

AutoDebias: Learning to Debias for Recommendation

Jiawei Chen^{1*}, Hande Dong^{1*}, Yang Qiu¹, Xiangnan He^{1†},
Xin Xin², Liang Chen³, Guli Lin⁴, Keping Yang⁴.

¹University of Science and Technology of China, ²University of Glasgow, ³Sun Yat-Sen University, ⁴Alibaba Group.
{cjwustc@,donghd@mail.,yangqiu@mail.,hexn@}ustc.edu.cn
x.xin.1@research.gla.ac.uk,chenliang6@mail.sysu.edu.cn,{guli.ling1,shaoyao}@taobao.com.

ABSTRACT

Recommender systems rely on user behavior data like ratings and clicks to build personalization model. However, the collected data is observational rather than experimental, causing various biases in the data which significantly affect the learned model. Most existing work for recommendation debiasing, such as the inverse propensity scoring and imputation approaches, focuses on one or two specific biases, lacking the universal capacity that can account for mixed or even unknown biases in the data.

Towards this research gap, we first analyze the origin of biases from the perspective of *risk discrepancy* that represents the difference between the expectation empirical risk and the true risk. Remarkably, we derive a general learning framework that well summarizes most existing debiasing strategies by specifying some parameters of the general framework. This provides a valuable opportunity to develop a universal solution for debiasing, e.g., by learning the debiasing parameters from data. However, the training data lacks important signal of how the data is biased and what the unbiased data looks like. To move this idea forward, we propose *AotoDebias* that leverages another (small) set of uniform data to optimize the debiasing parameters by solving the bi-level optimization problem with meta-learning. Through theoretical analyses, we derive the generalization bound for AutoDebias and prove its ability to acquire the appropriate debiasing strategy. Extensive experiments on two real datasets and a simulated dataset demonstrated effectiveness of AutoDebias. The code is available at <https://github.com/DongHande/AutoDebias>.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommendation, Bias, Debias, Meta-learning

* Jiawei Chen and Hande Dong contribute equally to the work.
† Xiangnan He is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '21, July 11–15, 2021, Virtual Event, Canada
© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8037-9/21/07...\$15.00
<https://doi.org/10.1145/3404835.3462919>

ACM Reference Format:

Jiawei Chen, Hande Dong, Yang Qiu, Xiangnan He, Xin Xin, Liang Chen, Guli Lin and Keping Yang. 2021. AutoDebias: Learning to Debias for Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3404835.3462919>

1 INTRODUCTION

Being able to provide personalized suggestions to each user, recommender systems (RS) have been widely used in countless online applications. Recent years have witnessed flourishing publications on recommendation, most of which aim at inventing machine learning model to fit user behavior data [17, 44, 52]. However, these models may be deteriorated in real-world RS, as the behavior data is often full of biases. In practice, the data is observational rather than experimental, and is often affected by many factors, including but not limited to self-selection of the user (*selection bias*) [19, 33], exposure mechanism of the system (*exposure bias*) [34, 47], public opinions (*conformity bias*) [25, 29] and the display position (*position bias*) [22, 23]. These biases make the data deviate from reflecting user true preference. Hence, blindly fitting data without considering the data biases would yield unexpected results, e.g., amplifying the long-tail effect [1] and previous-model bias [28].

Given the wide existence of data biases and their large impact on the learned model, we cannot emphasize too much the importance of properly debiasing for practical RS. Existing efforts on recommendation (or learning-to-rank) biases can be divided into three major categories: 1) data imputation [19, 42], which assigns pseudo-labels for missing data to reduce variance, 2) inverse propensity scoring (IPS) [39, 47], a counterfactual technique that reweights the collected data for an expectation-unbiased learning, and 3) generative modeling [27], which assumes the generation process of data and reduces the biases accordingly. Despite their effectiveness in some scenarios, we argue that they suffer from two limitations:

- **Lacking Universality.** These methods are designed for addressing one or two biases of a specific scenario, e.g., IPS for selection bias [39], click model for position bias [12]. Thus, when facing real data that commonly contains multiple types of biases, these methods will fall short.
- **Lacking Adaptivity.** The effectiveness of these methods is guaranteed only when the debiasing configurations (e.g., pseudo-labels, propensity scores, or data-generating process) are properly specified. However, obtaining such proper configurations is quite difficult, requiring domain expertise that thoroughly understands the biases in the data and how they affect the model. Even worse, the optimal configurations may evolve with time as

new users/items/interactions may change the data distribution, which has been a nightmare for practitioners to manually tune the configurations continually.

Considering the shortcomings of existing work, we believe it is essential to develop a universal debiasing solution, which not only accounts for multiple biases and their combinations, but also frees human efforts to identify biases and tune the configurations. To achieve this goal, we first review the common biases and debiasing strategies, offering two important insights: (1) Although different biases have different origins and properties, they all can be formulated as the *risk discrepancy* between the empirical risk and the true risk, resulting from the inconsistency between the distribution for which the training data is collected and the one used for unbiased test; (2) The success of most recent debiasing strategies can be attributed to their specific configurations to offset the discrepancy for model training. Based on the insights, we propose a general debiasing framework by reducing the risk discrepancy, which subsumes most debiasing strategies — each strategy can be recovered by specifying the parameters of the framework. This framework provides a valuable opportunity to develop a universal debiasing solution for recommendation — we can perform automatic debiasing by learning the debiasing parameters of the framework.

Now the question lies in how to optimize the debiasing parameters. Obviously, the biased training data lacks important signals of how the data is biased and what the unbiased data looks like. To deal with this problem, we propose to leverage another *uniform data* to supervise the learning of debiasing parameter. The uniform data is assumed to be collected by a random logging policy [39], reflecting user preference in an unbiased way. We make full use of this important evidence, optimizing the debiasing parameters by minimizing the loss on the uniform data. Specifically, we formulate the process as a bi-level optimization problem, where the debiasing parameters serve as the hyper-parameters for learning the recommender model, and optimize debiasing parameters by meta-learning technique [14]. We conduct theoretical analyses on the learning framework, proving that: (1) the optimum learned under such objective is approximate to the best case where biases are properly corrected; (2) it is able to learn a satisfactory debiasing strategy even if it is trained on a small uniform data.

Lastly, in terms of leveraging uniform data for recommendation, the most relevant work is the recently proposed KDRRec [28]. However, we argue that it does not sufficiently exploit the merits of uniform data. KDRRec trains a separate teacher model on the uniform data, and then transfers the model’s knowledge to the normal training on biased data. Since uniform data is collected at the expense of degrading user experience, its size is usually rather small. As such, the model trained on it suffers from high variance, decreasing the effectiveness of KDRRec. What’s more, it lacks theoretical guarantees, and the inherent mechanism of how teacher model offsets the bias is not entirely understood. Compared to KDRRec, our framework utilizes uniform data in a more theoretically sound way and yields significant empirical improvements.

In a nutshell, this work makes the following main contributions:

- Unifying various biases from the *risk discrepancy* perspective and developing a general debiasing framework that subsumes most debiasing strategies.

Table 1: Notations and Definitions.

Notations	Annotations
D_T	a biased training set with entries $\{(u_k, i_k, r_k)\}_{1 \leq k \leq D_T }$ collected from user interaction history
D_U	a unbiased uniform set with entries $\{(u_l, i_l, r_l)\}_{1 \leq l \leq D_U }$ collected with uniform logging policy
$p_T(u, i, r)$	the data distribution for which D_T collected
$p_U(u, i, r)$	the ideal unbiased data distribution
$f_\theta(\cdot, \cdot)$	RS model that maps a user-item pair into the prediction
$\hat{L}_T(f)$	the empirical risk of f on D_T
$L(f)$	the true risk of f
S_1	the subspace of $\mathcal{U} \times \mathcal{I} \times \mathcal{R}$ with constraints: $p_T(u, i, k) > 0, p_U(u, i, k) > 0$
S_0	the subspace of $\mathcal{U} \times \mathcal{I} \times \mathcal{R}$ with constraints: $p_T(u, i, k) = 0, p_U(u, i, k) > 0$

- Proposing a new method that leverages uniform data to learn optimal debiasing strategy with theoretical guarantees.
- Conducting experiments on three types of data (explicit and implicit feedback, and simulated data of list feedback) to validate the effectiveness of our proposal.

