

Review

Not peer-reviewed version

A Review of Deep Learning for Individualized Treatment Effect Estimation in Healthcare Time Series

Zi Cai^{*}, Yu Liu, [Tingting Zhu](#)

Posted Date: 14 April 2026

doi: 10.20944/preprints202604.0958.v1

Keywords: causal inference; individualized treatment effect; time-varying treatment effect; deep learning; healthcare



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

A Review of Deep Learning for Individualized Treatment Effect Estimation in Healthcare Time Series

Zi Cai *, Yu Liu  and Tingting Zhu 

Department of Engineering Science, University of Oxford, Oxford, UK

* Correspondence: zi.cai@lincoln.ox.ac.uk

Abstract

The proliferation of healthcare time series data presents a transformative opportunity for personalized healthcare. However, utilizing these observational data for clinical decision support requires moving beyond risk prediction to causal inference. In particular, estimating time-varying individual treatment effects is crucial for answering “what-if” questions in dynamic settings, where patient states and treatment efficacy evolve continuously over time. Traditional statistical and machine learning methods face well-documented difficulties in capturing the long-range temporal dependencies, high dimensionality, and non-linear treatment-response dynamics characteristic of such data. Deep learning has emerged as a powerful paradigm for addressing these challenges, motivating a rapidly growing body of work on time-varying causal inference from healthcare time series data. This review aims to provide a systematic overview of deep learning for time-varying ITE estimation in healthcare, connecting theoretical foundations with methodological practice. We first establish the theoretical foundations, then present a taxonomy of deep learning methodologies, followed by an analysis of current clinical applications and model evaluation practices. Finally, we discuss key unresolved challenges and outline a roadmap for future research on time-varying ITE estimation.

Keywords: causal inference; individualized treatment effect; time-varying treatment effect; deep learning; healthcare

1. Introduction

The paradigm of modern healthcare is undergoing a profound transformation, shifting from generalized treatment protocols to personalized healthcare [1]. In clinical decision-making, one of the ultimate objectives is the ability to predict how a specific patient will respond to a specific treatment at a specific time. Although traditional predictive modeling has excelled at estimating risk (e.g., predicting mortality given current vitals), it often fails to answer the critical counterfactual questions: “What would happen to this patient’s blood pressure if we administered drug A instead of drug B right now?” Answering this requires moving beyond association to causal inference [2]. Unlike correlation, which merely observes that two events occur together, causality asserts that the intervention (the treatment) is the mechanism responsible for the outcome. Fundamentally, this re-frames the learning objective: rather than mapping observed features to observed labels, the goal is to predict outcomes under treatments that may not have occurred. Since a patient can only receive one treatment sequence at a time, alternative outcomes (*counterfactuals*) are never observed, effectively turning causal inference into a complex missing data problem [3].

This need for individualized prediction has motivated a transition from estimating the Average Treatment Effect (ATE) to the Individual Treatment Effect (ITE). While the former is a measure of how a treatment affects a population on average, the latter seeks to quantify the causal effect of an intervention at the level of a single patient. Clinical heterogeneity means that a treatment beneficial for the “average” patient may be ineffective or harmful to a specific individual. Although early work on ITE estimation focused on static settings (a single decision made from baseline features) [4], this view is

misaligned with clinical reality. Healthcare is inherently dynamic and longitudinal. It is characterized by sequential decision-making, where treatments vary over time, patient states evolve in response to those treatments, and updated patient states influence future treatment decisions. For example, in cancer care, the clinician must decide at each stage whether to administer chemotherapy or radiation based on the patient's evolving tumor response, organ function, and tolerance to treatment. Each treatment decision alters the patient's state, which in turn shapes subsequent therapeutic choices. This temporal process is conceptualized in Figure 1, which tracks the history of observed patient states (i.e., covariates) and observed treatments leading up to a specific decision time t . Consequently, time-varying ITE estimation has emerged as a critical topic of research, with the aim of utilizing this history to estimate the causal impact in a future instance, specifically illustrated as a τ -step-ahead prediction. As visually distinguished in the diagram, the core objective is to quantify the treatment effect by comparing the factual outcome, the observed value following the prescribed treatment sequence, against the counterfactual outcome, the hypothetical value that would have been observed had an alternative treatment sequence been prescribed.

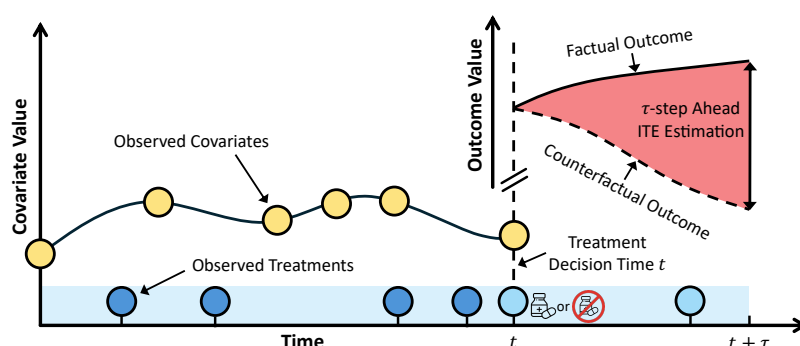


Figure 1. Illustration of τ -step-ahead time-varying ITE estimation. Observed covariates (yellow circles) and treatments (dark blue circles) recorded up to decision time t inform the assignment of future treatments (light blue circles) over the horizon $[t, t + \tau - 1]$. The τ -step-ahead ITE (shaded red region) is defined as the difference between the factual and counterfactual outcomes at $t + \tau$, corresponding to alternative treatment sequences initiated at time t .

The growing availability of longitudinal Electronic Health Records (EHRs) provides an unprecedented opportunity for estimating time-varying ITEs directly from routine clinical practice. EHRs capture time-stamped sequences of patient covariates (such as vital signs and laboratory results) alongside treatment administrations, forming rich healthcare time series data that naturally mirror the sequential treatment-response structure underlying time-varying ITE estimation. Classical approaches to this problem, notably Marginal Structural Models [5] and G-computation [6], provide strong theoretical foundations but face well-documented difficulties in scaling to the high-dimensional, irregular, and non-linear nature of such data [7]. These limitations have motivated a rapidly growing body of deep learning research that leverages expressive temporal architectures, including Recurrent Neural Networks (RNNs) [7,8] and Transformers [9], to learn latent representations of patient history and model complex interactions between patient covariates and assigned treatments over time.

Despite the growing body of work on time-varying ITE estimation, existing surveys do not adequately reflect the breadth of this expanding field. Most reviews on ITE estimation [13–15] focus predominantly on the static setting and do not account for the temporal complexities inherent in healthcare time series data. A comprehensive survey of this field requires coverage across several dimensions: causal theory specific to the time-varying setting, the deep learning architectures that underpin modern methods, clinical applications, and systematic model evaluation. As shown in Table 1, no existing review addresses all of these dimensions. Bica et al. [10] provide the broadest early coverage but predate many recent developments in time-varying ITE. Zhang et al. [11] discuss time-varying causal theory but do not cover deep learning approaches. Ghosheh et al. [12] address methodology and evaluation but omit theoretical foundations and clinical applications. This survey

Table 1. Comparison of recent surveys and reviews on ITE estimation. The table indicates whether each work accounts for deep learning methods (DL), causal inference and time-varying ITE theory, clinical applications, framework evaluation, and an approximate count of reviewed literature situated at the intersection of time-varying ITE estimation and deep learning.

Review	Year	DL	Theory	Application	Evaluation	#Literature [†]
Bica et al. [10]	2021	✓	✓	✓	✓	<10
Zhang et al. [11]	2024		✓			<10
Ghoshch et al. [12]	2025	✓			✓	<10
Our Work	2026	✓	✓	✓	✓	30+

[†] Literature count specifically refers to works addressing time-varying ITE estimation using deep learning architectures. Methods framed under the static setting or adopts statistical or traditional machine learning backbones are not included.

addresses these gaps, reviewing more than 30 methods to provide a unified and systematic guide to deep learning for time-varying ITE estimation in healthcare, spanning theoretical foundations, methodological taxonomy, clinical applications, and model evaluation. To support continued tracking of this rapidly evolving field, we maintain an up-to-date list of deep learning methods for time-varying ITE estimation in healthcare at https://github.com/caizicharles/TV_ITE_Review.git.

The remainder of this article is structured as follows: Section 2 establishes the necessary theoretical foundation, defines foundational concepts, illustrates the specific challenge of time-dependent confounding, and details the identifiability assumptions required for treatment effect estimation. Transitioning from theory to technical execution, Section 3 delineates the standard training and inference setups. Section 4 presents a comprehensive taxonomy of methodologies categorized by their adjustment mechanisms, followed by an examination of their clinical applications in Section 5. Section 6 addresses the critical issue of model validation, examining datasets, metrics, and baselines. Finally, we conclude by analyzing current challenges and proposing future research directions in Section 7 and Section 8, respectively.

Search Strategy

This survey considers the literature published up to December 2025. We conducted a primary search using Google Scholar, supplemented by targeted searches on PubMed and Scopus, and validated through recursive backward and forward citation tracking on seminal works. Our search strategy was structured around three keyword dimensions: causality (“causal inference,” “individualized treatment effect,” “conditional average treatment effect,” “heterogeneous treatment effect”, “counterfactual prediction”), temporality (“time-varying,” “time series”, “temporal,” “sequential”), and healthcare (“healthcare,” “electronic health records”). Search queries were constructed by combining at least one term from each dimension. Studies were included if they (i) proposed or substantially extended a deep learning-based method to estimate individualized treatment effects in time-varying or longitudinal settings, and (ii) evaluated their approach in at least one medical or healthcare-related dataset, whether real or synthetic. We excluded studies that only estimate population-level average treatment effects, those without a deep learning backbone, and theses or dissertations. We further exclude works that assume randomized treatment assignment or do not properly account for time-dependent confounding. The survey articles and editorials were consulted for completeness but were not included as primary literature.

2. Theoretical Foundation

This section presents the theoretical foundations for time-varying ITE estimation in healthcare time series data. We begin by categorizing ITE estimation settings based on whether patient covariates and treatments are static or time-varying, narrowing the focus of this survey to specific contexts. We then formalize the problem, introducing the relevant notation and the target quantity to be estimated. Next, we describe the core challenge that arises when estimating treatment effects from observational data, where treatment decisions are influenced by a patient’s evolving history. Finally, we present the

necessary assumptions under which this challenge can be addressed and describe how they enable the estimation of the time-varying ITE.

2.1. Problem Categorization

Problems in ITE estimation can be broadly categorized by the time-dependent nature of the available data. At a high level, these data consist of covariates (i.e., features describing the patient's state, such as age and vital signs) and treatments (i.e., interventions applied to the patient, such as drugs and medical procedures). Based on whether these elements are static (i.e., fixed at a single point) or time-varying (i.e., changing over time), we identify four distinct quadrants of analysis:

1. **Static Covariate & Static Treatment** The classic *static* ITE estimation problem in which baseline (pre-treatment) covariates are used to estimate the effect of a single fixed treatment. This is the fundamental scenario for answering questions, such as "Does Drug A work better than a placebo for patients given the baseline disease severity?"
2. **Time-Varying Covariate & Static Treatment** A single treatment is administered, but the decision to treat (or not) is based on a sequence of covariates. For example, deciding whether and when to perform an organ transplant is a one-time decision based on past lab results and health vitals.
3. **Static Covariate & Time-Varying Treatment** A sequence of treatments is administered based solely on baseline covariates, without adaptation to intermediate observations. Although uncommon in healthcare, this setting arises in industrial processes where, for example, a fixed sequence of temperature and pressure set points is predetermined from initial material properties and executed without adjustment during processing.
4. **Time-Varying Covariate & Time-Varying Treatment** Both covariates and treatments are subject to change and may influence one another over time. This is the most common and complex scenario, such as the management of a patient in the Intensive Care Unit (ICU), where a physician repeatedly measures vital signs, uses them to administer a dose of a blood pressure drug, and then observes the new vital signs to decide on the next dose, with the intention of optimizing a downstream outcome such as 30-day survival.

This review focuses on the complexities introduced by healthcare data series, with a focus on settings (2) and (4) for the estimation of ITE that varies over time. These two settings share the feature that covariates evolve over time, but differ in a key respect: setting (4) introduces a feedback loop in which each treatment alters the patient's future covariate trajectory, which in turn informs subsequent treatment decisions, whereas setting (2) involves a single treatment assigned after observing the full covariate history.

2.2. Formalization and Notation

To formalize time-varying ITE estimation, we adopt the Potential Outcomes (PO) framework (also known as the Neyman-Rubin Causal Model), originally developed for static treatment effect estimation [16–18], and extended to the time-varying setting [6,19]. We introduce the notation for the most general setting (4), from which the other settings can be viewed as special cases. Throughout, we adopt the following conventions: (i) Capital letters without an overbar (e.g., X_k , A_k , H_k) denote random variables or random vectors. (ii) Lowercase letters (e.g., x_k , a_k , h_k) denote specific, non-random values in the corresponding spaces (e.g., $x_k \in \mathcal{X}_k$, $a_k \in \mathcal{A}_k$). Depending on the context, these may represent either observed values or prescribed (hypothetical) treatment assignments. (iii) Calligraphic symbols (e.g., \mathcal{X}_k , \mathcal{A}_k , \mathcal{H}_k) denote non-random sets or spaces. All notations are summarized in Appendix A.

Consider a longitudinal study where our goal is to estimate a causal effect at a decision time t on an outcome that will be measured in τ time steps in the future, at time $t + \tau$. While time-varying ITE estimation can be formulated in continuous-time, we consider a discrete formalization for simplicity. At each discrete time step $k \in \{0, \dots, t, \dots\}$, we observe a sequence of data for a single individual (we omit the individual index i for notational simplicity in most cases):

- **(Pre-treatment) Covariates** $X_k \in \mathcal{X}_k$: A d_x -dimensional random vector of patient features at time k . The covariate space \mathcal{X}_k may comprise both continuous (e.g., vital signs, lab results) and discrete (e.g., diagnosis codes) components. Each covariate feature may be static (i.e., remaining constant across all time steps such as sex or ethnicity), or time-varying, changing in response to disease progression and prior treatments. A particular realized covariate vector is denoted $x_k \in \mathcal{X}_k$.
- **Treatment** $A_k \in \mathcal{A}_k$: The treatment assignment at time k is a random vector taking values in the treatment space \mathcal{A}_k . The treatment space may be discrete (e.g., a finite set of medication options), continuous (e.g., drug dosage), or multivariate when multiple treatments are administered concurrently. A realized treatment vector at time k is denoted by $a_k \in \mathcal{A}_k$.
- **Outcome** $Y_{t+\tau} \in \mathcal{Y}$: The outcome of interest, measured τ steps after the decision time t . The outcome space \mathcal{Y} may be binary (e.g., $\in \{0, 1\}$ for mortality) or continuous (e.g., $\in \mathbb{R}$ for blood pressure). A realized outcome is denoted $y_{t+\tau} \in \mathcal{Y}$.

In setting (2), with time-varying covariates and a single static treatment administered at the time of decision t , the treatment history \bar{A}_{t-1} consists of a sequence of null interventions, reflecting that no treatment has been administered prior to the current decision.

We use an overbar to denote the sequence of a variable up to a given time [8,9], for example,

$$\bar{X}_t = (X_0, \dots, X_t), \quad \bar{A}_t = (A_0, \dots, A_t).$$

At decision time t , before the treatment A_t is chosen, the *observed history* is denoted by a random vector

$$H_t = (\bar{X}_t, \bar{A}_{t-1}).$$

H_t represents the random vector of history, $h_t = (\bar{x}_t, \bar{a}_{t-1}) \in \mathcal{H}_t$ is a particular realization, where $\mathcal{H}_t = \prod_{k=0}^t \mathcal{X}_k \times \prod_{k=0}^{t-1} \mathcal{A}_k$ is the corresponding history space.

