Backdoor Adjustment via Group Adaptation for Debiased Coupon Recommendations

Junpeng Fang, Gongduo Zhang, Qing Cui, Caizhi Tang, Lihong Gu, Longfei Li, Jinjie Gu, Jun Zhou*

Ant Group, Hangzhou, China,

{junpeng.fjp,gongduo.zgd,cuiqing.cq,caizhi.tcz,lihong.glh,longyao.llf,jinjie.gujj}@antgroup.com; jun.zhoujun@antfin.com

Abstract

Accurate prediction of coupon usage is crucial for promoting user consumption through targeted coupon recommendations. However, in real-world coupon recommendations, the coupon allocation process is not solely determined by the model trained with the history interaction data but is also interfered with by marketing tactics desired to fulfill specific commercial goals. This interference creates an imbalance in the interactions, which causes the data to deviate from the user's natural preferences. We refer to this deviation as the matching bias. Such biased interaction data affects the efficacy of the model, and thus it is necessary to employ debiasing techniques to prevent any negative impact. We investigate the mitigation of matching bias in coupon recommendations from a causal-effect perspective. By treating the attributes of users and coupons associated with marketing tactics as confounders, we find the confounders open the backdoor path between user-coupon matching and the conversion, which introduces spurious correlation. To remove the bad effect, we propose a novel training paradigm named Backdoor Adjustment via Group Adaptation (BAGA) for debiased coupon recommendations, which performs intervened training and inference, i.e., separately modeling each user-coupon group pair. However, modeling all possible group pairs greatly increases the computational complexity and cost. To address the efficiency challenge, we further present a simple but effective dual-tower multi-task framework and leverage the Customized Gate Control (CGC) model architecture, which separately models each user and coupon group with a separate expert module. We instantiate BAGA on five representative models: FM, DNN, NCF, MASKNET, and DEEPFM, and conduct comprehensive offline and online experiments to demonstrate the efficacy of our proposed paradigm.

Introduction

With the development of e-commerce and m-payment, recommendation technology has been increasingly used in online marketing campaigns. For instance, a payment platform periodically organizes marketing campaigns designed to allocate online coupons to incentivize consumption and bolster customer retention (Li et al. 2020; Zhu et al. 2023; Fang et al. 2023). The key objective of modeling these campaigns is to estimate the conversion probability, which refers to the likelihood of a user using a coupon. As with other recommendation scenarios, researchers typically employ User-Item Click-Through-Rate (CTR) models (Sim and Lee 2014; Mutanen, Nousiainen, and Liang 2010; Huan et al. 2022; Huangfu et al. 2022) to handle this problem.

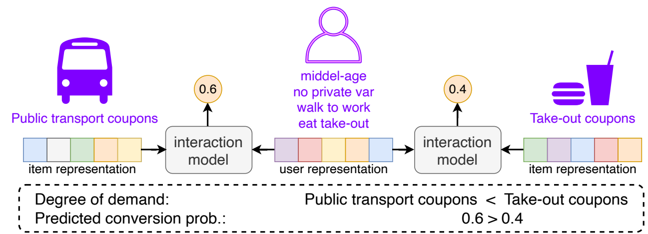

Figure 1: Empirical evidence of how imbalanced historical interaction data leads to imprecise estimations.

To achieve specific commercial objectives, such as maximizing daily active users (DAUs) or increasing the exposure of a particular coupon, the organizers may employ specific tactics to interfere with the coupon allocation process. These tactics can be manifested as limiting the likelihood of certain types of users receiving specific coupons or limiting the likelihood of certain types of coupons being exposed to particular users. By leveraging these tactics, organizers can effectively drive user behavior and achieve their desired goals. However, this interference can create an imbalance in user-coupon interaction data distribution, i.e., result in certain user-coupon pairs dominating the majority of historical interaction data, which causes the data to deviate from the user's natural preferences. This deviation is referred to as the *matching bias* (Fang et al. 2023). Models trained with such data tend to further exacerbate the imbalance by overemphasizing majority user-coupon pairs and distorting the estimation of conversion probability. Ultimately, this bias decreases the efficiency of the recommendation model.

Fig.1 shows the empirical evidence of the affection of matching bias in the marketing coupon recommendation scenario of a payment platform studied in this paper. Assuming there are two tactics employed [1]: public transport

*Corresponding author.

[1] Here, the tactics work by multiplying the predicted model scoring of coupons by a specific weighted score, and the resulting

coupons have a higher likelihood of being exposed to users without private cars, while take-out coupons have a higher likelihood of being exposed to younger users. For example, a middle-aged user who walks to work and occasionally eats take-out at work may have a greater demand for take-out coupons than public transport coupons. However, if the pair of users without private cars and public transport coupons occupies the majority in the training sample, while the pair of middle-aged people and take-out coupons are the minority, the model may learn this association and thus preferentially send public transport coupons to this user.

Matching bias is a prevalent issue in the recommendation, yet it has received less attention than biases caused by item popularity or position (Chen et al. 2023; Biega, Gummadi, and Weikum 2018; Zheng et al. 2021a; Zhang et al. 2021). To address this, we adopt the causal graph language (Pearl 2009) for qualitative analysis (see Fig.2 below). Through causal relationship examination, we observe that the attributes of users and coupons related to marketing tactics can impact both the conversion probability and user-coupon representations, resulting in a confounding bias in the causal inference literature (Grimes and Schulz 2002; VanderWeele, Hernán, and Robins 2008; Tang et al. 2023, 2022). Therefore, we transform the problem of mitigating matching bias to adjusting for the confounding effect, as described formally in Section **Methods**.

Inspired by (He et al. 2023; Zhang et al. 2023), we propose a novel de-confounded training paradigm called ***B**ackdoor **A**djustment via **G**roup **A**daptation* (BAGA) to eliminate the confounding effect caused by the grouping attributes. Specifically, by leveraging the principle of backdoor adjustment (Pearl 2009), BAGA cuts off the backdoor path in the causal graph which removes the matching bias and models the genuine causality relationship between user-coupon representations and conversion probabilities. Operationally, we segment the training data into different user and coupon groups based on marketing tactics in an adaptive manner. The underlying embedding is shared between groups, and for each group, we train a submodel to predict intra-group conversion probabilities using a multi-task hierarchical prediction approach. This allows for bottom model parameter sharing information across different adaptive groups while ensuring the variability of each group, bringing benefits for better representations and faster convergence. However, modeling all possible user-coupon group pairs greatly increases the computational complexity and cost. To address the efficiency challenge, we further present a simple but effective dual-tower multi-task framework and leverage the Customized Gate Control (CGC) (Tang et al. 2020) model architecture (*cf.* Fig.3), which separately models each user and coupon group with a separate expert module. With this approach, we retain the expressive power of the model at little additional cost.

