
- 1824 • *Decoder* predicts  $\tau$ -step outcomes using the representation generated by the encoder from last step.

1825 In our experiment, we replace the LSTM backbone of the encoder and decoder by transformers to ensure fairness.  
 1826 Training occurs in 3 phases: (1) First the two propensity networks are trained to predict the binary treatment probabilities  
 1827 using binary cross entropy losses. Then the stabilized weights can be computed for subsequent training of encoder and  
 1828 decoder. (2) Encoder is trained to predict the 1-step ahead outcome with the computed  $SW(\cdot, 1)$ -weighted MSE. (3) The  
 1829 decoder takes the latent representation from the encoder and then transform the vector into the latent representation fit  
 1830 for the decoder via a memory adapter (normally a fully connected linear layer). The decoder then predict the  $\tau$ -step  
 1831 outcome with the  $SW(\cdot, \tau)$ -weighted MSE.

In each phase, the training loss is aggregated through averaging the local loss at each time step  $0 \leq t \leq T - \tau$  (identical to our DeepBlip). At each local time step  $t$ , a representation is built with the transformer for the tasks. For more details please refer to Lim (2018).

### G.3 CAUSAL ESTIMATION VIA BALANCED REPRESENTATIONS

**Counterfactual recurrent network (CRN):** CRN (Bica et al., 2020) employs adversarial learning to create treatment-invariant representations. CRN heuristically addresses the time-varying confounding through creating temporal representations that are non-predictive of the treatment assignment. For this an LSTM encoder-decoder architecture with gradient reversal layers is built to enforce balanced representation:

- *Encoder:* LSTM produces hidden states  $\mathbf{h}_t$
- *Balanced Representation:*  $\Phi_t = \text{FC}(\mathbf{h}_t)$
- *Adversarial Loss:*

$$\mathcal{L} = \|\mathbf{Y}_{t+1} - G_Y(\Phi_t, \mathbf{A}_t)\|^2 - \lambda \sum_{j=1}^{d_a} \mathbf{1}_{[\mathbf{A}_t = a_j]} \log G_A(\Phi_t) \quad (71)$$

Here,  $G_Y$  predicts outcomes while  $G_A$  attempts (but fails due to gradient reversal) to predict treatments from  $\Phi_t$ . We fix  $\lambda = 1$  as in original implementations.

**Causal transformer (CT):** CT (Melnychuk et al., 2022) is the first transformer-based approach that estimates potential outcomes over time. CT employs a multi-input transformer architecture specifically designed to create history representations  $\Phi_t$  under time-varying treatment effect estimation setting. On top of  $\Phi_t$ , two additional networks are built:

- $G_Y$ : The outcome prediction network, which predicts the 1-step-ahead outcome  $\mathbf{Y}_{t+1}$  using  $\Phi_t$  and the current treatment  $\mathbf{A}_t$ .
- $G_A$ : The treatment classifier network, which attempts to predict the current treatment  $\mathbf{A}_t$  from  $\Phi_t$ .

Similar to CRN, CT attempts to create representations that are **predictive of outcomes** but **non-predictive of treatment assignments**, thereby mitigating confounding bias. To this, CT proposes counterfactual domain loss (CDC). The loss consists of: (1) Factual Outcome Loss:  $\mathcal{L}_{G_Y} = \|\mathbf{Y}_{t+1} - G_Y(\Phi_t, \mathbf{A}_t)\|^2$  This mean squared error ensures that  $\Phi_t$  is useful for predicting  $\mathbf{Y}_{t+1}$ . (2): Confusion Loss:  $\mathcal{L}_{\text{conf}} = -\sum_{j=1}^{d_a} \frac{1}{d_a} \log G_A(\Phi_t)$  This cross-entropy loss encourages  $G_A$  to output a uniform distribution over treatments, making  $\Phi_t$  treatment-invariant.

The training of CT involves optimizing two adversarial objectives:

- Optimize representation parameters ( $\theta_R$ ) and  $G_Y$  parameters ( $\theta_Y$ ):  $(\hat{\theta}_Y, \hat{\theta}_R) = \arg \min_{\theta_Y, \theta_R} \mathcal{L}_{G_Y}(\theta_Y, \theta_R) + \alpha \mathcal{L}_{\text{conf}}(\hat{\theta}_A, \theta_R)$
- Optimize the  $G_A$  parameters ( $\theta_A$ ):  $\hat{\theta}_A = \arg \min_{\theta_A} \alpha \mathcal{L}_{G_A}(\theta_A, \hat{\theta}_R)$

During inference, CT autoregressively predicts outcomes  $\hat{Y}_{t+k}$  using the balanced representations  $\Phi_{t+k}$  based on the observed history  $H_t$ .

