

Top-N Recommendation with Counterfactual User Preference Simulation

Mengyue Yang^{1,#}, Quanyu Dai^{2,#}, Zhenhua Dong², Xu Chen^{3,4,*}, Xiuqiang He², Jun Wang¹

¹ University College London ² Noah's Ark Lab, Huawei ³ Beijing Key Laboratory of Big Data Management and Analysis Methods ⁴ Gaoling School of Artificial Intelligence Renmin University of China

ABSTRACT

Top-N recommendation, which aims to learn user ranking-based preference, has long been a fundamental problem in a wide range of applications. Traditional models usually motivate themselves by designing complex or tailored architectures based on different assumptions. However, the training data of recommender system can be extremely sparse and imbalanced, which poses great challenges for boosting the recommendation performance. To alleviate this problem, in this paper, we propose to reformulate the recommendation task within the causal inference framework, which enables us to counterfactually simulate user ranking-based preferences to handle the data scarce problem. The core of our model lies in the counterfactual question: “what would be the user’s decision if the recommended items had been different?”. To answer this question, we firstly formulate the recommendation process with a series of structural equation models (SEMs), whose parameters are optimized based on the observed data. Then, we actively indicate many recommendation lists (called intervention in the causal inference terminology) which are not recorded in the dataset, and simulate user feedback according to the learned SEMs for generating new training samples. Instead of randomly intervening on the recommendation list, we design a learning-based method to discover more informative training samples. Considering that the learned SEMs can be not perfect, we, at last, theoretically analyze the relation between the number of generated samples and the model prediction error, based on which a heuristic method is designed to control the negative effect brought by the prediction error. Extensive experiments are conducted based on both synthetic and real-world datasets to demonstrate the effectiveness of our framework.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommender Systems, Bayesian Personalized Ranking, Structure Causal Model, Counterfactuals

Equal Contribution. * 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 ACM 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](https://permissions.acm.org).

CIKM '21, November 1–5, 2021, Virtual Event, Australia.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8446-9/21/11...\$15.00

<https://doi.org/10.1145/3459637.3482305>

ACM Reference Format:

Mengyue Yang^{1,#}, Quanyu Dai^{2,#}, Zhenhua Dong², Xu Chen^{3,4,*}, Xiuqiang He², Jun Wang¹. 2021. Top-N Recommendation with Counterfactual User Preference Simulation. In *Proceedings of the 30th ACM Int'l Conf. on Information and Knowledge Management (CIKM '21)*, November 1–5, 2021, Virtual Event, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3459637.3482305>

1 INTRODUCTION

Recommendation system basically aims to match a user with her most favorite items. In a typical recommendation process, the system firstly recommends an item list to a user, and then the user provides feedback on the recommendations. In real-world scenarios, people only access a small amount of items, which makes the observed dataset extremely sparse. Recommendation task is usually formulated as a ranking problem. Early models mostly base themselves on the simple matrix factorization method [26]. In order to achieve better performance and adapt different scenarios, recent years have witnessed much effort on neuralizing the recommendation models [18, 39]. However, the effectiveness of neural models usually depend on a large amount of training samples, which contradicts with the aforementioned sparse user behaviors.

In another research line, causal inference (CI) has been recently introduced into the machine learning community to augment the training data for more comprehensive model optimization [1]. The basic idea is firstly assuming an underlying structure causal model (SCM), and then learning the model parameters based on the observed data. At last, the new training samples are generated by actively changing the input variables (called intervention) and collecting the cared outputs. Such sample enrichment method has been successfully applied to the fields of neural language processing (NLP) [41] and computer vision (CV) [2, 6, 12]. In this paper, we adapt this method to the recommendation domain, which is expected to alleviate the contradiction between the more and more heavier neural recommender architectures and the sparse user behaviors. In a nutshell, the main building block of our idea lies in the counterfactual question: “what would be the user’s feedback if the recommendation list had been different?”. More specifically, we formulate the recommendation task by a causal graph including three nodes (see Figure 1(b)): \mathbf{U} represents the user profiles, \mathbf{R} is the recommendation list, and \mathbf{S} indicates the selection of the user from \mathbf{R} . The new training samples are generated by collecting the user feedback \mathbf{S} with different interventions on \mathbf{R} .

While the counterfactual idea seems to be promising, there are many challenges when applying it to the recommendation domain: to begin with, how to formally define the recommendation task by a structure causal model is still unclear. Then, the space of the candidate recommendation lists (\mathbf{R}) can be very large, and the samples induced from different recommendation lists (i.e., intervention)




Figure 1: (a) is the general framework of our idea. (b) is the structure causal model of the recommender simulator. (c) and (d) are two different intervention methods.

may vary on the effects of optimizing the target model [14, 31]. How to design an effective method to select \mathbf{R} remains to be an open problem. At last, the predefined structural equation models can be not perfect. How the prediction error influences the quality of the generated samples and how to lower the negative impact need to be carefully considered. For solving these challenges, in this paper, we propose a novel counterfactual personalized ranking framework (called **CPR**). In general, our framework is composed of two parts (see Figure 1(a)), i.e., the target ranking model and the recommender simulator. The ranking model is leveraged to provide the final recommendation list, and the recommender simulator aims to assist the optimization of the ranking model by generating additional training samples. When building the recommender simulator, we follow Pearl’s [23] counterfactual framework to define the recommendation process, where the structural equation models (SEMs) between \mathbf{U} , \mathbf{R} and \mathbf{S} are defined in a stochastic manner and learned by variational inference to capture the randomness in the recommender system. In order to handle the extremely large space of the recommendation list (i.e., \mathbf{R}), we design a learning-based method to select \mathbf{R} which can induce more informative training samples. Considering that the learned SEMs can be not perfect, we theoretically analyze the relation between its prediction error and the number of generated samples. Inspired by this theory, we propose a simple but effective strategy to control the quality of the generated samples.

The main contributions of this paper are summarized as follows:

- (1) We formulate the recommendation problem within Pearl’s causal inference framework, which allows us to generate more training samples for more sufficient model optimization.
- (2) We design a learning-based intervention method, which can lead to more informative training samples for optimizing the target ranking model.
- (3) We theoretically analyze the relation between the potential prediction error of the structural equation models and the number of generated samples.
- (4) Inspired by the above theory, we propose a heuristic method to control the quality of the generated samples.
- (5) We conduct extensive experiments based on both synthetic and real-world datasets to verify the effectiveness of our model.

2 PRELIMINARIES

In this section, we firstly recapitulate the key concepts and methodologies in Pearl’s causal inference framework. Then, we briefly introduce the ranking-based recommender models.

2.1 Pearl’s Causal Inference Framework

Pearl’s causal inference framework holds the promise of providing a complete and self-contained tool for studying causalities under both experimental and observational settings [25]. It is famous for the proposed three layer causal hierarchy, i.e., association, intervention and counterfactual. In Pearl’s causal inference framework, the first key concept is the structure causal model (SCM), which helps to formulate real-world problems with causal languages.

Definition 1 (Structure Causal Model [23]). A structural causal model M is composed of a causal graph G and a set of structure equation models (SEMs) F . G is usually a directed graph, where the edges indicate the causal relations between different variables. The nodes in G are classified into two groups: (i) **exogenous nodes** $U = \{u_1, \dots, u_{N_U}\}$, which are independent with each other, and summarize the environment when the data was generated, and (ii) **endogenous nodes** $X = \{x_1, \dots, x_{N_X}\}$, which corresponds the variables we need to model in the problem. $F = \{f_i\}_{i=1}^N$ basically parameterizes the node relations, that is, $x_i = f_i(PA_i, u_i)$, where PA_i is the parent of x_i in G .

