



Inverse Learning with Extremely Sparse Feedback for Recommendation

Guanyu Lin
Carnegie Mellon University
Pittsburgh, PA, United States
guanyul@andrew.cmu.edu

Chen Gao†
Tsinghua University
Beijing, China
chgao96@gmail.com

Yu Zheng
Tsinghua University
Beijing, China

Yinfeng Li
Kuaishou Technology
Beijing, China

Jianxin Chang
Kuaishou Technology
Beijing, China

Yanan Niu
Kuaishou Technology
Beijing, China

Yang Song
Kuaishou Technology
Beijing, China

Kun Gai
Unaffiliated
China

Zhiheng Li
Tsinghua University
Beijing, China

Depeng Jin
Tsinghua University
Beijing, China

Yong Li
Tsinghua University
Beijing, China

ABSTRACT

Modern personalized recommendation services often rely on user feedback, either explicit or implicit, to improve the quality of services. Explicit feedback refers to behaviors like ratings, while implicit feedback refers to behaviors like user clicks. However, in the scenario of full-screen video viewing experiences like Tiktok and Reels, the click action is absent, resulting in unclear feedback from users, hence introducing noises in modeling training. Existing approaches on de-noising recommendation mainly focus on positive instances while ignoring the noise in a large amount of sampled negative feedback. In this paper, we propose a meta-learning method to annotate the unlabeled data from loss and gradient perspectives, which considers the noises in both positive and negative instances. Specifically, we first propose an *Inverse Dual Loss* (IDL) to boost the true label learning and prevent the false label learning. Then we further propose an *Inverse Gradient* (IG) method to explore the correct updating gradient and adjust the updating based on meta-learning. Finally, we conduct extensive experiments on both benchmark and industrial datasets where our proposed method can significantly improve AUC by 9.25% against state-of-the-art methods. Further analysis verifies the proposed inverse learning framework is model-agnostic and can improve a variety of recommendation backbones. The source code, along with the best hyper-parameter settings, is available at this link: <https://github.com/tsinghua-fib-lab/WSDM24-InverseLearning>.

†Chen Gao is the corresponding author (chgao96@gmail.com).

KEYWORDS

De-noising, Meta Learning, Recommendation, Sparse Feedback

ACM Reference Format:

Guanyu Lin, Chen Gao†, Yu Zheng, Yinfeng Li, Jianxin Chang, Yanan Niu, Yang Song, Kun Gai, Zhiheng Li, Depeng Jin, and Yong Li. 2024. Inverse Learning with Extremely Sparse Feedback for Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*, March 4–8, 2024, Merida, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3616855.3635797>

1 INTRODUCTION

Recommender systems [23, 26, 30] aim to capture users' preferences based on their historical behaviors, with a focus on either explicit or implicit feedback. Explicit feedback, such as user ratings, provides direct indications of user preferences but is challenging to collect due to the need for active user participation [15, 21, 33]. In contrast, implicit feedback, including user clicks, purchases, and views, offers richer information and is more commonly utilized in modern recommender systems [4, 20]. In certain contexts like Micro-video platforms, users passively receive recommended items without actively engaging through actions like clicking or rating. Consequently, we encounter a scenario where the labeled feedback is extremely sparse, comprising predominantly quick-skip, long-stay, and a considerable number of slow-skip or short-stay videos with unclear feedback. Effectively leveraging this unlabeled feedback poses a significant challenge for recommendation systems.

The challenge of dealing with unclear feedback in recommender systems has led to various approaches that randomly sample unlabeled data and consider it as negative feedback, resulting in inevitable noise [3, 12, 14]. Typically, user-clicked data is treated as positive feedback, while unclicked data is sampled as negative feedback [3, 14]. However, this sampling strategy may include positive instances in the unlabeled data, leading to false-negative cases. Additionally, some studies have explored hard negative sampling



This work is licensed under a Creative Commons Attribution International 4.0 License.

WSDM '24, March 4–8, 2024, Merida, Mexico
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0371-3/24/03.
<https://doi.org/10.1145/3616855.3635797>

techniques, which reduce false-positive instances but increase false-negative instances [7, 8, 34]. Nevertheless, these methods often underperform when evaluated on true positive and negative data instead of the sampled negative data alone. Notably, a recent work called DenoisingRec [28] focuses on denoising positive feedback by manipulating the loss of false-positive instances but does not adequately address the issue of noisy negative feedback. Overall, existing approaches tend to concentrate solely on either the positive or negative perspective, without effectively tackling both aspects. Our early experimental analysis reveals two key observations emerge, serving as the motivation behind our proposed method: (1) **Full use of unlabeled data can boost the performance**, (2) **Labeled data can guide the learning on unlabeled data**.

To simultaneously tackle the unclear feedback problem from positive and negative perspectives, we propose a novel learning-based approach that employs Inverse Dual Loss (IDL) and Inverse Gradient (IG). Our method automatically annotates the unlabeled data and subsequently adjusts the falsely annotated labels. Formally speaking, we introduce IDL for unsupervised training on unlabeled data and leverage IG to guide the unlabeled data. Specifically, the IDL employs a well-designed loss function that leverages both positive and negative feedback, which can automatically annotate unlabeled data with unsupervised learning. We exploit the property that the loss associated with a false positive/negative instance exceeds that of a true positive/negative instance [28]. By assigning different weights to the positive and negative labels of unlabeled instances, calculated using the inverse dual loss, we effectively utilize true positive/negative instances while mitigating the noise introduced by false positive/negative instances. This approach allows us to fully capitalize on valuable information and enhance the quality of annotation. In addition, to adjust the false annotated labels and improve the robustness of IDL, we further propose an Inverse Gradient (IG) method. Here we build a meta-learning process [10, 18] and split the training data into training-train and training-test data. We first exploit training-train data to pre-train the model. Then we further use training-test data to validate the correctness of classification by IDL. In other words, supervising the proposed unsupervised IDL method via training-test data. Specifically, we calculate the gradient for the inverse dual loss of sampled instances as well as the additive inverse of the gradient. The model is optimized by either the direct gradient or the additive inverse of gradient, determined by the split training-test data. Experimental results illustrate that inverse gradient can truly improve the inverse dual loss. In summary, the main contributions of this paper are as follows:

- We take the pioneering step to address the unclear passive feedback in video feed recommendation, which is far more challenging than existing works that are either based on explicit or implicit active feedback.
- We propose Inverse Dual Loss (IDL) to annotate the labels for sampled instances in an unsupervised learning manner. Besides, we further propose Inverse Gradient to guide the unsupervised learning on unlabeled data and improve the robustness of IDL.
- We experiment on two real-world datasets, verifying the superiority of our method compared with state-of-the-art approaches. Further studies sustain the effectiveness of our proposed method in label annotation and convergence.