As established in settings (2) and (4), our interest lies in estimating the causal effect of treatments administered over the interval from decision time t to outcome measurement at $t + \tau$. During this interval, the patient continues to receive treatments at each time step, denoted as $(A_t, \dots, A_{t+\tau-1})$. A sequence of treatments from t to $t + \tau - 1$ is a fixed, non-random sequence of vectors defined as:

$$\bar{a}_{[t, t+\tau-1]} := (a_t, a_{t+1}, \dots, a_{t+\tau-1})$$

where each $a_k \in \mathcal{A}_k$ is a specified treatment. To reiterate, $\bar{a}_{[t, t+\tau-1]}$ is a prescribed (non-random) treatment sequence used to index hypothetical outcomes and need not correspond to any treatment sequence that was actually administered. In the PO framework, each individual is conceptualized as possessing a hypothetical outcome for every treatment sequence that could be administered. These are termed *potential outcomes*. Let $Y_{t+\tau}(\bar{a}_{[t, t+\tau-1]})$ denote the potential outcome that would be observed at time $t + \tau$, if the individual were administered the treatment sequence $\bar{a}_{[t, t+\tau-1]}$ over the interval $[t, t + \tau - 1]$. Note that the potential outcome indexes only the future treatment sequence from t onward. This reflects the fact that, at decision time t , the individual's history H_t has already been realized and is kept fixed. Only future treatments remain subject to intervention. For brevity, we write the potential outcomes as $Y_{t+\tau}^{\bar{a}}$ and $Y_{t+\tau}^{\bar{a}'}$, where the superscripts $\bar{a} = \bar{a}_{[t, t+\tau-1]}$ and $\bar{a}' = \bar{a}'_{[t, t+\tau-1]}$ denote two distinct prescribed treatment sequences.

We now define the causal *estimands*, the theoretical quantities that encode the treatment effects of interest. The most granular estimand of interest is the τ -step-ahead ITE:

Definition 1 (Time-Varying Individualized Treatment Effect (ITE)). Given a decision time t , a horizon $\tau \geq 1$, and two prescribed future treatment sequences \bar{a} and \bar{a}' from t to $t + \tau - 1$, the τ -step-ahead time-varying ITE for individual i is defined as

$$\text{ITE}_{i,t,\tau}(\bar{a}, \bar{a}') = Y_{i,t+\tau}^{\bar{a}} - Y_{i,t+\tau}^{\bar{a}'}. \quad (1)$$

The index i is reintroduced here to emphasize that the ITE is a quantity specific to a single unit, capturing the heterogeneity of treatment effects across individuals that is otherwise lost in population-level averages. However, identifying the ITE for a single individual is impossible due to the *fundamental problem of causal inference* [20]: at time $t + \tau$, only the factual outcome $Y_{i,t+\tau}$ corresponding to the sequence of treatment that the patient actually received is observed. The alternative outcome is a missing *counterfactual*. This is intuitive, as one cannot simultaneously “prescribe” and “not prescribe” a medication for an individual. This motivates the following estimable quantities that average over populations or sub-populations. Correspondingly, the time-varying Average Treatment Effect (ATE) compares two treatment sequences at the population level:

Definition 2 (Time-Varying Average Treatment Effect (ATE)). *The τ -step-ahead time-varying ATE at decision time t is the expectation over the population of individuals of the difference between potential outcomes:*

$$\text{ATE}_{t,\tau}(\bar{a}, \bar{a}') = \mathbb{E}\left[Y_{t+\tau}^{\bar{a}} - Y_{t+\tau}^{\bar{a}'}\right]. \quad (2)$$

Although this provides a useful population-level summary, the ATE does not capture the heterogeneity of the treatment effect, which is central to personalized healthcare. A more informative target is the time-varying Conditional Average Treatment Effect (CATE), which conditions on the observed history:

Definition 3 (Time-Varying Conditional Average Treatment Effect (CATE)). *The τ -step-ahead time-varying CATE at decision time t is the conditional expectation of the difference between potential outcomes given the observed history H_t . Formally, it is defined as*

$$\delta_{t,\tau}(h_t; \bar{a}, \bar{a}') = \mathbb{E}\left[Y_{t+\tau}^{\bar{a}} - Y_{t+\tau}^{\bar{a}'} \mid H_t = h_t\right]. \quad (3)$$

Importantly, $\delta_{t,\tau}(h_t; \bar{a}, \bar{a}')$ is a conditional average effect rather than an individual-level quantity. However, as conditioning h_t becomes sufficiently rich and high-dimensional, the conditional average can serve as a practically meaningful approximation for the individualized effect $\text{ITE}_{i,t,\tau}$. This perspective underlies much of the literature on ITE estimation in both static and time-varying settings [4,8], where the estimand in Equation (3) is often informally referred to as the time-varying ITE. We adopt this convention throughout, and when the treatment sequences (\bar{a}, \bar{a}') and times (t, τ) are fixed or clear from the context, we write $\delta(h_t)$ for brevity.

2.3. The Challenge of Time-Dependent Confounding

The estimation of time-varying ITE from healthcare time series data introduces a fundamental challenge. Such data are observational in nature: treatment decisions are informed by the patient’s evolving state, including both current and prior covariates and treatments. These temporal dependencies cause *time-dependent confounding*, where a covariate X_k acts as a common cause for both the immediate assignment of treatment A_k and the final outcome $Y_{t+\tau}$. As a result, patients receiving different treatment sequences may differ systematically in their prognoses, and naive comparisons of outcomes across sequences conflate the causal effect of treatment with pre-existing differences in patient states. Correcting for this *confounding bias*¹ is the central methodological challenge addressed by the methods reviewed in this survey. We illustrate these causal relationships using a directed acyclic graph (DAG). Figure 2 shows a 2-step-ahead setting ($t = 0, \tau = 2$) with a baseline covariate X_0 , prescribed treatments A_0 and A_1 , intermediate covariate X_1 , and outcome Y_2 . The intermediate covariate X_1 has a dual role that defines time-dependent confounding:

- *Mediator:* X_1 is caused by A_0 and, in turn, affects Y_2 , forming the mediated path ($A_0 \rightarrow X_1 \rightarrow Y_2$). This is a valid component of the total causal effect of treatment A_0 on Y_2 .

¹ This should be distinguished from selection bias, which arises from conditioning on a common effect (collider) of treatment and outcome [19].

- *Confounder*: X_1 is a common cause of A_1 and Y_2 , creating a spurious association between them through the path ($A_1 \leftarrow X_1 \rightarrow Y_2$), which does not reflect the causal effect of treatment A_1 .

To estimate the causal effect of the treatment sequence $\bar{a} = (A_0, A_1)$, we must block the confounding path ($A_1 \leftarrow X_1 \rightarrow Y_2$) without blocking the mediated causal path ($A_0 \rightarrow X_1 \rightarrow Y_2$). Standard conditioning on X_1 (e.g., via regression) blocks the confounding path but simultaneously blocks the mediated causal path, removing a valid component of the total causal effect of A_0 on Y_2 from the estimate. This dilemma is not specific to the 2-step-ahead case. For any prediction horizon $\tau > 1$, every intermediate covariate between decision time and the τ -step-ahead outcome exhibits the same dual role, and standard conditioning methods fail at each step for the same reason. Therefore, an adjustment procedure is required to obtain valid estimations of time-varying ITE.

Remark 1 (Adjustment). *In the time-varying setting, adjustment refers to any procedure that eliminates confounding bias of variables X_k at each step k , without blocking the causal effects of past treatments $A_{j < k}$ that are mediated through those same variables.*

Standard adjustment methods such as regression [21] and static propensity score matching [22] are insufficient in the time-varying setting, as they provide no mechanism to condition on a variable for confounding control while simultaneously preserving its role as a mediator [6,19]. Methods capable of correctly accounting for time-dependent confounding exist, but they rely on a set of assumptions about the data-generating process to recover the treatment effect from observational data.




Figure 2. Causal graph illustrating time-dependent confounding for a 2-step-ahead prediction ($t = 0$, $\tau = 2$). The baseline covariate X_0 represents the observed patient state at decision time $t = 0$ and acts as a confounder, causally influencing both treatment assignments and the outcome Y_2 . Treatments A_0 and A_1 are prescribed over the intervention horizon at time steps 0 and 1. The intermediate covariate X_1 exhibits the dual role that defines time-dependent confounding: it mediates (denoted as \textcircled{M}) the effect of A_0 on Y_2 through the path ($A_0 \rightarrow X_1 \rightarrow Y_2$) and simultaneously confounds (denoted as \textcircled{C}) the effect of A_1 on Y_2 through the path ($A_1 \leftarrow X_1 \rightarrow Y_2$). This dual role recurs at every intermediate covariate in longer horizons and prevents standard adjustment from producing unbiased estimates.

2.4. Identifiability of Time-Varying ITE

To address confounding and obtain unbiased estimates of the time-varying ITE, we must impose assumptions on the data-generating process to establish *identifiability*: the property that a causal estimand can, in principle, be expressed as a functional of the observed data distribution [19]. Recall that the estimand of interest at the decision time t and the horizon τ is the time-varying ITE in Equation (3): $\delta(h_t) = \mathbb{E} \left[Y_{t+\tau}^{\bar{a}} - Y_{t+\tau}^{\bar{a}'} \mid H_t = h_t \right]$. Because earlier treatments affect subsequent covariates, which in turn affect both future treatment decisions and outcomes, we require assumptions that are formulated sequentially over time:

1. **Consistency:** Let $\bar{a}_{[0, t+\tau-1]} = (\bar{a}_{t-1}, \bar{a}_{[t, t+\tau-1]})$ denote a fixed full treatment sequence, where \bar{a}_{t-1} is a particular realized past treatment sequence (a value of \bar{A}_{t-1}) and $\bar{a}_{[t, t+\tau-1]}$ is a prescribed future treatment sequence. For an individual, if the observed treatment sequence matches $\bar{a}_{[0, t+\tau-1]}$, then the observed outcome is equal to the corresponding potential outcome:

$$Y_{t+\tau} = Y_{t+\tau}(\bar{a}_{[0, t+\tau-1]}), \text{ if } \bar{A}_{t-1} = \bar{a}_{t-1} \text{ and } \bar{A}_{[t, t+\tau-1]} = \bar{a}_{[t, t+\tau-1]}. \quad (4)$$

This implies that there are no multiple versions of the treatment, i.e., consistency requires that the treatment definition be sufficiently precise so that technical variations do not change the potential

outcome. In the static ITE estimation setting, the corresponding assumption is the *Stable Unit Treatment Value Assumption (SUTVA)* [23], which consists of two components: (i) *consistency* as defined above, and (ii) *no interference* which states that an individual's potential outcomes do not depend on the treatment assignments of others. Both components remain necessary in the sequential setting. However, in modern deep learning literature, no interference is typically absorbed into the statistical setup rather than stated as a standalone assumption. Seminal works [8,24] explicitly list consistency as an assumption, but encode no interference implicitly by defining the dataset as a collection of N independent and identically distributed (i.i.d.) patient trajectories.

2. **Sequential Positivity:** For each time step $k \in \{t, \dots, t + \tau - 1\}$, every treatment value $a_k \in \mathcal{A}_k$ must be assignable with positive probability, conditional on any history $h_k \in \mathcal{H}_k$ with $\Pr(H_k = h_k) > 0$:

$$\Pr(A_k = a_k | H_k = h_k) > 0 \quad (5)$$

or, when A_k is continuous, $f_{A_k|H_k}(a_k|h_k) > 0$. More intuitively, the effect of a given treatment sequence cannot be identified for histories where that sequence (or parts of it) is impossible under the observed data. If a particular treatment a_k is never administered to patients with a history $H_k = h_k$, then we cannot use the observed data to learn the counterfactual outcome under that action in that stratum.

3. **Sequential Exchangeability (or Ignorability):** At each time step from baseline to end of the treatment horizon $k \in \{0, \dots, t + \tau - 1\}$, after conditioning on the history H_k , treatment A_k is independent of future potential outcomes under *any* treatment sequence. For the horizon- τ outcome, this can be written as

$$Y_{t+\tau}(\bar{a}_{[0, t+\tau-1]}) \perp\!\!\!\perp A_k | H_k \quad \forall k \in \{0, \dots, t + \tau - 1\}, \forall \bar{a}_{[0, t+\tau-1]} \in \prod_{j=0}^{t+\tau-1} \mathcal{A}_j. \quad (6)$$

This assumption must hold throughout history $k \in \{0, \dots, t + \tau - 1\}$, not only during the future treatment window $[t, t + \tau - 1]$. If unmeasured confounders influenced both treatment and covariates in any previous step $k < t$, then the observed history H_t itself depends on these confounders. Conditioning on H_t is therefore insufficient to make treatment independent of potential outcomes, and the resulting treatment effect estimates will be biased.

Under these assumptions, the time-varying ITE estimand $\delta(h_t)$ is formally identified. This *identification* result allows the estimand to be expressed as a functional of the observed data distribution. Any *estimation* of the desired estimand from finite data is therefore only valid as long as these assumptions hold in the setting under study.

3. From Theory to Practice

Building on the theoretical foundation established earlier, we provide an overview of the training and inference procedures for using observational healthcare time series data for ITE estimation, offering a better understanding of the ensuing methodologies. The observational dataset consists of time-varying covariates $x_{i,k}$, treatments $a_{i,k}$, and observed factual outcomes $y_{i,k+1}$ ² for the patient i , which are denoted as $\mathcal{D} = \left\{ (x_{i,k}, a_{i,k}, y_{i,k+1})_{k=0}^{T_i} \right\}_{i=1}^N$ for a total of N i.i.d. patients, each with a sequence length T_i . At any decision time t , the observed history $h_t = (\bar{x}_{i,t}, \bar{a}_{i,t-1})$ is constructed from the patient's observed covariate and treatment sequences up to time t . In practice, the covariate vector $x_{i,k}$ often subsumes past outcomes, since a previously observed outcome $y_{i,k}$ may influence both subsequent treatment decisions and future outcomes. Excluding it from $x_{i,k}$ would leave a confounder unadjusted, violating sequential exchangeability. Given this history h_t , the model estimates the expected outcome under a prescribed treatment sequence \bar{a} , from which the ITE estimate $\hat{\delta}(h_t)$ is obtained by contrasting

² Following standard practice in the literature, the outcome index is shifted to $k + 1$ to reflect the assumption that a treatment administered at time k affects the outcome observed at time $k + 1$.

predictions under two alternative sequences \bar{a} , \bar{a}' . Importantly, \bar{a} , \bar{a}' are prescribed sequences over the future horizon $[t, t + \tau - 1]$.

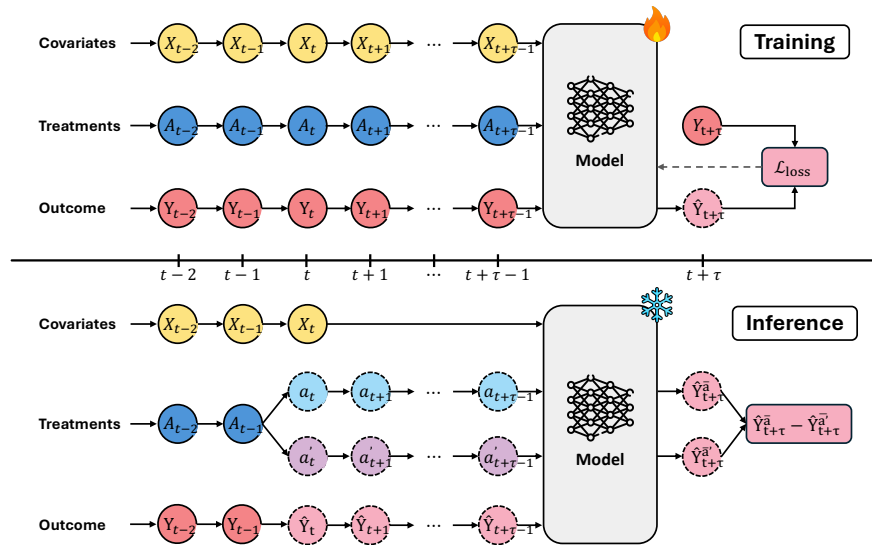




Figure 3. Visualization of the model training and inference stages for time-varying ITE estimation. The  and  icon separately denotes trainable and frozen model parameters. Solid and dashed circular outlines distinguishes between observed data and either prescribed or predicted values.

