SyncTwin: Treatment Effect Estimation with Longitudinal Outcomes

Zhaozhi Qian, Yao Zhang, Ioana Bica,
Angela Wood, Mihaela van der Schaar

van_der_Schaar
\ LAB
vanderschaar-lab.com

UNIVERSITY OF CAMBRIDGE

ITE Estimation – Longitudinal, Irregular, Point Treatment (LIP)

General Problem: Estimating the causal individual treatment effect (ITE) from observational data

Specific Scenario – LIP setting

- Medical observational study with EHR data

- Longitudinal outcomes
- Longitudinal covariates observed irregularly
- One-off treatment allocation

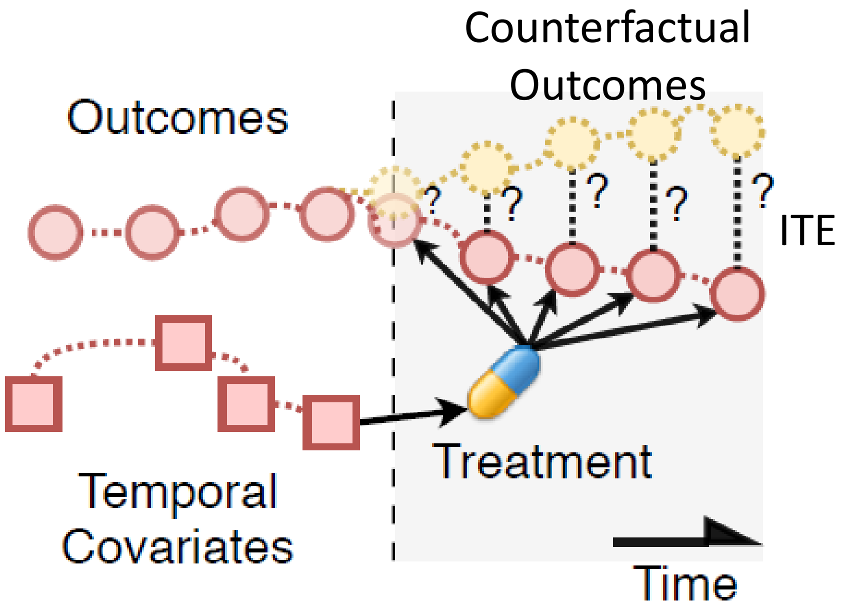

Importance of the LIP setting

1. EHR is longitudinal and irregular in nature

2. The treatment effect on clinical outcomes may be time-varying
   - Long term effect
   - Drug resistance
   - Adverse effect

3. Point treatment is widely applicable
   - Allocation fixed within a treatment regime
   - One off treatments (e.g. transplant)

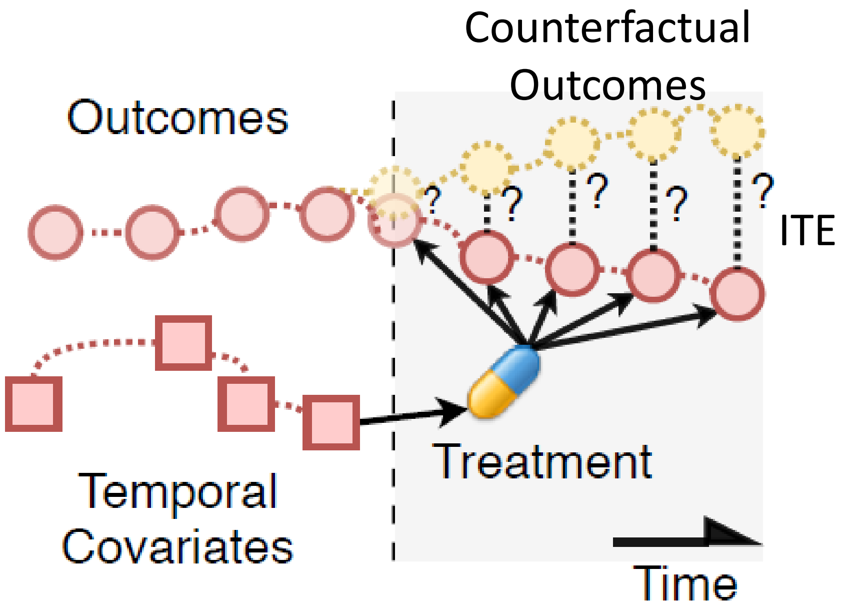

Methods for the LIP setting is underdeveloped

Table 1: **Problem settings considered in the literature**. "Static": observed (or allocated) only once; "Regular": observed (or allocated) over time at a regular frequency; "Irregular": observed over time irregularly; "-": not observed or modeled. * can be extended to Irregular. † can be extended to Regular. LIP: Longitudinal, Irregular, Point treatment.

