

On the Opportunity of Causal Learning in Recommendation Systems: Foundation, Estimation, Prediction and Challenges

Peng Wu^{1*}, Haoxuan Li^{1*}, Yuhao Deng¹, Wenjie Hu¹, Quanyu Dai², Zhenhua Dong^{2†}, Jie Sun³, Rui Zhang³ and Xiao-Hua Zhou^{1†}

¹Peking University

²Huawei Noah's Ark Lab

³Huawei Hong Kong Research Center
azhou@math.pku.edu.cn, dongzhenhua@huawei.com

Abstract

Recently, recommender system (RS) based on causal inference has gained much attention in the industrial community, as well as the states of the art performance in many prediction and debiasing tasks. Nevertheless, a unified causal analysis framework has not been established yet. Many causal-based prediction and debiasing studies rarely discuss the causal interpretation of various biases and the rationality of the corresponding causal assumptions. In this paper, we first provide a formal causal analysis framework to survey and unify the existing causal-inspired recommendation methods, which can accommodate different scenarios in RS. Then we propose a new taxonomy and give formal causal definitions of various biases in RS from the perspective of violating the assumptions adopted in causal analysis. Finally, we formalize many debiasing and prediction tasks in RS, and summarize the statistical and machine learning-based causal estimation methods, expecting to provide new research opportunities and perspectives to the causal RS community.

1 Introduction

In recent years, causal inference has attracted extensive attention from both academic and industrial communities. New theories, methods, and applications of causal inference are emerging at an alarming rate. Recommender system (RS) is a promising field for the development and application of causal inference. Many practical problems of interest in RS are essentially causal problems, such as post-view click-through rate prediction [Guo *et al.*, 2021], post-click conversion rate prediction [Zhang *et al.*, 2020; Guo *et al.*, 2021], and uplift modeling [Sato *et al.*, 2019; Sato *et al.*, 2020]. Causal recommendation approaches have several advantages over the traditional recommendation methods, including better interpretability and stability, higher accuracy, and generalization ability. Its effectiveness has been verified on both numeric

experiments and theoretical analyses across strands of literature [Wang *et al.*, 2019a].

Nevertheless, a unified causal analysis framework has not been established yet. On one hand, the existing causal approaches in RS lack a clear causal and mathematical formulation on the scientific questions of interest, inducing the prevalence of many nebulous causal concepts in this field and impeding the development of causal recommendation methods. Much confusion needs to be clarified: what exactly is being estimated, for what purpose, in which scenario, by which technique, and under what plausible assumptions. On the other hand, a conspicuous feature of the observational data in RS is the existence of various biases, which is the main obstacle to drawing causal conclusions. However, formal causal definitions of the biases in RS are still not clear, even though miscellaneous biases have been discovered and proposed in a descriptive way [Chen *et al.*, 2020]. Due to the lack of formal articulation of causal questions and types of biases, it is difficult to clearly discuss the theoretical properties, merits, and drawbacks of the debiasing approaches. And it is hard to clearly explain the assumptions underlying the methods. As a consequence, it is hard to develop new debiasing algorithms.

In this article, we aim to overcome the above limitations by proposing a causal analysis framework for RS within the potential outcome framework [Rubin, 1974; Imbens and Rubin, 2015], through which we survey and unify the existing causal-inspired recommendation methods. We provide a causal perspective on biases in RS by analyzing the causal assumptions that are potentially used but not discussed in existing studies, and then discuss the corresponding recommendation scenarios violating the above assumptions. In fact, a large number of recommendation tasks are rigorously clarified by applying the proposed causal framework. In addition, we overview statistical and machine learning-based causal estimation methods, providing many opportunities for innovative causal RS research. The main contributions of this paper are summarized as follows:

- Providing a guideline of how to define, recover and estimate a causal estimand in RS, thereby explicating the perplexing causal concepts within the potential outcome framework.
- Providing a new taxonomy and giving formal causal definitions of various biases in RS from the perspective

*Contributed equally

†Contact Author

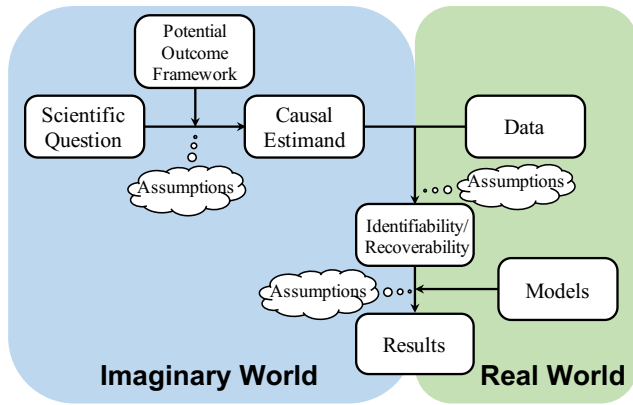


Figure 1: Causal analysis framework in RS.

of violating what assumptions are adopted in standard causal analysis.

- Revealing the key assumptions underlying various debiasing approaches, as well as distinguishing the selection bias and the confounding bias in RS.
- Applying the proposed causal analysis framework to the classical debiasing and prediction tasks in RS, and summarizing the statistical and machine learning-based causal estimation methods.
- Sharing and discussing several noteworthy open research directions for the causal RS community.