As shown in Figure 3, the ITE estimation model is first trained with patient data before performing inference. Without loss of generality, the description is provided in terms of a single patient. During training, the sequence of patient data is segmented and used to recursively generate 1-step-ahead predictions. The prediction and the ground truth value are used to compute the loss, which updates the weights of the model. The foundational training objective minimizes the factual outcome prediction loss \mathcal{L}_y , which measures the difference between predicted and observed outcomes across all patients and time steps, using *mean squared error (MSE)* for continuous outcomes and (*binary*) *cross-entropy* for discrete outcomes. Some methods have auxiliary objectives, such as predicting future covariates or treatment assignments.

Since observational data reflect confounded treatment assignments, the model must incorporate an adjustment mechanism to separate genuine treatment effects from spurious correlations. How this is achieved varies across methodologies (see Section 4). While some adjust through re-weighting the factual prediction loss, others introduce explicit regularization terms. The overall training objective \mathcal{L}_{loss} therefore encompasses both the factual outcome and auxiliary prediction losses, as well as method-dependent adjustment regularizations. In the inference stage, the data of a testing patient are partitioned according to a selected decision time t . The sequence prior to the decision time constitutes the observed history h_t . From the decision time, the model predicts the outcome recursively in a single step, given the prescribed treatment a_k at each time step up to the prediction horizon $t + \tau$. Finally, the estimate of the time-varying ITE is computed as the difference between the predicted outcomes $\hat{\delta}(h_t) = \hat{Y}_{t+\tau}^{\bar{a}} - \hat{Y}_{t+\tau}^{\bar{a}'}$, contrasting the predicted outcomes under alternative prescribed treatment sequences \bar{a} and \bar{a}' .

Table 2. Taxonomy of time-varying ITE estimation methodologies for healthcare time series data. The methods are grouped by the mechanism employed to adjust for confounders. For each method, we present the method name or abbreviation, publication year, outcome and treatment types (continuous and/or discrete), benchmarked healthcare datasets, and open-source availability where the repository is directly accessible by clicking on the tick mark.

Method	Year	Outcome	Treatment	Primary Datasets	Open
<i>Matching</i>					
SyncTwin [♠] [25]	2021	Continuous	Discrete	CPRD	✓
SCouT [♠] [26]	2023	Continuous	Discrete	CAMP, FA-COMS	✓
T4 [27]	2024	Continuous	Discrete	MIMIC-III, AUMCdb	✓
<i>G-Formula</i>					
G-Net [24]	2020	Both	Continuous	CVSim	
Su et al. [28]	2024	Both	Continuous	MIMIC-IV	
G-Transformer [29]	2024	Both	Both	CVSim, MIMIC-IV	
G-Transformer [30]	2024	Continuous	Discrete	PK-PD, MIMIC-III	✓
<i>Outcome Regression</i>					
CF-ODE ^{♠†} [31]	2022	Continuous	Discrete	CVM, PDM	✓
CURE [♠] [32]	2024	Discrete	Discrete	MS CCAE	✓
<i>Inverse Probability of Treatment Weighting</i>					
R-MSN [7]	2018	Continuous	Discrete	PK-PD	✓
DSW [33]	2020	Continuous	Discrete	MIMIC-III	✓
LipCDE [†] [34]	2023	Continuous	Discrete	MIMIC-III, COVID-19 DE	
LipSCDE [†] [35]	2023	Continuous	Discrete	MIMIC-III, COVID-19 DE	
Wu et al. [36]	2024	Continuous	Discrete	COVID-19 US	✓
SCIP-Net [†] [37]	2025	Continuous	Discrete	PK-PD, MIMIC-III	✓
VTNet [38]	2025	Continuous	Discrete	MIMIC-III, NACC	✓
<i>Doubly Robust</i>					
IVW-DR-learner [39]	2025	Continuous	Discrete	MIMIC-III	✓
<i>Representation Balancing</i>					
CRN [8]	2020	Continuous	Discrete	PK-PD, MIMIC-III	✓
CT [9]	2022	Continuous	Discrete	PK-PD, MIMIC-III	✓
TE-CDE [†] [40]	2022	Continuous	Discrete	PK-PD, MIMIC-III	✓
TCFimt [41]	2022	Continuous	Discrete	PK-PD	
CF-GODE [†] [42]	2023	Continuous	Discrete	PK-PD	
RCAB [43]	2023	Continuous	Continuous	PK-PD, MIMIC-III	
CAG-ODE [†] [44]	2024	Continuous	Discrete	PK-PD, COVID-19 US	✓
ACTIN [45]	2024	Continuous	Both	PK-PD, MIMIC-III	✓
Causal CPC [46]	2024	Continuous	Discrete	PK-PD, MIMIC-III	
OptAB [†] [47]	2024	Continuous	Discrete	MIMIC-IV	✓
CEET [48]	2025	Continuous	Discrete	PK-PD, MIMIC-III	
Liu et al. [49]	2025	Continuous	Discrete	PK-PD, MIMIC-III	
EDTS [50]	2025	Continuous	Discrete	PK-PD, EEG	✓
Mamba-CDSP [51]	2025	Continuous	Discrete	PK-PD, MIMIC-III	
Guski et al. [52]	2025	Continuous	Discrete	COVID-19 EU	✓
TV-SurvCaus [53]	2025	Continuous	Continuous	MIMIC-III	
<i>Representation Disentanglement</i>					
DCRN [54]	2021	Continuous	Discrete	PK-PD, MIMIC-III	
DR-CRN [55]	2022	Continuous	Discrete	MIMIC-III	✓

[♠] This superscript indicates methods that operate under problem setting (2) (see Section 2.1). All unmarked methods address setting (4).

[†] This superscript indicates methods formulated in continuous time. All unmarked methods adopts imputation and operate in discrete time.

4. Time-Varying ITE Estimation Methodology

This section reviews deep learning methodologies (see Table 2) for time-varying ITE estimation from healthcare time series data. We organized the methods by their *locus of adjustment*, i.e., the stage of the modeling pipeline at which confounding is addressed. As illustrated in Figure 4, adjustment strategies target distinct stages of the pipeline, including input space, model structure, latent representation, or training objective. Each subsection is structured as follows: We first define and explain the theory of how the adjustment mechanism mitigates confounding bias. We then review the deep learning methods that operationalize this adjustment mechanism to perform time-varying ITE estimation. Finally, we analyze the limitations of each mechanism along two axes specific to the time-varying setting. First, we identify the additional modeling requirements that each mechanism imposes for valid estimation beyond the shared identifiability assumptions in Section 2.4. These refer to the conditions necessary for valid adjustment, rather than to the architectural components common across methods. Second, we characterize the distinct failure modes that arise as the prediction horizon τ increases, which can undermine both theoretical validity and practical reliability in finite samples.

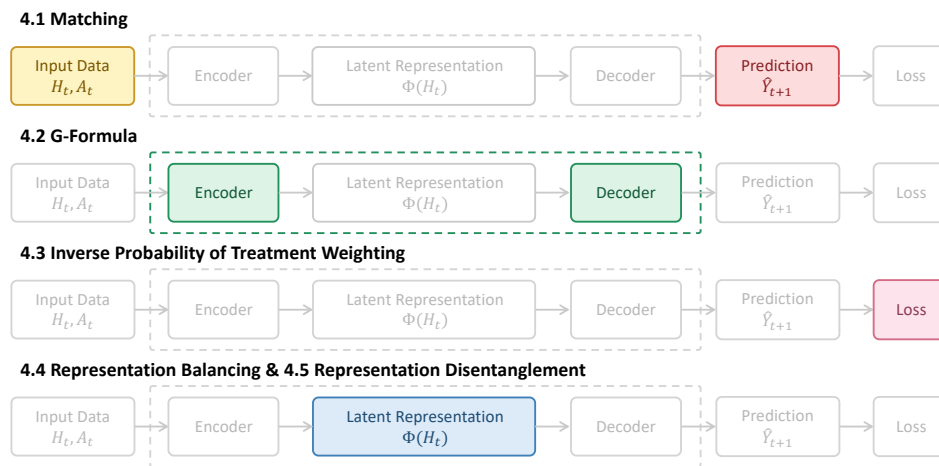


Figure 4. Figure depicting the locus of adjustment for different adjustment strategies on an exemplary time-varying ITE estimation pipeline. Each adjustment mechanism corresponds to a subsection in Section 4.

4.1. Matching

Matching adjusts for time-dependent confounding by estimating the unobserved counterfactual outcome from the outcomes of units with similar observed histories that followed a different treatment sequence. Under sequential exchangeability, units with sufficiently similar histories are exchangeable with respect to their potential outcomes, so that differences in observed outcomes can be attributed to the treatment rather than to pre-existing differences in patient state. To formalize, consider a unit i with an observed history $h_{i,t}$ for which we seek to estimate the potential outcome under a prescribed treatment sequence \bar{a}' . Matching draws from a *donor pool* that consists of units in the observational data \mathcal{D} whose observed treatment sequence over the horizon $[t, t + \tau - 1]$ corresponds to \bar{a}' . Given a donor pool, Matching constructs the counterfactual outcome as a weighted aggregation of donor outcomes:

$$\hat{Y}_{i,t+\tau}^{\bar{a}'} = \sum_{j \in \mathcal{D}, \bar{A}_{j,[t,t+\tau-1]} = \bar{a}'} w_j Y_{j,t+\tau}, \quad w_j \geq 0, \sum_j w_j = 1, \quad (7)$$

where weights w_j are determined by the similarity between unit i and each donor, measured on the raw histories or in a learned representation space, so that more similar donors receive larger weights. The time-varying ITE $\hat{\delta}(h_{i,t})$ is then obtained by contrasting the estimated outcomes in two prescribed sequences \bar{a} and \bar{a}' .

Methods adopting the matching adjustment mechanism differ in their choice of similarity measures and in the space in which they operate. For example, T4 [27] introduces a mini-batch matching strategy during training, in which propensity scores calculated from the observed history are used as a similarity measure to identify the neighbors most similar from alternative treatment groups. Identified similar patients augment each training batch with their observed outcomes, effectively adjusting for confounding in the learned encoder-decoder model. In the point treatment setting, SyncTwin [25] constructs a synthetic twin for each patient as a weighted combination of patients in the donor pool. This is accomplished by solving a constrained optimization problem that minimizes the difference between the latent representations of the targeted patient and the synthetic twin. Another approach, SCoUT [26], extends the synthetic control framework with nonlinear modeling. The model trains a spatiotemporal Transformer that takes the observed data of the target unit and the donor pool as input and autoregressively outputs the counterfactual outcome. The attention mechanism implicitly determines each donor's contribution at each time step, but the time-varying and context-dependent weights are learned end-to-end rather than derived from a predefined similarity measure.

Matching-based methods require the variable on which similarity is assessed to be sufficient for confounding adjustment. When operating in the original feature space, the measured covariates must capture all common causes of treatment assignment and outcome. When operating in a learned latent

space, this requirement transfers to the representation such that the encoder preserves all confounding information present in the original covariates. This ensures that at each time point, the matched donors can be conditionally exchanged with the target unit given the conditioning covariates. Thus, their observed outcomes serve as valid surrogates for the target's unobserved potential outcome [56]. As the prediction horizon increases, the dimensionality of the patient history used to match increases, making it increasingly difficult to find donors sufficiently similar in the pool [57]. When no donors can closely approximate the target's trajectory, the estimate is biased.

4.2. G-Formula

The *g-formula* [6] is a non-parametric identification result that adjusts for time-dependent confounding by explicitly modeling the forward evolution of intermediate covariates under a prescribed treatment sequence. Confounding bias is removed by marginalizing over the distribution of these intermediate covariates, producing an unbiased estimate of the expected potential outcome. Under the identifiability assumptions stated in Section 2.4, the expected potential outcome conditional on the observed history $H_t = h_t$ is identified by the *conditional g-formula*:

$$\mathbb{E}[Y_{t+\tau}^{\bar{a}} | H_t = h_t] = \int \cdots \int \underbrace{\mathbb{E}[Y_{t+\tau} | \bar{x}_{t+\tau-1}, \bar{a}_{t+\tau-1}]}_{\text{outcome model}} \prod_{k=t+1}^{t+\tau-1} \underbrace{p(x_k | \bar{x}_{k-1}, \bar{a}_{k-1})}_{\text{covariate model}} dx_{t+1} \cdots dx_{t+\tau-1}, \quad (8)$$

where $\bar{x}_{t+\tau-1} = (x_0, \dots, x_{t+\tau-1})$ and $\bar{a}_{t+\tau-1} = (a_0, \dots, a_{t+\tau-1})$ concatenate the observed history of h_t with future covariates and prescribed treatments, respectively. The formula requires two components: an *outcome model* that predicts the outcome given the full covariate and treatment trajectory ($\bar{x}_{t+\tau-1}$ and $\bar{a}_{t+\tau-1}$), and a *covariate model* that specifies the conditional distribution of each intermediate covariate x_k . The integral marginalizes over the future covariates, resolving the bias from time-dependent confounding. The *g-formula* is operationalized through *g-computation* [19,58]: Starting from the observed history h_t , the fitted covariate model sequentially samples future covariates from $\hat{p}(x_k | \bar{x}_{k-1}, \bar{a}_{k-1})$ under the prescribed treatments, generating a *counterfactual trajectory*, i.e., a simulated sequence of intermediate covariates representing the patient's projected state evolution. The fitted outcome model predicts $Y_{t+\tau}$ at the end of each trajectory, and the expected potential outcome is estimated by averaging the predictions of the independently sampled trajectories through Monte Carlo integration.

Within the *g-formula* framework, methods differ primarily in the architectures used to parameterize the covariate and outcome models. G-Net [24] utilizes an RNN to model time-series dependencies by factorizing the joint distribution of evolving covariates and outcomes into a sequential generation process. By conditioning each component on both historical representations and previously simulated covariate variables within the same time step, G-Net enables the simulation of counterfactual trajectories under dynamic regimes. The final causal estimate is then derived by averaging these simulated outcomes via Monte Carlo integration. Another work [28] adapts G-Net for sepsis by decoupling the process: G-Net simulates covariate evolution, which then serves as input to a standalone Long-Short-Term Memory (LSTM) classifier for the prediction of counterfactual results. G-Transformer [29] proposes the use of Transformers to capture long-range dependencies, where Transformer encoders are used to learn historical representations sequentially. Hess et al. [30] reformulate the *g-formula* as nested conditional expectations via the law of iterated expectations, eliminating the need for Monte Carlo sampling and reducing the estimation variance through backward recursive optimization from the observed outcome at the prediction horizon.

When the prescribed treatment sequence reduces to a single decision at time t with no intermediate time steps (Section 2.1 setting 2), the *g-formula* simplifies to its outcome model component, requiring no covariate simulation. This simplified form is termed *outcome regression*, which directly estimates the conditional expected outcome given the observed history and the assigned treatment [59], i.e., $\mathbb{E}[Y | H_t = h_t, A = a]$. CURE [32] pre-trains a domain-specific foundation model, which maps patient

history and assigned treatments into a joint latent representation, adopting outcome regression for adjustment to produce unbiased estimations.