2 PRELIMINARY

In this section, we formulate the recommendation task and review various biases in it from *risk discrepancy* perspective.

2.1 Task Formulation

Suppose we have a recommender system with a user set \mathcal{U} and an item set \mathcal{I} . Let u (or i) denotes a user (or an item) in \mathcal{U} (or \mathcal{I}). Let $r \in \mathcal{R}$ denotes the feedback (e.g., rating values, clicks, and retention time) given by a user to an item. The collected history behavior data D_T can be notated as a set of triplets $\{(u_k, i_k, r_k)\}_{1 \leq k \leq |D_T|}$ generated from an unknown distribution $p_T(u, i, r)$ over user-item-label space $\mathcal{U} \times \mathcal{I} \times \mathcal{R}$. The task of a recommendation system can be stated as follows: learning a recommendation model from D_T so that it can capture user preference and make a high-quality recommendation. Formally, let $\delta(\cdot, \cdot)$ denotes the error function between the prediction and the ground truth label. The goal of recommendation is to learn a parametric function $f_\theta : \mathcal{U} \times \mathcal{I} \rightarrow \mathcal{R}$ from the available dataset D_T to minimize the following *true risk*:

$$L(f) = \mathbb{E}_{p_U(u, i, r)} [\delta(f(u, i), r)], \quad (1)$$

where f_θ can be implemented by a specific recommendation model with a learnable parameters θ . We remark that in this paper we may simply omit subscript ‘ θ ’ in the notations for clear presentation. $p_U(u, i, r)$ denotes the ideal unbiased data distribution for model testing. This distribution can be factorized as the product of the user-item pair distribution $p_U(u, i)$ (often supposed as uniform) and the factual preference distribution for each user-item pair $p_U(r|u, i)$.

Since the true risk is not accessible, the learning is conducted on the training set D_T by optimizing the following *empirical risk*:

$$\hat{L}_T(f) = \frac{1}{|D_T|} \sum_{k=1}^{|D_T|} \delta(f(u_k, i_k), r_k). \quad (2)$$

If the empirical risk $\hat{L}_T(f)$ is an unbiased estimator of the true risk $L(f)$, i.e., $\mathbb{E}_{p_T}[\hat{L}_T(f)] = L(f)$, the PAC learning theory [16] states




Figure 1: An example of the biased empirical risk.

that the learned model will be approximately optimal if we have sufficiently large training data.

2.2 Biases in Recommendation

However, as various biases occur in real-world data collection, the training data distribution p_T is often inconsistent with the ideal unbiased distribution p_U . Training data only gives a skewed snapshot of user preference, making the recommendation model sink into sub-optimal result. Figure 1 illustrates this phenomenon. The red curve denotes the true risk function, while the blue curve denotes the expectation of the biased empirical risk function $\mathbb{E}_{p_T}[L_T(f)]$. As the two risks are expected over different distribution, they will behave rather differently even in their optimum (i.e., f^* versus f^T). It means that even if a sufficiently large training set is provided and the model arrives at empirical optimal point f^T , there still exists a gap ΔL between the optimum $L(f^*)$ and the empirical one $L(f^T)$. Above analysis reveals the impact of bias on recommendation: bias will incur the discrepancy between the true risk and the expected empirical risk. Blindly fitting a recommendation model without considering the risk discrepancy will result in inferior performance.

Recent work has identified various biases in recommendation. In this subsection, we will review these biases from risk discrepancy perspective to help the readers to better understand their properties and negative effect. We refer to [5] and categorize data biases into the following four classes:

Selection bias happens as users are free to choose which items to rate, so that the observed ratings are not a representative sample of all ratings [5]. Selection bias can be easily understood from the risk discrepancy perspective — it skews the user-item pair distribution $p_T(u, i)$ from the ideal uniform one $p_U(u, i)$. Typically, $p_T(u, i)$ inclines to the pairs with high rating values. Learning a recommendation model under such biased data will easily overestimate the preferences of users to items.

Conformity bias happens as users tend to behave similarly to the others in a group, even if doing so goes against their own judgment [5]. Conformity bias distorts label distribution $p_T(r|u, i)$ in conformity to the public opinions, making the feedback do not always signify user true preference, i.e., $p_T(r|u, i) \neq p_U(r|u, i)$.

Exposure bias happens in implicit feedback data as users are only exposed to a part of specific items [5]. It would be difficult to understand exposure bias from the above definition, but straightforward from risk discrepancy perspective. On the one hand, users generate behaviors on exposed items, making the observed user-item distribution $p_T(u, i)$ deviate from the ideal one $p_U(u, i)$. On the other hand, implicit feedback data only has positive feedback observed $p_T(r = 1|u, i) = 1$. Such positive-only data will cause ambiguity in the interpretation of unobserved interactions — they may be caused by non-exposure or by dislike.




Figure 2: An illustration of the fact that training data distribution only covers part of ideal data distribution. We need impute pseudo-data (marked as diamond) to the blank region S_0 , where the mark size reflects their weights.

Position bias happens as users tend to interact with items in higher position of the recommendation list [5]. Under position bias, training data distribution $p_T(u, i, r)$ will be sensitive to the item display position and fail to reflect user preference faithfully. On the one hand, the ranking position will affect the chance of the item exposure to the user [34], i.e., $p_T(u, i) \neq p_U(u, i)$. On the other hand, as users often trust the recommendation system, their judgments also will be affected by the position, i.e., $p_T(r|u, i) \neq p_U(r|u, i)$.

Generally speaking, above biases can be summarized as a type of risk discrepancy — they cause training data distribution $p_T(u, i, r)$ deviate from the ideal unbiased one $p_U(u, i, r)$. This insight motivates us to develop a powerful framework that directly conquers the risk discrepancy, enabling the elimination of the mixture of above biases or even unknown biases in the data.

3 A GENERAL DEBIASING FRAMEWORK

In this section, we provide a general debiasing framework that can account for various kinds of biases in recommendation data. We then discuss how it subsumes most existing debiasing strategies.

3.1 Debiasing Empirical Risk

The analyses presented in Section 2.2 show that various biases account for the risk discrepancy. For unbiased learning, we need to re-design the empirical risk function so that its expectation under biased training distribution is consistent with the true risk. By comparing $L(f)$ with $\mathbb{E}_{p_T}[\hat{L}_T(f)]$, we can see that the discrepancy origins from the data distribution, where the effect of each data has been skewed in empirical risk. To deal with it, we can re-weight training data and obtain a reweighted empirical risk function:

$$\hat{L}_T(f|w^{(1)}) = \frac{1}{|D_T|} \sum_{k=1}^{|D_T|} w_k^{(1)} L(f(u_k, i_k), r_k). \quad (3)$$

When the parameter $w^{(1)}$ is properly specified, i.e., $w_k^{(1)} = \frac{p_U(u_k, i_k, r_k)}{p_T(u_k, i_k, r_k)}$, the skewness of each sampled training data is corrected. Such empirical risk $\hat{L}_T(f|w^{(1)})$ seems to be an unbiased estimation of the true risk as:

$$\mathbb{E}_{p_T}[\hat{L}_T(f|w^{(1)})] = \sum_{(u, i, r) \in S_1} p_U(u, i, r) \delta(f(u, i), r), \quad (4)$$

where S_1 denotes the user-item-label subspace satisfying $p_T(u, i, r) > 0$, $p_U(u, i, r) > 0$, i.e., $S_1 \equiv \{(u, i, r) \in \mathcal{U} \times \mathcal{I} \times \mathcal{R} : p_T(u, i, r) > 0, p_U(u, i, r) > 0\}$. However, the equivalence of $\mathbb{E}_{p_T}[\hat{L}_T(f|w^{(1)})]$

with $L(f)$ is actually not held. As illustrated in Figure 2, the training data distribution p_T only covers a part of regions (*i.e.*, S_1) of the support of p_U , while does not have probability in other regions ($S_0 \equiv \{(u, i, r) \in U \times I \times R : p_T(u, i, k) = 0, p_U(u, i, k) > 0\}$). It means that even if a sufficiently large training data is available, it still provides partial user-item pairs only, while leaves the rest blank. Learning on S_1 only will suffer especially when S_0 and S_1 exhibit different patterns. This situation is common in practice: the system tends to display popular items [24], while a considerable number of unpopular items have little chance to be exposed. To deal with this problem, we need to impute pseudo-data to the blank region. In fact, the equivalent transformation of $L(f)$ is:

$$\begin{aligned} L(f) &= \sum_{u \in \mathcal{U}, i \in \mathcal{I}, r \in \mathcal{R}} p_U(u, i, r) \delta(f(u, i), r) \\ &= \sum_{(u, i, r) \in S_1} p_U(u, i, r) \delta(f(u, i), r) + \sum_{(u, i, r) \in S_0} p_U(u, i, r) \delta(f(u, i), r) \\ &= E_{p_T} \left[\frac{1}{|D_T|} \sum_{k=1}^{|D_T|} w_k^{(1)} \delta(f(u_k, i_k), r_k) \right] + \sum_{u \in \mathcal{U}, i \in \mathcal{I}} w_{ui}^{(2)} \delta(f(u, i), m_{ui}). \end{aligned} \quad (5)$$