To summarize, the contributions are listed as follows:

- **Causal Formulation of Matching Bias in Marketing Coupon Recommendations.** We provide a formal definition of the concept of matching bias, which is a common issue arising from the adoption of marketing tactics.
- **Causal-Effect Perspective Qualitative Analysis.** We investigate the removal of matching bias from a causal perspective and reveal the confounders that lead to the bias via the causal-directed acyclic graph framework.
- **Novel Training Paradigm.** We propose a novel de-confounded training paradigm BAGA to alleviate the matching bias by leveraging the backdoor adjustment.
- **Extensive Offline and Online Evaluations.** We instantiate BAGA on five representative models: FM (Rendle 2010), DNN (Goodfellow, Bengio, and Courville 2016), NCF (He et al. 2017), MASKNET (Wang, She, and Zhang 2021), and DEEPFM (Guo et al. 2017), and conduct extensive offline and online experiments to demonstrate the effectiveness of our proposed paradigm.

score is used to make decisions about coupon issuance.

Related Work

In this study, we investigate the alleviation of the matching bias from a causal-effect perspective which pertains to the areas of bias and debias in recommendations.

Bias is a pervasive issue in recommendations that arises from various factors, including clickbait (Wang et al. 2021a), feedback loop (Chaney, Stewart, and Engelhardt 2018), user choice (Marlin et al. 2007a), and uneven exposure (Liu et al. 2020), etc. (Chen et al. 2023) provides a comprehensive survey of different types of biases in recommender systems, among which we highlight the following three: **Selection bias** arises when users have unrestricted freedom to interact with items, leading to non-representative observed interaction data (Marlin et al. 2007b; Sun et al. 2023). **Popularity bias** occurs when popular items are recommended more frequently than their popularity would warrant (Abdollahpouri and Mansoury 2020). **Exposure bias** arises when users are only exposed to a subset of relevant items, leading to skewed observed interactions (Liu et al. 2020).

We briefly describe four representative types of debiasing methods commonly employed, including: The inverse probability weighting (IPW) or the inverse propensity scores (IPS) are leveraged to mitigate bias by shifting the distribution of exposure to a more uniform state (Seaman and White 2013; Schnabel et al. 2016; Mansournia and Altman 2016; Saito et al. 2020). Disentangling of embedding entails separating the interest representation from the bias representation to model a more robust causal relationship for debiasing (Wang et al. 2020; Zheng et al. 2021a,b). Double robust methods integrate imputed errors and propensities in a doubly robust manner to obtain unbiased estimation (Wang et al. 2019; Li, Zheng, and Wu 2022). Unbiased data can be added to train a new task with a small amount of unbiased data to obtain unbiased representations. Subsequently, the original representation distribution is constrained to unbiased representations, making the model more robust (Zhao et al. 2023; Wang et al. 2021b), etc.

Methods

In this section, we first trace the origin of the matching bias by means of the causal graph, followed by the detail of the proposed debiased approach BAGA.

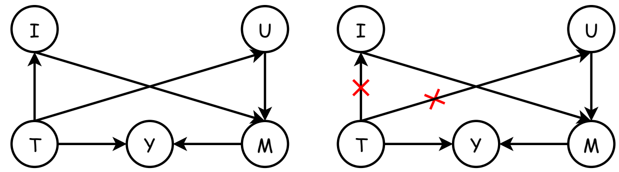

(a) Confounding effect of matching bias (b) Backdoor adjustment for matching bias

Figure 2: Causal graphs for marketing coupon recommendations, where T - confounder, U - user, I - coupon, M - matching result, and Y - conversion probability.

Causal View of Matching Bias

Causal graphs. In this study, we leverage causal directed acyclic graphs (DAG) (Pearl 2009) — comprising nodes and edges, with acyclic graphs and arrows indicating causal relations — to explicitly analyze the causal relationships. Random variables, particular values, and sample spaces are represented by capital letters (e.g., U), lowercase letters (e.g., $\boldsymbol{u}$), and calligraphic font (e.g., $\mathcal{U}$), respectively. Fig.2 illustrates the causal relations behind the marketing coupon recommendations studied in this paper, which consist of five random variables (U, I, M, T, Y). Specifically:

- Node U and I denote the user and coupon variable, including age, gender, coupon id, discount, etc.
- Node M denotes the real affinity score (degree of match) between a user U and a coupon I, reflecting what extent the coupon matches the preference of the user.
- Node T denotes the hierarchical attributes of users and coupons that are associated with the marketing tactics.
- Node Y denotes the conversion probability, or prediction score, i.e., whether the user used the coupon or not.
- Edges $(U, I) \rightarrow M$ capture the causal effect of the degree of match between users and coupons.
- Edges $T \rightarrow (U, I)$ indicate that the hierarchical attributes of users and coupons that are associated with the tactics will affect the recommendations. Due to the unbalanced distribution, recommender systems tend to favor the majority user-coupon pairs as shown in Fig.1.
- Edges $(T, M) \rightarrow Y$ indicate that the conversion probability is determined by the combination of the model scoring as well as the marketing tactics, implying that when two coupons equally match user interest, the coupon of the majority pair may be prioritized due to the biased distribution caused by the marketing tactic.