All of the above methods require sequential positivity, where every treatment sequence of interest has a positive probability of being observed for any patient history with positive support. In practice, certain covariate and treatment combinations may be rare or completely absent in the observational data, and strict positivity may not hold. CF-ODE [31] addresses this by explicitly representing limited data support as epistemic uncertainty in its counterfactual predictions. CF-ODE encodes pre-treatment covariates into a latent state and evolves this state forward under prescribed treatments via Neural Ordinary Differential Equations (ODEs). A variational posterior parameterized as a Neural Stochastic Differential Equation (SDE) captures trajectory-level uncertainty, such that regions of the covariate and treatment space with sparse data support produce high-variance predictions, signaling unreliable counterfactual estimates.

The g-formula demands a correct specification of both the outcome and the covariate models at each time step, as formalized in Equation (8). Misspecifying either of the models results in insufficient confounding adjustment. This dual requirement also shapes the mechanism's characteristic failure mode over long horizons. G-computation accumulates errors through its autoregressive forward simulation of intermediate covariates, where each sampled covariate serves as input for the next step [60]. Furthermore, the use of Monte Carlo sampling to approximate marginalization introduces additional variance in the final estimate. Reformulating the g-formula as nested expectations mitigates the variance from Monte Carlo sampling [30], although the dependence on correct covariate modeling at each step persists.

4.3. Inverse Probability of Treatment Weighting (IPTW)

IPTW adjusts for time-dependent confounding by re-weighting the observational data so that treatment assignment becomes independent of the covariates, creating a pseudo-population in which naive comparisons of outcomes across treatment sequences yield unbiased estimates. Each observation is weighted by the inverse of its probability of receiving the observed treatment given the history, so that patients whose treatment was strongly influenced by confounders contribute proportionally more, counteracting the original confounding pattern. Formally, the *unstabilized weights* are defined as the cumulative inverse probability of the observed treatment sequence given the history at each step:

$$W_{t,\tau} = \prod_{k=t}^{t+\tau-1} \frac{1}{\Pr(A_k = a_k | H_k = h_k)}, \quad (9)$$

where $\Pr(A_k = a_k | H_k = h_k)$ is the *propensity score* at step k . In practice, unstabilized weights can exhibit high variance when certain treatments are rarely observed for a given history. *Stabilized weights* mitigate this by introducing a numerator that conditions only on baseline history and previously prescribed treatments:

$$SW_{t,\tau} = \prod_{k=t}^{t+\tau-1} \frac{\Pr(A_k = a_k | \bar{A}_{[t,k-1]} = \bar{a}_{[t,k-1]}, H_t = h_t)}{\Pr(A_k = a_k | H_k = h_k)}. \quad (10)$$

The denominator adjusts for confounding, while the numerator reduces weight variance without reintroducing bias. By applying these weights to observed outcomes, the weighted average consistently estimates the expected potential outcome under a prescribed treatment sequence.

The IPTW methods differ in how propensity scores are estimated and how weights are incorporated into the learning objective. R-MSN [7] consists of a propensity network that estimates stabilized weights from the patient's history and a prediction network that learns a representation of the history to forecast outcomes over variable-length prediction horizons. The estimated weights are applied to the prediction loss to adjust for confounding. To address the challenges of estimating high-dimensional outcomes, Wu et al. [36] tackle the estimation of high-dimensional outcomes by applying weights

derived from IPTW to the loss of a generative model, which takes a noise vector and the full treatment sequence as input to generate counterfactual outcomes whose distribution matches the true counterfactual distribution. SCIP-Net [37] extends IPTW to continuous time by reformulating the intractable product integrals required for continuous time weighting as finite products and uses separate Neural Controlled Differential Equation (CDE) networks to estimate propensity weights and outcomes.

The aforementioned methods operate under the assumption of sequential exchangeability. A separate line of work relaxes this assumption by learning a latent representation that serves as a proxy for unmeasured confounders. This latent representation and the patient history at each time step are used to estimate propensity scores that account for both observed and inferred confounding, and the resulting inverse probability weights are applied to the loss of outcome prediction. One subset of methods targets multi-cause confounders, which are unmeasured variables that jointly influence multiple concurrent treatments. These methods relax sequential exchangeability to sequential *single* exchangeability, assuming the absence of single-cause confounders [61]. VTDNet [38] employs a Transformer-based variational autoencoder whose latent representation simultaneously serves as a confounder proxy and is used for covariate reconstruction and outcome prediction. LipCDE [34] and LipSCDE [35] further assume that multi-cause confounders can be decomposed into high- and low-frequency components. The low-frequency component that captures slowly varying trends and periodicity is considered to be approximately time-invariant. The high-frequency component captures fast variations and is assumed to be recoverable through observed covariates acting as noisy proxies. The two methods differ in their temporal backbone, with LipCDE using a Neural CDE and LipSCDE extending this with a stochastic CDE. Beyond the multi-cause setting, DSW [33] targets more general unmeasured confounders using an RNN backbone coupled with attention. The inferred representation is jointly trained to predict both treatment assignments and outcomes, which encourages it to capture sufficient confounding information for propensity estimation.

In contrast to the conditional g-formula, which relies on the correct specification of the outcome distribution conditional on the full history, the IPTW method requires an accurate model of treatment assignment given the time-varying covariates. Unfortunately, the complex underlying data generation process and the high dimensionality of time series observational data render the perfect specification of either model challenging. To mitigate the risk of model misspecification and robustly adjust for time-dependent confounders, doubly robust estimators are used [62]. *Doubly robust* methods combine these approaches by integrating both the outcome and the propensity score models into a unified estimator. This ensures that the overall estimate remains valid if *either* the outcome or the propensity score model is correctly specified [19,63]. The *IVW-DR-learner* [39] constructs a doubly robust pseudo-outcome by combining iterative g-computation estimates with IPTW. To mitigate high variance in low overlap settings where propensity scores are extreme, the final regressor minimizes a loss weighted by the inverse variance of the pseudo-outcome, down-weighting unreliable estimates.

IPTW relies on a correctly specified treatment model to yield unbiased estimates of the treatment effect. While outcome prediction models are employed in practice, their role is confined to prediction within the deconfounded pseudo-population created by the weights. In longer prediction horizons, IPTW is directly affected due to multiplicative weights over time steps. This cumulation indicates that a near-zero propensity score at any single step produces an extreme compounded weight [5]. Although stabilized weights have been proposed to reduce weight variability, this vulnerability persists when the relationship between the covariate and treatment is strong [64]. Weight truncation provides an additional mitigation by capping extreme values, but at the cost of introducing bias into the estimate [65]. Doubly robust estimation weakens the specification requirement so that the estimate is unbiased if the outcome or treatment model alone is correctly specified [19,63]. However, this combination also introduces unique failure modes that become increasingly likely in longer horizons. When both components are simultaneously misspecified, the bias can be amplified and yield estimates that are more biased than the g-formula or IPTW alone [66].

4.4. Representation Balancing

Representation balancing adjusts for time-dependent confounding by learning a representation of the patient's history $\Phi(H_k)$ that is statistically independent of treatment assignment A_k [4,8]. If the learned representation retains information about treatment, confounding bias persists in the outcome prediction. The methods reviewed in this section all enforce the same independence objective $\Phi(H_k) \perp A_k$, but differ in the mechanisms used to achieve it: predictive balancing enforces independence implicitly by training the representation to be non-predictive of treatment assignment, and distribution alignment enforces independence explicitly by measuring and minimizing a statistical divergence that quantifies the dependence between the representation and treatment.

4.4.1. Predictive Balancing

Predictive balancing enforces $\Phi(H_k) \perp\!\!\!\perp A_k$ through adversarial learning between the historical representation encoder and a treatment classifier [8]. The treatment classifier minimizes its prediction loss to accurately predict treatment from the learned representation, while the encoder simultaneously maximizes this loss. The objective takes the form of a min-max game:

$$\min_{\theta_h, \theta_y} \max_{\theta_a} \sum_i \mathcal{L}_y^i(\theta_h, \theta_y) - \lambda \mathcal{L}_a^i(\theta_h, \theta_a), \quad (11)$$

where θ_h , θ_y , and θ_a denote the learnable parameters of the history encoder, outcome predictor, and treatment classifier, respectively. \mathcal{L}_y^i and \mathcal{L}_a^i separately denote the outcome and treatment prediction losses for patient i . At equilibrium, the treatment classifier performs no better than random guessing, indicating that $\Phi(H_k)$ retains no information about A_k .

In practice, the adversarial objective is most commonly implemented via a gradient reversal layer (GRL), which leaves the forward pass unchanged but negates the treatment prediction gradient for the encoder during backpropagation. CRN [8] introduces this approach for time-varying ITE estimation, using an LSTM encoder with a GRL-based treatment classifier. CEET [48] replaces the recurrent backbone with a Transformer encoder-decoder and introduces relative position encoding via Toeplitz matrices, motivated by the observation that relative time intervals are more clinically meaningful than absolute time points in longitudinal health data. Both CRN and CEET assume discrete treatments at each time step, but clinical practice often departs from this. Treatments may be administered concurrently, involving individual causal contributions with interaction effects. TCFimt [41] addresses this by combining GRL with contrastive learning to decompose the total effect into individual and synergistic components. Furthermore, treatment dosages are often continuous rather than categorical, preventing the treatment classifier from operating on a finite set of labels. RCAB [43] replaces the GRL with a Generative Adversarial Network (GAN)-based architecture. This is trained with a continuous imbalance loss that aggregates contributions from samples with similar dosage values. At longer horizons, multi-step-ahead prediction entails a combinatorial explosion of possible treatment sequences, destabilizing adversarial training. EDTS [50] trains a GAN where the generator produces counterfactual outcomes and a discriminator distinguishes them from the observed outcomes. By constraining the counterfactual space to first-degree relatives of the factual sequence, it limits the number of counterfactuals that the treatment and dosage discriminator needs to discern.

A separate line of work extends predictive balancing to the continuous-time setting, retaining GRL-based balancing while changing the temporal backbone. TE-CDE [40] replaces the recurrent encoder with a Neural CDE backbone operating under continuous-time sequential randomization, encoding an irregular history into a continuous latent state that is then evolved under hypothetical treatments. To predict future outcomes, the decoder evolves this state conditional on hypothetical future treatments. CF-GODE [42] and CAG-ODE [44] extend this setting to multi-agent dynamic systems where patients co-evolve and interfere with one another over a graph, strengthening the sequential exchangeability assumption to require independence from both treatment assignment and neighbor interference given the history and graph structure. CF-GODE encodes baseline covariates in an initial latent state and

evolves them through a Treatment-Induced GraphODE. By coupling GNNs with ODEs over a fixed interaction graph, it injects neighbor treatment effects into the dynamics. Building on this, CAG-ODE further models the effect of simultaneous treatment and evolving interactions using a spatio-temporal GNN encoder. The encoder initializes a coupled GraphODE that jointly models the evolution of nodes and dynamic edges.

4.4.2. Distribution Alignment

Distribution alignment quantifies the statistical dependence between the learned representation $\Phi(H_k)$ and the treatment assignment A_k through an explicit divergence measure $D(\cdot, \cdot)$ and minimizes this quantity as a regularization objective during training. When the divergence is driven to zero, the representation achieves statistical independence from treatment assignment, $\Phi(H_k) \perp\!\!\!\perp A_k$. The independence condition can be expressed through two equivalent distributional formulations. The first measures the discrepancy between the joint distribution of the representation and treatment and the product of their marginal distributions:

$$\min_{\theta_h} D\left(\Pr(\Phi(H_k), A_k), \Pr(\Phi(H_k)) \Pr(A_k)\right). \quad (12)$$

When this divergence equals zero, the joint distribution factorizes into the product of its marginals and achieves statistical independence. The second formulation compares the conditional representation distributions across all pairs of treatment values:

$$\min_{\theta_h} \sum_{\substack{a, a' \in \mathcal{A}_k \\ a \neq a'}} D\left(\Pr(\Phi(H_k)|A_k = a), \Pr(\Phi(H_k)|A_k = a')\right). \quad (13)$$

When this discrepancy equals zero for all treatment pairs, the representation distribution no longer varies with treatment assignment, again yielding $\Phi(H_k) \perp\!\!\!\perp A_k$. The equivalence follows from the observation that this independence holds if and only if $\Pr(\Phi(H_k)|A_k = a) = \Pr(\Phi(H_k)) \forall a \in \mathcal{A}_k$, which holds if and only if the conditional distributions are identical across all treatment values. By enforcing independence through either formulation, the outcome predictions derived from $\Phi(H_k)$ no longer suffer confounding bias, enabling unbiased estimation of the time-varying ITE in Equation (3).

As an early instantiation of the distribution alignment framework, the Causal Transformer [9] introduces the Counterfactual Domain Confusion loss, which measures the cross entropy between the treatment distribution predicted from $\Phi(H_k)$ and a uniform reference distribution over treatment categories. The encoder and the treatment prediction head are trained adversarially, where the treatment head minimizes its prediction loss, while the encoder simultaneously minimizes the cross-entropy toward the uniform reference. This forces the predicted treatment distribution to converge to uniformity, making the representation uninformative about the assignment of the treatment. Causal CPC [46], on the other hand, frames balancing as minimization of mutual information $I(\Phi(H_k), A_k)$, directly targeting the divergence in Equation (12). Mutual information is estimated through CLUB [67], parameterized by an encoder that minimizes this estimate and a treatment classifier that maximizes its predictive accuracy, thus establishing an adversarial objective. To prevent the encoder from discarding relevant information and introducing hidden confounders, Causal CPC additionally enforces invertibility of the learned representation by maximizing mutual information between encoded past and future patient histories. ACTIN [45] also targets Equation (12) but estimates the divergence through a generator-discriminator architecture trained adversarially. The encoder serves as the generator, producing historical representations that form two sets of paired samples: real pairs drawn from the joint distribution $\Pr(\Phi(H_k), A_k)$ and synthetic pairs drawn from the product of the marginals $\Pr(\Phi(H_k)) \Pr(A_k)$. The discriminator is trained to distinguish between these two sets, while the generator is updated to deceive the discriminator, driving the joint and marginal product distributions toward coincidence. Methods that instead instantiate Equation (13) compare representation

distributions directly across treatment groups using integral probability metrics [4]. For example, TV-SurvCaus [53] employs the maximum mean discrepancy, and Liu et al. [49] take advantage of the Wasserstein-1 distance to yield theoretically tighter counterfactual prediction error bounds. The latter further introduces Sub-treatment Group Alignment, which posits that treatment groups comprise latent subgroups corresponding to distinct patient profiles (e.g., patient demographics). Rather than enforcing global invariance across entire treatment groups, this approach uses Gaussian mixture models for clustering and aligns the corresponding subgroups individually.

Enforcing full distributional invariance can be too aggressive when the information needed for accurate outcome prediction is also removed. Mamba-CDSP [51] addresses this by minimizing only the cross-covariance between the historical representation and the treatment assignment through a penalty on Mamba’s selective parameters. This removes linear statistical dependence, which for binary treatments is equivalent to matching the first moments of the representations across treatment groups while preserving their variance structure. Under Gaussian covariate and linear model assumptions, the authors prove that this yields tighter error bounds than adversarial balancing.

The independence objective shared by the aforementioned methods for representation balancing introduces an inherent tension between confounding removal and outcome prediction. Because confounders influence both treatment assignment and the outcome, suppressing treatment-predictive information in the representation to satisfy the independence constraint inevitably removes information that is also relevant for outcome prediction [54]. This distinguishes representation balancing from other adjustment mechanisms that converge to the true estimand if modeling requirements are satisfied. In addition, a valid adjustment using representation balancing requires choosing an appropriate dimension. The low-dimensional representation risks losing confounding information that results in insufficient adjustment [68]. Representation balancing methods face unique compounding challenges over longer prediction horizons. In discrete-time architectures, balanced representations in autoregressive decoders are updated using predicted outcomes and the past balanced representation. In continuous-time architectures, the balanced representation serves as the initial condition for differential equation integration and evolves forward. Outcome prediction errors, any residual imbalance, or information loss in the learned representation could be carried over and amplified in subsequent time steps.

