# Causal Debiasing in Recommender Systems: Principles and Prospects

*Xiaoyu* Zhang*

Dundee International Institute of Central South University, Central South University, Changsha, Hunan, China

**Abstract.** The recommender system (RS) directly influences users' experience in e-commerce, video, social media and service platforms. However, it performs restrictively due in part to the inevitable various biases. Although traditional debiasing methods such as Inverse Propensity Score (IPS) and Double Robustness (DR) can mitigate the problems to some extent, new glitches (idealization of exposure mechanism and high variation) still have an effect on the performance of RS. Nowadays, causal inference, which can facilitate the robustness and fairness of RS models, has become one crucial method to handle such bias problems. This paper initially outlines the connotation, significance and restrictions of RS and causal inference respectively. Depending on the procedure of RS, it then classifies and analyzes some popular causal debiasing methods in the past three years. Finally, potential future prospects are provided according to the characteristics of those popular methods, aiming at giving a guide to subsequent research and study.

## 1 Introduction

With the booming trend of e-commerce platforms (Amazon, Alibaba), video sharing platforms (YouTube, TikTok), social platforms (Instagram, Facebook), and service platforms (Meituan, Yelp) in recent years, a great recommender system can not only bring convenience to the user from the overloaded information, but also increase the potential profits of platforms. However, the model has various biases such as exposure bias (some items are more frequently displayed), position bias (the first few items attract more attention) and popularity bias (vogue items are more likely to be recommended) due to the dependence on historical interaction logs [1]. The robustness and accuracy will decrease on account of those biases, thereby affecting the efficiency of searching and the user's satisfaction.

For such problems, traditional statistical methods (IPS, DR estimation, etc.) mitigate the biases by amending or weighing the observed data, which requires the propensity score to be sufficiently precise [2]. The aim is to revise the statistical data to be a debiased one.

Nowadays, causal inference is found to be a more effective way of handling the biases in the RS. Unlike traditional methods, causal inference relies on intervention and counterfactual reasoning problems: what if this item were available to the user? Through the causal model,

* Corresponding author's email: 7805230224@csu.edu.cn

the factors that truly drive the users are separated from spurious correlations. Therefore, the RS outputs a more accurate preference of the users [1].

In this context, this paper aims to clarify the available causal debiasing methods in the RS. Initially, the contents, importance, and limitations of the RS and causal inference are presented to demonstrate their value and research significance. Next, this paper selects some popular research from the past three years and clarifies them according to the stage at which the method works. Finally, possible future prospects are provided for the readers, including... Through this paper, the author hopes to introduce a knowledge map for beginners and offer inspiration and references for related workers.

# 2 Recommender System and Causal Inference

## 2.1 Recommender system

Recommender System (RS), which is widely used in platforms such as e-commerce, video sharing platforms and social media, is one of the most crucial techniques on the internet platform. It can predict users' possible interests due to historical information and preferences from a huge amount of data.

It can personalize the exposure content for different users to improve user stickiness, efficiency and satisfaction. In terms of the platform, a good RS can provide more click-through rates and conversion rates, thereby increasing the profits of the platform providers. A rational recommendation system can also take into account both popular and unpopular content, maintaining the diversity and fairness of the system.

Nevertheless, due to the influence of factors like the trend of popular content and the system display mechanism, the recommendation system unavoidably encounters deviation issues. As a result, the data obtained may not precisely mirror the genuine preferences of users. If this deviation is not properly managed, the model is prone to misinterpreting this false data as user preferences, thereby exerting an adverse effect on the consistency between online performance and offline metrics [1].

From the aspect of system implementation, in accordance with the official guiding paper of RecSys2022, the industrial process of the recommendation system can be clearly divided into the following six phases: feature engineering and preprocessing, a retrieval model for candidate generation (coarse screening), filtering (rule-based filtering (duplicate removal, constraints)), a feature store query (supplementing context features from the feature library), a ranking model for scoring (main model scoring), and an ordering stage (ranking display/strategy adjustment stage) [3].

From the perspectives of learning and reasoning, by integrating multiple papers, a standard process for recommendation systems can be constructed, which includes five stages: data construction, representation learning, model training, strategy optimization, and evaluation. This paper mainly discusses based on the stages from the perspectives of learning and reasoning [1, 4].

## 2.2 Causal inference

Causal Inference, differing from correlation analysis, concentrates on cause and effect. For instance, it concerns whether the modification of one variable will lead to the change of another one.

