Time Series Deconfounder: Estimating Treatment Effects over Time in the Presence of Hidden Confounders

Ioana Bica, Ahmed M. Alaa, Mihaela van der Schaar

International Conference on Machine Learning 2020

Introduction

- **Aim:** Estimate the individualized effects of time-dependent treatments.

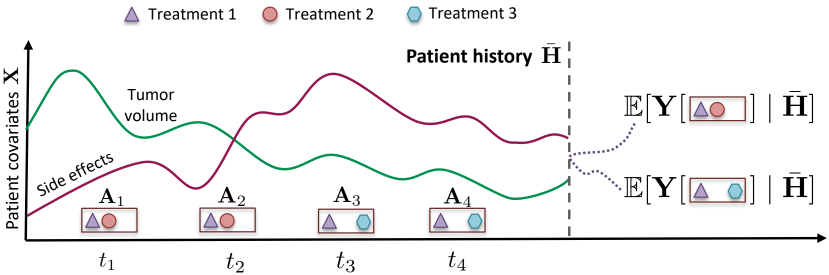

Introduction

- **Aim:** Estimate the individualized effects of time-dependent treatments.

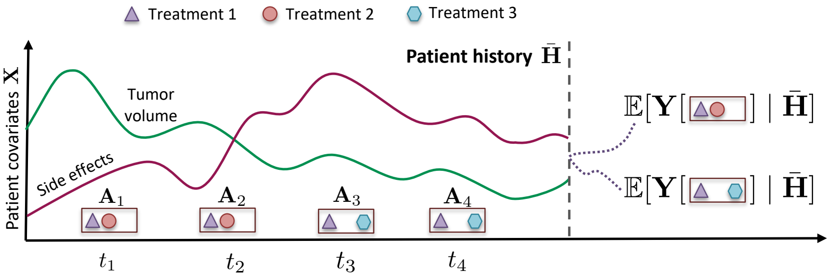

- All existing methods for estimating treatment effects over time assume that there are no hidden confounders.

Hidden confounders

- Hidden confounders introduce bias when estimating treatment effects over time.

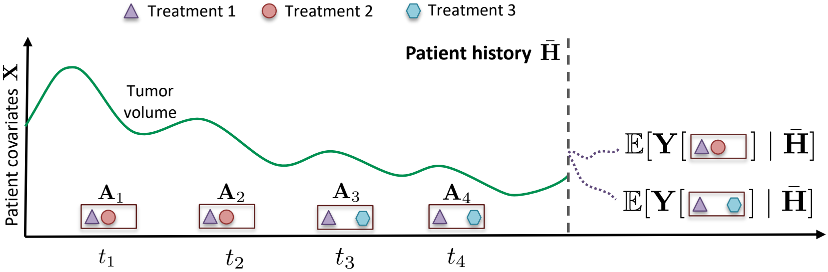

Hidden confounders

- Hidden confounders introduce bias when estimating treatment effects over time.

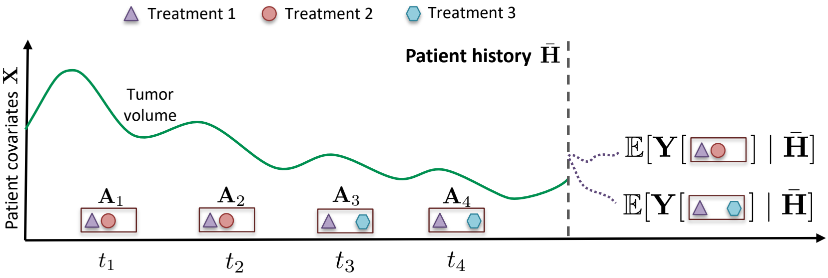

- **Proposed solution:** infer latent variables that capture the dependencies in the treatment assignments over time and can be used as substitutes for the hidden confounders.

Problem formalism

- Observational data for each patient:
  - Time-dependent patient covariates:

$$\bar{\mathbf{X}}_t = (\mathbf{X}_1, \ldots, \mathbf{X}_t)$$

  - Time-dependent treatments:

$$\bar{\mathbf{A}}_t = (\mathbf{A}_1, \ldots, \mathbf{A}_t), \text{ where } \mathbf{A}_t = [A_{t1} \ldots A_{tk}]$$

  - Observed patient outcome given history of covariates $\bar{\mathbf{X}}_t$ and treatments $\bar{\mathbf{A}}_t$: $\mathbf{Y}_{t+1}$.