The last equation holds when the parameters $\phi \equiv \{w^{(1)}, w^{(2)}, m\}$ are set as:

$$\begin{aligned} w_k^{(1)} &= \frac{p_U(u_k, i_k, r_k)}{p_T(u_k, i_k, r_k)} \\ w_{ui}^{(2)} &= \sum_{r \in \mathcal{R}} p_U(u, i, r) \mathbf{I}[p_T(u, i, r) = 0] \\ m_{ui} &= E_{p_U(r|u, i)} [r \mathbf{I}[p_T(u, i, r) = 0]], \end{aligned} \quad (6)$$

where $\mathbf{I}[\cdot]$ denotes indicator function. Here we absorb the expectation over $\delta(\cdot, \cdot)$ into the pseudo-label, *i.e.*, $\mathbb{E}[\delta(\cdot, \cdot)] = \delta(\cdot, \mathbb{E}[\cdot])$. It holds for many commonly-used loss functions, such as L2, L1 and cross-entropy. Thus, we define the following empirical risk function as:

$$\hat{L}_T(f|\phi) = \frac{1}{|D_T|} \sum_{k=1}^{|D_T|} w_k^{(1)} \delta(f(u_k, i_k), r_k) + \sum_{u \in \mathcal{U}, i \in \mathcal{I}} w_{ui}^{(2)} \delta(f(u, i), m_{ui}), \quad (7)$$

which is an unbiased estimator of the true risk when the parameters are properly specified. We remark that there may exist multiple solutions of ϕ but at least one (*i.e.*, Equation (6)) that makes equation $\hat{L}_T(f|\phi) = L(f)$ hold. This generic empirical risk function provides a valuable opportunity to develop a universal solution – achieving automatic debiasing by learning the debiasing parameters ϕ .

3.2 Link to Related Work

To show the universality of our framework, we review representative debiasing strategies and discuss how the framework subsumes them.

Selection bias. Existing methods are mainly three types:

(1) Inverse Propensity Score (IPS) [39] reweights the collected data for an unbiased learning, defining the empirical risk as:

$$\hat{L}_{IPS}(f) = \frac{1}{|\mathcal{U}||\mathcal{I}|} \sum_{k=1}^{|D_T|} \frac{1}{q_{u_k i_k}} \delta(f(u_k, i_k), r_k), \quad (8)$$

where $q_{u_k i_k}$ is defined as propensity, which estimates the probability of the data to be observed. The framework recovers it by setting $w_k^{(1)} = \frac{|D_T|}{q_{u_k i_k} |\mathcal{U}||\mathcal{I}|}$, $w_k^{(2)} = 0$.

(2) Data Imputation [19] assigns pseudo-labels for missing data and optimizes the following risk function:

$$\hat{L}_{IM}(f) = \frac{1}{|\mathcal{U}||\mathcal{I}|} \left(\sum_{k=1}^{|D_T|} \delta(f(u_k, i_k), r_k) + \sum_{u \in \mathcal{U}, i \in \mathcal{I}} \lambda \delta(f(u, i), m_{ui}) \right), \quad (9)$$

where m_{ui} denotes the imputed labels which can be specified heuristically [21] or inferred by a dedicated model [7, 32, 37]. λ is to control the contribution of the imputed data. It is a special case of our framework if we set $w_k^{(1)} = \frac{|D_T|}{|\mathcal{U}||\mathcal{I}|}$, $w_{ui}^{(2)} = \frac{\lambda}{|\mathcal{U}||\mathcal{I}|}$.

(3) Doubly Robust [48] combines the above models for more robustness – the capability to remain unbiased if either the imputed data or propensities are accurate. It optimizes:

$$\hat{L}_{DR}(f) = \frac{1}{|\mathcal{U}||\mathcal{I}|} \sum_{u \in \mathcal{U}, i \in \mathcal{I}} \left(\delta(f(u, i), m_{ui}) + \frac{O_{ui} d_{ui}}{q_{ui}} \right), \quad (10)$$

where $d_{ui} = \delta(f(u, i), r_{ui}^o) - \delta(f(u, i), r_{ui}^i)$ denotes the difference between the predicted error and imputed error. r_{ui}^o denotes the observed rating values. O_{ui} denotes whether the interaction of (u, i) is observed. Our framework can recover this one by setting $w_k^{(1)} = \frac{|D_T|}{q_{u_k i_k} |\mathcal{U}||\mathcal{I}|}$, $w_{ui}^{(2)} = \frac{1}{|\mathcal{U}||\mathcal{I}|} - \frac{O_{ui}}{q_{ui} |\mathcal{U}||\mathcal{I}|}$.

Conformity bias. The conformity effect can be offset by optimizing [29, 31]:

$$\hat{L}_{OFF}(f) = \frac{1}{|D_T|} \sum_{k=1}^{|D_T|} (\alpha r_k + (1 - \alpha) b_{u_k i_k} - f(u_k, i_k))^2, \quad (11)$$

where $b_{u_k i_k}$ denotes the introduced bias term, which can be specified as the average rating over all users [29] or social friends [31]. α controls the effect of conformity. Our framework subsumes it by setting $w_k^{(1)} = 1$, $w_{ui}^{(2)} = O_{ui}(1 - \lambda)$, $m_{ui} = -b_{ui}$.

Exposure bias. Existing methods are mainly two types:

(1) Negative Weighting, treats unobserved interactions as negative and down-weights their contributions [21]:

$$\hat{L}_{NW}(f) = \frac{1}{|D_T|} \sum_{k=1}^{|D_T|} \delta(f(u_k, i_k), r_k) + \sum_{(u, i) \in \mathcal{U} \times \mathcal{I}: O_{ui}=0} a_{ui} \delta(f(u, i), 0), \quad (12)$$

where parameter a_{ui} indicates how likely the item is exposed to a user, which can be specified heuristically [21] or by an exposure model [6, 9, 27]. We can recover this method by setting $w_k^{(1)} = 1$, $w_{ui}^{(2)} = a_{ui}(1 - O_{ui})$, $m_{ui} = 0$.

(2) IPS Variant, reweights observed data and imputes missing data [38]:

$$\hat{L}_{IPSV}(f) = \sum_{k=1}^{|D_T|} \frac{1}{q_{u_k i_k}} \delta(f(u_k, i_k), r_k) + \sum_{u \in \mathcal{U}, i \in \mathcal{I}} \left(1 - \frac{O_{ui}}{q_{ui}} \right) \delta(f(u, i), 0). \quad (13)$$

We can recover it by setting $w_k^{(1)} = \frac{|D_T|}{q_{u_k i_k} |\mathcal{U}||\mathcal{I}|}$, $w_{ui}^{(2)} = \frac{1}{|\mathcal{U}||\mathcal{I}|} - \frac{O_{ui}}{q_{ui} |\mathcal{U}||\mathcal{I}|}$, $m_{ui} = 0$. It is worth mentioning that the unbiasedness of

Variant IPS conditions on that the training distribution p_T covers the whole support of p_U , however it seldom holds in practice.

