IPW for Temporality not Treatment

Temporal Inverse Probability Weighting for Causal Discovery in Controlled Before–After Studies: Discovering ADEs in Generics

Aubrey Barnard · **David Page** · **Peggy Peissig**

University of Wisconsin–Madison · Duke University · Marshfield Clinic Research Institute

Temporal Confounding Solved by IPW

Problem Temporal confounding (in longitudinal causal discovery).

Solution Controlled before–after study analyzed with:

- DFC differential classification old good
- DFP differential prediction old better
- TIPW temporal IPW new best!

Application Discover adverse drug events (ADEs) in generic drugs.

Study Compare brand to generic drugs over time in electronic health records (EHR) data.

Causal Discovery Methods hypothesize effects to discover unknown ADEs.

Causal CTBN for EHR Timelines

Continuous-time Bayesian network causal model generates synthetic EHR timelines for brand v. generic.

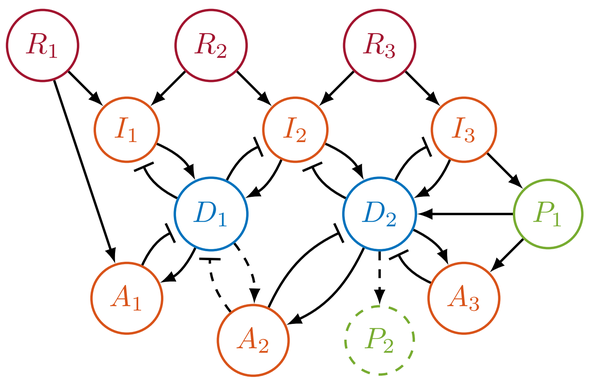

Figure 1: *(R)isk factor, (I)ndication, (D)rug, (P)rocedure, (A)DE. P_2 is introduced at the same time as generic D_1, midway through the timelines. Generic D_1 causes A_2 whereas brand does not. The dashed lines indicate these temporal differences. Perpendicular arrowheads $\dashv$ mark inhibitors.*

Controlled Before–After Studies

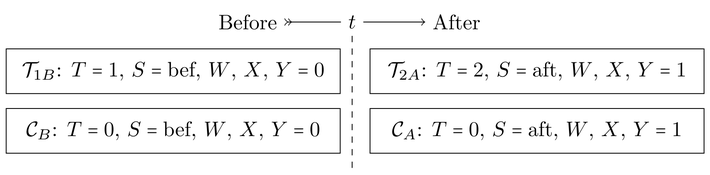

Figure 2: *Controlled before–after study for two treatments T, with time spans S, unit weight W, covariates X, and outcome Y. (Outcome is only for illustration.)*

Study Analysis

Difference in Differences as Binary Classification

Conceptually, the analysis is difference in differences, but we set it up as a binary classification task.

$$(f(\mathcal{T}_{2A}) - f(\mathcal{T}_{1B})) - (f(\mathcal{C}_A) - f(\mathcal{C}_B)) \tag{1}$$

$$\text{Classes: } +\mathcal{T}_{2A},\ -\mathcal{T}_{1B},\ -\mathcal{C}_A,\ +\mathcal{C}_B \tag{2}$$

TIPW extends this concept to a new algorithm.

TIPW Optimizes Relative Risk

TIPW finds events $E = f(X)$ that maximize RR.

$$\overset{\text{RR}}{\frac{\mathbb{P}(E \mid T=2)}{\mathbb{P}(E \mid T=1)}} = \overset{\mathcal{M}_\mathcal{T}}{\frac{\mathbb{P}(T=2 \mid E)}{\mathbb{P}(T=1 \mid E)}} \overset{\mathcal{M}_\mathcal{C}^{-1}}{\frac{\mathbb{P}(T=1)}{\mathbb{P}(T=2)}} \tag{3}$$

$$\frac{\mathbb{P}(T=1)}{\mathbb{P}(T=2)} = \frac{\mathbb{P}(S=b)}{\mathbb{P}(S=a)} = 1 \Big/ \overset{\text{IPW}}{\frac{\mathbb{P}(S=a)}{\mathbb{P}(S=b)}} \tag{4}$$

TIPW Dominates DFC

TIPW is more statistically efficient than DFC in terms of true positive rate (TPR).

Proposition 1. *Let TIPW and DFC have TPR α on the treateds. Let DFC have TPR β on the controls. Then, in terms of overall TPRs,*

$$\alpha \geq \beta \iff \mathit{TPR}_{\mathit{TIPW}} \geq \mathit{TPR}_{\mathit{DFC}}. \tag{5}$$

Temporal IPW Algorithm

Solves the study analysis binary classification problem in a novel way by using IPW to correct for temporality not treatment. Models the temporal trends in the controls, reweights to remove those trends from the treateds, and then models the treateds to discover causal effects.