The structure causal model builds the basis of studying the three layer causal hierarchy. In this paper, we mainly focus on counterfactual estimation. In general, this tool aims to predict the outcome if several input variables in the causal graph had been different. We briefly introduce the estimation process in the following definition:

Definition 2 (Counterfactual Estimation [23]). Suppose we have two variable sets $Y, Z \subset X$, and we use small letters y and z to denote their instantiations. The notation $p(y_{Z=z}^M(z'))$ defines the distribution of Y if Z had been set as z in structural causal model M given that Z is currently observed as z' . This distribution corresponds the counterfactual question “what would be Y if Z had been set as z given its current observation z' ?”, which is typically derived according to the following three steps: (i) **Abduction**: deriving the posterior of the exogenous variables $p(U|Z = z')$ based on the prior $p(U)$ and the observation $Z = z'$. (ii) **Action**: modifying G by removing the edges going into Z and set $Z = z$ (called intervention) to derive $p(y|Z = z, U)$. (iii) **Prediction**: computing the distribution $P(y_{Z=z}^M(z'))$ by $\int_U p(y|Z = z, U)p(U|Z = z')dU$.

The key motivation of the above “abduction-action-prediction” procedure is to make sure that *the new (counterfactual) and observed data enjoy the same generation environment, so that they can be compatible and just like to be produced at the same time* [25]. Since Pearl’s causal inference framework is not the focus of this paper, we refer the readers to [23] for more details. Based on the tool of counterfactual estimation, one can generate additional training samples to assist the downstream tasks, where a remaining problem is how to indicate z to obtain effective samples.

2.2 Ranking-based Recommender Models

Ranking-based recommender models are powerful tools for solving the Top-N recommendation task, which can be optimized in either

pair-wise or point-wise manners. In the pair-wise method, there are usually three inputs, i.e., a user u and a positive-negative item pair (i, j) . The goal is to maximize the user preference margin between the positive and negative items. Suppose the user preference is estimated from a function $f(\cdot)$, then the optimization target is:

$$L_1(O) = - \sum_{(u,i,j) \in O} \log \sigma(f(u, i) - f(u, j)), \quad (1)$$

where $\sigma(\cdot)$ is the sigmoid function to avoid trivial solutions. $O = \{(u, i, j) | i \in \mathcal{I}_u^+, j \in \mathcal{I} \setminus \mathcal{I}_u^+\}$ is the set of training samples. \mathcal{I} is the whole item set, and \mathcal{I}_u^+ indicates the set of items which have received u 's positive feedback. For the point-wise method, the input is a user-item pair (u, i) , and the user preference estimation is treated as a classification problem, where f is optimized by the following cross entropy objective:

$$L_2(O) = - \sum_{(u,i) \in O^+} \log f(u, i) - \sum_{(u,i) \in O^-} (1 - \log(f(u, i))), \quad (2)$$

where $O^+ = \{(u, i) | i \in \mathcal{I}_u^+\}$ and $O^- = \{(u, i) | i \in \mathcal{I} \setminus \mathcal{I}_u^+\}$ are the sets of positive and negative samples. $O = O^+ \cup O^-$ denotes the complete training set. Ranking-based recommender models differentiate themselves by the various implementation of f . Simple f can be realized by matrix factorization [26], and more advanced f includes multi-layer perceptron [18, 21], attentive neural network [17] and graph convolutional network [16], among others.

3 THE CPR FRAMEWORK

Our framework is composed of two parts (see Figure 1(a)): one is the recommender simulator, which is responsible for generating new training samples. The other is the target ranking model, which is learned based on both of the observed and generated data and leveraged to provide the final recommendation list. Our framework can be applied to any ranking-based recommender models, and we detail the recommender simulator in the following sections.

3.1 Recommender Simulator

To begin with, we reformulate the recommendation problem by a structure causal model $M = \{G, F\}$. In specific, the causal graph G is defined as follows (see Figure 1(b)): (1) \mathbf{U} , \mathbf{R} and \mathbf{S} are the nodes representing the user, the recommendation list and the positive items¹ selected by the user. (2) $\mathbf{U} \rightarrow \mathbf{R}$ encodes the fact that the recommendation list is generated according to the user preference. (3) $\mathbf{U} \rightarrow \mathbf{S}$ and $\mathbf{R} \rightarrow \mathbf{S}$ indicate that the positive items are jointly determined by the recommendation list and user preference.

The structure equation models F are defined in a stochastic manner as follows:

$$F : \begin{cases} \mathbf{R} \sim p_R(\mathbf{R} | \mathbf{U}, \boldsymbol{\alpha}) \\ \mathbf{S} \sim p_S(\mathbf{S} | \mathbf{U}, \mathbf{R}, \boldsymbol{\beta}) \\ \boldsymbol{\alpha}, \boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \end{cases} \quad (3)$$

where \mathbf{U} , \mathbf{R} and \mathbf{S} are endogenous variables. $\boldsymbol{\alpha} \in \mathbb{R}^{|\mathcal{I}|}$ and $\boldsymbol{\beta} \in \mathbb{R}^K$ are exogenous variables. p_R is the probability of recommending item list \mathbf{R} . p_S is the probability of selecting \mathbf{S} as the positive feedback. $\mathcal{N}(\mathbf{0}, \mathbf{I})$ is the standard Gaussian distribution.

¹In practice, the positive feedback can be click, purchase, download, and etc.

Remark. (1) In order to consider the potential noisy information and randomness in the recommender system, the structure equation models (i.e., F) are defined in a stochastic manner, which helps to learn user preference in a more accurate and robust manner. (2) The exogenous variables in the recommendation problem can be explained as the conditions (e.g., system status, user habit, etc.) which induces the currently observed data. Take movie recommendation as an example, the users may be more likely to watch movies in their leisure time. So if the data is collected at night, then the movie click probability should be higher. Such temporal condition is latent but can influence the data generation process, which is expected to be captured by the exogenous variables, and recovered by inferring the corresponding posteriors.

3.1.1 Specification of F . We implement F by considering the unique characters of the recommender system. Recall that p_R is the probability of recommending item list \mathbf{R} . Suppose we have $|\mathcal{I}|$ items in the system, then the number² of candidate item lists is $C_{|\mathcal{I}|}^{|\mathbf{R}|}$, which is extremely large in the recommendation task. To alleviate this problem, we ignore the combinatorial effect between different items, and the probability of recommending an item list can be factorized into the multiplication of the single item recommendation probabilities, that is:

$$p_R(\mathbf{R} = \mathbf{r} | \mathbf{U} = u, \boldsymbol{\alpha}) = \prod_{k=1}^K p(R_k = r_k | \mathbf{U} = u, \boldsymbol{\alpha}) \\ = \prod_{k=1}^K \frac{\exp(P_u^T \mathbf{Q}_{r_k} + w_{r_k}^R \alpha_{r_k})}{\sum_{j=1}^{|\mathcal{I}|} \exp(P_u^T \mathbf{Q}_j + w_j^R \alpha_j)} \quad (4)$$

where $\mathbf{r} = \{r_1, r_2, \dots, r_K\}$ is the recommendation list. For the single item recommendation probability, we predict it by a softmax operator. We use $\mathbf{P} \in \mathbb{R}^{|\mathcal{U}| \times d_R}$ and $\mathbf{Q} \in \mathbb{R}^{|\mathcal{I}| \times d_R}$ to represent the embeddings of the users and items. They can be initialized by the ID or profile information of the users and items. d_R is the embedding size, and $w^R \in \mathbb{R}^{|\mathcal{I}|}$ is a weighting parameter.

p_S represents the probability of selecting item set \mathbf{S} from \mathbf{R} , which is specified as:

$$p_S(\mathbf{S} = \mathbf{s} | \mathbf{U} = u, \mathbf{R} = \mathbf{r}, \boldsymbol{\beta}) = \prod_{t=1}^M \frac{\exp(X_u^T Y_{s_t} + w_t^S \beta_t)}{\sum_{j=1}^K \exp(X_u^T Y_{r_j} + w_j^S \beta_j)} \quad (5)$$

where $\mathbf{s} = \{s_1, s_2, \dots, s_M\}$ is the set of selected items. $w^S \in \mathbb{R}^K$ is a weighting parameter. The user and item properties are encoded by the metrics of $\mathbf{X} \in \mathbb{R}^{|\mathcal{U}| \times d_S}$ and $\mathbf{Y} \in \mathbb{R}^{|\mathcal{I}| \times d_S}$, respectively, and d_S is the embedding size. In p_R and p_S , we represent the users and items with different embedding metrics. This can more flexibly characterize the system impression model (i.e., $\mathbf{U} \rightarrow \mathbf{R}$ in Figure 1(b)) and the user response model (i.e., $\mathbf{U} \rightarrow \mathbf{S}$ in Figure 1(b)), which usually hold different promises in real recommender systems.