4.5. Representation Disentanglement

Representation disentanglement decomposes the learned representation of patient history into functionally distinct latent representations, each defined by its causal relationship to treatment and outcome. Rather than adjusting the entire representation as in representation balancing, this decomposition enables selective adjustment, where confounding bias is removed by balancing only the representation that jointly influences treatment and outcome, while the remaining representations preserve their full predictive power. The decomposition assigns each representation a role according to its predictive relationship with treatment assignment A_k and outcome $Y_{t+\tau}$. In its most granular form, the representation is partitioned into three representations,

$$\Phi(H_t) \longrightarrow \{\Phi_I(H_t), \Phi_C(H_t), \Phi_P(H_t)\}, \quad (14)$$

where $\Phi_I(H_t)$ denotes the instrumental representation that is only predictive of treatment, $\Phi_C(H_t)$ denotes the confounding representation that is predictive of both treatment and outcome, and $\Phi_P(H_t)$ denotes the prognostic representation that is only predictive of outcome but not of treatment. The number and definition of representations may vary across methods, with some adopting coarser two representation decompositions. Regardless of granularity, the shared principle is that confounding bias resides exclusively in the representation that influences both treatment and outcome. Adjustment applied to this representation alone eliminates confounding while preserving the information in the remaining representations. The time-varying ITE $\delta(h_t)$ in Equation (3) is then estimated by

using only the outcome-relevant representations to predict the outcomes under alternative prescribed treatment sequences.

Among the methods that operationalize representation disentanglement, DCRN [54] adopts the three representation decompositions in Equation (14). To adjust for confounding, propensity scores are derived from the confounding representation and used to construct inverse probability weights in a manner analogous to IPTW (Section 4.3). To ensure that the prognostic component remains free of confounding information, DCRN additionally aligns the representation distributions across treatment groups using the maximum mean discrepancy, connecting to the distributional balance objective in Equation (13). DR-CRN [55] adopts a coarser two-representation decomposition, partitioning patient histories into a treatment-predictive representation and an outcome-predictive representation. Separate RNNs are used to learn each representation. To remove residual confounding, DR-CRN minimizes a mutual information upper bound between the two representations, enforcing statistical independence so that neither representation carries information about the other's prediction target.

Representation disentanglement in the time-varying setting was originally proposed to remediate the over-regularization problem in representation balancing [54]. It requires that the patient's history can be decomposed into components with distinct causal roles, as formalized in Equation (14). Whether this decomposition is identifiable from observational data alone remains an open question [69]. This is a fundamental concern because it questions whether the learned components correspond to the intended causal roles at all. In the case of an extended prediction horizon, representation disentanglement shares similar concerns with representation balancing when predicted outcomes and previous representations are used for representation updates. The aforementioned identifiability issue is further exacerbated if the learned representation becomes incomplete. Furthermore, the causal role assumed by each covariate may vary over time due to the varying state of the patient and the clinical protocol. A variable may be a confounder at one time, but it becomes only predictive of the outcome later. Failure to capture this variability over time renders the estimate biased.

5. Clinical Application Scenarios

Estimation of time-varying ITE in healthcare is critical for decision-making, as therapeutic strategies must adapt dynamically to the evolving physiological state of a patient. Unlike static interventions, optimal patient care is often defined by sequential treatment updates, which require precise decisions regarding three key parameters: type of treatment, dosage, and timing. The primary advantage of time-varying ITE estimation models in this context is their ability to handle treatment-covariate feedback loops, where past treatments influence future patient states and subsequently influence future treatment decisions. By simulating unobservable counterfactual trajectories, these models allow clinicians to estimate the specific causal effect of a potential intervention sequence before it is administered. To structure this landscape, we categorize applications by their distinct temporal and systemic demands, ranging from the high-frequency volatility of critical care to the longitudinal disease progression of secondary care, and finally to the interdependent network dynamics of public health.

Critical Care

In the high-stakes environment of the ICU, the primary challenge is the volatility of patient stability and the narrow window for effective intervention. A reliable time-varying ITE model simulates future states under different treatment plans, effectively moving clinical care from reactive management to proactive prevention. For example, in the treatment of sepsis, recent methodologies using critical care datasets [70,71] demonstrate the potential to dynamically refine clinical decisions based on the patient's evolving state. In the context of antibiotic stewardship, these models assist clinicians to determine the optimal treatment and timing, balancing mortality reduction against risks such as toxicity accumulation and drug resistance, while avoiding specific contraindications such as drug-induced kidney injury [28,47]. Similarly, for hemodynamic management, they facilitate granular decisions on fluid resuscitation versus vasopressor use, allowing clinicians to predict adverse events such as pulmonary edema under liberal versus conservative strategies [29,72]. Beyond outcome

prediction, these approaches improve clinical trust by decomposing the influence of specific covariates on treatment response. This interpretability facilitates personalized healthcare by stratifying patients into subgroups of high and low-responders [38], ensuring that aggressive interventions are strictly aimed at those who will derive a causal benefit.

Chronic Care

Moving beyond rapid physiological fluctuations in the ICU, the estimation of time-varying ITE is equally critical for managing the long-term evolutionary progression of chronic diseases. In oncology, tumors develop drug resistance over time, altering the treatment effect at different stages of progression [73]. A static model might correctly identify a chemotherapy agent as effective initially, but it fails to detect the inflection point at which the causal benefit diminishes due to acquired resistance or where cumulative toxicity outweighs efficacy. Time-varying ITE estimation models allow clinicians to determine the optimal treatment stopping time by continuously estimating counterfactual trajectories, thus preventing over-treatment when the effect plateaus. To validate these approaches, studies [7,8,24,74] frequently utilize the pharmacokinetic-pharmacodynamic (PK-PD) synthetic dataset [75] to simulate tumor volume trajectories in response to different sequences of chemotherapy and radiotherapy, demonstrating how iterative estimation can identify schedules that optimally balance tumor reduction with patient tolerance. In neurodegenerative disorders such as Alzheimer's Disease and Friedreich's Ataxia, the outcomes of interest shift from tumor volume to disease onset risk or the trajectory of irreversible functional decline. A critical challenge in this domain is isolating the subtle, time-dependent benefits of therapy from the noisy background of natural progression. Initial studies have applied time-varying ITE estimation to the onset of Alzheimer's disease under concurrent medications [38,76] and to the functional progression in Friedreich's Ataxia under calcitriol use [26,77], producing individualized counterfactual predictions in both settings.

Public Health

The necessity of time-varying ITE estimation scales from the individual physiology of the patient to the collective epidemiology of public health. In this domain, the clinical objective shifts from reversing organ failure or tumor growth to alleviating viral transmission in the community. Just as a physician must adapt drug dosages to the evolving vital signs of a patient, public health officials must continuously adjust the intensity of policies, such as social distancing or vaccination mandates, in response to changing infection rates [42,44,52]. Recent applications [52] demonstrate that these models can estimate the effective number of reproductions under dynamic conditions, proving beneficial not only for the retrospective evaluation of policy effectiveness, but also for prospective planning to recommend optimal strategies for emerging pandemics. Furthermore, these approaches address the critical challenge of network interference, where interventions are not isolated to a single unit but spread to affect connected populations [42,44]. For example, mobility restrictions or vaccination mandates in one region causally mitigate transmission risks in neighboring areas. These studies demonstrate that time-varying ITE estimation can account for spatial interference and estimate the impact of policy interventions that vary in time between interconnected populations.

6. Evaluation

This section illustrates the necessary components for empirically evaluating time-varying ITE estimation methods, including datasets, evaluation metrics, and the predominantly used baselines for comparison.

Table 3. Clinical datasets used for time varying ITE estimation evaluation. The name, evaluation type (synthetic, semi-synthetic, or real-world), specific clinical domain, examples of covariates, examples of treatments, examples of outcomes, the raw/original dataset sample size, and the range of sample sizes used for experimentation in existing works. N/A stands for not applicable, and “-” represents unspecified.

Dataset	Type	Clinical Domain	Covariate Example	Treatment Example	Outcome Example	Raw [♣] Size	Range [♣]
PK-PD [75]	Syn.	NSCLC	Tumor volume history	Chemo, RT	Tumor volume	N/A	1200–21,000
CVSim [78]	Syn.	Cardiovascular	Hemodynamics	Fluids, Vaso	MAP, PO	N/A	11,702–13,000
CVM [31,79,80]	Syn.	Cardiovascular	Hemodynamics	Fluid intake	ABP	N/A	-
PDM [31,81,82]	Syn.	Immunology	Type I IFNs, T Cells	Dexamethasone	Type I IFNs	N/A	-
MIMIC-II [83]	Real	Nephrology	Creatinine history	Dialysis	Creatinine	~800	231–428
MIMIC-III [70]	Real	ICU	Vitals, Labs	Vaso, MV	DBP	46,467	2,313–11,715
MIMIC-III [70]	Semi	ICU	Vitals, Labs	-	-	46,467	1,000–11,715
MIMIC-IV [71]	Real	ICU [◇]	Vitals, Labs	Fluids	Fluid overload	180,733	3,553–8,934
AUMCdb [84]	Real	Sepsis	Vitals, Labs	Antibiotics	SOFA score	20,109	3,136–6,285
CAMP [85]	Real	Asthma	Lung function metrics	No treatment	PreFF	1,041	275
FA-COMS [77]	Real	FA	Quality of life scales	Calcitriol	FARs, 9-HPT, ADL	410	25
NACC [76]	Real	AD	Vital signs, Biomarkers	Statins, NSAIDs	AD Risk	3,410 [◇]	31,059
EEG [86]	Semi	Alcoholism	Electrode values	-	-	122 [†]	1,000 [†]
COVID-19 Penn [36]	Real	COVID	Deaths, Symptoms	Mask mandate	New cases	3,219 [†]	3,219 [†]
COVID-19 US [44]	Real	COVID	Deaths, Recoveries	Stay-at-home	Confirmed cases	47 [†]	47 [†]
COVID-19 DE [87]	Real	COVID	-	-	Active cases	401 [†]	401 [†]
COVID-19 EU [52]	Real	COVID	COVID tests	Stay-at-home	Effective reproduction	7 [†]	7 [†]
CPRD [88]	Real	CVD	HDL, Creatinine	Statins	LDL	~6.9% [◇]	125,784
MS CCAE [89]	Real	Stroke	Medical history	Aspirin, Rivaroxaban	Stroke	>10 ⁸	2,955,399

[♣] The default units are #simulations for synthetic dataset and #patients for semi-synthetic and real-world datasets.

[◇] Dataset Clarification: The MIMIC-IV dataset contains both ICU admission data and general hospital admissions; the NACC database contains 3,410 patients at its initial launch in 2007, but the data have been accumulating over the years; the EEG dataset contains 122 patients, each with 120 trials; the CPRD dataset has raw size ~6.9% UK population.

[†] Sample Size Unit Clarification: EEG dataset uses #trials; COVID-19 Penn uses #counties; COVID-19 US uses #states; COVID-19 DE uses districts; COVID-19 EU uses countries.

6.1. Datasets

The evaluation of time-varying ITE estimation focuses primarily on the prediction of factual and counterfactual outcomes. Although a factual evaluation is performed by comparing the predicted values with the ground truth, the fundamental problem of causal inference renders the evaluation of counterfactual predictions on real-world datasets impossible. Consequently, the existing literature relies on (semi-)synthetic datasets for counterfactual prediction evaluation and uses real-world datasets for factual prediction evaluation. The most commonly used datasets are summarized in Table 3.

6.1.1. Synthetic Datasets

The PK-PD [7,75] and CVSim [24,78] datasets are synthetic datasets that are primarily used to model biological processes using known differential equations. The equations underlying the PK-PD dataset simulate the treatment of non-small cell lung cancer, where the tumor cell number changes in response to chemotherapy and radiation. This dataset is adapted for time-varying ITE benchmarking by incorporating a biased treatment assignment policy. The covariates consist of tumor volume history and patient subgroups, the treatments are binary applications of fixed-dose chemotherapy and radiotherapy (4 choices), and the outcome is continuous-valued tumor volume.

Crucially, time-dependent confounding is induced by modeling the probability of treatment assignment as a sigmoid function of the patient’s tumor diameter history. This mechanism ensures that more severe patients with larger tumors are more likely to receive treatment, thereby creating a realistic observational bias that is controlled by a tunable parameter. In comparison, CVSim simulates the cardiovascular system as an electrical circuit using a lumped-parameter model. To evaluate time-varying ITE estimation, the simulator was extended with stochastic elements and a treatment

mechanism, which resulted in the S-CVSim dataset. The covariates comprise 18 hemodynamic features, including heart rate, pressures, flows, and volumes, the treatments are continuous dosages of fluids and vasopressors, and the outcomes are a subset of patient covariates such as mean arterial pressure, central venous pressure, and pulmonary edema. Time-dependent confounding is introduced by designing observational regimes. The probability of assigning a treatment is explicitly calculated as a function of covariates.

6.1.2. Semi-Synthetic Datasets

Fully synthetic datasets provide accessible counterfactual ground truths but suffer from a distributional shift, as the dynamics of differential equations fail to capture the complexity of real clinical data. To bridge this gap, the literature utilizes semi-synthetic datasets that combine real-world covariates with simulated causal mechanisms. In these setups, time-varying patient features are extracted from observational datasets (e.g., MIMIC-III [70], FA-COMS [77], EEG [86], and COVID-19 Penn [36]) to preserve the complex structure of patient health variations. To establish a ground truth for evaluation, researchers discard the observed treatments and outcomes, instead overlaying a simulated probabilistic treatment policy (to induce controlled time-dependent confounding) and generating synthetic outcomes via a known response mapping. Using the widely adopted MIMIC-III dataset as an example, we describe the simplified semi-synthetic dataset construction steps below [9,37,51]. The notation $f(\cdot)$, $g(\cdot)$, and $h(\cdot)$ denote arbitrary functions. Indices k and j separately denote the indices of time and treatment choice.

1. **Extract Covariates X_k :** Real-world healthcare time series data are used to ensure that the complexity of the feature space is realistic. Vital signs and demographics of the patient are extracted from MIMIC-III to serve as time-varying and static covariates, respectively.
2. **Simulate Untreated Outcomes Z_k :** The untreated outcome represents the outcome that would be observed if no treatment were administered. It is generated by combining a longitudinal trend $f(k)$ that captures the patient's endogenous progression over time, a covariate-driven component $g(k)$ representing exogenous fluctuations caused by current physiological states, and additive noise ϵ_k .

$$Z_k = \underbrace{f(k)}_{\text{trend}} + g(X_k) + \epsilon_k \quad (15)$$

3. **Simulate Treatment Vector A_k :** We introduce time-dependent confounding by modeling the conditional distribution of each treatment component $A_{k,j}$ based on the patient's history. This creates a biased assignment mechanism in which the intensity or probability of the treatment depends on the severity of the past.

$$\theta_{k,j} = \sigma \left(\underbrace{h(\bar{Y}_{k-1}, X_k)}_{\text{history representation}} \right) \quad (16)$$

where $\sigma(\cdot)$ denotes the logistic sigmoid function, mapping the history representation to the probability interval $(0, 1)$. The realization $a_{k,j}$ is sampled from a distribution specific to the domain A_j . For binary treatments (i.e., $A_j = \{0, 1\}$): $\theta_{k,j}$ represents Bernoulli probabilities and $a_{k,j} \sim \text{Bernoulli}(\theta_{k,j})$.

