Estimating counterfactual treatment outcomes over time through adversarially balanced representations

Ioana Bica*,‡, Ahmed M. Alaa†, James Jordon*, & Mihaela van der Schaar+,†,‡

*University of Oxford, ‡The Alan Turing Institute, †UCLA, +University of Cambridge

This work was supported by the Alan Turing Institute, Office of Naval Research and the National Science Foundation.

Introduction

Problem

- Estimate the effects of time-dependent treatments from observational data.
- Need to handle the bias from time-dependent confounders: patient covariates that are affected by past treatments which then influence future treatments and outcomes.

Main Ideas

- Construct treatment invariant representations at each timestep in order to reduce the bias from time-dependent confounders.
- Use domain adversarial training (Ganin et al., 2016) to trade-off between building balancing representations and predicting patient outcomes.

Using counterfactuals for cancer treatment planning:

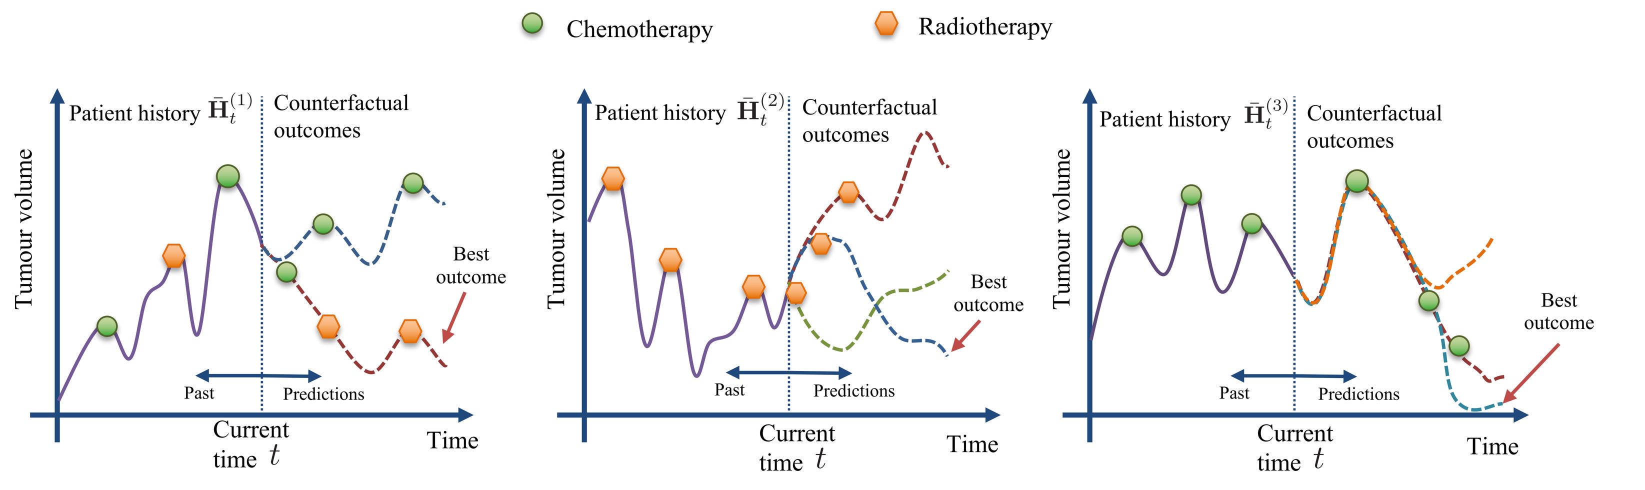

(a) Decide treatment plan (b) Decide optimal time of treatment (c) Decide when to stop treatment

Problem Formulation

Observational dataset

$$\mathcal{D} = \{\{\mathbf{x}_t^{(i)}, \mathbf{a}_t^{(i)}, \mathbf{y}_{t+1}^{(i)}\}_{t=1}^{T^{(i)}} \cup \{\mathbf{v}^{(i)}\}\}_{i=1}^{N}$$

for N independent patients consisting of:

- **Time-dependent covariates** $\mathbf{X}_t \in \mathcal{X}_t$ and **baseline features** $\mathbf{V} \in \mathcal{V}$.
- **Time-dependent treatments** $\mathbf{A}_t \in \{A_1, \dots A_K\}$.
- **Observed outomes** $\mathbf{Y}_{t+1} \in \mathcal{Y}_{t+1}$.