3.1.2 Learning F . Suppose the training set is $O = \{(u_i, \mathbf{r}_i, \mathbf{s}_i)\}_{i=1}^N$, where u_i is the target user, \mathbf{r}_i is the recommendation list, \mathbf{s}_i is the item set selected by u_i from \mathbf{r}_i , N is the sample number. For p_R , since the number of items (i.e., $|\mathcal{I}|$) can be very large, it is hard to directly optimize the softmax term in equation (4). Thus we resort

²Here, if we consider the order of the recommended item, the candidate item space can be even larger.

to the negative sampling technique [18, 26], which induces the following optimization target:

$$L_R = \sum_{i=1}^N \sum_{k \in r_i} \log(\sigma(\mathbf{P}_{u_i}^T \mathbf{Q}_k + w_k^R \alpha_k)) + \sum_{k^- \in r_i^-} \log(1 - \sigma(\mathbf{P}_{u_i}^T \mathbf{Q}_{k^-} + w_{k^-}^R \alpha_{k^-})) \quad (6)$$

where $\{P, Q, w^R\}$ are the parameters to be learned. r_i^- is the set of negative samples, which are randomly sampled from the non-recommended items. $\alpha = [\alpha_i]$ is sampled³ from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. With this objective, the parameters are optimized to maximize the recommended item probability and simultaneously minimize the ones which are not presented to the users.

For p_S , since the length of the recommendation list (i.e., K) is usually not large, we directly maximize the following likelihood without approximation:

$$L_S = \sum_{i=1}^N \sum_{l=1}^M \log\left(\frac{\exp(X_{u_i}^T Y_{s_{il}} + w_l^S \beta_l)}{\sum_{j=1}^K \exp(X_{u_i}^T Y_{r_{ij}} + w_j^S \beta_j)}\right) \quad (7)$$

where $\{X, Y, w^S\}$ are the parameters to be optimized. s_{il} and r_{ij} are the l th and j th items in s_i and r_i , respectively. $\beta = [\beta_i]$ is sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. It seems that this objective only maximizes the probability of the items which are selected by the user, but with the softmax operator, the probabilities of the non-selected items are automatically lowered.

3.1.3 Counterfactual estimation based on F . Once we have learned p_R and p_S , we follow Pearl’s “abduction-action-prediction” procedure [24] to conduct counterfactual estimation. In the step of abduction, we estimate the posterior of α and β given the observed dataset O . For β , the posterior can be derived based on the following Bayesian rules:

$$p(\beta|O) \propto p(\beta, O) = p(\beta)p(O|\beta) \quad (8)$$

where $p(O|\beta)$ can be easily derived based on equation (5). Recall that our goal is to sample from the posteriors. However, the result of equation (8) is too complex for sampling, which motivates us to use variational inference [3] for approximation. In specific, we firstly define a Gaussian distribution $q_\phi(\beta) \sim \mathcal{N}(\mu, \sigma)$, where $\phi = \{\mu, \sigma\}$ is set of learnable parameters. Then we optimize ϕ by minimizing the KL-divergence between $q_\phi(\beta)$ and $p(\beta|O)$, where the evidence lower bound (ELBO) [3] we need to maximize is:

$$\text{ELBO} = \mathbb{E}_{q_\phi(\beta)} [\log p(\beta, O)] - \mathbb{E}_{q_\phi(\beta)} [\log q_\phi(\beta)] \quad (9)$$

Similarly, we can learn a variational distribution for $p(\alpha|O)$. Due to the space limitation, readers are referred to [3] for more technical details of the variational inference.

In the action step, we select a user \hat{u} , and set $R = \hat{r}$. When making prediction, we firstly sample $\hat{\beta}$ from $q_\phi(\beta)$, and then compute the probability of item \hat{r}_k by:

$$p_S(\hat{r}_k|U = \hat{u}, R = \hat{r}, \hat{\beta}) = \frac{\exp(X_{\hat{u}}^T Y_{\hat{r}_k} + w_k^S \hat{\beta}_k)}{\sum_{j=1}^K \exp(X_{\hat{u}}^T Y_{\hat{r}_j} + w_j^S \hat{\beta}_j)} \quad (10)$$

At last, \hat{s} is predicted by collecting M items with the highest probabilities.

³For improving the accuracy, α is sampled multiple times for implementation, and the loss is averaged across different α ’s.

3.2 Learning-based Intervention

Based on the above designed recommender simulator, one can attempt different \hat{r} ’s and derive the corresponding \hat{s} ’s to form new (counterfactual) training samples. A nature question is how to set \hat{r} . Straightforwardly, one can explore \hat{r} in a random manner (e.g., show the users with random recommendation lists). However, as mentioned before, the space of \hat{r} can be extremely large, and it is well known that different training samples are not equally important for model optimization [31], thus the random method can be less effective in hitting the optimal samples. In order to achieve better optimization results, we design a learning-based method to select \hat{r} . Our key idea is to make the generated samples more informative for the target ranking model. It has been studied in the previous work [12, 13] that the samples with larger loss can usually provide more knowledge for the model to learn (i.e., more informative). They can well challenge the model and bring more inspirations to improve the performance. Following these studies, we use the loss of the target ranking model as the reward, and build a learning-based method to generate \hat{r} .

Formally, suppose the target ranking model is f , and we denote its loss function as L_f , which can be specified as equation (1) or (2). The goal of the agent is to conduct actions (generating recommendation lists) which can lead to larger L_f . Considering that the action space can be very large, we follow the previous work [40] to learn a Gaussian policy to generate the continuous item center of \hat{r} , after which the discrete item IDs are recovered based on equation (4). The final learning objective is:

$$L_{\text{Agent}}(\theta) = \mathbb{E}_{\hat{r}_t \in \pi(\cdot|\hat{u}, \theta)} \left[\sum_{t=1}^T L_f(C(\hat{r}_t)) \right] \quad (11)$$

where π is the Gaussian policy implemented as a two-layer fully connected neural network with ReLU as the activation function. \hat{u} indicates the target user. $L_f(C(\hat{r}_t))$ denotes the loss of the generated training samples $C(\hat{r}_t)$. We generate T sets of training samples, and each one is derived from \hat{r}_t based on the following steps:

- Deriving \hat{r} by selecting K items near the item center \hat{r}_t according to the scores $\{(\hat{r}_t^T \mathbf{Q}_k + w_k^R \alpha_k)\}_{k=1}^I$.
- Deriving \hat{s} by selecting M items with the largest probabilities of $p_S(\hat{r}_k|U = \hat{u}, R = \hat{r}, \hat{\beta})$ ($k \in [1, K]$).
- For objective (1), $C(\hat{r}_t) = \{(\hat{u}, i, j) | i \in \hat{s}, j \in \hat{r} \setminus \hat{s}\}$. For objective (2), $C(\hat{r}_t) = \{(\hat{u}, i) | i \in \hat{s}\} \cup \{(\hat{u}, i) | i \in \hat{r} \setminus \hat{s}\}$.

