Causality-driven User Modeling for Sequential Recommendations over Time

Xingming Chen
The Hong Kong Polytechnic University
Hong Kong, China
xingming.chen@connect.polyu.hk

Qing Li
The Hong Kong Polytechnic University
Hong Kong, China
qing-prof.li@polyu.edu.hk

ABSTRACT

Contemporary sequential recommendation systems predominantly leverage statistical correlations derived from user interaction histories to predict future preferences. However, these correlations often mask implicit challenges. On the one hand, user data is frequently plagued by implicit, noisy feedback, misdirecting users towards items that fail to align with their actual interests, which is magnified in sequential recommendation contexts. On the other hand, prevalent methods tend to over-rely on similarity-based attention mechanisms across item pairs, which are prone to utilizing heuristic shortcuts, thereby leading to suboptimal recommendation.

To tackle these issues, we put forward a causality-driven user modeling approach for sequential recommendation, which pivots towards a causal perspective. Specifically, we involves the application of a causal graph to identify confounding factors that give rise to spurious correlations and to isolate conceptual variables that causally encapsulate user preferences. By learning the representation of these disentangled causal variables at the conceptual level, we can distinguish between causal and non-causal associations while preserving the inherent sequential nature of user behaviors. This enables us to ascertain which elements are critical and which may induce unintended biases. The framework of our method can be compatible with various mainstream sequential models, which offers a robust foundation for reconstructing more accurate and meaningful user and item representations driven by causality.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; • **Information systems** → **Users and interactive retrieval**.

KEYWORDS

Sequential recommendation, Causality learning, Bias alleviation

ACM Reference Format:
Xingming Chen and Qing Li. 2024. Causality-driven User Modeling for Sequential Recommendations over Time. In *Companion Proceedings of the ACM Web Conference 2024 (WWW '24 Companion), May 13–17, 2024, Singapore, Singapore.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3589335.3651896

WWW '24 Companion, May 13–17, 2024, Singapore, Singapore

ACM ISBN 979-8-4007-0172-6/24/05.
https://doi.org/10.1145/3589335.3651896

1 INTRODUCTION

Sequential recommendation systems have become indispensable in filtering and personalizing the content for users across various digital platforms. These methods are traditionally designed to predict future interactions by learning from the sequence of historical data that encapsulates user actions, item properties, and other contextual features, such as user profiles, item attributes, and social relationships [3, 13]. The effectiveness of these systems has been hinged on the premise of accurately representing this rich information to foresee user preferences.

Recent advancements in sequential recommendation have yielded various methods that exploit the rich statistical associations present in users' historical interactions. These techniques aim to effectively model user preferences, thereby improving the accuracy of recommendations. Attention mechanisms [21] and Recurrent Neural Networks (RNNs) [5] are frequently used to capture various aspects of user engagement and sequential dependencies. Meanwhile, Graph Neural Networks (GNNs) [16] and hybrid models [17] are employed to discern complex co-occurrence relations and high-order structural dependencies inherent in sequential data. Nevertheless, these sophisticated approaches often share a common shortfall: they presume that the occurrence of user profiles and item exposures are independent within the observational data. This assumption, while convenient for certain analyses, does not adequately reflect the complex dynamics of user behavior across time. Interactions within a user's history are interconnected, influenced by a variety of factors that can introduce confounding biases or spurious correlations. Sequential data, especially when spanning extensive periods, may be riddled with implicit, noisy feedback, leading users to interact with items that do not genuinely pique their interests, such as due to the influence of item popularity or caption biases [4], thereby potentially distorting the recommendation outcomes. In addition, many state-of-the-art models utilize similarity-based attention mechanisms [1, 14] that concentrate on the correlation between pairs of items. This focus can inadvertently cause an over-reliance on what might be termed 'shortcut paths', neglecting the intricate sequence of interactions and the consequential dependencies. Consequently, these models might fail to discern the actual reasons behind a user's interaction with an item, ultimately leading to suboptimal recommendations.

To tackle these challenges, our innovative approach utilizes a causal graph to model sequential recommendations from a causal perspective. The causal graph is instrumental in identifying and separating the exogenous variables that can lead to undesired correlations and biases. By discerning the true causal factors that influence user interests and item exposure over time, we can better understand and predict user behaviors. Specifically, a causality-driven user modeling technique is built upon the causal graph. By treating user interaction histories as observational data, we autonomously identify valid causal factors at the conceptual level and mitigates the influence of potential biases without depending on explicitly predefined features. Furthermore, our method leverages the latent causal factors that underlie user preferences throughout the interaction sequence, as well as the immediate item interactions. The former aids in mitigating bias from short-cut features and capturing pertinent information across the sequence, while the latter maintains essential context that causally influences item interactions at each timestep, without exacerbating potential real-time confounding biases. Last but not least, with these identified causal concepts, we can generate causality-driven representation for user preference modeling, which can be integrate into the base matching model to provide refined predictions for user interactions. The principal contributions of this paper are as follows:

- We implement a causal graph to model sequential recommendations, enhancing the understanding of user behavior and improving prediction accuracy.
- We develop a causality-driven technique for user modeling over time, which can autonomously identify causal factors from user interaction data, reducing reliance on predefined features and mitigating biases.
- Our learned latent causal concepts can be utilized to mitigate confounding bias and preserve critical context, leading to more refined predictions of user interactions.