2 Causal Analysis Framework in RS

The proposed causal analysis framework refers to a unified workflow of investigating causal problems in RS, which consists of three steps: (1) Define a causal estimand to answer the scientific question; (2) Discuss the recoverability of the estimand given the data; (3) Build models to obtain the consistent estimator of the estimand. Figure 1 depicts the causal analysis framework.

One of the main contributions that causal inference has brought is a focus on clearly defined estimands before building models [Daniel *et al.*, 2016; Vansteelandt and Dukes, 2021]. Through formalizing the scientific question into a well-defined causal estimand in an imaginary world, we can answer the following questions: what exactly is being estimated and for what purpose. Yet, little literature in RS has an explicit statement of the estimand of interest.

After defining estimands, we proceed to consider whether a consistent/unbiased estimator, under suitable assumptions, can be derived from the observed data. It is equivalent to theoretically discussing the recoverability property of the estimand [Mohan and Pearl, 2021]. If the available RS data does not support deriving a consistent estimator under plausible assumptions, then we should consider how to collect new data; If the estimand is recoverable, we can explicitly present and assess the recoverability assumptions in different RS scenarios underlying the estimation approaches and then build models to estimate it.

As shown in Figure 1, we need a variety of assumptions to climb from association (data and model) to causation (causal

estimand and causal conclusion) at each stage of the causal analysis framework. Violating these assumptions may result in various biases. This perspective provides a unified way to discuss the different biases in RS. Table 1 clearly presents the most common assumptions in causal inference and establishes their connections to the biases in RS, from which we can define the descriptive biases in RS formally using the rigorous syntax of causal inference. In addition, it also provides an opportunity to apply the existing causal inference methods to RS. For example, the non-compliance problem and interference bias have been intensively studied in causal inference literature, while rarely being discussed in RS.

We emphasize that different causal problems in RS correspond to different causal estimands, and may suffer from different types of biases given the collected data. Then, various assumptions and models are needed to estimate the estimands and answer the scientific problems.

3 From Scientific Question to Causal Estimand

In this section, we provide a guideline of how to formalize a causal problem by using the potential outcome framework. The workflow of translating a scientific problem into a meaningful causal estimand is summarized as follows: (1) Define the unit; (2) Define the treatment, feature, outcome, and potential outcomes corresponding to the scientific question under study; (3) Define the target population; (4) Define the causal estimand.

The unit is the most fine-grained research subject. Unit is a terminology that is often overlooked in RS. However, a clear explanation of it is very important to define the causal estimand. In RS, a unit usually corresponds to a user-item pair [Guo *et al.*, 2021]; sometimes it is a user [Liang *et al.*, 2020] or an item [Deldjoo *et al.*, 2021]. The variety and vagueness of the unit stem from the fact that RS involves two entangled populations: users and items. Therefore, an explicit statement of the unit is helpful to eliminate ambiguity in interpretation.

For each unit, we have a treatment T , an outcome Y , and possibly a feature vector X . Usually, X is the attributes or feature embedding of the unit. However, T , X , and Y are insufficient to define a causal estimand. The potential outcome is a general syntax to formalize causal estimand, thereby translating the meaningful causal problems into causal parameters [Goetghebeur *et al.*, 2020].

Definition 1 (Potential outcome). *A potential (or counterfactual) outcome $Y(t)$ for $t \in \mathcal{T}$ is the outcome that would be observed if T had been set to t .*

The stable unit treatment value assumption (SUTVA, [Rubin, 1980]) is necessary to ensure the well-definedness of potential outcome $Y(t)$.

Assumption 1 (SUTVA). *(a) No multiple versions of treatment, only a single version of the treatment and a single version of the control; (b) Non-interference, the potential outcomes of a unit are not affected by the treatment status of the other units.*

	Assumptions	Biases in causal inference	Biases in [Chen <i>et al.</i> , 2020]
Define causal estimands	SUTVA(a) SUTVA(b)	undefined interference bias	position bias conformity bias
Recoverability	consistency positivity exchangeability conditional exchangeability random sampling	noncompliance undefined confounding bias hidden confounding bias selection bias	undefined exposure bias popularity bias undefined user/model selection bias, exposure bias
Model	model specification	model mis-specification	inductive bias

Table 1: New perspective of biases in RS.

In RS, even though there is already a certain amount of studies based on a causal perspective, most of them tacitly assume that SUTVA holds without discussion. However, in many scenarios, the SUTVA assumption does not necessarily hold. For example, the position bias can be seen as a violation of SUTVA(a). In the task of click-through rate prediction, suppose a unit is a user-item pair. Define $Y_{u,i}(1)$ as the click behavior if the item i is exposed to the user u . Then $Y_{u,i}(1)$ will rely on the position of exposure and multiple versions of treatment occur. In addition, the conformity bias means that users tend to rate similarly with others in a group. It violates SUTVA(b), i.e., non-interference. This is because the conformity phenomenon may lead to the value of $Y_m(t)$ depending on the treatment value T_j for unit $j \neq m$. Violations of SUTVA(b) in RS should receive more attention due to the existence of users’ social networks.

To clarify the population of interest, we need to specify a target population, which is the population that we want to make an inference on. We denote \mathbb{P} and \mathbb{E} as the distribution and expectation on the target population. In RS, the target population is usually the population consisting of all user-item pairs, or all users, or all items. Based on the target population, we can define the causal estimand as follows.