Potential outcomes

- Use the potential outcomes framework (Rubin (1978), Neyman (1923), Robins & Hernan (2008)).
- Estimate **individualized treatment effects**, i.e. potential outcomes under treatment plan $\bar{\mathbf{a}}_{\geq t}$ conditional on patient history at timestep t:

$$\mathbb{E}[\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t]$$

- Assume consistency and positivity.

Potential outcomes and hidden confounders

- Estimate individualized treatment effects, i.e. potential outcomes under treatment plan $\bar{\mathbf{a}}_{\geq t}$ conditional on patient history at timestep t:

$$\mathbb{E}[\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t]$$

- All existing methods assume that there are **no hidden confounders**:

$$\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \perp\!\!\!\perp \mathbf{A}_t \mid \bar{\mathbf{X}}_t, \bar{\mathbf{A}}_{t-1} \text{ for all } \bar{\mathbf{a}}_{\geq t} \text{ and for all } t,$$

which is **untestable** in practice.

Hidden confounders - from static to temporal setting

- The Blessing of Multiple Causes (Wang & Blei, 2019):
  - Static causal inference setting.
  - Hidden confounders introduce dependencies in the treatment assignments.
  - Infer latent variables that capture these dependencies and render the treatments conditionally independent.

- In the temporal setting, the hidden confounders may change over time and may be affected by past treatments and covariates.

Time Series Deconfounder - Main ideas

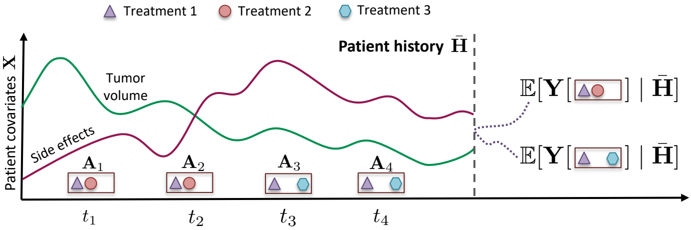

- Hidden confounders may vary over time and may be affected by previous treatments and covariates.

Time Series Deconfounder - Main ideas

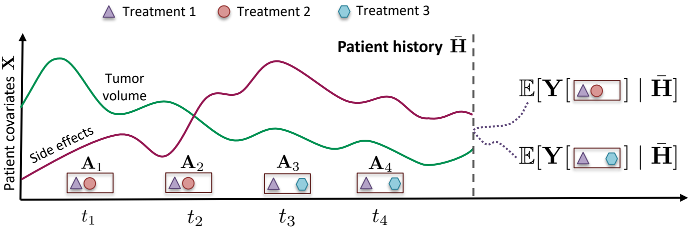

- Take advantage of the way multiple treatments are assigned over time to infer substitutes for the hidden confounders.

$$\bar{\mathbf{Z}}_t = (\mathbf{Z}_1, \ldots, \mathbf{Z}_t)$$

- Augment the observational dataset with $\bar{\mathbf{Z}}_t$ and use an outcome model to obtain unbiased estimates of the treatment effects.

Time Series Deconfounder - Factor model

Step 1: Fit **factor model over time** to infer substitutes for the hidden confounders.

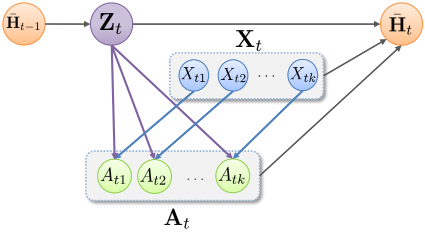

- At time t construct the **latent variable $\mathbf{Z}_t$** as a function of history $\bar{\mathbf{H}}_{t-1} = (\bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_{t-1}, \bar{\mathbf{Z}}_{t-1})$, such that:

$$p(A_{t1}, \ldots, A_{tk} \mid \mathbf{Z}_t, \mathbf{X}_t) = \prod_{j=1}^{k} p(A_{tj} \mid \mathbf{Z}_t, \mathbf{X}_t).$$

Time Series Deconfounder - Factor model

Step 1: Fit **factor model over time** to infer substitutes for the hidden confounders.