### G.4 G-COMPUTATION-BASED NETWORK (G-NET, G-LSTM)

Both methods are based on the G-computation formula (Robins, 1986; Robins et al., 1999):

$$\mathbb{E}[Y_{t+\tau}^{(a_{t:t+\tau})} | H_t = h_t] = \mathbb{E} \left\{ \mathbb{E} \left[ \dots \mathbb{E} \left\{ \mathbb{E}[Y_{t+\tau} | H_{t+\tau}, A_{t:t+\tau} = a_{t:t+\tau}] \right. \right. \right. \\ \left. \left. \left. \left| H_{t+\tau-1}, A_{t:t+\tau-1} = a_{t:t+\tau-2} \right\} \right. \right. \\ \left. \left. \left. \dots \left| H_{t+1}, A_{t:t+1} = a_{t:t+1} \right| \right| H_t = \bar{h}_t, A_t = a_t \right] \right\} \quad (72)$$

Due to the nested structure of G-computation in 72, the model misspecification error get easily propagated through the iterative steps, leading to potential bias and high variance.

**G-Net (Simulation-based)** G-Net(Rui Li et al., 2021) is the first neural network model that uses G-computation to estimate CAPO over time. It make the Monte-Carlo simulation based on the integration:

$$\int_{\mathbb{R}^{d_x \times (\tau)} \times \mathbb{R}^{d_y \times (\tau)}} \mathbb{E}[Y_{t+\tau} | H_{t+\tau} = h_{t+\tau-1}, A_{t:t+\tau} = a_{t:t+\tau}] \times \prod_{\delta=1}^{\tau} p(x_{t+\delta}, y_{t+\delta} | H_t = h_t, x_{t+1:t+\delta-1}, y_{t:t+\delta-1}, a_{t:t+\delta-1}) d(x_{t+1:t+\tau-1}, y_{t:t+\tau-1}) \quad (73)$$

G-Net predicts the potential outcome in two steps:

1. Estimate the conditional distribution  $p(x_{t+\delta}, y_{t+\delta} | H_t = h_t, x_{t+1:t+\delta-1}, y_{t:t+\delta-1}, a_{t:t+\delta-1})$
2. Compute the empirical sum of M (M=50) trajectories via Monte-Carlo simulation to estimate the integration in 73.

G-Net originally uses LSTM to encode the history into a hidden state, with an extra linear transformation layer mapping the hidden state to the latent representation. We replace the LSTM architecture by a transformer in our experiment to ensure fairness. The representation is fed into output computing head (Fully connected networks with respective activation output) to predict the conditional distribution. For fairness, we replace the LSTM encoder with the encoder of the transformer.

**G-transformer (regression-based)** Unlike G-Net, which uses Monte-Carlo simulation, G-transformer (**GT**) (Hess et al., 2025) is built via a iterative regression on the pseudo outcomes defined as:

1.  $G_{t+\tau}^a = Y_{t+\tau}$
2.  $G_{t+\delta}^a = \mathbb{E}[G_{t+\delta+1}^a | \bar{H}_{t+\delta}^t, A_{t:t+\delta} = a_{t:t+\delta}]$  for  $\delta = 0, \dots, \tau - 1$

GT predicts the outcome in a masked transformer encoder. It uses transformer to encode the history at  $t + \delta$  into a representation  $Z_{t+\delta} = z_\phi(h_{t+\delta})$ , and then predicts the pseudo outcome at step  $\delta$  as:

$$g_\phi^\delta(Z_{t+\delta}^a, A_{t+\delta}) \approx G_{t+\delta}^a$$

And then the CAPO is predicted as  $g_\phi^0(Z_t, a_t)$ . The encoder  $z_\phi(\cdot)$  is instantiated by a transformer architecture similar to CT (For more details see Hess et al. (2025)).

## G.5 DYNAMIC DML

The Dynamic DML method (heterogeneous version) is implemented as originally designed (See link to the code (Supplemental material)), using sequential g-estimation to solve the blip predictors. Since the original method was developed for static covariates, we adapt it to our time-series setting by concatenating all variables in the patient history  $H_t$  into a single, high-dimensional feature vector. This vector then serves as the conditioning variable for the Dynamic DML model. To ensure fairness, the nuisance residuals  $(\tilde{Y}, \tilde{A})$  are taken directly from our trained Stage 1 nuisance network, providing both methods with inputs of the same quality. However, in its second stage, Dynamic DML is limited because it can only accept a history variable from a fixed time step. This limitation prohibits the use of a complex neural network in scenarios where one needs to take dynamic patient history as input. Hence, we follow the original implementation and use the LASSO regression to model the blip coefficient predictor.