Setting	Example	Pre-treatment $\mathbf{X}$	Pre-treatment $\mathbf{y}^-$	Treatment a	Post-treatment $\mathbf{y}$	Nonlinear f: $\mathbf{y} = f(\mathbf{X})$
Static	[41]	Static*	-	Static	Static†	✓
DT	[11]	Regular*	-	Regular	Regular	✓
SC	[2]	Regular	Regular	Static	Regular	×
LIP (This work)	This work	Irregular	Regular	Static	Regular	✓

Key insight: leveraging pre-treatment outcomes

How to leverage the pre-treatment outcomes to inform the estimation?

Existing methods

- Ignore pre-treatment outcomes
- Treat them as temporal covariates

SyncTwin

- Explicitly model the outcome time series
- Extension of Synthetic Control

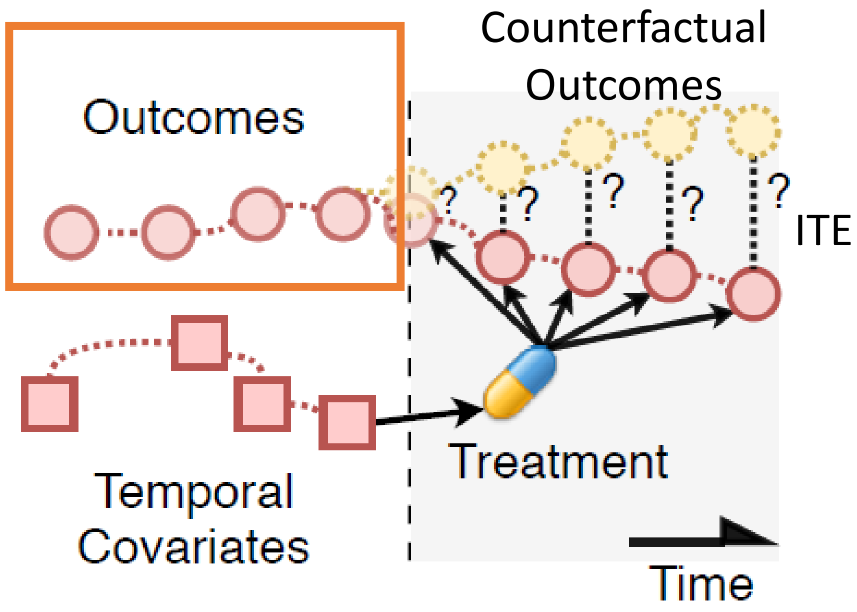

Additional features: (1) individualized error control

1. *Individualized error control*: point out when the models *does not* work for a particular individual

Existing methods

- No individual-level error guarantees

SyncTwin

- Uses pre-treatment outcomes for error control

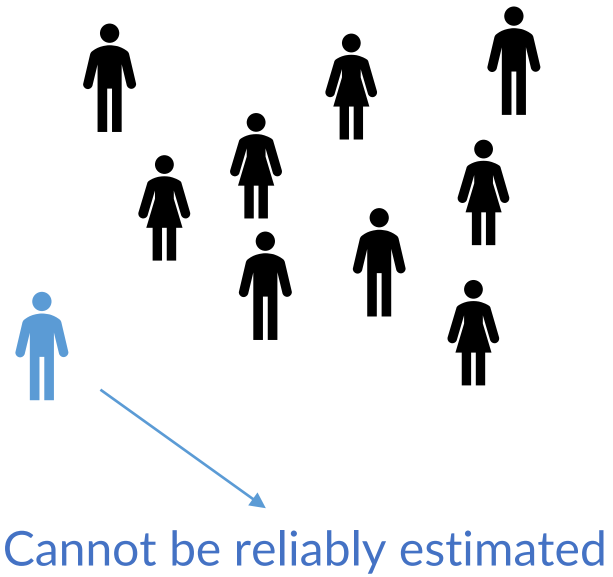

Additional features: (2) explainability by examples

2. *Explainability by examples:* explain the estimation based on a small subset of "contributors"