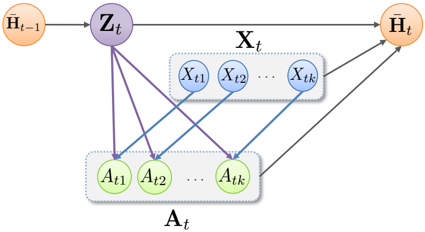

- Factor model of the assigned treatments has joint distribution:

$$p(\theta_{1:k}, \bar{\mathbf{x}}_T, \bar{\mathbf{z}}_T, \bar{\mathbf{a}}_T) = p(\theta_{1:k}) p(\bar{\mathbf{x}}_T) \cdot \prod_{t=1}^{T} \Big(p(\mathbf{z}_t \mid \bar{\mathbf{h}}_{t-1}) \prod_{j=1}^{k} p(a_{tj} \mid \mathbf{z}_t, \mathbf{x}_t, \theta_j) \Big).$$

Time Series Deconfounder - Factor model

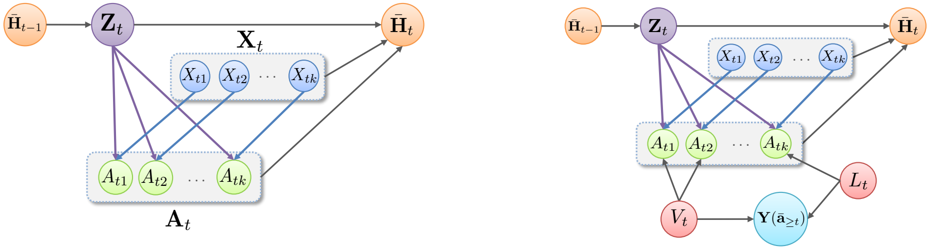

Assumption (Sequential single strong ignorability)

$$\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \perp\!\!\!\perp A_{tj} \mid \mathbf{X}_t, \bar{\mathbf{H}}_{t-1},$$

$\forall \bar{\mathbf{a}}_{\geq t}$ and $\forall t \in \{0, \ldots, T\}$ and $\forall j \in \{1, \ldots, k\}$.

Time Series Deconfounder - Sequential strong ignorability

Theorem

If the distribution of the assigned causes $p(\bar{\mathbf{a}}_T)$ can be written as the factor model $p(\theta_{1:k}, \bar{\mathbf{x}}_T, \bar{\mathbf{z}}_T, \bar{\mathbf{a}}_T)$, we obtain sequential ignorable treatment assignment:

$$\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \perp\!\!\!\perp (A_{t1}, \ldots, A_{tk}) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \bar{\mathbf{Z}}_t,$$

for all $\bar{\mathbf{a}}_{\geq t}$ and for all $t \in \{0, \ldots, T\}$.

Evaluate factor model

- Use predictive checks (Rubin, 1984) to asses how well the factor model captures the distribution of treatments at each timestep.
- The inferred substitutes for the hidden confounders $\mathbf{Z}_t$ also need to satisfy positivity, i.e.

$$P(\mathbf{A}_t = \mathbf{a}_t \mid \bar{\mathbf{A}}_{t-1} = \bar{\mathbf{a}}_{t-1}, \bar{\mathbf{Z}}_t = \bar{\mathbf{z}}_t, \bar{\mathbf{X}}_t = \bar{\mathbf{x}}_t) > 0.$$

Time Series Deconfounder - Outcome model

Step 2: Sample $\hat{\bar{\mathbf{Z}}}_t = [\hat{\mathbf{Z}}_1 \dots \hat{\mathbf{Z}}_t]$ from the factor model and fit an outcome model to estimate:

$$\mathbb{E}[\mathbf{Y} \mid \bar{\mathbf{a}}_{\geq t}, \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \hat{\bar{\mathbf{Z}}}_t] = \mathbb{E}[\mathbf{Y}(\bar{\mathbf{a}}_{\geq t}) \mid \bar{\mathbf{A}}_{t-1}, \bar{\mathbf{X}}_t, \hat{\bar{\mathbf{Z}}}_t].$$

Example outcome models: *Marginal Structural Models* (Robins et al. 2000), *Recurrent Marginal Structural Networks* (Lim et al., 2018).

Proposed factor model architecture

- Proposed architecture for the factor model: recurrent neural network (RNN) with **multitask output** and **variational dropout**.

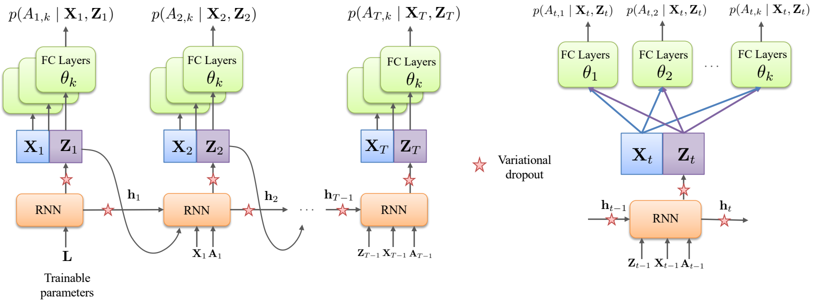

$$\mathbf{Z}_1 = \text{RNN}(\mathbf{L}) \qquad \mathbf{Z}_t = \text{RNN}(\bar{\mathbf{Z}}_{t-1}, \bar{\mathbf{X}}_{t-1}, \bar{\mathbf{A}}_{t-1}, \mathbf{L})$$

$$A_{tj} = \text{FC}(\mathbf{X}_t, \mathbf{Z}_t; \theta_j), \text{ for all } j = 1, \ldots k$$

