

# Would You Like Your Data to Be Trained? A User Controllable Recommendation Framework

Lei Wang<sup>1</sup>, Xu Chen<sup>1,2\*</sup>, Zhenhua Dong<sup>3</sup>, Quanyu Dai<sup>3</sup>

<sup>1</sup>Gaoling School of Artificial Intelligence, Renmin University of China

<sup>2</sup>Beijing Key Laboratory of Big Data Management and Analysis Methods

<sup>3</sup>Noah's Ark Lab, Huawei

wanglei154@ruc.edu.cn, xu.chen@ruc.edu.cn, dongzhenhua@huawei.com, daiquanyu@huawei.com

## Abstract

Recommender systems have a significant impact on various real-world applications, shaping people's daily lives and enhancing productivity. Traditional recommender models aim to collect extensive user information to accurately estimate user preferences. However, in practical scenarios, users may not want all their behaviors to be included in the model training process. This paper introduces a novel recommendation paradigm that allows users to indicate their "willingness" regarding which data should contribute to model training. The models are then optimized to maximize utility, which considers the trade-off between recommendation performance and respecting user preferences. The recommendation problem is formulated as a multiplayer game, with each user acting as a player and using a selection vector to indicate their willingness to include specific interacted items in training. To efficiently solve this game, an influence function-based model is proposed to approximate recommendation performances for different actions without re-optimizing the model. Furthermore, an enhanced model leveraging multiple anchor actions for the influence function is introduced to improve performance approximation accuracy. The convergence rate of the algorithm is theoretically analyzed, and the advantages of incorporating multiple anchor actions are demonstrated. Extensive experiments on both simulated and real-world datasets validate the effectiveness of the proposed models in balancing recommendation quality and user willingness. To promote this research direction, we have released our project at <https://paitesanshi.github.io/IFRQE/>.

## Introduction

Recommender systems are critical components in real-world applications, providing significant business value and user convenience. Collaborative filtering (CF) has been a fundamental assumption in recommender systems, assuming that users who behaved similarly in the past will continue to behave similarly in the future. To accurately estimate user similarities, recommender systems traditionally collect extensive user behavior data. However, in real-world scenarios, users may not want all their data to be used for model training. As exemplified in Figure 1, a male user may occasionally click on items that do not reflect his true preferences,

and a female user may purchase items for others that are not indicative of her own preferences. In such cases, users may prefer that these irrelevant items not be used for training the model. This motivates the need for a recommender system that allows users to actively indicate their willingness regarding which items should be leveraged for training.

To address this challenge, we propose a novel, user-controllable recommendation paradigm. In this paradigm, each user specifies a willingness vector, which indicates their preference for excluding certain items from model training. We formulate the problem as a multiplayer game, where each player represents a user and their action is a selection vector representing a subset of their interacted items. The recommender model is then optimized based on the selected items from different users, striking a balance between recommendation performance and user preferences. However, solving this game requires repeated optimization of the recommender model to compute rewards for different actions, which may not be feasible in practice.

Faced with the aforementioned challenge, we utilize the concept of influence functions to approximate recommendation performance without re-optimizing the model. Specifically, we start by setting an anchor selection vector to train the recommender model. Then, for different actions, we directly compute the recommendation performance using the influence function (Koh and Liang 2017) and the previously obtained model trained with anchor selection vector. Our method is referred to as **influence function based recommendation quality exploration** (called IFRQE for short). To achieve more accurate performance approximation, we extend this approach by employing multiple anchor selection vectors. In addition to these models, we provide a theoretical analysis of the convergence rate of our learning algorithm and demonstrate the advantages of increasing the number of anchor selection vectors. Extensive experiments are conducted using both synthetic and real-world datasets to validate the effectiveness of our models.

It is worth noting that there have been previous works on federated recommendation (Yang et al. 2020), recommendation unlearning (Chen et al. 2022a), and controllable recommendation. However, these studies differ fundamentally from our approach as they neither allow users to actively indicate data nor incorporate user willingness into the optimization target. Furthermore, while (Chen et al. 2022b) pro-

\*Corresponding author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

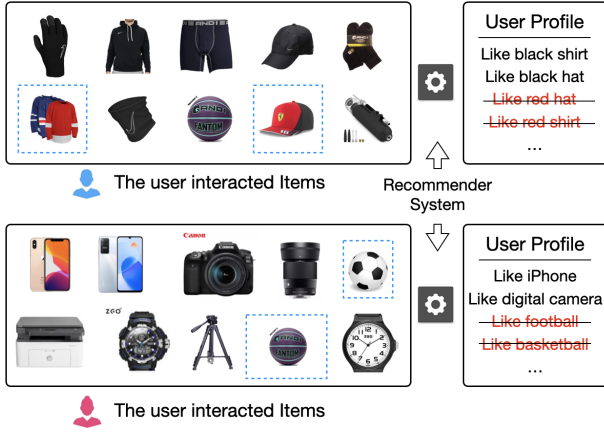


Figure 1: Motivating examples. In the upper case, the user does not want to leverage the occasionally clicked items to train the model. In the bottom case, the items purchased for the other people are expected to be omitted by the model.

posed a model for balancing user privacy and recommendation performance, we improve upon it by introducing more flexible user willingness and utilizing influence functions to enhance optimization efficiency.

The main contributions of this paper are summarized as follows: (1) We introduce a novel recommendation paradigm where users can explicitly indicate their willingness to leverage different items for model training. (2) To address this problem, we formulate the recommendation task as a multiplayer game and propose two influence function-based models for efficient game solving. (3) We provide theoretical analysis by deriving the convergence rate of our learning algorithm and showcasing the superiority of multiple anchor selection vectors. (4) Extensive experiments are conducted on synthetic and real-world datasets to demonstrate the effectiveness of our models.

## Problem Formulation

Suppose we have a user set  $\mathcal{U}$  and an item set  $\mathcal{V}$ . Let  $\mathcal{O}^u$  be the set of items interacted by user  $u$ , and we separate it into a training set  $\mathcal{S}^u$ , a validation set  $\mathcal{T}^u$  and a testing set  $\mathcal{D}^u$ . In our problem, each user  $u$  can specify a willingness vector  $\beta^u = \{\beta_1^u, \beta_2^u, \dots, \beta_{|\mathcal{S}^u|}^u\}$ . Intuitively, for different items, the users may have various willingness on leveraging them to train the model. Thus, we allow the users to indicate a continuous value to describe their willingness flexibly, that is,  $\beta_k^u \in [0, 1], \forall k \in [1, |\mathcal{S}^u|]$ . Suppose  $\mathcal{S}^u = \{s_1^u, s_2^u, \dots, s_{|\mathcal{S}^u|}^u\}$ , then the larger  $\beta_k^u$  is, the more the user do not want  $s_k^u$  to join into the model training process. Our primary task involves selecting an appropriate subset from  $\mathcal{S} = \mathcal{S}^u | u \in \mathcal{U}$  for training the recommender model, ensuring both satisfactory recommendation performance and adherence to user willingness. Increasing the number of selected items can enhance user understanding and improve recommendation performance. However, this approach may violate user willingness, as items with larger  $\beta_k^u$  values may be included. On the other hand, selecting only a few items

to train the model may better satisfy user willingness but potentially lead to reduced recommendation performance. The key challenge lies in effectively balancing recommendation performance and user willingness.

