

User-Controllable Recommendation via Counterfactual Retrospective and Prospective Explanations

Juntao Tan^{1,*}, Yingqiang Ge¹, Yan Zhu², Yinglong Xia², Jiebo Luo³, Jianchao Ji¹ and Yongfeng Zhang¹

¹Rutgers University, ²Meta Platforms, Inc., ³University of Rochester
{juntao.tan, yingqiang.ge}@rutgers.edu, yzhuuu@gmail.com, yxia@fb.com, jluo@cs.rochester.edu,
{jianchao.ji, yongfeng.zhang}@rutgers.edu,

Abstract. Modern recommender systems utilize users' historical behaviors to generate personalized recommendations. However, these systems often lack user controllability, leading to diminished user satisfaction and trust in the systems. Acknowledging the recent advancements in explainable recommender systems that enhance users' understanding of recommendation mechanisms, we propose leveraging these advancements to improve user controllability. In this paper, we present a user-controllable recommender system that seamlessly integrates explainability and controllability within a unified framework. By providing both retrospective and prospective explanations through counterfactual reasoning, users can customize their control over the system by interacting with these explanations.

Furthermore, we introduce and assess two attributes of controllability in recommendation systems: the complexity of controllability and the accuracy of controllability. Experimental evaluations on MovieLens and Yelp datasets substantiate the effectiveness of our proposed framework. Additionally, our experiments demonstrate that offering users control options can potentially enhance recommendation accuracy in the future. Source code and data are available at <https://github.com/chrisjtan/ucr>.

1 Introduction

Controllability is crucial in developing trustworthy recommender systems. Studies have demonstrated that offering users the ability to control recommender systems can enhance satisfaction and trust in the recommendation results [21, 33, 17, 18, 7]. However, current recommender systems typically offer limited user control, leaving users to passively receive and accept the recommendations provided. If users are dissatisfied with the suggested items, there is often no way for them to intervene, adjust, or modify the recommendation lists in a better direction. Moreover, users are uncertain about what the recommendation models have learned about them and based on what information. Such uncertainties can raise privacy concerns and further impair users' trust in recommender systems.

Therefore, some research has sought to increase controllability in recommendations by enabling users to explicitly express their preferences on pre-defined aspects or features, which is a common approach in the industry [22, 14, 32, 3]. However, users may not always be able to accurately address their preferences [16], and their preferences may evolve over time, rendering past preferences outdated [15, 1, 16, 17]. Besides, due to the opaque nature of the rec-

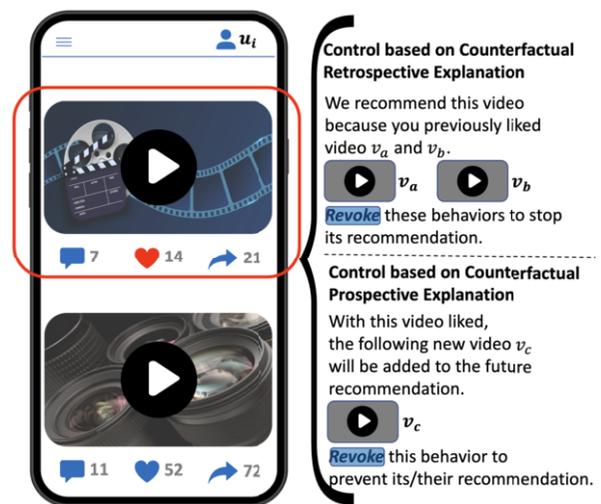


Figure 1. A toy example illustrates the concept of user-controllable recommendations. For the recommended video highlighted by a red bounding box, the retrospective explanation (displayed in the top-right) informs users of the past behaviors that resulted in the video's recommendation, while the prospective explanation (shown in the bottom-right) conveys the impact of liking this video. Users have the option to click the "revoke" button to rescind related behaviors.

ommendation mechanism, users are not directly aware of the effects or consequences of their expressed preferences [35, 29]. Alternative approaches allow users to implicitly update their preferences by interacting with the recommendation results [12, 23]. In these studies, the lack of understanding remains a primary concern, as users are not informed about how their interactions update their preferences or what the implications of their feedback are in the future.

Upon examining the aforementioned research, we notice that the challenge of achieving optimal user controllability comes from the inherent opacity and black-box nature of recommendation systems. Users are often unaware of how their past behaviors influence present recommendations or how their current behavior will affect future recommendations. Allowing users to manipulate AI systems without a comprehensive understanding of the underlying mechanisms may lead to more harm than benefit [4]. Achieving effective controllability extends beyond simply offering an interface. A user-controllable recommender system should be capable of explaining the reasoning

* Corresponding Author. Email: juntao.tan@rutgers.edu.

behind its recommendations and the consequences of each user action. Therefore, we believe that harnessing the explainability of these systems has the potential to enhance user controllability.

In this paper, we propose a user-controllable recommendation framework with two essential properties. The first property, **retrospective controllability** with **retrospective explanation**, enables users to identify how their previous behaviors, such as clicks, purchases, and likes, contribute to current recommendations. The system should also provide the ability to delete or disregard specific past behaviors. This concept aligns with users' *right to be forgotten (RTBF)* on the internet [25]. The second property, **prospective controllability** with **prospective explanation**, offers users insight into how their interaction with the currently recommended item(s) will influence future recommendations. By operating in an informed manner, users can effectively control their future recommendations toward desired directions.

To achieve these objectives, we investigate the integration of explainability and controllability in a unified recommender system. Recent advancements in counterfactual reasoning have enabled accurate explanations for black-box machine learning models [27, 29, 10, 28, 30, 24, 8, 9]. The core concept of counterfactual reasoning aligns well with user-controllable recommendation goals. Both retrospective and prospective controllability involve addressing counterfactual "what if" questions: 1) What recommendations would have been made if the user had (not) interacted with specific items previously? and 2) How will future recommendations change if the user interacts with certain items presently? Providing counterfactual retrospective and prospective explanations empowers users with a comprehensive understanding of the consequences of their past and current behaviors, allowing them to selectively exclude such behaviors when necessary. Thus, we propose a model-agnostic framework that generates counterfactual retrospective and prospective explanations, informing users of their past and future recommendation impact and enabling informed decision-making and intervention in the recommendation process through engagement with these explanations.