Definition 2 (Causal estimand). *Causal estimand is a functional of the joint distribution of treatment, feature and potential outcomes on the target population, providing a recipe for answering the scientific question of interest from any hypothetical data whenever it is available [Pearl, 2019].*

It should be noted that the definition of the causal estimand does not involve the data collected and the model adopted. More detailed examples are provided in Section 5.

4 Recoverability: From Causal Estimand to Consistent Estimation

In this section, we discuss the recoverability of the estimand, and relate the associated assumptions with the biases in RS.

Definition 3 (Recoverability of target quantity Q [Mohan and Pearl, 2021]). *Let \mathcal{A} denote the set of assumptions about the data generation process and let Q be any functional of the underlying distribution $\mathbb{P}(X, T, \{Y(t), t \in \mathcal{T}\})$. Q is recoverable if there exists a procedure that computes a consistent estimator of Q for all strictly positive observed-data distributions.*

Recoverability is a crucial ingredient in causal inference, while it is rarely discussed in RS. The significance of discussing recoverability is at least twofold: First, we can ascertain whether a consistent estimator of the counterfactual estimand can be obtained from the data available under some reasonable assumptions. Second, if the estimand is recoverable, we can explicitly present the recoverability assumptions underlying the estimation approaches. This provides a desirable perspective to evaluate the debiasing methods by assessing the assumptions and provides an opportunity to develop new approaches by weakening the assumptions.

4.1 Common Assumptions for Recoverability

Consistency and positivity are two indispensable assumptions required in most causal inference approaches for recovering the causal estimand.

Assumption 2 (Consistency). $Y(t) = \sum_{t^* \in \mathcal{T}} I(t^* = t)Y$ for any $t \in \mathcal{T}$.

Assumption 3 (Positivity). $\mathbb{P}(T = t \mid X = x) > 0$ for any t and x .

The consistency assumption implies that $Y_m(t) = Y_m$ if $T_m = t$ for each unit m . It links the potential outcomes in the hypothetical world to the observed outcomes in reality. Positivity ensures that units have a positive probability to take each treatment, and this assumption is sometimes also called “overlap” to depict the features of units overlapping in different treatment groups [Hernán and Robins, 2020].

In RS, exposure bias results from that a user is only exposed to a part of specific items. That is, some users are exposed to some items with zero probability. Therefore, exposure bias can be viewed as a violation of the positivity assumption. Noncompliance problem is prevalent in RS, while rarely discussed. A typical example is the exposure-click-conversion model, where we assume that the exposure affects conversion only through click. If the effect of exposure on conversion is of interest, the inconsistency between exposure and click, called noncompliance, would violate the consistency assumption and pose a big challenge in estimating the causal effect. More discussions are provided in Section 5.5.

Confounding bias is often caused by violation of the conditional exchangeability assumption defined as follows.

Assumption 4 (Conditional exchangeability). $Y(t) \perp T \mid X$, for any $t \in \mathcal{T}$. A stronger version is exchangeability: $Y(t) \perp T$, for any $t \in \mathcal{T}$.

Conditional exchangeability is also called ignorability or unconfoundedness [Rosenbaum and Rubin, 1983]. In the language of the causal graphical model, conditional exchangeability means that X blocks every back-door path between T and Y [Pearl and Mackenzie, 2018; Hünernmund and Bareinboim, 2019]. Besides, an underlying assumption in causal inference is that the observed samples can reflect the target population.

Assumption 5 (Random sampling). $\mathbb{P}(x, t, y) = \mathbb{P}_{\mathcal{O}}(x, t, y)$, where \mathbb{P} represents the target population distribution and $\mathbb{P}_{\mathcal{O}}$ represents the observed sample distribution.

The combination of Assumptions 1-5 can obtain the recoverability of most causal quantities by using observed samples. For example, if $\mathbb{E}[Y(t) | X = x]$ is of interest, we can reformulate it as

$$\mathbb{E}[Y(t)|X = x] = \mathbb{E}[Y(t)|X = x, T = t] = \mathbb{E}[Y|X = x, T = t], \quad (1)$$

where the first identity relies on the positivity and conditional exchangeability assumptions, and the second identity requires the consistency assumption. Based on (1), under random sampling assumption, $\mathbb{E}[Y|X = x, T = t]$ can be estimated consistently from the observed data, which has been implemented with satisfactory performance by a large number of RS literature, then the recoverability is realized.

These assumptions can be divided into associational assumptions and causal assumptions [Pearl, 2009]. The former (e.g. model specification of propensity score), is testable in principle. In contrast, the latter, such as Assumptions 1–5, cannot be directly verified from data, unless one resorts to experimental control. In addition, Assumptions 1–5 might not be guaranteed in observational studies, which require different sets of assumptions, such as introducing an instrumental variable or using the front-door criterion to achieve recoverability. In practice, whether the causal assumptions hold need to be discussed by expert’s knowledge (e.g. drawing causal graphs) for each specific problem.

4.2 Selection Bias and Confounding Bias in RS

In causal inference, selection bias and confounding bias are two of the most common barriers to achieving causal estimates [Correa *et al.*, 2019], which can be formally defined as violations of exchangeability and random sampling assumptions respectively. In RS, it is worth emphasizing that the research studies should first distinguish the two biases.