---

## 1944 H ALGORITHM

---

### 1946 Algorithm 1 DeepBlip training algorithm

---

1948 1: **Input:** Dataset  $\mathcal{D}$ , horizon  $\tau$ , feature map  $\phi$ , number of folds  $W$ , learning rates  $\eta_N, \eta_B$

1949 2: **Output:** Blip prediction network  $(\mathcal{E}_{\theta_B}^B, \mathbf{gb}_{\theta_B}^k) \hat{\psi}_{\theta_B}$

1950 3: 

---

1951 4: **First Stage: Train Nuisance Network**  $(\mathcal{E}_{\theta_N}^N, \mathbf{gp}_{\theta_N}^k, \mathbf{gq}_{\theta_N}^{j,k})$

1952 5: Partition  $\mathcal{D}$  into  $W$  folds:  $S_1, S_2, \dots, S_W$

1953 6: Partition  $\mathcal{D}$  into  $W$  folds:  $S_1, S_2, \dots, S_W$

1954 7: **for** each fold  $w = 1$  to  $W$  **do**

1955 8:     Train nuisance network on  $\mathcal{D} \setminus S_w$

1956 9:     **for** each  $t = 1$  to  $T - \tau$ ,  $k = 0$  to  $\tau$  **do**

1957 10:          $z_{t+k}^N \leftarrow \mathcal{E}_{\theta}^N(h_{t+k})$

1958 11:          $\hat{p}_k(z_{t+k}^N) \leftarrow \mathbf{gp}_{\theta_N}^k(z_{t+k}^N)$

1959 12:          $\mathcal{L}_{t,k}^p = (Y_{t+\tau} - \hat{p}_k(z_{t+k}^N))^2$

1960 13:         **for**  $j = k$  to  $\tau$  **do**

1961 14:              $\hat{q}_{j,k}(z_{t+k}^N) \leftarrow \mathbf{gq}_{\theta_N}^{j,k}(z_{t+k}^N)$

1962 15:              $Q_{t,j} = \phi(X_{t+1:t+j}, A_{t:t+j}) - \phi(X_{t+1:t+j}, (A_{t:t+j-1}, 0))$

1963 16:              $\mathcal{L}_{t,j,k}^q = BCE(\hat{q}_{j,k}(z_{t+k}^N), Q_{t,j})$  (if binary treatment)

1964 17:         **end for**

1965 18:     **end for**

1966 19:      $\mathcal{L}_N \leftarrow \frac{1}{T-\tau} \sum_{t=1}^{T-\tau} \left( \frac{1}{\tau+1} \sum_{k=0}^{\tau} \mathcal{L}_{t,k}^p + \frac{2}{(\tau+1)(\tau+2)} \sum_{\tau \geq j \geq k \geq 0} \mathcal{L}_{t,j,k}^q \right)$

1967 20:     {Compute the gradient update the nuisance network parameter  $\theta_N$ }

1968 21:      $\theta_N \leftarrow \theta_N - \eta_N \nabla_{\theta_N} \mathcal{L}_N$

1969 22:     **for** each sample in  $S_w$  **do**

1970 23:         Compute the residuals:

1971 24:          $\tilde{Y}_{t,k} = Y_{t+\tau} - \mathbf{gp}_{\theta_N}^k(z_{t+k}^N)$

1972 25:          $\tilde{A}_{t,j,k} = A_j - \mathbf{gq}_{\theta_N}^{j,k}(z_{t+k}^N)$

1973 26:     **end for**

1974 27: **end for**

1975 28: 

---

1976 29: **Second Stage: Train Blip Prediction Network**  $(\mathcal{E}_{\theta_B}^B, \mathbf{gb}_{\theta_B}^k)$

1977 30: **for** each  $t = 1$  to  $T - \tau$  **do**

1978 31:     Compute the hidden state of the encoder:  $z_t^B \leftarrow \mathcal{E}_{\theta_B}^B(h_t)$

1979 32:     **for**  $k = 0$  to  $\tau$  **do**

1980 33:          $\hat{\psi}_k^1(z_t^B) \leftarrow \mathbf{gb}_{\theta_B}^k(z_t^B)$

1981 34:          $\hat{\psi}_k^2(z_t^B) \leftarrow \mathbf{gb}_{\theta_B}^k(z_t^B)$

1982 35:          $\mathcal{L}_{\text{blip},t,k} = \left( \tilde{Y}_{t,k} - \sum_{j=k+1}^{\tau} \hat{\psi}_j^2(z_t^B)' \tilde{A}_{t,j,k} - \hat{\psi}_k^1(z_t^B)' \tilde{A}_{t,k,k} \right)^2$

1983 36:     **end for**

1984 37: **end for**

1985 38:  $\mathcal{L}_{\text{blip}} \leftarrow \frac{1}{T-\tau} \sum_{t=1}^{T-\tau} \left( \frac{1}{\tau+1} \sum_{k=0}^{\tau} \mathcal{L}_{\text{blip},t,k} \right)$

1986 39: Compute the gradient update the blip prediction network parameter  $\theta_B$

1987 40:  $\theta_B \leftarrow \theta_B - \eta_B \nabla_{\theta_B} \mathcal{L}_{\text{blip}}$

1988 41: 

---

1990 *Note:* We use  $\leftarrow$  operation for value assignment in the computation graph while  $\leftarrow\leftarrow$  operations are detached from the computation graph.