Another significant contribution of this research is the incorporation of a quantitative evaluation of controllability in recommender systems, an aspect often overlooked in previous controllable recommendation studies. We introduce distinct evaluation protocols for retrospective and prospective controllability. For retrospective controllability, we assess the framework from two unique perspectives: 1) complexity of controllability, which quantifies the cost of intervening in past behaviors to modify current recommendations, and 2) accuracy of controllability, measuring the system's effectiveness in adjusting recommendations according to provided explanations when users implement suggested modifications. In terms of prospective controllability, we simulate user behaviors in a real-world context and demonstrate that the inclusion of control options has the potential to yield more accurate recommendations in the future. This assertion is supported by standard evaluation metrics in recommendation.

In summary, our work contributes the following:

- We propose the first user-controllable recommendation framework utilizing counterfactual explanations.
- We define two categories of counterfactual explanations for user-controllable recommendations: counterfactual retrospective explanations and counterfactual prospective explanations.
- We introduce two perspectives for evaluating controllability and, based on these, develop two metrics to enable the standardized assessment of controllable recommendation systems.
- For the first time, we quantitatively demonstrate the potential for

improving recommender system accuracy by providing users with additional control options.

2 Related Works

Most of the existing research on controllable recommendation is conducted from a user study perspective rather than a machine learning perspective. Early works in controllable recommendation focus on explicitly defining users' preferences regarding specific content. For instance, in [22], an interaction strategy was introduced where users can indicate the weights they assign to each item attribute to express their preferences. Similarly, [14] implemented a music recommender system that enables users to intervene in the recommendation process by selecting preferred item categories or editing their profiles. The authors observed that when the system provides highly precise recommendations, user intervention leads to increased satisfaction. In [32], a recommender system was proposed for recommending menus at restaurants, allowing users to update their preferences and provide information about allergies towards food ingredients. However, these methods necessitate users to explicitly answer questions, which may not always be feasible as users might face difficulty accurately expressing their preferences.

Some other studies incorporate implicit updating of user preferences through dynamic feedback. In these studies, the terms "interactive" and "controllable" are often used interchangeably. In [12], users were given the ability to directly interact with the recommendation results by re-ranking the generated recommendation lists to align with their expectations. Upon re-ranking, the system automatically updates the users' preferences for two attributes: popularity and age of the items. However, [12] has certain limitations, such as the inability for users to remove disliked items and a limited number of item attributes. [5] proposed an interactive recommender system that utilizes a three-column display. On the left side, users can adjust the weights assigned to items and content that contribute to the recommendation results, and they can cease adjusting once they are satisfied with the recommended items. A similar interface was created in [26], allowing users to add, delete, and re-rate items on the left side while observing the resulting changes in recommendations. In these studies, users are aware of the impacts of their interactions. However, when compared to our proposed method, the aforementioned works have several limitations: 1) Users are not provided with explanations for each individual recommendation, forcing them to guess the relationship between their actions and the recommendations, often leading to trials and errors. 2) The longer users engage with the recommender system, the more complex their accumulated behavior records become, making it increasingly challenging for users to make accurate assumptions and control the system. Our method distinguishes itself from previous works in that we employ machine learning approaches to generate counterfactual explanations for controllability. Additionally, our method is not restricted to a specific recommendation model; it can be applied to any recommendation model that relies on user behaviors for making recommendations.

3 User Controllable Recommendation

In this section, we introduce the two types of controllable counterfactual explanations in Section 3.1 and Section 3.2. Subsequently, we mathematically formulate the user-controllable recommendation framework in Section 3.3. Lastly, in Section 3.4, we discuss and define two fundamental aspects of user-controllable recommendation: the complexity of controllability and the accuracy of controllability.

3.1 Counterfactual Retrospective Explanation

Counterfactual retrospective explanations are derived from users' past behaviors and address the question: "If past behaviors were different, would this item still be recommended?" These explanations enable users to understand which historical actions influenced current recommendations. They take the following form:

We recommend this item because you interacted with [item(s)]. Revoke these behaviors to stop its recommendation.

The term "interacted with" can be substituted with expressions such as "liked," "viewed," "purchased," "clicked," and others, depending on the specific form of interactions within the platform. Control via counterfactual retrospective explanation in recommendation is illustrated in the top-right part of Figure 1.

3.2 Counterfactual Prospective Explanation

Counterfactual prospective explanations are generated based on users' present behaviors. These explanations offer insights into future outcomes. When users engage in a specific behavior, prospective explanations ask the question: "If this behavior is carried out, how will future recommendations differ compared to not engaging in it?". They take the following form:

With the current interaction, [item(s)] will be added to future recommendations. Revoke this behavior to prevent their recommendation.

Similarly, the choice of a specific interaction type depends on the real-world context. Control via counterfactual prospective explanation is illustrated in the bottom-right part of Figure 1.

3.3 Definition of User Controllable Recommendation

Technically, the proposed framework is applicable to any recommender system that generates recommendations based on users' historical behaviors. In this paper, we specifically introduce the controllable recommendation problem within the context of sequential recommendation. This choice is motivated by the fact that sequential recommender systems generate recommendations by considering user-item interactions in chronological order, and the time-aware nature of such systems aligns seamlessly with the concept of retrospective and prospective controllability.

Suppose a user set $\mathcal{U} = \{u_1, u_2, \dots, u_m\}$ with m users, and an item set $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ with n items. For each item v_j , e_j represents the related d -dimensional item embedding. A recommendation model learns the embedding matrix $\mathbf{E} \in \mathbb{R}^{n \times d}$ for all items.

For each user u_i with last T timestamps of the interaction history $\{v_t^i \mid t = 1, \dots, T\}$, $\mathbf{S}_i = [e_t^i \mid t = 1, \dots, T]$ denotes the concatenated item embeddings. T is a hyper parameter defined by the system, which indicates the maximum length of user history to be considered when making a recommendation.

A recommender system predicts ranking score r_{ij} for a user u_i on an item v_j based on \mathbf{S}_i , which can be generally represented as:

$$r_{ij} = f(v_j, \mathbf{S}_i \mid \Theta) \quad (1)$$

where Θ are the model parameters. The function f depends on the method used in the base recommendation model. In the rest of the paper, Eq. (1) is simplified as $r_{ij} = f(v_j, \mathbf{S}_i)$.