To solve the above task, we regard each user as a player, and formulate the recommendation problem as a multiplayer game. For each user  $u$ , the action is a binary selection vector  $\mathbf{o}^u = \{o_1^u, \dots, o_{|\mathcal{S}^u|}^u\}$ , where  $o_k^u = 1$  means  $s_k^u$  is selected to train the model, otherwise  $o_k^u = 0$ . The reward for user  $u$  is designed as follows:

$$\bar{z}_u(\mathbf{o}^u, \mathbf{o}^{-u}) = -L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o}^u, \mathbf{o}^{-u})) - \lambda \sum_{k=1}^{|\mathcal{S}^u|} o_k^u \beta_k^u, \quad (1)$$

where  $\mathbf{o}^{-u} = \{\mathbf{o}^1, \dots, \mathbf{o}^{u-1}, \mathbf{o}^{u+1}, \dots, \mathbf{o}^N\}$  is the joint selection vectors of all the user except  $u$ .  $f$  is a recommender model.  $\hat{\theta}(\mathbf{o}^u, \mathbf{o}^{-u})$  are the parameters of  $f$  learned based on the training samples selected by  $\{\mathbf{o}^u, \mathbf{o}^{-u}\}$ .  $-L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o}^u, \mathbf{o}^{-u}))$  is the negative validation loss based on  $\hat{\theta}(\mathbf{o}^u, \mathbf{o}^{-u})$ , which is leveraged to measure the recommendation performance. The second term evaluates the violation of the user willingness. If the items that the user want to omit (e.g.,  $\beta_k^u$  is large) are selected to train the model, then  $o_k^u \beta_k^u$  is large, which lowers the reward.  $\lambda$  is a pre-defined balancing parameter.

Instead of learning the specific selection vectors, we formulate our problem in a more general manner, where we assume that there is potential distribution (or called strategy) on the selection vectors for each user. Let the strategy of user  $u$  be  $\alpha^u = [\alpha_{\mathcal{O}^u}^u] \in \Delta(2^{|\mathcal{S}^u|})$ , which is a discrete distribution, and  $\alpha_{\mathcal{O}^u}^u$  is the probability of leveraging the items indicated by  $\mathcal{O}^u$  to train the model. For example, suppose the training set is  $\mathcal{S}^u = \{0, 1, 2\}$ , then  $\alpha^u$  is a distribution on  $\{A_1, A_2, A_3, A_4, A_5, A_6, A_7, A_8\}$ , where  $A_1$  denotes  $\{0, 0, 0\}$ ,  $A_2$  denotes  $\{0, 0, 1\}$ , and so on. Furthermore,  $A_5$  denotes  $\{0, 1, 1\}$  and  $\alpha_{A_5}^u$  represents the probability of using items  $\{1, 2\}$  for model training.

We aim to learn the optimal joint strategy  $\alpha^* = \{\alpha^{1*}, \alpha^{2*}, \dots, \alpha^{N*}\}$ , such that the corresponding expected reward of each user  $u$  is the largest when the other user strategies are fixed, that is:

$$z_u(\alpha^{u*}, \alpha^{-u*}) \geq z_u(\alpha^u, \alpha^{-u*}), \forall u \in [N], \alpha^u \in \Delta(2^{|\mathcal{S}^u|}), \quad (2)$$

where  $[N] = [1, 2, \dots, N]$ .  $\alpha^{-u}$  is the joint strategy of all the users except  $u$ .  $z_u(\alpha^u, \alpha^{-u})$  is the expected reward for user  $u$  under  $\{\alpha^u, \alpha^{-u}\}$ , that is,

$$z_u(\alpha^u, \alpha^{-u}) = E_{\mathbf{o}^u \sim \alpha^u, \mathbf{o}^{-u} \sim \alpha^{-u}}[\bar{z}_u(\mathbf{o}^u, \mathbf{o}^{-u})]. \quad (3)$$

After obtaining the optimal strategy  $\alpha^*$ , we firstly sample the selection vectors  $\mathbf{o} = \{\mathbf{o}^u, \mathbf{o}^{-u}\}$  from  $\alpha^*$ . Then, the training samples are generated by filtering  $\mathcal{S}$  with  $\mathbf{o}$ , which are leveraged to optimize the final recommender model.

*Remark.* (i) In the above formulation, we do not impose any assumption on  $f$  and  $L_f$ , which makes our framework applicable for any recommender model and loss function. (ii) Although the action space in objective (2) may appear large, potentially impacting training efficiency, there are strategies

to mitigate this issue. In our experiments, we initialize  $\alpha$  with an informative prior and focus the search for the optimal solution around this prior to accelerate the training process. Additionally, in practice, one can assign the same value  $o_k^u$  to items in the same category or limit user indications to the most important or recent parts of their interactions.

## The IFRQE Model

In this section, we firstly introduce a basic model for solving the above game based on influence function. Then, we design an improved model to enhance the accuracy for approximating the recommendation performance. At last, we theoretically analyze the learning algorithm of our model. In the following, we describe our models more in detail.

### The Basic Model

To solve the game defined based on (1) and (2), one has to derive the validation loss  $L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o}))$  for different  $\mathbf{o}$ 's. A straightforward method is firstly training  $f$  for each  $\mathbf{o}$  to obtain the corresponding parameter  $\hat{\theta}(\mathbf{o})$ , and then the loss is computed based on the validation set  $\mathcal{T}^u$  and  $\hat{\theta}(\mathbf{o})$ . However, such method is infeasible, since repeatedly training the recommender model is quite time-consuming. Fortunately, the previous studies on influence function (Koh and Liang 2017) may shed some lights on approximating the validation loss without retraining the model. In specific, we first define an anchor selection vector  $\tilde{\mathbf{o}} = \{\tilde{\mathbf{o}}^1, \dots, \tilde{\mathbf{o}}^N\}$ , and then train the recommender model  $f$  based on  $\tilde{\mathbf{o}}$  to obtain the parameters  $\tilde{\theta}$ . At last, for a selection vector  $\mathbf{o}$ , we approximate  $L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o}))$  based on the following theory:

**Theorem 1.** *Given the model parameters  $\theta$ , we define the validation loss by  $L_f(\mathcal{T}^u, \theta) = \sum_{y \in \mathcal{T}^u} l_f(y, \theta)$ , where  $l_f(y, \theta)$  is the loss of sample  $y$  based on  $\theta$ . Suppose  $|\nabla^2 l_f(s_k^v, \tilde{\theta})| \leq B$  and  $\sum_v \sum_{k=1}^{|\mathcal{S}^v|} (\tilde{o}_k^v - o_k^v) B$  is a small value, then the validation loss for  $\mathbf{o}$  is:*

$$L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o})) \approx L_f(\mathcal{T}^u, \tilde{\theta}) - \frac{1}{Z} \sum_{y \in \mathcal{T}^u} \sum_{v \in \mathcal{U}} \sum_{k=1}^{|\mathcal{S}^v|} o_k^v \nabla l_f(y, \tilde{\theta}) H_{\tilde{\theta}}^{-1} \nabla l_f(s_k^v, \tilde{\theta}) \quad (4)$$

where  $\tilde{\theta} = \arg \min_{\theta} \frac{1}{Z} \sum_{v \in \mathcal{U}} \sum_{k=1}^{|\mathcal{S}^v|} \tilde{o}_k^v l_f(s_k^v, \theta)$  are the parameters learned based on the anchor selection vector.  $Z$  is the total number of training samples.  $\nabla l_f(\cdot)$  is the gradient of a sample loss.  $H_{\tilde{\theta}} = \frac{1}{Z} \sum_v \sum_{k=1}^{|\mathcal{S}^v|} \tilde{o}_k^v \nabla^2 l_f(s_k^v, \tilde{\theta})$  is the Hessian matrix of the training loss.