2 PROBLEM DEFINITION

We will formulate the problem here. The recommendation task aims to model relevance score $\hat{y}_{ui}^\theta = f(u, i|\theta)$ of user u towards item i under parameters θ . The *LogLoss* function [35, 36] function to learn ideal parameters θ^* is as:

$$\mathcal{L}_{\mathcal{D}^*}(\theta) = \frac{1}{|\mathcal{D}^*|} \sum_{(u, i, y_{ui}^*) \in \mathcal{D}^*} \ell(\hat{y}_{ui}^\theta, y_{ui}^*), \quad (1)$$

where $\ell(\hat{y}_{ui}^\theta, y_{ui}^*) = - (y_{ui}^* \log(\hat{y}_{ui}^\theta) + (1 - y_{ui}^*) \log(1 - \hat{y}_{ui}^\theta))$, $y_{ui}^* \in \{0, 1\}$ is the feedback of user u towards item i . $\mathcal{D}^* = \{(u, i, y_{ui}^*)\}$, $u \in \mathcal{U}, i \in \mathcal{I}$ is the reliable interaction data between all user-item pairs. Indeed, due to the limited collected feedback, the model training is truly formalized as follows: $\hat{\theta} = \arg \min_{\theta} \mathcal{L}_{\mathcal{D}^l}(\theta) + \mathcal{L}_{\mathcal{D}^u}(\theta)$, where $\mathcal{D}^l \sim \mathcal{D}^*$ is the collected labeled data, and $\mathcal{D}^u = \{(u, i, \bar{y}_{ui}) \mid u \in \mathcal{U}, i \in \mathcal{I}\}$ is the sampled unlabeled data where $\bar{y}_{ui} = 0$ is often assumed in existing recommenders for negative sampling. However, such a strategy will inevitably introduce noise because there are some positive unlabeled instances in the sampled data. As a consequence, a model (i.e., $\hat{\theta}$) trained with noisy data tends to exhibit suboptimal performance. Thus, our goal is to construct a denoising recommender approximating to the ideal recommender θ^* as:

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\mathcal{D}^l}(\theta) + \mathcal{L}_{\mathcal{D}^u}^{\text{denoise}}(\theta), \quad (2)$$

where $\mathcal{L}_{\mathcal{D}^u}^{\text{denoise}}(\theta)$ indicates the loss on unlabeled data with all samples annotated correctly, *i.e.* denoising sampling.

3 METHODOLOGY

In this section, we will first perform an in-depth analysis of existing solutions and their limitations. Then we will propose inverse dual loss to address the limitations of existing works for easy samples. Finally, we further propose inverse gradient to address the limitation of inverse dual loss and make it capable of not only easy samples but also hard samples that are misclassified.

3.1 Inverse Dual Loss

In this section, we first analyze the characteristics of existing solutions on the sampled unlabeled data. Then we introduce the proposed inverse dual loss solution to denoise sampled data.

3.1.1 Analysis of Existing Approach. We first explain the data sparsity problem in recommender systems from the perspective of classification boundary, based on which we will introduce the existing solutions. As shown in Figure 1 (a), in recommender systems, labeled data tends to be extremely sparse compared with a large number of unlabeled data. A recommendation model is prone to overfitting if it is only trained based on the sparsely labeled data, compared with the ground-truth in Figure 1 (h).

In practice, existing recommenders often sample from unlabeled data and treat all the sampled data as negative feedback. Such an approach introduces false negative, which fails to retrieve items that users may be interested in, as shown by the classification boundary in Figure 1 (b). That is, there exists noise in the sampled negative data. However, existing denoising approaches mainly focus on the

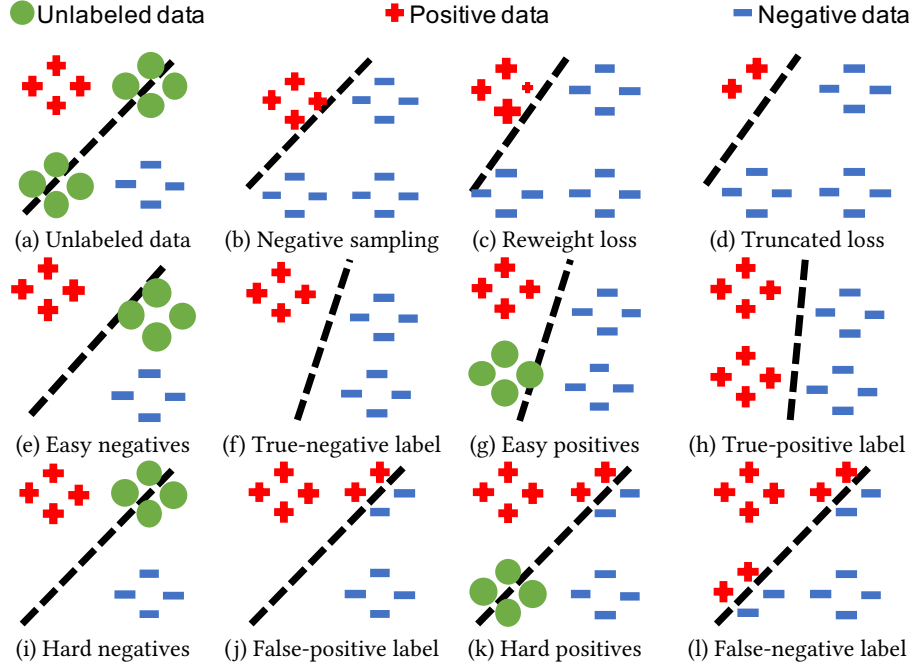


Figure 1: Illustrations of existing solutions and our inverse dual loss’s effectiveness and limitation. (a)-(d) are the illustrations of existing solutions: (a) illustrates there are a lot of unlabeled data; (b) illustrates the traditional negative sampling approach; (c) illustrates the reweighted loss of DenoisingRec; (d) illustrates the truncated loss of DenoisingRec adapted on the false-negative instance. (e)-(f) are the illustrations of our inverse dual loss’s effectiveness with easy sampling: (e) illustrates the easy negative instances are sampled; (f) illustrates labeling the sampled instances as true negative; (g) illustrates the easy positive instances are sampled; (h) illustrates labeling the sampled instances as true positive and approximate to ground-truth. (i)-(l) are the illustrations of our inverse dual loss’s limitation with hard sampling: (i) illustrates the hard negative instances are sampled; (j) illustrates labeling part of the sampled instances as false positive; (k) illustrates the hard positive instances are sampled; (l) illustrates labeling part of the sampled instances as false negative.