Existing methods

- Black box models

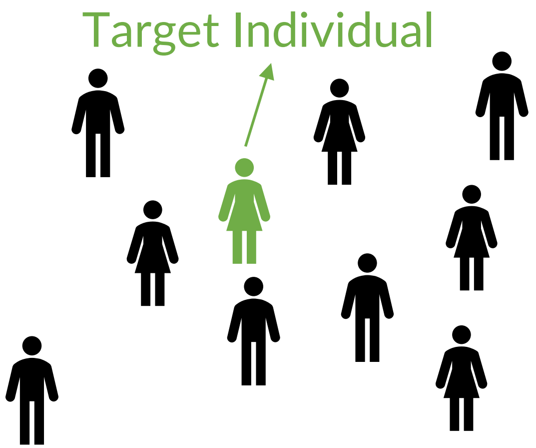

Additional features: (2) explainability by examples

2. *Explainability by examples:* explain the estimation based on a small subset of "contributors"

Existing methods

- Black box models

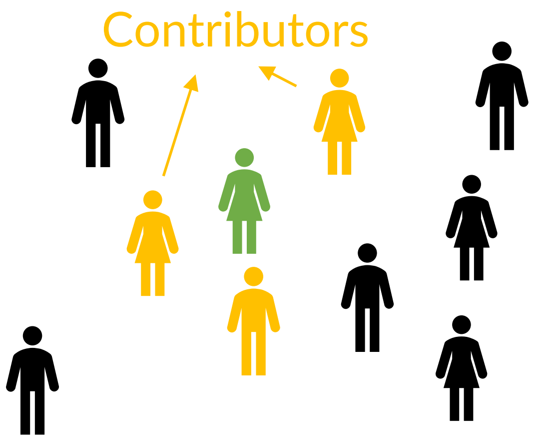

SyncTwin: Encode, Synthetize, Estimate

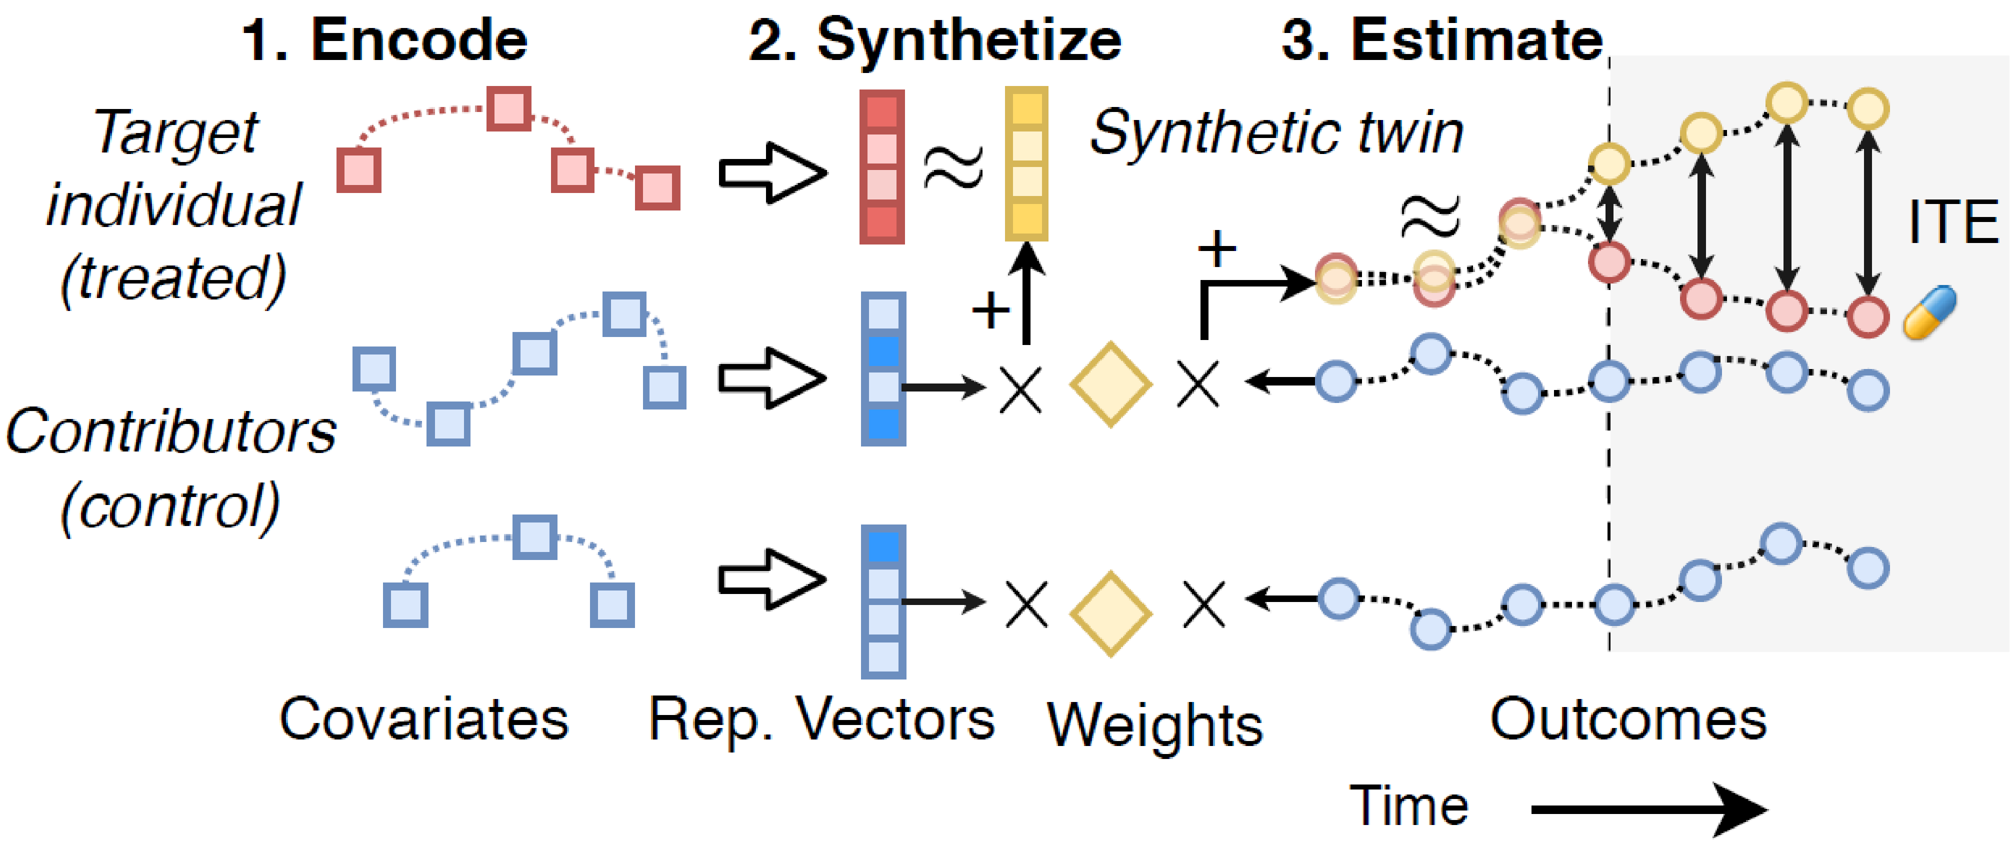

SyncTwin: (1) Encode