---

1998  
1999  
2000  
2001  
2002  
2003  
2004  
2005  
2006  
2007  
2008  
2009  
2010  
2011  
2012  
2013  
2014  
2015  
2016  
2017  
2018  
2019  
2020  
2021  
2022  
2023  
2024  
2025  
2026  
2027  
2028  
2029  
2030  
2031  
2032  
2033  
2034  
2035  
2036  
2037  
2038  
2039  
2040  
2041  
2042  
2043  
2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051

---

**Algorithm 2** DeepBlip inference algorithm

---

- 1: **Input:** Patient history  $H_t = h_t$ , treatment sequence  $a^*$ , baseline treatment sequence  $b^*$
  - 2: **Output:** CATE estimate  $\mathbb{E}[Y_{t+\tau}^{(a^*)} - Y_{t+\tau}^{(b^*)} \mid H_t = h_t]$
  - 3: **Step 1: Compute the hidden state of the encoder**
  - 4:  $z_t^B \leftarrow \mathcal{E}_{\theta_B}^B(h_t)$
  - 5: **Step 2: Predict the blip coefficients**
  - 6:  $\hat{\psi}_k(z_t^B) \leftarrow \text{gb}_{\theta_B}^k(z_t^B)$
  - 7: **Step 3: Estimate CATE**
  - 8: (when  $\phi$  only depends the last treatment:  $\phi(x_{t+1:t+k}, a_{t:t+k}) = \phi(a_{t+k})$ )
  - 9:  $\text{CATE} = \sum_{k=0}^{\tau} \hat{\psi}_k(z_t^B)'(\phi(a_k^*) - \phi(b_k^*))$
-

## I IMPLEMENTATION DETAILS

In subsection I.2, we introduce the runtime of our system. We provide a runtime comparison of the different methods in subsection I.2. Then in subsection I.4 and I.3, we introduce how to instantiate our DeepBlip with LSTM and transformers respectively. In subsection I.5 and subsection I.6, we report the details of hyperparameter tuning.

### I.1 SOFTWARE & SYSTEM

Our DeepBlip is implemented using **PyTorch** (Paszke et al., 2019), combined with **PyTorch Lightning** (Falcon et al., 2019) for streamlined code organization and model training. Our code is available at <https://anonymous.4open.science/t/DeepBlip-A39B>. Training and inference are performed on a single GeForce RTX 4080 with 18GB memory.

### I.2 RUNTIME COMPARISON

In this section we report: (1) The runtime for training a single epoch on the training set, (2) The runtime for inference on the test set, and (3) The theoretical runtime required to identify the optimal treatment offline. The last metric is crucial for efficient offline treatment planning in practical medical scenarios. During the test, we choose transformer as the main backbone for all the selected methods. For two stage methods like RMSNs and DeepBlip, the first stage uses a simpler LSTM instantiation instead for higher efficiency. The results are reported in table 3.

Baselines like HA-TRM, RMSNs, CRN, CT estimate one conditional average potential outcome (CAPO) over time within a single inference. Therefore it takes  $N(\tau)$  inferences to identify the treatment sequence with best outcome, where  $N(\tau)$  denotes the number of possible treatment combinations within the time window  $[t, t + \tau]$ . G-Net computes the CAPO over time by performing Monte Carlo simulations, generating  $M$  trajectories for each possible treatment sequence. G-transformer only predicts the CAPO under a fixed intervention sequence, which means it requires additional  $N(\tau) - 1$  re-trainings to calculate all the possible outcomes. Our DeepBlip, however, estimates CATE over time via predicting blip coefficients to identify the treatment combination with the optimal effect. By Eq. 4, only one forward pass to predict the blip coefficients suffices to identify the optimal treatment. Hence, although our DeepBlip is not the most efficient method during training or inference, **it is significantly faster in offline treatment planning**, which is critical for personalized medicine.

### I.3 TRANSFORMER INSTANTIATION FOR DEEPLIP

Our DeepBlip uses the transformer architecture to encode history into a latent representation. The transformer is built upon the multi-input transformer from causal transformer (Melnychuk et al., 2022). The transformer is specially designed for medical application where the inputs include the outcomes, covariates and treatments. Each type of variable has a corresponding sub-transformer. The three sub-transformers perform not only the classic multi-headed self-attention within themselves, but also the cross-attention in-between. This ensures that information is shared across these sub-transformers.