noise in positive samples (false positive). For example, DenoisingRec [28] attempts to achieve denoising for false positive instances as:

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\mathcal{D}^l \cup \mathcal{D}^{noise}}^{\text{denoise}}(\theta) + \mathcal{L}_{\mathcal{D}^u}(\theta) \quad (3)$$

where $\mathcal{D}^{noise} = \{(u, i, 1) \mid u \in \mathcal{U}, i \in \mathcal{I}, y_{ui}^* = 0\}$ is the noisy false positive data they introduce in experiments. For example, R-CE (Reweight Cross-Entropy) of DenoisingRec assigns lower weight on false-positive instances with large loss (Figure 1 (c)), and T-CE (Truncated Cross-Entropy) of DenoisingRec discards those false positive instances with large loss (Figure 1 (d)). Though achieving denoising for false-positive instances, they ignore false-negative instances and fail to address the noise of negative sampling.

In fact, the collected labeled data is much cleaner than sampled unlabeled data, and the number of false-positive instances is limited in real-world recommender systems. On the contrary, the noise brought by negative sampling is far more harmful. In other words, the noise level of positive unlabeled data incorrectly sampled as negative feedback is much higher than that of negative samples wrongly regarded as positive feedback.

To sum up, existing solutions either introduce noise or perform incomplete denoising, which motivates us to further propose a denoising solution for unlabeled data from both positive and negative perspectives.

3.1.2 Labeling with Inverse Dual Loss. As an existing attempt in DenoisingRec, we have discovered that the false positive instances are with a greater loss. It is also an apparent phenomenon in machine learning. For example, if we have a positive instance and a well-trained model, the loss of classifying it as negative will be greater than that of classifying it as positive. Otherwise, if we have a negative instance and a well-trained model, the loss of classifying it as positive will be greater. Hence, we can assume the sampled unlabeled instances are both possibly positive and negative and then exploit this inherent characteristic to automatically weigh more on the true positive or negative instances while weighing less on the false ones.

DEFINITION 1. (Inverse Dual Loss) The inverse dual loss is defined as denoising loss to automatically classify the unlabeled data as:

$$\mathcal{L}_{\mathcal{D}^u}^{\text{dual}}(\theta) = \frac{1}{|\mathcal{D}^u|} \sum_{(u, i, \hat{y}_{ui}) \in \mathcal{D}^u} w^1 \ell(\hat{y}_{ui}^{\theta}, 1) + w^0 \ell(\hat{y}_{ui}^{\theta}, 0), \quad (4)$$

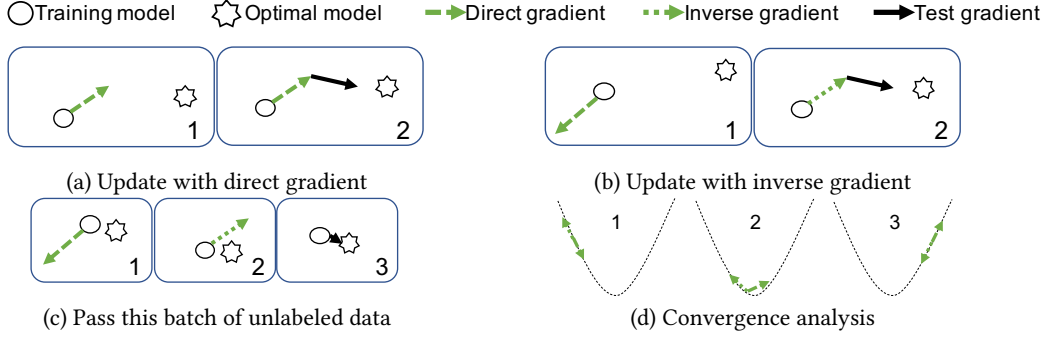


Figure 2: Illustrations of inverse gradient adaption.

where $w^1 = \frac{\text{stopgrad}(\ell(\hat{y}_{ui}^0, 0))}{z_w \text{stopgrad}(\ell(\hat{y}_{ui}^0, 1))}$, $w^0 = \frac{\text{stopgrad}(\ell(\hat{y}_{ui}^0, 1))}{z_w \text{stopgrad}(\ell(\hat{y}_{ui}^0, 0))}$ are the weights for positive loss and negative loss, respectively. $z_w = \frac{\text{stopgrad}(\ell(\hat{y}_{ui}^0, 0))}{\text{stopgrad}(\ell(\hat{y}_{ui}^0, 1))} + \frac{\text{stopgrad}(\ell(\hat{y}_{ui}^0, 1))}{\text{stopgrad}(\ell(\hat{y}_{ui}^0, 0))}$ is the normalization parameter. \mathcal{D}^u is the sampled unlabeled data and θ is the model parameters to be learned. Here **stopgrad** is a stop-gradient operation.

The pros of inverse dual loss are as shown in (e)-(h) of Figure 1: (e) when the easy negative instances are sampled, the loss of classifying them as positive will be greater than that of classifying them as negative, and thus inverse dual loss will assign more weights on the negative loss; (f) gradually assigning more and more weights on the negative loss, the negative unlabeled instances will eventually be classified as negative; (g) likewise, when the easy positive instances are sampled, the inverse dual loss will assign more weights on the positive loss; (h) the positive unlabeled instances will eventually be classified as positive and approximate to the ground-truth.

3.1.3 Limitation of Inverse Dual Loss. When sampling the easy positive or negative instances, our inverse dual loss can boost learning by correctly labeling the sampled instances. However, it may be an obstacle when there are some hard positive or negative instances. As shown in (i)-(l) of Figure 1, given some hard positive or negative instances, the classification boundary will be prevented from the ground-truth: (i) the hard negative instances are sampled; (j) half of the negative unlabeled instances will be classified as positive and become noise; (k) the hard positive instances are sampled; (l) half of the positive unlabeled instances will be classified as negative and become noise. That is to say, our inverse dual loss relies heavily on the current training classification boundary and the difficulty of sampled data, requiring us to further improve its robustness.