Position bias. The most popular strategy is IPS [3], which reweighs the collected data with a position-aware score \hat{q}_{t_k} :

$$\hat{L}_{IPS}(f) = \frac{1}{|D_T|} \sum_{k=1}^{|D_T|} \frac{\delta(f(u_k, i_k), r_k)}{q_{t_k}}. \quad (14)$$

It can be recovered by setting $w_k^{(1)} = \frac{1}{q_{t_k}}$, $w_{ui}^{(2)} = 0$.

4 AUTODEBIAS ALGORITHM

We now consider how to optimize the aforementioned framework. Since the training data lacks important signals of how the data is biased and what the unbiased data looks like, it is impossible to learn proper debiasing parameters from such data. To deal with this problem, another *uniform data* D_U is introduced to supervise the learning of debiasing parameters. The uniform data contains a list of triplets $\{(u_l, i_l, r_l)\}_{1 \leq l \leq |D_U|}$, which is assumed to be collected by a random logging policy, providing a gold standard evidence on the unbiased recommendation performance. We make full use of this evidence and optimize ϕ towards better performance on uniform data. Specifically, the learning process can be formulated as a *meta learning* process with:

Base learner: The base recommendation model is optimized on the training data with current debiasing parameters ϕ :

$$\theta^*(\phi) = \arg \min_{\theta} \hat{L}_T(f_{\theta}|\phi). \quad (15)$$

where debiasing parameters ϕ can be seen as the hyper-parameters of the base learner.

Meta learner: Given the learned base recommendation model $\theta^*(\phi)$ from training data with the hyper-parameters ϕ , ϕ is optimized towards better recommendation performance on uniform data:

$$\phi^* = \arg \min_{\phi} \frac{1}{|D_U|} \sum_{l=1}^{|D_U|} \delta(f_{\theta^*(\phi)}(u_l, i_l), r_l). \quad (16)$$

For better description, the empirical risk on uniform data is marked as $\hat{L}_U(f_{\theta})$, i.e., $\hat{L}_U(f_{\theta}) = \frac{1}{|D_U|} \sum_{l=1}^{|D_U|} \delta(f_{\theta}(u_l, i_l), r_l)$.

As the uniform data is often collected in a small scale, directly learning all parameters in ϕ from it will incur over-fitting. To deal with this problem, ϕ can be re-parameterized with a concise *meta model*. This treatment can reduce the number of parameters and encode useful information (e.g., user id, observed feedback) into debiasing. In this paper, we simply choose a linear model for implementation as:

$$\begin{aligned} w_k^{(1)} &= \exp(\varphi_1^T [\mathbf{x}_{u_k} \circ \mathbf{x}_{i_k} \circ \mathbf{e}_{r_k}]) \\ w_{ui}^{(2)} &= \exp(\varphi_2^T [\mathbf{x}_u \circ \mathbf{x}_i \circ \mathbf{e}_{O_{ui}}]) \\ m_{ui} &= \sigma(\varphi_3^T [\mathbf{e}_{r_{ui}} \circ \mathbf{e}_{O_{ui}}]), \end{aligned} \quad (17)$$

where \mathbf{x}_u and \mathbf{x}_i denote the feature vectors (e.g., one-hot vector of its id) of user u and item i , respectively; \mathbf{e}_r , $\mathbf{e}_{O_{ui}}$ are one-hot vectors of r and O_{ui} ; the mark \circ denotes the concatenation operation; $\varphi \equiv \{\varphi_1, \varphi_2, \varphi_3\}$ are surrogate parameters to be learned; $\sigma(\cdot)$ denotes the activated function controlling the scale of the imputation values, e.g., $\tanh(\cdot)$. One might concern that modeling ϕ with a meta model

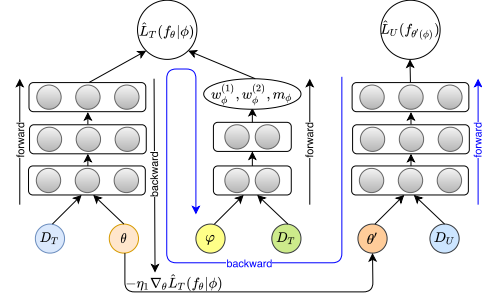


Figure 3: The working flow of AutoDebias, consists of three steps: (1) tentatively updating θ to θ' on the training data D_T with current ϕ (black arrows); (2) updating ϕ based on θ' on the uniform data (blue arrows); (3) actually updating θ with the updated ϕ (black arrows).

potentially induce inductive bias, restricting the flexibility of ϕ and making it fail to arrive at the global optimum. In fact, our framework is relatively robust to such inductive bias, which has been validated in Section 5.

Model learning. Note that obtaining optimal ϕ^* involves nested loops of optimization — updating ϕ a step forward requires a loop of full training of θ , which is expensive. To deal with this problem, we consider to update θ and ϕ alternately in a loop. That is, for each training iteration, we make a tentative updating of recommendation model with current ϕ and inspect its performance on the uniform data. The loss on uniform data will give feedback signal to update meta model. To be exact, as illustrated in Figure 3, we perform the following training procedure in each iteration:

- **Assumed update of θ .** As the black arrows in Figure 3, we make an assumed updating of θ :

$$\theta'(\phi) = \theta - \eta_1 \nabla_{\theta} \hat{L}_T(f_{\theta}|\phi), \quad (18)$$

where we update θ using gradient descent with learning rate η_1 .

- **Update of ϕ (φ).** As the blue arrows shown in figure 3, we test $\theta'(\phi)$ on the uniform data with \hat{L}_U . The loss function gives a feedback signal (gradient) to update the meta model φ :

$$\varphi \leftarrow \varphi - \eta_2 \nabla_{\varphi} \hat{L}_U(f_{\theta'}(\phi)). \quad (19)$$

The gradient can be calculated by using the back-propagation along the chain $\hat{L}_U(f_{\theta'}(\phi)) \rightarrow \theta'(\phi) \rightarrow \nabla_{\theta} \hat{L}_T(f_{\theta}|\phi) \rightarrow \phi \rightarrow \varphi$.

- **Update of θ .** Given the updated ϕ , we update θ actually:

$$\theta \leftarrow \theta - \eta_1 \nabla_{\theta} \hat{L}_T(f_{\theta}|\phi). \quad (20)$$

While this alternative optimization strategy is not guaranteed to find the global optimum, it empirically works well for bi-level optimization problem [14].

5 THEORETICAL ANALYSIS

In this section, we conduct theoretical analyses to answer the following questions: (1) Can AutoDebias acquire the optimal ϕ making $E_{p_T}[\hat{L}_T(f)]$ consistent with $L(f)$? (2) How does inductive bias in meta model affect recommendation performance?

Addressing question (1). We first give some notations. Let Φ and \mathbb{F} denote the hypothesis space of the meta model and the

recommendation model; f^* denotes the optimum function that minimizes $L(f)$; $\hat{f}(\phi)$ denotes empirical optimum under specific ϕ by optimizing i.e., $\hat{f}(\phi) = \arg \min_{f \in \mathbb{F}} \hat{L}_T(f|\phi)$; ϕ^* , the optimal debiasing parameters making $E_{p_T}[L_T(f|\phi^*)] = E[L(f)]$ hold, is supposed to be contained in Φ in this question, then the general case will be discussed in next question. ϕ^o denotes surrogate optimum from Φ that minimizes $\phi^o = \arg \min_{\phi \in \Phi} L(\hat{f}(\phi))$; $\hat{\phi}$ denotes the empirical ϕ from the empirical risk on uniform data $L_U(\hat{f}(\phi))$. As the empirical risk $\hat{L}_T(f|\phi^*)$ is an unbiased estimation of the true risk, we have:

LEMMA 1. Generalization bound of $L(\hat{f}(\phi^*))$. For any finite hypothesis space of recommendation models $\mathbb{F} = \{f_1, f_2, \dots, f_{|\mathbb{F}|}\}$, with probability of at least $1 - \eta$, we have the following generalization error bound of the learned model \hat{f} :

$$L(\hat{f}(\phi^*)) \leq L(f^*) + \sqrt{\frac{S_{w^{(1)}} \rho^2}{2|D_T|} \log \frac{2|\mathbb{F}|}{\eta}}, \quad (21)$$

where $S_{w^{(1)}}$ denotes mean-square of $w^{(1)}$ with $S_{w^{(1)}} = \frac{|D_T|}{\sum_{k=1}^{|D_T|} (w_k^{(1)})^2}$. ρ denotes the bound of loss function δ .