The proof of this theorem is presented in the Appendix<sup>1</sup>. Based on this theory, we can compute  $L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o}))$  without retraining the model via  $\hat{\theta}(\mathbf{o}) = \arg \min_{\theta} \frac{1}{Z} \sum_v \sum_{k=1}^{|\mathcal{S}^v|} o_k^{t,v} l_f(s_k^v, \theta)$ . By bring-

<sup>1</sup>Refer to Appendix: Proof of Theorem 1. Our Appendix is available at <https://github.com/Paitesanshi/IFRQE>.

ing (4) into (1),  $\bar{z}_u(\mathbf{o}^u, \mathbf{o}^{-u})$  can be written as:

$$\bar{z}_u(\mathbf{o}^u, \mathbf{o}^{-u}) = -L_f(\mathcal{T}^u, \tilde{\theta}) + \frac{1}{Z} \sum_{y \in \mathcal{T}^u} \sum_{v \in \mathcal{U}} \sum_{k=1}^{|\mathcal{S}^v|} o_k^v \nabla l_f(y, \tilde{\theta}) H_{\tilde{\theta}}^{-1} \nabla l_f(s_k^v, \tilde{\theta}) - \lambda \sum_{k=1}^{|\mathcal{S}^u|} o_k^u \beta_k^u. \quad (5)$$

Accordingly, the expected reward  $z_u(\alpha^u, \alpha^{-u})$  is:

**Theorem 2.** *Let  $\mathbf{g}_y^v = [g(s_1^v, y), \dots, g(s_{|\mathcal{S}^v|}^v, y)]$ , where  $g(s_k^v, y) = \nabla_{\theta} l_f(y, \tilde{\theta})^T H_{\tilde{\theta}}^{-1} \nabla_{\theta} l_f(s_k^v, \tilde{\theta})$ , then:*

$$z_u(\alpha^u, \alpha^{-u}) = \sum_{\mathbf{o}^u} \alpha_{\mathbf{o}^u}^u \left[ \left( \sum_{y \in \mathcal{T}^u} \frac{\mathbf{g}_y^u}{N} - \lambda \beta^u \right)^T \mathbf{o}^u \right] + C, \quad (6)$$

where  $C = -L(\mathcal{T}^u, \tilde{\theta}) + \sum_{v \neq u} E_{\mathbf{o}}[(\mathbf{o}^v)^T \sum_{y \in \mathcal{T}^u} \frac{\mathbf{g}_y^v}{Z}]$ .

The deriving process of the above theorem is presented in the Appendix. To solve objective (2), we need to find an optimal  $\alpha^u$  for the following optimization problem:

$$\max_{\alpha^u} \sum_{\mathbf{o}^u} \alpha_{\mathbf{o}^u}^u A(\mathbf{o}^u), \text{ s.t. } \sum_{\mathbf{o}^u} \alpha_{\mathbf{o}^u}^u = 1, \alpha_{\mathbf{o}^u}^u \geq 0 \quad (7)$$

where we denote  $(\sum_{y \in \mathcal{T}^u} \frac{\mathbf{g}_y^u}{N} - \lambda \beta^u)^T \mathbf{o}^u$  by  $A(\mathbf{o}^u)$ .

### The Improved Model

The key of the above method lies in the accurate approximation of  $L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o}))$ . However, if the current selection vector  $\mathbf{o}$  is too much different from the anchor vector  $\tilde{\mathbf{o}}$ , then the assumption  $\sum_v \sum_{k=1}^{|\mathcal{S}^v|} (\tilde{o}_k^v - o_k^v) B$  is a small value in theory 1 may not hold, which can lead to larger approximation errors. To alleviate this problem, we propose to set multiple anchor vectors  $\{\tilde{\mathbf{o}}^1, \tilde{\mathbf{o}}^2, \dots, \tilde{\mathbf{o}}^T\}$ , where  $\tilde{\mathbf{o}}^t = \{\tilde{\mathbf{o}}^{t,1}, \dots, \tilde{\mathbf{o}}^{t,N}\}$  is the  $t$ th anchor vector and  $\tilde{o}_k^{t,v}$  is the  $k$ th element of  $\tilde{\mathbf{o}}^{t,v}$ . For approximating  $L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o}))$ , we select the anchor vector nearest to  $\mathbf{o}$ , where we have the following theory.

**Theorem 3.** *For a candidate selection vector  $\mathbf{o}$ , suppose  $t = \arg \min_{i \in [1, T]} \sum_{v=1}^N D(\tilde{\mathbf{o}}^{i,v}, \mathbf{o}^v)$ , where  $D$  is the hamming distance between two vectors, and  $\tilde{\theta}^t = \arg \min_{\theta} \frac{1}{Z} \sum_v \sum_{k=1}^{|\mathcal{S}^v|} \tilde{o}_k^{t,v} l_f(s_k^v, \theta)$ , then:*

$$L_f(\mathcal{T}^u, \hat{\theta}(\mathbf{o})) \approx L_f(\mathcal{T}^u, \tilde{\theta}^t) - \frac{1}{Z} \sum_{y \in \mathcal{T}^u} \sum_{v \in \mathcal{U}} \sum_{k=1}^{|\mathcal{S}^v|} o_k^v \nabla l_f(y, \tilde{\theta}^t) H_{\tilde{\theta}^t}^{-1} \nabla l_f(s_k^v, \tilde{\theta}^t), \quad (8)$$

where  $H_{\tilde{\theta}^t} = \frac{1}{Z} \sum_v \sum_{k=1}^{|\mathcal{S}^v|} \tilde{o}_k^{t,v} \nabla^2 l_f(s_k^v, \tilde{\theta}^t)$ .

Based on (8), the expected reward can be derived based on the following theory.

**Theorem 4.** *Let  $A_t = \{\mathbf{o} \mid \sum_{v=1}^N D(\tilde{\mathbf{o}}^{t,v}, \mathbf{o}^v) \leq \sum_{v=1}^N D(\tilde{\mathbf{o}}^{t',v}, \mathbf{o}^v), \forall t' \neq t\}$ ,  $\mathbf{g}_y^{t,v} = [g(s_k^v, y, t)]_{k=1}^{|\mathcal{S}^v|}$ , where  $g(s_k^v, y, t) = \nabla l_f(y, \tilde{\theta}^t) H_{\tilde{\theta}^t}^{-1} \nabla l_f(s_k^v, \tilde{\theta}^t)$  and  $\mathbf{g}^{t,v} = \sum_{y \in \mathcal{T}^u} \mathbf{g}_y^{t,v}$ . We define  $\bar{z}_u(\mathbf{o}^u, \mathbf{o}^{-u}, t) = -L_f(\mathcal{T}^u, \tilde{\theta}^t) +$*

---

**Algorithm 1:** Learning algorithm for  $\alpha_m^u$ 


---

- 1 Indicate the learning rate  $\gamma$ .
  - 2 Let  $\alpha^{-u} = \alpha_{m-1}^{-u}$  and  $\alpha_1^u = \alpha_{m-1}^u$ .
  - 3 **for**  $l$  in  $[1, L]$  **do**
  - 4     Sample  $\mathbf{o}^u, \mathbf{o}^{-u}$  according to  $\alpha_l^u, \alpha^{-u}$ .
  - 5     Compute the gradient  $\hat{g}_{\mathbf{o}^u}$  based on (11).
  - 6     Let  $[\hat{g}]_s = 1(s = \mathbf{o}^u)\hat{g}_{\mathbf{o}^u}$ .
  - 7     Update  $\alpha_{l+1}^u = \Pi_{\Delta}[\alpha_l^u + \gamma\hat{g}]$ , where  $\Pi_{\Delta}$  means projecting a vector into a simplex.
  - 8 **end**
  - 9 Return  $\alpha^u = \frac{1}{L} \sum_{l=1}^L \alpha_l^u$ .
- 