With the above ranking function, for each user u_i , the recommender system will generate a top- K recommendation, denoted as $\mathcal{R}(i, K)$. Under the the retrospective controllable recommendation scenario, if u_i is not satisfied with some items in $\mathcal{R}(i, K)$, the system should allow the user to change their history \mathbf{S}_i by removing certain previous actions, such that the recommendation results will be changed to the desired direction.

Suppose the intervention is represented by a binary vector $\Delta_i \in \{0, 1\}^T$, which will be applied to \mathbf{S}_i to form a counterfactual user history \mathbf{S}_i^* such that:

$$\mathbf{S}_i^* = \mathbf{S}_i \odot (1 - \Delta_i) \quad (2)$$

where \odot stands for the element-wise product. Each item embedding $e_t^i \in \mathbf{S}_i$ will be multiplied by the related $1 - \delta_t$, where $\delta_t \in \Delta_i$ is the t^{th} element in Δ_i . When $\delta_t = 0$, the item embedding e_t^i at timestamp t remains unchanged, which means the corresponding behavior of user u_i at timestamp t is preserved. When $\delta_t = 1$, e_t^i will become the padding embedding after multiplication, which means the related behavior is removed. After changing these behaviors, the base recommender will generate a new recommendation list $\mathcal{R}^*(i, K)$, which should better fit the user's expectation.

Under the prospective controllable recommendation scenario, the control option is more straight-forward. Suppose user i interacts with a new item v_x and generates a new user history embedding $[\mathbf{S}_i, e_x]$. The user will be informed of the change in the future (i.e., the new items that will be added to the future recommendation). Users are allowed to decide whether this particular item, v_x , should be considered by the system or not in future recommendations.

3.4 Complexity and Accuracy of Controllability

Prior to the design and evaluation of a user-controllable recommendation model, it is crucial to establish the desired properties of an effective controllable method. In this paper, we propose two properties of controllability within the controllable recommendation scenario.

The first property is the complexity of controllability, which measures the cost for the user to control the recommendation results. Given that our framework relies on user behavior, complexity can be defined as the proportion of a user's past behaviors that must be eliminated in order to alter the recommendation results towards the intended direction. With the above notation, for a user u_i , when the change Δ_i is applied to the history embedding \mathbf{S}_i , the control complexity can be defined as:

$$\text{Complexity} = \frac{\|\Delta_i\|_0}{T}, \quad (3)$$

where $\|\Delta_i\|_0$ is the number of ones in Δ_i , i.e., the number of removed behaviors of user u_i . We note that in reality, the length of some user histories may be smaller than T . In this case, it should be replaced by the real length of the user history. This applies to all the equations in the rest of the paper.

The second property is the accuracy of controllability. The control options offered by the system should accurately eliminate the undesired items from the recommendation list without inadvertently removing other relevant items. Suppose u_i is unsatisfied with a subset of items in the recommendation list $\mathcal{R}_g(i, K) \subseteq \mathcal{R}(i, K)$, in an ideal situation, the item(s) removed from $\mathcal{R}(i, K)$ should only be the item(s) in $\mathcal{R}_g(i, K)$, which can be expressed as:

$$\text{Accurate control} \Leftrightarrow \mathcal{R}(i, K) \setminus \mathcal{R}^*(i, K) = \mathcal{R}_g(i, K) \quad (4)$$

However, a 100% accurate control may be impossible to acquire. Thus, we further define the accuracy of the controllability as the Jacard similarity (also known as Intersection over Union (IoU)) between the set of removed items and the set of undesired items. This can be formulated as follows:

$$\text{Accuracy} = \frac{|\mathcal{R}(i, K) \setminus \mathcal{R}^*(i, K) \cap \mathcal{R}_g(i, K)|}{|\mathcal{R}(i, K) \setminus \mathcal{R}^*(i, K) \cup \mathcal{R}_g(i, K)|} \quad (5)$$

The accuracy of controllability should be maximized to ensure that only the undesired items are removed, while preserving the other items in the recommendation list. Simultaneously, the number of altered behaviors (i.e., the complexity) in the explanation should be minimized to enhance the comprehensibility of the explanation and facilitate the users in easily implementing the provided control options.

We note that these two properties not only guide the design of user controllable recommendation systems but are also highly suitable for quantitatively evaluating the controllabilities.

4 Generate Controllable Explanation

In this section, we introduce the algorithms for generating the two types of explanations.

4.1 Generate Counterfactual Retrospective Explanation

To generate counterfactual retrospective explanations, we propose a constrained optimization problem that incorporates the previously defined complexity and accuracy of controllability.

The objective of the algorithm is to generate effective counterfactual retrospective explanations, ensuring that by revoking the past behaviors, the undesired items are removed from the recommendation list. This serves as a strict constraint in the optimization process. Meanwhile, as discussed in Section 3.4, the objective part of the optimization aims to minimize the complexity of the generated explanation while maximizing its accuracy. Considering the generation of an explanation to remove an item v_j recommended to user u_i , and based on the mathematical definitions of complexity and accuracy, we formulate the optimization problem as follows:

$$\begin{aligned} & \underset{\Delta_{i,j}}{\text{minimize}} \quad \frac{\|\Delta_{i,j}\|_0}{T} - \gamma \left(\frac{|\mathcal{R}(i, K) \setminus \mathcal{R}^*(i, K) \cap \{v_j\}|}{|\mathcal{R}(i, K) \setminus \mathcal{R}^*(i, K) \cup \{v_j\}|} \right) \\ & \text{s.t., } v_j \notin \mathcal{R}^*(i, K) \end{aligned} \quad (6)$$

where γ is a hyper-parameter to control the two terms on the same scale. Since $\frac{\|\Delta_{i,j}\|_0}{T} \sim \|\Delta_{i,j}\|_0$ and v_j is guaranteed to be removed as in the constraint, Eq.(6) is equivalent to:

$$\begin{aligned} & \underset{\Delta_{i,j}}{\text{minimize}} \quad \|\Delta_{i,j}\|_0 - \gamma \frac{1}{|\mathcal{R}(i, K) \setminus \mathcal{R}^*(i, K)|} \\ & \text{s.t., } v_j \notin \mathcal{R}^*(i, K) \end{aligned} \quad (7)$$