Counterfactual estimation

- Use potential outcomes framework (Neyman, 1923; Rubin, 1978; Robins & Hernán, 2008).
- **Estimate patient outcomes** for any sequence of treatments: $\bar{\mathbf{a}}(t, t+\tau-1) = [\mathbf{a}_t, \dots \mathbf{a}_{t+\tau-1}]$:
$$\mathbb{E}(\mathbf{Y}_{t+\tau}[\bar{\mathbf{a}}(t, t+\tau-1)] | \bar{\mathbf{H}}_t),$$
given patient history $\bar{\mathbf{H}}_t = (\bar{\mathbf{X}}_t, \bar{\mathbf{A}}_{t-1}, \mathbf{V})$.
- **Assumptions:** consistency, positivity, sequential strong ignorability (no hidden confounders).
- History $\bar{\mathbf{H}}_t$ contains **time-varying confounders** $\bar{\mathbf{X}}_t$ that bias the treatment assignment $\mathbf{A}_t$.

Training CRN

Adversarial training procedure for each timestep t:

- **Treatment (domain) loss:**
$$\mathcal{L}_{t,a}^{(i)}(\theta_r, \theta_a) = -\sum_{j=1}^{K} \mathbb{I}_{\{\mathbf{a}_t^{(i)}=a_j\}} \log(G_a^j(\Phi(\bar{\mathbf{H}}_t;\theta_r);\theta_a))$$
- **Outcome loss:**
$$\mathcal{L}_{t,y}^{(i)}(\theta_r, \theta_y) = \|\mathbf{Y}_{t+1}^{(i)} - (G_y(\Phi(\bar{\mathbf{H}}_t;\theta_r),\theta_y))\|^2$$
- **Overall objective:**
$$\mathcal{L}_t^{(i)}(\theta_r, \theta_y, \theta_a) = \sum_{i=1}^{N} \mathcal{L}_{t,y}^{(i)}(\theta_r, \theta_y) - \lambda \mathcal{L}_{t,a}^{(i)}(\theta_r, \theta_a)$$
where λ controls this trade-off between treatment discrimination and outcome prediction.

Related Work

- Use **inverse probability of treatment weighting** (Robins et al., 2000; Lim et al., 2018) as part of the training objective.
- Creates a pseudo-population where the probability of treatment does not depend on the time-varying confounders.
- **Problems**: numerically unstable due to extreme weights, high variance.

Counterfactual Recurrent Network (CRN)