We omit the lemma proof due to the limit space. In fact, the proof is a variant of the proof of Corollary 4.6 in [40], where we use our debiasing empirical risk and a more strict Hoeffding inequality. For convenient, $\sqrt{\frac{S_{w^{(1)}} \rho^2}{2|D_T|} \log \frac{2|\mathbb{F}|}{\eta}}$ is marked as ε_1 .

Note that when the space Φ contains optimum ϕ^* , the relations $L(f^*) \leq L(\hat{f}(\phi^o)) \leq L(\hat{f}(\phi^*)) \leq L(f^*) + \varepsilon_1$ hold, we have:

COROLLARY 1. When the meta hypothesis space Φ contains optimum ϕ^* , with probability of at least $1 - \delta$, the differences between $L(f^*)$, $L(\hat{f}(\phi^o))$, $L(\hat{f}(\phi^*))$ are bounded with:

$$|L(f^*) - L(\hat{f}(\phi^o))| < \varepsilon_1, |L(\hat{f}(\phi^o)) - L(\hat{f}(\phi^*))| < \varepsilon_1. \quad (22)$$

This corollary suggests that when the training data is sufficiently large, the surrogate parameter ϕ^o is approximately correct, making the recommendation model f trained in an unbiased manner and finally arrive at the approximate optimum.

Note that $L_U(\hat{f}(\phi))$ is an unbiased estimation of $L(\hat{f}(\phi))$, similar to lemma 1, we have:

LEMMA 2. For any finite hypothesis space of meta model $\Phi = \{\phi_1, \phi_2, \dots, \phi_{|\Phi|}\}$, with probability of at least $1 - \eta$, we have:

$$L(\hat{f}(\hat{\phi})) \leq L(\hat{f}(\phi^o)) + \sqrt{\frac{\rho^2}{2|D_U|} \log \frac{2|\Phi|}{\eta}}. \quad (23)$$

This lemma details the relation of $\hat{\phi}$ and ϕ^o . By combining lemma 2 and corollary 1, we have the following generalization bound of the framework:

LEMMA 3. Generalization error bound of AutoDeBias. For any recommendation model f and meta model ϕ with finite hypothesis space \mathbb{F} and Φ , when Φ contains ϕ^* , with probability of at least $1 - \eta$, the generalization error of the model that is trained with AutoDebias, is bounded by:

$$L(\hat{f}(\hat{\phi})) < L(f^*) + \sqrt{\frac{S_{w^{(1)}} \rho^2}{2|D_T|} \log \frac{4|\mathbb{F}|}{\eta}} + \sqrt{\frac{\rho^2}{2|D_U|} \log \frac{4|\Phi|}{\eta}}. \quad (24)$$

Also, $L(\hat{f}(\hat{\phi}))$ does not deviate from $L(\hat{f}(\phi^*))$ more than:

$$|L(\hat{f}(\hat{\phi})) - L(\hat{f}(\phi^*))| < \sqrt{\frac{S_{w^{(1)}} \rho^2}{2|D_T|} \log \frac{4|\mathbb{F}|}{\eta}} + \sqrt{\frac{\rho^2}{2|D_U|} \log \frac{4|\Phi|}{\eta}}. \quad (25)$$

This lemma proves that when sufficient training data and uniform data are provided, AutoDeBias can acquire an almost optimal debiasing strategy, making the recommendation model arrive at the approximate optimum.

Addressing question (2). Inductive bias occurs in a small-capacity meta model, making the hypothesis space Φ do not contain ϕ^* and Equation (22) do not hold accordingly. Nevertheless, we have:

LEMMA 4. Generalization error bound of AutoDeBias with inductive bias. For any recommendation model f and meta model ϕ with finite hypothesis space \mathbb{F} and Φ , the generalization error of the model that is trained with AutoDebias, is bounded by:

$$L(\hat{f}(\hat{\phi})) < L(f^*) + \sqrt{\frac{S_{w^{(1)}}^o \rho}{2|D_T|} \log \frac{4|\mathbb{F}|}{\eta}} + \sqrt{\frac{\rho^2}{2|D_U|} \log \frac{4|\Phi|}{\eta}} + (\Delta_{(1)} + \Delta_{(2)})\rho + \Delta_m \rho_d. \quad (26)$$

where $S_{w^{(1)}}^o$ denotes the mean-square of the surrogate $w^{o(1)}$ in ϕ^o with $S_{w^{(1)}}^o = \frac{|D_T|}{\sum_{k=1}^{|D_T|} (w_k^{o(1)})^2}$; and $\Delta_{(1)}$, $\Delta_{(2)}$, Δ_m denotes the expected deviation between ϕ^o and ϕ^* caused by inductive bias in terms of $w^{(1)}$, $w^{(2)}$, m respectively, i.e., $\Delta_{(1)} = \sum_{k=1}^{|D_T|} p_U(u_k, i_k, r_k) \left| \frac{w_k^{*(1)} - w_k^{o(1)}}{w_k^{*(1)}} \right|$, $\Delta_{(2)} = \sum_{u \in U, i \in I} |w_{ui}^{*(2)} - w_{ui}^{o(2)}|$, $\Delta_m = \sum_{u \in U, i \in I} w_{ui}^{*(2)} |m_{ui}^* - m_{ui}^o|$; ρ_d denotes the bound of the gradient $\nabla \delta(\cdot, \cdot)$.

PROOF. Due to the limited space, here we just give brief proof. For convenience, we mark $(\Delta_{(1)} + \Delta_{(2)})\rho + \Delta_m \rho_d$ as ε_3 . Let $L_s(f|\phi)$ denote the surrogate risk with parameter ϕ , that is, $L_s(f|\phi) = E_{p_T}[\hat{L}_T(f|\phi)]$. Naturally, $L_s(f|\phi)$ deviates from the true risk $L(f)$. But their difference, which can be calculated from Equation (5), is bounded by ε_3 . Let $f^*(\phi)$ denote the optimum of the surrogate risk $L_s(f|\phi)$ with the parameters ϕ . The inequality $L_s(f^*(\phi^o)|\phi^o) < L(f^*) + \varepsilon_3$ holds. Further, note that $L_T(f|\phi^o)$ is an unbiased estimation of $L_s(f|\phi^o)$. The difference between $L(\hat{f}(\phi^o))$ and $L_s(\hat{f}(\phi^o))$ is bounded. Combining these conclusions with the lemma 3, lemma 4 is proved. \square

Lemma 4 tells us that although the inductive bias occurs, the generalization error is bounded. In fact, constraining the capacity of the meta model with inductive bias, often can improve model performance. On the one hand, the size of hypothesis space Φ is reduced; On the other hand, the surrogate optimum ϕ^o in the constraint space Φ often has smaller variance, reducing the mean-square $S_{w^{(1)}}^o$ and further tightening the bound.

6 EXPERIMENTS

In this section, we conduct experiments to evaluate the performance of our proposed AutoDebias. Our experiments are intended to address the following research questions:

RQ1: Does AutoDebias outperform SOTA debiasing methods?

Table 2: Statistics of the datasets.

Dataset	Users	Items	Training	Uniform	Validation	Test
Yahoo!R3	15,400	1,000	311,704	2,700	2,700	48,600
Coat	290	300	6,960	232	232	4,176
Simulation	500	500	12,500	3,750	3,750	67,500

RQ2: How do different components (*i.e.*, learning of $w^{(1)}$, $w^{(2)}$ or m) affect AutoDebias performance?

RQ3: How do learned debiasing parameters ϕ correct data bias?

RQ4: Is AutoDebias universal to handle various biases?

6.1 Experimental Setup

In order to validate the universality of AutoDebias, our experiments are on three types of data: explicit feedback, implicit feedback, and the feedback on recommendation list.