In Eq. (7), the second term is, in fact, maximizing the intersection between the original and counterfactual recommendation lists. Consequently, for the purpose of enhancing simplicity and feasibility of optimization, we reformulate the equation as follows:

Algorithm 1 Greedy Search

Input: user u_i , target item v_j , history length T , user interaction history embedding \mathbf{S}_i , heuristic function h , length of recommendation list K .
Output: the optimized explanation set $\mathcal{E}_{i,j}$.

```

1: Initialize explanation set  $\mathcal{E}_{i,j} \leftarrow \{\}$ 
2:  $\mathbf{S}_i^* = \mathbf{S}_i$ 
3: while  $v_j$  in the top- $K$  recommendation AND  $|\mathcal{E}_{i,j}| < T$  do
4:   // Choose single behavior in a greedy manner.
5:   lowest_heuristic  $\leftarrow \infty$ 
6:   for each past behavior  $\notin \mathcal{E}_{i,j}$  do
7:     Generate one-hot vector  $\mathbf{p}$ 
8:     if  $h(\mathbf{p}) < \text{lowest\_heuristic}$  then
9:       update the best behavior according to  $\mathbf{p}$ 
10:      update lowest heuristic
11:    end if
12:  end for
13:  Add the best behavior to  $\mathcal{E}_{i,j}$ 
14: end while
15: return  $\mathcal{E}_{i,j}$ 

```

$$\underset{\Delta_{i,j}}{\text{minimize}} \quad \|\Delta_{i,j}\|_0 - \gamma \sum_{v_q \in \mathcal{R}(u_i, K), v_q \neq v_j} \mathbb{1}(v_q, \mathcal{R}^*(i, K)) \quad (8)$$

$$\text{s.t., } v_j \notin \mathcal{R}^*(i, K)$$

where $\mathbb{1}(v_q, \mathcal{R}^*(i, K)) = 1$ if $v_q \in \mathcal{R}^*(i, K)$, otherwise 0.

After solving $\Delta_{i,j}$, for $\forall \delta_t \in \Delta_{i,j}$, if $\delta_t = 1$, the related interacted item v_t^i will be included in the generated explanation set $\mathcal{E}_{i,j}$. However, searching for an optimal set of past behaviors is exponential. For a behavior set with size x , the searching space contains $\mathcal{O}(2^x)$ subsets. Therefore, inspired by a similar approach in [10], we provide two options to approximate the optimization problem. One based on greedy search, and the other utilizing continuous relaxation.

Greedy Search: In the greedy search approach, we iteratively choose a single behavior to revoke at each time step, until the constraint in Eq. (8) is satisfied, i.e., the accumulative effects remove item v_j from the recommendation list $\mathcal{R}(i, K)$. At each step, we choose the behavior that fits both the constraint and objective in Eq. (8) best. More specifically, instead of globally optimizing $\Delta_{i,j}$, we choose each item at one time by learning a one-hot vector $\mathbf{p}_{i,j}^{(r)}$ at each iterative step r . By removing the related item from user history, the ranking score of the target item v_j should be as small as possible while the ranking scores of the other items in $\mathcal{R}(i, K)$ should be as large as possible. This strategy is guided by minimizing a heuristic function $h(\mathbf{p}_{i,j}^{(r)})$:

$$\underset{\mathbf{p}_{i,j}^{(r)}}{\text{minimize}} \quad h(\mathbf{p}_{i,j}^{(r)}) = f(v_j, \mathbf{S}_i^{*(r)}) - \gamma_1 \sum_{v_q \in \mathcal{R}(i, K), v_q \neq v_j} f(v_q, \mathbf{S}_i^{*(r)})$$

$$\text{s.t., } \mathbf{p}_{i,j}^{(r)} = \{0, 1\}^T; \quad \|\mathbf{p}_{i,j}^{(r)}\|_1 = 1$$

$$\text{where } \mathbf{S}_i^{*(r)} = \mathbf{S}_i^{*(r-1)} \odot (\mathbb{1} - \mathbf{p}_{i,j}^{(r)}); \quad \mathbf{S}_i^{*(0)} = \mathbf{S}_i \quad (9)$$

According to Eq. (9), we select one behavior from all the unchanged past behaviors to revoke at each step until v_j is removed from the recommendation list, and return the change vector $\Delta_{i,j} = \sum_r \mathbf{p}_{i,j}^{(r)}$. If all the past behaviors have been changed but the item v_j is still in the recommendation list, then the algorithm fails to generate retrospective explanation for this user-item pair. This process is illustrated in Algorithm 1.

Continuous Relaxation: We provide another option to solve Eq. (8) by transforming the original combinatorial problem into a continuous optimization problem. Since both the constraint part and the

objective part are not differentiable, in order to learn the changing vector $\Delta_{i,j}$, we relax each component in Eq. (8) into continuous space to provide approximate solutions.

For the complexity objective, we relax the binary vector $\Delta_{i,j} \in \{0, 1\}^T$ as a continuous vector $\Delta_{i,j}^c \in [0, 1]^T$. $\|\Delta_{i,j}^c\|_1$ is used to approximate $\|\Delta_{i,j}\|_0$.

For the constraint, $v_j \notin \mathcal{R}^*(i, K)$. It is equivalent to the ranking score of v_j is smaller than the ranking score of the item at the K th position of the recommendation list. Thus, we further relax it as a pair-wise loss between these two items, which is:

$$R_1 = \max(0, \alpha_1 + f(v_j, \mathbf{S}_i^*) - f(v_K, \mathbf{S}_i^*)) \quad (10)$$

where v_K is the item at the K th position of the top- K recommendation list, which may be updated at each step.

For the accuracy objective $\gamma \sum_{v_q \in \mathcal{R}(i, K), v_q \neq v_j} \mathbb{1}(v_q, \mathcal{R}^*(i, K))$, we would like the ranking score of an item v_q to remain as large as possible. Thus, we introduce another loss that measures the average ranking score of the items that shouldn't be removed, which is:

$$R_2 = -\frac{1}{k-1} \sum_{v_q \in \mathcal{R}(i, K), v_q \neq v_j} f(v_q, \mathbf{S}_i^*) \quad (11)$$

After the relaxations, the final equation of the optimization problem is reformulated as follows:

$$\underset{\Delta_{i,j}^c}{\text{minimize}} \|\Delta_{i,j}^c\|_1 + \lambda(R_1 + \gamma_2 R_2) \quad (12)$$

Eq. (12) is differentiable and can be solved by machine learning methods such as Stochastic Gradient Descent (SGD). After optimization, a threshold will be applied to $\Delta_{i,j}^c$ to select the interacted items that will be included in the generated explanation $\mathcal{E}_{i,j}$. In this paper, we set the threshold to 0.5. For continuous relaxation approach, the explanations should be generated only when they can successfully remove the target items. Thus, a post-checking process is required.