Matching bias. According to Fig.2(a), there exist backdoor paths $(U, I) \leftarrow T \rightarrow Y$, which cause spurious correlations between (U, I) and Y. However, the non-debiasing methods ignore the confounding effect caused by T and directly estimate the conditional probability $P(Y|U, I)$, which leads to imprecise estimations. Formally, the conditional probability $P(Y|U = \boldsymbol{u}, I = \boldsymbol{i})$ are derived (See **Appendix** for detailed derivations) as:

$$
\begin{aligned}
&P(Y|U=\boldsymbol{u}, I=\boldsymbol{i}) &(1)\\
&\overset{(1a)}{=} \sum_{\boldsymbol{t}}\sum_{\boldsymbol{m}} P(Y, \boldsymbol{t}, \boldsymbol{m}, \boldsymbol{u}, \boldsymbol{i}) / \sum_{\boldsymbol{t}}\sum_{\boldsymbol{m}} P(\boldsymbol{t}, \boldsymbol{m}, \boldsymbol{u}, \boldsymbol{i})\\
&\overset{(1b)}{=} P(Y|\boldsymbol{t}_{\boldsymbol{ui}}, M(\boldsymbol{u}, \boldsymbol{i})) * P(\boldsymbol{t}_{\boldsymbol{ui}}|\boldsymbol{u}, \boldsymbol{i}).
\end{aligned}
$$

From Eq.($1b$), it is worth noting that $\boldsymbol{t}$ affect not only $(\boldsymbol{u}, \boldsymbol{i})$, but also Y, causing a spurious correlation between them, i.e., $P(\boldsymbol{t}_{\boldsymbol{ui}}|\boldsymbol{u}, \boldsymbol{i})$ force $\boldsymbol{u}$ and $\boldsymbol{i}$ to favor the majority user-coupon pairs, leading to overestimate Y via $P(Y|\boldsymbol{t}_{\boldsymbol{ui}}, \boldsymbol{m}_{\boldsymbol{ui}})$. The preceding analysis suggests that to precisely estimate conversion probability, we have to block the confounding effect and focus on the causal relationships between these variables so that spurious correlations can be avoided. Therefore, we adopt *do*-operator $do(\cdot)$ (Pearl 2009) to model the causal effect of (U, I) on Y, denoted by $P\big(Y|do(U = \boldsymbol{u}), do(I = \boldsymbol{i})\big)$.

De-Matching Bias in BAGA

In this section, we leverage backdoor adjustment (Pearl 2009) to de-confound matching bias. The key of BAGA lies in deriving and computing the causal effect, and we provide the theoretical proof and architectures form later.

Backdoor adjustment. We can remove the effect of spurious association by cutting off the backdoor paths in the causal DAG as illustrated in Fig.2(b). Formally, the expression of $P\big(Y|do(U = \boldsymbol{u}), do(I = \boldsymbol{i})\big)$ are derived as follow:

$$
\begin{aligned}
&P\big(Y|do(U=\boldsymbol{u}), do(I=\boldsymbol{i})\big) &(2)\\
&\overset{(2a)}{=} \sum_{\boldsymbol{t}} P\Big(Y|\boldsymbol{t}, M\big(do(\boldsymbol{u}), do(\boldsymbol{i})\big)\Big) * P\big(\boldsymbol{t}|do(\boldsymbol{u}), do(\boldsymbol{i})\big)\\
&\overset{(2b)}{=} \sum_{\boldsymbol{t}} P\big(Y|\boldsymbol{t}, M(\boldsymbol{u}, \boldsymbol{i})\big) * P(\boldsymbol{t}),
\end{aligned}
$$

where the derivation of Eq.($2a$) is similar to that of Eq.($1b$), which follows the law of total probability and Bayes rule with *do*-calculus. Eq.($2b$) utilizes *insertion/deletion of actions* and *action/observation exchange* in Theorem 3.4.1 of (Pearl 2009) to simplify the formula.

Eq.(2) sheds light on the design to de-confound the effect of confounder (hierarchical attributes): one can separately estimate $P\big(Y|\boldsymbol{t}, M(\boldsymbol{u}, \boldsymbol{i})\big)$ for the case of $P(\boldsymbol{t}) = 1$ and then output the results corresponding to the case of $P(\boldsymbol{t}) = 1$ at the inference stage. In this paper, we propose to discretize the hierarchical attribute distribution into adaptive disjoint groups and fit a group-wise coupon usage prediction model to complete the estimation.

Approximation. We now present a simple but effective dual-tower multi-task framework that leverages the Customized Gate Control (CGC) model architecture, as illustrated in Fig.3.

The high-level idea is to segment users and coupons based on the hierarchical attributes utilized in marketing tactics. Then constructs group-wise estimation to de-confounded coupon usage. Operationally, to cut off edges $T \rightarrow (U, I)$, we adaptively segment training samples into J user groups