Let the three sub-transformers be  $z_{\theta}^k(\cdot)$ ,  $k = 1, 2, 3$ . Suppose for  $k = 1$ , the input sequence is the  $U^k = (X_1, \dots, X_t)$ . Then the sub-transformer  $z_{\phi}^k(\cdot)$  generate the representation in the following steps.

Methods	Training ( $t_T$ )	Inference ( $t_I$ )	Offline optimal treatment identification
HA-TRM (Frauen et al., 2025)	$2.5 \pm 0.4$	$0.16 \pm 0.04$	$t_I \cdot N(\tau)$
RMSNs (Lim, 2018)	$7.2 \pm 0.9$	$0.30 \pm 0.06$	$t_I \cdot N(\tau)$
CRN (Bica et al., 2020)	$5.4 \pm 0.6$	$0.28 \pm 0.06$	$t_I \cdot N(\tau)$
CT (Melnychuk et al., 2022)	$3.4 \pm 0.5$	$0.20 \pm 0.04$	$t_I \cdot N(\tau)$
G-Net (Rui Li et al., 2021)	$2.8 \pm 0.4$	$0.24 \pm 0.03$	$t_I \cdot N(\tau) \cdot M$ †
GT (Hess et al., 2025)	$2.6 \pm 0.4$	$0.20 \pm 0.04$	$(N(\tau) - 1) \cdot t_T \cdot n_e + t_I \cdot N(\tau)$ ††
<b>DeepBlip (ours)</b>	$3.6 \pm 0.5$	$0.17 \pm 0.06$	$t_I$

†: G-Net performs  $M$  simulations to infer the potential outcome.

††: G-transformer requires re-training additionally  $N(\tau) - 1$  times.

Table 3: Runtime report for the methods instantiated by transformers.  $t_T$  (in min) is the training time on the training dataset per epoch.  $t_I$  (in min) is the inference time on the test dataset. In the last column we report the time needed to identify the optimal treatment.  $N(\tau)$  is the number of all possible combination of treatments during the future horizons  $[t : t + \tau]$ .

**Input transformation:** The raw input first goes through a linear transformation:

$$H_0^k = \text{Linear}_k(U) \in \mathbb{R}^{t \cdot d_{\text{lm}}} \quad (74)$$

, where  $H_0^k$  is the input of the first transformer blocks.

**Transformer blocks:** The multi-input sub-transformer has  $B$  transformer blocks. Within each transformer block, the multi-headed self-attention and cross-attention are performed. Then the output goes through a feed forward network. Here are the details of a transformer block at layer  $b$ .

(1) Multi-headed self/cross-attentions First the keys, queries and values, namely  $K, Q, V$  for  $n_h$  attention heads, are computed from the output from last transformer block:

$$\begin{aligned} Q_b^{k,i} &= H_b^{k,i} W_Q^{k,i} + \mathbf{b}_Q^{k,i} \in \mathbb{R}^{t \cdot d_{qkv}}, \\ K_b^{k,i} &= H_b^{k,i} W_K^{k,i} + \mathbf{b}_W^{k,i} \in \mathbb{R}^{d_{t \cdot qkv}}, \\ V_b^{k,i} &= H_b^{k,i} W_V^{k,i} + \mathbf{b}_V^{k,i} \in \mathbb{R}^{d_{t \cdot qkv}}, \end{aligned} \quad (75)$$

where  $i \in [1, n_h]$  is the index of a single attention head. Then the  $i$ -th attention is computed via a softmax over the scaled dot-product:

$$\text{Attn}_b^{k,i} = \text{Softmax}\left(\frac{Q_b^{k,i} (K_b^{k,i})^T}{\sqrt{d_{qkv}}}\right) V_b^{k,i} \quad (76)$$

The multi-head self-attention is then calculated by concatenating the attention heads:

$$\text{MHA}(Q_b^k, K_b^k, V_b^k) = \text{Concat}(\text{Attn}_b^{k,1}, \dots, \text{Attn}_b^{k,n_h}) \in \mathbb{R}^{t \cdot d_{\text{lm}}} \quad (77)$$

Since  $\text{MHA}(Q_b^k, K_b^k, V_b^k)$  has the same dimension as  $H_b^k$ , then we could generate a new set of queries for the cross-attentions with residual connection:

$$\tilde{Q}_b^k = H_b^k + \text{MHA}(Q_b^k, K_b^k, V_b^k) \quad (78)$$

The cross-attentions are put on top of the self-attention layers and uses the queries and keys from the other two sub-transformers. For example, the multi-head cross-attention of sub-transformer  $k$  with sub-transformer  $m$  is:

$$\text{MHA}(\tilde{Q}_b^k, K_b^m, V_b^m) = \text{Concat}(\text{CrossAttn}_b^{k,m,1}, \dots, \text{CrossAttn}_b^{k,m,n_h}), \quad (m \neq k), \quad (79)$$

where the CrossAttn is defined as:

$$\text{CrossAttn}_b^{k,m,i} = \text{Softmax}\left(\frac{\tilde{Q}_b^k (K_b^{m,i})^T}{\sqrt{d_{qkv}}}\right) V_b^{m,i}. \quad (80)$$

