Addressing Hidden Confounding with Heterogeneous Observational Datasets for Recommendation

Yanghao Xiao [1,3] Haoxuan Li[2] Yongqiang Tang [3] Wensheng Zhang [4]

[1]University of Chinese Academy of Sciences [2]Peking University [3]Institute of Automation, Chinese Academy of Sciences [4]Guangzhou University

Contributions

- This is the first work to employ heterogeneous observational datasets to address hidden confounding in debiased recommendations, wherein some data is subject to hidden confounding while the remaining is not.
- In this study, we relax the reliance on expensive randomized controlled trial (RCT) data in previous data fusion methods.
- We propose a meta-learning based debiasing method called MetaDebias to explicitly estimate the oracle error imputation and hidden confounding bias, employing bi-level optimization for model training.
- We conduct extensive experiments on three public datasets, and our method achieves state-of-the-art performance in the presence of hidden confounding, regardless of the availability of RCT data.

Preliminaries

- *Unit*: a user-item pair (u,i).
- *Target population*: the set of all user-item pairs $\mathcal{D}=\{(u,i) \mid u \in \mathcal{U}, i \in \mathcal{I}\}$.
- *Feature*: $x_{u,i}$, the observed feature of user u and item i.
- *Treatment*: $o_{u,i} \in \{0,1\}$ is the exposure indicator of (u,i).
- *Outcome*: $r_{u,i}$, the feedback of user-item pair (u,i).
- *Data Source Indicator*: $g_{u,i} \in \{0,1\}$ indicates whether hidden confounding exists.

Motivation

- Existing methods for mitigating hidden confounding are challenging to be applied in real-world scenarios, as they either rely on **strong assumptions** on hidden confounding strength or depend on the **costly RCT data**.

Sensitivity Analysis

- Sensitivity analysis based approach assumes the true propensity $p_{u,i}$ is near and can be bounded by the estimated propensity $\hat{p}_{u,i}$, i.e., given bound $\Gamma \geq 1$

$$\frac{1}{\Gamma} \leq \frac{(1-\hat{p}_{u,i})\, p_{u,i}}{\hat{p}_{u,i}\,(1-p_{u,i})} \leq \Gamma.$$

However, above strong assumption on hidden confounding strength is hard to be satisfied in real world, and such method fails when the assumption is violated.

Model Calibration with RCT data

- Recent works propose to leverage a few unbiased RCT data for model calibration, where biased propensity and imputation models can be corrected using such unbiased loss, for instance, with the help of additive residual models or multiplicative reweighting models.
- However, collecting RCT data requires users to rate items randomly, which indicates the acquisition cost of RCT data is prohibitively high, posing challenges to the practical implementation of such methods in real-world settings.

Proposed Method

Key Idea

- We propose to explicitly estimate the prediction error $e_{u,i}=L(\hat{r}_{u,i}, r_{u,i})$ on all user-item pairs $\mathcal{D}$ and hidden confounding bias, where $\hat{r}_{u,i}=f(x_{u,i};\theta)$ is the prediction model with parameter θ, and $L(\cdot,\cdot)$ is a loss function.
- The goal is to accurately estimate the oracle error imputation $\mathbb{E}_{\mathcal{D}}\left[e_{u,i} \mid x_{u,i}\right]$.

Propensity Score and Naive Imputation

- We define the identifiable propensity score $\pi(x,g)$ to model the two types of missing mechanisms for both absence and presence of hidden confounding:

$$\pi(x,g)=\mathbb{P}\left(o_{u,i}=1 \mid x_{u,i}=x, g_{u,i}=g\right).$$

- We define the naive error imputation $m(x,g)$ on target population $\mathcal{D}$:

$$m(x,g)=\mathbb{E}[o_{u,i}\cdot e_{u,i} \mid x_{u,i}=x, g_{u,i}=g].$$

Estimation of Oracle Error Imputation

- Based on the propensity and naive imputation, we have:

$$m(x,g)=\{\mathbb{E}\left[e_{u,i} \mid x_{u,i}=x\right]+(1-g)\eta(x)\}\cdot \pi(x,g),$$

where $\eta(x)=\mathbb{E}\left[e_{u,i} \mid x_{u,i}=x, g_{u,i}=0, o_{u,i}=1\right]-\mathbb{E}\left[e_{u,i} \mid x_{u,i}=x\right]$ is the bias introduced by hidden confounding.
- Further incorporating more information, we can achieve a robust estimation:

$$o_{u,i}\cdot e_{u,i}-m(x,g)=\{\mathbb{E}\left[e_{u,i} \mid x_{u,i}=x\right]+(1-g)\eta(x)\}\cdot\{o_{u,i}-\pi(x,g)\}+\xi,$$

where $\xi=o_{u,i}\cdot\{e_{u,i}-\{\mathbb{E}[e_{u,i} \mid x_{u,i}=x]+(1-g)\eta(x)\}\}$ with $\mathbb{E}[\xi \mid x,g]=0$ can be regarded as a noise due to its zero-mean property.