Definition 4 (Selection bias). *Selection bias means that the sample distribution is different from that of target population [Hernán and Robins, 2020], i.e.,*

$$\mathbb{P}(x, t, y) \neq \mathbb{P}_{\mathcal{O}}(x, t, y).$$

Definition 5 (Confounding bias). *Confounding bias refers to the association (T and Y) created due to the presence of factors affecting both the treatment and the outcome [Correa *et al.*, 2019], i.e., $\exists t \in \mathcal{T}, Y(t) \not\perp T$. Usually it will lead to*

$$\mathbb{E}[Y(t)] \neq \mathbb{E}[Y(t)|T = t].$$

Selection bias abounds in RS. For example, the system aims to recommend items that the user may like by filtering out items with low predicted ratings, and this kind of

selection is previously called **model selection bias** [Yuan *et al.*, 2019]; users tend to rate recommended items that they like and rarely rate recommended items that they dislike, which is the **user self-selection bias** [Saito, 2020]. In such cases, the data-gathering process will reflect a distortion in the sample’s proportions, since the data is no longer a faithful representation of the target population, and biased estimates will be produced regardless of the size of samples collected. Interestingly, the exposure bias will also lead to $\mathbb{P}(x, t, y) \neq \mathbb{P}_{\mathcal{O}}(x, t, y)$, hence belonging to the selection bias. Similar insight is founded in [Chen *et al.*, 2021].

Confounding bias differs fundamentally from the selection bias. Selection bias comes from the systematic bias during the collection of units into the sample [Bareinboim *et al.*, 2014]. A well-designed sampling procedure can reduce selection bias, such as recommending items to users randomly to obtain unbiased data, but the cost is extraordinarily expensive. In contrast, confounding bias stems from the systematic bias inherently determined by the causal mechanism (relations) among features, treatment, and outcome, irrespective of the data collection process. Randomization of treatment assignment can eliminate the effect of (unmeasured) confounding bias, but cannot remove the influence of selection bias [Correa *et al.*, 2019].

5 Applying the Proposed Causal Analysis Framework to Recommendation Tasks

In this section, we illustrate how to apply the proposed causal analysis framework to the classic recommendation tasks. Throughout, denote with $u \in \mathcal{U}$ the users and with $i \in \mathcal{I}$ the items, and denote $\mathcal{D} = \mathcal{U} \times \mathcal{I}$ as the set of all user-item pairs. Let $Y_{u,i}$ be the outcome of interest for user-item pair (u, i) , $T_{u,i}$ be the treatment, and $X_{u,i}$ be the corresponding feature embedding of user u and item i . For general treatment, define

$$\mu_t(x) = \mathbb{E}[Y(t) | X = x], \quad t \in \mathcal{T}, \quad (2)$$

and for binary treatment, define

$$\tau(x) = \mathbb{E}[Y(1) - Y(0) | X = x], \quad (3)$$

which are the common causal estimands in RS. Based on the proposed causal analysis framework, it can be divided into the following scenarios according to different causal estimands and research perspectives, claiming that most of the recommendation tasks can be included into the following scenarios.

5.1 Missing Not at Random (MNAR)

This scenario focuses on the problem of missing outcome data. Consider the case of movie rating websites [Wang *et al.*, 2019b; Wang *et al.*, 2020a]. Applying the proposed causal analysis framework, a unit is a user-item pair, the feature $X_{u,i}$ is the attributes of user u and movie i , the outcome $Y_{u,i}$ is the true rating of user u for movie i . However, the outcome suffers a problem of missing, and selection bias is induced due to the fact that the users incline to rate the movies they like. Usually, the missing mechanism is MNAR. Let $O_{u,i}$ be the observing indicator of $Y_{u,i}$. We consider the observing indicator as the treatment, then $Y_{u,i}(1)$ denotes the true rating of

user u for movie i if $O_{u,i} = 1$. The target population consists of all user-item pairs. The goal of the MNAR debiasing task is to estimate $Y_{u,i}(1)$ using feature $X_{u,i}$, i.e., the causal estimand of interest is $\mu_1(x)$.

Many studies have tried to give an accurate estimate of the $\mu_1(x)$ from the perspective of causality. [Schnabel *et al.*, 2016] use the inverse propensity score (IPS) method to recover the distribution of the target population by weighting the non-missing units with propensity score, and further introduce the self-normalized IPS (SNIPS) estimator to reduce the large variance problem caused by extremely small estimated propensity scores. [Wang *et al.*, 2019a] propose the doubly robust (DR) method and the joint learning optimization technique. Based on the DR estimator, [Guo *et al.*, 2021] propose a more robust doubly robust (MRDR) estimator to further control the variance while retaining its double robustness. In addition, [Wang *et al.*, 2020b] propose the counterfactual variational information bottleneck (CVIB) approach, and the core idea is to separate the task-aware mutual information term into factual and counterfactual parts and balance them. [Liu *et al.*, 2021] propose debiased information bottleneck (DIB) based on the debiased representation from the causal diagrams and information theory. [Wu *et al.*, 2022] propose a doubly robust collaborative targeted learning that makes the recommendation model more accurate and robust.

5.2 Binary Treatment (e.g. CTR Prediction)

This scenario discusses the case of binary treatment. Different from the MNAR scenario, this scenario has no missing outcome. An typical example is advertising recommendation [Gopalan *et al.*, 2015; Liang *et al.*, 2020], where a unit is a user-item pair, the target population consists of all user-item pairs, and the outcome $Y_{u,i}$ is the indicator of a click event, i.e., $Y_{u,i} = 1$ if user u clicks item i , $Y_{u,i} = 0$ otherwise. The treatment $T_{u,i} = 1$ if item i is exposed to user u , $T_{u,i} = 0$ otherwise, the potential outcomes $Y_{u,i}(1)$ and $Y_{u,i}(0)$ denote the indicator of click event if the item is/isn't exposed to the user u . The estimand of interest is $\mu_1(x)$ denoting the click-through rate (CTR), or $\tau(x)$ denoting the uplift of CTR.