Finally, the output is achieved by adding self-attention's output and cross-attention's output together:

$$Z_b^k = \tilde{Q}_b^k + \sum_{m \neq k} \text{MHA}(\tilde{Q}_b^k, K_b^m, V_b^m) + S. \quad (81)$$

The static covariates  $S$  are added when pooling different cross-attention outputs (see section 4.1 from Melnychuk et al. (2022)). The operations ensures that the information is shared across sub-transformers. Layer normalizations are added for all the self- and cross attentions (Vaswani et al., 2017).

(2) Feed-forward networks The output of the attention mechanism is processed through a position-wise feed-forward network (FFN) applied to each time step:

$$H_b^{k,\text{ff}} = \text{Linear}(\text{ReLU}(\text{Linear}(Z_b^k))), \quad (82)$$

where linear layers are followed by a dropout (Melnychuk et al., 2022). The output  $H_b^{k,\text{ff}}$  serves as the input ( $H_{b+1}^k$ ) to the next transformer block. After processing through all  $B$  blocks, the final representation for sub-transformer  $k$  becomes  $H_B^k$ . These final representations are then aggregated via a element-wise average pooling followed by a linear transformation and a exponential linear unit (ELU):

$$H_B = \text{ELU}(\text{Linear}(\text{Pool}(H_B^1, H_B^2, H_B^3))) \in \mathbb{R}^{t \cdot d_{\text{hr}}}. \quad (83)$$

Hence, the transformer-based sequential encoder  $\mathcal{E}_\theta$  outputs the latent vector:  $\mathcal{E}_\theta(H_t) = \mathbf{hr}_t = H_{B,t} \in \mathbb{R}^{d_{\text{lm}}}$  at time step  $t$ . We omit the positional encoding for better clarity. For more relevant details see section 4.2 from Melnychuk et al. (2022).

Our DeepBlip has two stages: (1) Nuisance network, and (2) Blip prediction network. Each stage uses a separate sequential encoder:  $\mathcal{E}_{\theta_N}^N$  for (1) and  $\mathcal{E}_{\theta_B}^B$  for (2). To accelerate the training, we only instantiate  $\mathcal{E}_{\theta_B}^B$  by the transformer, and use LSTM to instantiate (1). The underlying reason is that the demand for accuracy in Stage (1) is not as high as Stage (2), as the  $L^2$ -moment loss is Neyman orthogonal over the nuisance functions.

**Nuisance Network:** At each time step  $t$ , the nuisance network Estimates the nuisance functions:

$$q_{t,j,k}(h_{t+k}) := \mathbb{E}[Y_{t+\tau} \mid H_{t+k} = h_{t+k}], \quad 1 \leq t \leq T - \tau, 0 \leq k \leq \tau \quad (84)$$

$$q_{t,j,k}(h_{t+k}) := \mathbb{E}[Q_{t,j} \mid H_{t+k} = h_{t+k}], \quad 1 \leq t \leq T - \tau, 0 \leq k \leq j \leq \tau \quad (85)$$

Therefore a **multi-head output layer** is added on top of the encoder  $\mathcal{E}_{\theta_N}^N$ . The heads receive the representation  $\mathbf{hr}_t^N$  from the encoder and transform the input with multi-layer perceptron networks:

1.  $\hat{p}_{t,k}(h_{t+k}) = \text{MLP}_p^{(k)}(\mathbf{hr}_{t+k}^N) \approx p_{t,k}(h_{t+k}), \quad k = 0, 1, \dots, \tau$
2.  $\hat{q}_{t,j,k}(h_{t+k}) = \text{MLP}_q^{(j,k)}(\mathbf{hr}_{t+k}^N) \approx q_{t,j,k}(h_{t+k}), \quad 0 \leq k \leq j \leq \tau$

where  $\text{MLP}_p^{(k)}$  and  $\text{MLP}_q^{(j,k)}$  are fully connected networks with ReLU activations. Each MLP has input size of  $d_{\text{hr}}$ , a single hidden layer with  $d_{\text{hidden}}$  perceptrons and an single output transformation that maps the hidden layer to a 1-dimensional output. For binary output in some  $q_{t,j,k}$ , an additional sigmoid activation is added to contract the range into  $[0, 1]$ .