2 RELATIVE WORK

2.1 Causal Recommendation

Recent research has begun leveraging causal analysis in recommendation systems to address confounding bias, with popular efforts aimed at disentangling specific biases in practical applications [8, 9, 19]. For example, some works [2] treat exposure rates as confounding factors that influence both the propensity of user engagement and the assessment of user satisfaction to mitigate exposure bias. Causal interventions have been used to decouple item popularity from user representations [20], addressing popularity-induced confounding bias. Other methods employ instrumental variables from external querying contexts [10], regressing user representations on these contexts to sidestep confounding effects. However, these strategies often rely on predefined objectives and assumed causal variables within standard recommendation frameworks, highlighting the need for more adaptive techniques to uncover implicit causal variables in sequential recommendation.

3 METHODOLOGY

3.1 Problem Formulation

In sequential recommendation systems, we define a set of users $U = \{u_1, u_2, \ldots, u_{|U|}\}$ and a set of items $I = \{i_1, i_2, \ldots, i_{|I|}\}$, where $u_j \in U$ represents an individual user and $i_k \in I$ represents an individual item. For any given user u, we represent their interaction sequence as $S_u = (i_1^u, i_2^u, \ldots, i_n^u)$, where the sequence of n interacted items is ordered chronologically and $i_t^u \in I$ signifies the item

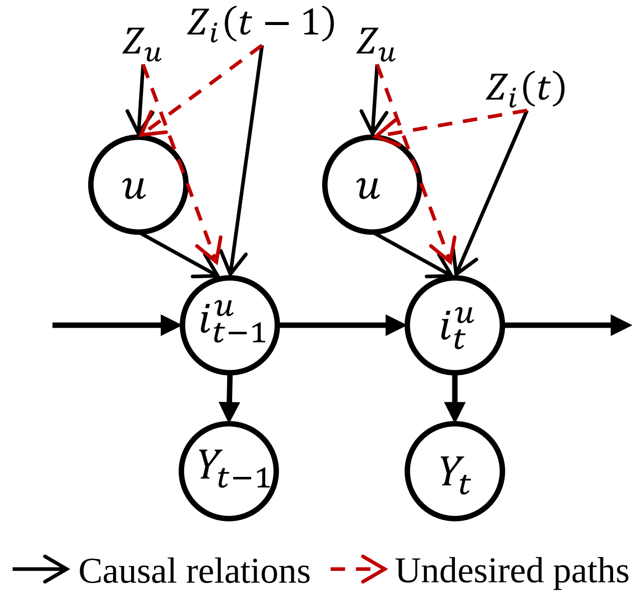

Figure 1: Causal graph of sequential recommendation.

interacted with by user u at time t. Both users and items are typically initialized as id embeddings in a d-dimensional space. The primary objective of sequential recommendation is to predict the next item, i_{t+1} that a user u is likely to engage with. More precisely, we aim to compute the conditional probability $Y_{t+1} = P(i_{t+1}^u|S_u)$ for the subsequent item. Items are then recommended to the user in descending order of this predicted probability.

3.2 Viewing from Causal Graph

To predict the next item a user will interact with, it is crucial to understand the reasons behind a user's interaction at each timestep, thereby learning accurate and robust representations of user behavior. Existing approaches often rely on user-item correlation matching, operating under the assumption that the co-occurrence of users and exposed items are independent within the observational data. For instance, at a given timestep t, conventional thinking suggests that the user's behavior of interacting with item i_t^u is exclusively driven by their preferences, implying the absence of a confounding back-door path between user u and item i_t. However, this assumption often fails to hold in complex, real-world scenarios.

By framing the sequential recommendation task from a causal perspective, we can address these discrepancies. We propose the use of a causal graph, as depicted in Figure 1, for modeling user preferences. This graph helps to identify exogenous variables that create undesired correlations between user preferences and the items they interact with instantaneously. Specifically, the path $i_t^u \leftarrow Z_i(t) \rightarrow u$ can be seen as conveying bias information that misguides the user to interact with item i_t^u at timestamp t, such as popularity bias, while $i_t^u \leftarrow Z_u \rightarrow u$ may represent a shortcut bias that overlooks sequential dependencies and intermediate interactions. Our aim is to discern the true causal factors, Z_u and Z_i, which directly influence user interests and the real-time exposure of items.