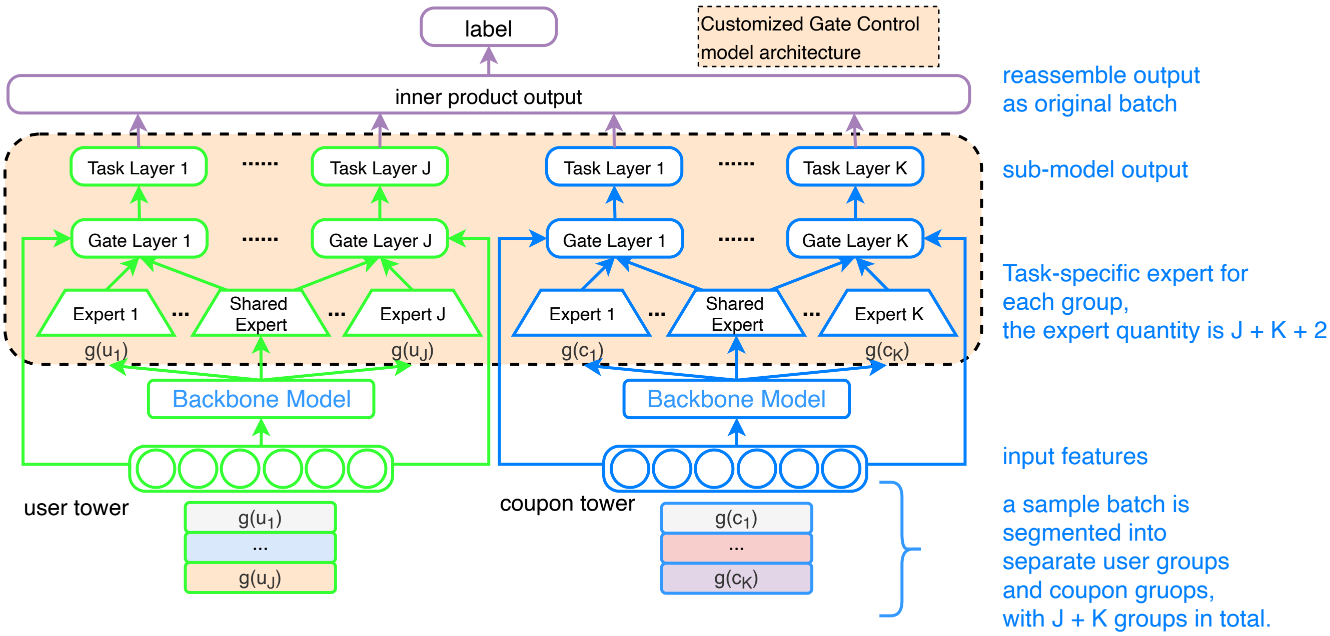

Figure 3: Illustration of the dual-tower multi-task learning framework that leverages the CGC model architecture, where $\{g(u_1), ..., g(u_J)\}$ and $\{g(c_1), ..., g(c_K)\}$ denotes different user and coupon groups, respectively. The left and right parts of the network represent the user and coupon towers, respectively. Each task is trained with only the corresponding group of samples, and then the outputs of the corresponding samples are reassembled into the original batch to perform the inner product.

and K coupon groups based on hierarchical attributes respectively, which discretizes the distribution $P(\boldsymbol{t})$ into disjoint components. Let $\{\mathcal{U}_m\}_{m=1}^J$ be the user groups and $\{\mathcal{I}_n\}_{n=1}^K$ be the coupon groups, conceptually, the sample space of $\boldsymbol{t}$ is $J * K$. Continuing the derivation in Eq.(2), we estimate the conversion probability $P\big(Y|do(U = \boldsymbol{u}), do(I = \boldsymbol{i})\big)$ via the following approximation:

$$\begin{aligned} & P\big(Y|do(U = \boldsymbol{u}), do(I = \boldsymbol{i})\big) \qquad (3) \\ & = \sum\nolimits_{\boldsymbol{t}} P\big(Y|\boldsymbol{t}, M(\boldsymbol{u}, \boldsymbol{i})\big) * P(\boldsymbol{t}) \\ & \approx \sum\nolimits_{m=1}^{J} \sum\nolimits_{n=1}^{K} P\big(Y|\boldsymbol{t}, M(\boldsymbol{u} \in \mathcal{U}_m, \boldsymbol{i} \in \mathcal{I}_n\big) * P(\boldsymbol{t}_{mn}) \\ & \triangleq \sum\nolimits_{m=1}^{J} \sum\nolimits_{n=1}^{K} \mathbf{1}\{\boldsymbol{u} \in \mathcal{U}_m \wedge \boldsymbol{i} \in \mathcal{I}_n\} f_{mn}\big(\boldsymbol{u}, \boldsymbol{i}\big), \end{aligned}$$

where for each group $\{\mathcal{U}_m\}$ and $\{\mathcal{I}_n\}$, $f_{mn}\big(\boldsymbol{u}, \boldsymbol{i}\big)$ is the sub-model fitted on samples $\{(\boldsymbol{u}, \boldsymbol{i}, Y)|\text{where } \boldsymbol{u} \in \mathcal{U}_m\ ,\ \boldsymbol{i} \in \mathcal{I}_n\}$, and $\mathbf{1}(\cdot)$ is the indicator function.

We hereby provide an intuitive explanation of why such an adaptively data segmentation procedure can mitigate the matching bias as illustrated in Fig.2(b). In prediction models such as DNN (Goodfellow, Bengio, and Courville 2016) or FM (Rendle 2010), majority group samples weight more in the gradient updating, leading to poor performance on minority group samples. By segmenting data and group-wisely fitting the model, we alleviate the interference from majority group samples on minority group samples during model training. However, this data segmentation training approach raises computational complexity concerns when fitting an individual sub-model f_{mn} for each user-coupon group pair. The model size will grow unacceptably large, with a complexity of $\mathcal{O}(J * K)$, which is not practical for real production systems. To address this concern, we propose a strategy that separately segments users and coupons into different groups and adopts a dual-tower multi-task framework as illustrated in Fig.2(b). This approach allows us to alleviate matching bias and maintain an acceptable model complexity (e.g., $\mathcal{O}(J + K)$) for scalability.

Segmentation and Overall BAGA Training

Segmentation Methodology. As illustrated in Fig.1, supposing a marketing scenario has only two specific tactics, i.e., for users, the segmentation is only related to whether they own a private car and the age group they are in. Then users can be segmented into $2 * 3 = 6$ [2] strata. The coupon segmentation procedure is the same as described above.

Overall Training. Moving on, we describe how we adapt a dual-tower multi-task framework that leverages the CGC architecture to train the prediction model. Concretely, the BAGA framework (*cf.* Fig.3) includes:

- *Backbone Model*, which aims to learn a representation for user $\boldsymbol{u}$ and coupon $\boldsymbol{i}$, respectively.
- *CGC Model*, which contains task-specific experts, shared experts, gate and task layers. Here, each sub-task-specific network models a group of samples individually, while

[2] With or without a private car; low, middle, and high age group.

Data	#Users	#Items	#UGs	#IGs	#Samples
Ds_A	156,401	41,002	8	20	1,094,807
Ds_B	565,495	2,002	15	9	1,809,584
Ds_C	90,039	12,303	8	8	6,302,730
Ds_D	1,141,980	1,342	12	16	10,277,820
Ds_E	8,591,785	326	8	21	60,611,596

Table 1: Characteristics of the experimental datasets.

the shared expert network takes full advantage of the global data optimization by aggregating the group samples modeled by each sub-task to ensure performance.

- *Inner Product*, which reassembles the task-related group samples' output of the CGC Model as the original batch [3] and performs inner product to map the output to the conversion probability.

Accordingly, $f_{mn}(\boldsymbol{u}, \boldsymbol{i})$ can be deduced by:

$$f_{mn}(\boldsymbol{u}, \boldsymbol{i}) = \big(f(\boldsymbol{u}) \cdot f(\boldsymbol{i})\big), \tag{4}$$

where $f(\cdot)$ is a recombination of the output of each task.