Input: sets of data $\mathcal{C}_B$, $\mathcal{C}_A$, $\mathcal{T}_{1B}$, $\mathcal{T}_{2A}$ from a controlled before–after study
Output: model $\mathcal{M}_\mathcal{T}$ that discriminates treatments while controlling for changes over time

$$\texttt{ipw}(m, t, x) = 1/\mathbb{P}_{\mathcal{M}=m}(T = t \mid X = x) \tag{IPW}$$

$\mathcal{M}_\mathcal{C} \leftarrow \texttt{fit}(-\mathcal{C}_B \veebar +\mathcal{C}_A)$; `// Fit model of temporal trends`
for $\mathcal{T}$ **in** $[\mathcal{T}_{1B}, \mathcal{T}_{2A}]$ **do** `// Remove trends by reweighting treateds`
  $\mathcal{T}' \leftarrow [(t, s, x, w \cdot \texttt{ipw}(\mathcal{M}_\mathcal{C}, t, x))$ **for** $(T = t, S = s, X = x, W = w)$ **in** $\mathcal{T}]$;
$\mathcal{M}_\mathcal{T} \leftarrow \texttt{fit}(-\mathcal{T}'_{1B} \veebar +\mathcal{T}'_{2A})$; `// Fit model of treatments`

EHR Timelines

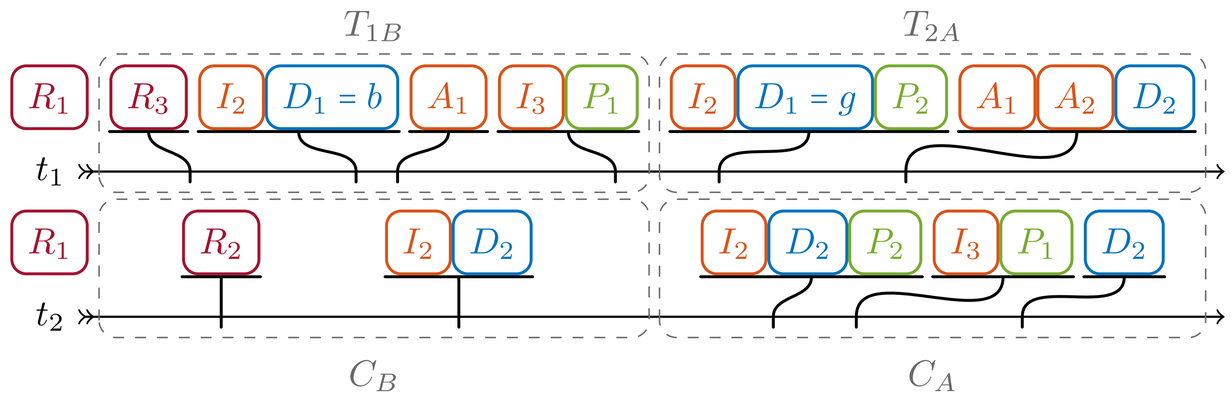

Figure 3: *Example EHR timelines, divided into study groups. The treatment is D_1.*

Binary Classifiers for Controlled Before–After Studies

Method	Model	Before–After Study Classification Setup
DFP, 2-model	$\mathbb{P}(Y = 1 \mid T = 1, X) \veebar \mathbb{P}(Y = 1 \mid T = 0, X)$	$\mathcal{M}_{\text{DFP2}}(\mathcal{M}_\mathcal{C}(-\mathcal{C}_B \veebar +\mathcal{C}_A) \veebar \mathcal{M}_\mathcal{T}(-\mathcal{T}_{1B} \veebar +\mathcal{T}_{2A}))$
DFP, 1-model	$\mathbb{P}(Z \mid X)$ where $Z = (T=0)(Y=0) + (T=1)(Y=1)$	$\mathcal{M}_{\text{DFP1}}(-\mathcal{C}_A \veebar -\mathcal{T}_{1B} \veebar +\mathcal{C}_B \veebar +\mathcal{T}_{2A})$
DFC	$\mathbb{P}(Z \mid X, Y)$ where $Z = (T=0)(S=b) + (T=1)(S=a)$	$\mathcal{M}_{\text{DFC}}((-\mathcal{C}_A \cup -\mathcal{T}_{1B}) \veebar (+\mathcal{C}_B \cup +\mathcal{T}_{2A}))$
temporal IPW	$\mathbb{P}(S \mid T=0, W, X) \overset{\text{IPW}}{\to} \mathbb{P}(T \mid T \neq 0, W', X)$	$\mathcal{M}_{\text{TIPW}}(\mathcal{M}_\mathcal{C}(-\mathcal{C}_B \veebar +\mathcal{C}_A) \overset{\text{IPW}}{\to} \mathcal{M}_\mathcal{T}(-\mathcal{T}'_{1B} \veebar +\mathcal{T}'_{2A}))$