Dataset. For explicit feedback, we refer to [28] and use two public datasets (*Yahoo!R3* and *Coat*) for our experiments. Both datasets contain a set of biased data collecting the normal interactions of users in the platform, and a small set of unbiased data from stochastic experiment where items are assigned randomly. Following [28], we regard the biased data as training set D_T , while split the unbiased data into three parts: 5% for uniform set D_U to help training, 5% for validation set D_V to tune the hyper-parameters, and 90% for test set D_{Te} to evaluate the model. The ratings are binarized with threshold 3. That is, the observed rating value larger than 3 is labeled as positive ($r = 1$), otherwise negative $r = -1$.

To generate implicit feedback data, we also use the aforementioned datasets, but abandon the negative feedback of the training data. this treatment is the same as the recent work [38].

We also adopt a synthetic dataset *Simulation*, which simulates user feedback on recommendation lists to validate that AutoDebias can handle the situation when both position bias and selection bias occur. Here we use synthetic dataset as there is no public dataset that contains ground-truth of unbiased data for model evaluation. The simulation process consists of four steps: (1) we follow [45] to generated the ground-truth preference scores r_{ui} of 500 users on 500 items. For each user, we randomly split their feedback into two parts: 150 for unbiased data set (D_1) and 350 remaining for simulating biased data (D_2). (2) For each user, 25 items are sampled randomly from D_2 to train a recommendation model M . M returns recommendation lists containing top-25 items for each user, which induces selection bias in data. (3) We then refer to [15] to simulate user feedback on recommendation lists, where position bias has been considered. Specifically, the click is generated with the probability of $\min(\frac{r_{u,i}}{2 * p^{1/2}}, 1)$. The statistics of those datasets is in Table 2.

Evaluation metrics. We adopt the following metrics to evaluate recommendation performance:

- **NLL** evaluates the performance of the predictions with:

$$NLL = -\frac{1}{|D_{Te}|} \sum_{(u,i,r) \in D_{Te}} \log \left(1 + e^{-r * f_{\theta}(u,i)} \right), \quad (27)$$

- **AUC** evaluates the performance of rankings:

$$AUC = \frac{\sum_{(u,i) \in D_{Te}^+} \hat{Z}_{u,i} - (|D_{Te}^+| + 1)(|D_{Te}^+|)/2}{(|D_{Te}^+|) * (|D_{Te}| - |D_{Te}^+|)}, \quad (28)$$

where $|D_{Te}^+|$ denotes the number of positive data in D_{Te} , $\hat{Z}_{u,i}$ denotes the rank position of a positive feedback (u, i).

- **NDCG@k** measures the quality of recommendation through discounted importance based on position:

$$DCG_u@k = \sum_{(u,i) \in D_{Te}} \frac{I(\hat{Z}_{u,i} \leq k)}{\log(\hat{Z}_{u,i} + 1)} \quad (29)$$

$$NDCG@k = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \frac{DCG_u@k}{IDCG_u@k},$$

where $IDCG_u@k$ is the ideal $DCG_u@k$.

Implementation details. Matrix Factorization (MF) has been selected as a benchmark recommendation model for experiments, and it would be straightforward to replace it with more sophisticated models such as Factorization Machine [35], or Neural Network [18]. SGD has been adopted for optimizing base model and Adam for meta model. Grid search is used to find the best hyper-parameters based on the performance on validate set. We optimize the base model with SGD and meta model with Adam. The search space of learning rate and weight decay are [1e-4, 1e-3, 1e-2, 1e-1].

6.2 Performance Comparison on Explicit Feedback (RQ1)

Baseline. (1) MF(biased), MF(uniform) and MF(combine): the basic matrix factorization model that trained on D_T , D_U , and $D_T + D_U$, respectively; (2) Inverse propensity score (IPS) [39]: a counterfactual technique that reweighs the collected data. We follow [39] and calculate the propensity with naive bayes; (3) Doubly robust (DR) [48]: combining data imputation and inverse propensity score; (4) KD-Label [28]: the state-of-the-art method that transfers the unbiased information with a teacher model. We refer to [28] and choose the best label-based distillation for comparison. (5) CausE [28]: distilling unbiased information with an extra alignment term.

Performance comparison. Table 3 presents the recommendation performance of the compared methods in terms of three evaluation metrics. The boldface font denotes the winner in that column. We have the following observations: (1) Overall, AutoDebias outperforms all compared methods on all datasets for all metrics. Especially in the dataset Yahoo!R3, the improvements are rather impressive – 5.6% and 11.2% in terms of NLL and NDCG@5. This result validates that our proposed AutoDebias can learn better debiasing configurations than compared methods. (2) MF(uniform) that is directly trained on uniform data performs quite terribly. The reason is that uniform data is often in a quite small scale. Training a recommendation model on uniform data will suffer from serious over-fitting. This phenomenon also affect the performance of KD-based methods (KD-label and CausE). Thus, they perform worse than AutoDebias with a certain margin.

To further validate the high efficiency of AutoDebias on utilizing uniform data, we test these methods with varying size of uniform data as shown in figure 4. We can find AutoDebias consistently outperforms compared methods. More impressively, AutoDebias still perform well even if a quite small scale of uniform data is provided (*e.g.*, Ratio=1%), while KD-label falls short in this situation – performs close to or even worse than MF(biased).

Table 3: Performance comparisons on explicit feedback data. The boldface font denotes the winner in that column.

Method	Yahoo!R3			Coat		
	NLL	AUC	NDCG@5	NLL	AUC	NDCG@5
MF(biased)	-0.587	0.727	0.550	-0.539	0.747	0.500
MF(uniform)	-0.513	0.573	0.449	-0.623	0.580	0.358
MF(combine)	-0.580	0.730	0.554	-0.538	0.750	0.504
IPS	-0.448	0.723	0.549	-0.515	0.759	0.509
DR	-0.444	0.723	0.552	-0.512	0.765	0.521
CausE	-0.579	0.731	0.551	-0.529	0.762	0.500
KD-Label	-0.632	0.740	0.580	-0.593	0.748	0.504
AutoDebias	-0.419	0.741	0.645	-0.512	0.766	0.522




Figure 4: Performance comparisons with varying uniform data size. Here we sample different percentages of total unbiased data as uniform data.

6.3 Ablation Study (RQ2)

We conduct ablation study to explore whether the three components $w^{(1)}$, $w^{(2)}$ and m are all desired to be introduced. Table 4 shows the performance of the ablated model where different components are removed. From this table, we can conclude that all the parameters are important for debiasing. Introducing $w^{(1)}$ and m consistently boost the performance. Despite introducing $w^{(2)}$ harms the performance a bit on Coat, it brings an impressive improvement on Yahoo!R3. It is worth emphasising that even if imputation modular is removed, AutoDebias-w1 still outperforms IPS. This results validates the effectiveness of the meta-learning algorithm on finding the proper $w^{(1)}$.

6.4 Exploratory Analysis (RQ3)

To answer question RQ3, we now explore the learned meta parameters from the dataset Yahoo!R3. This exploration provides insights into how AutoDebias corrects data bias.

Exploring imputation values m : Table 5 presents the imputation values for user-item pairs with different kinds of feedback. We can find the imputation value for the positive user-item pairs is larger than missing pairs, while the value for missing pairs is large than negative pair. More interestingly, we can find the optimal imputation value for missing pair is rather small (e.g., -0.913). This phenomenon is consistent with the finding in [32] that a user dislikes a random-selected item in a high probability.

Exploring weights $w^{(1)}$: Note that $w^{(1)}$ is implemented with a linear model while the feature vector consists of three one hot vectors in terms of user id, item id and the feedback. The linear parameter φ_1 can be divided into three parts $\varphi_1 = [\varphi_{11} \circ \varphi_{12} \circ \varphi_{13}]$ and $w^{(1)}$ can be decomposed as a product of user-dependent weights

Table 4: Ablation study in terms of NDCG@5.

Method	$w^{(1)}$	m	$w^{(2)}$	Yahoo!R3	Coat
MF	×	×	×	0.550	0.500
AutoDebias-w1	✓	×	×	0.573	0.510
AutoDebias-w1m	✓	✓	×	0.581	0.526
AutoDebias	✓	✓	✓	0.645	0.521

Table 5: The learned parameters m for different feedback.

Positive ($r = 1$)	Missing	Negative ($r = 1$)
-0.992	-0.913	0.227

Table 6: The learned parameters $w_r^{(1)}$ in different models.