3.2 Inverse Gradient

To improve the robustness of Inverse Dual Loss on the hard sampled data towards the current training model, in this section, we further propose inverse gradient to adjust the gradient of false annotated data, inspired by the meta-learning framework [10, 18]. Then we analyze the convergence of the proposed Inverse Gradient.

3.2.1 Learning to Label with Inverse Gradient. In this part, we introduce our solution for tackling the false annotated instances of Inverse Dual Loss.

DEFINITION 2. (Inverse Gradient) We define the gradient and additive inverse of gradient calculated by (4) w.r.t. $\nabla \mathcal{L}_{\mathcal{D}^u}^{\text{dual}}(\theta)$ and $-\nabla \mathcal{L}_{\mathcal{D}^u}^{\text{dual}}(\theta)$ as direct gradient and inverse gradient, respectively, of the loss for unlabeled data \mathcal{D}^u .

THEOREM 1. Given learning rate $\alpha \in \mathbb{R}, \alpha \neq 0$, assume the temporal model parameters updated by the direct gradient and inverse gradient, respectively, are as $\theta^d = \theta - \alpha \circ \nabla \mathcal{L}_{\mathcal{D}^u}^{\text{dual}}(\theta)$ and $\theta^i = \theta + \alpha \circ \nabla \mathcal{L}_{\mathcal{D}^u}^{\text{dual}}(\theta)$. Then, the relationship between the loss of them and the model with parameter θ on data \mathcal{D}^l will be either $\mathcal{L}_{\mathcal{D}^l}(\theta^d) > \mathcal{L}_{\mathcal{D}^l}(\theta) > \mathcal{L}_{\mathcal{D}^l}(\theta^i)$ or $\mathcal{L}_{\mathcal{D}^l}(\theta^d) < \mathcal{L}_{\mathcal{D}^l}(\theta) < \mathcal{L}_{\mathcal{D}^l}(\theta^i)$.

Based on this theorem, we can have the following gradient updating strategies. Generally, we will first split the training data into training-train data and training-test data, where we pre-train the model on the training-train data. Then we will calculate the direct gradient of inverse dual loss on the sampled unlabeled data, which can further result in the following three cases:

- When the sampled data is easy, we can exploit the direct gradient to update the model, and it will gain a smaller test loss on the training-test data, as shown in Figure 2 (a);
- When the sampled data is hard, exploiting the inverse gradient to update the model will gain a smaller test loss on the training-test data, and thus we exploit inverse gradient here as shown in Figure 2 (b);
- When the model is approximately optimal, either direct gradient or inverse gradient will prevent it from ground-truth, and we discard this batch of unlabeled data as shown in Figure 2 (c);

3.2.2 Algorithm. We present the procedure of exploring these three cases as Algorithm 1. The algorithm first pre-trains the model using the split training-train data. Then the model will update with direct gradient or inverse gradient or even not update, determined by the validation on the split training-test data.

More specifically, the inputs of our proposed algorithm are labeled data \mathcal{D}^l , unlabeled data \mathcal{D}^u , and learning rate γ, α . The first iteration aims to pre-train the model using the split training-train data. The second iteration aims to explore the direct gradient and inverse gradient on the loss for sampled unlabeled data, where three strategies are explored here as lines 13-15 of Algorithm 1 with

Algorithm 1: Inverse Gradient Adaptation

input : Labeled data \mathcal{D}^l , unlabeled data \mathcal{D}^u , learning rate γ , α
output: θ

- 1 Initialize θ , split \mathcal{D}^l into training-train data $\text{train}(\mathcal{D}^l)$ and training-test data $\text{test}(\mathcal{D}^l)$;
- 2 **while** not done **do**
- 3 **for** $t = 1$ **to** T **do**
- 4 Sample batch of training labeled data $\text{train}(\mathcal{D}_t^l) \sim \mathcal{D}^l$;
- 5 $\mathcal{L}_{\text{train}(\mathcal{D}_t^l)}(\theta) = \frac{1}{|\text{train}(\mathcal{D}_t^l)|} \sum_{(u,i,y_{ui}^*) \in \text{train}(\mathcal{D}_t^l)} \ell(\hat{y}_{ui}^\theta, \hat{y}_{ui}^*)$;
- 6 $\theta = \theta - \gamma \circ \nabla \mathcal{L}_{\text{train}(\mathcal{D}_t^l)}(\theta)$;
- 7 **for** $t = 1$ **to** T **do**
- 8 Sample batch of training unlabeled data $\text{train}(\mathcal{D}_t^u) \sim \mathcal{D}^u$;
- 9 $\mathcal{L}_{\text{train}(\mathcal{D}_t^u)}^{\text{dual}}(\theta) = \frac{1}{|\text{train}(\mathcal{D}_t^u)|} \sum_{(u,i,y_{ui}) \in \text{train}(\mathcal{D}_t^u)} w^1 \ell(\hat{y}_{ui}^\theta, 1) + w^0 \ell(\hat{y}_{ui}^\theta, 0)$;
- 10 $\theta^d = \theta - \alpha \circ \nabla \mathcal{L}_{\text{train}(\mathcal{D}_t^u)}^{\text{dual}}(\theta)$;
- 11 $\theta^i = \theta + \alpha \circ \nabla \mathcal{L}_{\text{train}(\mathcal{D}_t^u)}^{\text{dual}}(\theta)$;
- 12 Sample batch of test labeled data $\text{test}(\mathcal{D}_t^l) \sim \mathcal{D}^l$;
- 13 $\mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta^d) = \frac{1}{|\text{test}(\mathcal{D}_t^l)|} \sum_{(u,i,y_{ui}^*) \in \text{test}(\mathcal{D}_t^l)} \ell(\hat{y}_{ui}^{\theta^d}, \hat{y}_{ui}^*)$;
- 14 $\mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta^i) = \frac{1}{|\text{test}(\mathcal{D}_t^l)|} \sum_{(u,i,y_{ui}^*) \in \text{test}(\mathcal{D}_t^l)} \ell(\hat{y}_{ui}^{\theta^i}, \hat{y}_{ui}^*)$;
- 15 $\mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta) = \frac{1}{|\text{test}(\mathcal{D}_t^l)|} \sum_{(u,i,y_{ui}^*) \in \text{test}(\mathcal{D}_t^l)} \ell(\hat{y}_{ui}^\theta, \hat{y}_{ui}^*)$;
- 16 $\theta = \arg \min_{\{\theta^d, \theta, \theta^i\}} \{ \mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta^d), \mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta), \mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta^i) \}$;
- 17 $\mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta) = \min \{ \mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta^d), \mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta), \mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta^i) \}$;
- 18 $\theta = \theta - \gamma \circ \nabla \mathcal{L}_{\text{test}(\mathcal{D}_t^l)}(\theta)$;

θ^d , θ^i and θ as the model parameters updated by the direct gradient, inverse gradient and without being updated by the gradient on the loss for sampled unlabeled data, respectively. Finally, the explored updated direction with minimal test loss on the training-test data will be selected to update the model for this iteration.