$\frac{1}{Z} \sum_{y \in \mathcal{T}^u} \sum_{v \in \mathcal{U}} \sum_{k=1}^{|\mathcal{S}^v|} o_k^v \nabla l_f(y, \tilde{\theta}^t) H_{\tilde{\theta}^t}^{-1} \nabla l_f(s_k^v, \tilde{\theta}^t) - \lambda \sum_{k=1}^{|\mathcal{S}^u|} o_k^u \beta_k^u$ , Then, we have:

$$\begin{aligned} z_u(\alpha^u, \alpha^{-u}) &= E_{\mathbf{o}} \left[ \sum_{t=1}^T 1(\mathbf{o} \in A_t) \bar{z}_u(\mathbf{o}^u, \mathbf{o}^{-u}, t) \right] \\ &= \sum_{t=1}^T \sum_{\mathbf{o}} 1(\mathbf{o} \in A_t) \alpha_{\mathbf{o}^u}^u \alpha_{\mathbf{o}^{-u}}^{-u} \{-L(\mathcal{T}^u, \tilde{\theta}^t) + \frac{1}{Z} \sum_{v \neq u} (\mathbf{o}^v)^T \mathbf{g}^{t,v} \\ &\quad + \frac{1}{Z} (\mathbf{o}^u)^T \mathbf{g}^{t,u} - \lambda (\mathbf{o}^u)^T \beta^u\}, \end{aligned} \quad (9)$$

where  $1(c) = 1$  if the condition  $c$  is true, otherwise  $1(c) = 0$ . And  $\alpha_{\mathbf{o}^{-u}}^{-u} = \alpha_{\mathbf{o}^1}^1 \dots \alpha_{\mathbf{o}^{u-1}}^{u-1} \alpha_{\mathbf{o}^{u+1}}^{u+1} \dots \alpha_{\mathbf{o}^N}^N$ .

The proof of the above theorem is presented in the Appendix.

Finally, we derive  $\alpha^u$  by solving the following problem:

$$\begin{aligned} \max_{\alpha^u} \sum_{t=1}^T \sum_{\mathbf{o}} 1(\mathbf{o} \in A_t) \alpha_{\mathbf{o}^u}^u \alpha_{\mathbf{o}^{-u}}^{-u} B(\mathbf{o}^u, \mathbf{o}^{-u}, t), \\ \text{s.t. } \sum_{\mathbf{o}^u} \alpha^u(\mathbf{o}^u) = 1, \quad \alpha^u(\mathbf{o}^u) > 0 \end{aligned} \quad (10)$$

where the notation  $B(\mathbf{o}^u, \mathbf{o}^{-u}, t) = -L(\mathcal{T}^u, \tilde{\theta}^t) + \frac{1}{Z} \sum_{v \neq u} (\mathbf{o}^v)^T \mathbf{g}^{t,v} + \frac{1}{Z} (\mathbf{o}^u)^T \mathbf{g}^{t,u} - \lambda (\mathbf{o}^u)^T \beta^u$ .

Since it is hard to efficiently obtain a closed-form solution for  $\alpha^u$ , we learn it based on the projected gradient descent method (Calamai and Moré 1987). More specifically, the gradient of  $z_u(\alpha^u, \alpha^{-u})$  w.r.t  $\alpha_{\mathbf{o}^u}^u$  is:

$$\begin{aligned} \frac{\partial z_u(\alpha^u, \alpha^{-u})}{\partial \alpha_{\mathbf{o}^u}^u} &\stackrel{\text{def}}{=} g_{\mathbf{o}^u} \\ &= E_{\mathbf{o}^{-u}} \left[ \sum_{t=1}^T B(\mathbf{o}^u, \mathbf{o}^{-u}, t) \mathbf{1}(\mathbf{o}^{-u} \in S_t(\mathbf{o}^u)) \right], \end{aligned} \quad (11)$$

where  $S_t(\mathbf{o}^u) = \{\mathbf{o}^{-u} | [\mathbf{o}^u, \mathbf{o}^{-u}] \in A_t\}$ . Let  $\mathbf{g} = [g_{\mathbf{o}^u}]_{\mathbf{o}^u}$  be the gradient of  $z_u(\alpha^u, \alpha^{-u})$  w.r.t  $\alpha^u$ , which is a  $2^{|\mathcal{S}^u|}$  dimensional vector.  $\hat{g}_{\mathbf{o}^u}$  is the stochastic gradient by sampling  $\mathbf{o}^{-u}$  from  $\alpha^{-u}$ . Then the complete training process can be seen in Algorithm 2.

---

**Algorithm 2:** Learning Algorithm with Multiple  $\theta_t$ 


---

- 1 Initialize  $\{\alpha^1, \alpha^2, \dots, \alpha^N\}$  and let  $\alpha_0^u = \alpha^u$  ( $u \in [1, N]$ ).
  - 2 Indicate the max iteration number  $M$  and threshold  $\kappa$ .
  - 3 Indicate the anchor vectors  $\{\tilde{\theta}^1, \tilde{\theta}^2, \dots, \tilde{\theta}^T\}$ .
  - 4 Obtain the model parameters  $\tilde{\theta}^t$  for each  $\tilde{\theta}^t$ .
  - 5 **for**  $m$  in  $[1, M]$  **do**
  - 6     **for**  $u$  in  $[1, N]$  **do**
  - 7         Let  $\alpha_{m-1}^{-u} = \{\alpha_{m-1}^1, \dots, \alpha_{m-1}^{u-1}, \alpha_{m-1}^{u+1}, \dots, \alpha_{m-1}^N\}$ .
  - 8         Obtain  $\alpha_m^u$  by inputting  $\alpha_{m-1}^{-u}$  and  $\alpha_{m-1}^u$  into Algorithm 1.
  - 9     **end**
  - 10     **if**  $|\alpha_m^u - \alpha_{m-1}^u| < \kappa, \forall u \in [1, N]$  **then**
  - 11         Break.
  - 12     **end**
  - 13 **end**
  - 14 Output  $\alpha^{u*} = \alpha_m^u$  ( $u \in [1, N]$ ).
- 

**Theoretical Analysis**

In this section, we provide theoretical analysis on the convergence of our algorithm and also present the advantage of multiple anchor selection vectors comparing with the single one in terms of the validation loss approximation. For the convergence analysis, we have the following theory:

**Theorem 5.** Based on the definition of  $\mathbf{g}$ , we have  $z_u(\alpha^u, \alpha^{-u}) = (\alpha^u)^T \mathbf{g}$ . Suppose  $\hat{\mathbf{g}}$  is an unbiased estimation of  $\mathbf{g}$ , and  $\|\hat{\mathbf{g}}(\alpha^u)\|_2^2 \leq G$ , then the solution obtained from Algorithm 1 is larger than the optimal solution minus a bounded value, that is:

$$E[z_u(\hat{\alpha}^u, \alpha^{-u})] \geq \max_{\alpha^u} E[z_u(\alpha^u, \alpha^{-u})] - \left( \frac{1}{L\gamma} + \gamma^2 G^2 \right), \quad (12)$$

where  $\hat{\alpha}^u$  is obtained based on Algorithm 1.