The uplift modeling in RS [Sato *et al.*, 2019] are closely related to the binary treatment scenario, which aims to predict the change of feedback value caused by the increment of the treatment, and there are many studies that have successfully applied causal inference techniques to estimate the causal effect in uplift models accurately [Gutierrez and Gérardy, 2017; Athey and Imbens, 2015a; Hitsch and Misra, 2018].

5.3 Multi-valued Treatment (e.g. Position Bias)

Consider the scenario of multi-valued treatment, which corresponds to the contextual bandit in RS with finite action space [Li *et al.*, 2010]. We also consider the task of CTR prediction and assume it suffers from the problem of position bias [Yuan *et al.*, 2020]. If the exposed locations are available, assuming that there are K different positions, then this is a problem of multi-valued treatment. Specifically, the treatment $T_{u,i}$ has K levels, where $T_{u,i} = j$ means that the item i is exposed to user u at j -th position. Correspondingly, there are K potential outcomes for each unit. The causal estimand can be $\mu_j(x)$ or $\tau_{j,k}(x) = \mathbb{E}[Y_{u,i}(j) - Y_{u,i}(k) \mid X_{u,i} = x]$. The former

represents the CTR at position j , while the latter indicates the change in CTR from exposure to position k to position j .

5.4 Continuous Treatment (e.g. Cash Reward)

This section further extends the scenarios in Sections 5.2 and 5.3 to continuous treatment, which corresponds to the contextual bandit in RS with infinite action space [Li *et al.*, 2010]. In fact, in order to increase CTR or post-click conversion rate (CVR), many RS (like TikTok and Kuaishou) carry cash rewards when users click on an advertisement, in which the reward is usually a continuous variable. In this situation, we treat the user-item pairs as units representing the scenario of user u clicking the ads of item i , reward $T_{u,i}$ as the continuous exposure treatment, and the corresponding profit $Y_{u,i}$ as outcomes. If the goal of RS is to estimate the cumulative net income, the estimand of interest is $\mu_t(x) - t = \mathbb{E}[Y_{u,i}(t) - t \mid X_{u,i} = x]$, where t is the treatment value, representing the cash rewards for u clicking the ads of item i .

5.5 Compliance (e.g. Exposure-Click-Conversion)

The compliance scenario involves two variables C and Y measured after the treatment T , and the causal relationship among them is $T \rightarrow C \rightarrow Y$, i.e., T affects Y only through C . Consider an example of online advertising. The units are the user-item pairs, the target population is all the user-item pairs, and the feature $X_{u,i}$ is the embedding of user u and item i . $T_{u,i}$, $C_{u,i}$, and $Y_{u,i}$ are indicators of the user exposed on the item advertising, click on the item and the conversion on the item. If we treat T as treatment, then $\mathbb{E}[C(1) \mid X = x]$ denotes the CTR. If we treat C as treatment, then $\mathbb{E}[Y(1) \mid X = x]$ denotes the CVR. However, if we want to detect the effect of T on Y , then the estimand is more complicated. In such a case, we regard T as the treatment and let $C(0)$ and $C(1)$ be the potential click behaviors if $T = 1$ and $T = 0$, respectively. The definition of potential conversion depends on both T and C . Let $Y(t, c)$ be the potential conversion if $T = t$ and $C(t) = c$. Then the estimand

$$\tilde{\tau}(x) = \mathbb{E}[Y(1, C(1)) - Y(0, C(0)) \mid X = x] \quad (4)$$

measures the causal effect of T on Y .

The idea of compliance is widely used in causal inference, but it is rarely discussed in RS. [Gu *et al.*, 2021] considers the compliance framework from the perspective of advertising demanders. Since click advertising requires payment, we prefer to push advertising to user-item pairs with causal effects rather than free rider (referring to user-item pairs that will always be converted regardless of whether advertising is recommended or not), and this phenomenon is very common in popular products.

5.6 Recommendation Policy Evaluation and Learning

This scenario treats the recommendation problem as a policy learning problem, and no longer pays attention to the estimands $\mu_t(x)$ and $\tau(x)$. For reinforcement learning in RS, it is always focusing on the evaluation or optimization of recommendation strategies [Afsar *et al.*, 2021; Munemasa *et al.*,

2018]. The counterfactual framework can be further used to deal with the delayed feedback [Zhang *et al.*, 2021], which acts as a common research direction in RS. From the causal perspective, suppose that there are a total of I items and U users. A unit is a user, the feature X_u is the attribute of user u , the treatment T_u has I levels, denoted as $\mathcal{T} = \{1, 2, \dots, I\}$, where $T_u = i$ means that item i is exposed to user u . The reward caused by user u exposed to the item i as the potential outcome is denoted by $Y_u(i)$. The target population is all the users. The observed data consist of U observations of features X_u , treatment T_u , and reward $Y_u = Y_u(T_u)$. And the target quantity is the optimal policy defined by

$$\pi_0^* = \arg \max_{\pi \in \Pi_0} V(\pi),$$

where Π_0 is a policy class, $V(\pi)$ is the policy value, refers to the expectation of the reward under the policy π , i.e.,

$$V(\pi) = \mathbb{E} \left[\sum_{t \in \mathcal{T}} \pi(t|X) Y(t) \right] = \mathbb{E} \left[\sum_{t \in \mathcal{T}} \pi(t|X) \mu_t(X) \right],$$

where $\mu_t(x)$ is defined in (2).