Then the parameters Θ of BAGA are optimized by:

$$\Theta = \arg\min_{\Theta} \sum_{m=1}^{J} \sum_{n=1}^{K} \mathbf{1}\{\boldsymbol{u} \in \mathcal{U}_m \wedge \boldsymbol{i} \in \mathcal{I}_n\} l\big(f_{mn}(\boldsymbol{u}, \boldsymbol{i}), Y\big), \tag{5}$$

where $Y \in \{0, 1\}$ indicates whether $\boldsymbol{u}$ has interacted with $\boldsymbol{i}$ (i.e., $\boldsymbol{y}_{ui} = 1$) or not (i.e., $\boldsymbol{y}_{ui} = 0$) and $l(\cdot)$ is the loss function, e.g., cross-entropy loss (Shore and Johnson 1981) or focal loss (Lin et al. 2017).

Experiments

We constructed our model with TensorFlow (Abadi et al. 2015) and assessed its performance through rigorous offline experiments. Additionally, we conducted online A/B tests in a real-world marketing coupon recommendation scenario to answer the following key questions.

- **Q1**: How does BAGA compare to the original (bias) model in terms of performance? (see **A1**)
- **Q2**: How does BAGA compare to other SOTA debiasing methods in terms of debiasing ability? (see **A2**)
- **Q3**: How does the CGC model architecture contribute to our proposed methods? (see **A3**)
- **Q4**: How does BAGA compare to other SOTA debiasing methods in terms of complexity? (see **A4**)
- **Q5**: How does BAGA perform in real-world marketing coupon recommendation applications? (see **A5**)

[3]Here the reassemble is in the form of corresponding positional splicing, i.e., the group samples corresponding to each sub-task are reorganized into the form of input batch according to the indexes.

Experimental Setup

Datasets Although our approach is conceived based on marketing coupon recommendations, it is essentially a click-through rate (CTR) model. Consequently, commonly available click-through rate datasets can also be utilized for experimental analysis as long as items are treated as coupons and the samples are filtered to simulate marketing tactics. We collected five experimental datasets as follows:

- Ds_A Amazon Dataset (Bengio, Ducharme, and Vincent 2000) contains transaction information from Amazon.
- Ds_B Ele.me Data (Tianchi 2022) is a set of recommendation-related data provided by Ele.me [4].
- Ds_C. MovieLens data (Harper and Konstan 2015) is a rating (a continuous value ranging from 0 to 5) data set collected from MovieLens website. To make it suitable for the CTR prediction task, we transform it into binary classification data by labeling the samples with ratings of 4 and 5 as positive and the rest as negative.
- Ds_D. Alibaba Ads Click Dataset (Tianchi 2018) is collected from the online advertising system in Alibaba.
- Ds_E. This dataset is non-open source data from a real-world marketing coupon recommendation scenario [5].

For the first four datasets, we filtered the samples to simulate marketing tactics. For the non-open-source dataset Ds_E, we directly conduct experiments on it because its scenario itself has already applied marketing tactics. Detailed characteristics are summarized in Table 1, where #UG and #IG indicate the number of user groups and coupon groups, respectively.

Benchmark Methods As the backbone model of BAGA is model-agnostic, we compare our method with five state-of-the-art model-agnostic debiasing methods. In particular, for **IPW** (Saito et al. 2020), we use the reciprocal of the group sample share as the propensity score; for **FAIRCO** (Morik et al. 2020), we calculate the error term based on the score list sorted by relevance; for **PDA** (Zhang et al. 2021), the hyper-parameter γ is set to 0.1 by gird search; for **DICE** (Zheng et al. 2021a), the hyper-parameters α and β are set to 0.1 and 0.01 as recommended, respectively; and for **DMBR** (Fang et al. 2023) the confounding effect of the imbalanced distribution of users/items over each other is eliminated. In addition, we use FM (Rendle 2010), DNN (Goodfellow, Bengio, and Courville 2016), NCF (He et al. 2017), MASKNET (Wang, She, and Zhang 2021), and DEEPFM (Guo et al. 2017) as vanilla models, respectively.

Evaluation Metrics

To evaluate the debiasing capabilities of the comparative method, we conducted experiments on unbiased testing data with reference to existing studies (Wei et al. 2021; Bonner and Vasile 2018; Zheng et al. 2021a). To accomplish this, we sample instances from a uniform distribution to simulate

[4]Ele.me is the top take-out platform in China.

[5]This dataset is desensitized and encrypted and thus does not contain any Personal Identifiable Information (PII). In addition, this dataset is only used for academic research, it does not represent any real business situation.

Methods	Ds_A	Ds_B	Ds_C	Ds_D	Ds_E
Vanilla	0.697±0.002	0.726±0.004	0.653±0.003	0.839±0.003	0.919±0.002
IPW	0.699±0.002	0.725±0.003	0.653±0.004	0.837±0.003	0.920±0.004
FAIRCO	0.701±0.003	0.731±0.005	**0.675±0.005**	0.841±0.004	0.927±0.006
PDA	**0.708±0.001**	0.733±0.003	0.668±0.004	0.844±0.004	**0.929±0.006**
DICE	0.700±0.001	0.728±0.005	0.658±0.005	0.856±0.003	0.920±0.005
DMBR	**0.707±0.001**	0.735±0.003	0.672±0.003	0.843±0.002	**0.930±0.004**
BAGA	**0.710±0.002**	**0.749±0.004**	**0.681±0.002**	**0.869±0.004**	**0.933±0.003**
BAGA-*mlp*	0.697±0.001	0.724±0.002	0.663±0.005	0.836±0.004	0.921±0.002

Table 2: Offline results of each comparing method in terms of AUC (the larger the better), where the best performance on each data set is shown in boldface (pairwise t-test at 0.05 significance level).

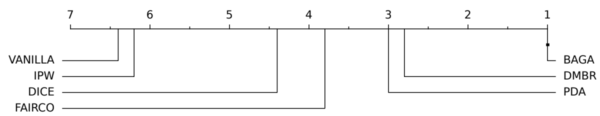

Figure 4: Comparison of BAGA (control method) against other methods with *Bonferroni-Dunn test* in terms of AUC.