Efficiency and Scalability Discussion: We note that in this framework, the complexity of the optimization space is linearly related to the fixed length T of the user history and is irrelevant to the number of users and items in the system. This is attributed to the underlying mechanism of counterfactual optimization. Empirically, when $T = 100$ and the algorithm is executed on a single Nvidia A5000 GPU, it takes approximately 4 seconds for the greedy search method to generate an explanation for a single recommendation, while the relaxation approach takes approximately 7 seconds. The inference time does **not** encounter any scalability issues as the size of the dataset increases.

4.2 Generate Counterfactual Prospective Explanation

In comparison to searching for past behaviors that lead to a specific recommendation, it is more straightforward to determine the future consequences of the current behavior. This is due to the fact that when the recommender system is accessible, regardless of its transparency, the outcomes of any given current behavior are known.

Suppose a user u_i interacts with an upcoming item v_x . The generation of counterfactual prospective explanations relies on assessing how the future recommendation would be altered by incorporating this particular user behavior into their history. Therefore, we update \mathbf{S}_i with $\mathbf{S}_i' = \mathbf{S}_i \cup \{v_x\}$. $\mathcal{R}(i, K)'$ is the updated recommendation list in the future, which is generated based on \mathbf{S}_i' .

Our user-controllable recommendation method looks for the items that are newly added to the recommendation list to form the counterfactual prospective explanation. For each item v_j :

$$v_j \in \mathcal{E}_{i,j}, \quad \text{if } v_j \in (\mathcal{R}'(i, K) \setminus \mathcal{R}(i, K)) \quad (13)$$

If a user is dissatisfied with the change in the future recommendation list, they have the option to revoke the current interaction, thereby preventing the system from considering this behavior in the future. As the ‘‘harmful’’ interactions are eliminated dynamically, it is reasonable to assume that the recommender system can benefit from prospective controllability and deliver more precise recommendations in the future. In Section 5.5, we conduct experiments to validate this assumption.

5 Experiments

Table 1. Statistics of the datasets.

Dataset	#User	#Item	#Interaction	Density
ML-1M	6,040	3,416	987,540	0.031 %
Yelp	5,342	17,746	515,128	0.543 %

Table 2. Ranking accuracy of the base models

Models	ML-1M		Yelp	
	NDCG@10	HT@10	NDCG@10	HT@10
SASRec	0.1203	0.2265	0.0327	0.0518
GRU4Rec	0.1150	0.2242	0.0145	0.0294

5.1 Dataset Description

The experiments are conducted on two widely used public datasets: 1) MovieLens-1M¹ [11]: containing movie ratings from the MovieLens website. It has approximately 1 million ratings; 2) Yelp²: containing users’ interactions and ratings on various businesses. For the Yelp dataset, following previous works [36, 31], we remove the users with less than 20 interactions and items with less than 10 interactions. Table 1 shows the overall statistics of the two datasets.

5.2 Base Recommenders

We choose two popular sequential recommendation models as the base models:

- GRU4Rec[13]: It applies Recurrent Neural Network (RNN) with Gated Recurrent Units (GRU) to model user click sequences for session-based recommendation.
- SASRec[19]: It uses the self-attention mechanism to recommend next items.

5.3 Implementation Details

Training Base Models: When training these two models, we follow the same setup as in their original papers. For GRU4Rec, we apply 1 layer of GRU units. For SASRec, we apply 2 self-attention blocks. The item embedding size is 100 for both models. For both datasets, we set the maximum length T of user sequences to 100, which means

¹ <https://files.grouplens.org/datasets/movielens>

² <https://www.yelp.com/dataset>

Table 3. Complexity and accuracy of controllability in retrospective evaluation. All numbers are percentages. UCR_Relax and UCR_Search are our proposed methods, which represent the relaxation approach and greedy search approach, respectively. All numbers are percentages.

SASRec	ML-1M						Yelp					
	Complexity %↓			Accuracy %↑			Complexity %↓			Accuracy %↑		
	@3	@5	@10	@3	@5	@10	@3	@5	@10	@3	@5	@10
Dynamic[26]	29.28	29.80	32.26	55.55	42.03	24.26	31.20	32.20	35.32	52.17	39.03	22.61
EFM[36]	<u>10.22</u>	11.88	13.29	58.21	46.16	33.30	<u>12.96</u>	14.57	17.22	56.88	47.29	<u>33.29</u>
UCR_Relax	12.16	<u>10.54</u>	<u>9.67</u>	<u>67.02</u>	<u>53.14</u>	<u>36.40</u>	16.22	<u>14.54</u>	<u>14.02</u>	<u>64.27</u>	<u>50.70</u>	32.26
UCR_Search	4.99	4.58	4.67	76.79	65.28	48.02	8.23	7.92	7.88	78.94	68.02	49.52

GRU4Rec	ML-1M						Yelp					
	Complexity %↓			Accuracy %↑			Complexity %↓			Accuracy %↑		
	@3	@5	@10	@3	@5	@10	@3	@5	@10	@3	@5	@10
Dynamic[26]	26.89	29.67	32.49	60.61	41.52	25.27	46.47	49.51	48.83	57.92	40.72	24.87
EFM[36]	17.35	15.48	17.14	62.33	45.59	29.54	33.19	32.78	34.11	61.83	45.65	27.33
UCR_Relax	<u>14.73</u>	<u>13.84</u>	<u>13.25</u>	<u>74.49</u>	<u>52.93</u>	<u>30.52</u>	17.12	15.61	15.49	<u>75.57</u>	<u>55.57</u>	<u>32.13</u>
UCR_Search	5.51	5.69	6.40	77.66	59.38	38.46	<u>19.59</u>	<u>19.87</u>	<u>20.39</u>	77.38	59.95	37.33

Table 4. Fidelity of the relaxation method. All numbers are percentages.