3.2.3 Convergence Analysis. As shown in Figure 2 (d), the first case illustrates when the unlabeled data is ideally sampled, the updating with direct gradient will lead to a smaller loss and better convergence on the training-test data, while the third case with poorly sampled data supposes to update with inverse gradient. However, when the model approximates convergence on the training-test data, updating with either direct gradient or inverse gradient may be poorer than no updating, as shown in stage 2 of Figure 2 (d).

To avoid the gradient ascent problem for the second case, we can set the learning rate for the inverse dual loss to be smaller than that

Table 1: Data statistics for processed Micro Video dataset and ML1M dataset.

| Dataset | Micro Video | ML1M | |
|-----------------|-----------------|-----------|-----------|
| Users | 37,692 | 6,041 | |
| Items | 131,690 | 3,953 | |
| Feedback | Positive | 4,915,745 | 836,478 |
| | Negative | 4,546,747 | 163,731 |
| | Total | 9,462,492 | 1,000,209 |
| Density | 0.19% | 4.19% | |

for the test loss, i.e., $\alpha < \gamma$ in Algorithm 1. In this way, the scale of updating by the gradient for inverse dual loss will be within the scale of updating by the gradient for test loss. That is to say, the gradient ascent problem is less likely to occur on the inverse dual loss for unlabeled data than the loss for labeled data.

4 EXPERIMENTS

In this section, we perform experiments on two real-world datasets, targeting four research questions (RQs):

- **RQ1:** How does the proposed method perform compared with the state-of-the-art denoising recommenders? What is the effect of two proposed components, i.e., Inverse Gradient (IG) and Inverse Dual Loss (IDL)?
- **RQ2:** How does our proposed inverse dual loss identify the unlabeled data?
- **RQ3:** What is the effect of the inverse gradient on convergence?
- **RQ4:** What is the optimal ratio between the learning rates for inverse dual loss and training-test?

4.1 Experimental Setup

4.1.1 Datasets. To practice and verify the effectiveness of our proposed method, we conduct experiments on an industrial Micro Video dataset and a public benchmark ML1M dataset, which is widely used in existing work for recommender systems [6, 19]. Micro Video is an extremely sparse dataset where users are passive in receiving the feed videos and have rare active feedback. The details of them are as Table 1.

4.1.2 Baselines and Evaluation Metrics. To demonstrate the effectiveness of our proposed inverse learning on unlabeled data, we compared the performance of recommenders trained by our inverse gradient (IG) with recommenders trained by inverse dual loss (IDL) and normal training by standard loss or negative sampling (NS) [3, 14, 25]. Besides we also compare our inverse learning method with the state-of-the-art methods for denoising recommender systems. Specifically, we also compare two adaptive denoising training strategies, T-CE and R-CE, of DenoisingRec [28]. Following DenoisingRec [28], we select GMF and NeuMF [14] as backbones, which are neural Collaborative Filtering models.

We adopt widely-used AUC and GAUC as accuracy metrics [13]. Besides, two widely-used ranking metrics [2], MRR and NDCG@10, are also adopted for evaluation.

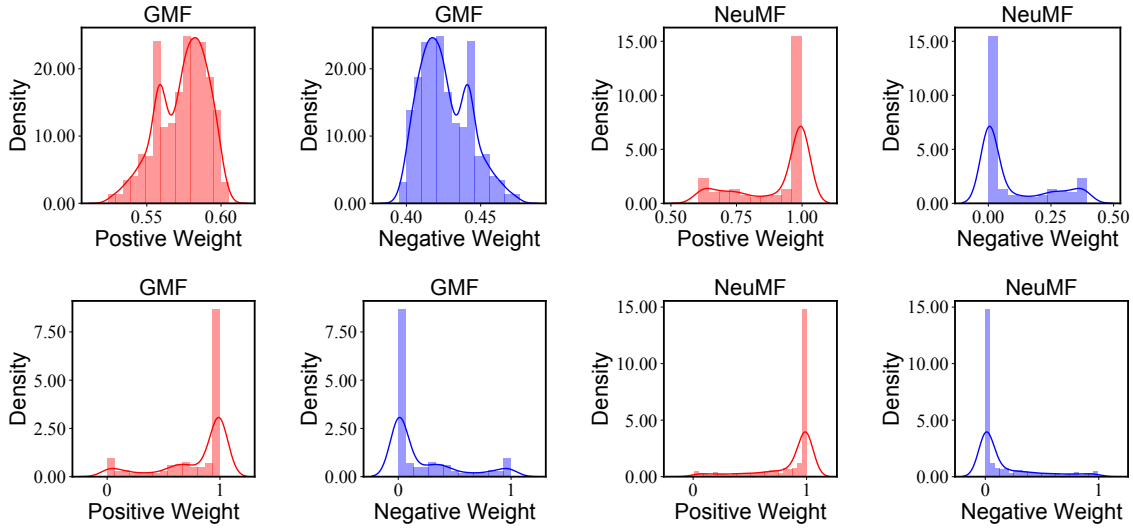


Figure 3: Positive and negative weight distributions for dual loss on ML1M at first (up) and final (bottom) epochs.