We summarize the complete learning process in Algorithm 1 and 2. In specific, the target ranking model is firstly trained based on the original dataset. And then, we generate many counterfactual training samples based on the Gaussian policy. At last, the target ranking model is retrained based on the generated datasets.

3.3 Theoretical Analysis

In this section, we theoretically analyze the proposed framework within the probably approximately correct (PAC) learning framework. We focus on the pair-wise learning objective, and the conclusions can be easily extended to the point-wise case.

3.3.1 Theoretical insights on the learning-based intervention method. The key motivation of the learning-based method is to generate harder samples, so that the target model can be more informed and

Algorithm 1: Learning algorithm of CPR

- 1 Learn the SEMs based on objective (6) and (7).
 - 2 Train the target model f based on the original dataset.
 - 3 Pre-train the Gaussian policy π according to Algorithm 2.
 - 4 **for** Iteration number k in $[1, K]$ **do**
 - 5 **for** each user \hat{u} **do**
 - 6 Select an action $\hat{\tau}_t$ according to $\pi(\hat{\tau}_t|\hat{u}, \theta)$.
 - 7 Derive $C(\hat{\tau}_t)$ based on $\hat{\tau}_t$.
 - 8 Train f based on $C(\hat{\tau}_t)$.
 - 9 **end**
 - 10 **end**
 - 11 Making recommendation based on the ranking model f .
-

Algorithm 2: Learning algorithm of π

- 1 Initialize the parameter set θ in the Gaussian policy.
 - 2 **for** episode number in $[1, K]$ **do**
 - 3 **for** t in $[1, T]$ **do**
 - 4 Sample a user \hat{u} .
 - 5 Select an action $\hat{\tau}_t$ according to $\pi(\hat{\tau}_t|\hat{u}, \theta) + N_t$,
 where N_t is an exploration noise.
 - 6 Sample α and β from their posteriors.
 - 7 Derive $C(\hat{\tau}_t)$ and the reward $L_f(C(\hat{\tau}_t))$ based on $\hat{\tau}_t$.
 - 8 **end**
 - 9 $\theta \leftarrow \theta + \alpha[\sum_{t=1}^T L_f(C(\hat{\tau}_t))\nabla_{\theta} \log(\pi(r_t|\theta))]$.
 - 10 **end**
-

achieve better performance. The effectiveness of this idea has been empirically demonstrated in the previous work [13, 31]. Here, we provide a theoretical justification based on the following theory:

Theorem 1. *Suppose the users’ feedback on an item pair (i, j) is probabilistic, and we can observe $i > j$ and $i < j$ with the probabilities of η and $1 - \eta$, respectively. Then, η can actually measure the hardness of the sample, when η is closer to $\frac{1}{2}$, then the sample is harder, since the propensity between the items is more ambiguous. Suppose $\epsilon, \delta \in (0, 1)$, and we use a simple voting mechanism to determine the relation between i and j , then we need to have $\frac{\log \frac{1}{\delta}}{2(1-2\eta)^2}$ samples on (i, j) to ensure that the prediction error is smaller than δ .*

This theory suggests that we need to generate more samples (i.e., larger $\frac{\log \frac{1}{\delta}}{2(1-2\eta)^2}$) for the harder (i.e., smaller η) item pairs. It agrees with our goal in learning-based intervention method, where the produced samples are expected to be hard and can well challenge the target model. We present the proof of this theory in the Appendix.

3.3.2 Imperfection of SEMs. Since the new samples are generated based on the learned SEMs, careful readers may have concerns on how would the approximation error of F influence the sample generation. To shed lights on this problem, we theoretically analyze the relation between the prediction error of F and the number of generated samples. To begin with, we assume that the recommender simulator can recover the true ranking of the item pairs with a noisy parameter $\zeta \in (0, 0.5)$, i.e., suppose the true triplet is (u, i, j) , then the recommender simulator generates the true (i.e.,

(u, i, j)) and wrong (i.e., (u, j, i)) samples with the probabilities of $1 - \zeta$ and ζ , respectively. As the special cases, $\zeta = 0$ means the recommender simulator is perfect, and there is no noisy information in the produced samples. $\zeta = 0.5$ means the recommender simulator is a totally random predictor, and the generated data is nothing but noise. Then, we have the following theory:

Theorem 2. *Suppose $\epsilon, \delta \in (0, 1)$, and $f \in \mathcal{F}$ is the target ranking model. If the number of training samples is larger than $\frac{2 \log(\frac{2|F|}{\delta})}{\epsilon^2(1-2\zeta)^2}$, then the estimation error of f in terms of ranking prediction is smaller than ϵ with the probability larger than $1 - \delta$.*

The proof of this theory is similar to that of theory 1 in [34]. From this theory, we can see, as the noisy parameter ζ becomes larger, more samples (i.e., $\frac{2 \log(\frac{2|F|}{\delta})}{\epsilon^2(1-2\zeta)^2}$) are needed to achieve sufficiently well performance (the prediction error is smaller than ϵ). Extremely, when $\zeta \rightarrow 0.5$, we need to produce infinite training samples. This implies that if the recommender simulator is completely random, then we can not expect well performance by training on the generated data, which is aligned with the intuition.

3.3.3 Controlling the noisy information. Inspired by this theory, we propose a simple but effective method to control the noisy information. The general idea is to remain the samples with higher confidence. In specific, for a given recommendation list \hat{r} , we use the selection probability $p_S(\hat{r}_k|U = \hat{u}, R = \hat{r}, \hat{\beta})$ to measure the confidence of the recommender simulator. We denote by \hat{s}_+^k and \hat{s}_-^k the sets of k items with the largest and smallest selection probabilities. Then the training set for the pair-wise objective is built as $C(\hat{\tau}_t) = \{(\hat{u}, i, j) | i \in \hat{s}_+^k, j \in \hat{s}_-^k\}$. For the point-wise objective, $C(\hat{\tau}_t)$ is set as $\{(\hat{u}, i) | i \in \hat{s}_+^k\} \cup \{(\hat{u}, i) | i \in \hat{s}_-^k\}$. In this method, if k is small, then the model has more confidence on the generated samples, which may reduce the noisy information, but at the same time, less samples can be generated, which may lead to insufficient model optimization. If k is large, more samples will be generated, but the noisy level can also be high. In this sense, k is actually a parameter to trade-off the noisy information and the number of generated samples.

4 EXPERIMENTS

In this section, we conduct extensive experiments to verify the effectiveness of our framework. In the following, we begin with the experiment setup, and then report and analyze the results.

4.1 Experiment Setup

4.1.1 Datasets. Our experiments are based on both synthetic and real-world datasets. With the synthetic dataset, we aim to evaluate our framework in a controlled manner under clean environment. By the real-world dataset, we try to verify our framework’s effectiveness in real-world settings.

We follow the previous work [42] to build the synthetic dataset, where we simulate $N_U (= 600)$ users and $N_I (= 300)$ items. For each user i (or item j), the preferences $\mathbf{p}_i \in \mathbb{R}^d$ (or properties $\mathbf{q}_j \in \mathbb{R}^d$) are generated from a multi-variable Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, where d and \mathbf{I} represent the vector size and unit matrix, respectively. In order to generate the recommendation list for user

i , we firstly compute the score of recommending item j based on a neural network: $r_{ij} = 1 - \sigma(\mathbf{a}^T \kappa_1(\kappa_2([\mathbf{p}_i, \mathbf{q}_j])) + b)$, where $\mathbf{a} \in \mathbb{R}^{2d}$ is specified as an all-one vector, and b is set as zero. We follow [42] to set $\kappa_1(\cdot)$ and $\kappa_2(\cdot)$ as piecewise functions, and $\kappa_1(x) = x - 0.5$ if $x > 0$, otherwise $\kappa_1(x) = 0$, $\kappa_2(x) = x$ if $x > 0$, otherwise $\kappa_2(x) = 0$. Given r_{ij} , the probability of recommending item j is $\frac{e^{r_{ij}}}{\sum_{k=1}^{N_I} e^{r_{ik}}}$. For each user, we generate 25 impression lists, each of which is composed of 5 items.