5.7 Existing Causal Debiasing Methodologies

The methods developed in the MNAR scenario can be applied for the estimation of $\mu_t(x)$ and $\tau(x)$ in Sections 5.2 to 5.6. Besides, this section further reviews other existing debiasing methods under Assumptions 1-5. Throughout, let $\mu(x) = \mathbb{E}[Y|X = x]$, $\pi(x) = \mathbb{P}(T = 1|X = x)$, $\mu_t(x) = \mathbb{E}[Y(t)|X = x] = \mathbb{E}[Y|X = x, T = t]$ for $t = 0, 1$.

Many statistical methods can be used to estimate the causal effect, including S-learner [Hill, 2011], T-learner [Hansotia and Rukstales, 2002], U-learner [Nie and Wager, 2021], R-learner [Nie and Wager, 2021], X-learner [Kunzel *et al.*, 2019], IPW-learner [Horvitz and Thompson, 1952] and DR-learner [Kennedy, 2020]. For example, the U-learner obtains the estimator of $\tau(x)$ by regressing $(Y - \mu(X))/\pi(X)$ on X , and IPW-learner by regressing $TY/\pi(X) - (1 - T)Y/(1 - \pi(X))$ on X . Clearly, these methods rely on the model specifications of the nuisance parameters $\pi(X)$, $\mu_t(X)$ or $\mu(X)$. There are many other machine learning methods that are designed to estimate $\tau(x)$ directly, such as causal tree [Athey and Imbens, 2015b], causal forest [Wager and Athey, 2018; Lechner, 2019; Athey *et al.*, 2019; Oprescu *et al.*, 2020], causal BART [Carvalho, 2020], causal boosting and causal MARS [Powers *et al.*, 2018], balancing counterfactual regression [Johansson *et al.*, 2016], Generative Adversarial Nets [Yoon *et al.*, 2018], Causal Effect Variational Autoencoder [Louizos *et al.*, 2017], and local similarity preserved SITE [Yao *et al.*, 2018].

6 Open Research Directions

Recently, more and more researchers in RS are trying to apply causal inference methods to handle RS tasks such as CTR/CVR prediction, delayed feedback, etc., nevertheless, there are still many challenges and opportunities. By matching the existing research with the causal analysis framework discussed above, we have identified the following open research directions, which are rarely formalized in the potential outcome framework for RS to conduct research.

Data Fusion. A typical scenario involves a combination of a large biased (observational or non-uniform) dataset and a small unbiased (experimental or uniform) dataset. The biased data and unbiased data have complementary characteristics. The biased data is inevitable to suffer from the problem of hidden/unmeasured confounders, which will distort the causal conclusions even the sample size is infinity [Kallus *et al.*, 2018]. In comparison, collected through a carefully designed experiment, the unbiased data has no (hidden) confounding bias, and it provides the gold standard for evaluating the debiasing approaches. In summary, data fusion is a promising strategy to improve the quality of RS.

Sequential Recommendation. By modeling the user behavior sequence, such as the sequence of purchasing items, RS can learn the change of user interest and predict the user's next behavior. From the perspective of causality, it can be considered that the assignment mechanism and potential outcomes of RS are changing with the time series, and the goal is to dynamically capture the changes of users' interests, so as to achieve more accurate recommendations.

Fairness in RS. Many literature define group fairness and individual fairness through counterfactual causality [Kusner *et al.*, 2017; Nabi and Shpitser, 2018; Chiappa, 2019]. However, how to formalize the fairness in RS with causal framework is still vague, especially when the user has a social network, which will violate SUTVA(b), and bring greater challenges to the well-definedness of causal fairness. In addition, there is still a lot of research space on how to modify the traditional causal recommendation model to achieve the balance of accuracy and fairness.

Interference. Even though almost all articles on causality inspired recommendation acquiesce in the SUTVA assumption, as discussed in Section 3, the SUTVA assumption will be violated in many cases, resulting in biased estimation. Another form of interference is between potential outcomes of different units. For example, a user's purchase behavior will affect the purchase behavior of other users in his social network, which is often encountered in social recommendation.

7 Conclusion

Causality offers new opportunities for robust and outstanding performance of algorithms for debiasing and prediction tasks in RS. This article reviews related research by providing a unified causal analysis framework for RS, revealing and discussing in detail the validity of the always neglected but equally important causal assumptions. New interpretations of various biases in RS are provided from the perspective of violating causal assumptions. The proposed causal RS analysis framework is applied to rigorously formulate a large number of RS tasks, such as the non-compliance problem, interference bias, and policy learning, which have been intensively studied in causal inference literature.

The paper concludes with an overview of causal estimation methods that will hopefully provide new research opportunities in the field of causal recommendation including but not limited to debiasing and prediction tasks. In addition, it is expected to develop new methods with weakening or substituting the common assumptions in RS studies.