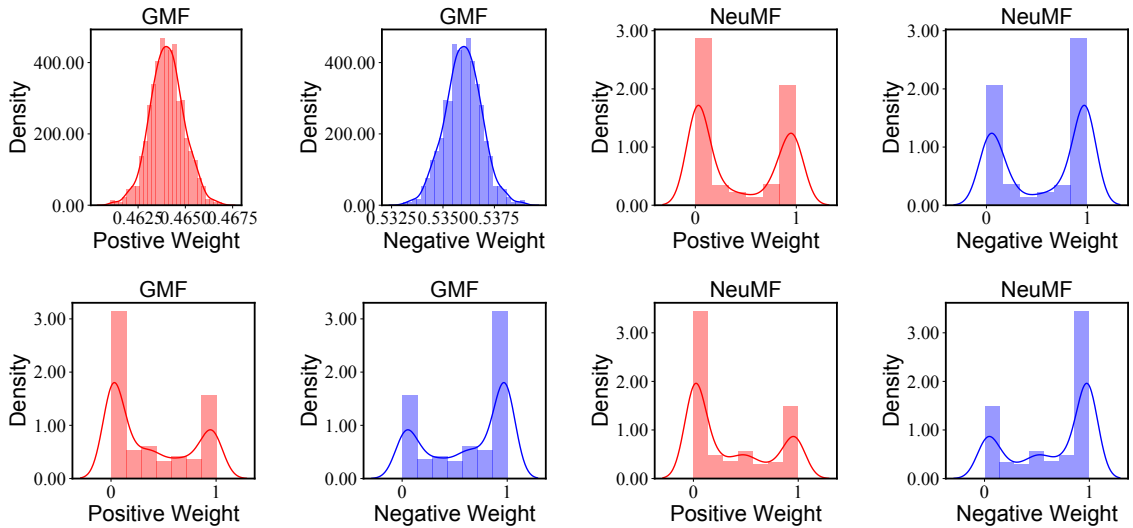


Figure 4: Positive and negative weight distributions for dual loss on Micro Video at first (up) and final (bottom) epochs.

4.1.3 Hyper-parameter Settings. For the two denoising strategies [28], we followed their default settings and verified the effectiveness of our methods under the same conditions. The embedding size and batch size of all models are set as 32 and 1,024, respectively. Besides, we adopt Adam [17] to optimize all the model parameters with the learning rate γ initialized as 0.0001 and 0.00001 for labeled data on ML1M and Micro Video datasets, respectively, while the learning rate for sampled data is set as $\alpha = 0.1\gamma$. As for the inverse gradient, we split 90% of the training data as training-train data, and the left is used as training-test data. The sampling rate is set as 1. The provided code has included the best hyper-parameters.

4.2 Overall Performance (RQ1)

The performance comparison is shown in Table 2, from which we have the following observations.

- **Our inverse gradient performs best.** Our inverse gradient (IG) method achieves the best performance compared with four baselines and our inverse dual loss (IDL) for three metrics. Specifically, our IG improves the backbone sharply, which shows the ability of our proposed method to well classify the unlabeled data and achieve effective data augmentation to resolve the data sparsity problem of existing recommenders. Note that on apart from IDL and IG, GMF is better than NeuMF in general. But NeuMF

Table 2: Performance comparisons with GMF and NeuMF backbones on two datasets. Bold and underline refer to the best and second best results, respectively. Here IG includes the IDL method.

| Model | Method | ML1M | | | Micro Video | | |
|-------|--------|---------------|---------------|---------------|---------------|---------------|---------------|
| | | NDCG | AUC | GAUC | NDCG | AUC | GAUC |
| GMF | None | 0.9285 | <u>0.7671</u> | 0.6919 | 0.7365 | <u>0.8024</u> | <u>0.7558</u> |
| | NS | <u>0.9400</u> | <u>0.7639</u> | 0.7200 | 0.7049 | <u>0.7802</u> | <u>0.7247</u> |
| | T-CE | 0.9349 | 0.7612 | 0.7163 | 0.6994 | 0.6486 | 0.6956 |
| | R-CE | 0.9391 | 0.7632 | <u>0.7192</u> | 0.7069 | 0.7820 | 0.7260 |
| | IDL | 0.8996 | 0.7304 | <u>0.6400</u> | 0.7272 | 0.7858 | 0.7335 |
| | IG | 0.9521 | 0.8318 | 0.7642 | 0.7773 | 0.8033 | 0.7593 |
| NeuMF | None | 0.9214 | 0.7524 | 0.6856 | 0.6649 | <u>0.7504</u> | <u>0.6916</u> |
| | NS | 0.9298 | 0.7495 | 0.7088 | 0.6350 | 0.7191 | 0.6689 |
| | T-CE | 0.9351 | 0.7587 | <u>0.7158</u> | <u>0.6725</u> | 0.7469 | 0.6911 |
| | R-CE | <u>0.9349</u> | 0.7521 | 0.7117 | 0.6288 | 0.7038 | 0.6568 |
| | IDL | <u>0.9212</u> | <u>0.7962</u> | 0.6984 | 0.6466 | 0.7174 | 0.6647 |
| | IG | 0.9449 | 0.8253 | 0.7569 | 0.7809 | 0.8198 | 0.7689 |

with IG can significantly improve the performance, which is even better than GMF with IG. This is because NeuMF is a deep-based model, which will overfit when data is less or noise. Besides, IG outperforms the existing negative sampling (NS) method, which means there is truly a large number of positive unlabeled data, and directly treating them all as negative feedback will confuse the model. Finally, IG also outperforms existing state-of-art denoising methods, T-CE and R-CE, showing the importance of tackling the noise from both positive and negative feedback.

- **Inverse gradient can improve inverse dual loss to be more robust.** Inverse dual loss (IDL) only outperforms the NeuMF on ML1M dataset for AUC and GAUC, which shows the inferior robustness of IDL since it depends heavily on the training model and sampled data. Besides, it is even outperformed by GMF on ML1M and NeuMF on Micro Video, which means the poor sampling of unlabeled data will have a negative impact on the model training. These results show the significance of improving the robustness of IDL and confirm that it is necessary to exploit IG to adjust IDL.