**Blip prediction network's output** Likewise a multi-head output layer is added on top of the sequential encoder  $\mathcal{E}_{\theta_B}^B$  to predict the blip coefficients. For each horizon  $k \in \{0, 1, \dots, \tau\}$ :

$$\hat{\psi}_k(h_t) = \text{MLP}_{\theta_B}^{(k)}(\mathbf{hr}_t^B) \approx \psi_{t,k}(h_t)$$

where  $\text{MLP}_{\theta_B}^{(k)}$  maps  $\mathbf{hr}_t^B$  to blip coefficients.  $\text{MLP}_{\theta_B}^{(k)}$  has the same structure as the MLPs in stage (1).

#### I.4 LSTM INSTANTIATION FOR DEEPBLIP

DeepBlip-LSTM uses LSTM as the sequential encoders in two stages: 1. Stage (1) (Nuisance Network): Estimates residuals for blip function estimation. 2. Stage (2) (Blip Prediction Network): Predicts blip parameters  $\psi_{t,k}(h_t)$ .

**Sequential Encoding via LSTM** Each stage shares an individual LSTM to encode patient history. Let  $\mathbf{hz}_t$  be the hidden state at time  $t$ . At each time step  $t$ , the input vector  $\mathbf{v}_t \in \mathbb{R}^{d_{\text{input}}}$  concatenates:

1. Static features  $X_s \in \mathbb{R}^{d_{\text{static}}}$  (remain constant across time)
2. Current time-varying covariates  $X_d \in \mathbb{R}^{d_{\text{dynamic}}}$  ( $d_{\text{dynamic}} + d_{\text{static}} = d_x$ )
3. Previous treatments  $a_{t-1} \in \mathbb{R}^{d_a}$
4. Previous outcomes  $Y_{t-1} \in \mathbb{R}$

to form the input vector:

$$\mathbf{v}_t = \text{Concat}(X_s, X_d, A_{t-1}, Y_{t-1}) \in \mathbb{R}^{d_{\text{input}}}, \quad d_{\text{input}} = d_x + d_a + 1.$$

Then the LSTM processes  $\{\mathbf{v}_t\}_{t=1}^T$  into hidden state  $\mathbf{hs}_t$  and cell state  $c_t$ :

$$(\mathbf{hs}_t, c_t) = \text{LSTM}(\mathbf{v}_t, (\mathbf{hs}_{t-1}, c_{t-1})),$$

where  $\mathbf{h}_t \in \mathbb{R}^{d_{\text{hidden}}}$ . Next we take the hidden state and derive the temporal representation of the patient state at time  $t$ :

$$\mathbf{hr}_t = \text{ELU}(\mathbf{W}_{\text{hr}} \cdot \text{dropout}(\mathbf{hs}_t) + b_{\text{hr}}),$$

where  $\mathbf{W}_{\text{hr}} \in \mathbb{R}^{d_{\text{hr}} \times d_{\text{hidden}}}$  is a learnable projection matrix.

Above is the whole process of generating the temporal representation at time  $t$ . Since the DeepBlip has two stages, therefore we need to train **two separate** sequential encoders:  $\mathcal{E}_{\theta_N}^N(h_t) = \mathbf{hr}_t^N$  for the nuisance network and  $\mathcal{E}_{\theta_B}^B(h_t) = \mathbf{hr}_t^B$  for the blip prediction network. The output layer of both stages remain the same as the DeepBlip-TRM.

## I.5 HYPERPARAMETER TUNING FOR TUMOR GROWTH DATASET

Table 4: Hyperparameter Tuning for Methods on Tumor Growth Data: We perform a random grid search with 20 iterations and choose the best hyperparameters for each task.  $C = d_{\text{input}}$  is the dimension of the input. For the causal transformer, the CDC coefficient is  $\alpha = 0.01$  (Melnychuk et al., 2022). The number of MC samples of G-Net is 50 (Lim, 2018). Models are either instantiated by LSTMs or transformers. Here we display the search ranges of the *original* instantiations for each method. In the experiments, we adopt a universal instantiation type for all the baselines for fairness. Hence, the performance comparison in experiment section is fair.