Methods	ML-1M			Yelp		
	@3 %	@5 %	@10 %	@3 %	@5 %	@10 %
SASRec	62.97	59.76	55.14	52.37	49.06	46.55
GRU4Rec	82.00	80.70	78.31	54.30	52.26	50.30

only the last 100 interactions in the user history are taken into consideration when making predictions.

We use the same training strategy when training these two models on the two datasets. First, we follow the generic leave-one-out strategy [2, 34] to split the last item in the user sequences as the test data. We note that besides the test item, we also remove a certain number of last interacted items from the user history for prospective evaluation purpose, which will be introduced in Section 5.5. This doesn’t affect the training of the base models. During training, we use Adam optimizer[20] with 0.001 learning rate. The batch size is 128. We set dropout rate to 0.2 to prevent overfitting. Then, we train the model until they get the highest validation performance. We report the performance of the trained base models in Table 2. We note that the ranking accuracy of the base models are not the focus of this work. After training, the learned parameters in the recommendation models will be fixed for evaluating the controllability.

Generating Retrospective Explanations: For the relaxation approach in Eq.(12), we set the hype-parameter $\gamma_2 = 1$, and $\lambda = 10$. We also set the $\gamma_1 = 1$ in the greedy approach in Eq. (9). These parameters have varies effects on the generated explanations, which are shown in the ablation study in Section 5.4. For solving the relaxation optimization problem, we use Adam optimizer with learning rate equals to 0.01 and learn the change vector for 500 steps.

Generating Prospective Explanations: Given that prospective explanations aim to inform users about the potential outcomes of their ongoing interaction, we employ the next one item following the training sequence to simulate this process. To be precise, for each user, we append the next item to the end of the training sequence and observe the resulting “future” changes. All evaluations are conducted based on this procedure, which we elaborate on in Section 5.5.

5.4 Retrospective Evaluation

We randomly sample 1000 users as the test set and generate retrospective explanations for them. When evaluating the controllability, we revoke the interactions included in the generated controllable ex-

planations by changing the related item embeddings to the padding embedding in the user interaction sequences.

Then, we evaluate the controllability according to the complexity and accuracy as introduced in Section 3.4.

Comparable Baselines: To the best of our knowledge, there is no such paper that proposes an optimization algorithm to control the recommender system based on the user behaviors. Thus, we adjust two existing works in controllable recommendation and explainable recommendation to make them comparable: 1) **Dynamic Feedback** [26]: It proposes a manual process to provide controllability, allowing the users to randomly remove the past behaviors until they are satisfied with the recommendation list. We create a random search algorithm to simulate this manual process; 2) **EFM** [36]: An explainable recommendation framework grounded in the concept of matching, a common approach in explainable recommendation [31, 6]. If an item is recommended, similar items in the user’s history may be the reason for that recommendation. In each step, we remove the best-aligned item (according to the learned latent embeddings) until the target item is removed.

Quantitative Analysis: In Table 3, we evaluate the complexity and accuracy of the controllability for each baseline. First, we observe that both the greedy search and continuous relaxation methods provide significantly better controllability than the other methods. On average, based on the SASRec model, the relaxation method outperforms the second-best baseline by 2.20% in terms of complexity and by 9.45% in terms of accuracy. The greedy search method outperforms the second best baseline by 52.31% in terms of complexity and by 41.49% in terms of accuracy. While based on GRU4Rec model, the relaxation method outperforms the second-best baseline by 33.96% in terms of complexity and by 16.74% in terms of accuracy. The greedy search method outperforms the second-best baseline by 52.45% in terms of complexity and by 29.68% in terms of accuracy. It is worth noting that complexity is defined as the ratio between the number of items included in a generated explanation and the length of the user sequence. On average, the number of generated items is 3.48 for the searching method and 5.66 for the relaxation method. Although the relaxation and searching approaches produce comparable results when generating controllable explanations for the Yelp dataset based on GRU4Rec, according to all other results, the searching approach is generally the better option. Furthermore, the relaxation approach is unable to generate sufficient explanations for all the recommended items, for which we report the fidelity (i.e., the percentage of successful generation) in Table 4. Therefore, we sug-

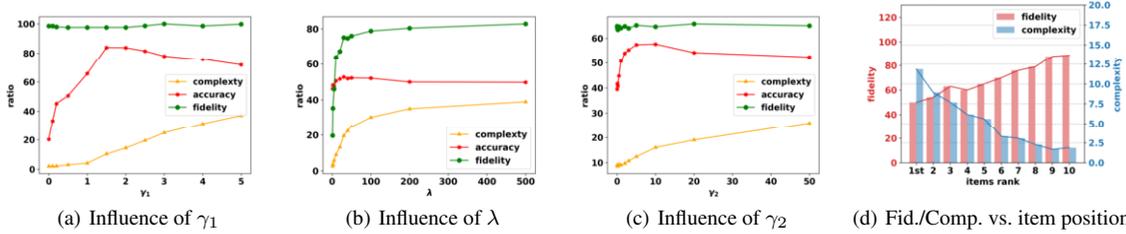


Figure 2. Influence of the hyper-parameters according to the experiments on SASRec with MovieLens dataset. (a) Influence of γ_1 in greedy approach. (b) Influence of λ in relaxation approach. (c) Influence of γ_2 in relaxation approach. (d) Fidelity and Complexity to remove items at different ranking positions.

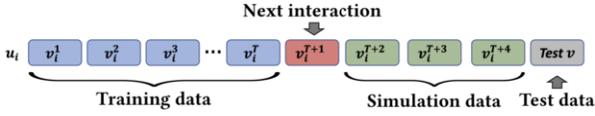


Figure 3. Data splitting in prospective evaluation.

gest always starting with the greedy search method when applying the proposed framework to new datasets or models.

Next, as shown in Figure 2, we conduct ablation studies to verify the influences of γ_1 in the greedy search method, as well as λ and γ_2 in the continuous relaxation method. First, Figure 2 (a) shows the influence of γ_1 in the greedy search method. When γ_1 is small, the control accuracy is highly positively related to it. However, when $\gamma_1 > 1$, the complexity increases dramatically. Therefore, choosing $\gamma_1 = 1$ is the optimal option. For the relaxation approach, as shown in Figure 2 (b), the λ controls the trade-off between fidelity and complexity. When λ increases, more effective explanations can be generated, but at the same time, the complexity also increases. Figure 2 (c) shows the influence of γ_2 on the relaxation method. It is very similar to the influence of γ_1 in the searching approach. When γ_2 is small, the control accuracy improves as γ_2 increases. However, when $\gamma_2 > 10$, choosing a larger γ_2 starts to negatively impact the model’s performance. Furthermore, the complexity of controllability also steadily increases as γ_2 increases.

