

Enhanced Bayesian Personalized Ranking for Robust Hard Negative Sampling in Recommender Systems

Kexin Shi
The Hong Kong University of Science
and Technology
kshiaf@connect.ust.hk

Jing Zhang
The Hong Kong University of Science
and Technology
jzhanggy@connect.ust.hk

Linjiajie Fang
The Hong Kong University of Science
and Technology
lfangad@connect.ust.hk

Wenjia Wang
The Hong Kong University of Science
and Technology (Guangzhou)
wenjiawang@hkust-gz.edu.cn

Bingyi Jing
Southern University of Science
and Technology
jingby@sustech.edu.cn

ABSTRACT

In implicit collaborative filtering, hard negative mining techniques are developed to accelerate and enhance the recommendation model learning. However, the inadvertent selection of false negatives remains a major concern in hard negative sampling, as these false negatives can provide incorrect information and mislead the model learning. To date, only a small number of studies have been committed to solve the false negative problem, primarily focusing on designing sophisticated sampling algorithms to filter false negatives. In contrast, this paper shifts its focus to refining the loss function. We find that the original Bayesian Personalized Ranking (BPR), initially designed for uniform negative sampling, is inadequate in adapting to hard sampling scenarios. Hence, we introduce an enhanced Bayesian Personalized Ranking objective, named as Hard-BPR, which is specifically crafted for dynamic hard negative sampling to mitigate the influence of false negatives. This method is simple yet efficient for real-world deployment. Extensive experiments conducted on three real-world datasets demonstrate the effectiveness and robustness of our approach, along with the enhanced ability to distinguish false negatives.

CCS CONCEPTS

• Information systems → Collaborative filtering; • Computing methodologies → Learning from implicit feedback.

KEYWORDS

Recommender Systems; Collaborative Filtering; Negative Sampling; False Negatives

ACM Reference Format:

Kexin Shi, Jing Zhang, Linjiajie Fang, Wenjia Wang, and Bingyi Jing. 2018. Enhanced Bayesian Personalized Ranking for Robust Hard Negative Sampling in Recommender Systems. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Implicit collaborative filtering (CF) in recommender systems is to leverage users' implicit feedback, such as clicks, purchases or viewing history, to improve user experiences by offering personalized content recommendations [13, 21, 23, 25, 26]. Prior efforts in model architecture design, such as matrix factorization (MF) [7, 17, 19, 20, 24, 30] and graph neural network (GNN) based methods [11, 14, 33, 35, 36], have effectively improved personalized recommendation accuracy. Nevertheless, a gap remains in the data domain, specifically in addressing the challenge presented by the lack of explicit negative feedback during recommendation model training. To address this, Bayesian Personalized Ranking (BPR) [29] formulates the training objective as a pairwise ranking task, which aims to maximize the probability of user preferring observed positive items than randomly sampled negative items. In BPR, the logistic sigmoid function is used to model the individual preference probability, which almost remains a fixed paradigm in subsequent studies. However, random sampling to generate negative samples presents low convergence due to gradients vanishing under the BPR loss. Subsequently, hard negative sampling methods are developed to adaptively choose challenging items as negatives, compelling the model to achieve faster and more effective learning [3, 8, 15, 27, 28, 31, 37–39]. Various hard mining techniques, such as dynamic negative sampling [38], Generative Adversarial Network (GAN) based sampling [27, 34] and importance sampling based hard mining [4, 22], emerge. Among hard sampling methods focusing on sampling progress, many adopt the BPR loss for model optimization. Conversely, for hard sampling methods that improve loss function, it is prevalent to apply importance sampling to the BPR loss but keep the core BPR unchanged. The question is: *Does the original BPR effectively accommodate hard sampling scenarios?*

In practice, the application of hard sampling in optimizing the recommendation model under the original BPR objective has been reported to potentially result in overfitting [31, 40]. This is because

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

false negatives, inadvertently selected during hard negative sampling, contribute significant gradients to model updating and potentially mislead the learning direction. To address the false negative problem and improve model's robustness, SRNS [6] introduces a variance-based sampling strategy to identify real hard negatives and GDNS [40] considers the expectation gain between two training iterations as an indicator. Both methods concentrate on crafting sophisticated sampling processes to mitigate the risk of sampling false negatives. However, achieving this task solely through a single metric, such as variance or expectation gain, is challenging.