When generating the feedback of a user i on an item j , we follow the previous work [42] to explore both linear and non-linear user response models, that is:

$$\begin{aligned} x_1 &= \mathbf{a}^T \kappa_1(\kappa_2([\mathbf{p}_i, \mathbf{q}_j])) + b \\ x_2 &= \mathbf{a}^T \kappa_3(\kappa_2([- \mathbf{p}_i, - \mathbf{q}_j])) + b \\ s_{ij}^{\text{lin}} &= \mathbb{I}(\sigma(x_1) + N_y - 0.5) \\ s_{ij}^{\text{nonlin}} &= \mathbb{I}(\sigma(x_1 + x_1 \cdot x_2) + N_y - 0.5) \end{aligned} \quad (12)$$

where $\kappa_3(x) = x + 0.5$ if $x < 0$, otherwise $\kappa_3(x) = 0$. N_y is used to set the noisy level of the datasets, and its default value is 0. $\mathbb{I}(x)$ is an indicator function, which is 1 if $x > 0$, and 0 otherwise.

The real-world experiments are based on the recently released MIND⁴ dataset [35]. This data is collected from the user behavior log of Microsoft News, and we use the MIND-small dataset for experiments, where we are provided with 156,965 recommendation lists and 234,468 interactions between 50,000 users and 20,288 items.

4.1.2 Baselines. In order to evaluate the effectiveness of our proposed framework, we compare our model with the following representative methods: **ItemPop** is a non-personalized model, and the ranking score of an item is based on its popularity in the training set. **ItemKNN** [30] is an item-based k-nearest-neighborhood (KNN) model. **BPR** [26] is a well-known recommender model based user implicit feedback. **GMF**, **MLP** and **NeuMF** [18] are neural recommender models, among which NeuMF is a combination between GMF and MLP. **CDAE** [36] is a method based on denoising auto-encoder, which is proved to be the generalization of several well-known collaborative filtering models. **LightGCN** [16] is a graph-based recommender model, where the user-item structure information is considered in the modeling process.

In order to verify the generality of our framework, we apply it to MF, GMF, MLP, NeuMF and LightGCN, respectively, which leads to the models of **CPR-MF**, **CPR-GMF**, **CPR-MLP**, **CPR-NeuMF** and **CPR-LightGCN**. To demonstrate the necessity of learning-based intervention method, we also evaluate the performance our framework with random intervention, where the recommendation list \mathbf{R} is assigned in a random manner. We denote the corresponding baselines as **CPR-MF-r**, **CPR-GMF-r**, **CPR-MLP-r**, **CPR-NeuMF-r** and **CPR-LightGCN-r**.

4.1.3 Implementation details. We follow the previous work [17, 18, 26] to employ standard leave-one-out protocol for evaluation. 10 items are recommended from each model to compare with the users' actually interacted ones. We leverage the widely used metrics including Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG) to compare different models. Among these metrics,

HR aims to measure the overlapping between the predicted and ground truth items, while NDCG further takes the ranking of the prediction results into consideration, and higher ranked accurate items contribute more to the final score. The implementations of our model and all the baselines are based on PyTorch [22] with mini-batch Adam [19] optimizer. For ItemKNN and CDAE, we follow the settings of the original paper. For the other baselines, we determine their hyper-parameters based on grid search, and the search ranges for the embedding size, batch size, regularization coefficient and learning rate are set as $\{32, 64, 128, 256, 512\}$, $\{256, 512, 1024, 2048\}$, $\{2 \times 10^{-2}, 10^{-2}, 5 \times 10^{-3}, 2 \times 10^{-3}, 10^{-3}, 5 \times 10^{-4}, 10^{-4}, 10^{-5}\}$ and $\{5 \times 10^{-3}, 2 \times 10^{-3}, 10^{-3}\}$, respectively. The negative samples are selected from the whole item set or the impression list. When optimizing p_R and p_S , we empirically set the learning rate as 0.001, and the user/item embedding sizes are both tuned in $\{16, 32, 64, 128, 256, 512\}$. The length of the impression list (i.e., $|\mathbf{R}|$) is set as 5, and the size of S (i.e., k) is determined in $\{1, 2, 3\}$. The Gaussian policy is implemented as a two-layers fully-connected neural network, where the hidden dimension is searched in $\{16, 32, 64\}$.

4.2 Experiments with Synthetic Dataset

In this subsection, we present and analyze the experimental results on the synthetic dataset.

4.2.1 Overall performance. In the experiments, we explore different user/item representation dimensions as well as both linear and non-linear user response models. The overall comparison results are shown in Table 1. It is interesting to see that the relative simple baselines (e.g., ItemPop and ItemKNN) are competitive in many cases. This observation is consistent with the previous work [9], and suggests that simple methods may still be useful in some recommendation scenarios.

It is encouraging to see that our framework can consistently improve the performance of different target models. The improvement is consistent across different settings. In specific, our framework can on average improve the performance of the target model by about 10.08% and 11.17% on HR and NDCG, respectively. This result demonstrates the effectiveness of our framework. The reason can be that the recommender simulator enables us to explore the potential user preferences, which provides useful signals to widen the model visions and improve the performance. The importance of the learning-based intervention method is evidenced in the lowered performance when we use random strategy as a replacement, which verifies our claims in the introduction. It is interesting to see that the performance gains on neural models (e.g. MLP, NeuMF and LightGCN) are usually larger than that of the matrix factorization method. In many cases, the performance is even lowered for GMF and BPR after applying our framework, e.g., when we use non-linear user response model and set $d = 32$. This observation agrees with our expectation, that is, neural models need more training samples to display its strong expressiveness, and our framework should be more effective for them. Since most of the promising recommender models in practice are based on neural architectures, this observation demonstrates the potential of our framework in real-world settings.

⁴<https://msnews.github.io/>

Table 1: Performance comparison based on the synthetic dataset. We present the relative improvements of our framework over the target model in the parentheses.

Feedback Function	Linear				Non-Linear			
	d=16		d=32		d=16		d=32	
User/Item Dimension	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10
Metrics								
ItemPop	0.1450	0.0695	0.1219	0.0597	0.3522	0.1849	0.3289	0.1727
ItemKNN	0.1351	0.0632	0.1188	0.0544	0.3477	0.1663	0.3289	0.1626
CDAE	0.1179	0.0550	0.1071	0.0507	0.1869	0.0917	0.1421	0.0661
BPR	0.1567	0.0852	0.1388	0.0712	0.3799	0.2162	0.3636	0.2052
CPR-BPR-r	0.1632	0.0861	0.1465	0.0788	0.3748	0.2118	0.3615	0.2017
CPR-BPR	0.1652 (+5.4%)	0.0937 (+9.9%)	0.1482 (+6.7%)	0.0760 (+6.7%)	0.4024 (+5.9%)	0.2268 (+4.9%)	0.3705 (+1.8%)	0.2028 (-1.1%)
GMF	0.1498	0.0796	0.1304	0.0646	0.3654	0.1927	0.3487	0.1909
CPR-GMF-r	0.1573	0.0760	0.1373	0.0640	0.3830	0.2014	0.3606	0.1886
CPR-GMF	0.1624 (+8.4%)	0.0744 (-6.5%)	0.1344 (+3.0%)	0.0642 (-0.6%)	0.3808 (+4.2%)	0.1944 (+0.8%)	0.3576 (+2.5%)	0.1891 (-0.9%)
MLP	0.1421	0.0705	0.1174	0.0582	0.3603	0.1886	0.3286	0.1718
CPR-MLP-r	0.1422	0.0641	0.1197	0.0593	0.3842	0.1949	0.3626	0.1861
CPR-MLP	0.1592 (+12.0%)	0.0802 (+13.7%)	0.1331 (+13.3%)	0.0608 (+4.4%)	0.3831 (+6.3%)	0.1940 (+2.8%)	0.3668 (+11.6%)	0.1897 (+10.4%)
NeuMF	0.1524	0.0758	0.1203	0.0598	0.3679	0.1917	0.3350	0.1754
CPR-NeuMF-r	0.1731	0.0937	0.1582	0.0815	0.3915	0.2122	0.3701	0.1971
CPR-NeuMF	0.1851 (+21.4%)	0.0963 (+27.0%)	0.1667 (+38.5%)	0.0835 (+39.6%)	0.4007 (+8.9%)	0.2166 (+12.9%)	0.3724 (+11.1%)	0.2019 (+15.1%)
LightGCN	0.1503	0.0725	0.1262	0.0605	0.3732	0.1915	0.3451	0.1744
CPR-LightGCN-r	0.1600	0.0834	0.1317	0.0643	0.3759	0.2008	0.3482	0.1817
CPR-LightGCN	0.1677 (+11.6%)	0.0879 (+21.2%)	0.1551 (+22.9%)	0.0836 (+38.2%)	0.3829 (+2.5%)	0.2016 (+5.3%)	0.3576 (+3.6%)	0.1862 (+6.7%)