Moreover, Figure 2 (d) displays the fidelity and complexity of the generated controllable explanations using the relaxation method, with respect to removing items at different positions in the recommendation list. It is evident that removing items ranked higher in the recommendation list is more challenging. This aligns with the notion that higher-ranked items have stronger reasons for their recommendation, thus more behaviors have to be revoked for the removing.

5.5 Prospective Evaluation

With prospective controllability, users have the ability to revoke their current behavior if they are dissatisfied with the ongoing changes, in hopes of receiving better recommendations in the future. In this section, we simulate this real-world scenario by introducing a novel data splitting strategy to validate this assumption. As illustrated in Figure 3, the basic idea is as follows: in addition to excluding the test item, we also omit the last M items in the user sequences as a "simulation set" that will not be observed during training. When a user interacts with a new item (i.e., adds the next interaction to the original user sequence), we examine the changes between the new recommendation list and the original recommendation list. If none of the newly added items are present in the simulation set, we can reasonably assume that in a real-world scenario, the user may not prefer this prospective change. Consequently, they would revoke this behavior. Then, for all user sequences that satisfy this condition, we

test whether removing the current behavior from the user sequences will improve the accuracy on the test data.

During the experiments, we set $M = 20$. We note that we concatenate the original user sequence, regardless of whether the current item is removed or not, with the simulation subsequence when making predictions. This is because we would like to test the long-term benefits of the prospective control option, rather than relying on short-term trivial observations. As shown in Table 5, we observe some intriguing results. When we remove the next interaction for **all** users and compare it to making predictions with the unchanged user history, the ranking accuracy drops significantly. This outcome is expected because removing an item from the user sequence results in the loss of valuable information. However, if we only remove the next interacted items for the users selected by the aforementioned simulation process, the overall ranking accuracy for this subset of users actually improves. This proves that, in addition to enhancing the transparency of the recommender system, providing system users with prospective controllability significantly increases the likelihood of receiving more accurate recommendations in the future.

Table 5. Change of ranking accuracy

SASRec	ML-1M		Yelp	
	NDCG@10	HT@10	NDCG@10	HT@10
All users	0.1203	0.2265	0.0327	0.0518
All users (revoke)	0.0987 (↓)	0.1905 (↓)	0.0244 (↓)	0.0423 (↓)
Target users	0.1078	0.1995	0.0301	0.0479
Target users (revoke)	0.1118 (↑)	0.2104 (↑)	0.0322 (↑)	0.0496 (↑)
GRU4Rec	ML-1M		Yelp	
	NDCG@10	HT@10	NDCG@10	HT@10
All users	0.1150	0.2242	0.0145	0.0294
All users (revoke)	0.0973 (↓)	0.1914 (↓)	0.0136 (↓)	0.0278 (↓)
Target users	0.1070	0.2074	0.0138	0.0283
Target users (revoke)	0.1082 (↑)	0.2094 (↑)	0.0140 (↑)	0.0283 (-)

6 Conclusion and Future Work

In this paper, we present a framework that integrates explainability and controllability for user-controllable recommendation. By incorporating counterfactual retrospective and prospective explanations, our proposed framework enables users to control over the recommender system based on the user behaviors. As a first attempt, we introduce two fundamental properties of controllability: the complexity of controllability and the accuracy of controllability. These properties serve as general evaluation metrics for assessing the user-controllable recommendation problem. Controllability is important not only for recommendation, but also for many other AI sub-fields. In the future, we expect the exploration of this method beyond recommendation systems, extending its application to diverse AI systems, such as vision-based and language-based intelligent systems.

Acknowledgement This work was supported in part by NSF IIS 1910154, 2007907, 2046457, 2127918. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