3.3 The Base Matching Model

Our proposed method is model-agnostic in a way, and it can be integrated into multiple existing matching models for sequential recommendation. Since many models for sequential recommendation share similar architecture, we regard it as the base model and implement our method over it. Generally speaking, a matching model includes a user encoder $f_S(u, S_u) \in \mathbb{R}^d$, which takes the user profile u and the user interaction sequence $S_u = (i_1^u, i_2^u, \ldots, i_n^u)$ as input and outputs one d-dimensional vector to represent the user interaction sequence, and an item encoder $f_Y(i) \in \mathbb{R}^d$ representing the item in the same vector space as $f_S(u, S_u)$. The matching score is generally calculated by maximizing the likelihood of the next interacted item given the interaction sequence:

$$L_{base} = -\frac{1}{|U|} \sum_{u \in U} \log P_\theta(i_{n+1}^u | f_s(u, S_u)), \tag{1}$$

3.4 Causality-Driven User Modeling

Existing de-biasing techniques in the realm of sequential recommendation commonly rely on precisely identified confounding factors or supplementary contexts [9, 10, 20], such as user search queries, to serve as proxy variables. These proxies are instrumental in uncovering the genuine causal factors that reflect user interests. In this section, we treat users' sequential interaction histories as observational data. Our objective is to autonomously discern valid causal factor across both the interaction sequence and the immediate interaction, at the conceptual level. Additionally, we aim to mitigate the influence of potential back-door paths present at each timestep, and we strive to achieve this without depending on explicitly predefined biases. Our causality-driven user modeling method can be depicted as Figure 2.

Initially, we leverage Graph Attention Networks (GATs) [11] to encode user and item interaction representations, aiming to extract sophisticated features from all observed interaction sequences:

$$\mathrm{E}_{\mathcal{G}} = \mathrm{f}_{\mathrm{readout}}(\mathrm{Attn}(\mathrm{WI}, \mathrm{A})), \tag{2}$$

where W is a trainable weight matrix and A denotes the adjacency matrix that encapsulates the structure of interaction sequences. Attn(.) signifies a multi-head attention mechanism [12] which serves to assess the significance of neighboring behaviors and enhance the representation of item interactions at each timestep. And $f_{readout}$ refers to a readout function that propagates item-level features $\boldsymbol{i}_t^u$ to generate a graph-level representation $\mathrm{E}_{\mathcal{G}}$, and the representation $\boldsymbol{u}$ for encapsulating user preference can be obtained by inputting the corresponding interaction sequence from the graph.

3.4.1 Causal Concept Representation over the Interaction Sequence. To effectively represent the latent variables Z_u that causally impact user interests at a conceptual level, our primary approach is to identify aspects of the historical interaction sequence that meet two criteria: first, they must be relevant to the user representation $\boldsymbol{u}$; second, they should be exclusive to the real-time item engagement $\boldsymbol{i}_t^u$ except through $\boldsymbol{u}$.

Assuming there are K latent conceptual prototypes $Z_u \in \mathbb{R}^{K \times d}$, We first address the relevance of these prototypes to user interests. Inspired by the work of causal inference [18], to isolate potential aspects of the sequence S_u correlating with the user representation $\boldsymbol{u}$ for inclusion in the prototype matrix Z_u, we employ a variational distribution $q_{\theta_{Z_u}}(\boldsymbol{u}|Z_u)$ to approximate the conditional distribution of $\boldsymbol{u}$ given Z_u. This approximation is facilitated by a two-layer multi-layer perceptron (MLP). The variational distribution $q_{\theta_{Z_u}}(\boldsymbol{u}|Z_u)$ is derived by maximizing the likelihood estimation:

$$\mathcal{L}_{Z_u}^r = -\frac{1}{|U|} \sum_{u \in U} \log q_{\theta_{Z_u}}(\boldsymbol{u}|W_{ru} Z_u), \tag{3}$$

where W_{ru} denotes a linear transformation of the causal concepts Z_u. Subsequently, by means of maximizing the mutual information between $\boldsymbol{u}$ and Z_u, the causal connection of the conceptual prototypes on user interests can be strengthened:

$$\mathcal{L}_{Z_u}^{r'} = -\frac{1}{|U|} \sum_{u \in U} \frac{1}{L} \sum_{j=1}^{L} (\log q_{\theta_{Z_u}}(\boldsymbol{u}|W_{ru} Z_u) - \log q_{\theta_{Z_u}}(\boldsymbol{u'}|W_{ru} Z_u)), \tag{4}$$

where u denotes the positive sample, whereas u'_j signifies a negative sample that has been randomly selected, and L is the number of negative samples selected. By maximizing the discrepancy between positive and negative samples, we can meet the relevance condition for causal concepts in relation to user interests, thereby enabling a more precise depiction of user preferences for items.