According to the theory that the effects of a specific sampling strategy on model learning can be equivalently achieved by adjusting the loss function [9], our focus shifts from designing intricate sampling process for filtering false negatives towards crafting the loss function. We introduce an enhanced BPR scheme, named Hard-BPR, to mitigate the influence of false negatives in hard sampling. Compared to BPR, only the function to estimate the individual preference probability is modified in Hard-BPR, where three coefficients are added to lower the gradient magnitudes of excessive hard samples when applying stochastic gradient descent. The framework is straightforward, employing the hard negative sampling method DNS [38] to sample hard negatives and utilizing Hard-BPR as the optimization objective. Since the negative sampling process, the primary time-consuming element, remains DNS, our algorithm achieves superior time efficiency compared to other hard sampling approaches. Moreover, a recent study discloses that employing DNS under BPR is an exact estimator of One-way Partial AUC (OPAUC) [32]. We show that equipping DNS under the proposed Hard-BPR also serves as an exact OPAUC estimator. We conduct experiments on three real-world datasets, evaluating the effects of Hard-BPR comprehensively. Our results illustrate its enhanced ability to distinguish false negatives from real hard negatives, along with its effectiveness and efficiency in recommendation model training. In addition, a parameter study on three critical coefficients in Hard-BPR reveals that only two coefficients require fine-tuning, providing valuable guidance for algorithm deployment.

In summary, the contributions of this paper are:

- We are, to the best of our knowledge, the first to address the false negative problem in implicit CF from the perspective of BPR redesign.
- We introduce an optimization criterion, Hard-BPR, an enhanced Bayesian Personalized Ranking for dynamic hard negative sampling (DNS) in recommender systems.
- Our experiments empirically show the proposed method's robustness, efficiency and effectiveness. Moreover, parameter study offers valuable insights for its implementation.

2 PRELIMINARIES

In this section, we present the necessary background about the derivation of BPR in implicit collaborative filtering and dynamic hard negative sampling (DNS).

2.1 Bayesian Personalized Ranking (BPR)

In implicit collaborative filtering, only the implicit feedback $\mathcal{S} \subseteq \mathcal{U} \times \mathcal{I}$ is observed (such as clicks, views, or purchases), where \mathcal{U} denotes the user set and \mathcal{I} the item set. The user-item interaction pairs in implicit feedback $(u, i) \in \mathcal{S}$ are positive observations because they generally indicate users' preferences. The remaining unobserved user-item combinations $(\mathcal{U} \times \mathcal{I}) \setminus \mathcal{S}$ consist of actual negative feedback (indicating dislikes) and the potential positive feedback (implying future purchase interests).

For each user u , we define $\mathcal{I}^+(u) := \{i : (i, u) \in \mathcal{S}\}$. The set $\mathcal{I}^+(u)$ contains all items that interact with user u . BPR assumes that user u prefers item $i \in \mathcal{I}^+(u)$ over $j \in \mathcal{I} \setminus \mathcal{I}^+(u)$ (symbolized as $i >_u j$). The objective of BPR is to learn a scoring function $f(\cdot|\Theta) : \mathcal{I} \times \mathcal{U} \rightarrow \mathbb{R}$ to capture users' above pairwise preferences, where $f(i|u, \Theta)$ is the score of item i given by user u and Θ is the model parameter. We denote $\mathcal{D}_S = \{(u, i, j) \mid i \in \mathcal{I}^+(u), j \in \mathcal{I} \setminus \mathcal{I}^+(u)\}$. The likelihood function of observing all user pairwise preferences in \mathcal{D}_S given the model parameters Θ is as follows:

$$\prod_{(u, i, j) \in \mathcal{D}_S} P(i >_u j | \Theta), \quad (1)$$

where $P(i >_u j | \Theta)$ is the probability that user u prefers item i over item j . The individual preference probability $P(i >_u j | \Theta)$ is modeled as:

$$P(i >_u j | \Theta) = \sigma(\hat{x}_{uij}), \quad (2)$$

where \hat{x}_{uij} represents the estimated preference difference between item i and item j , which can be calculated as $f(i|u, \Theta) - f(j|u, \Theta)$, and σ is the logistic sigmoid function:

$$\sigma(x) := \frac{1}{1 + e^{-x}}. \quad (3)$$

The logistic sigmoid function σ is used to transform the estimated preference difference between two items into a probability between 0 and 1. This non-linear transformation can help capture the complex nature of user preferences in recommender systems.

Then, the BPR loss is derived as the negative log-likelihood of the pairwise preferences in \mathcal{D}_S :

$$\mathcal{L}_{\text{BPR}} = - \sum_{(u, i, j) \in \mathcal{D}_S} \ln \sigma(\hat{x}_{uij}), \quad (4)$$

In the context of optimizing the Bayesian Personalized Ranking (BPR) objective, Stochastic Gradient Descent (SGD) is employed, wherein triplets (u, i, j) are uniformly drawn from \mathcal{D}_S to facilitate the updating of model parameters. Specifically, for each positive user-item interaction (u, i) , a negative item j is randomly selected from the set $\mathcal{I} \setminus \mathcal{I}^+(u)$. The gradient of the BPR loss with respect to a model parameter $\theta \in \Theta$ is as follows:

$$\frac{\partial \mathcal{L}_{\text{BPR}}}{\partial \theta} = - \sum_{(u, i, j) \in \mathcal{D}_S} (1 - \sigma(\hat{x}_{uij})) \frac{\partial (\hat{x}_{uij})}{\partial \theta}, \quad (5)$$