We prove this theory in Appendix. It provides foundations for Algorithm 1, which demonstrates that the learned strategies are nearly optimal. For the advantage of using multiple anchor selection vectors, we have the following theory:

**Theorem 6.** For two sets of anchor selection vectors  $\mathbf{P} = \{\tilde{\theta}_P^1, \tilde{\theta}_P^2, \dots, \tilde{\theta}_P^{T_P}\}$  and  $\mathbf{Q} = \{\tilde{\theta}_Q^1, \tilde{\theta}_Q^2, \dots, \tilde{\theta}_Q^{T_Q}\}$ , suppose  $A_t^P = \{\mathbf{o} | \sum_{v=1}^N D(\tilde{\theta}^{t,v}, \mathbf{o}^v) \leq \sum_{v=1}^N D(\tilde{\theta}^{t',v}, \mathbf{o}^v), \forall \tilde{\theta}^{t'} \in \mathbf{P}\}$ ,  $A_t^Q = \{\mathbf{o} | \sum_{v=1}^N D(\tilde{\theta}^{t,v}, \mathbf{o}^v) \leq \sum_{v=1}^N D(\tilde{\theta}^{t',v}, \mathbf{o}^v), \forall \tilde{\theta}^{t'} \in \mathbf{Q}\}$ . Based on theory 1, we consider the following upper bounds of the approximation error for the validation loss:

$$\begin{aligned} \text{err}(\mathbf{P}) &= \sum_{t=1}^{T_P} \sum_{\mathbf{o} \in A_t^P} \left[ \sum_v D(\tilde{\theta}^{t,v}, \mathbf{o}^v) \right] B, \\ \text{err}(\mathbf{Q}) &= \sum_{t=1}^{T_Q} \sum_{\mathbf{o} \in A_t^Q} \left[ \sum_v D(\tilde{\theta}^{t,v}, \mathbf{o}^v) \right] B. \end{aligned} \quad (13)$$

We have if  $\mathbf{P} \subseteq \mathbf{Q}$ , then  $\text{err}(\mathbf{P}) \geq \text{err}(\mathbf{Q})$ .

The proof of this theorem is provided in the Appendix, indicating that including the single anchor vector in the multiple anchor vectors reduces the upper bound of the approximation error for the validation loss. However, it is important to note that introducing more anchor vectors inevitably requires additional retraining of the recommender model, consequently resulting in a decrease in efficiency. Nevertheless, the improved model allows for a balance between approximation accuracy and model efficiency.

## Related Work

Recommender system is a rapidly developing research field, with a large amount of models developed each year (Ricci, Rokach, and Shapira 2015; Wang et al. 2023). Early recommender models primarily relied on matrix factorization techniques (Koren, Bell, and Volinsky 2009). However, with the advancements in deep learning, numerous deep recommender models have emerged. These models encompass various approaches, such as sequential (Kang and McAuley 2018; Chen et al. 2018; Wang et al. 2021) and graph-based algorithms (Wang et al. 2019; He et al. 2020). In an effort to protect user privacy, federated recommender models have been proposed (Lin et al. 2020; Wu et al. 2021; Yang et al. 2020; Chen et al. 2020). These models aim to shift the training process from the server to the clients to prevent access to raw user behaviors and ensure privacy. While our model can also contribute to privacy protection, it does not assume an adversarial server. Instead, our focus is on preventing the leakage of privacy information through recommendation unlearning, which seeks to eliminate the influences of user-item interactions from trained recommender models (Chen et al. 2022a; Li et al. 2022). However, these models do not account for user active disclosure behaviors or incorporate user willingness into optimization targets. Controllable recommendation has also received attention in recent studies (Wang et al. 2022; Parra and Brusilovsky 2015). However, these approaches primarily aim to break filter bubbles rather than aligning with user-indicated willingness. Game theory-based recommender models have also been explored (Ben-Porat and Tennenholtz 2018; Xu et al. 2018; Halkidi and Koutsopoulos 2011). Although our work shares some similarities with these models, we distinguish ourselves by constructing the game from the user’s perspective, not limiting our framework to rating behaviors, and leveraging influence functions to enhance recommendation performance estimation, a novel approach in this context.

## Experiments

In this section, we conduct comprehensive experiments to demonstrate the effectiveness of our models, where we focus on the following research questions: **RQ1**: Whether our framework can achieve better trade-off between the recommendation performance and user willingness? **RQ2**: Whether the validation loss can be properly approximated by our framework? **RQ3**: What is the influence of the number of anchor selection vectors? **RQ4**: How does the dataset sparsity influence the performance of our framework? **RQ5**:

Whether our framework is effective for different weighting parameter  $\lambda$ ’s? **RQ6**: Whether our framework is effective for the user willingness with different characters? **RQ7**: How does our framework perform if the recommender model is trained from scratch without using the influence function? In the following, we first introduce the experiment setup, followed by a presentation and analysis of the empirical results to address the above questions.

## Experiment Setup

**Datasets.** We base our experiments on both synthetic and real-world datasets. For the synthetic dataset, we generated it based on a well known recommendation simulator called RecSim<sup>2</sup>. In specific, we generate 1000 users and 1000 items. For each user or item, we simulate its profiles by four random features, which is denoted by  $e_u \in \mathbb{R}^4$  or  $e_v \in \mathbb{R}^4$ . The dataset is generated according to the observational parameter  $\tau_{uv} = \mathbf{1}(\sigma(e_u^T e_v) \geq \eta)$ , where  $\mathbf{1}(\cdot)$  is the indicator function and  $\eta$  is a threshold for setting the hardness of generating the interaction. If  $\tau_{uv} = 1$ , then the user-item pair  $(u, v)$  is observed in the dataset. For the willingness of user  $u$  on item  $v$ , we simulate it by:  $\beta_v^u = \sigma(s_u g(e_u, e_v) + b_u)$ , where  $s_u$  is sampled from a uniform distribution in the range of  $[a_1, a_2]$ , controlling the sensitivity of the user willingness. Smaller  $s_u$  means that the user has similar willingness on leveraging different items to train the model, while larger  $s_u$  means that the user willingness on different items are more diverse.  $b_u$  is sampled from a Gaussian distribution  $\mathcal{N}(0, a_3)$ , indicating the average level of the user willingness.  $g$  is a one-layer fully connected neural network with randomly assigned parameters. In the experiments, the threshold  $\eta$  and  $(a_1, a_2, a_3)$  are initially set as 0.5 and  $(0.5, 1, 1)$ , respectively. Then, we tune them to study the influence of different dataset sparsities and user willingness characters on our framework performance. For the real-world experiments, we evaluate different models based on three publicly available datasets from different domains, including **Diginetica**<sup>3</sup>, **Steam**<sup>4</sup> and **Amazon Video**<sup>5</sup>. Diginetica and Amazon Video are e-commerce datasets, where we are provided with the user-item purchasing records. Steam is a game dataset, which includes the interactions between the users and games. The statistics of the above datasets are concluded in Appendix.

**Baselines.** To demonstrate the generality of our idea, we implement the base model  $f$  with the following recommender algorithms<sup>6</sup>: **MF** (Koren, Bell, and Volinsky 2009) is the traditional matrix factorization method. **NeuMF** (He et al. 2017) is a well known neural collaborative filtering model. **LightGCN** (He et al. 2020) is a graph-based recommender model. We compare our framework to various methods. Firstly, **Random** uses user behaviors randomly in training (*i.e.*, selection vector  $\mathbf{o}$  is randomly sampled). Sec-

<sup>2</sup><https://github.com/google-research/recsim>

<sup>3</sup>[https://darel13712.github.io/rs\\_datasets/Datasets/diginetica/](https://darel13712.github.io/rs_datasets/Datasets/diginetica/)

<sup>4</sup><https://steam.internet.byu.edu/>

<sup>5</sup><https://jmcauley.ucsd.edu/data/amazon/>

<sup>6</sup>Experiments with more other base models can be found in the Appendix.