Evaluation metric	τ_F	critical value
AUC	50.9999	2.1620

Table 3: Friedman statistics τ_F in terms of AUC as well as the critical value at 0.05 significance level ((# comparing methods $k = 7$; # data sets $N = 25$).

an unbiased scenario. In the marketing coupon recommendation scenarios, only one coupon is offered to the user at a time, resulting in a binary classification problem that predicts the likelihood of a user using a coupon. Ranking performance metrics commonly used in other recommendation scenarios would be inappropriate for this problem, alternative evaluation metrics were utilized. In this paper, we use Area Under Curve (AUC) to measure the efficiency of alleviating matching bias in a simulated uniform distribution. Conceptually, a higher AUC score means that the debiasing method is more effective.

Debiasing Performance Comparison (A1 & A2)

Table 2 demonstrates the AUC performance of each compared approach in the simulated uniform (unbiased) distribution, where the vanilla model is DEEPFM. Five-fold cross-validation is conducted, and the mean metric values and standard deviations for each approach are recorded. A pairwise t-test at 0.05 significance level is also performed for the five-fold cross-validation to show whether the performance of BAGA is significantly different from the compared approach. The conclusions of the experiments on other vanilla molds are generally consistent with DEEPFM.

In addition, the widely-used Friedman test (Demšar 2006) is utilized to statistically compare multiple methods on multiple datasets [6]. The Friedman statistic critical value τ_F is 2.1620 according to (Demšar 2006). Table 3 reports the Friedman statistics τ_F of evaluation metrics at a 0.05 significance level. The τ_F value exceeds the critical value, indicating that the null hypothesis of "equal" performance among comparing methods should be clearly rejected. Therefore, *Bonferroni-Dunn test* (Demšar 2006) is employed as the post-hoc statistical test to analyze the relative performance improvement. Here, the difference between the average ranks of the control method (i.e. BAGA) and one comparing method is calibrated with the *critical difference* (CD), and the performance difference is considered significant if their average ranks differ by at least one CD, here $CD = 1.6118$.

Fig.4 illustrates the CD diagram on the AUC metric, with BAGA serving as the control method. The average rank of each method is marked on the axis, with lower ranks being to the right. Comparing methods whose average ranks are within one CD of BAGA are interconnected with thick lines. Otherwise, it is considered to have a significantly different performance against BAGA.

Based on the reported experimental results, the following observations can be made:

- IPW achieves comparable performance over the vanilla models on all datasets. It can be deduced that simply re-weighting cannot precisely estimate the effect of the matching bias in marketing coupon recommendations.
- FAIRCO, DICE, DMBR, and PDA achieve marginal improvements over the vanilla models on almost all datasets. In particular, for Ds_C, FAIRCO and BAGA achieve the best performance, for Ds_A and Ds_D, DMBR, PDA, and BAGA achieve the best performance.
- BAGA effectively improves the performance over the vanilla models. Although BAGA does not significantly outperform some of the comparing methods in a pairwise t-test at the 0.05 significance level, it still outperformed all of the comparing methods in terms of relative performance improvement.
- As shown in Fig.4, BAGA has the lowest average rank on the AUC evaluation metric. That is the performance of BAGA is statistically globally optimal.

In summary, the above results validate the effectiveness of BAGA for debiased marketing coupon recommendations.

[6]We tested five vanilla models on five datasets, so the number of statistics datasets here is equivalent to 25.

Methods	Training	Testing
DeepFm	$\mathcal{O}\big(\mathcal{F}_D(m,d)\big)$	$\mathcal{O}\big(\mathcal{F'}_D(d)\big)$
Ipw	$\mathcal{O}\big(\mathcal{F}_D(m,d)\big)$	$\mathcal{O}\big(\mathcal{F'}_D(d)\big)$
FairCo	$\mathcal{O}\big(\mathcal{F}_D(m,d)\big)$	$\mathcal{O}\big(\mathcal{F'}_D(d)\big)$
Pda	$\mathcal{O}\big(\mathcal{F}_D(m,d)\big)$	$\mathcal{O}\big(\mathcal{F'}_D(d)\big)$
Dice	$\mathcal{O}\big(2*\mathcal{F}_D(m,d)\big)$	$\mathcal{O}\big(\mathcal{F'}_D(d)\big)$
Dmbr	$\mathcal{O}\big(\mathcal{F}_D(m,d+t)+\mathcal{F}_F(m,d,2t)\big)$	$\mathcal{O}\big(\mathcal{F'}_D(d+t)+\mathcal{F'}_F(d,2t)\big)$
Baga	$\mathcal{O}\big(\mathcal{F}_D(m,d)+2n*\mathcal{F}_M(\frac{m}{n},d)\big)$	$\mathcal{O}\big(\mathcal{F'}_D(d)+2*\mathcal{F}_M(d)\big)$

Table 4: Summary of the algorithmic complexity of comparing methods with vanilla model DeepFm.

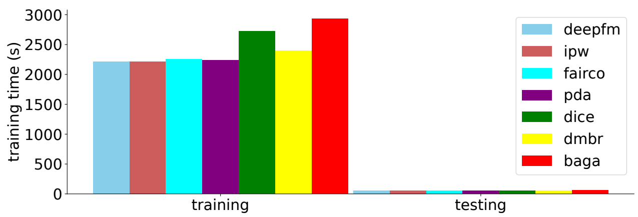

Figure 5: Mean of five repeated running times (training/testing) for each method on Ds_C under the same environmental settings. The height of the bar graph indicates running times, and the y-axis represents running time in seconds.

Futher Analysis (A3 & A4)

Ablation study In this section, we dissect Baga to answer the question of how the CGC model architecture contributes to it. Therefore, we study the effect of replacing components to provide additional insights. Specifically, we replace the CGC model with multi-task MLP layers denoted as Baga-mlp and conduct experiments on it to demonstrate the effect of the CGC model architecture. The experiment results are shown in Table 2. We can observe that Baga significantly outperforms Baga-mlp in all datasets. In the Baga-mlp approach, the network for sub-tasks with smaller sample sizes will be poorly optimized. Whereas in the CGC model, with the help of the gating network and shared expert network, the network of sub-tasks with smaller samples can enjoy the gain from global optimization and thus get better optimization.