As previously noted, a user's historical sequence of interactions over time can create shortcut paths. For example, a user might directly jump from an initial item i_1^u to a distant item i_t^u, bypassing intermediate items. If the model relies solely on similarity-based statistical correlations, it may overstate the direct connection between two items and discard sequential dependencies. This shortcut perspective could prevent the model from capturing a true representation of user preference through the paths $u \leftarrow Z_u \rightarrow i_1^u$ and $u \leftarrow Z_u \rightarrow i_t^u$, making the reason behind a user's item choice more obscure. Therefore, it is crucial to block the back-door path $u \leftarrow Z_u \rightarrow i_t^u$ over time to mitigate the influence of these shortcut views. This can be accomplished by excluding the concepts Z_u associated with the interaction at a specific time, thereby extracting valid insights that accurately reflect real user preferences based on sequential dependencies, rather than an immediate interaction's shortcut view. To satisfy the exclusion condition, we can apply methods similar to those used for maintaining the relevance condition. Initially, the variational distribution $q_{\phi_{Z_u}}(\boldsymbol{i}_t^u|Z_u)$ is used to approximated the conditional distribution of each i_t^u over time given Z_u. The corresponding loss function for negative maximum log-likelihood estimation is:

$$\mathcal{L}_{Z_u}^e = -\frac{1}{|U|} \sum_{u \in U} \sum_{t=1}^{n} \log q_{\phi_{Z_u}}(\boldsymbol{i}_t^u|W_{eu} Z_u), \tag{5}$$

where W_{eu} denotes the trainable weights. Then, to ensure that the causal concept representations Z_u from the user's perspective are not influenced by the shortcut views of items at one or two time steps, it is imperative to minimize the mutual information between i_t^u and Z_u. This can be optimized by minimizing the difference in mutual information between the positive sample (u, i_t^u) and the

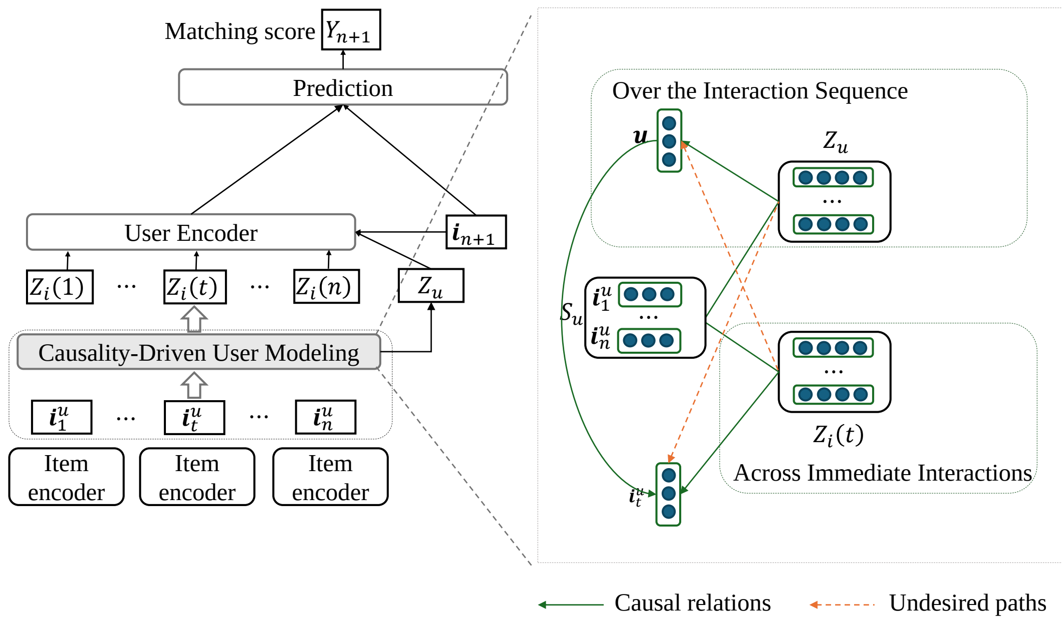

Figure 2: The architecture of our framework. Left: integration of Causality-Driven User Modeling within the standard sequential recommendation matching Scheme. Right: detailed implementation of Causality-Driven User Modeling.

negative sample $(u', i_t'^u)$:

$$\mathcal{L}_{Z_u}^{e'} = \frac{1}{|U|}\sum_{u\in U}\sum_{j=1}^{L}\sum_{t=1}^{n}(sim(\boldsymbol{u},\boldsymbol{u}')\times (\log q_{\phi_{Z_u}}(\boldsymbol{i}_t^u|W_{eu}Z_u) - \log q_{\phi_{Z_u}}(\boldsymbol{i}_t'^u|W_{eu}Z_u))). \tag{6}$$