4. **Calculate the Cumulative Treatment Effect $E(k)$:** Treatments are modeled to have a delayed and decaying impact (e.g., inverse-square decay) rather than merely an immediate influence. The influence is scaled by the assignment probability to correlate treatment efficacy with patient severity.

$$E(k) = \sum_{u \leq k} \lambda(k-u) \cdot s(a_u, \theta_u) \quad (17)$$

where $\lambda(\cdot)$ is a decay function (e.g., inverse-square), and $s(\cdot)$ computes the influence of treatment over time u given the realized treatment vector a_u and the probability vector $\theta_u = (\theta_{u,1}, \dots, \theta_{u,d_u})$.

5. **Simulate Treated Outcomes** Y_k : Combine the untreated outcome with the cumulative treatment effects to obtain the treated outcome.

$$Y_k = Z_k + E(k) \quad (18)$$

6.1.3. Real-World Datasets

Complementing these synthetic and semi-synthetic datasets, real-world datasets are employed to evaluate factual prediction performance—the model’s ability to accurately predict observed outcomes given the observed history. Although these datasets lack counterfactual ground truth, they provide a critical evaluation of the model’s capacity to accurately capture complex clinical patterns. The MIMIC-III database [70] remains the standard for this purpose in critical care. However, recent literature has expanded validation to diverse clinical environments, such as neurodegenerative disorders using NACC [76] and cardiovascular diseases using CPRD [88]. Furthermore, studies focusing on learning the interactions between agents in the system have utilized infectious disease datasets, specifically COVID-19 Penn [36] and COVID-19 DE [87]. A typical formulation of the MIMIC-III dataset for factual outcome evaluation [46,51,61,90] is as follows: vital signs and lab tests serve as time-varying covariates, with patient demographics (e.g., gender, age) as static covariates, treatments are binary applications of vasopressors and mechanical ventilation (4 choices), and the outcome is diastolic blood pressure.

6.2. Evaluation Metrics

Evaluating time-varying ITE models is challenging because counterfactual outcomes are unobservable in real-world datasets. Consequently, evaluation strategies depend on the availability of ground-truth counterfactuals and the component of the causal inference pipeline that is being assessed. All metrics are computed for a predetermined decision time t and a prediction horizon τ . For notational clarity, we suppress these indices in the definitions below. The definitions for the metrics are presented in Table 4.

Table 4. Summary of evaluation metrics for time-varying ITE estimation. For each metric, we present its name or abbreviation, the evaluation target, applicable data type(s), outcome type, and its equation.

Metric	Target	Data Type	Outcome	Equation
RMSE	Outcome prediction	Both	Continuous	$\sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}$ (19)
AUROC	Outcome prediction	Both	Discrete	$\int_0^1 \text{TPR}(\text{FPR}^{-1}(t)) dt$ (20)
AUPRC	Outcome prediction	Both	Discrete	$\int_0^1 \text{Precision}(\text{TPR}^{-1}(t)) dt$ (21)
Accuracy	Treatment prediction	Synthetic	Discrete	$\frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{a}_i^* = a_i^*)$ (22)
PEHE	Effect estimation	Synthetic	Continuous	$\frac{1}{N} \sum_{i=1}^N (\delta(h_i) - \hat{\delta}(h_i))^2$ (23)
IF-PEHE	Effect estimation	Both	Continuous	—
t-SNE	Representation	Both	—	—

6.2.1. Outcome Prediction

The prerequisite for accurate time-varying ITE estimation is the ability to learn a reliable representation of the patient’s history and outcome mechanism. The outcome prediction compares the model’s predicted outcome \hat{y}_i against the ground truth outcome y_i . These metrics are applicable to both real-world datasets for evaluating factual outcomes and (semi-)synthetic datasets for evaluating counterfactual outcomes. For continuous outcomes, such as tumor volume in the PK-PD dataset or diastolic blood pressure in MIMIC-III, the *Root Mean Squared Error (RMSE)* is the standard metric of measurement [7–9,51]. In scenarios where outcomes are discrete or binary, such as the occurrence of stroke or pulmonary edema, existing work relies on threshold-independent metrics such as the *Area Under the Receiver Operating Characteristic Curve (AUROC)* or the *Area Under the Precision-Recall Curve (AUPRC)* to evaluate discriminative capability [28,29,32].

6.2.2. Treatment Prediction

Beyond evaluating the precision of predicted outcomes, it is important to assess whether a model can identify the optimal treatment strategy. *Accuracy* is utilized to evaluate the selection of discrete treatment types and optimal administration timing [8,40,41]. Notably, because counterfactual outcomes are unobserved in real-world scenarios, this metric is evaluated using synthetic datasets where an *oracle* provides ground truth outcomes for all potential treatment sequences. For both treatment type and timing, the optimal choice a_i^* for patient i is the choice that maximizes the patient's health utility (e.g., minimizing tumor volume) at the prediction horizon, as shown in Equation (22). When applied to treatment type, a_i^* is simply the treatment that produces the best outcome from a discrete set of options. When applied to treatment timing, the evaluation treats the temporal dimension as a discrete set of intervention choices. The model predicts the outcome that would result from initiating treatment at each possible delay $\delta \in \{0, \dots, \tau - 1\}$ within the prediction horizon, and the predicted optimal timing a_i^* is the start time yielding the best final outcome. In both cases, accuracy measures the frequency with which the model's selection coincides with the oracle-optimal choice. Achieving high accuracy indicates the model's capacity to translate predictive estimates into actionable clinical recommendations.

6.2.3. Treatment Effect Estimation

To validate that the model has successfully disentangled the causal effect from confounding bias, *Precision in Estimation of Heterogeneous Effect (PEHE)* [91] is used to measure the MSE between the ground truth and estimated ITEs [27,31,33]. This metric requires oracle counterfactuals and is therefore restricted to synthetic and semi-synthetic datasets. Variations of the PEHE metric have also been developed to address specific reporting needs and data limitations. For example, IF-PEHE [32,92] is introduced to calculate the error between the estimated treatment effects and the approximated true effects derived via influence functions that *do not* require oracle information.

6.2.4. Representation Evaluation

The prevalence of leverage representation balancing for confounder adjustment necessitated the assessment of the quality of balancing. *t-SNE* [93] is frequently used to visualize the latent space [8,42,44,45]. In a successfully balanced representation, the treated and control samples should overlap significantly (indicating de-confounding), rather than forming distinct, separable clusters based on treatment assignment.

6.3. Baselines

To rigorously assess the contribution of novel time-varying ITE methods, it is vital to benchmark their performance against established methods. The predominantly used baselines can be categorized into classical statistical methods, discrete-time deep learning models, and continuous-time deep learning models.

Classical Statistical Baseline

MSM [5] is the preferred classical statistical baseline in the literature. Using pooled logistic regression, MSM estimates the time-varying propensity scores to employ IPTW for adjustment. Despite their linearity, MSMs remain the requisite statistical baseline. It is worth mentioning that while linear regression has been used as a baseline for confounded performance, the need for adjustment for counterfactual outcome prediction has been repeatedly validated [7,8,33].

Discrete-time Deep Learning Baseline

For discrete-time deep learning comparisons, existing work predominantly utilizes R-MSN [7], G-Net [24], CRN [8], and CT [9] for representing distinct approaches to de-confounding: R-MSN implements RNN-based IPTW adjustment; G-Net employs the g-formula with covariate evolution modeling; and CRN and CT perform representation balancing adjustment by establishing an adversarial game.

Continuous-Time Baselines

A critical distinction in modern evaluation is the handling of irregular sampling and continuous dynamics. Consequently, TE-CDE [40] has emerged as the standard baseline for continuous-time ITE estimation evaluation. Unlike discrete-time models that require data discretization or imputation, TE-CDE operates under the assumption of an underlying continuous latent path. It is, therefore, a necessary comparator to demonstrate that a proposed model correctly handles the irregular nature of real-world clinical events.

Emerging Standards for Evaluation

Although the aforementioned baselines provide a robust foundation for comparative analysis, the rapid evolution of causal architecture requires an expanded scope for rigorous evaluation. Recent literature has introduced diverse modeling strategies that challenge previous benchmarks: IVW-DR-learner [39] adapts meta-learning frameworks to the time-varying ITE estimation domain, and ACTIN [45] advances representation balancing through a generator-discriminator framework. In the continuous-time domain, innovation has moved beyond standard Neural CDEs, with SCIP-Net [37] coupling IPTW with Neural CDEs to strictly enforce de-confounding, and INSITE [94] performing explicit ODE discovery for interpretable estimation.

7. Current Challenges

Despite significant methodological advances in time-varying ITE estimation, the adoption of these approaches from theoretical formulation to clinical deployment is hindered by substantial difficulties. This section critically examines the primary challenges facing the field. We first analyze the complexities introduced by the intrinsic properties of observational data, specifically irregularity and continuous-time dynamics. Secondly, we discuss the crisis of evaluation—the inability to validate models against ground truth and the subsequent reliance on imperfect simulations.

7.1. Challenges From Data

The practical application of time-varying ITE estimation is further constrained by the properties of healthcare time series data. Such data are recorded at irregular intervals driven by clinical workflows rather than a fixed sampling schedule, exhibit high rates of missingness, and reflect continuously evolving physiological processes [95]. These properties introduce challenges that extend beyond standard missing data handling, as they interact directly with the causal assumptions underlying ITE estimation.

7.1.1. Continuous Time Dynamics

Physiological processes and pharmacological effects evolve continuously, but standard ITE frameworks typically operate in discrete time. When continuous observations are aggregated into fixed-width time bins, the temporal ordering of treatment and response within a bin can be lost. For example, in an ICU setting, a vasopressor administered at 10:15 am may cause an increase in blood pressure at 10:20 am. If both events fall within the same hourly bin, cause and effect appear simultaneously, and the causal direction becomes unidentifiable from the binned data. More generally, coarser discretization attenuates the precise temporal signal on which valid treatment effect estimation depends.

To circumvent the limitations of discretization, recent frameworks leverage Neural ODEs/CDEs/SDEs to model the evolution of the patient state continuously [40,42,44]. By defining the hidden state dynamics via a parameterized derivative, these methods naturally accommodate irregular time intervals and allow for the estimation of treatment effects at arbitrary points in time. However, standard formulations of these models impose Lipschitz continuity on the learned dynamics to guarantee numerical stability and solution uniqueness [96]. This smoothness constraint may limit the capacity to capture stiff physiological systems characterized by extended periods of relative stability punctuated by rapid transitions, such as the precipitous hemodynamic deterioration observed in septic

shock [97,98]. Hence, developing continuous-time architectures that can accommodate such abrupt dynamics without sacrificing stability remains an open challenge.

7.1.2. Irregularities

Observational healthcare data are characterized by high sparsity and irregularity. However, the challenge is not merely technical (i.e., missing data handling) but fundamentally causal: the missingness mechanism is informative, which violates the standard Missing At Random assumptions [99,100]. In clinical practice, the sampling mechanism, the decision to observe a variable, is often correlated with the underlying state of the patient. A patient in critical condition will have clinical measurements recorded frequently, while a stable patient will be monitored sparsely. Consequently, the intensity of the sampling manifests itself as a proxy of the hidden confounder that simultaneously affects both the treatment and the outcome [101].

Existing work typically employs the Last Observation Carried Forward (LOCF) imputation to convert irregularly sampled observations into regularly spaced sequences for discrete-time architectures [9,37]. However, by deterministically copying previous values forward, LOCF conflates the lack of measurement with clinical stability, making uniformly valued segments ambiguous between true stability and prolonged non-observation. This attenuates the sampling intensity signal and simultaneously induces a structured, time-varying measurement error, the magnitude is correlated with the latent patient state and past treatments. Hence, the proxy signal that could be leveraged to learn hidden time-dependent confounding is diminished while introducing additional biases to the data. Initial work has begun to address this gap. TESAR-CDE [102] formalizes informative sampling as a covariate shift problem in the context of treatment outcome forecasting. The proposed method leverages inverse intensity-weighting, analogous to IPTW for treatment assignment, to correct for the selection bias in observational data. By explicitly modeling when observations occur and why, such observation-process-aware frameworks also offer the potential to move beyond naive imputation to account for informative missingness.

7.2. Evaluation Challenges

The evaluation of time-varying ITE frameworks remains inherently challenging. First, only the factual outcomes are accessible, while the counterfactual outcomes are inherently unobservable, eliminating the direct ground truth for individual-level effect evaluation. Second, limited reproducibility, which comes from incomplete public access to code and preprocessing pipelines—makes it difficult to independently verify results and compare methods consistently across studies.

7.2.1. Unobservable Counterfactuals

The primary challenge in evaluation is the reliance on the proxy ground truth. As explained in Section 6.1, the standard approach utilizes synthetic or semi-synthetic datasets, where counterfactual outcomes are generated via known functions. These synthetic ground truths enable ITE estimation error metrics (e.g., PEHE) to be computed and allow sensitivity analysis to be conducted across varying confounding strengths and degrees of overlap. In addition, because the simulator that generates the data is fully specified, multiple independent realizations can be drawn for repeated evaluation, supporting experimental reproducibility. However, synthetic evaluation is an imperfect substitute for clinical validity [103], because simulators are often domain-specific and depend on compartmental simplifications and structural rigidity [75,78]. This fundamentally misrepresents the stochasticity and complexity of real physiological processes, where multi-system interactions, drug resistance, and additional unknown factors defy closed-form representations. Furthermore, these simulations fail to capture informative sampling and irregular missingness in healthcare time series data. As a result, models that excel on these benchmarks may simply be overfitting to the inductive biases of differential equations rather than learning robust causal representations [10,104]. A natural path toward strengthening evaluation is to improve the fidelity of synthetic and semi-synthetic benchmarks that provide oracle access to counterfactual outcomes. In the static treatment effect estimation literature,

generative modeling has emerged as a promising approach to producing more realistic evaluation benchmarks. Parikh et al. [104] generate realistic synthetic data using variational autoencoders trained on real data but embedded with user-specified treatment effects and confounding bias, while Neal et al. [105] generate realistic semi-synthetic data by fitting generative models to approximate the conditional distributions of real treatments and outcomes. Adapting these evaluations to the time-varying setting, where temporal structure, time-dependent confounding, and informative sampling introduce additional dimensions of complexity, represents a concrete direction for future work.

7.2.2. Reproducibility

Advancements in digital healthcare are highly dependent on the transparency and reproducibility of reported methods. In time-varying ITE estimation specifically, a meaningful proportion of published methods summarize their implementation but are not accompanied by end-to-end code releases. Among the works that release code, it is common for reproducibility to be limited to synthetic datasets, while the procedures for conducting real-world experiments are not fully captured. As time-varying ITE estimation depends on subtle identifiability and data-generating assumptions, offering a complete and reproducible implementation is equally important to developing the underlying theoretical foundation and a prerequisite for reliable scientific accumulation.

8. Future Directions

Following the summarization and analysis of the boundaries of existing research, we propose three critical frontiers for future investigation: the integration of survival analysis for time-to-event outcomes, the leveraging of multi-modal data to mitigate hidden confounding, and the development of digital twins for robust evaluation.

8.1. Time-to-Event Outcomes in Time-varying ITE

Current time-varying ITE research predominantly focuses on continuous scalar outcomes (e.g., blood pressure, tumor size) or binary classifications (e.g., mortality at a fixed horizon). However, in many clinical scenarios, the quantity of interest is not just *if* an event occurs, but *when* it occurs. For example, if a binary mortality outcome is used to measure the effectiveness of a drug for aggressive glioblastoma, all patients may eventually die within that window, and the drug will be deemed ineffective. However, the time-to-event analysis reveals the true clinical benefit by demonstrating that the drug extends the median survival time, capturing the crucial value of delayed progression. This requires the integration of survival analysis, which models the time duration until an event of interest occurs.

Unlike the standard regression setup, time-to-event analysis must account for *censoring*—instances where the observation window ends or a patient is lost to follow-up before the event occurs. This complexity is compounded by *competing risks*, where the occurrence of one event (e.g., death) precludes the occurrence of another (e.g., hospital discharge). Future research must extend counterfactual estimation to model the hazard function that jointly represents the instantaneous rate of occurrence for the competing event under the prescribed treatments. By combining deep survival models with ITE estimation [106], researchers can derive patient-specific survival curves and compute the difference in survival probability or mean survival time under alternative treatment sequences.