References

- [1] Xavier Amatriain, Josep Pujol, Nava Tintarev, and Nuria Oliver, ‘Rate it again: Increasing recommendation accuracy by user re-rating’, pp. 173–180, (01 2009).
- [2] Ting Bai, Lixin Zou, Wayne Xin Zhao, Pan Du, Weidong Liu, Jian-Yun Nie, and Ji-Rong Wen, ‘Ctrec: A long-short demands evolution model for continuous-time recommendation’, in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 675–684, (2019).
- [3] Krisztian Balog, Filip Radlinski, and Shushan Arakelyan, ‘Transparent, scrutable and explainable user models for personalized recommendation’, in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 265–274, (2019).
- [4] Gabriel Diniz Junqueira Barbosa and Simone Diniz Junqueira Barbosa, ‘You should not control what you do not understand: the risks of controllability in ai’, *Human Computer Interaction and Emerging Technologies: Adjunct Proceedings from*, 231, (2020).
- [5] Svetlin Bostandjiev, John O’Donovan, and Tobias Höllerer, ‘Tasteweights: a visual interactive hybrid recommender system’, in *Proceedings of the sixth ACM conference on Recommender systems*, pp. 35–42, (2012).
- [6] Tong Chen, Hongzhi Yin, Guanhua Ye, Zi Huang, Yang Wang, and Meng Wang, ‘Try this instead: Personalized and interpretable substitute recommendation’, in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 891–900, (2020).
- [7] Yingqiang Ge, Shuchang Liu, Zuohui Fu, Juntao Tan, Zelong Li, Shuyuan Xu, Yunqi Li, Yikun Xian, and Yongfeng Zhang, ‘A survey on trustworthy recommender systems’, *arXiv preprint arXiv:2207.12515*, (2022).
- [8] Yingqiang Ge, Juntao Tan, Yan Zhu, Yinglong Xia, Jiebo Luo, Shuchang Liu, Zuohui Fu, Shijie Geng, Zelong Li, and Yongfeng Zhang, ‘Explainable fairness in recommendation’, in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 681–691, (2022).
- [9] Azin Ghazimatin, Oana Balalau, Rishiraj Saha Roy, and Gerhard Weikum, ‘Prince: Provider-side interpretability with counterfactual explanations in recommender systems’, in *Proceedings of the 13th International Conference on Web Search and Data Mining*, pp. 196–204, (2020).
- [10] Yash Goyal, Ziyang Wu, Jan Ernst, Dhruv Batra, Devi Parikh, and Stefan Lee, ‘Counterfactual visual explanations’, in *International Conference on Machine Learning*, pp. 2376–2384. PMLR, (2019).
- [11] F Maxwell Harper and Joseph A Konstan, ‘The movielens datasets: History and context’, *Acm transactions on interactive intelligent systems (tiis)*, 5(4), 1–19, (2015).
- [12] F. Maxwell Harper, Funing Xu, Harmanpreet Kaur, Kyle Condiff, Shuo Chang, and Loren Terveen, ‘Putting users in control of their recommendations’, *RecSys ’15*, p. 3–10, New York, NY, USA, (2015). Association for Computing Machinery.
- [13] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk, ‘Session-based recommendations with recurrent neural networks’, *arXiv preprint arXiv:1511.06939*, (2015).
- [14] Yoshinori Hijikata, Yuki Kai, and Shogo Nishida, ‘The relation between user intervention and user satisfaction for information recommendation’, in *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, pp. 2002–2007, (2012).
- [15] Yifan Hu, Yehuda Koren, and Chris Volinsky, ‘Collaborative filtering for implicit feedback datasets’, in *2008 Eighth IEEE International Conference on Data Mining*, pp. 263–272, (2008).
- [16] D. Jannach, Sidra Naveed, and Michael Jugovac, ‘User control in recommender systems: Overview and interaction challenges’, in *EC-Web*, (2016).
- [17] Dietmar Jannach, Michael Jugovac, and Ingrid Nunes, ‘Explanations and user control in recommender systems’, in *Proceedings of the 23rd International Workshop on Personalization and Recommendation on the Web and Beyond*, eds., Mirjam Augstein, Eelco Herder, Wolfgang Würndl, and Enes Yigitbas, p. 31, New York, NY, (2019). ACM.
- [18] Michael Jugovac and Dietmar Jannach, ‘Interacting with recommenders—overview and research directions’, *ACM Transactions on Interactive Intelligent Systems (TiIS)*, 7(3), 1–46, (2017).
- [19] Wang-Cheng Kang and Julian McAuley, ‘Self-attentive sequential recommendation’, in *2018 IEEE international conference on data mining (ICDM)*, pp. 197–206. IEEE, (2018).
- [20] Diederik P Kingma and Jimmy Ba, ‘Adam: A method for stochastic optimization’, *arXiv preprint arXiv:1412.6980*, (2014).
- [21] Bart P. Knijnenburg, Svetlin Bostandjiev, John O’Donovan, and Alfred Kobsa, ‘Inspectability and control in social recommenders’, in *Proceedings of the Sixth ACM Conference on Recommender Systems, RecSys ’12*, p. 43–50. Association for Computing Machinery, (2012).
- [22] Bart P Knijnenburg, Niels JM Reijmer, and Martijn C Willemsen, ‘Each to his own: how different users call for different interaction methods in recommender systems’, in *Proceedings of the fifth ACM conference on Recommender systems*, pp. 141–148, (2011).
- [23] B. Lamche, Ummahan Adıgüzel, and W. Würndl, ‘Interactive explanations in mobile shopping recommender systems’, *CEUR Workshop Proceedings*, 1253, 14–21, (01 2014).
- [24] Ramaravind K Mothilal, Amit Sharma, and Chenhao Tan, ‘Explaining machine learning classifiers through diverse counterfactual explanations’, in *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 607–617, (2020).
- [25] European Parliament and Council of the European Union. General data protection regulation (gdpr), art. 17, right to erasure (‘right to be forgotten’), 2016. <https://gdpr.eu/article-17-right-to-be-forgotten/>.
- [26] James Schaffer, Tobias Hollerer, and John O’Donovan, ‘Hypothetical recommendation: A study of interactive profile manipulation behavior for recommender systems’, in *The Twenty-Eighth International Flairs Conference*, (2015).
- [27] Iliia Stepin, Jose M Alonso, Alejandro Catala, and Martín Pereira-Fariña, ‘A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence’, *IEEE Access*, 9, 11974–12001, (2021).
- [28] Juntao Tan, Shijie Geng, Zuohui Fu, Yingqiang Ge, Shuyuan Xu, Yunqi Li, and Yongfeng Zhang, ‘Learning and evaluating graph neural network explanations based on counterfactual and factual reasoning’, in *Proceedings of the ACM Web Conference 2022*, pp. 1018–1027, (2022).
- [29] Juntao Tan, Shuyuan Xu, Yingqiang Ge, Yunqi Li, Xu Chen, and Yongfeng Zhang, ‘Counterfactual explainable recommendation’, in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pp. 1784–1793, (2021).
- [30] Juntao Tan and Yongfeng Zhang, ‘Explainablefold: Understanding alphafold prediction with explainable ai’, *arXiv preprint arXiv:2301.11765*, (2023).
- [31] Nan Wang, Hongning Wang, Yiling Jia, and Yue Yin, ‘Explainable recommendation via multi-task learning in opinionated text data’, in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 165–174, (2018).
- [32] Rainer Wasinger, James Wallbank, Luiz Pizzato, Judy Kay, Bob Kummerfeld, Matthias Böhmer, and Antonio Krüger, ‘Scrutable user models and personalised item recommendation in mobile lifestyle applications’, in *International Conference on User Modeling, Adaptation, and Personalization*, pp. 77–88. Springer, (2013).
- [33] Bo Xiao and Izak Benbasat, ‘E-commerce product recommendation agents: Use, characteristics, and impact’, *MIS quarterly*, 137–209, (2007).
- [34] Hamed Zamani and W Bruce Croft, ‘Learning a joint search and recommendation model from user-item interactions’, in *Proceedings of the 13th International Conference on Web Search and Data Mining*, pp. 717–725, (2020).
- [35] Yongfeng Zhang and Xu Chen, ‘Explainable recommendation: A survey and new perspectives’, *Foundations and Trends in Information Retrieval*, (2020).
- [36] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma, ‘Explicit factor models for explainable recommendation based on phrase-level sentiment analysis’, in *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pp. 83–92, (2014).