Table 2: Effects of the noisy control parameter k , the best performance for each method and setting are labeled by bold fonts.

	N_y	$\mathcal{N}(0, 0.2)$				0			
		Dimension	p=16		p=32		p=16		p=32
	Metrics	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10
CPR-BPR	$k = 1$	0.3591	0.1876	0.3119	0.1780	0.4024	0.2268	0.3663	0.1971
	$k = 2$	0.3609	0.2011	0.3098	0.1762	0.3936	0.2250	0.3705	0.2028
	$k = 3$	0.3580	0.1891	0.3011	0.1665	0.3877	0.2090	0.3647	0.1914
CPR-GMF	$k = 1$	0.3437	0.1827	0.2994	0.1593	0.3725	0.1983	0.3470	0.1869
	$k = 2$	0.3402	0.1762	0.2941	0.1575	0.3808	0.1944	0.3576	0.1891
	$k = 3$	0.3354	0.1709	0.2948	0.1588	0.3564	0.1841	0.3800	0.1945
CPR-MLP	$k = 1$	0.3580	0.1812	0.3051	0.1560	0.3592	0.1868	0.3625	0.1853
	$k = 2$	0.3537	0.1794	0.2878	0.1465	0.3820	0.1937	0.3611	0.1840
	$k = 3$	0.3574	0.1798	0.2886	0.1469	0.3831	0.1940	0.3668	0.1897
CPR-NeuMF	$k = 1$	0.3559	0.1800	0.3046	0.1607	0.3865	0.2005	0.3631	0.1851
	$k = 2$	0.3511	0.1799	0.2870	0.1412	0.3939	0.2115	0.3724	0.2019
	$k = 3$	0.3552	0.1793	0.2990	0.1644	0.4007	0.2166	0.3643	0.1883
CPR-LightGCN	$k = 1$	0.3512	0.1805	0.2935	0.1556	0.3739	0.1952	0.3550	0.1886
	$k = 2$	0.3365	0.1764	0.2794	0.1499	0.3818	0.2014	0.3528	0.1876
	$k = 3$	0.3489	0.1791	0.2962	0.1569	0.3583	0.1874	0.3462	0.1813

4.2.2 *Effects of the noisy control parameter k .* In this section, we investigate the effectiveness of the noisy control method proposed in section 3.3. By this method, we hope our framework can adaptively tune itself to accommodate different noisy recommendation scenarios. In order to evaluate such capability, we set $N_y = 0$ or sample it from $\mathcal{N}(0, 0.2)$ to simulate the settings where the data is observed with different noisy levels. We base the experiment on the non-linear user response model, and we tune k in the range of $\{1, 2, 3\}$. The results are presented in table 2. It is not surprising to see that the performance of the same model is lowered when the noisy level

is increased from 0 to sampling from $\mathcal{N}(0, 0.2)$. When N_y is sampled from $\mathcal{N}(0, 0.2)$, the best performance is mostly achieved when $k = 1$, while if we set $N_y = 0$, $k = 2$ or $k = 3$ can usually lead to more favored results. We speculate that: when the dataset is noisy, the recommender simulator can be not trained well, thus the quality of the generated samples are not high. In such a scenario, removing the noisy samples seems to be more important. On the contrary, if the dataset is noise-free, then the recommender simulator can be more accurate. In this case, the sample quality is not the main issue, and more samples are favored to achieve sufficient model optimization and better performance. For different noisy-level datasets, our

Table 3: Overall results on the real-world dataset.

Model	H@10	NDCG@10
ItemPop	0.3374	0.1608
ItemKNN	0.2403	0.1299
CDAE	0.3403	0.1637
BPR	0.3440	0.1685
CPR-BPR-r	0.3524	0.1718
CPR-BPR	0.3526 (+2.50%)	0.1810 (+7.41%)
GMF	0.3453	0.1651
CPR-GMF-r	0.3525	0.1686
CPR-GMF	0.3519 (+1.91%)	0.1713 (+3.75%)
MLP	0.3469	0.1662
CPR-MLP-r	0.3560	0.1702
CPR-MLP	0.3526 (+1.64%)	0.1747 (+5.11%)
NeuMF	0.3381	0.1630
CPR-NeuMF-r	0.3466	0.1662
CPR-NeuMF	0.3512 (+3.87%)	0.1755 (+7.66%)
LightGCN	0.3451	0.1658
CPR-LightGCN-r	0.3521	0.1719
CPR-LightGCN	0.3688 (+6.86%)	0.1830 (+10.4%)

framework can always adapt itself to achieve better performance, which demonstrates the effectiveness of the noisy control method.

4.3 Experiments with Real-World Dataset

The above synthetic experiments suggest that our idea is effective under ideal environment. In this section, we experiment with the real-world dataset, which is more challenging but more practical.

4.3.1 Overall comparison. To begin with, we compare our model with the baselines introduced in section 4.1.2. The results are presented in Table 3, from which we can see: similar to the synthetic dataset, our CPR framework can consistently lead to improved performances comparing with the target models, and the learning-based intervention method outperforms the random strategy in most cases. The performance gains on neural models are usually larger than that of the shallow ones, which again demonstrates the potential of our model for deep recommender models. Remarkably, the improvement of our framework is not that large comparing with the results on the synthetic dataset. We speculate that the real-world dataset can be noisy. Training based on it can lead to imperfect recommender simulator, which lowers the quality of the generated samples and limits the final recommendation performance.

4.3.2 Performance on the items with different coldness. Cold start has long been a fundamental problem in the recommendation domain. In this section, we study the effectiveness of our framework on the items with different coldnesses. More specifically, we first run the models based on the original training set, and then evaluate their performances on three different testing sets varying on the item coldness. For the first testing set (denoted as “Low”), each item has less than 5 interactions with the users in the training set, which means the model is trained in a quite insufficient manner on these items. In the second testing set (denoted as “Middle”), each item has 5 to 15 interactions, thus the model is provided with more knowledge during the training process. In the last testing set (denoted

as “High”), each item has more than 15 interactions, which is the “warmest” setting in the experiment. The model parameters are set as their optimal values tuned in the above experiments.