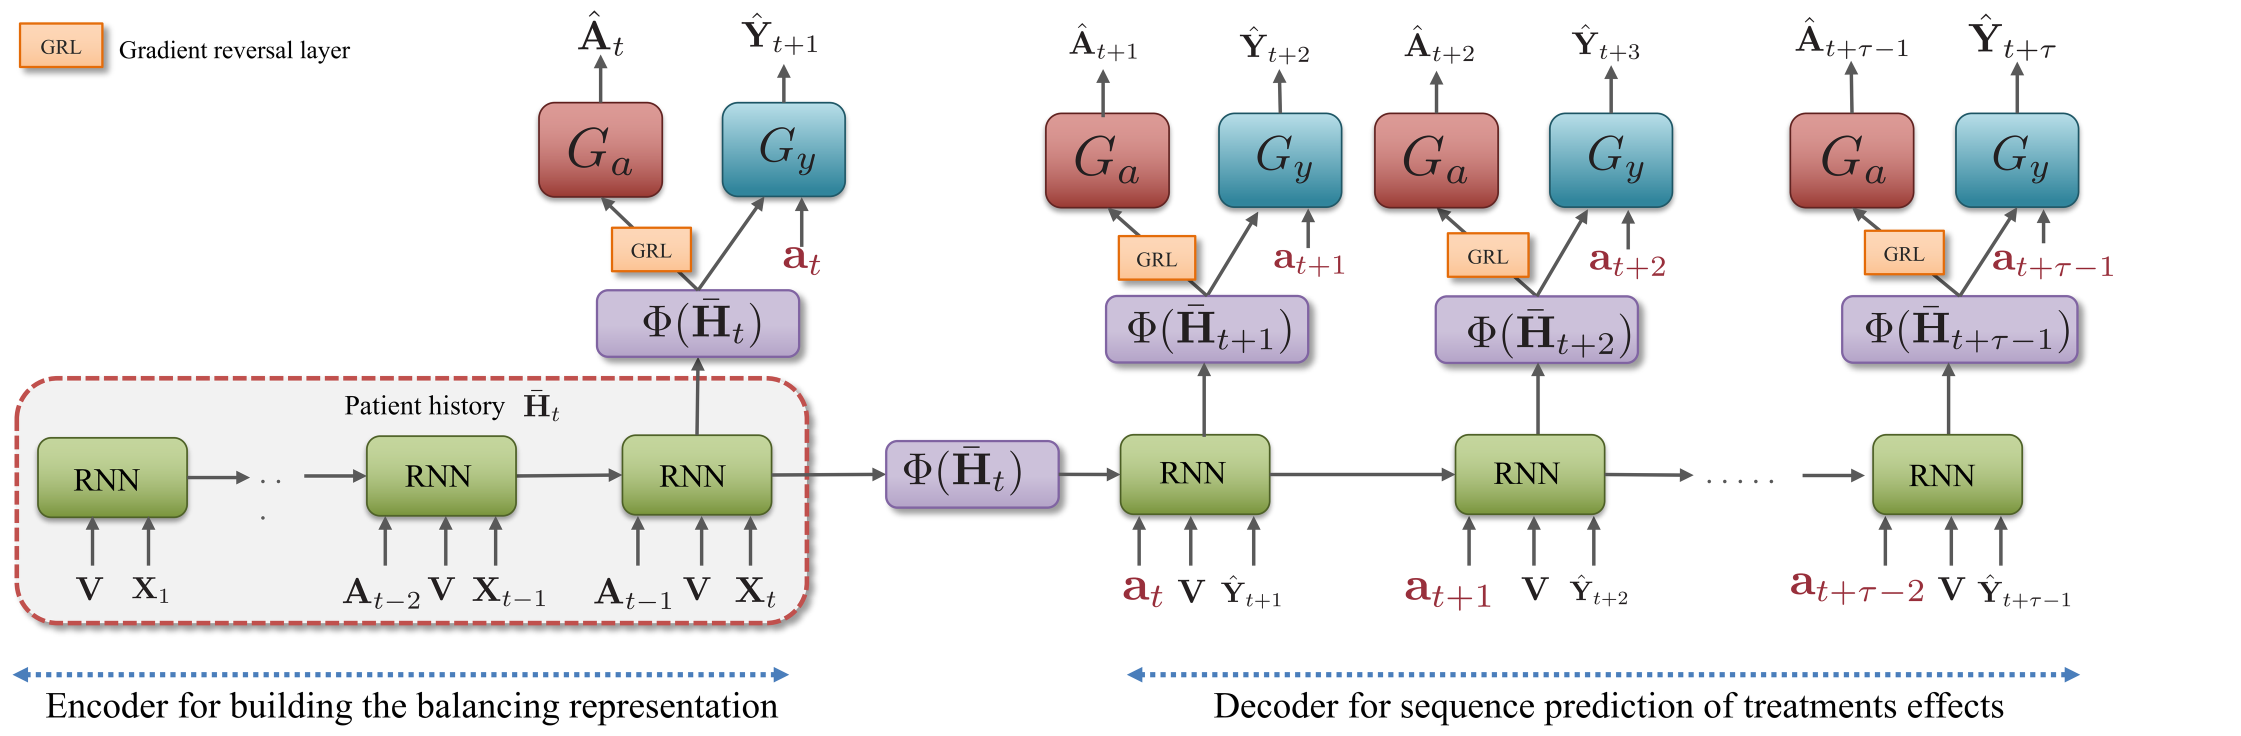

- **Encoder** builds a representation of the patient history $\Phi(\bar{\mathbf{H}}_t)$ that is both invariant to the treatment given at timestep t, $\mathbf{A}_t$, and that achieves low error in estimating the outcome $\mathbf{Y}_{t+1}$.
- **Decoder** updates the balancing representation while estimating counterfactual outcomes for a sequence of future treatments.

Model of Tumour Growth

The volume of tumour $t+1$ days after diagnosis is modelled as follows:

$$V(t+1) = \Big(1 + \underbrace{\rho \log\Big(\frac{K}{V(t)}\Big)}_{\text{Tumor growth}} - \underbrace{\beta_c C(t)}_{\text{Chemotherapy}} - \underbrace{(\alpha_r d(t) + \beta_r d(t)^2)}_{\text{Radiotherapy}} + \underbrace{e_t}_{\text{Noise}}\Big) V(t).$$

Time-varying confounding is introduced by modelling chemotherapy and radiotherapy assignments as Bernoulli random variables, with probabilities p_c and p_r depending on the tumour diameter:

$$p_c(t) = \sigma\Big(\frac{\gamma_c}{D_{\max}}(\bar{D}(t) - \delta_c)\Big) \qquad p_r(t) = \sigma\Big(\frac{\gamma_r}{D_{\max}}(\bar{D}(t) - \delta_r)\Big)$$

The parameters γ_c, γ_r control the amount of time-dependent confounding.