The impact of each triplet (u, i, j) on the model's update is quantified by the gradient magnitude $\Delta_\sigma(\hat{x}_{uij})$, expressed as:

$$\Delta_\sigma(\hat{x}_{uij}) := (1 - \sigma(\hat{x}_{uij})) = (1 - P(i >_u j | \Theta)) \quad (6)$$

Rendle and Freudenthaler [28] suggests that uniform sampling strategy results in slow convergence of the model because most of

samples yield small gradient magnitudes. Specifically, the scoring function $f(\cdot|\Theta)$ generally gives a larger score to positive pair (u, i) than pair (u, j) and the pairwise preference probability $P(i >_u j|\Theta)$ is close to 1.

2.2 Dynamic Hard Negative Sampling

To accelerate the model's convergence and enhance its ability to differentiate between positive and negative items, various non-uniform sampling methods have been proposed. Among these, hard negative samplers demonstrate notable efficiency. Hard negative sampling is to choose negative instances that are challenging for model to distinguish from positive examples, thereby enhancing the model's capacity to learn fine-grained features and improving its discriminative ability.

Dynamic Negative Sampling (DNS) [38], supported by various studies [6, 15, 31], stands as an advanced hard negative sampling method in recommender systems. Shi et al. [32] has proven that the BPR loss equipped with DNS serves as an exact estimator of OPAUC. DNS involves a two-step process. This two-step process begins by selecting $H(\geq 1)$ random items from the set $\mathcal{I} \setminus \mathcal{I}^+(u)$ for each positive user-item pair (u, i) , creating a negative candidate pool $\mathcal{P} = \{j_1, j_2, \dots, j_H\}$. Then, the item from \mathcal{P} that is assigned the highest score by the scoring function $f(\cdot|u, \Theta)$ is selected as the hard negative sample. This is represented as:

$$j = \arg \max_{k \in \mathcal{P}} f(\cdot|u, \Theta), \quad (7)$$

The selection of a hard negative item results in a relatively large gradient magnitude, $\Delta_\sigma(\hat{x}_{uij})$, for the triplet (u, i, j) since the estimated preference difference between the positive item i and the hard negative item j is small. Consequently, this leads to a more significant contribution to the model's updating process. By efficiently choosing hard negative samples, DNS accelerates training processes and improves prediction accuracy.

3 METHOD: HARD-BPR

In this section, we introduce the false negative problem for hard negative sampling in implicit CF. Subsequently, we propose an enhanced Bayesian Personalized Ranking, termed as Hard-BPR, to address the challenges posed by false negatives.

3.1 Challenges of False Negatives

A critical challenge in hard negative sampling within recommender systems is the risk of inadvertently including false negatives, which may induce overfitting, potentially compromising the model's generalization ability. False negatives in recommender systems are items that the user has not interacted with, but the user would have liked or found interesting. As the hardness of the negatives sampled increases, the probability of encountering false negatives correspondingly increases [1, 2, 31]. To be specific, let $H(j)$ denote the hardness of a negative item j , and $P_{\text{FN}}(j)$ be the probability that item j is a false negative. As $H(j)$ increases, $P_{\text{FN}}(j)$ also rises. This relationship poses a significant limitation to hard negative sampling methods, as it is challenging to distinguish between real hard negatives and false negatives.

The corruption of false negatives not only reduces the accuracy of personalized recommendations but also exacerbates overfitting




Figure 1: (a) The function curves of $\sigma(\cdot)$ and $g(\cdot)$. (b) The function curves of $\Delta_\sigma(\cdot)$ and $\Delta_g(\cdot)$.

during the model training [31, 40]. Avoiding false negatives in hard negative sampling is crucial, yet only a limited number of studies focus on developing robust hard sampling methods.

3.2 Hard-BPR

The Equivalence Theory. The foundational principle guiding our methodology is derived from the equivalence theory between loss functions and sampling strategies. This theory indicates that the effects induced by a specific sampling strategy on model learning can be equivalently achieved through modifying the loss function [9].

In implicit CF task, previous works focus on designing complex and sophisticated sampling algorithms to avoid false negatives in hard negative sampling. The equivalence theory suggests that an alternate loss function, when optimized with a simple sampling method (e.g., DNS), can yield an expected gradient that is similar to the one obtained with a complex false-negative-exclusive hard sampling approach. The equivalence theory can be mathematically expressed in terms of the expected gradient of the loss function, as follows:

$$E_{(u,i,j) \sim f_{\text{complex}}} [\nabla L_{\text{BPR}}(u, i, j)] = E_{(u,i,j) \sim f_{\text{DNS}}} [\nabla L_{\text{mod}}(u, i, j)]. \quad (8)$$

Here, $E_{(u,i,j) \sim f_{\text{complex}}} [\nabla L_{\text{BPR}}(u, i, j)]$ denotes the expected gradient of BPR loss under the sophisticatedly designed false-negative-exclusive hard sampling, and $E_{(u,i,j) \sim f_{\text{DNS}}} [\nabla L_{\text{mod}}(u, i, j)]$ represents the expected gradient of a modified loss function under DNS.