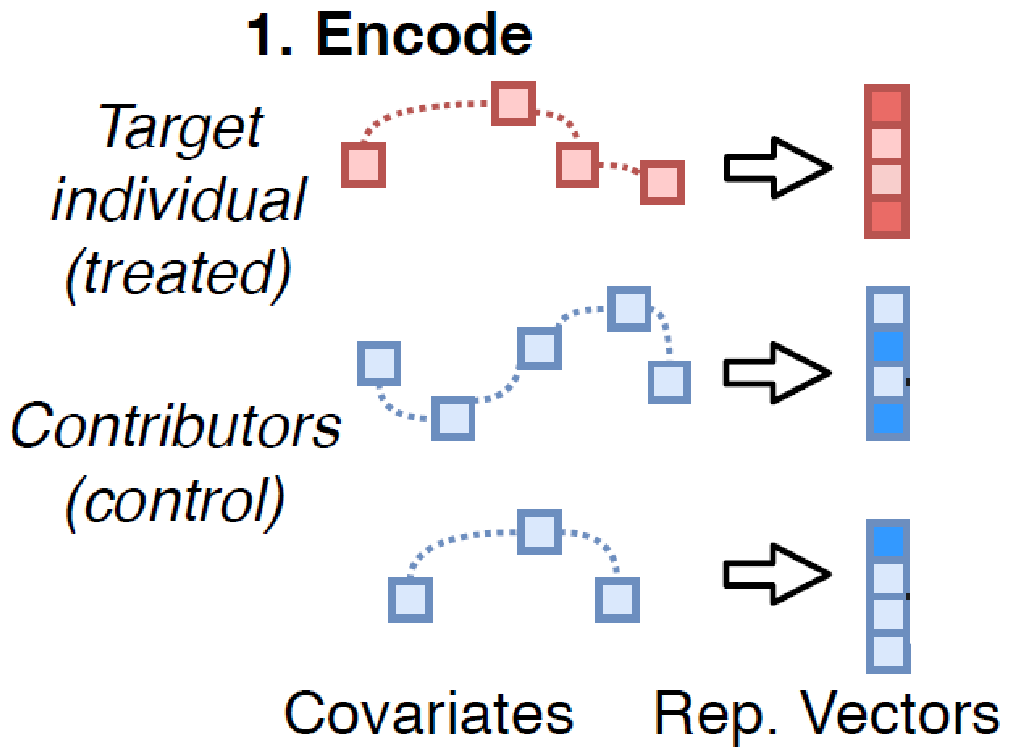

Neural network encoder

- RNN + Attention
- Outcome prediction loss

$$\mathcal{L}_s(\mathcal{D}_0) = \sum_{i \in \mathcal{D}_0} ||\tilde{\mathbf{y}}_i(0) - \mathbf{y}_i(0)||^2$$

- Covariate reconstruction loss

$$\mathcal{L}_r(\mathcal{D}_0, \mathcal{D}_1) = \sum_{i \in \mathcal{D}_0 \cup \mathcal{D}_1} ||(\tilde{\mathbf{X}}_i - \mathbf{X}_i) \odot \mathbf{M}_i||^2$$