Unlike that in keeping relevance condition, the causal concept representations can not be absolutely irrelevant with the immediate interaction, for the reason that they are actually indirectly causally connected through the path $Z_u \to u \to i_t^u$. The exclusion requirement should be fulfilled only when conditioning on the user preference. Therefore, $sim(u, u')$ is employed to describe the similarity between the positive and negative ones, i.e., u and u'. When they are similar to each other, the weight of the discrepancy about log-likelihood between the positive and negative samples should be large. Otherwise, the weight should be small. In other words, users with similar interaction sequences should be concentrated more on minimizing the mutual information. And the similarity $sim(u, u')$ is calculated by:

$$sim(\boldsymbol{u},\boldsymbol{u}') = \mathrm{softmax}(e^{-\|\boldsymbol{u}-\boldsymbol{u}'\|}). \tag{7}$$

3.4.2 Causal Concept Representation across Immediate Interactions. While mitigating the impact of shortcut features within the causal concepts throughout the interaction sequence is crucial for reducing bias, it's important to recognize that these features are still valuable. They encompass contextual information that has a causal effect on item interactions at each timestep, which is vital for enhancing the performance of recommendations. Consequently, it's necessary to also discern the causal concepts Z_i from the viewpoint of the immediate interaction. Given that the causal concepts underlying each item interaction tend to evolve over time, we introduce a set of K concept prototype embeddings $Z_i(t) \in \mathbb{R}^{K\times d}$, designated to capture the causal concepts relevant to each timestep.

In a manner analogous to the method employed for learning Z_u, we utilize the variational distribution $q_{\phi_{Z_i}}(\boldsymbol{i}_t^u|Z_i(t))$ to uncover the latent information that causally influences the immediate item interactions, which can be achieved by minimizing the negative log-likelihood function:

$$\mathcal{L}_{Z_i}^{r} = -\frac{1}{|U|}\sum_{u\in U}\frac{1}{n}\sum_{t=1}^{n}\log q_{\theta_{Z_i}}(\boldsymbol{i}_t^u|W_{ri}Z_i(t)), \tag{8}$$

where W_{ri} is a trainable weight matrix. Subsequently, the distance between the log-likelihood expectations of positive sample i_t^u and a negative sample i_t' are required to be far away as possible. By minimizing E.q. (9), the relevance between the real-time item choice i_t^u and its causal concepts $Z_i(t)$ can be kept to a further extent:

$$\mathcal{L}_{Z_i}^{r'} = -\frac{1}{|U|}\sum_{u\in U}\frac{1}{n}\sum_{t=1}^{n}(\log q_{\theta_{Z_i}}(\boldsymbol{i}_t^u|W_{ri}Z_i(t)) - \log q_{\theta_{Z_i}}(\boldsymbol{i}_t'|W_{ri}Z_i(t))), \tag{9}$$

through the optimization of which, we preserve the critical shortcut features within the immediate causal concept representation $Z_i(t)$.

However, we must also be vigilant of the potential back-door path $i_t^u \leftarrow Z_i(t) \to u$, which could introduce confounding bias. Such bias might mislead users into clicking on items that do not align with their genuine interests. For instance, the popularity of an item might increase its likelihood of being exposed to users, prompting clicks on the item due to its popularity rather than its relevance to the users' preferences. Therefore, we need remove the spurious correlation with the user representation $\boldsymbol{u}$ to learn instrumental Z_i. To be specific, the relevance between $\boldsymbol{u}$ and $Z_i(t)$ can also be captured with the variational approximation via estimating the maximum likelihood of the distribution:

$$\mathcal{L}^{e}_{Z_i} = -\frac{1}{|U|}\sum_{u\in U}\frac{1}{n}\sum_{t=1}^{n}\log q_{\theta_{Z_i}}(\boldsymbol{u}|W_{ei}Z_i(t)), \tag{10}$$

where W_{ei} represents a linear transformation.Then, similar to the operation in E.q. (6), the immediate causal concepts $Z_i(t)$ are only allowed to be correlative with the user representation $\boldsymbol{u}$ when regressing on $\boldsymbol{i}^u_t$:

$$\begin{aligned}\mathcal{L}^{e'}_{Z_i} = \frac{1}{|U|}\sum_{u\in U}\sum_{t=1}^{n}\big(sim(\boldsymbol{i}^u_t, \boldsymbol{i}'^u_t)\times \\ \big(\log q_{\theta_{Z_i}}(\boldsymbol{u}|W_{ei}Z_i(t)) - \\ \log q_{\theta_{Z_i}}(\boldsymbol{u}'|W_{ei}Z_i(t))\big)\big),\end{aligned} \tag{11}$$

where $\boldsymbol{i}'^u_t$ and $\boldsymbol{u}'$ denotes the representations of randomly selected negative samples, and $sim(\boldsymbol{i}^u_t, \boldsymbol{i}'^u_t)$ is implemented with the softmax function as E.q. (7).