This theory suggests a new perspective on addressing the challenges of hard negative sampling in implicit CF. Instead of creating complex sampling algorithms to discern and exclude false negatives, the equivalence theory encourages a shift towards revising the loss function.

Limitations of BPR in Hard Negative Sampling. As mentioned in Section 2.1, BPR employs the logistic sigmoid function $\sigma(\cdot)$ to calculate the individual preference probability $P(i >_u j|\Theta)$. When employing BPR, $P(i >_u j|\Theta)$ approaches 0 if \hat{x}_{uij} is small (e.g., ≤ -4). In this situation, the gradient magnitude $\Delta_\sigma(\hat{x}_{uij})$ is near 1, and thus the triplet (u, i, j) can provide useful information for model optimization via SGD. Conversely, a large \hat{x}_{uij} results in a negligible $\Delta_\sigma(\hat{x}_{uij})$, rendering the triplet (u, i, j) useless in parameter updating.

Rendle and Freudenthaler [28] states that \hat{x}_{uij} is generally a large value when item j is sampled uniformly from $\mathcal{I} \setminus \mathcal{I}^+(u)$, and

only a few (u, i, j) triplets may yield small \hat{x}_{uij} s. In this case, the BPR loss is suitable, since it can assign large gradient magnitudes $\Delta_\sigma(\hat{x}_{uij})$ to these few triplets, thereby expediting the model convergence.

However, when adopting hard negative sampling strategies such as DNS for negative selection, the \hat{x}_{uij} of triplet (u, i, j) is typically small, and items j yielding smaller values are more likely to be false negatives, as discussed in Section 3.1. In this case, the logistic sigmoid function in BPR loss is inappropriate because it easily generates a large $\Delta_\sigma(\hat{x}_{uij})$ for false negatives and misleading the parameter updating.

The Proposed Method. The proposed framework is shown in Algorithm 1. To alleviate the negative effects of false negatives on model learning, we propose Hard-BPR, an enhanced Bayesian Personalized Ranking. This is achieved by substituting $\sigma(\cdot)$ with $g(\cdot)$ for the estimation of individual preference probability. The function $g(\cdot)$ is a transformation of $\sigma(\cdot)$ and its mathematical expression is as below:

$$g(x) = \frac{1}{1+a} [\sigma(cx+b) + a], \quad (9)$$

where a, b and c are constant coefficients with $a \geq 0$ and $c > 0$. The Hard-BPR loss is defined as:

$$L_{\text{Hard-BPR}} = - \sum_{\substack{(u,i) \in \mathcal{S} \\ j \sim f_{\text{DNS}}}} \ln g(\hat{x}_{uij}). \quad (10)$$

Here $j \sim f_{\text{DNS}}$ represents that item j is sampled as a negative by hard negative sampling strategy DNS. The Hard-BPR loss proposed in this study is an exact estimator of OPAUC when equipped with DNS, because $-\ln g(\cdot)$ ($a \geq 0, c > 0$) is a convex, differentiable and monotonically decreasing function, satisfying the sufficient condition to be consistent for OPAUC maximization according to works [10, 32]. OPAUC denotes the partial area under the ROC curve, wherein the false positive rate (FPR) is constrained within the range $[0, \beta]$ ($0 < \beta \leq 1$). According to Shi et al. [32], optimizing OPAUC can lead to improved Top-K recommendation performance when the parameter β tuned appropriately. This optimization strategy allows for emphasizing the ranking of top-ranked items instead of the whole ranking list.

For the purpose of model optimization under SGD, the gradient of Hard-BPR loss is calculated as:

$$\begin{aligned} \frac{\partial L_{\text{Hard-BPR}}}{\partial \theta} &= - \sum_{\substack{(u,i) \in \mathcal{S} \\ j \sim f_{\text{DNS}}}} \Delta_g(\hat{x}_{uij}) \frac{\partial (\hat{x}_{uij})}{\partial \theta}, \\ \text{with } \Delta_g(\hat{x}_{uij}) &= \frac{c\sigma(c\hat{x}_{uij}+b)(1-\sigma(c\hat{x}_{uij}+b))}{\sigma(c\hat{x}_{uij}+b)+a}. \end{aligned} \quad (11)$$

Properties. The coefficient a in $g(\cdot)$ is the most critical coefficient, which is essential for converting the magnitude gradient function from a monotonically decreasing pattern to unimodal symmetry, as illustrated in Figure 1(a). $\Delta_g(\cdot)$ is bell-shaped under the condition $a > 0$. Additionally, the coefficients b and c are for $\Delta_g(\cdot)$ translation and stretching. The function $\Delta_g(\cdot)$ derived from our new individual preference probability estimator $g(\cdot)$ has following properties:

- The limit of $\Delta_g(x)$ as x approaches either positive or negative infinity is equal to 0.
- $\Delta_g(x)$ is unimodal and symmetric with respect to the maximum point located at x (see Appendix for more details).
- $\Delta_g(x)$'s maximum point x_{\max} and maximum value $\Delta_{g,\max}$ are calculated as:

$$x_{\max} = \frac{-b + \ln(\frac{\sqrt{a}}{\sqrt{1+a}})}{c}, \quad (12)$$

$$\Delta_{g,\max} = \frac{\sqrt{1+a}c}{2\sqrt{a} + 2a^{3/2} + \sqrt{1+a} + 2a\sqrt{1+a}}. \quad (13)$$

Mechanism Analysis. We analyze the resilience of Hard-BPR to false negatives in hard negative sampling. We illustrate this by comparing the functions $\sigma(\cdot)$ and $g(\cdot)$ in Figure 1(a), along with their respective gradient magnitudes $\Delta_\sigma(\cdot)$ and $\Delta_g(\cdot)$ in Figure 1(b). For $g(\cdot)$, we set parameters a, b and c to 1, -1 and 0.8.

The foundational principle of designing $g(\cdot)$ is to moderate the weight assigned to excessively hard negatives in model learning. Specifically, if a negative item j chosen by DNS for the user-item pair (u, i) is excessively hard (characterized by a very small \hat{x}_{uij}), we avoid modeling $P(i >_u j | \Theta)$ as nearly zero, as is done with $\sigma(\cdot)$. Instead, using $g(\cdot)$, we assign a small positive value to $P(i >_u j | \Theta)$. In this way, the gradient magnitude $\Delta_g(u, i, j)$ is reduced to a relatively small value, thereby decreasing its influence on model updating in cases where the hard negative instance is a false negative.

The Hard-BPR loss allows more flexibility than the BPR loss. The gradient magnitude curve derived from $g(\cdot)$ is single-peaked, by adjusting the constants a, b and c to change the position of peak value, the model can focus on learning from a specific hardness level of selected negatives. Note that setting a, b and c to $(0, 0, 1)$ reduces the Hard-BPR loss to the original BPR loss.

Algorithm 1 Optimizing models with Hard-BPR

Input: Training set $\mathcal{S} = \{(u, i)\}$, the scoring function $f(\cdot | \Theta)$ with learnable parameter Θ , the size of negative candidate pool H , constant coefficients a, b and c in the Hard-BPR loss.

```

1: for  $t = 1, 2, \dots, T$  do
2:   Sample a mini-batch  $\mathcal{S}_{\text{batch}} \in \mathcal{S}$ 
3:   for each  $(u, i) \in \mathcal{S}_{\text{batch}}$  do
4:     Randomly sample  $H$  items from  $\mathcal{I} \setminus \mathcal{I}_u$  to form the
5:     negative candidate pool  $\mathcal{P}$  for user  $u$ ;
6:     Choose the hard negative item  $j$  by DNS;
7:     Update  $L_{\text{Hard-BPR}}$ .
8:   end for
9:   Optimize  $\Theta$  based on  $L_{\text{Hard-BPR}}$  with SGD algorithm.
10: end for
```

3.3 Complexity analysis

Time Complexity. We conduct a time complexity comparison between the proposed algorithm and a set of state-of-the-art negative sampling approaches, as outlined in Table 1. The proposed method requires lower time complexity among the considered baselines.

Table 1: Comparison of the proposed method and baselines, where \mathcal{S} is the training set containing user-item interactions, \mathcal{U} is the user set, \mathcal{I} is the item set, H_1 and H_2 denote the negative candidate pool size, T is the time to compute user-item score and T_Q is the sampling time for a proposal distribution in AdaSIR.

Learning Objective		Time Complexity	Robustness to FN
RNS	BPR	$O(\mathcal{S} T)$	\times
AdvIR	adversarial loss	$O(\mathcal{S} H_1T)$	\times
IRGAN	adversarial loss	$O(\mathcal{S} \mathcal{I} T)$	\times
DNS	BPR	$O(\mathcal{S} H_1T)$	\times
SRNS	BPR	$O(\mathcal{S} (H_1 + H_2)T)$	\checkmark
GDNS	group-wise ranking loss	$O(\mathcal{U} (2H_1T + H_1^2))$	\checkmark
MixGCF	BPR	$O(\mathcal{S} H_1T)$	\times
AdaSIR	importance sampling BPR	$O(\mathcal{U} H_1(T_Q + T) + \mathcal{S} H_2T)$	\times
Ours	Hard-BPR	$O(\mathcal{S} H_1T)$	\checkmark

Table 2: Statistics of datasets.