Complexity and Time Analysis Table 4 summarizes the algorithmic complexity of comparing methods with vanilla model DeepFm, w.r.t. several common factors, i.e., m (# training examples), d (# features), and t (# additional features). Furthermore, to account for specific algorithmic components employed by each method, the following terms are introduced in complexity characterization: $\mathcal{F}_D(m,d)$ and $\mathcal{F'}_D(d)$ represent the training and testing complexity for vanilla model DeepFm; $\mathcal{F}_M(m,d)$ and $\mathcal{F'}_M(d)$ represent the training and testing complexity for the expert, gate, task layer combination; $\mathcal{F}_F(m,d)$ and $\mathcal{F'}_F(d)$ represent the training and testing complexity for a Fm layer.

Furthermore, Fig.5 illustrates the training and testing time of each method on Ds_C. Generally, the neurons of each layer in the vanilla model DeepFm and the CGC model are 128*32*1 and 64, respectively. We implement it with Tensorflow (Abadi et al. 2015) and use Adam (Kingma and Ba 2014) optimizer with the default setting to train the model on a PC with Intel i7 six cores 2.6GHz CPU and 16GB memory, and the operating environment is Python 3.7.12. As shown in Fig.5, the empirical training and testing time of Baga is longer compared to other methods, but within an acceptable range. The gap will be further narrowed when performing large-scale distributed training in a real-world environment.

Methods	Send-Number	Use-Number	Use-Rate
Baseline	14,600,310	69,556	4.76%
Baga	14,609,694	70,673[†]	4.83%[†]
%improv.		**+1.60%**	**+1.54%**

Table 5: Online results between Baga and the baseline, where † means better performance.

Online Results (A5)

To assess the effectiveness of our proposed debiasing method (Baga based on DeepFm) in real-world scenarios, we conducted an online A/B test in a marketing coupon recommendation campaign, which two metrics are measured: the number of coupons used and the use rate. Due to budget constraints, we only compared Baga with the baseline DeepFm in the online coupon recommendation system, which randomly and evenly divided all candidates into two buckets. The experimental results are summarized in Table 5, that revealed our proposed approach Baga achieved a 1.54% increase in use rate and a 1.60% increase in the number of coupons used, demonstrating a significant improvement in the real-world marketing coupon recommendation campaign. These results further validate the effectiveness of our proposed debiasing method Baga.

Conclusion

In this paper, we explore how to alleviate the matching bias in marketing coupon recommendations and present a new de-confounded training paradigm Baga. To be specific, Baga segments the training data into different user and coupon groups based on marketing tactics in an adaptive manner, and separately models each sub-task to capture the genuine causality relationship between user-coupon representations. The results of both offline experiments and online A/B testing demonstrate the effectiveness of Baga in alleviating matching bias and improving accuracy.

References

Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Goodfellow, I.; Harp, A.; Irving, G.; Isard, M.; Jia, Y.; Jozefowicz, R.; Kaiser, L.; Kudlur, M.; Levenberg, J.; Mané, D.; Monga, R.; Moore, S.; Murray, D.; Olah, C.; Schuster, M.; Shlens, J.; Steiner, B.; Sutskever, I.; Talwar, K.; Tucker, P.; Vanhoucke, V.; Vasudevan, V.; Viégas, F.; Vinyals, O.; Warden, P.; Wattenberg, M.; Wicke, M.; Yu, Y.; and Zheng, X. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org.

Abdollahpouri, H.; and Mansoury, M. 2020. Multi-sided Exposure Bias in Recommendation.

Bengio, Y.; Ducharme, R.; and Vincent, P. 2000. A neural probabilistic language model. *Advances in neural information processing systems*, 13.

Biega, A. J.; Gummadi, K. P.; and Weikum, G. 2018. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*, 405–414.

Bonner, S.; and Vasile, F. 2018. Causal embeddings for recommendation. In *Proceedings of the 12th ACM conference on recommender systems*, 104–112.

Chaney, A. J.; Stewart, B. M.; and Engelhardt, B. E. 2018. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM conference on recommender systems*, 224–232.

Chen, J.; Dong, H.; Wang, X.; Feng, F.; Wang, M.; and He, X. 2023. Bias and Debias in Recommender System: A Survey and Future Directions. *ACM Trans. Inf. Syst.*, 41(3).

Demšar, J. 2006. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine learning research*, 7: 1–30.

Fang, J.; Cui, Q.; Zhang, G.; Tang, C.; Gu, L.; Li, L.; Gu, J.; Zhou, J.; and Wu, F. 2023. Alleviating Matching Bias in Marketing Recommendations. In *Proceedings of the 46rd International ACM SIGIR Conference on Research and Development in Information Retrieval*.

Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.

Grimes, D. A.; and Schulz, K. F. 2002. Bias and causal associations in observational research. *The lancet*, 359(9302): 248–252.

Guo, H.; TANG, R.; Ye, Y.; Li, Z.; and He, X. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 1725–1731.

Harper, F. M.; and Konstan, J. A. 2015. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4): 1–19.

He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, 173–182.

He, X.; Zhang, Y.; Feng, F.; Song, C.; Yi, L.; Ling, G.; and Zhang, Y. 2023. Addressing confounding feature issue for causal recommendation. *ACM Transactions on Information Systems*, 41(3): 1–23.

Huan, Z.; Zhang, G.; Zhang, X.; Zhou, J.; Wu, Q.; Gu, L.; Gu, J.; He, Y.; Zhu, Y.; and Mo, L. 2022. An Industrial Framework for Cold-Start Recommendation in Zero-Shot Scenarios. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, 3403–3407. ACM.

Huangfu, Z.; Zhang, G.; Wu, Z.; Wu, Q.; Zhang, Z.; Gu, L.; Zhou, J.; and Gu, J. 2022. A Multi-Task Learning Approach for Delayed Feedback Modeling. *Companion Proceedings of the Web Conference 2022*.

Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *CoRR*, abs/1412.6980.

Li, C.; Chu, M.; Zhou, C.; and Zhao, L. 2020. Two-period discount pricing strategies for an e-commerce platform with strategic consumers. *Computers & Industrial Engineering*, 147: 106640.

Li, H.; Zheng, C.; and Wu, P. 2022. StableDR: Stabilized Doubly Robust Learning for Recommendation on Data Missing Not at Random. In *The Eleventh International Conference on Learning Representations*.

Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.