To date, we have pinpointed the latent causal concepts Z_u and Z_i that pertain to the interaction sequence and real-time item interactions, respectively. The representation driven these causal concepts aids in deciphering the user preferences driven by causality, where biases, including those from shortcut and popularity bias, have been mitigated. This approach is instrumental in achieving superior recommendation outcomes.

3.5 Model Optimization

The latent causal concept representations Z_u and Z_i are integrated into the base matching model to forecast the user's subsequent interactions. These predictions take the form $f_S(Z_u, [Z_i(1), \ldots, Z_i(t)])$ for the sequence. The comprehensive optimization objective that we strive to minimize is expressed as follows:

$$L_{overall} = L_{base} + \alpha(\mathcal{L}^{r}_{Z_u} + \mathcal{L}^{r'}_{Z_u} + \mathcal{L}^{e}_{Z_u} + \mathcal{L}^{e'}_{Z_u}) + \beta(\mathcal{L}^{r}_{Z_i} + \mathcal{L}^{r'}_{Z_i} + \mathcal{L}^{e}_{Z_i} + \mathcal{L}^{e'}_{Z_i}), \tag{12}$$

where the hyperparameters α and β balance the significance of the variational approximation loss and the mutual information constraint loss, respectively. These parameters are fine-tuned using the validation set. The loss function of the base model is prioritized as the primary objective, given that our ultimate aim is to accurately predict the user's next item of interaction based on their profile and interaction history. This base model framework is versatile and can be adapted to most current mainstream models with similar structures.

4 EXPERIMENT

4.1 Datasets and Experimental Settings

To evaluate the effectiveness of our proposed method for sequential recommendation, we conducted experiments using three publicly available datasets: *Diginetica*[1], *MovieLens*[2], *Books-Crossing*[3]:

- **Diginetica.** This dataset originates from an e-commerce platform and includes instances of users' purchase records.
- **MvoieLens.** This is a popular dataset for movie recommendations, featuring data on users' movie-watching behavior.
- **Books-Crossing.** A classic dataset containing book ratings from the Book-Crossing community.

In our experiment, we arrange the items within the interaction sequences chronologically. To augment the number of instances for each user, we use each interaction sequence multiple times. For example, from the sequence of interactions $i_1 \rightarrow i_2 \rightarrow i_3 \rightarrow i_4$, we generate three training instances: i_1, $i_1 \rightarrow i_2$, and $i_1 \rightarrow i_2 \rightarrow i_3$. The labels for these instances are the subsequent items with which the user interacted next. In alignment with established protocols [15, 22], we reserve users' last interactions for testing, the second-to-last for validation, and use remaining data for training. We evaluate our model using NDCG and F1, comparing the top 5 recommended items with actual user interactions. For tuning, grid search refines hyperparameters on the validation set: α and β for loss components at 0.1 and 0.15, and we set latent concepts K to 12 and negative samples L to 8 for learning Z_u. Baseline model parameters are set to previously reported optimal values.

4.2 Baselines

Our proposed framework is designed to enhance existing models for sequential recommendation that have a base structure analogous to the one described in Section 3.3. To demonstrate the efficacy of our method, we have implemented and integrated it with four established baseline models:

- **STAMP [7]** A model leverages an attention mechanism to emphasize short-term user preferences while also capturing their long-term interests.
- **SASRec [15].** A model based on self-attention mechanisms that captures the relationships between items in the entire sequence.
- **SRGNN [6].** A GNN-based model that represents each sequence as a graph to effectively model the complex transitions between items.
- **GES-SASRec [22].** A framework that applies graph convolutions in the context of embedding smoothing for sequential recommendation, leveraging semantic item relationships derived from inherent item attributes.

4.3 Main Results

Our experimental results encompass four baseline models across three datasets, with the findings presented in Table 1. We have enhanced the baseline models STAMP, SASRec, SRGNN, and GES-SASRec, by integrating our method, denoting the augmented versions as CD-STAMP, CD-SASRec, CD-SRGNN, and CD-GES-SASRec, respectively. Here, "CD-" stands for "Causality-Driven".

Our examination of the results reveals that, compared to the original baselines, the enhanced versions exhibit improvements across all datasets. This underscores the efficacy of the causal concepts we have incorporated for learning accurate representations. The datasets span a broad spectrum of recommendation domains, and the observed enhancements underscore the generality and robustness of our causality-driven representations.

[1]https://competitions.codalab.org/competitions/11161
[2]https://grouplens.org/datasets/movielens/
[3]http://www2.informatik.uni-freiburg.de/ cziegler/BX/

We note that the performance gains with CD-GES-SASRec are less substantial than those achieved by the other enhanced models. We attribute this to the fact that GES-SASRec necessitates additional item attributes as inputs. Our approach has not yet disentangled the causal concepts related to these item attributes, which might include implicit noisy feedback, leading to potential negative effects. While our method is capable of extracting potential causal concepts from item attributes, doing so could necessitate modifications to the base model's framework, which is beyond our current scope.