Dataset	#User	#Item	#Train	#Val	#Test	Density
Taobao	22,976	29,149	351,444	43,930	43,931	0.00066
Tmall	10,000	14,965	369,520	46,189	46,190	0.00309
Gowalla	29,858	40,981	821,896	102,737	102,737	0.00084

For each training epoch, the proposed method takes a computational cost of $O(|\mathcal{S}|HT)$, where $|\mathcal{S}|$ denotes the number of positive user-item observations in set \mathcal{S} , H represents the number of negative candidates selected for each positive user-item pair, and T denotes the runtime for user-item score computation. The time complexity of the proposed method is the same with DNS, since only the function for estimating pairwise preferences has been modified. Both DNS and our method stand out as two of the most efficient algorithms in hard negative mining. In comparison, other methods such as IRGAN, SRNS, GDNS, and AdaSIR, which design specific sampling procedures, demand relatively higher computational resources.

Space Complexity. In addition to model parameters, the space complexity of the proposed method comes from storing negative candidates for each positive user-item pair, incurring a space complexity of $O(|\mathcal{S}|H)$.

4 EXPERIMENTS

In this section, we evaluate the performance of our approach on three real-world datasets, with MF and LightGCN as the base recommendation models. Additionally, we study the proposed method's ability to discern false negatives and investigate its behavior under different parameter settings.

4.1 Experimental Settings

Datasets. We assess our method on three real-world datasets: Taobao¹, Tmall² and Gowalla. Taobao and Tmall, sourced from e-commerce platforms, capture diverse user behaviors such as clicks, cart additions, and purchases. To address data sparsity, varied interaction behaviors are considered positive labels in the prediction task. The 10-core setting is applied by randomly selecting users with a

minimum of 10 interaction records. Gowalla consists of check-in records with user locations and we directly utilize the pre-processed dataset in work [36] for model performance evaluation. For each dataset, we construct training, validation, and test sets by sorting interactions based on timestamps and retaining the latest 10% of records for testing, with the remaining interactions divided into an 80/10 ratio for training and validation (see details in Table 2).

Metrics. The main aim of personalized recommendation systems is to provide a ranked list containing the Top- K items with the highest scores for each user. This study uses two widely-adopted metrics, Recall@ K and NDCG@ K , to evaluate the model's ability to understand users' preferences. The work by Shi et al. [32] reports the connection between K in Top- K metrics and the size of the negative candidate pool (H) in DNS, which suggests that the lower the K , the larger the H when the training curve reaches its maximum values. To control memory usage related to storing negative candidates, we set K to 50 as done in previous works [4, 22, 32].

Baselines. A variety of negative samplers in recommender systems are chosen as baselines to assess the effectiveness of the proposed method. The selected baselines are detailed as follows:

- **RNS** [29] selects negative instances by a uniform distribution, which is widely employed in implicit CF studies.
- **AdvIR** [27] incorporates Generative Adversarial Nets to produce challenging negative instances, and utilizes virtual adversarial training to enhance model performance.
- **IRGAN** [34] is also a GAN-based hard negative sampler and optimizes model parameters through a min-max game between the generative retrieval network and the discriminative network.
- **DNS** [38] first randomly samples a set of negative candidates for each positive instance, then chooses the highest-scored candidate as the hard negative sample.
- **SRNS** [6] considers the variance of item scores during training to avoid false negatives in the hard sampling procedure.
- **GDNS** [40] employs a gain-aware sampler to reduce the probability of introducing false negatives during hard negative sampling.

¹<https://tianchi.aliyun.com/dataset/dataDetail?dataId=649>

²<https://tianchi.aliyun.com/dataset/dataDetail?dataId=121045>

- **MixGCF** [15] introduces hop-mixing and positive mixing techniques to synthesized hard negative samples for GNN-based recommender models.
- **AdaSIR** [4] reuses informative negatives and adopts importance resampling to approximate the softmax distribution for effectively selecting hard negatives. AdaSIR(U) and AdaSIR(P) represent sampling negative candidate using a uniform distribution and a popularity distribution respectively.

Implementation Details. We perform experiments using PyTorch on a single Linux server featuring an AMD EPYC 7543 processor, 128GB RAM, and an NVIDIA GeForce RTX 4090 GPU. Adam [16] is employed as the optimizer, and a fixed batch size of 2048 is utilized across all experiments. Without loss of generality, our underlying recommendation models include matrix factorization [18] (MF) and LightGCN [12], where LightGCN is a state-of-the-art GNN-based framework. For the proposed method, we first conduct grid search to determine optimal hyperparameters for DNS. Subsequently, Bayesian optimization is employed to explore coefficients a, b, c in $L_{\text{Hard-BPR}}$. User and item embedding dimensions are tuned from the set $\{8, 16, 32, 64\}$, while the negative candidate pool size for our method is searched in $\{8, 16, 32, 64\}$. L_2 regularization is adjusted within $\{0, 0.1, 0.01, 0.001, 0.0001\}$, and the learning rate is explored within $\{0.0005, 0.001, 0.005, 0.01\}$. Coefficients a, b and c in $L_{\text{Hard-BPR}}$ are uniformly searched in the ranges $[0, 10]$, $[-10, 10]$, and $(0, 5]$, respectively. Hyperparameters are tuned to optimize Recall@50 on the validation set, and final results are reported on the test set. To ensure a fair comparison, all baselines are finely tuned based on the best Recall@50 achieved on the validation set. Additionally, two GAN-based samplers are trained from a uniformly pretrained model.