SyncTwin: (2) Synthetize

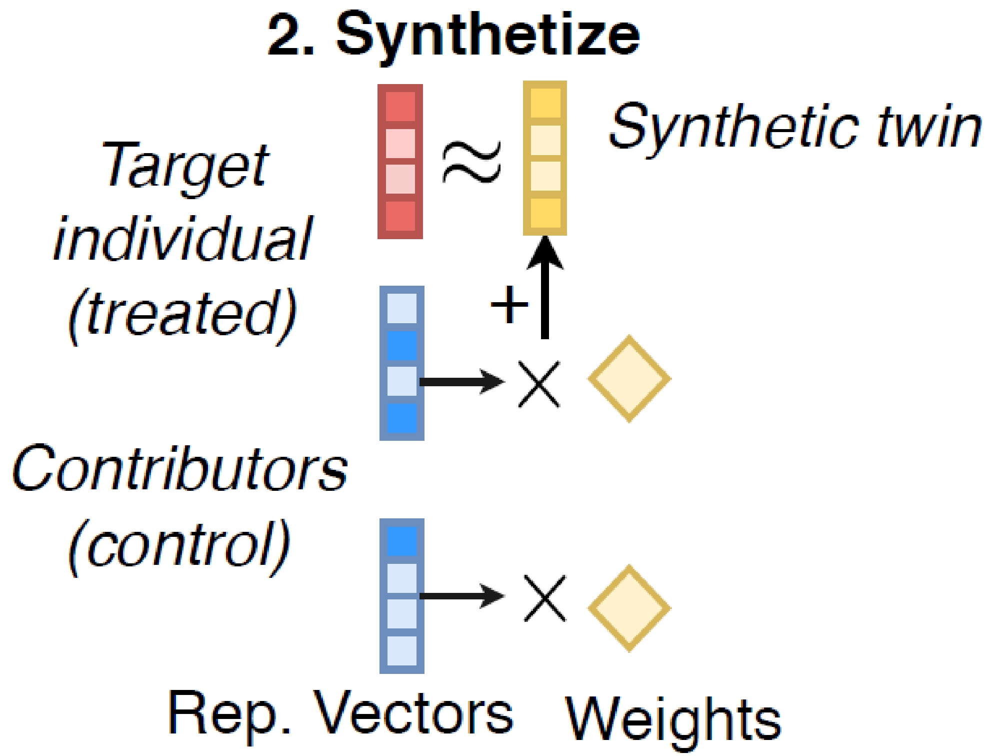

Construct synthetic twin

- Convex combination of contributors in the control group
- Reconstruct target representation

$$\boldsymbol{b}_i = \arg\min_{\tilde{\boldsymbol{b}}_i} \|\tilde{\mathbf{c}}_i - \sum_{j \in \mathcal{I}_0} \tilde{b}_{ij}\tilde{\mathbf{c}}_j\|^2$$

$$\text{s.t. } \tilde{b}_{ij} \geq 0,\ \forall j \in \mathcal{I}_0 \text{ and } \sum_{j \in \mathcal{I}_0} \tilde{b}_{ij} = 1$$