Despite this, it is evident that our causality-driven method positively influences the performance, even though the causal concepts within item attributes have not been fully disentangled. This suggests that our approach contributes beneficially, albeit with room for further refinement in terms of causal concept analysis within item attributes.

Table 1: Comparative main results between the baseline models and our causality-driven methods.

Datasets	Diginetica		MvoieLens		Books-Crossing	
metric (@5)	NDCG	F1	NDCG	F1	NDCG	F1
STAMP	14.67	7.01	9.07	4.50	1.65	0.78
CD-STAMP	15.89	7.82	9.33	4.82	2.60	1.06
SASRec	16.29	8.03	8.57	4.50	1.69	0.80
CD-SASRec	16.96	8.62	8.80	4.62	1.92	1.12
SRGNN	16.55	8.23	8.75	4.72	1.74	0.91
CD-SRGNN	16.98	8.71	8.99	4.83	1.96	1.15
GES-SASRec	17.14	8.66	9.02	4.79	1.92	1.02
CD-GES-SASRec	17.35	8.94	9.13	4.86	2.03	1.12

4.4 Confounding Bias Alleviation

To evaluate the effectiveness of our causality-driven representations in reducing confounding bias, we examine the role of item popularity. Popularity bias exists because frequently exposed items tend to receive more user engagement, regardless of the users' true preferences. We quantify item popularity by the item interaction frequency and divide the training and evaluation sets accordingly, labeling the top 20% items as popular and the rest as unpopular.

We present comparative results using the NDCG@5 and F1@5 metric on the Diginetica dataset to underscore the differences in performance between popular and unpopular items, as shown in Figure 3 and Figure 4. A notable disparity in performance favoring popular items would suggest a significant presence of popularity bias. Note that GES-SASRec is excluded from this comparison because it necessitates additional item attributes. The results indicate that incorporating our causality-driven modeling method significantly narrows the performance gap between popular and unpopular items, which shows the effectiveness of our method in alleviating popularity bias. In contrast, the original baseline models all demonstrate a marked decline in performance when moving from the popular to the unpopular group. In some instances, the performance within the popular group even exceeds that prior to data partitioning, suggesting that popularity bias is prevalent in sequential recommendation systems. Our causality-driven user modeling methods can alleviate such implicit bias by uncovering the latent causal concepts within user interaction sequences.

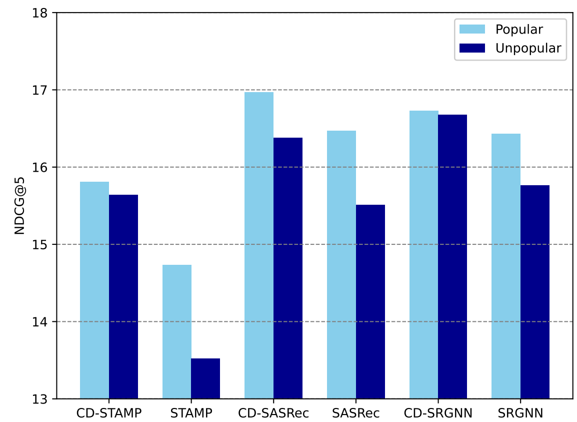

Figure 3: Results highlighting popularity bias (NDCG@5).

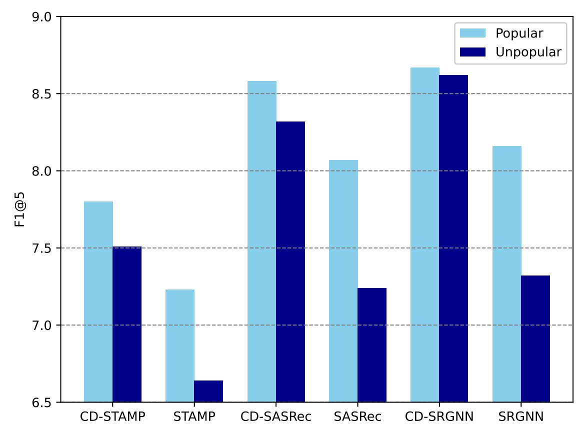

Figure 4: Results highlighting popularity bias (F1@5).

5 CONCLUSION

In conclusion, our causality-driven user modeling approach for sequential recommendation notably improves prediction accuracy and reduces bias. Unlike conventional methods based on statistical correlations, our approach employs a causal graph to redefine the recommendation process. By analyzing user interaction histories as observational data, our method independently identifies causal factors that reflect user preferences at the conceptual level, bypassing the need for predefined features. Additionally, we leverage latent causal factors within user interactions to eliminate back-door paths that could lead to spurious correlations. This ensures the preservation of pertinent information across user sequences. Similarly, we also captures the causal concepts across immediate item interactions, thus mitigating real-time confounding biases. By incorporating these causal concepts into user preference modeling, we enhance various prevalent matching model, leading to more precise predictions of user interactions.