4.2 Performance comparison

To compare different hard mining methods, we conduct experiments on Taobao, Tmall and Gowalla, employing MF or LightGCN as the base recommendation model, respectively. The results, presented in Table 3, yield the following observations:

- Across both recommendation models, MF and LightGCN, our method consistently outperforms all baselines on the three datasets in terms of Recall@50 and NDCG@50. This highlights the effectiveness of the proposed Hard-BPR loss in a hard negative sampling scenario. Besides, DNS acts as a competitive approach, achieving a top performance in two out of six comparisons among baselines.
- The average relative improvement of our method over the best baseline under MF is 7.40% for NDCG@50, surpassing the improvement under LightGCN, which is 3.07%. This difference arises because Hard-BPR facilitates MF to learn more fine-grained features, thereby unleashing the predictive potential inherent in this simple model.

Furthermore, we present the training curves of RNS, DNS, and our method under MF and LightGCN, shown in Figure 2. The key findings are outlined as follows:

- Compared to RNS, the hard negative sampling method DNS converges (or reaches its maximum value if overfitting occurs) faster and achieves a higher Recall@50. However, the

application of DNS in recommender systems may lead to overfitting, especially when LightGCN serves as the underlying model. As studied in previous works [31], overfitting may result from the inclusion of false negatives during hard negative sampling. In addition, overfitting is more severe under LightGCN compared to MF. This is because the optimal hyperparameter L_2 regularization under MF is larger than the one under LightGCN in experiments, thus preventing overfitting to some extent.

- By replacing $\sigma(\cdot)$ with $g(\cdot)$, overfitting is effectively eliminated as expected. It verifies that Hard-BPR exhibits the capacity to mitigate the influence of false negatives in hard mining process. Our method is thereby demonstrated as a robust hard negative sampling method in implicit CF.
- The training efficiency of our method is comparable to that of DNS, both of which exhibit the lowest time complexity among baselines as shown in Table 1. Notably, when MF acts as the base recommendation model, our method even converges faster than DNS. This suggests our method's efficiency, which facilitates its deployment in realistic scenarios.