Liu, D.; Cheng, P.; Dong, Z.; He, X.; Pan, W.; and Ming, Z. 2020. A general knowledge distillation framework for counterfactual recommendation via uniform data. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 831–840.

Mansournia, M. A.; and Altman, D. G. 2016. Inverse probability weighting. *Bmj*, 352.

Marlin, B. M.; Zemel, R. S.; Roweis, S.; and Slaney, M. 2007a. Collaborative Filtering and the Missing at Random Assumption. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, UAI'07, 267–275. Arlington, Virginia, USA: AUAI Press. ISBN 0974903930.

Marlin, B. M.; Zemel, R. S.; Roweis, S.; and Slaney, M. 2007b. Collaborative filtering and the missing at random assumption. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, 267–275.

Morik, M.; Singh, A.; Hong, J.; and Joachims, T. 2020. Controlling fairness and bias in dynamic learning-to-rank. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 429–438.

Mutanen, T.; Nousiainen, S.; and Liang, H. 2010. Personalized recommendation on discount coupons. In *Annu. Int. Conf. Comput. Sci. Educ. Innov. Technol*, 127–132.

Pearl, J. 2009. *Causality*. Cambridge university press.

Rendle, S. 2010. Factorization machines. In *2010 IEEE International conference on data mining*, 995–1000. IEEE.

Saito, Y.; Yaginuma, S.; Nishino, Y.; Sakata, H.; and Nakata, K. 2020. Unbiased recommender learning from missing-not-at-random implicit feedback. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, 501–509.

Schnabel, T.; Swaminathan, A.; Singh, A.; Chandak, N.; and Joachims, T. 2016. Recommendations as treatments: Debiasing learning and evaluation. In *international conference on machine learning*, 1670–1679. PMLR.

Seaman, S. R.; and White, I. R. 2013. Review of inverse probability weighting for dealing with missing data. *Statistical methods in medical research*, 22(3): 278–295.

Shore, J.; and Johnson, R. 1981. Properties of cross-entropy minimization. *IEEE Transactions on Information Theory*, 27(4): 472–482.

Sim, W.-I.; and Lee, S.-Y. 2014. Customized Coupon Recommendation Model based on Fuzzy AHP Reflecting User Preference. *Journal of Digital Convergence*, 12(5): 395–401.

Sun, Y.-X.; Zhang, Y.-L.; Wang, W.; Li, L.; and Zhou, J. 2023. Treatment Effect Estimation across Domains. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2352–2361.

Tang, C.; Wang, H.; Li, X.; Cui, Q.; Li, L.; and Zhou, J. 2023. Difference-in-Differences Meets Tree-based Methods: Heterogeneous Treatment Effects Estimation with Unmeasured Confounding. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 33792–33803. PMLR.

Tang, C.; Wang, H.; Li, X.; Cui, Q.; Zhang, Y.-L.; Zhu, F.; Li, L.; Zhou, J.; and Jiang, L. 2022. Debiased Causal Tree: Heterogeneous Treatment Effects Estimation with Unmeasured Confounding. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 5628–5640. Curran Associates, Inc.

Tang, H.; Liu, J.; Zhao, M.; and Gong, X. 2020. Progressive layered extraction (ple): A novel multi-task learning (mtl) model for personalized recommendations. In *Proceedings of the 14th ACM Conference on Recommender Systems*, 269–278.

Tianchi. 2018. Alibaba Ads Click Dataset.

Tianchi. 2022. RecELE.ME.

VanderWeele, T. J.; Hernán, M. A.; and Robins, J. M. 2008. Causal directed acyclic graphs and the direction of unmeasured confounding bias. *Epidemiology (Cambridge, Mass.)*, 19(5): 720.

Wang, W.; Feng, F.; He, X.; Zhang, H.; and Chua, T.-S. 2021a. Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1288–1297.

Wang, X.; Jin, H.; Zhang, A.; He, X.; Xu, T.; and Chua, T.-S. 2020. Disentangled graph collaborative filtering. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, 1001–1010.

Wang, X.; Zhang, R.; Sun, Y.; and Qi, J. 2019. Doubly robust joint learning for recommendation on data missing not at random. In *International Conference on Machine Learning*, 6638–6647. PMLR.

Wang, X.; Zhang, R.; Sun, Y.; and Qi, J. 2021b. Combating selection biases in recommender systems with a few unbiased ratings. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, 427–435.

Wang, Z.; She, Q.; and Zhang, J. 2021. MaskNet: Introducing Feature-Wise Multiplication to CTR Ranking Models by Instance-Guided Mask. *ArXiv*, abs/2102.07619.

Wei, T.; Feng, F.; Chen, J.; Wu, Z.; Yi, J.; and He, X. 2021. Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 1791–1800.

Zhang, Y.; Bai, Y.; Chang, J.; Zang, X.; Lu, S.; Lu, J.; Feng, F.; Niu, Y.; and Song, Y. 2023. Leveraging Watch-Time Feedback for Short-Video Recommendations: A Causal Labeling Framework. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, 4952–4959. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701245.

Zhang, Y.; Feng, F.; He, X.; Wei, T.; Song, C.; Ling, G.; and Zhang, Y. 2021. Causal intervention for leveraging popularity bias in recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 11–20.

Zhao, H.; Cui, Q.; Li, X.; Bao, R.; Li, L.; Zhou, J.; Liu, Z.; and Feng, J. 2023. MDI: A Debiasing Method Combining Unbiased and Biased Data. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 3280–3284.

Zheng, Y.; Gao, C.; Li, X.; He, X.; Li, Y.; and Jin, D. 2021a. Disentangling User Interest and Conformity for Recommendation with Causal Embedding. In *Proceedings of the Web Conference 2021*, 2980–2991.

Zheng, Y.; Gao, C.; Li, X.; He, X.; Li, Y.; and Jin, D. 2021b. Disentangling User Interest and Conformity for Recommendation with Causal Embedding. In *Proceedings of the Web Conference 2021*, WWW '21, 2980–2991. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383127.

Zhu, F.; Zhong, M.; Yang, X.; Li, L.; Lu, Y.; Zhang, T.; Zhou, J.; Chen, C.; Wu, F.; Liu, G.; et al. 2023. DCMT: a direct entire-space causal multi-task framework for post-click conversion estimation. In *The 39th IEEE International Conference on Data Engineering*.