References

- [Afsar *et al.*, 2021] M Mehdi Afsar, Trafford Crump, and Behrouz Far. Reinforcement learning based recommender systems: A survey. *arXiv:2101.06286*, 2021.
- [Athey and Imbens, 2015a] Susan Athey and Guido W Imbens. Machine learning methods for estimating heterogeneous causal effects. *stat*, 1050:1–26, 2015.
- [Athey and Imbens, 2015b] Susan Athey and Guido W Imbens. Recursive partitioning for heterogeneous causal effects. *PNAS*, 113:7353–7360, 2015.
- [Athey *et al.*, 2019] Susan Athey, Julia Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47:1148–1178, 2019.
- [Bareinboim *et al.*, 2014] Elias Bareinboim, Jin Tian, and Judea Pearl. Recovering from selection bias in causal and statistical inference. In *AAAI*, 2014.
- [Carvalho, 2020] P. Richard Hahn. Jared S. Murray. Carlos M. Carvalho. Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Journal of the American Statistical Association*, 15:965–1056, 2020.
- [Chen *et al.*, 2020] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. Bias and debias in recommender system: A survey and future directions. *arXiv:2010.03240*, 2020.
- [Chen *et al.*, 2021] Jiawei Chen, Hande Dong, Yang Qiu, Xiangnan He, Xin Xin, Liang Chen, Guli Lin, and Keping Yang. Autodebias: Learning to debias for recommendation. In *SIGIR*, 2021.
- [Chiappa, 2019] Silvia Chiappa. Path-specific counterfactual fairness. In *AAAI*, 2019.
- [Correa *et al.*, 2019] Juan D. Correa, Jin Tian, and Elias Bareinboim. Identification of causal effect in the presence of selection bias. In *AAAI*, 2019.
- [Daniel *et al.*, 2016] Rhian M Daniel, Bianca L de Stavola, and Stijn Vansteelandt. Commentary: The formal approach to quantitative causal inference in epidemiology: misguided or misrepresented? *International Journal of Epidemiology*, 45:1817–1829, 2016.
- [Deldjoo *et al.*, 2021] Yashar Deldjoo, Vito Walter Anelli, Hamed Zamani, and et al. A flexible framework for evaluating user and item fairness in recommender systems. *User Modeling and User-Adapted Interaction*, 31:457–511, 2021.
- [Goetghebeur *et al.*, 2020] Els Goetghebeur, Saskia le Cessie, Bianca De Stavola, and et al. Formulating causal questions and principled statistical answers. *Statistics In Medicine*, 39:4922–4948, 2020.
- [Gopalan *et al.*, 2015] Prem Gopalan, Jake M. Hofman, and David M. Blei. Scalable recommendation with hierarchical poisson factorization. In *UAI*, 2015.
- [Gu *et al.*, 2021] Tiankai Gu, Kun Kuang, Hong Zhu, and et al. Estimating true post-click conversion via group-stratified counterfactual inference. In *ADKDD*, 2021.
- [Guo *et al.*, 2021] Siyuan Guo, Lixin Zou, Yiding Liu, et al. Wang, Hechang Chen, Dawei Yin, and Yi Chang. Enhanced doubly robust learning for debiasing post-click conversion rate estimation. In *SIGIR*, 2021.
- [Gutierrez and Gérardy, 2017] Pierre Gutierrez and Jean-Yves Gérardy. Causal inference and uplift modelling: A review of the literature. In *International Conference on Predictive Applications and APIs*, pages 1–13, 2017.
- [Hansotia and Rukstales, 2002] Behram Hansotia and Brad Rukstales. Incremental value modeling. *Journal of Interactive Marketing*, 16:35–46, 2002.
- [Hernán and Robins, 2020] Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Boca Raton: Chapman and Hall/CRC, 2020.
- [Hill, 2011] Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20:217–240, 2011.
- [Hitsch and Misra, 2018] Günter J Hitsch and Sanjog Misra. Heterogeneous treatment effects and optimal targeting policy evaluation. Available at SSRN 3111957, 2018.
- [Horvitz and Thompson, 1952] Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47:663–685, 1952.
- [Hünermund and Bareinboim, 2019] Paul Hünermund and Elias Bareinboim. Causal inference and data-fusion in econometrics. *arXiv:1912.09104*, 2019.
- [Imbens and Rubin, 2015] Guido W. Imbens and Donald B. Rubin. *Causal Inference For Statistics Social and Biomedical Science*. Cambridge University Press, 2015.
- [Johansson *et al.*, 2016] Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *ICML*, 2016.
- [Kallus *et al.*, 2018] Nathan Kallus, Ahlad Manas Puli, and Uri Shalit. Removing hidden confounding by experimental grounding. In *NIPS*, 2018.
- [Kennedy, 2020] Edward H. Kennedy. Optimal doubly robust estimation of heterogeneous causal effects. *arXiv:2004.14497v1*, 2020.
- [Kunzel *et al.*, 2019] Soren R Kunzel, Jasjeet S Sekhon, Peter J Bickel, and Bin Y. Metalearners for estimating heterogeneous treatment effects using machine learning. *PNAS*, 116:4156–4165, 2019.
- [Kusner *et al.*, 2017] Matt J Kusner, Joshua R Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *arXiv preprint arXiv:1703.06856*, 2017.
- [Lechner, 2019] Michael Lechner. Modified causal forests for estimating heterogeneous causal effects. *arXiv:1812.09487*, 2019.
- [Li *et al.*, 2010] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *WWW*, 2010.