REFERENCES

[1] Jianxin Chang, Chen Gao, Yu Zheng, Yiqun Hui, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. 2021. Sequential recommendation with graph neural networks. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 378–387.

[2] Konstantina Christakopoulou, Madeleine Traverse, Trevor Potter, Emma Marriott, Daniel Li, Chris Haulk, Ed H Chi, and Minmin Chen. 2020. Deconfounding User Satisfaction Estimation from Response Rate Bias. In *Fourteenth ACM Conference on Recommender Systems*. 450–455.

[3] Ruining He, Wang-Cheng Kang, Julian J McAuley, et al. 2018. Translation-based Recommendation: A Scalable Method for Modeling Sequential Behavior. In *IJCAI*. 5264–5268.

[4] Yue He, Zimu Wang, Peng Cui, Hao Zou, Yafeng Zhang, Qiang Cui, and Yong Jiang. 2022. Causpref: Causal preference learning for out-of-distribution recommendation. In *Proceedings of the ACM Web Conference 2022*. 410–421.

[5] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent neural networks with top-k gains for session-based recommendations. In *Proceedings of the 27th ACM international conference on information and knowledge management*. 843–852.

[6] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.

[7] Qiao Liu, Yifu Zeng, Refuoe Mokhosi, and Haibin Zhang. 2018. STAMP: short-term attention/memory priority model for session-based recommendation. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1831–1839.

[8] Rishabh Mehrotra, Prasanta Bhattacharya, and Mounia Lalmas. 2020. Inferring the Causal Impact of New Track Releases on Music Recommendation Platforms through Counterfactual Predictions. In *Fourteenth ACM Conference on Recommender Systems*. 687–691.

[9] Guoshun Nan, Rui Qiao, Yao Xiao, Jun Liu, Sicong Leng, Hao Zhang, and Wei Lu. 2021. Interventional video grounding with dual contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2765–2775.

[10] Zihua Si, Xueran Han, Xiao Zhang, Jun Xu, Yue Yin, Yang Song, and Ji-Rong Wen. 2022. A model-agnostic causal learning framework for recommendation using search data. In *Proceedings of the ACM Web Conference 2022*. 224–233.

[11] Weiping Song, Zhiping Xiao, Yifan Wang, Laurent Charlin, Ming Zhang, and Jian Tang. 2019. Session-based social recommendation via dynamic graph attention networks. In *Proceedings of the Twelfth ACM international conference on web search and data mining*. 555–563.

[12] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. 2017. Graph attention networks. *stat* 1050, 20 (2017), 10–48550.

[13] Chenyang Wang, Min Zhang, Weizhi Ma, Yiqun Liu, and Shaoping Ma. 2019. Modeling item-specific temporal dynamics of repeat consumption for recommender systems. In *The world wide web conference*. 1977–1987.

[14] Shoujin Wang, Liang Hu, Yan Wang, Longbing Cao, Quan Z. Sheng, and Mehmet A. Orgun. 2019. Sequential Recommender Systems: Challenges, Progress and Prospects. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*. 6332–6338.

[15] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 346–353.

[16] Hao Xu, Bo Yang, Xiangkun Liu, Wenqi Fan, and Qing Li. 2022. Category-aware Multi-relation Heterogeneous Graph Neural Networks for session-based recommendation. *Knowledge-Based Systems* 251 (2022), 109246.

[17] Zuoxi Yang, Shoubin Dong, and Jinlong Hu. 2021. GFE: General knowledge enhanced framework for explainable sequential recommendation. *Knowledge-Based Systems* 230 (2021), 107375.

[18] Junkun Yuan, Anpeng Wu, Kun Kuang, Bo Li, Runze Wu, Fei Wu, and Lanfen Lin. 2022. Auto iv: Counterfactual prediction via automatic instrumental variable decomposition. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 16, 4 (2022), 1–20.

[19] Shengyu Zhang, Dong Yao, Zhou Zhao, Tat-Seng Chua, and Fei Wu. 2021. Causerec: Counterfactual user sequence synthesis for sequential recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 367–377.

[20] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal intervention for leveraging popularity bias in recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 11–20.

[21] Yihu Zhang, Bo Yang, Haodong Liu, and Dongsheng Li. 2023. A time-aware self-attention based neural network model for sequential recommendation. *Applied Soft Computing* 133 (2023), 109894.

[22] Tianyu Zhu, Leilei Sun, and Guoqing Chen. 2021. Graph-based embedding smoothing for sequential recommendation. *IEEE Transactions on Knowledge and Data Engineering* 35, 1 (2021), 496–508.