4.3 Annotation on Unlabeled Data (RQ2)

To study the ability of our proposed method to annotate the unlabeled data, we visualize the distribution of weights for positive loss and negative loss in the ML1M and Micro Video datasets, respectively, as Figure 3 and 4. In specific, the upper part represents the weights at the first epoch after pre-training on the training-train data, and the bottom part represents the weights at the convergence epoch. From the figure, we can observe that:

- **Both deep and linear models are improved maximally by our method.** After convergence, both GMF and NeuMF can well classify the unlabeled data and capture similar patterns. In the ML1M dataset, most sampled instances are labeled as positive with more positive weights. In the Micro Video dataset, sampled instances are labeled in a more balanced manner. It means the deep model and linear model have the same upper bound on these two datasets, which is also consistent with the results of IG in

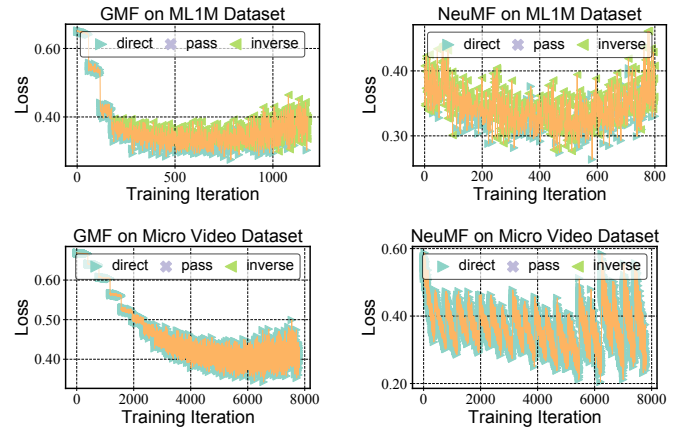


Figure 5: The loss on training-test data with the adapted gradient from the dual loss of unlabeled data. Each point of the loss curve is marked by its updated gradient direction.

overall performance where these two methods have competitive performance. This also shows the robustness of our proposed IG for improving both linear and deep models to their upper bound.

- **Deep model annotates the unlabeled data faster than the linear model.** It is obvious that at the beginning, the linear GMF model fails to well classify the unlabeled data but the deep NeuMF model has well captured the pattern of unlabeled data, which shows the generalization ability of deep learning model [5].

4.4 Convergence Analysis (RQ3)

To investigate the convergence of our proposed inverse gradient, we also plot the loss curve of training-test data for the ML1M and Micro Video datasets on the upper and bottom parts in Figure 5, respectively. Based on the results, we can discover that:

- **Inverse Gradient can promote convergence.** For the ML1M dataset, at the early stage, GMF model is updated with direct gradient, then with a hybrid of direct and inverse gradients, and finally with inverse gradient. In NeuMF model, we can discover more hybrid gradients in the valley of the loss curve. This is because the test gradient is more likely to ascent at the valley where our Inverse Gradient inverts the gradient for dual loss to adjust its direction for better convergence.
- **Proper learning rate can prevent gradient ascent.** For the Micro Video dataset, the models are always updated with direct gradient. This is because the learning rate is relatively low here (as analyzed in the Section 4.5), leading to almost no gradient ascent problem. Most importantly, we can discover that there is no case of passing gradient in the descent procedures, which supports our analysis at Section 3.2.3 that setting a smaller value of learning rate α can avoid the gradient ascent problem for inverse dual loss.

4.5 Hyper-parameter Study (RQ4)

As discussed in Section 3.2.3, we can avoid the gradient ascent problem by setting the learning rate for inverse dual loss smaller




Figure 6: The loss of GMF model on training-test data with different learning rates. Each point of the loss curve is marked by its updated gradient direction.

than that for test loss, i.e., $\alpha < \gamma$ in Algorithm 1. To experimentally study this convergence analysis at Section 3.2.3 and investigate the impact of the learning rate for convergence, we vary the learning rate for inverse dual loss with 1, 10, 50, and 100 times the learning rate for test loss. Here we study the loss curve of GMF on Micro Video as Figure 6, where we can discover that:

- **Smaller learning rate for inverse dual loss can avoid gradient ascent.** When α is smaller than 10 times γ , the gradients are often direct gradients. However, when α grows up to 50 times γ , the inverse gradients appear. Moreover, when α grows up to 100 times γ , there even appear pass gradients which means the occurrence of gradient ascent. Thus it is consistent to our analysis at Section 3.2.3 that we can avoid the gradient ascent problem by limiting the learning rate for inverse dual loss according to the learning rate for test loss.
- **Greater learning rate for inverse dual loss can speed up the convergence but more fluctuation.** With the growth of learning rate α , the loss convergence becomes faster. However, it also results in fluctuation as there appear more inverse gradients and pass gradients. This observation is consistent to our analysis of Section 3.2.3 for convergence.

5 RELATED WORK

Implicit Feedback with Negative Sampling. Existing recommenders are generally based on implicit feedback data, where the collected data is often treated as positive feedback, and negative sampling [3, 14, 25] is exploited to balance the lack of negative instances. However, the negative sampling strategy will introduce noise because there are some positive unlabeled [1, 9, 27] data in the sampled instances. To improve existing implicit feedback recommendation, the identification of negative experiences [11, 16] has grabbed the researchers’ attention. However, these methods collect either the various user feedback (e.g., dwell time [16] and skip [29]) or the item characteristics [24], requiring additional feedback and

manual labeling, e.g., users are supposed to actively provide their satisfaction. Besides, the evaluation of items relies heavily on manual labeling and professional knowledge [24]. Thus in practice, these methods are too expensive to implement in real-world recommenders. In addition, hard negative sampling is adopted to improve the negative sampling [7, 8, 34]. However, with fewer false positive samples, the hard negative instances also bring more false-negative samples. Our meta-learning method elegantly annotates the unlabeled instances based on the sparsely labeled instances.

Denosing Recommender Systems. One intuitive approach to reduce noise is to directly include more accurate feedback [22, 31], such as dwell time [32] and skip [29]). However, forcefully requiring additional feedback from users may harm user experiences. To address this problem, DenosingRec [28] achieves denosing recommendation for implicit feedback without any additional data. More specifically, they perform denosing on the false positive instances via truncating or reweighting the samples with larger loss. However, they only consider the positive feedback without further addressing the noise brought by negative sampling. Our work considers the sampled instances as possibly positive and negative, then achieve denosing data augmentation from both positive and negative perspectives.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a novel method that automatically annotated the unlabeled data and adjusted the false annotated labels. Such exploration not only addressed the unavoidable noise brought by widely used negative sampling but also improved the current denosing recommenders. Specifically, we proposed inverse learning from both loss and gradient perspectives. The first one was the Inverse Dual Loss that assumed the sampled data to be possibly positive or negative and automatically annotated them. If the positive loss was greater than the negative loss (difficult to label the data as positive), the Inverse Dual Loss would inversely assign more weights to the negative loss and vice versa. Since the Inverse Dual Loss depended heavily on the current training model and the quality of sampled data, we further proposed Inverse Gradient which made Inverse Dual Loss more robust by adjusting the gradient for those falsely annotated instances. Here we designed a meta-learning method splitting the training data into training-train data and training-test data, which pre-trained the model and determined the updated gradient, respectively.