Method	$r = -1$	$r = 1$
AutoDebias-w1	3.37	1.56
AutoDebias	0.29	4.77

$w_u^{(1)}$, the item-dependent weights $w_i^{(1)}$ and the label-dependent weights $w_r^{(1)}$, i.e., $w_k^{(1)} = \exp(\varphi_{11}^T x_{u_k}) \exp(\varphi_{12}^T x_{i_k}) \exp(\varphi_{13}^T x_{r_k}) = w_{u_k}^{(1)} w_{i_k}^{(1)} w_{r_k}^{(1)}$. Here we explore the distribution of the learned $w_u^{(1)}$, $w_i^{(1)}$, $w_r^{(1)}$ to provide the insights.

Table 6 shows the value of learned $w_r^{(1)}$ for ablated AutoDebias-w1 and AutoDebias. We can find that $w_r^{(1)}$ for $r = -1$ is larger than $w_r^{(1)}$ for $r = 1$ in AutoDebias-w1, while $w_r^{(1)}$ for $r = -1$ is smaller than $w_r^{(1)}$ for $r = 1$ in AutoDebias. This interesting phenomenon can be explained as following. AutoDebias-w1 that does not use data imputation, will up-weight the negative instances to account for such selection bias. For AutoDebias, as it has given negative imputed labels for missing data, it does not need to address such under-represented problem and turn to learn other aspects such as data confidence. This result validates the adaptation of our AutoDebias.

We also explore the distribution of $w_i^{(1)}$ with the item popularity as presented in Figure 5(a). We can find: (1) The weights decreases as their popularity increases. This phenomenon is consistent with our intuition that the popular items are likely to be exposed (aka. popularity bias) and their feedback is over-represented and need down-weighting to mitigate bias. (2) The variance of the weights of the popular items is larger than unpopular. The reason is that some popular items are high-quality, on which the feedback is indeed valuable, while the feedback of others may be caused by the conformity effect and is unreliable. As a result, the weights exhibit more diverse.

To further validate the superiority of our model on addressing such popularity bias. We draw the plots of the performance of different methods in terms of popular and unpopular items in Figure 5(a). The popular items are defined as the top 20% items according to their positive feedback frequency, while the unpopular items are for the rest. We can find AutoDebias provides more exposure opportunities to unpopular items, and the improvement mainly comes from unpopular items. This result validates that AutoDebias can generate fairer recommendation results.




Figure 5: (a) The learned $w_i^{(1)}$ for each item i with its popularity; (b) Recommendation results in terms of popular items; (c) Recommendation results in terms of unpopular items.

Table 7: The performance on implicit feedback data.

Method	Yahoo!R3-Im		Coat-Im	
	AUC	NDCG@5	AUC	NDCG@5
WMF	0.635	0.547	0.749	0.521
RI-MF	0.673	0.554	0.696	0.527
AWMF	0.675	0.578	0.614	0.505
AutoDebias	0.730	0.635	0.746	0.527

Table 8: The performance in the list feedback of compared methods on Simulation dataset.

	NLL	AUC	NDCG@5
MF(biased)	-0.712	0.564	0.589
DLA	-0.698	0.567	0.593
HeckE	-0.688	0.587	0.648
AutoDebias	-0.667	0.634	0.707

6.5 Performance Comparison on Implicit and Feedback on Lists (RQ4)

To validate the universality of AutoDebias, we also conduct experiments on implicit feedback and the feedback on lists.

Implicit feedback. We compare AutoDebias with three SOTA methods on exposure bias: WMF, RL-MF [38], and AWMF [8]. Table 7 shows the performance comparison. As this table shows, AutoDebias performs much better than others, which validates our AutoDebias can handle the exposure bias in implicit feedback.

The feedback on recommendation list. In this scenario, we add position information in modeling $w_k^{(1)}$, i.e., $w_k^{(1)} = \exp(\varphi_1^T [\mathbf{x}_{u_k} \circ \mathbf{x}_{i_k} \circ \mathbf{e}_{r_k} \circ \mathbf{e}_{p_k}])$. We compare AutoDebias with two SOTA methods: Dual Learning Algorithm (DLA) [3] which mitigate position bias, and Heckman Ensemble (HeckE) [34] which mitigates both position bias and selection bias by ensemble. Table 8 shows the results. we can find AutoDebias still outperform compared methods. this result verifies that AutoDebias can handle various biases, even their combination.

7 RELATED WORK

Bias in recommendation. Besides the data bias that has been detailed in section 2&3, two important biases in recommendation results have been studied: (1) When a recommender model is trained on a long-tailed data, popular items are recommended even more frequently than their popularity would warrant [2], raising so-called *popularity bias*. The long-tail phenomenon is common in RS data,

and ignoring the popularity bias will incur many issues, e.g., hurting user serendipity, making popular items even more popular. To deal with this problem, [1, 11, 49] introduced regularizers to guide the model to give more balanced results; [50, 54] disentangled the effect of user interest and item popularity with causal inference to address popularity bias. (2) The system systematically and unfairly discriminates against certain individuals or groups of individuals in favor others, raising so-called *unfairness*. Unfairness happens as different user (or item) groups are usually unequally represented in data. When training on such unbalanced data, the models are highly likely to learn these over-represented groups, and potentially discriminates against other under-represented groups [13, 43]. There are four types of strategies on addressing unfairness, including re-balancing[41], regularizer[51], adversarial learning [4], causal inference [26]. We encourage the readers refer to the survey [5] for more details.

Meta learning in recommendation. Meta learning is an automatic learning algorithms that aims at using metadata to improve the performance of existing learning algorithms or to learn (induce) the learning algorithm itself [14, 20, 46]. There is also some work introducing meta learning in RS. For example, [10, 36] leveraged meta learning to learn a finer-grained regularization parameters; [53] leveraged meta-learning to guide the re-training of a recommendation model towards better performance; Meta-learning also has been utilized to address the cold-start problem [30].

8 CONCLUSION AND FUTURE WORK

This paper develops a universal debiasing framework that not only can well account for multiple biases and their combinations, but also frees human efforts to identify biases and tune the configurations. We first formulate various data biases as a case of risk discrepancy, and then derive a general learning framework that subsumes most debiasing strategies. We further propose a meta-learning-based algorithm to adaptively learn the optimal debiasing configurations from uniform data. Both theoretical and empirical analyses have been conducted to validate the effectiveness of our proposal.

One interesting direction for future work is to explore more sophisticated meta model, which could capture more complex patterns and potentially achieve better performance than linear meta model. Also, note that in real-world biases are usually dynamic rather than static. it will be valuable to explore how bias evolves with the time goes by and develop a universal solutions for dynamic biases.

ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China (U19A2079, 61972372), National Key Research and Development Program of China (2020AAA0106000), USTC Research Funds of the Double First-Class Initiative (WK210000019), and the Alibaba Group through Alibaba Innovative Research Program.

REFERENCES

- [1] Himan Abdollahpour, Robin Burke, and Bamshad Mobasher. 2017. Controlling Popularity Bias in Learning-to-Rank Recommendation. In *RecSys*. 42–46.
- [2] Himan Abdollahpour and Masoud Mansoury. 2020. Multi-sided exposure bias in recommendation. *arXiv preprint arXiv:2006.15772* (2020).
- [3] Qingyao Ai, Keping Bi, Cheng Luo, Jiafeng Guo, and W. Bruce Croft. 2018. Unbiased Learning to Rank with Unbiased Propensity Estimation. In *SIGIR*. 385–394.