From Table 4 we can see: our framework can enhance the performance of the target models in most cases, which demonstrates its effectiveness under different cold start settings. An interesting observation is that the performance improvement on the “Low” and “Middle” testing sets are much larger than that of the “High” dataset. For example, CPR-LightGCN can produce a remarkable 151.74% and 162.75% improvements over LightGCN on HR@10 and NDCG@10 for the “Low” dataset, and the improvements on the “Middle” dataset is about 40%-50%. However, on the “High” dataset, the performance is even lowered on the metric of HR@10, and the improvement on NDCG@10 is also limited. We speculate that, in the “Low” dataset, the testing items are trained in a quite insufficient manner. A large amount of knowledge has been ignored by the target model, which provides more opportunities for the generated samples to introduce useful signals and improve the performance. In real-world recommendation scenarios, cold-start is a notorious problem obsessing people for a long time, this experiment demonstrates the potential of our framework in alleviating this problem.

5 RELATED WORK

In this section, we compare our model with the previous work to highlight the contributions of this paper.

Relation with causal recommendation. Recent years have witnessed many contributions on incorporating causal inference into the recommendation domain [5, 11, 27, 28, 37]. For example, [28] explains the recommendation problem by a treatment-effect model, and designs a re-weighting method to remove the bias in the observed data. However, this method is based on user explicit feedback, and the loss function is restricted to root mean square error (RMSE). In order to handle user implicit feedback, [27, 37] extend this method by incorporating cross-entropy loss and designing tailored debiasing models. In addition, [20] proposes a general knowledge distillation framework to debias the training data. [5] provides a thorough discussion on the recent progress on debiased recommendation. [4] leverages uniform data to learn causal user/item embeddings for more fair and unbiased recommendation. These methods have achieved many successes in the recommendation domain. However, they mostly leverage causal inference to debias the training data, which is different from our data augmentation purpose. In addition, previous work mostly follow Rubin’s potential outcome framework. However, our idea is inspired from Pearl’s structure causal models, where we can explicitly model the data generated environment to ensure compatibleness between the generated and observed data.

Relation with counterfactual data augmentation. Counterfactual data augmentation stems from the human introspection behavior. It has been recently leveraged to alleviate the training data insufficiency problem in the machine learning community. In the past few years, this idea has achieved great successes in the fields of neural language processing (NLP) [41] and computer vision (CV) [2, 6, 12]. In this paper, we apply it to the top-N recommendation task, which, to the best of our knowledge, is the first time in this field. The major differences between our framework and

Table 4: Performances on the items with different coldnesses.

Popularity	Low		Middle		High	
	HR@10	NDCG@10	HR@10	NDCG@10	HR@10	NDCG@10
BPR	0.0367	0.0113	0.4165	0.1426	0.8966	0.4937
CPR-BPR	0.0380 (3.54%)	0.0121 (7.08%)	0.4494 (7.90%)	0.1602 (12.34%)	0.8930 (-0.40%)	0.5027 (1.82%)
GMF	0.0079	0.0023	0.3715	0.1163	0.9157	0.4793
CPR-GMF	0.0092 (16.46%)	0.0027 (17.39%)	0.4028 (8.43%)	0.1270 (9.20%)	0.9152 (-0.05%)	0.4810 (0.35%)
MLP	0.0154	0.0046	0.3531	0.1130	0.9240	0.4845
CPR-MLP	0.0196 (27.27%)	0.0059 (28.26%)	0.4091 (15.86%)	0.13656 (20.85%)	0.9261 (0.23%)	0.4870 (0.52%)
NeuMF	0.0246	0.0075	0.3010	0.0970	0.9018	0.4782
CPR-NeuMF	0.0277 (12.60%)	0.0085 (13.33%)	0.3940 (30.90%)	0.1271 (31.03%)	0.9182 (1.82%)	0.4872(1.88%)
LightGCN	0.0172	0.0051	0.3385	0.1061	0.9171	0.4781
CPR-LightGCN	0.0433 (151.74%)	0.0134 (162.75%)	0.4769 (40.89%)	0.1601 (50.90%)	0.9115 (-0.61%)	0.4860 (1.65%)

the previous work is: (i) we design a leaning-based intervention method to encourage the informativeness of the generated samples. (ii) We theoretically analyze the noisy information in the generated samples, and design a tailored noise control method.

Relation with ranking based recommendation model. Ranking based recommender models have been widely studied for the Top-N recommendation task. The most famous model in early years is Bayesian personalized ranking (BPR), which is optimized by maximizing the user preference margin between the positive and negative items. This method has inspired many promising models. For example, CKE [38] proposes a hybrid model to integrate collaborative filtering and knowledge base for recommendation. [10] improves BPR with a better negative sampler which leverages additional data in E-commerce. AMF [17] applies adversarial training method to enhance the performance of BPR. With the ever prospering of deep neural network, recent years have witnessed the surge of neural recommender models [7, 8, 15, 18, 32]. For example, NeuMF [18] is designed to capture the non-linear user-item correlations. NGCF [33] and LightGCN [16] regard the user-item interactions as a graph, and explicitly incorporate the structure information to enhance the utilization of the collaborative filtering signals. By looking back the history of ranking based recommendation models, it is evident that the model architectures are becoming deeper and heavier. While the comprehensive parameters can indeed lead to improved recommendation performance, an unprecedented problem is also emerging, that is, more training data is needed to satisfy the heavy architectures, which contradicts with the sparse user behaviors in realities. In this paper, we propose a solution to this problem based on causal inference, which is parallel with the previous model-drive research.

6 CONCLUSION

In this paper, we propose to enhance Top-N recommendation with Pearl’s causal inference framework, where we can simulate user preference and generate counterfactual samples for alleviating the training data insufficiency problem. We design a learning-based intervention method for discovering the informative samples, and conduct theoretical analysis to reveal the relation between the number of generated samples and the potential model prediction error. Extensive experiments based on both synthetic and real-world datasets are conducted to verify our model’s effectiveness.

This paper actually opens the door of incorporating Pearl’s causal inference framework into the recommendation domain. There is still much room for improvement. For example, one can design more advanced exogenous node structures to introduce more reasonable prior knowledge for different recommendation scenarios. In addition, it should be also interesting to incorporate side information into the structure equation models, which can help to capture more comprehensive user preference and obtain more accurate recommender simulators.

7 APPENDIX

7.1 Proof of the Theories in Section 3.3

7.1.1 Proof of theory 1.

PROOF. Without loss of generality, we suppose $\eta > \frac{1}{2}$. We define the random variable X_i as the i th observation of the relation between i and j . We use $X_i = 1$ to represent $i > j$, and $X_i = 0$ to indicate $i < j$. Suppose, we have N observations, and we define $\bar{X} = \frac{1}{N} \sum_{i=1}^N X_i$. Following the voting mechanism, if $\bar{X} < \frac{1}{2}$, then the prediction is wrong, since it contradicts with $\eta > \frac{1}{2}$. Obviously,

$$E[\bar{X}] = \frac{1}{N} \sum_{i=1}^N E[X_i] = \eta. \quad (13)$$

According to the Hoeffding’s inequality [29], we have:

$$p(\bar{X} < \frac{1}{2}) = p(\bar{X} - E[\bar{X}] < \frac{1}{2} - \eta) < \exp(-2N(\frac{1}{2} - \eta)^2). \quad (14)$$

By setting $\delta = \exp(-2N(\frac{1}{2} - \eta)^2)$, we have when N is larger than $\frac{\log \frac{1}{\delta}}{2(1-2\eta)^2}$, $p(\bar{X} < \frac{1}{2}) < \delta$, which means the prediction error is smaller than δ . \square

7.1.2 Proof of theory 2.

PROOF. For a hypothesis $f \in \mathcal{F}$, suppose its prediction error is s (i.e., $\sum \mathbb{I}(g_f \neq g) = s$)⁵, then the mis-matching probability between the observed and predicted results comes from two parts:

- The observed result is true, but the prediction is wrong, that is, $s(1 - \zeta)$.
- The observed result is wrong, but the prediction is right, that is $(1 - s)\zeta$.

⁵ g_f is the prediction from f , and g is the ground truth.