SyncTwin: (3) Estimate

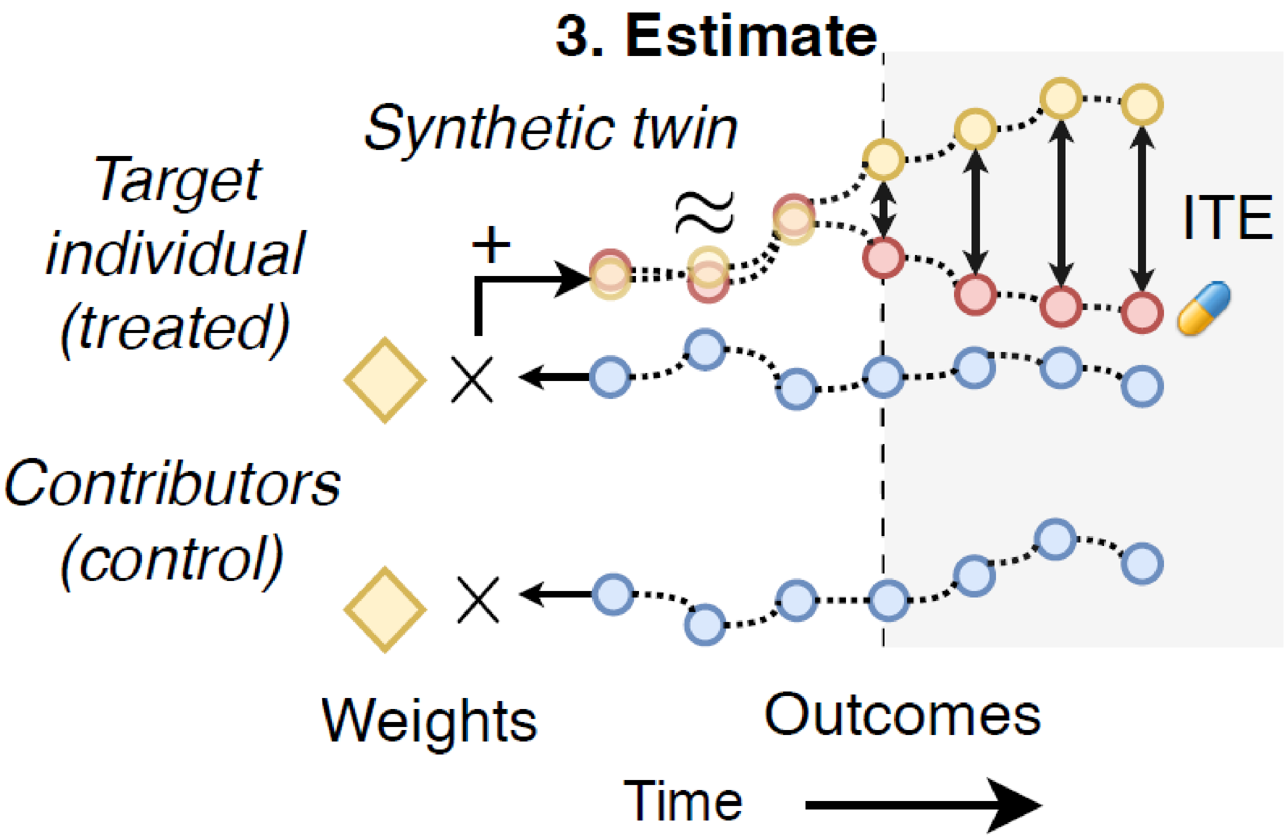

Estimate the potential outcome

- Using the learned weights

$$\hat{y}_{it}(0) = \sum_{j \in \mathcal{I}_0} b_{ij} y_{jt}(0) = \sum_{j \in \mathcal{I}_0} b_{ij} y_{jt}$$

When will SyncTwin work?

Assumptions

- *Consistency*
- *No anticipation*
- *Data generating assumption* (latent factor model)

$$\mathrm{y}_{it}(0) = \mathbf{q}_t^\top \mathbf{c}_i + \xi_{it}$$

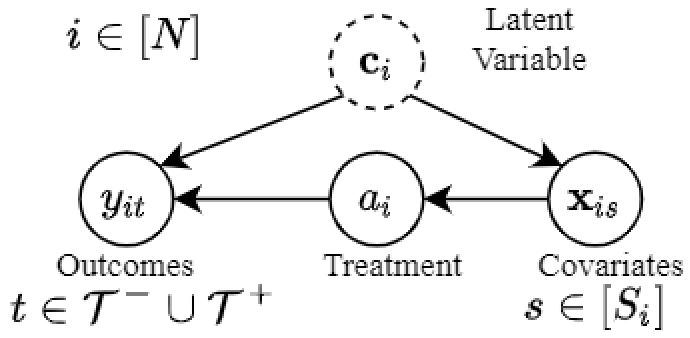

* Data generating assumption is needed to establish the theoretical results. Experiments show SyncTwin works well even if the data is not directly generated from the assumed model.

When will SyncTwin work?

Individualized error control

- *Proposed metric*: error in pre-treatment outcome

$$\mathrm{d}_i^y = \|\hat{\mathbf{y}}_i^-(0) - \mathbf{y}_i^-(0)\|_1$$

Proposition 3 (Error control under no hidden confounders). *Given any target error threshold $\delta > 0$, define the acceptance group of treated individuals as*

$$\mathcal{A}_\delta = \left\{ i \in \mathcal{I}_1 | \mathrm{d}_i^y \leq \delta |\mathcal{T}^-| / |\mathcal{T}^+| \right\}.$$

Under the assumptions in Section 3.1, the post-treatment estimation error $|\mathbb{E}[\hat{\mathbf{y}}_i(0)] - \mathbb{E}[\mathbf{y}_i(0)]| \leq \delta$, $\forall i \in \mathcal{A}_\delta$.

Successfully reproduces a large-scale clinical trial

Prior clinical trial: Heart protection study (HPS) estimates the statins' LDL lowering effect. It reports an average treatment effect of **-1.26** mmol/L (SD=0.06) in the first-year follow-up.

Method: SyncTwin applied to a cohort with matching selection criterion. Observational data from EHR (CPRD)

Findings: SyncTwin estimates the average treatment effect to be **-1.25** mmol/L (SD 0.01). Baseline methods fail to reproduce the findings **(-0.72** mmol/L, SD 0.01)

Reference

Z. Qian, Y. Zhang, I. Bica, A. M. Wood, M. van der Schaar, SyncTwin: Treatment Effect Estimation with Longitudinal Outcomes, Neurips 2021

Code: https://github.com/ZhaozhiQIAN/

Lab website: https://www.vanderschaar-lab.com/

Personal profile: https://www.linkedin.com/in/qianzhaozhi/