Dataset	Simulation			Diginetica			Steam			Amazon Video		
Metric	$F_1 \uparrow$	$wv \downarrow$	reward $\uparrow$	$F_1 \uparrow$	$wv \downarrow$	reward $\uparrow$	$F_1 \uparrow$	$wv \downarrow$	reward $\uparrow$	$F_1 \uparrow$	$wv \downarrow$	reward $\uparrow$
MF	<b>1.75</b> <sub>(.022)</sub>	2.12 <sub>(.007)</sub>	-2.25 <sub>(.032)</sub>	4.72 <sub>(.044)</sub>	2.01 <sub>(.053)</sub>	-2.65 <sub>(.072)</sub>	<b>10.4</b> <sub>(.043)</sub>	2.62 <sub>(.014)</sub>	-2.99 <sub>(.093)</sub>	1.70 <sub>(.012)</sub>	2.36 <sub>(.021)</sub>	-2.43 <sub>(.022)</sub>
w/ Random	1.73 <sub>(.031)</sub>	2.01 <sub>(.006)</sub>	-2.13 <sub>(.014)</sub>	4.18 <sub>(.014)</sub>	1.96 <sub>(.013)</sub>	-2.04 <sub>(.006)</sub>	9.40 <sub>(.009)</sub>	2.61 <sub>(.017)</sub>	-2.82 <sub>(.015)</sub>	1.53 <sub>(.011)</sub>	2.24 <sub>(.022)</sub>	-2.30 <sub>(.023)</sub>
w/ Threshold	1.70 <sub>(.036)</sub>	2.01 <sub>(.032)</sub>	-2.20 <sub>(.021)</sub>	4.70 <sub>(.012)</sub>	1.97 <sub>(.011)</sub>	-2.06 <sub>(.040)</sub>	9.25 <sub>(.048)</sub>	2.35 <sub>(.023)</sub>	-2.57 <sub>(.017)</sub>	1.50 <sub>(.026)</sub>	2.12 <sub>(.027)</sub>	-2.20 <sub>(.021)</sub>
w/ Proactive	1.23 <sub>(.025)</sub>	<b>1.56</b> <sub>(.013)</sub>	-2.25 <sub>(.007)</sub>	2.21 <sub>(.018)</sub>	<b>1.30</b> <sub>(.024)</sub>	-1.99 <sub>(.016)</sub>	9.93 <sub>(.019)</sub>	<b>2.08</b> <sub>(.022)</sub>	-2.43 <sub>(.025)</sub>	1.55 <sub>(.035)</sub>	1.94 <sub>(.026)</sub>	-2.07 <sub>(.005)</sub>
w/ IFRQE	1.50 <sub>(.015)</sub>	1.97 <sub>(.003)</sub>	-2.09 <sub>(.017)</sub>	<b>5.23</b> <sub>(.008)</sub>	2.01 <sub>(.014)</sub>	-2.18 <sub>(.006)</sub>	<b>10.4</b> <sub>(.012)</sub>	2.14 <sub>(.012)</sub>	-2.43 <sub>(.015)</sub>	1.57 <sub>(.015)</sub>	<b>1.70</b> <sub>(.016)</sub>	-2.10 <sub>(.005)</sub>
w/ IFRQE++	1.63 <sub>(.007)</sub>	1.82 <sub>(.003)</sub>	<b>-1.92</b> <sub>(.012)</sub>	4.02 <sub>(.011)</sub>	1.79 <sub>(.019)</sub>	<b>-1.92</b> <sub>(.022)</sub>	9.83 <sub>(.017)</sub>	2.13 <sub>(.032)</sub>	<b>-2.42</b> <sub>(.014)</sub>	<b>1.72</b> <sub>(.012)</sub>	1.87 <sub>(.032)</sub>	<b>-1.98</b> <sub>(.056)</sub>
NeuMF	0.87 <sub>(.017)</sub>	2.13 <sub>(.013)</sub>	-2.32 <sub>(.011)</sub>	<b>4.22</b> <sub>(.013)</sub>	2.01 <sub>(.015)</sub>	-2.11 <sub>(.003)</sub>	<b>10.2</b> <sub>(.007)</sub>	2.62 <sub>(.012)</sub>	-2.76 <sub>(.007)</sub>	1.65 <sub>(.014)</sub>	2.36 <sub>(.013)</sub>	-2.47 <sub>(.010)</sub>
w/ Random	0.89 <sub>(.009)</sub>	2.02 <sub>(.005)</sub>	-2.19 <sub>(.011)</sub>	2.38 <sub>(.018)</sub>	1.91 <sub>(.011)</sub>	-2.11 <sub>(.016)</sub>	7.73 <sub>(.005)</sub>	2.47 <sub>(.019)</sub>	-2.59 <sub>(.011)</sub>	1.80 <sub>(.014)</sub>	2.31 <sub>(.006)</sub>	-2.41 <sub>(.010)</sub>
w/ Threshold	1.50 <sub>(.032)</sub>	2.01 <sub>(.048)</sub>	-2.18 <sub>(.031)</sub>	4.01 <sub>(.012)</sub>	1.97 <sub>(.044)</sub>	-2.13 <sub>(.006)</sub>	7.50 <sub>(.019)</sub>	2.35 <sub>(.020)</sub>	-2.86 <sub>(.010)</sub>	1.55 <sub>(.002)</sub>	2.12 <sub>(.018)</sub>	-2.21 <sub>(.017)</sub>
w/ Proactive	<b>1.83</b> <sub>(.005)</sub>	2.01 <sub>(.023)</sub>	-2.18 <sub>(.017)</sub>	1.30 <sub>(.008)</sub>	<b>1.69</b> <sub>(.011)</sub>	-2.14 <sub>(.012)</sub>	9.55 <sub>(.025)</sub>	2.22 <sub>(.032)</sub>	-2.42 <sub>(.005)</sub>	1.43 <sub>(.015)</sub>	<b>1.96</b> <sub>(.006)</sub>	-2.41 <sub>(.015)</sub>
w/ IFRQE	1.40 <sub>(.014)</sub>	1.91 <sub>(.013)</sub>	-2.48 <sub>(.013)</sub>	2.92 <sub>(.137)</sub>	2.01 <sub>(.017)</sub>	-2.17 <sub>(.019)</sub>	8.17 <sub>(.033)</sub>	<b>2.16</b> <sub>(.019)</sub>	-2.45 <sub>(.027)</sub>	<b>1.83</b> <sub>(.015)</sub>	1.99 <sub>(.014)</sub>	-2.24 <sub>(.012)</sub>
w/ IFRQE++	0.70 <sub>(.015)</sub>	<b>1.80</b> <sub>(.003)</sub>	<b>-2.11</b> <sub>(.011)</sub>	3.00 <sub>(.017)</sub>	1.91 <sub>(.013)</sub>	<b>-2.08</b> <sub>(.023)</sub>	8.87 <sub>(.021)</sub>	2.21 <sub>(.014)</sub>	<b>-2.37</b> <sub>(.027)</sub>	1.22 <sub>(.031)</sub>	2.00 <sub>(.014)</sub>	<b>-2.17</b> <sub>(.011)</sub>
LightGCN	0.90 <sub>(.005)</sub>	2.13 <sub>(.008)</sub>	-2.82 <sub>(.012)</sub>	<b>7.50</b> <sub>(.009)</sub>	2.01 <sub>(.012)</sub>	-2.71 <sub>(.009)</sub>	10.1 <sub>(.010)</sub>	2.62 <sub>(.009)</sub>	-3.13 <sub>(.014)</sub>	<b>2.15</b> <sub>(.021)</sub>	2.36 <sub>(.008)</sub>	-3.05 <sub>(.007)</sub>
w/ Random	0.82 <sub>(.013)</sub>	2.02 <sub>(.013)</sub>	-2.71 <sub>(.022)</sub>	7.48 <sub>(.011)</sub>	1.95 <sub>(.005)</sub>	-2.64 <sub>(.007)</sub>	10.0 <sub>(.027)</sub>	2.49 <sub>(.026)</sub>	-3.07 <sub>(.019)</sub>	1.75 <sub>(.003)</sub>	2.24 <sub>(.009)</sub>	-2.93 <sub>(.015)</sub>
w/ Threshold	0.96 <sub>(.038)</sub>	2.01 <sub>(.008)</sub>	-2.70 <sub>(.003)</sub>	7.13 <sub>(.017)</sub>	1.97 <sub>(.040)</sub>	-2.60 <sub>(.011)</sub>	9.40 <sub>(.048)</sub>	2.35 <sub>(.027)</sub>	-2.82 <sub>(.046)</sub>	1.53 <sub>(.033)</sub>	2.12 <sub>(.022)</sub>	-2.30 <sub>(.048)</sub>
w/ Proactive	<b>1.01</b> <sub>(.013)</sub>	1.78 <sub>(.011)</sub>	-2.48 <sub>(.017)</sub>	5.51 <sub>(.038)</sub>	<b>1.42</b> <sub>(.007)</sub>	-2.11 <sub>(.018)</sub>	<b>10.3</b> <sub>(.021)</sub>	<b>2.12</b> <sub>(.022)</sub>	-2.81 <sub>(.015)</sub>	1.28 <sub>(.013)</sub>	1.95 <sub>(.016)</sub>	-2.64 <sub>(.018)</sub>
w/ IFRQE	0.97 <sub>(.008)</sub>	2.12 <sub>(.010)</sub>	-2.82 <sub>(.014)</sub>	4.88 <sub>(.019)</sub>	1.47 <sub>(.011)</sub>	-2.16 <sub>(.013)</sub>	9.02 <sub>(.016)</sub>	2.49 <sub>(.025)</sub>	-2.82 <sub>(.008)</sub>	1.90 <sub>(.007)</sub>	<b>1.75</b> <sub>(.018)</sub>	-2.13 <sub>(.013)</sub>
w/ IFRQE++	0.75 <sub>(.022)</sub>	<b>1.75</b> <sub>(.008)</sub>	<b>-2.44</b> <sub>(.016)</sub>	5.26 <sub>(.012)</sub>	1.32 <sub>(.014)</sub>	<b>-2.01</b> <sub>(.009)</sub>	8.90 <sub>(.018)</sub>	2.48 <sub>(.018)</sub>	<b>-2.80</b> <sub>(.005)</sub>	1.62 <sub>(.013)</sub>	<b>1.75</b> <sub>(.009)</sub>	<b>-2.03</b> <sub>(.023)</sub>