As for future work, we plan to apply our inverse learning with more recommendation models as the backbones to further verify the generalization of our proposed methods.

ACKNOWLEDGEMENT

This work is supported in part by National Natural Science Foundation of China under 62272262 and U23B2030. This work is also supported in part by National Key Research and Development Program of China under 2022YFB3104702. This work is supported in part by a grant from the Guoqiang Institute, Tsinghua University under 2021GQG1005. This work is supported in part by Beijing National Research Center for Information Science and Technology. This work is also supported by Kuaishou.

REFERENCES

- [1] Jessa Bekker, Pieter Robberechts, and Jesse Davis. 2019. Beyond the selected completely at random assumption for learning from positive and unlabeled data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 71–85.
- [2] Jianxin Chang, Chen Gao, Yu Zheng, Yiqun Hui, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. 2021. Sequential Recommendation with Graph Neural Networks. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 378–387.
- [3] Chong Chen, Min Zhang, Yiqun Liu, and Shaoping Ma. 2019. Social attentional memory network: Modeling aspect-and friend-level differences in recommendation. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 177–185.
- [4] Chong Chen, Min Zhang, Weizhi Ma, Yiqun Liu, and Shaoping Ma. 2020. Efficient Non-Sampling Factorization Machines for Optimal Context-Aware Recommendation. In *Proceedings of The Web Conference 2020*. 2400–2410.
- [5] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhya, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 7–10.
- [6] Weiyu Cheng, Yanyan Shen, and Linpeng Huang. 2020. Adaptive factorization network: Learning adaptive-order feature interactions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 3609–3616.
- [7] Jingtao Ding, Yuhan Quan, Xiangnan He, Yong Li, and Depeng Jin. 2019. Reinforced Negative Sampling for Recommendation with Exposure Data. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 2230–2236. <https://doi.org/10.24963/ijcai.2019/309>
- [8] Jingtao Ding, Yuhan Quan, Quanming Yao, Yong Li, and Depeng Jin. 2020. Simplify and Robustify Negative Sampling for Implicit Collaborative Filtering. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 1094–1105.
- [9] Charles Elkan and Keith Noto. 2008. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 213–220.
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*. PMLR, 1126–1135.
- [11] Steve Fox, Kuldeep Karnawat, Mark Mydland, Susan Dumais, and Thomas White. 2005. Evaluating implicit measures to improve web search. *ACM Transactions on Information Systems (TOIS)* 23, 2 (2005), 147–168.
- [12] Chen Gao, Yu Zheng, Wenjie Wang, Fuli Feng, Xiangnan He, and Yong Li. 2024. Causal inference in recommender systems: A survey and future directions. *ACM Transactions on Information Systems* (2024).
- [13] Asela Gunawardana and Guy Shani. 2015. Evaluating Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, 265–308. https://doi.org/10.1007/978-1-4899-7637-6_8
- [14] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [15] Dietmar Jannach, Lukas Lerche, and Markus Zanker. 2018. Recommending based on implicit feedback. In *Social Information Access*. Springer, 510–569.
- [16] Youngho Kim, Ahmed Hassan, Ryen W White, and Imed Zitouni. 2014. Modeling dwell time to predict click-level satisfaction. In *Proceedings of the 7th ACM international conference on Web search and data mining*. 193–202.
- [17] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [18] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835* (2017).
- [19] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1754–1763.
- [20] Dawen Liang, Laurent Charlin, James McInerney, and David M Blei. 2016. Modeling user exposure in recommendation. In *Proceedings of the 25th international conference on World Wide Web*. 951–961.
- [21] Feng Liang, Weike Pan, and Zhong Ming. 2021. Fedrec++: Lossless federated recommendation with explicit feedback. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 4224–4231.
- [22] Chao Liu, Ryen W White, and Susan Dumais. 2010. Understanding web browsing behaviors through Weibull analysis of dwell time. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. 379–386.
- [23] Jiahui Liu, Peter Dolan, and Elin Rønby Pedersen. 2010. Personalized news recommendation based on click behavior. In *Proceedings of the 15th international conference on Intelligent user interfaces*. 31–40.
- [24] Hongyu Lu, Min Zhang, and Shaoping Ma. 2018. Between clicks and satisfaction: Study on multi-phase user preferences and satisfaction for online news reading. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 435–444.
- [25] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [26] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. Introduction to recommender systems handbook. In *Recommender systems handbook*. Springer, 1–35.
- [27] Yuta Saito, Suguru Yaginuma, Yuta Nishino, Hayato Sakata, and Kazuhide Nakata. 2020. Unbiased Recommender Learning from Missing-Not-At-Random Implicit Feedback. (2020), 501–509.
- [28] Wenjie Wang, Fuli Feng, Xiangnan He, Liqiang Nie, and Tat-Seng Chua. 2021. Denoising Implicit Feedback for Recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. ACM. <https://doi.org/10.1145/3437963.3441800>
- [29] Hongyi Wen, Longqi Yang, and Deborah Estrin. 2019. Leveraging post-click feedback for content recommendations. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 278–286.
- [30] Hong-Jian Xue, Xinyu Dai, Jianbing Zhang, Shujian Huang, and Jiajun Chen. 2017. Deep Matrix Factorization Models for Recommender Systems.. In *IJCAI*, Vol. 17. Melbourne, Australia, 3203–3209.
- [31] Byoungju Yang, Sangkeun Lee, Sungchan Park, and Sang-goo Lee. 2012. Exploiting various implicit feedback for collaborative filtering. In *Proceedings of the 21st International Conference on World Wide Web*. 639–640.
- [32] Xing Yi, Liangjie Hong, Erheng Zhong, Nanthan Nan Liu, and Suju Rajan. 2014. Beyond clicks: dwell time for personalization. In *Proceedings of the 8th ACM Conference on Recommender systems*. 113–120.
- [33] An Zhang, Jingnan Zheng, Xiang Wang, Yancheng Yuan, and Tat-Seng Chua. 2023. Invariant Collaborative Filtering to Popularity Distribution Shift. In *WWW*.
- [34] Weinan Zhang, Tianqi Chen, Jun Wang, and Yong Yu. 2013. Optimizing Top-n Collaborative Filtering via Dynamic Negative Item Sampling (*SIGIR '13*). Association for Computing Machinery, New York, NY, USA, 785–788. <https://doi.org/10.1145/2484028.2484126>
- [35] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. In *AAAI* 5941–5948.
- [36] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *KDD*. 1059–1068.