Experiments

Table 1. Recommendation performances in terms of AUC, Recall@5 (R@5), NDCG@5 (N@5) on Coat and Yahoo! R3.

	Coat			Yahoo! R3		
Method	AUC	R@5	N@5	AUC	R@5	N@5
Naive	0.698	0.478	0.444	0.705	0.638	0.489
IPS	0.717	0.483	0.446	0.699	0.642	0.492
DR	0.725	0.485	0.448	0.709	0.643	0.498
Stable-DR	0.734	0.486	0.452	0.715	0.656	0.515
TDR	0.736	0.492	0.458	0.717	0.669	0.525
ESCM^2-IPS	0.730	0.484	0.451	0.713	0.666	0.520
ESCM^2-DR	0.737	0.492	0.458	0.715	0.670	0.521
BRD-IPS	0.733	0.490	0.462	0.712	0.659	0.515
BRD-DR	0.739	0.494	0.464	0.714	0.663	0.516
KD-Label	0.735	0.488	0.461	0.712	0.664	0.517
AutoDebias	0.736	<u>0.501</u>	0.465	0.710	0.667	0.520
Bal-IPS	0.733	0.486	0.462	0.708	0.665	0.515
Bal-DR	0.735	0.490	0.464	0.708	0.668	0.517
Res-IPS	0.738	0.494	0.465	0.718	0.675	0.534
Res-DR	<u>0.740</u>	0.498	<u>0.467</u>	<u>0.720</u>	<u>0.678</u>	<u>0.538</u>
MetaDebias(ours)	**0.746**	**0.510**	**0.473**	**0.722**	**0.688**	**0.544**

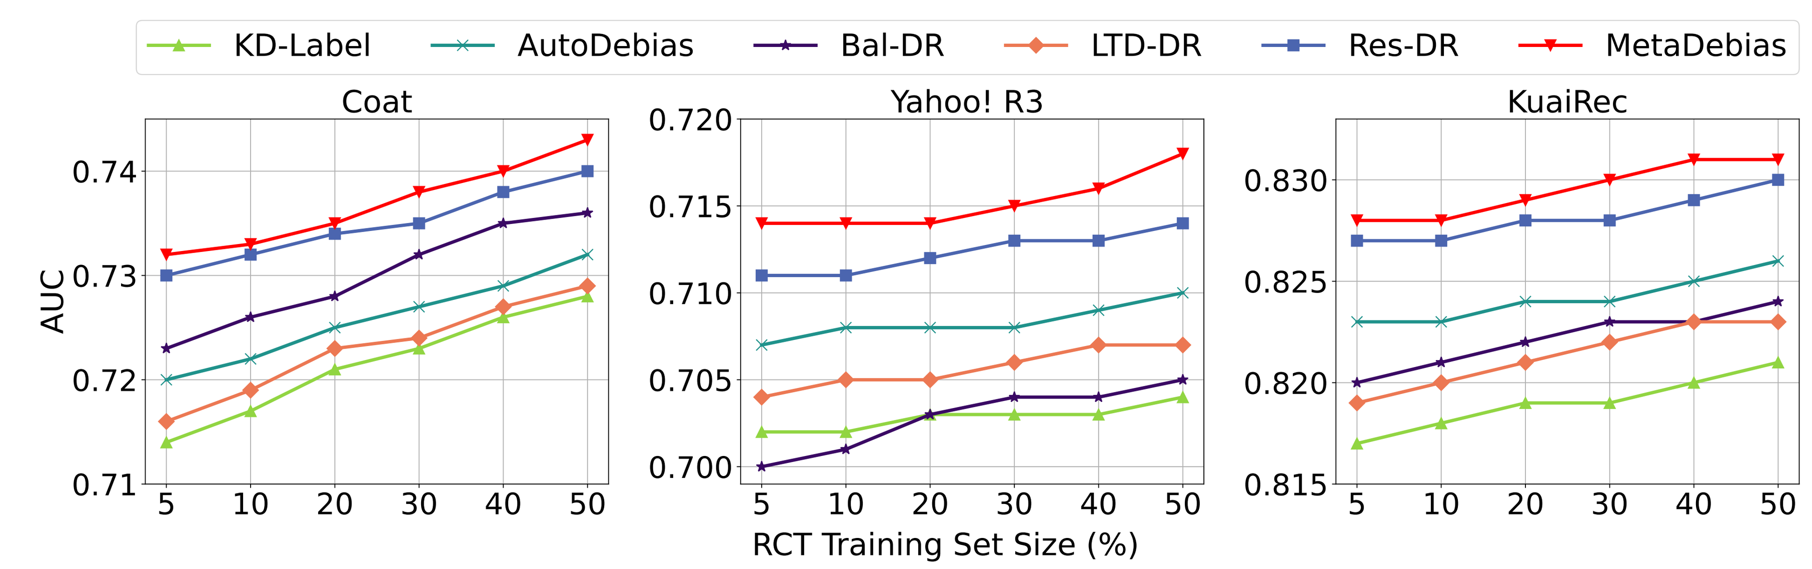

Figure 1. Effects of varying RCT training set size on AUC on three datasets.