8.2. Multi-Modality for Time-Varying ITE

The prevailing dependency on structured tabular data (e.g., vital signs and lab tests) in time-varying ITE estimation methods creates a fundamental validity challenge, notwithstanding their significant strides in modeling temporal dynamics. Since clinical decision-making is intrinsically multi-modal, often driven by medical imaging, unstructured clinical notes, and waveform data, models developed solely with tabular data inevitably fail to capture the complete patient history, resulting in a structural violation of the sequential exchangeability assumption. To address this, future efforts should be devoted to the development of multi-modal representation learning frameworks capable

of mapping heterogeneous data sources into a unified latent space for the ensuing adjustment and treatment effect estimation.

Realizing this potential presents substantial technical challenges beyond multi-modal fusion. From a general modeling perspective, architectures must account for data sparsity and asynchronous temporal alignment, as high-dimensional modalities such as MRI or pathology reports are acquired far less frequently than bedside vitals. In the causal context, additional problems are introduced with respect to a valid representation. Multi-modal data exacerbates the curse of dimensionality, resulting in violations of the sequential positivity assumption [107]. A potential resolution to this tension is to learn a lower-dimensional representation that is sufficient to adjust for confounders while mitigating positivity violations. Furthermore, as discussed in Section 7.1, the presence of informative sampling must be considered to preserve the integrity of the identifiability assumptions. By incorporating the full spectrum of information available to the clinician, multi-modal frameworks simultaneously reduce confounding in data and learn a more comprehensive representation of the patient, thereby minimizing confounding bias and enabling reliable time-varying ITE estimation in complex healthcare settings.

8.3. Digital Twin in Time-Varying ITE

A healthcare digital twin is a patient-specific computational model that integrates mechanistic physiological knowledge with data-driven learning to dynamically simulate an individual's health trajectory over time [108]. The capabilities of a digital twin extend beyond the existing time-varying ITE estimation methods in two aspects. First, existing methods learn input-output mappings $h_t, \bar{a}_{[t, t+\tau-1]} \mapsto \hat{Y}_{t+\tau}^{\bar{a}}$ from data without encoding the mechanistic structure, while a digital twin embeds known physiological relationships that constrain predictions to remain biologically plausible even under treatment sequences not represented in the training data. Second, current models estimate the time-varying ITE by conditioning on observed patient history, but this estimate remains a subgroup-level quantity. A digital twin advances this by serving as a personalized generative simulation engine that begins by capturing global physiological dynamics and treatment responses across a large cohort. It subsequently finetunes to an individual's continuous data stream to update structural priors, effectively narrowing to a patient-specific model.

Realizing such a system for ITE estimation, however, poses several open research challenges. A foundational requirement is the specification of the causal structure. In the time-varying setting, the validity of treatment effect estimates depends on correctly identifying which variables are confounders, instruments, or prognostic factors, as this influences the adjustment strategy. A digital twin, therefore, requires a formal causal graph that encodes these relationships as the structural outline for its simulation engine. This scaffolding has traditionally been elicited from domain experts, a process that is both labor-intensive and difficult to scale. Recent work has shown that large language models can generate plausible causal graphs from variable metadata, offering a potential means to partially automate this specification bottleneck [109].

Beyond causal specification, a digital twin requires a dynamic simulation engine capable of faithfully propagating the patient state forward in continuous time under arbitrary treatment sequences. A key requirement is that generated counterfactual trajectories remain biologically plausible even outside the distribution of the training data, a property that purely data-driven models cannot guarantee. This motivates research into hybrid architectures that integrate mechanistic physiological models with deep learning. In the ITE literature, initial steps toward mechanistic ITE models exist, where ODEs at the population-level have been discovered from observational data and subsequently refined into patient-specific models through individualized parameter calibration [94]. Extending such approaches to richer physiological systems and coupling them with flexible neural components is a non-trivial challenge. Additionally, reliable simulations require the ability to provide calibrated uncertainty estimates that account for both the limitations of the model itself and the inherent variability of the data, offering a principled mechanism for flagging when a simulated treatment effect should not be trusted [110,111]. Finally, digital twins could partially address the evaluation challenge identified in

Section 6 by generating mechanistically grounded counterfactual trajectories. However, the digital twin itself must be validated, and this circularity remains an inherent limitation that must be acknowledged.

9. Conclusions

This literature review serves to bridge the gap between theoretical advancements in time-varying ITE estimation and practical applications in clinical settings. By consolidating the underpinning theoretical foundations, surveying deep learning methodologies, mapping current clinical applications, and analyzing evaluation frameworks, we have outlined the trajectory of the field alongside its key challenges and future research directions. This highlights a fundamental paradigm shift: moving from static, population-level protocols towards dynamic, individualized treatment planning. As deep learning architectures continue to advance in their capacity to model the complex, non-linear relationships inherent in modern healthcare data, they offer growing potential for reliable causal estimation, bringing the field closer to the realization of truly personalized healthcare.

Acknowledgments: Tingting Zhu was supported by the Royal Academy of Engineering under the Research Fellowship scheme.

Appendix A. Notations

Table A1. Summary of notation used in the theoretical framework (Section 2 and 3). Capital letters denote random variables, lowercase letters denote their realizations, and calligraphic symbols denote sets or spaces.

Symbol	Domain	Description
<i>Indices, Constants, and Spaces</i>		
i	$\{1, \dots, N\}$	Patient index
k	\mathbb{N}_0	Discrete time step
t	\mathbb{N}_0	Decision time (current time)
τ	$\mathbb{N}_{\geq 1}$	Estimation horizon (number of future time steps)
N	$\mathbb{N}_{\geq 1}$	Number of patients
T_i	$\mathbb{N}_{\geq 1}$	Number of observed time steps for patient i
d_x, d_a	$\mathbb{N}_{\geq 1}$	Dimensions of covariate and treatment spaces
$\mathcal{X}, \mathcal{A}, \mathcal{Y}$	$\mathbb{R}^{d_x}, \mathbb{R}^{d_a}, \mathbb{R}$ or $\{0, 1\}$	Spaces of covariates, treatments, and outcomes, respectively
\mathcal{H}_t	$\prod_{k=0}^t \mathcal{X}_k \times \prod_{k=0}^{t-1} \mathcal{A}_k$	Joint space of all covariates and treatments observed up to time t
<i>Observed Variables</i>		
X_k, x_k	\mathcal{X}	Covariates at time k (random / realized)
\bar{X}_t, \bar{x}_t	\mathcal{X}^{t+1}	Covariate sequence up to time t : (X_0, \dots, X_t)
A_k, a_k	\mathcal{A}	Treatment at time k (random / realized)
$\bar{A}_{t-1}, \bar{a}_{t-1}$	\mathcal{A}^t	Treatment sequence up to time $t-1$: (A_0, \dots, A_{t-1})
H_t, h_t	\mathcal{H}_t	Observed history of covariates and treatments prior to decision at t
$Y_{t+\tau}, y_{t+\tau}$	\mathcal{Y}	Observed (factual) outcome at time $t + \tau$
\mathcal{D}	—	Observational dataset of N patient trajectories
<i>Potential Outcomes and Causal Estimands</i>		
$\bar{a} = \bar{a}_{[t:t+\tau-1]}$	\mathcal{A}^τ	Prescribed future treatment sequence; \bar{a}' denotes an alternative
$\bar{a}_{[0:t+\tau-1]}$	$\mathcal{A}^{t+\tau}$	Complete treatment sequence concatenating observed past and prescribed future
$Y_{t+\tau}^{\bar{a}}$	\mathcal{Y}	Potential outcome at $t + \tau$ under \bar{a}
$\hat{Y}_{t+\tau}^{\bar{a}}$	\mathcal{Y}	Model-predicted outcome at $t + \tau$ under \bar{a}
$\delta(h_t)$	\mathbb{R}	ITE/CATE estimand: $\mathbb{E}[Y_{t+\tau}^{\bar{a}} - Y_{t+\tau}^{\bar{a}'} \mid H_t = h_t]$
$\hat{\delta}(h_t)$	\mathbb{R}	Estimated ITE: $\hat{Y}_{t+\tau}^{\bar{a}} - \hat{Y}_{t+\tau}^{\bar{a}'}$
$ATE_{t,\tau}(\bar{a}, \bar{a}')$	\mathbb{R}	Population ATE: $\mathbb{E}[Y_{t+\tau}^{\bar{a}} - Y_{t+\tau}^{\bar{a}'}]$

References

1. Marques, L.; Costa, B.; Pereira, M.; Silva, A.; Santos, J.; Saldanha, L.; Silva, I.; Magalhães, P.; Schmidt, S.; Vale, N. Advancing precision medicine: a review of innovative in silico approaches for drug development, clinical pharmacology and personalized healthcare. *Pharmaceutics* **2024**, *16*, 332.
2. Pearl, J. *Causality*, 2 ed.; Cambridge University Press, 2009.
3. Rubin, D.B. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American statistical Association* **2005**, *100*, 322–331.

4. Shalit, U.; Johansson, F.D.; Sontag, D. Estimating individual treatment effect: generalization bounds and algorithms. In Proceedings of the ICML. PMLR, 2017, pp. 3076–3085.
5. Robins, J.M.; Hernan, M.A.; Brumback, B. Marginal structural models and causal inference in epidemiology, 2000.
6. Robins, J. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling* **1986**, *7*, 1393–1512.
7. Lim, B. Forecasting treatment responses over time using recurrent marginal structural networks. *NeurIPS* **2018**, *31*.
8. Bica, I.; Alaa, A.M.; Jordon, J.; Van Der Schaar, M. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. *ICLR* **2020**.
9. Melnychuk, V.; Frauen, D.; Feuerriegel, S. Causal transformer for estimating counterfactual outcomes. In Proceedings of the ICML. PMLR, 2022, pp. 15293–15329.
10. Bica, I.; Alaa, A.M.; Lambert, C.; Van Der Schaar, M. From real-world patient data to individualized treatment effects using machine learning: current and future methods to address underlying challenges. *Clinical Pharmacology & Therapeutics* **2021**, *109*, 87–100.
11. Zhang, Y.; Kreif, N.; Gc, V.S.; Manca, A. Machine learning methods to estimate individualized treatment effects for use in health technology assessment. *Medical Decision Making* **2024**, *44*, 756–769.
12. Ghosheh, G.O.; Gögl, M.; Zhu, T. A perspective on individualized treatment effects estimation from time-series health data. *JAMIA* **2025**, p. ocae323.
13. Munroe, E.S.; Spicer, A.; Castellvi-Font, A.; Zalucky, A.; Dianti, J.; Linck, E.G.; Talisa, V.; Urner, M.; Angus, D.C.; Baedorf-Kassis, E.; et al. Evidence-based personalised medicine in critical care: a framework for quantifying and applying individualised treatment effects in patients who are critically ill. *The Lancet Respiratory Medicine* **2025**, *13*, 556–568.
14. Feuerriegel, S.; Frauen, D.; Melnychuk, V.; Schweisthal, J.; Hess, K.; Curth, A.; Bauer, S.; Kilbertus, N.; Kohane, I.S.; van der Schaar, M. Causal machine learning for predicting treatment outcomes. *Nature Medicine* **2024**, *30*, 958–968.
15. Shi, J.; Norgeot, B. Learning causal effects from observational data in healthcare: a review and summary. *Frontiers in Medicine* **2022**, *9*, 864882.
16. Neyman, J. On the application of probability theory to agricultural experiments. Essay on principles. *Ann. Agricultural Sciences* **1923**, pp. 1–51.
17. Rubin, D.B. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* **1974**, *66*, 688.
18. Imbens, G.W.; Rubin, D.B. *Causal inference in statistics, social, and biomedical sciences*; Cambridge university press, 2015.
19. Hernan, M.; Robins, J. *Causal Inference: What If*; Chapman & Hall/CRC Monographs on Statistics & Applied Probab, CRC Press, 2025.
20. Holland, P.W. Statistics and causal inference. *Journal of the American statistical Association* **1986**, *81*, 945–960.
21. Angrist, J.D.; Pischke, J.S. *Mostly harmless econometrics: An empiricist's companion*; Princeton university press, 2009.
22. Rosenbaum, P.R.; Rubin, D.B. The central role of the propensity score in observational studies for causal effects. *Biometrika* **1983**, *70*, 41–55.
23. Rubin, D.B. Randomization analysis of experimental data: The Fisher randomization test comment. *Journal of the American statistical association* **1980**, *75*, 591–593.
24. Li, R.; Hu, S.; Lu, M.; Utsumi, Y.; Chakraborty, P.; Sow, D.M.; Madan, P.; Li, J.; Ghalwash, M.; Shahn, Z.; et al. G-net: a recurrent network approach to g-computation for counterfactual prediction under a dynamic treatment regime. In Proceedings of the Machine Learning for Health. PMLR, 2021, pp. 282–299.
25. Qian, Z.; Zhang, Y.; Bica, I.; Wood, A.; van der Schaar, M. Synctwin: Treatment effect estimation with longitudinal outcomes. *NeurIPS* **2021**, *34*, 3178–3190.
26. Dedhia, B.; Balasubramanian, R.; Jha, N.K. SCouT: synthetic counterfactuals via spatiotemporal transformers for actionable healthcare. *ACM Transactions on Computing for Healthcare* **2023**, *4*, 1–28.
27. Liu, R.; Hunold, K.M.; Caterino, J.M.; Zhang, P. Estimating treatment effects for time-to-treatment antibiotic stewardship in sepsis. *Nature machine intelligence* **2023**, *5*, 421–431.
28. Su, M.; Hu, S.; Xiong, H.; Kassis, E.B.; Lehman, L.w.H. Counterfactual Sepsis Outcome Prediction Under Dynamic and Time-Varying Treatment Regimes. *AMIA Summits on Translational Science Proceedings* **2024**, *2024*, 285.