Thus, the total mis-matching probability is $\zeta + s(1 - 2\zeta)$. Suppose the prediction error of f (i.e., s) is larger than ϵ . Then, at least one of the following statements hold:

- The empirical mis-matching rate of f is smaller than $\zeta + \frac{\epsilon(1-2\zeta)}{2}$.
- The empirical mis-matching rate of the optimal f^* is larger than $\zeta + \frac{\epsilon(1-2\zeta)}{2}$.

These statements are easy to understand, since if both of them do not hold, we can conclude that the empirical loss of f is larger than that of f^* , which does not agree with the ERM definition. However, according to the Hoeffding inequality [29], both of the above statements hold with the probability smaller than δ , which implies that the prediction error of f is smaller than ϵ with the probability larger than $1 - \delta$. \square

REFERENCES

- [1] Rupam Acharyya, Shouman Das, Ankani Chattoraj, and Md Iftekhar Tanveer. 2020. FairyTED: A Fair Rating Predictor for TED Talk Data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 338–345.
- [2] Oron Ashual and Lior Wolf. 2019. Specifying object attributes and relations in interactive scene generation. In *Proceedings of the IEEE International Conference on Computer Vision*. 4561–4569.
- [3] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. 2017. Variational inference: A review for statisticians. *Journal of the American statistical Association* 112, 518 (2017), 859–877.
- [4] Stephen Bonner and Flavian Vasile. 2018. Causal embeddings for recommendation. In *Proceedings of the 12th ACM conference on recommender systems*. 104–112.
- [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] Long Chen, Hanwang Zhang, Jun Xiao, Xiangnan He, Shiliang Pu, and Shih-Fu Chang. 2019. Counterfactual critic multi-agent training for scene graph generation. In *Proceedings of the IEEE International Conference on Computer Vision*. 4613–4623.
- [7] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ipsir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 7–10.
- [8] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.
- [9] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are we really making much progress? A worrying analysis of recent neural recommendation approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys 2019, Copenhagen, Denmark, September 16-20, 2019*. 101–109.
- [10] Jingtao Ding, Fuli Feng, Xiangnan He, Guanghui Yu, Yong Li, and Depeng Jin. 2018. An improved sampler for bayesian personalized ranking by leveraging view data. In *Companion Proceedings of the The Web Conference 2018*. 13–14.
- [11] Zhenhua Dong, Hong Zhu, Pengxiang Cheng, Xinhua Feng, Guohao Cai, Xiquang He, Jun Xu, and Jirong Wen. 2020. Counterfactual learning for recommender system. In *Fourteenth ACM Conference on Recommender Systems*. 568–569.
- [12] Tsu-Jui Fu, Xin Eric Wang, Matthew F Peterson, Scott T Grafton, Miguel P Eckstein, and William Yang Wang. 2020. Counterfactual Vision-and-Language Navigation via Adversarial Path Sampler. In *European Conference on Computer Vision*. Springer, 71–86.
- [13] Hongchang Gao and Heng Huang. 2018. Self-paced network embedding. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1406–1415.
- [14] Yash Goyal, Ziyang Wu, Jan Ernst, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Counterfactual visual explanations. *arXiv preprint arXiv:1904.07451* (2019).
- [15] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguang Li, Xiuqiang He, and Zhenhua Dong. 2018. Deepfm: An end-to-end wide & deep learning framework for CTR prediction. *arXiv preprint arXiv:1804.04950* (2018).
- [16] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yong-Dong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*. ACM, 639–648.
- [17] Xiangnan He, Zhankui He, Xiaoyu Du, and Tat-Seng Chua. 2018. Adversarial personalized ranking for recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 355–364.
- [18] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. International World Wide Web Conferences Steering Committee, 173–182.
- [19] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.).
- [20] Dugang Liu, Pengxiang Cheng, Zhenhua Dong, Xiuqiang He, Weiwei Pan, and Zhong Ming. 2020. A general knowledge distillation framework for counterfactual recommendation via uniform data. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 831–840.
- [21] Wei Niu, James Caverlee, and Haokai Lu. 2018. Neural personalized ranking for image recommendation. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 423–431.
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. 8024–8035.
- [23] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [24] Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. 2016. *Causal inference in statistics: A primer*. John Wiley & Sons.
- [25] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. 2017. *Elements of causal inference: foundations and learning algorithms*. The MIT Press.
- [26] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*. AUAI Press, 452–461.
- [27] Yuta Saito. 2019. Unbiased Pairwise Learning from Implicit Feedback. In *NeurIPS 2019 Workshop on Causal Machine Learning*.
- [28] Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. 2016. Recommendations as treatments: Debiasing learning and evaluation. *arXiv preprint arXiv:1602.05352* (2016).
- [29] Shai Shalev-Shwartz and Shai Ben-David. 2014. *Understanding machine learning: From theory to algorithms*. Cambridge university press.
- [30] Jun Wang, Arjen P. de Vries, and Marcel J. T. Reinders. 2006. Unifying user-based and item-based collaborative filtering approaches by similarity fusion. In *SIGIR 2006: Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Seattle, Washington, USA, August 6-11, 2006*. ACM, 501–508.
- [31] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. 2017. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*. 515–524.
- [32] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. 2017. Deep & cross network for ad click predictions. In *Proceedings of the ADKDD'17*. 1–7.
- [33] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *SIGIR*.
- [34] Zhenlei Wang, Jingsen Zhang, Hongteng Xu, Xu Chen, Yongfeng Zhang, Wayne Xin Zhao, and Ji-Rong Wen. 2021. Counterfactual Data-Augmented Sequential Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 347–356.
- [35] Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, and Ming Zhou. [n.d.]. MIND: A Large-scale Dataset for News Recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*. 3597–3606.
- [36] Yao Wu, Christopher DuBois, Alice X. Zheng, and Martin Ester. 2016. Collaborative Denoising Auto-Encoders for Top-N Recommender Systems. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, San Francisco, CA, USA, February 22-25, 2016*. ACM, 153–162.
- [37] Longqi Yang, Yin Cui, Yuan Xuan, Chenyang Wang, Serge Belongie, and Deborah Estrin. 2018. Unbiased offline recommender evaluation for missing-not-at-random implicit feedback. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 279–287.
- [38] Fuzheng Zhang, Nicholas Jing Yuan, Defu Lian, Xing Xie, and Wei-Ying Ma. 2016. Collaborative Knowledge Base Embedding for Recommender Systems. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, Balaji Krishnapuram, Mohak Shah, Alexander J. Smola, Charu C. Aggarwal, Dou Shen, and Rajeev Rastogi (Eds.). ACM, 353–362.
- [39] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Kaiyuan Li, Yushuo Chen, Yujie Lu, Hui Wang, Changxin Tian, Xingyu Pan, Yingqian Min, Zhichao Feng, Xinyan Fan, Xu Chen, Pengfei Wang, Wendi Ji, Yaliang Li, Xiaoling Wang,

- and Ji-Rong Wen. 2020. RecBole: Towards a Unified, Comprehensive and Efficient Framework for Recommendation Algorithms. *CoRR* abs/2011.01731 (2020). arXiv:2011.01731 <https://arxiv.org/abs/2011.01731>
- [40] Xiangyu Zhao, Liang Zhang, Long Xia, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2017. Deep reinforcement learning for list-wise recommendations. *arXiv preprint arXiv:1801.00209* (2017).
- [41] Ran Zmigrod, Sabrina J Mielke, Hanna Wallach, and Ryan Cotterell. 2019. Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology. *arXiv preprint arXiv:1906.04571* (2019).
- [42] Hao Zou, Kun Kuang, Boqi Chen, Peixuan Chen, and Peng Cui. 2019. Focused Context Balancing for Robust Offline Policy Evaluation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, Ankur Teredesai, Vipin Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis (Eds.). ACM, 696–704. <https://doi.org/10.1145/3292500.3330852>