Table 1: *Analysis methods. DFP: differential prediction, DFC: differential classification, $\veebar$: versus.*

TIPW Recovers the Known ADE in Synthetic EHR Data

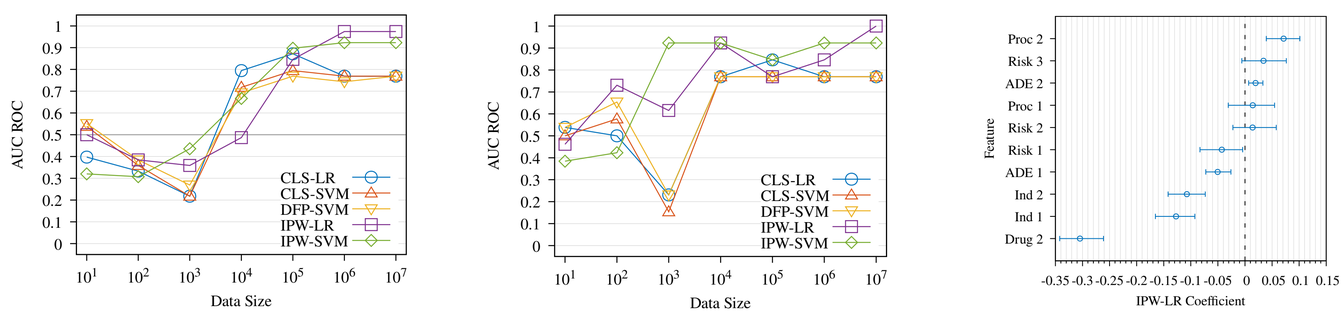

(a) Default hyperparameters (b) Best avg. ADE Discovery (c) Top 10 IPW-LR Features

Figure 4: *Results of experiments on the synthetic data: learning curves and top features from "default" IPW-LR on data size 10^5. Positive coefficients favor generic; negative coefficients favor brand. CLS: differential classification, DFP: differential prediction, IPW: temporal IPW, LR: logistic regression, SVM: support vector machine with a linear kernel.*

TIPW Resists False Discoveries in Real EHR Data

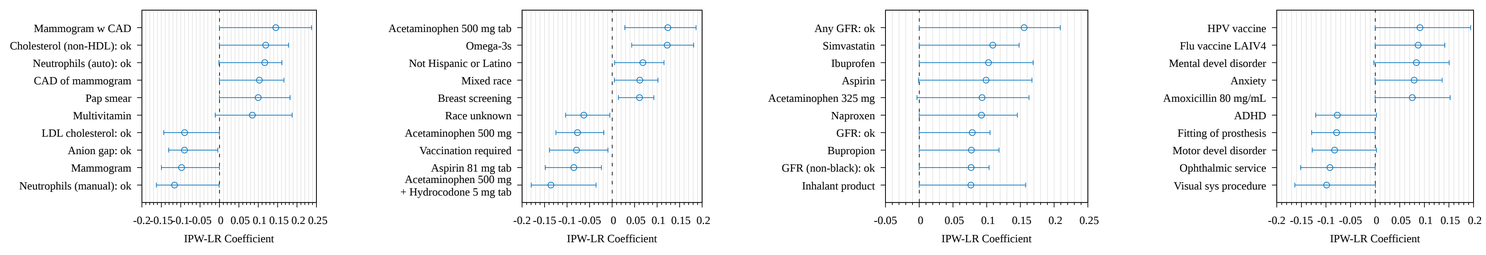

(a) Bupropion (b) Duloxetine (c) Gabapentin (d) Methylphenidate

Figure 5: *Top 10 features from IPW-LR by LR coefficient magnitude with bootstrapped 99% confidence intervals. Positive coefficients favor generic; negative coefficients favor brand.*

Causal Discovery & Hypothesizing Effects

Possible effects E are unknown, so can't model $\mathbb{P}(E \mid X)$. Instead, model the treatment $\mathbb{P}(T \mid X)$, based on data after treatment. Detecting differences between treateds and controls $\rightarrow$ finding possible effects.

Acknowledgments This work was supported by the US FDA (FDA BAA13-00119) (AB, DP), and the Computation and Informatics in Biology and Medicine Training Program (NLM training grant 5T15LM007359) (AB). The views are the authors' own; they do not speak on behalf of the FDA or NLM.