29. Xiong, H.; Wu, F.; Deng, L.; Su, M.; Shahn, Z.; Lehman, L.w.H. G-transformer: Counterfactual outcome prediction under dynamic and time-varying treatment regimes. *PMLR* **2024**, 252, <https-proceedings>.
30. Hess, K.; Frauen, D.; Melnychuk, V.; Feuerriegel, S. G-transformer for conditional average potential outcome estimation over time. *arXiv preprint arXiv:2405.21012* **2024**.
31. De Brouwer, E.; Gonzalez, J.; Hyland, S. Predicting the impact of treatments over time with uncertainty aware neural differential equations. In Proceedings of the AISTATS. PMLR, 2022, pp. 4705–4722.
32. Liu, R.; Chen, P.Y.; Zhang, P. CURE: A deep learning framework pre-trained on large-scale patient data for treatment effect estimation. *Patterns* **2024**, 5.
33. Liu, R.; Yin, C.; Zhang, P. Estimating individual treatment effects with time-varying confounders. In Proceedings of the ICDM. IEEE, 2020, pp. 382–391.
34. Cao, D.; Enouen, J.; Wang, Y.; Song, X.; Meng, C.; Niu, H.; Liu, Y. Estimating treatment effects from irregular time series observations with hidden confounders. In Proceedings of the AAI, 2023, Vol. 37, pp. 6897–6905.
35. Cao, D.; Enouen, J.; Liu, Y. Estimating treatment effects in continuous time with hidden confounders. *arXiv preprint arXiv:2302.09446* **2023**.
36. Wu, S.; Zhou, W.; Chen, M.; Zhu, S. Counterfactual generative models for time-varying treatments. In Proceedings of the KDD, 2024, pp. 3402–3413.
37. Hess, K.; Feuerriegel, S. Stabilized neural prediction of potential outcomes in continuous time. *ICLR* **2025**.
38. Dai, H.; Huang, Y.; Liu, Y.; He, X.; Guo, J.; Prospero, M.; Bian, J. Variational temporal deconfounder network for individualized treatment effect estimation with longitudinal observational data. *Journal of Biomedical Informatics* **2025**, p. 104880.
39. Frauen, D.; Hess, K.; Feuerriegel, S. Model-agnostic meta-learners for estimating heterogeneous treatment effects over time. *ICLR* **2025**.
40. Seedat, N.; Imrie, F.; Bellot, A.; Qian, Z.; van der Schaar, M. Continuous-time modeling of counterfactual outcomes using neural controlled differential equations. *ICML* **2022**.
41. Xi, P.; Wang, G.; Hu, Z.; Xiong, Y.; Gong, M.; Huang, W.; Wu, R.; Ding, Y.; Lv, T.; Fan, C.; et al. TCFimt: Temporal Counterfactual Forecasting from Individual Multiple Treatment Perspective. *arXiv preprint arXiv:2212.08890* **2022**.
42. Jiang, S.; Huang, Z.; Luo, X.; Sun, Y. CF-GODE: Continuous-time causal inference for multi-agent dynamical systems. In Proceedings of the KDD, 2023, pp. 997–1009.
43. Zhu, P.; Li, Z.; Ogino, M. Recurrent Continuous Adversarial Balancing for Estimating Individual Treatment Effect Over Time. *International Journal of Computational Intelligence Systems* **2023**, 16, 82.
44. Huang, Z.; Hwang, J.; Zhang, J.; Baik, J.; Zhang, W.; Wodarz, D.; Sun, Y.; Gu, Q.; Wang, W. Causal graph ode: Continuous treatment effect modeling in multi-agent dynamical systems. In Proceedings of the WWW, 2024, pp. 4607–4617.
45. Wang, X.; Lyu, S.; Yang, L.; Zhan, Y.; Chen, H. A dual-module framework for counterfactual estimation over time. In Proceedings of the ICML, 2024.
46. El Bouchattaoui, M.; Tami, M.; Lepetit, B.; Cournède, P.H. Causal contrastive learning for counterfactual regression over time. *NeurIPS* **2024**, 37, 1333–1369.
47. Wendland, P.; Schenkel-Haeger, C.; Wenningmann, I.; Kschischo, M. An optimal antibiotic selection framework for Sepsis patients using Artificial Intelligence. *npj Digital Medicine* **2024**, 7, 343.
48. Zhang, W.; Peng, H.; He, H. CEET: Time-Series Casual Effect Estimation via Transformer Based Representation Learning. In Proceedings of the ISCAIT. IEEE, 2025, pp. 369–373.
49. Liu, Y.; Dong, J.; Fu, C.; Shi, W.; Jiang, Z.; Hua, Z.; Long, B.; Carlson, D. Counterfactual Outcome Estimation in Time Series via Sub-treatment Group Alignment and Random Temporal Masking **2025**.
50. Norimatsu, Y.; Imaizumi, M. Encode-Decoder-based GAN for Estimating Counterfactual Outcomes under Sequential Selection Bias and Combinatorial Explosion. In Proceedings of the CLear, 2025.
51. Wang, H.; Li, H.; Zou, H.; Chi, H.; Lan, L.; Huang, W.; Yang, W. Effective and Efficient Time-Varying Counterfactual Prediction with State-Space Models. In Proceedings of the ICLR, 2025.
52. Guski, J.; Botz, J.; Fröhlich, H. Estimating the causal impact of non-pharmaceutical interventions on COVID-19 spread in seven EU countries via machine learning. *Scientific Reports* **2025**, 15, 9203.
53. Abraich, A. TV-SurvCaus: Dynamic Representation Balancing for Causal Survival Analysis. *arXiv preprint arXiv:2505.01785* **2025**.
54. Berrevoets, J.; Curth, A.; Bica, I.; McKinney, E.; van der Schaar, M. Disentangled counterfactual recurrent networks for treatment effect inference over time. *arXiv preprint arXiv:2112.03811* **2021**.

55. Chu, J.; Zhang, Y.; Huang, F.; Si, L.; Huang, S.; Huang, Z. Disentangled representation for sequential treatment effect estimation. *Computer Methods and Programs in Biomedicine* **2022**, *226*, 107175.
56. Thomas, L.E.; Yang, S.; Wojdyla, D.; Schaubel, D.E. Matching with time-dependent treatments: a review and look forward. *Statistics in medicine* **2020**, *39*, 2350–2370.
57. Abadie, A.; L'hour, J. A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association* **2021**, *116*, 1817–1834.
58. Taubman, S.L.; Robins, J.M.; Mittleman, M.A.; Hernán, M.A. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *International journal of epidemiology* **2009**, *38*, 1599–1611.
59. Imbens, G.W. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics* **2004**, *86*, 4–29.
60. Chen, X.; Hu, L.; Li, F. A flexible Bayesian g-formula for causal survival analyses with time-dependent confounding: X. Chen et al. *Lifetime Data Analysis* **2025**, *31*, 394–421.
61. Bica, I.; Alaa, A.; Van Der Schaar, M. Time series deconfounder: Estimating treatment effects over time in the presence of hidden confounders. In Proceedings of the ICML. PMLR, 2020, pp. 884–895.
62. Bang, H.; Robins, J.M. Doubly robust estimation in missing data and causal inference models. *Biometrics* **2005**, *61*, 962–973.
63. Robins, J.M.; Rotnitzky, A.; Zhao, L.P. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association* **1994**, *89*, 846–866.
64. Xiao, Y.; Moodie, E.E.; Abrahamowicz, M. Comparison of approaches to weight truncation for marginal structural Cox models. *Epidemiologic Methods* **2013**, *2*, 1–20.
65. Léger, M.; Chatton, A.; Le Borgne, F.; Pirracchio, R.; Lasocki, S.; Foucher, Y. Causal inference in case of near-violation of positivity: comparison of methods. *Biometrical Journal* **2022**, *64*, 1389–1403.
66. Testa, L.; Chiaromonte, F.; Roeder, K. Rescuing double robustness: safe estimation under complete misspecification. *arXiv preprint arXiv:2509.22446* **2025**.
67. Cheng, P.; Hao, W.; Dai, S.; Liu, J.; Gan, Z.; Carin, L. Club: A contrastive log-ratio upper bound of mutual information. In Proceedings of the ICML. PMLR, 2020, pp. 1779–1788.
68. Melnychuk, V.; Frauen, D.; Feuerriegel, S. Bounds on representation-induced confounding bias for treatment effect estimation. *ICLR* **2024**.
69. Welch, R.; Zhang, J.; Uhler, C. Identifiability guarantees for causal disentanglement from purely observational data. *NeurIPS* **2024**, *37*, 102796–102821.
70. Johnson, A.E.; Pollard, T.J.; Shen, L.; Lehman, L.W.H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Anthony Celi, L.; Mark, R.G. MIMIC-III, a freely accessible critical care database. *Scientific data* **2016**, *3*, 1–9.
71. Johnson, A.E.; Bulgarelli, L.; Shen, L.; Gayles, A.; Shammout, A.; Horng, S.; Pollard, T.J.; Hao, S.; Moody, B.; Gow, B.; et al. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data* **2023**, *10*, 1.
72. Kuzmanovic, M.; Hatt, T.; Feuerriegel, S. Deconfounding Temporal Autoencoder: estimating treatment effects over time using noisy proxies. In Proceedings of the Machine learning for health. PMLR, 2021, pp. 143–155.
73. Vasan, N.; Baselga, J.; Hyman, D.M. A view on drug resistance in cancer. *Nature* **2019**, *575*, 299–309.
74. Wang, X.; Lyu, S.; Luo, C.; Zhou, X.; Chen, H. Variational counterfactual intervention planning to achieve target outcomes. In Proceedings of the ICML, 2025.
75. Geng, C.; Paganetti, H.; Grassberger, C. Prediction of treatment response for combined chemo-and radiation therapy for non-small cell lung cancer patients using a bio-mathematical model. *Scientific reports* **2017**, *7*, 13542.
76. Beekly, D.L.; Ramos, E.M.; Lee, W.W.; Deitrich, W.D.; Jacka, M.E.; Wu, J.; Hubbard, J.L.; Koepsell, T.D.; Morris, J.C.; Kukull, W.A.; et al. The National Alzheimer's Coordinating Center (NACC) database: the uniform data set. *Alzheimer Disease & Associated Disorders* **2007**, *21*, 249–258.
77. Regner, S.R.; Wilcox, N.S.; Friedman, L.S.; Seyer, L.A.; Schadt, K.A.; Brigatti, K.W.; Perlman, S.; Delatycki, M.; Wilmot, G.R.; Gomez, C.M.; et al. Friedreich ataxia clinical outcome measures: natural history evaluation in 410 participants. *Journal of child neurology* **2012**, *27*, 1152–1158.
78. Heldt, T.; Mukkamala, R.; Moody, G.B.; Mark, R.G. CVSim: an open-source cardiovascular simulator for teaching and research. *The open pacing, electrophysiology & therapy journal* **2010**, *3*, 45.
79. Zenker, S.; Rubin, J.; Clermont, G. From inverse problems in mathematical physiology to quantitative differential diagnoses. *PLoS computational biology* **2007**, *3*, e204.
80. Linial, O.; Ravid, N.; Eytan, D.; Shalit, U. Generative ode modeling with known unknowns. In Proceedings of the Proceedings of the Conference on Health, Inference, and Learning, 2021, pp. 79–94.

81. Dai, W.; Rao, R.; Sher, A.; Tania, N.; Musante, C.J.; Allen, R. A prototype QSP model of the immune response to SARS-CoV-2 for community development. *CPT: pharmacometrics & systems pharmacology* **2021**, *10*, 18–29.
82. Qian, Z.; Zame, W.; Fleuren, L.; Elbers, P.; van der Schaar, M. Integrating expert ODEs into neural ODEs: pharmacology and disease progression. *NeurIPS* **2021**, *34*, 11364–11383.
83. Saeed, M.; Lieu, C.; Raber, G.; Mark, R.G. MIMIC II: a massive temporal ICU patient database to support research in intelligent patient monitoring. In Proceedings of the Computers in cardiology. IEEE, 2002, pp. 641–644.
84. Thorat, P.J.; Peppink, J.M.; Driessen, R.H.; Sijbrands, E.J.; Kompanje, E.J.; Kaplan, L.; Bailey, H.; Kesecioglu, J.; Cecconi, M.; Churpek, M.; et al. Sharing ICU patient data responsibly under the society of critical care medicine/European society of intensive care medicine joint data science collaboration: the Amsterdam university medical centers database (AmsterdamUMCdb) example. *Critical care medicine* **2021**, *49*, e563–e577.
85. Group, C.A.M.P.R.; et al. The childhood asthma management program (CAMP): design, rationale, and methods. *Controlled clinical trials* **1999**, *20*, 91–120.
86. Zhang, X.L.; Begleiter, H.; Porjesz, B.; Wang, W.; Litke, A. Event related potentials during object recognition tasks. *Brain research bulletin* **1995**, *38*, 531–538.
87. Steiger, E.; Mußgnug, T.; Kroll, L.E. Causal analysis of COVID-19 observational data in German districts reveals effects of mobility, awareness, and temperature. *medRxiv* **2020**, pp. 2020–07.
88. Herrett, E.; Gallagher, A.M.; Bhaskaran, K.; Forbes, H.; Mathur, R.; Van Staa, T.; Smeeth, L. Data resource profile: clinical practice research datalink (CPRD). *International journal of epidemiology* **2015**, *44*, 827–836.
89. Merative. MarketScan Research Databases, 2023.
90. Hatt, T.; Feuerriegel, S. Sequential deconfounding for causal inference with unobserved confounders. In Proceedings of the Causal Learning and Reasoning. PMLR, 2024, pp. 934–956.
91. Hill, J.L. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics* **2011**, *20*, 217–240.
92. Alaa, A.; Van Der Schaar, M. Validating causal inference models via influence functions. In Proceedings of the ICML. PMLR, 2019, pp. 191–201.
93. Maaten, L.v.d.; Hinton, G. Visualizing data using t-SNE. *Journal of machine learning research* **2008**, *9*, 2579–2605.
94. Kacprzyk, K.; Holt, S.; Berrevoets, J.; Qian, Z.; van der Schaar, M. ODE discovery for longitudinal heterogeneous treatment effects inference. *ICLR* **2024**.
95. Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; Liu, Y. Recurrent neural networks for multivariate time series with missing values. *Scientific reports* **2018**, *8*, 6085.
96. Chen, R.T.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D.K. Neural ordinary differential equations. *NeurIPS* **2018**, *31*.
97. Kim, S.; Ji, W.; Deng, S.; Ma, Y.; Rackauckas, C. Stiff neural ordinary differential equations. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **2021**, *31*.
98. Singer, M.; Deutschman, C.S.; Seymour, C.W.; Shankar-Hari, M.; Annane, D.; Bauer, M.; Bellomo, R.; Bernard, G.R.; Chiche, J.D.; Coopersmith, C.M.; et al. The third international consensus definitions for sepsis and septic shock (Sepsis-3). *Jama* **2016**, *315*, 801–810.
99. Sisk, R.; Lin, L.; Sperrin, M.; Barrett, J.K.; Tom, B.; Diaz-Ordaz, K.; Peek, N.; Martin, G.P. Informative presence and observation in routine health data: a review of methodology for clinical risk prediction. *JAMIA* **2021**, *28*, 155–166.
100. Rubin, D.B. Inference and missing data. *Biometrika* **1976**, *63*, 581–592.
101. Agniel, D.; Kohane, I.S.; Weber, G.M. Biases in electronic health record data due to processes within the healthcare system: retrospective observational study. *BMJ* **2018**, *361*.
102. Vanderschueren, T.; Curth, A.; Verbeke, W.; Van Der Schaar, M. Accounting for informative sampling when learning to forecast treatment outcomes over time. In Proceedings of the ICML. PMLR, 2023, pp. 34855–34874.
103. Gentzel, A.; Garant, D.; Jensen, D. The case for evaluating causal models using interventional measures and empirical data. *NeurIPS* **2019**, *32*.
104. Parikh, H.; Varjao, C.; Xu, L.; Tchetgen, E.T. Validating causal inference methods. In Proceedings of the ICML. PMLR, 2022, pp. 17346–17358.
105. Neal, B.; Huang, C.W.; Raghupathi, S. Realcause: Realistic causal inference benchmarking. *arXiv preprint arXiv:2011.15007* **2020**.
106. Gögl, M.; Liu, Y.; Yau, C.; Watkinson, P.; Zhu, T. DoseSurv: Predicting Personalized Survival Outcomes under Continuous-Valued Treatments. 2025.

107. D'Amour, A.; Ding, P.; Feller, A.; Lei, L.; Sekhon, J. Overlap in observational studies with high-dimensional covariates. *Journal of Econometrics* **2021**, *221*, 644–654.
108. Corral-Acero, J.; Margara, F.; Marciniak, M.; Rodero, C.; Loncaric, F.; Feng, Y.; Gilbert, A.; Fernandes, J.F.; Bukhari, H.A.; Wajdan, A.; et al. The 'Digital Twin'to enable the vision of precision cardiology. *European heart journal* **2020**, *41*, 4556–4564.
109. Kiciman, E.; Ness, R.; Sharma, A.; Tan, C. Causal reasoning and large language models: Opening a new frontier for causality. *TMLR* **2023**.
110. Deng, L.; Xiong, H.; Wu, F.; Kapoor, S.; Ghosh, S.; Shahn, Z.; Lehman, L.w.H. Uncertainty Quantification for Conditional Treatment Effect Estimation under Dynamic Treatment Regimes. *PMLR* **2024**, *259*, 248.
111. Hess, K.; Melnychuk, V.; Frauen, D.; Feuerriegel, S. Bayesian neural controlled differential equations for treatment effect estimation. *ICLR* **2023**.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.