Method	Component	Hyperparameter	Tuning Range
HA-TRM (Frauen et al., 2025)	(end-to-end)	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 10
CRN (Bica et al., 2020)	Encoder	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) LSTM dropout rate ( $p$ ) Balanced representation size ( $d_r$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [32, 64, 128, 256] [0.1, 0.2, 0.3] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] 10
	Decoder	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) Balanced representation size ( $d_r$ ) FC hidden units ( $d_{\text{fc}}$ ) LSTM dropout rate ( $p$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [256, 512, 1024] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.1, 0.2, 0.3] 20
CT (Melnychuk et al., 2022)	(end-to-end)	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 20
R-MSNs (Lim, 2018)	Treatment/History Propensity Network	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) FC hidden units ( $d_{\text{fc}}$ ) LSTM dropout rate ( $p$ ) Max gradient norm Number of epochs ( $n_e$ )	1 [0.1, 0.01, 0.001] [64, 128, 256] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] [0.1, 0.2] [0.5, 1.0, 2.0] 10
	Encoder / Decoder	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) FC hidden units ( $d_{\text{fc}}$ ) LSTM dropout rate ( $p$ ) Max gradient norm Number of epochs ( $n_e$ )	1 [0.01, 0.001, 0.0001] [256, 512, 1024] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] [0.1, 0.2, 0.4] [0.5, 1.0, 2.0] 20
G-Net (Rui Li et al., 2021)	(end-to-end)	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) LSTM output size ( $d_o$ ) Feed-forward hidden units ( $d_{\text{ff}}$ ) LSTM dropout rate ( $p$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [64, 128, 256] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.1, 0.2] 20
GT (Hess et al., 2025)	(end-to-end)	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 10
DeepBlip (ours)	Nuisance Network	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) LSTM dropout rate ( $p$ ) hidden representation size ( $d_r$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [32, 64, 128, 256] [0.1, 0.2] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] 10
	Blip Prediction Network	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 20

## I.6 HYPERPARAMETER TUNING FOR MIMIC-III DATASET

Table 5: Hyperparameter Tuning for Methods on MIMIC-III semi-synthetic Data: We perform a random grid search with 20 iterations and choose the best hyperparameters for each task.  $C = d_{\text{input}}$  is the dimension of the input. For the causal transformer, the CDC coefficient is  $\alpha = 0.01$  (Melnichuk et al., 2022). The number of MC samples of G-Net is 50 (Lim, 2018). Models are either instantiated by LSTMs or transformers. Here we display the search ranges of the *original* instantiations for each method. In the experiments, we adopt a universal instantiation type for all the baselines for fairness. Hence, the performance comparison in experiment section is fair.

Method	Component	Hyperparameter	Tuning Range
HA-TRM (Frauen et al., 2025)	(end-to-end)	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 10
CRN (Bica et al., 2020)	Encoder	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) LSTM dropout rate ( $p$ ) Balanced representation size ( $d_r$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [32, 64, 128, 256] [0.1, 0.2, 0.3] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] 10
	Decoder	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) Balanced representation size ( $d_r$ ) FC hidden units ( $d_{\text{fc}}$ ) LSTM dropout rate ( $p$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [256, 512, 1024] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.1, 0.2, 0.3] 20
CT (Melnichuk et al., 2022)	(end-to-end)	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 20
R-MSNs (Lim, 2018)	Treatment/History Propensity Network	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) FC hidden units ( $d_{\text{fc}}$ ) LSTM dropout rate ( $p$ ) Max gradient norm Number of epochs ( $n_e$ )	1 [0.1, 0.01, 0.001] [64, 128, 256] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] [0.1, 0.2] [0.5, 1.0, 2.0] 10
	Encoder / Decoder	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) FC hidden units ( $d_{\text{fc}}$ ) LSTM dropout rate ( $p$ ) Max gradient norm Number of epochs ( $n_e$ )	1 [0.01, 0.001, 0.0001] [256, 512, 1024] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] [0.1, 0.2, 0.4] [0.5, 1.0, 2.0] 20
G-Net (Rui Li et al., 2021)	(end-to-end)	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) LSTM output size ( $d_o$ ) Feed-forward hidden units ( $d_{\text{ff}}$ ) LSTM dropout rate ( $p$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [64, 128, 256] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.1, 0.2] 20
GT (Hess et al., 2025)	(end-to-end)	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 10
DeepBlip (ours)	Nuisance Network	LSTM layers ( $l$ ) Learning rate ( $\eta$ ) LSTM hidden units ( $d_h$ ) LSTM dropout rate ( $p$ ) hidden representation size ( $d_r$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2, 3] [0.1, 0.01, 0.001] [32, 64, 128, 256] [0.1, 0.2] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] [0.5 $d_h$ , 1 $d_h$ , 2 $d_h$ ] 10
	Blip Prediction Network	Transformer blocks ( $B$ ) Learning rate ( $\eta$ ) Batch size Attention heads ( $n_h$ ) Transformer units ( $d_{\text{attn}}$ ) hidden representation ( $d_{\text{hr}}$ ) FC hidden units ( $d_{\text{fc}}$ ) Number of epochs ( $n_e$ )	[1, 2] [0.01, 0.001, 0.0001] [64, 128] 2 [0.5C, 1C] [0.5C, 1C] [0.5 $d_{\text{hr}}$ , 1 $d_{\text{hr}}$ ] 20