Experiments on synthetic data

Build synthetic dataset using $p-$order autoregressive processes:

$$X_{t,j} = \frac{1}{p} \sum_{i=1}^{p} (\alpha_{i,j} X_{t-i,j} + \omega_{i,j} A_{t-i,j}) + \eta_t,$$

$$Z_t = \frac{1}{p} \sum_{i=1}^{p} (\beta_i Z_{t-i} + \sum_{j=1}^{k} \lambda_{i,j} A_{t-i,j}) + \epsilon_t,$$

$$\pi_{tj} = \gamma_A \hat{Z}_t + (1 - \gamma_A) \hat{X}_{tj}, \qquad A_{tj} \mid \pi_{tj} \sim \text{Bernoulli}(\sigma(\lambda \pi_{tj})),$$

$$\mathbf{Y}_{t+1} = \gamma_Y Z_{t+1} + (1 - \gamma_Y) \Big(\frac{1}{k} \sum_{j=1}^{k} X_{t+1,j} \Big).$$

Experiments on synthetic data

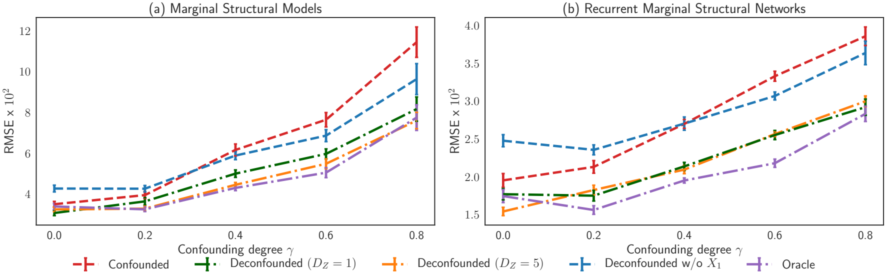

- Root mean squared error (RMSE) obtained for one-step ahead estimation of treatment effects.
- The parameters $\gamma = \gamma_A = \gamma_Y$ control the amount of hidden confounding.

Experiments on MIMIC III

- Dataset with 6256 patients, with 25 covariates (lab tests and vital signs) per person and trajectories up to 50 days.
- Estimate the effect of antibiotics, vassopressors and mechanical ventilator on patient covariates.
- Hidden confounding is present in the dataset as patient comorbidities and several lab tests were not included.

	White blood cell count		Blood pressure		Oxygen saturation	
Outcome model	MSM	R-MSN	MSM	R-MSN	MSM	R-MSN
Confounded	3.90 ± 0.00	2.91 ± 0.05	12.04 ± 0.00	10.29 ± 0.05	2.92 ± 0.00	1.74 ± 0.03
$D_Z = 1$	3.55 ± 0.05	2.62 ± 0.07	11.69 ± 0.14	9.35 ± 0.11	2.42 ± 0.02	1.24 ± 0.05
$D_Z = 5$	3.56 ± 0.04	2.41 ± 0.04	11.63 ± 0.10	9.45 ± 0.10	2.43 ± 0.02	1.21 ± 0.07
$D_Z = 10$	3.58 ± 0.03	2.48 ± 0.06	11.66 ± 0.14	9.20 ± 0.12	2.42 ± 0.01	1.17 ± 0.06
$D_Z = 20$	3.54 ± 0.04	2.55 ± 0.05	11.57 ± 0.12	9.63 ± 0.14	2.40 ± 0.01	1.28 ± 0.08

Discussion and limitations

- The Time Series Deconfounder enables the estimation of treatment effects over time using weaker assumptions than existing methods.
- Identifiability of the potential outcomes using the deconfounder framework may represent an issue:
  - non-identifiability will be indicated by the high variance of the estimated outcomes.

More from van der Schaar Lab...

- Read our publications:
  `https://www.vanderschaar-lab.com/publications/`
- Use our software:
  `https://www.vanderschaar-lab.com/software/`

Thank you for listening!