Causal Inference identifies and eliminates noise, making the data more accurate to reflect users' genuine preferences. Moreover, causal methods can grasp the inner logics in the RS and keep the performance stable when the distribution of data shifts. It can also estimate the real long-term performance of the strategy in reinforcing learning or recommendation in the long run [1].

Theoretically, causal inference depends on the accuracy of the cause-and-effect diagram and observability of potential noise. Otherwise, the result will be biased. Randomized controlled trials (RCTS) commonly used in causal identification are costly and time-consuming, and are not suitable for large-scale and frequent use. Most of the time, the observational data used have a large amount of deviation. At the application level, if the data of new users or new things is insufficient, causal models may fail to identify stable and accurate causal relationships [1].

# 3 Principal Methods

The methods are classified according to the procedure from the perspectives of learning and reasoning. The variety of biases, the logic of using causal inference and some crucial methods are mentioned in each stage.

## 3.1 Data construction stage

In this stage, common biases include exposure bias and position bias, leading to the mistaken thought that one may not be interested in the items that are infrequently demonstrated to them since the lack of interaction information.

To solve this problem, structural causal logic and counterfactual generation can simulate those missing unpopular items, filling up the interaction samples. Three corresponding prevalent methods are given below.

The first method is Counterfactual Data Augmentation (CoDA), which generates unexposed training samples through its algorithm [5]. The second one is CAIAC, which refers to Causal Action Influence Aware Counterfactual Data Augmentation. It can create viable synthetic transitions from static offline data and is illustrated to be more robust in offline learning algorithms when facing distributional shift [6]. SCM4SR, Structural Causal Model-based Data Augmentation for Robust Session-based Recommendation, provides a structural causal model (SCM) based on a neural network (NN) to generate counterfactual conversations, completing the logs [7].

## 3.2 Representation learning stage

During the learning stage of recommender systems, conformity bias and popularity bias are two common issues. Users tend to click on or choose items that are already popular, which in turn further amplifies the exposure of these items, creating a self-reinforcing cycle where "popular items become even more popular."

To better capture genuine user preferences, it is necessary to analyze the causal logic behind user behaviors and distinguish between conformity-driven actions and true personal interests. By identifying the underlying cause of a user's choice, the system can more accurately infer what the user genuinely likes.

A series of recent works aims to alleviate these biases with the proposed methods. For example, the DICE model decomposes user behaviors into "interest component" and a "conformity component" to identify the actual preference, and the DCCL model uses contrastive learning to discriminate the causal features from the spurious features to alleviate the influences of conformity bias and popularity bias respectively [4, 8].

### 3.3 Model training stage

During the model training stage, two typical forms of bias are popularity bias and confounding bias. The former causes the system to overemphasize popular items, while the latter arises from irrelevant user–item attributes that influence exposure and click rates, preventing the model from accurately capturing users' true interests.

Researchers commonly use such front-door and back-door criteria to control the confounding effect and reliably identify the underlying causalities that dictate users' behaviors when such methodological problems are encountered.

For example, CausalInt applies the front-door criterion to model causal pathways in multi-scenario recommendation tasks, enabling a clearer understanding of inter-scenario causal structures [9]. Meanwhile, CausalEPP models the causal chain from popularity to exposure and then to clicks, using intervention mechanisms to neutralize the influence of popularity and better recover intrinsic user preferences [10].

### 3.4 Strategy optimization stage

During the optimization phase for strategy selection, the system might be excessively biased towards the long-term reward. Certain contents might generate relatively higher short-term CTR, but actually may cause users to be unhappy in the long term, which might result in selections with a bias towards the short-term.

People hereafter address the challenge that, to accommodate the temporal dynamics and delayed feedback, counterfactual trajectories for a user to imitate its long-term behavioral tendency should be built, which could also be employed in policy optimization.

Typical examples include CDT4Rec, which leverages decision transformers to learn offline causal policies balancing short- and long-term outcomes, and PGCR, which employs causal state representations to guide policy generation and achieve more stable long-term performance [11-12].

## 4 Future Works

In recent years, an increasing number of studies have begun to explore cross-stage and integrated causal recommendation frameworks. Some works combine causal representation learning with data augmentation. For example, SCM4SR uses a simulated causal graph to synthesize new training samples, where the simulation inherently captures part of the causal representation.

Other approaches take a plug-in perspective that links several stages of the pipeline. One such example is CIPHER, which first learns disentangled causal embeddings during the data preprocessing and representation learning phases, and then feeds these embeddings into downstream recommendation models [13].