- [4] Ghazaleh Beigi, Ahmadrza Mosallanezhad, Ruocheng Guo, Hamidreza Alviri, Alexander Nou, and Huan Liu. 2020. Privacy-Aware Recommendation with Private-Attribute Protection using Adversarial Learning. In *WSDM*. 34–42.
- [5] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2020. Bias and Debias in Recommender System: A Survey and Future Directions. *arXiv preprint arXiv:2010.03240* (2020).
- [6] Jiawei Chen, Yan Feng, Martin Ester, Sheng Zhou, Chun Chen, and Can Wang. 2018. Modeling Users’ Exposure with Social Knowledge Influence and Consumption Influence for Recommendation. In *CIKM*. 953–962.
- [7] Jiawei Chen, Can Wang, Martin Ester, Qihao Shi, Yan Feng, and Chun Chen. 2018. Social Recommendation with Missing Not at Random Data. In *ICDM*. 29–38.
- [8] Jiawei Chen, Can Wang, Sheng Zhou, Qihao Shi, Jingbang Chen, Yan Feng, and Chun Chen. 2020. Fast Adaptively Weighted Matrix Factorization for Recommendation with Implicit Feedback. In *AAAI*. 3470–3477.
- [9] Jiawei Chen, Can Wang, Sheng Zhou, Qihao Shi, Yan Feng, and Chun Chen. 2019. SamWalker: Social Recommendation with Informative Sampling Strategy. In *WWW*. 228–239.
- [10] Yihong Chen, Bei Chen, Xiangnan He, Chen Gao, Yong Li, Jian-Guang Lou, and Yue Wang. 2019. λ Opt: Learn to Regularize Recommender Models in Finer Levels. In *KDD*. 978–986.
- [11] Zhihong Chen, Rong Xiao, Chenliang Li, Gangfeng Ye, Haochuan Sun, and Hongbo Deng. 2020. ESAM: Discriminative Domain Adaptation with Non-Displayed Items to Improve Long-Tail Performance. In *SIGIR*. 579–588.
- [12] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *WSDM*. 87–94.
- [13] Michael D. Ekstrand, Mucun Tian, Mohammed R. Imran Kazi, Hoda Mehrpouyan, and Daniel Kluber. 2018. Exploring author gender in book rating and recommendation. In *RecSys*. 242–250.
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *ICML*. 1126–1135.
- [15] David F. Gleich and Lek-Heng Lim. 2011. Rank aggregation via nuclear norm minimization. In *KDD*. 60–68.
- [16] David Haussler. 1990. *Probably approximately correct learning*. University of California, Santa Cruz, Computer Research Laboratory.
- [17] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*. 639–648.
- [18] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *WWW*. 173–182.
- [19] José Miguel Hernández-Lobato, Neil Houlsby, and Zoubin Ghahramani. 2014. Probabilistic Matrix Factorization with Non-random Missing Data. In *ICML*. 1512–1520.
- [20] Sepp Hochreiter, A. Steven Younger, and Peter R. Conwell. 2001. Learning to Learn Using Gradient Descent. In *ICANN*. 87–94.
- [21] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *ICDM*. 263–272.
- [22] Thorsten Joachims, Laura A. Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2017. Accurately Interpreting Clickthrough Data as Implicit Feedback. In *SIGIR*. 4–11.
- [23] Thorsten Joachims, Laura A. Granka, Bing Pan, Helene Hembrooke, Filip Radlinski, and Geri Gay. 2007. Evaluating the accuracy of implicit feedback from clicks and query reformulations in Web search. *TOIS* 25, 2 (2007), 7.
- [24] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2014. Correcting Popularity Bias by Enhancing Recommendation Neutrality. In *RecSys*, Vol. 1247.
- [25] Sanjay Krishnan, Jay Patel, Michael J Franklin, and Ken Goldberg. 2014. A methodology for learning, analyzing, and mitigating social influence bias in recommender systems. In *RecSys*. 137–144.
- [26] Matt J. Kusner, Joshua R. Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. In *NeurIPS*. 4066–4076.
- [27] Dawen Liang, Laurent Charlin, James McInerney, and David M. Blei. 2016. Modeling User Exposure in Recommendation. In *WWW*. 951–961.
- [28] Dugang Liu, Pengxiang Cheng, Zhenhua Dong, Xiuqiang He, Weike Pan, and Zhong Ming. 2020. A General Knowledge Distillation Framework for Counterfactual Recommendation via Uniform Data. In *SIGIR*. 831–840.
- [29] Yiming Liu, Xuezhai Cao, and Yong Yu. 2016. Are You Influenced by Others When Rating? Improve Rating Prediction by Conformity Modeling. In *RecSys*. 269–272.
- [30] Yuanfu Lu, Yuan Fang, and Chuan Shi. 2020. Meta-learning on heterogeneous information networks for cold-start recommendation. In *KDD*. 1563–1573.
- [31] Hao Ma, Irwin King, and Michael R. Lyu. 2009. Learning to recommend with social trust ensemble. In *SIGIR*. 203–210.
- [32] Benjamin M. Marlin and Richard S. Zemel. 2009. Collaborative prediction and ranking with non-random missing data. In *RecSys*. 5–12.
- [33] Benjamin M. Marlin, Richard S. Zemel, Sam T. Roweis, and Malcolm Slaney. 2007. Collaborative Filtering and the Missing at Random Assumption. In *UAI*. 267–275.
- [34] Zohreh Ovaisi, Ragib Ahsan, Yifan Zhang, Kathryn Vasilaky, and Elena Zheleva. 2020. Correcting for Selection Bias in Learning-to-rank Systems. In *WWW*. 1863–1873.
- [35] Steffen Rendle. 2012. Factorization machines with libfm. *TIST* 3, 3 (2012), 57:1–57:22.
- [36] Steffen Rendle. 2012. Learning recommender systems with adaptive regularization. In *WSDM*. 133–142.
- [37] Yuta Saito. 2020. Asymmetric Tri-training for Debiasing Missing-Not-At-Random Explicit Feedback. In *SIGIR*. 309–318.
- [38] Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. 2020. Unbiased Recommender Learning from Missing-Not-At-Random Implicit Feedback. In *WSDM*. 501–509.
- [39] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. 2016. Recommendations as Treatments: Debiasing Learning and Evaluation. In *ICML*. 1670–1679.
- [40] Shai Shalev-Shwartz and Shai Ben-David. 2014. *Understanding machine learning: From theory to algorithms*. Cambridge university press.
- [41] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of Exposure in Rankings. In *KDD*. 2219–2228.
- [42] Harald Steck. 2013. Evaluation of recommendations: rating-prediction and ranking. In *RecSys*. 213–220.
- [43] Ana-Andreea Stoica, Christopher J. Riederer, and Augustin Chaintreau. 2018. Algorithmic Glass Ceiling in Social Networks: The effects of social recommendations on network diversity. In *WWW*. 923–932.
- [44] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *CIKM*. 1441–1450.
- [45] Wenlong Sun, Sami Khenissi, Olfa Nasraoui, and Patrick Shafto. 2019. Debiasing the Human-Recommender System Feedback Loop in Collaborative Filtering. In *WWW*. 645–651.
- [46] Sebastian Thrun and Lorien Y. Pratt (Eds.). 1998. *Learning to Learn*. Springer.
- [47] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to Rank with Selection Bias in Personal Search. In *SIGIR*. 115–124.
- [48] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. 2019. Doubly robust joint learning for recommendation on data missing not at random. In *ICML*. 6638–6647.
- [49] Jacek Wasilewski and Neil Hurley. 2016. Incorporating Diversity in a Learning to Rank Recommender System. In *FLAIRS*. 572–578.
- [50] Tianxin Wei, Fuli Feng, Jiawei Chen, Chufeng Shi, Ziwei Wu, Jinfeng Yi, and Xiangnan He. 2020. Model-Agnostic Counterfactual Reasoning for Eliminating Popularity Bias in Recommender System. *arXiv preprint arXiv:2010.15363* (2020).
- [51] Sirui Yao and Bert Huang. 2017. Beyond Parity: Fairness Objectives for Collaborative Filtering. In *NeurIPS*. 2921–2930.
- [52] Fajie Yuan, Xiangnan He, Alexandros Karatzoglou, and Liguang Zhang. 2020. Parameter-Efficient Transfer from Sequential Behaviors for User Modeling and Recommendation. In *SIGIR*. 1469–1478.
- [53] Yang Zhang, Fuli Feng, Chenxu Wang, Xiangnan He, Meng Wang, Yan Li, and Yongdong Zhang. 2020. How to retrain recommender system? A sequential meta-learning method. In *SIGIR*. 1479–1488.
- [54] Yu Zheng, Chen Gao, Xiang Li, Xiangnan He, Depeng Jin, and Yong Li. 2021. Disentangling User Interest and Conformity for Recommendation with Causal Embedding. In *WWW*.