Table 1: Overall comparison between different models. We use bold fonts to label the best performance for each dataset, evaluation metric and base model. “( )” indicates the standard error. The results of  $F_1$  are percentage values with “%” omitted. For the metrics,  $\uparrow$  means the larger the better, while  $\downarrow$  means the lower the better. The performance improvements of our model against the baselines are significant under paired t-test.

only, **Threshold** sets  $o_u^k = 0$  if  $\beta_k^u > 0.5$ . We also compare with **Proactive**(Chen et al. 2022b), a model designed to consider both recommendation accuracy and user privacy. For the sake of fairness, we use the same parameter  $\lambda$  to balance these two aspects in objective. Our basic and improved methods are named **IFRQE** and **IFRQE++**, respectively.

**Implementation Details.** For each user, we use 20% and 10% of interactions for testing and validation, respectively, with the rest for training. All models employ binary cross entropy as the loss function. Recommendation performance is evaluated using  $F_1$  (Chicco and Jurman 2020) with five items recommended for comparison to the ground truth. To assess adherence to user willingness, we calculate overall willingness violation,  $wv = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} \sum_{k=1}^{|\mathcal{S}^u|} o_k^u \beta_k^u$ , lower  $wv$  indicating higher satisfaction. Model performance in balancing recommendation quality and user willingness is compared via **reward** (eq. 1). Our framework and baselines are optimized for the same reward, testing different  $\lambda$ ’s. Anchor selection vectors in IFRQE and IFRQE++ are randomly selected<sup>7</sup>. Model hyperparameters are set via grid search. For instance, the number of anchor selection vectors is searched in [1, 2, ..., 10], learning rate and batch size in [0.0001, 0.001, 0.005, 0.01] and [1024, 2048, 4096], respectively. Anchor selection vectors are binomially distributed, with a mean of 0.9. For real-world datasets, we simulate user willingness with random assignments. To control for randomness, we repeat experiments ten times, reporting average and standard error. More implementation details<sup>8</sup> and

complete experiment results can be found in the Appendix.

## Overall Comparison

The overall comparison results are presented in Table 1. From the user willingness perspective, our model consistently outperforms the base model and heuristic methods, indicating a finer consideration of user preferences. This improvement is expected, as the base model utilizes all available items for training, while the heuristic methods adopt a coarser approach to user willingness. In terms of recommendation performance, our models exhibit only a moderate decrease compared to the base model, with an average drop of approximately 8.02%. Notably, in certain cases, such as with the Diginetica and Amazon Video datasets using MF as the base model, our approach even demonstrates improved performance. This observation suggests that the removal of items with higher  $\beta_k^u$  values, which may not be essential for predicting items in the testing set, leads to higher rewards by aligning with user willingness and enhances performance by focusing on more informative training samples. From the reward perspective, our framework consistently achieves the best performances across most datasets and base models, illustrating the effectiveness and generalizability of our approach in striking a balance between recommendation quality and user willingness. On average, our IFRQE++ method improves the base model by approximately 12.4%, 11.1%, 14.6%, and 21.3% on the simulation, Diginetica, Steam, and Amazon Video datasets, respectively. The superior performance of our framework compared to the random method underscores the non-trivial nature of the problem and the limitations of random selection. When compared to Proactive, our framework demonstrates comparable performance

<sup>7</sup>IFRQE++’s anchor selection vectors include those of IFRQE

<sup>8</sup>MindSpore implementation at <https://github.com/mindspore-lab/models/tree/master/research/huawei-noah/IFRQE>.






Figure 2: (a) Approximation error on the validation loss. (b) Influence of  $T$ .




Figure 3: Model comparisons on the datasets with differently characterized user willingness.

improvements, while significantly reducing the training time required. Specifically, our framework achieves similar results to Proactive in a fraction of the time, suggesting its practicality and feasibility in real-world scenarios. Furthermore, in the comparison between IFRQE and IFRQE++, the latter consistently outperforms the former in a variety of cases. We hypothesize that the incorporation of multiple anchor vectors enables a more accurate approximation of the validation loss, facilitating the learning of optimal strategies and resulting in better rewards.

### Approximation of the Validation Loss

The key of our framework for speeding up the training efficiency lies in the approximation of the validation loss. In this section, we study how large is approximation error induced by our framework. We present the performance of MF on datasets Diginetica and Amazon Video, with the complete results provided in the Appendix. We firstly generate ten selection vectors and train the base model to obtain the parameters and the corresponding real validation loss. Then, we approximate the validation loss by equation (4) and (8), where we use two anchor vectors in **IFRQE++**. At last, the (approximated) validation losses are reported by averaging over all the selection vectors, and we compute the approximation error by  $\frac{|\text{real validation loss} - \text{approximated validation loss}|}{\text{real validation loss}}$ . The results are presented in Figure 2(a), where we can see: the approximation accuracy of IFRQE and IFRQE++ are satisfied. Especially, on Amazon Video, the approximation error of **IFRQE++** is only 1.2%. By leveraging more anchor vectors, **IFRQE++** is more accurate than **IFRQE** on both datasets. In specific, the approximation error is reduced from 19.2% to 15.2% and 9.4% to 1.2% on the datasets of Diginetica and Amazon Video, respectively.