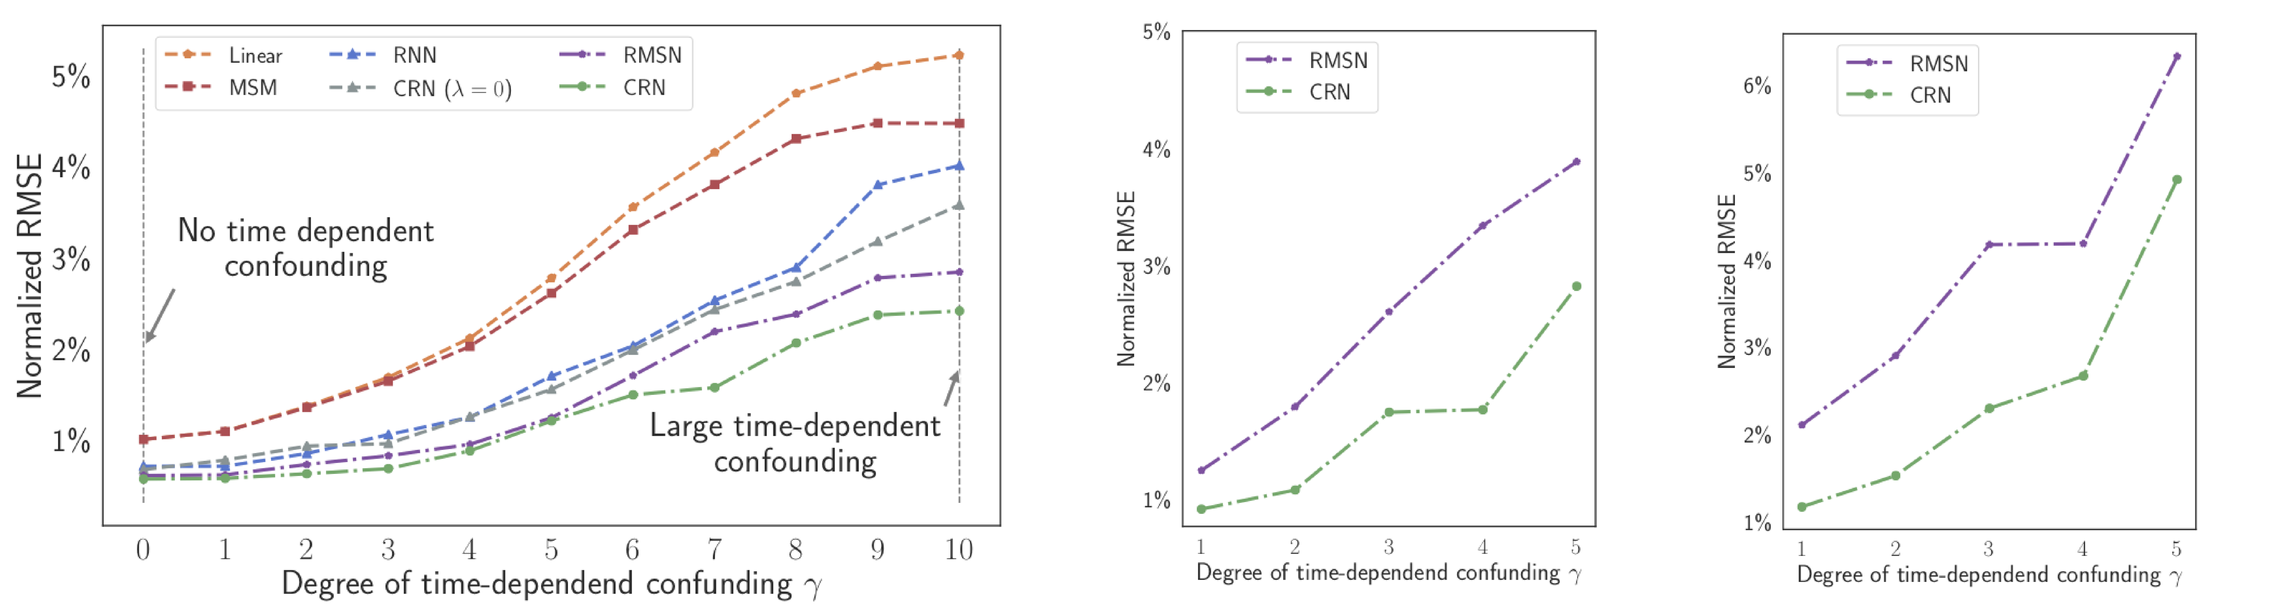

(a) **One**-step ahead prediction. (b) **Four**-step ahead prediction. (c) **Six**-step ahead prediction.

Recommending the Right Treatment and Timing

	τ	$\gamma_c = 5, \gamma_R = 5$ CRN	RMSN	MSM	$\gamma_c = 5, \gamma_R = 0$ CRN	RMSN	MSM	$\gamma_c = 0, \gamma_R = 5$ CRN	RMSN	MSM
Treatment Accuracy	3	**83.1%**	75.3%	73.9%	**83.2%**	78.6%	77.1%	**92.9%**	87.3%	74.9%
	4	**82.5%**	74.1%	68.5%	**81.3%**	77.7%	73.9%	**85.7%**	83.8%	74.1%
	5	**73.5%**	72.7%	63.2%	**78.3%**	77.2%	72.3%	**83.8%**	82.1%	72.8%
	6	**69.4%**	66.7%	62.7%	**79.5%**	76.3%	71.8%	**78.6%**	69.7%	64.5%
Treatment Timing Accuracy	3	**79.6%**	78.1%	67.6%	**80.5%**	76.8%	77.5%	**79.8%**	75.7%	60.6%
	4	**73.9%**	70.3%	63.1%	**79.0%**	77.2%	73.4%	**75.4%**	71.4%	58.2%
	5	**69.8%**	68.6%	62.4%	**78.3%**	73.3%	63.6%	**66.9%**	31.3%	29.5%
	6	**66.9%**	66.2%	62.6%	**73.5%**	72.1%	63.9%	**65.8%**	24.2%	15.5%