While the community is still far from agreeing upon a single"end-to-end" causal recommendation architecture, existing literature clearly moves away from single stage approaches in favor of: having the representation generated by an earlier stage helping with policy optimization, having causal samples generated earlier helping with both model training and testing, and having an approach which starts to connect causal methods from the data preparation and modeling stage to strategy fine-tuning and user response, respectively.

## 5 Conclusion

The contributions of this paper are twofold. On one hand, present the comprehensive surveys about the recent advances in causal debiasing methods for recommender systems, including four steps: construction of data, representation learning, modeling training and strategy optimization. On the other hand, summarize the major categories of bias and the proposed causal solutions based on the bias in different steps. In contrast to general statistical correction methods, causal inference provides a more profound view of the underlying relationship between intervention and outcome, which leads to a more robust, fair, and interpretable recommendation model. Also, recent researches take the steps from single-stage approaches to multi-stage approaches, such as cross-stage approaches, e.g., CIPHER, which is a representative of the tendency of considering holistic causal modeling in the research of recommendation.

However, there are still a few drawbacks. First, many causal structure identification methods rely on experts' assumptions and so the generality is restricted. Second, causal inference is still costly computationally and experimentally in large-scale online environments. Third, the trade-off between recommendation performance and fairness still exists in a real-life dynamic environment and is challenging. Future work can be studied in the directions of creating unified end-to-end models, interpretable causal representations, and adaptive debiasing to combine causal induction with RL, to foster more robust, interpretable and sustainable RS systems.

## References

1. J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, X. He, Bias and debias in recommender system: A survey and future directions. ACM Trans. Inf. Syst. 41(3), 1–39 (2023)
2. H. Li, Y. Xiao, C. Zheng, P. Wu, P. Cui, Propensity matters: Measuring and enhancing balancing for recommendation, in Proceedings of the Int. Conf. on Machine Learning (ICML 2023), 20182–20194 (2023)
3. R. Ak, B. Schifferer, S. Rabhi, G. De Souza Pereira Moreira, Training and deploying multi-stage recommender systems, in Proceedings of the 16th ACM Conf. on Recommender Systems (RecSys 2022), 706–707 (2022)
4. Y. Zheng, C. Gao, X. Li, X. He, Y. Li, D. Jin, Disentangling user interest and conformity for recommendation with causal embedding, in Proceedings of the Web Conference 2021 (WWW 2021), 2980–2991 (2021)
5. S. Pitis, E. Creager, A. Garg, Counterfactual data augmentation using locally factored dynamics. Adv. Neural Inf. Process. Syst. 33, 3976–3990 (2020)
6. N.A. Urpí, M. Bagatella, M. Vlastelica, G. Martius, Causal action influence aware counterfactual data augmentation. arXiv preprint arXiv:2405.18917 (2024)
7. M. Gupta, P. Gupta, J. Narwariya, L. Vig, G. Shroff, SCM4SR: Structural causal model-based data augmentation for robust session-based recommendation, in Proceedings of the 47th Int. ACM SIGIR Conf. on Research and Development in Inf. Retrieval (SIGIR 2024), 2609–2613 (2024)
8. W. Zhao, D. Tang, X. Chen, D. Lv, D. Ou, B. Li, K. Gai, Disentangled causal embedding with contrastive learning for recommender system. in Companion Proceedings of the ACM Web Conference 2023 (WWW 2023), 406–410 (2023)
9. Y. Wang, H. Guo, B. Chen, W. Liu, Z. Liu, Q. Zhang, R. Tang, CausalINT: Causal inspired intervention for multi-scenario recommendation, in Proceedings of the 28th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD 2022), 4090–4099 (2022)

10. S. Tan, D. Li, R. Jiang, Z. Wang, X. Yu, M. Okumura, Taming recommendation bias with causal intervention on evolving personal popularity, in Proceedings of the 31st ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD 2025) V.2, 2790–2800 (2025)
11. S. Wang, X. Chen, D. Jannach, L. Yao, Causal decision transformer for recommender systems via offline reinforcement learning, in Proceedings of the 46th Int. ACM SIGIR Conf. on Research and Development in Inf. Retrieval (SIGIR 2023), 1599–1608 (2023)
12. S. Wang, X. Chen, L. Yao, Policy-guided causal state representation for offline reinforcement learning recommendation, in Proceedings of the ACM on Web Conference 2025 (WWW 2025), 402–412 (2025)
13. S. Lee, H. Jo, Y. Jeon, B.K. Jeon, CIPHER: Causal intent plug-in framework for the mitigation of historical exposure bias in recommender systems (2025)