- [Liang *et al.*, 2020] Dawen Liang, Laurent Charlin, and David M. Blei. Causal inference for recommender systems. In *RecSys*, 2020.
- [Liu *et al.*, 2021] Dugang Liu, Pengxiang Cheng, Hong Zhu, and et al. Mitigating confounding bias in recommendation via information bottleneck. In *RecSys*, 2021.
- [Louizos *et al.*, 2017] Christos Louizos, Uri Shalit, Joris Mooij, and et al. Causal effect inference with deep latent-variable models. *arXiv:1705.08821*, 2017.
- [Mohan and Pearl, 2021] Karthika Mohan and Judea Pearl. Graphical models for processing missing data. *Journal of the American Statistical Association*, 2021.
- [Munemasa *et al.*, 2018] Isshu Munemasa, Yuta Tomomatsu, Kunioki Hayashi, and et al. Deep reinforcement learning for recommender systems. In *ICOACT*, 2018.
- [Nabi and Shpitser, 2018] Razieh Nabi and Ilya Shpitser. Fair inference on outcomes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [Nie and Wager, 2021] Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108:4156–4165, 2021.
- [Oprescu *et al.*, 2020] Miruna Oprescu, Vasilis Syrgkanis, and Zhiwei Steven Wu. Orthogonal random forest for causal inference. *arXiv:1806.03467*, 2020.
- [Pearl and Mackenzie, 2018] Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. New York: Basic Books, 2018.
- [Pearl, 2009] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, second edition, 2009.
- [Pearl, 2019] Judea Pearl. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62:54–60, 2019.
- [Powers *et al.*, 2018] Scott Powers, Junyang Qian, Kenneth Jung, and et al. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in medicine*, 37:1767–1787, 2018.
- [Rosenbaum and Rubin, 1983] Paul R. Rosenbaum and Donald B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.
- [Rubin, 1974] Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational psychology*, 66:688–701, 1974.
- [Rubin, 1980] Donald B. Rubin. Discussion of randomization analysis of experimental data in the fisher randomization test by basu. *Journal of the American Statistical Association*, 75:591–593, 1980.
- [Saito, 2020] Yuta Saito. Asymmetric tri-training for debiasing missing-not-at-random explicit feedback. In *SIGIR*, pages 309–318, 2020.
- [Sato *et al.*, 2019] Masahiro Sato, Janmajay Singh, Sho Takemori, and et al. Uplift-based evaluation and optimization of recommenders. In *RecSys*, 2019.
- [Sato *et al.*, 2020] Masahiro Sato, Sho Takemori, Janmajay Singh, and Tomoko Ohkuma. Unbiased learning for the causal effect of recommendation. In *RecSys*, 2020.
- [Schnabel *et al.*, 2016] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *ICML*, 2016.
- [Vansteelandt and Dukes, 2021] Stijn Vansteelandt and Oliver Dukes. Discussion of kallus and mo, qi, and liu: New objectives for policy learning. *Journal of the American Statistical Association*, 116:675–679, 2021.
- [Wager and Athey, 2018] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113:1228–1242, 2018.
- [Wang *et al.*, 2019a] Xiaojie Wang, Rui Zhang, Yu Sun, and Jianzhong Qi. Doubly robust joint learning for recommendation on data missing not at random. In *ICML*, 2019.
- [Wang *et al.*, 2019b] Yixin Wang, Dawen Liang, Laurent Charlin, and David M. Blei. The deconfounded recommender: A causal inference approach to recommendation. *arXiv:1808.06581*, 2019.
- [Wang *et al.*, 2020a] Yixin Wang, Dawen Liang, Laurent Charlin, and David M. Blei. Causal inference for recommender systems. In *RecSys*, 2020.
- [Wang *et al.*, 2020b] Zifeng Wang, Xi Chen, Rui Wen, and et al. Information theoretic counterfactual learning from missing-not-at-random feedback. *NeurIPS*, 2020.
- [Wu *et al.*, 2022] Peng Wu, Haoxuan Li, Yan Lyu, and Xiaohua Zhou. Doubly robust collaborative targeted learning for recommendation on data missing not at random. *arXiv:2203.10258*, 2022.
- [Yao *et al.*, 2018] Liuyi Yao, Sheng Li, Yaliang Li, and et al. Representation learning for treatment effect estimation from observational data. *NeurIPS*, 31, 2018.
- [Yoon *et al.*, 2018] Jinsung Yoon, James Jordon, and Michaela Van Der Schaar. Ganite: Estimation of individualized treatment effects using generative adversarial nets. In *ICLR*, 2018.
- [Yuan *et al.*, 2019] Bowen Yuan, Juiyang Hsia, Mengyuan Yang, and et al. Improving ad click prediction by considering non-displayed events. In *CIKM*, 2019.
- [Yuan *et al.*, 2020] Bowen Yuan, Yaxu Liu, Jui-Yang Hsia, and et al. Unbiased ad click prediction for position-aware advertising systems. In *RecSys*, 2020.
- [Zhang *et al.*, 2020] Wenhao Zhang, Wentian Bao, Xiaoyang Liu, and et al. Large-scale causal approaches to debiasing post-click conversion rate estimation with multi-task learning. In *WWW*, 2020.
- [Zhang *et al.*, 2021] Xiao Zhang, Haonan Jia, Hanjing Su, Wenhan Wang, Jun Xu, and Ji-Rong Wen. Counterfactual reward modification for streaming recommendation with delayed feedback. In *SIGIR*, pages 41–50, 2021.