### Influence of the Number of Anchor Vectors

In this section, we investigate the influence of parameter  $T$  on the final performance by tuning its value within the

range of  $[1, 10]$ . Notably,  $T = 1$  corresponds to the method IFRQE. Following the experimental settings described earlier, we present the reward and time cost for each  $T$  in Figure 2(b), with comprehensive results provided in the Appendix. From the results, we observe that the reward achieved by our framework is sub-optimal when  $T = 1$ , which aligns with the findings presented in Table 1. By increasing the number of anchor vectors, the reward improves. In most cases, as  $T$  increases, the reward initially rises before reaching an optimal point where it stabilizes. We speculate that initially, additional anchor vectors contribute to more accurate validation approximation, leading to better rewards. However, once a sufficient number of anchor vectors are included, further increasing  $T$  may not provide significant improvements as the validation approximation primarily relies on key anchor vectors. Regarding efficiency, as  $T$  increases, our model incurs a longer time cost. This outcome is expected since larger values of  $T$  necessitate more rounds of base recommender model retraining. Notably, these observations indicate that IFRQE++ offers an opportunity to balance the effectiveness and efficiency of the framework.

### Influence of Different User Willingness

Here we study whether our model is always effective on differently characterized user willingness. We base the experiments on the simulation datasets, where the willingness of a user is controlled by the sensitive parameter  $s_u$  and average willingness level  $b_u$ . Considering that the distributions of  $s_u$  and  $b_u$  are determined by  $a_1, a_2$  and  $a_3$ , we build four datasets by specifying  $(a_1, a_2, a_3)$  with  $A = (0.6, 2, 1.2)$ ,  $B = (0.4, 2, 0.8)$  and  $C = (0.6, 1, 1.5)$  and  $D = (0.5, 1, 1)$ , respectively, to simulate different user willingness characters. We present the comparisons of different models on these datasets in Figure 3, where we can see: the random method is less competitive, and sometimes, it is even worse than the base model (e.g., Datasets A and B with LightGCN as the base model). For different datasets and base models, IFRQE and IFRQE++ can usually obtain the second best and best performances. This result agrees with the observations in Table 1 and manifests that our methods are generally effective for different user willingness characters.

### Conclusion and Future Work

This paper improved traditional system-centered recommendation paradigm by allowing users to participate into the model optimization process via specifying their willingness. We built two models based on influence function to efficiently solve the above problem, and also provided theoretical guarantees. Extensive experiments were conducted to demonstrate the effectiveness of our models. However, there are some limitations can be studied in the future. We do not consider the dynamic nature of the recommender system, therefore, an interesting direction is to study the temporal user willingness, for example, under online settings. Additionally, we assume that the user overall utility is a linear combination between the recommendation quality and user willingness. One can extend it to the non-linear case, where more efforts are needed to solve the multiplayer game.

## Acknowledgements

This work is supported in part by National Natural Science Foundation of China (No. 62102420), Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098, Intelligent Social Governance Platform, Major Innovation & Planning Interdisciplinary Platform for the "Double-First Class" Initiative, Renmin University of China, Public Computing Cloud, Renmin University of China, fund for building world-class universities (disciplines) of Renmin University of China, Intelligent Social Governance Platform. The work is sponsored by Huawei Innovation Research Programs. We gratefully acknowledge the support from Mindspore<sup>9</sup>, CANN (Compute Architecture for Neural Networks) and Ascend Ai Processor used for this research.

## References

- Ben-Porat, O.; and Tennenholtz, M. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems*, 31.
- Calamai, P. H.; and Moré, J. J. 1987. Projected gradient methods for linearly constrained problems. *Mathematical programming*, 39(1): 93–116.
- Chen, C.; Sun, F.; Zhang, M.; and Ding, B. 2022a. Recommendation Unlearning. *arXiv preprint arXiv:2201.06820*.
- Chen, C.; Zhang, J.; Tung, A. K.; Kankanhalli, M.; and Chen, G. 2020. Robust federated recommendation system. *arXiv preprint arXiv:2006.08259*.
- Chen, X.; Xu, H.; Zhang, Y.; Tang, J.; Cao, Y.; Qin, Z.; and Zha, H. 2018. Sequential recommendation with user memory networks. In *Proceedings of the eleventh ACM international conference on web search and data mining*, 108–116.
- Chen, Z.; Sun, F.; Tang, Y.; Chen, H.; Gao, J.; and Ding, B. 2022b. Proactively Control Privacy in Recommender Systems. *arXiv preprint arXiv:2204.00279*.
- Chicco, D.; and Jurman, G. 2020. The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC genomics*, 21(1): 1–13.
- Halkidi, M.; and Koutsopoulos, I. 2011. A game theoretic framework for data privacy preservation in recommender systems. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 629–644. Springer.
- He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. Lightgen: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, 639–648.
- He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, 173–182.
- Kang, W.-C.; and McAuley, J. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*, 197–206. IEEE.
- Koh, P. W.; and Liang, P. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*, 1885–1894. PMLR.
- Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer*, 42(8): 30–37.
- Li, Y.; Zheng, X.; Chen, C.; and Liu, J. 2022. Making Recommender Systems Forget: Learning and Unlearning for Erasable Recommendation. *arXiv preprint arXiv:2203.11491*.
- Lin, G.; Liang, F.; Pan, W.; and Ming, Z. 2020. Fedrec: Federated recommendation with explicit feedback. *IEEE Intelligent Systems*, 36(5): 21–30.
- Parra, D.; and Brusilovsky, P. 2015. User-controllable personalization: A case study with SetFusion. *International Journal of Human-Computer Studies*, 78: 43–67.
- Ricci, F.; Rokach, L.; and Shapira, B. 2015. Recommender systems: introduction and challenges. In *Recommender systems handbook*, 1–34. Springer.
- Wang, L.; Zhang, J.; Chen, X.; Lin, Y.; Song, R.; Zhao, W. X.; and Wen, J.-R. 2023. RecAgent: A Novel Simulation Paradigm for Recommender Systems. *arXiv preprint arXiv:2306.02552*.
- Wang, W.; Feng, F.; Nie, L.; and Chua, T.-S. 2022. User-controllable Recommendation Against Filter Bubbles. *arXiv preprint arXiv:2204.13844*.
- Wang, X.; He, X.; Cao, Y.; Liu, M.; and Chua, T.-S. 2019. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 950–958.
- Wang, Z.; Zhang, J.; Xu, H.; Chen, X.; Zhang, Y.; Zhao, W. X.; and Wen, J.-R. 2021. Counterfactual data-augmented sequential recommendation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 347–356.
- Wu, C.; Wu, F.; Cao, Y.; Huang, Y.; and Xie, X. 2021. Fedgcn: Federated graph neural network for privacy-preserving recommendation. *arXiv preprint arXiv:2102.04925*.
- Xu, L.; Jiang, C.; Chen, Y.; Ren, Y.; and Liu, K. R. 2018. User participation in collaborative filtering-based recommendation systems: a game theoretic approach. *IEEE transactions on cybernetics*, 49(4): 1339–1352.
- Yang, L.; Tan, B.; Zheng, V. W.; Chen, K.; and Yang, Q. 2020. Federated recommendation systems. In *Federated Learning*, 225–239. Springer.

<sup>9</sup><https://www.mindspore.cn>