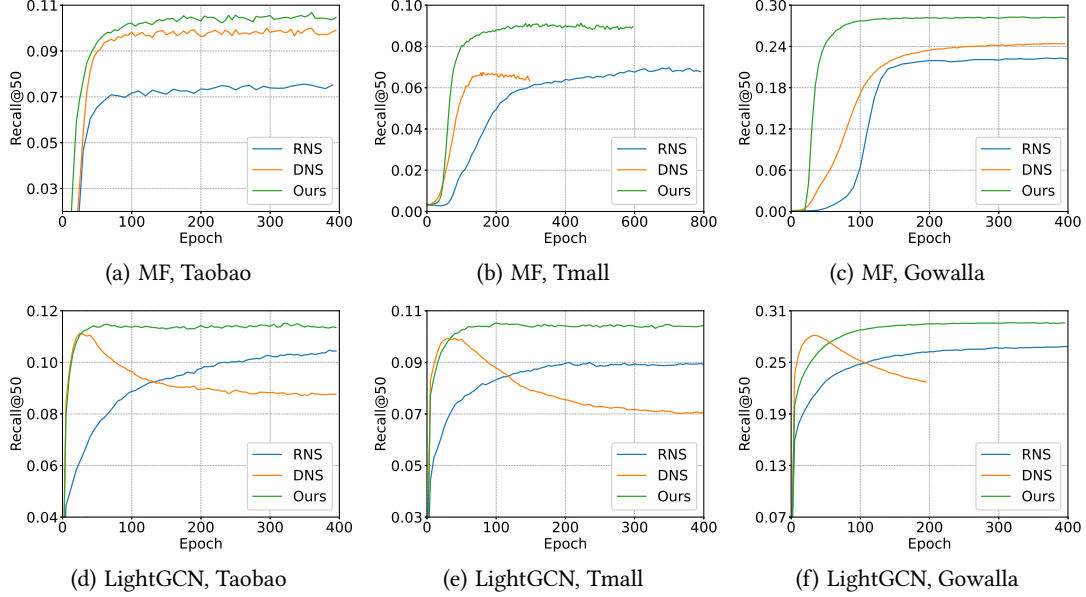
4.3 False negative analysis

False negatives generally indicate a user's potential preferences. Accurately identifying false negatives is critical for recommending items that the user would have liked or interacted with. In the preceding section, we validate that Hard-BPR effectively prevents the model from being misled by incorrect information resulting from false negatives. In this section, we evaluate the recommender system's capability to distinguish false negatives from true negatives. Experiments on Taobao employ MF and LightGCN as the base models, comparing three negative sampling methods: RNS, DNS and our method. For each user in Taobao, we treat their interacted items in the test set as false negatives in model training process and the remaining non-interacted items as true negatives. Models are saved at the point of achieving the best validation Recall@50. For each user, we compute his (or her) matching scores with true negative instances (yellow) and false negative instances (blue), plotting their distributions in Figure 3. To clearly manifest the discriminative capability of three negative sampling methods, we apply Gaussian kernel density estimation to estimate the probability distributions of true negative scores and false negative scores. We then calculate the Kullback-Leibler (KL) divergence between these two probability distributions, as presented in Table 4.

We find that the KL divergence under LightGCN averages 1.44278 across RNS, DNS and our method, surpassing that under MF, which is 1.36869. This observation indicates the superior ability of LightGCN to distinguish false negatives compared to MF. With MF as the base model, our method exhibits a 6.20% relative improvement in KL divergence over DNS, while with LightGCN, our method achieves a relative improvement of 10.79% over DNS. It verifies the capacity of Hard-BPR in false negative identification during hard mining process. Notably, when employing LightGCN, the KL divergence between the two distributions is merely 1.35639 under DNS, even smaller than that under RNS (1.46923). This discrepancy may be attributed to the adverse impact of false negatives on model

Table 3: Performance comparison. The best results are presented in bold and the second best results are underlined. The relative improvements of the best results compared to the second best results are listed in the last row.

	MF						LightGCN					
	Taobao		Tmall		Gowalla		Taobao		Tmall		Gowalla	
	Recall	NDCG	Recall	NDCG	Recall	NDCG	Recall	NDCG	Recall	NDCG	Recall	NDCG
RNS	0.0751	0.0268	0.0692	0.0333	0.2230	0.1418	0.1041	0.0378	0.0895	0.0476	0.2712	0.1804
AdvIR	0.0911	0.0327	0.0792	0.0426	0.2603	0.1620	0.1052	0.0382	0.0917	0.0487	0.2710	0.1801
IRGAN	0.0859	0.0314	0.0771	0.0415	0.2421	0.1508	0.1049	0.0380	0.0904	0.0480	0.2799	0.1847
DNS	<u>0.0999</u>	<u>0.0360</u>	0.0688	0.0320	0.2490	0.1582	0.1107	<u>0.0410</u>	0.0970	0.0504	<u>0.2848</u>	<u>0.1878</u>
SRNS	0.0968	0.0349	0.0801	0.0433	0.2588	0.1620	0.1086	0.0392	0.0967	0.0509	0.2819	0.1866
GDNS	0.0970	0.0352	<u>0.0816</u>	<u>0.0440</u>	0.2611	0.1638	0.1089	0.0393	0.0975	0.0516	0.2835	0.1865
MixGCF	-	-	-	-	-	-	<u>0.1110</u>	0.0403	0.0963	0.0508	0.2840	0.1872
AdaSIR(U)	0.0958	0.0336	0.0800	0.0429	<u>0.2711</u>	<u>0.1712</u>	0.1100	0.0401	0.1009	0.0530	0.2816	0.1857
AdaSIR(P)	0.0955	0.0333	0.0810	0.0438	0.2702	0.1700	0.1080	0.0389	<u>0.1016</u>	<u>0.0538</u>	0.2808	0.1853
Ours	0.1055	0.0386	0.0876	0.0480	0.2825	0.1813	0.1150	0.0419	0.1054	0.0556	0.2976	0.1947
	5.61%	7.22%	7.35%	9.09%	4.21%	5.90%	3.60%	2.20%	3.74%	3.35%	4.49%	3.67%

**Figure 2: The training curves of RNS, DNS and the proposed method on the test sets of three datasets. (a)(b)(c) With MF as the base model. (d)(e)(f) With LightGCN as the base model.****Table 4: The Kullback-Leibler divergence between the distributions of predicted true negative scores and predicted false negative scores.**

	MF			LightGCN		
	RNS	DNS	Ours	RNS	DNS	Ours
KL divergence	1.32216	1.35012	<u>1.43378</u>	1.46923	1.35639	<u>1.50273</u>

updating with DNS under LightGCN, which not only gives rise to

severe overfitting as shown in Figure 2(d)–(f) but also diminishes the model’s discriminative ability for false negatives.

4.4 Parameter study

While the implementation of our method alleviates the need for tuning sophisticated sampling process, there remains a necessity to finely adjust coefficients a , b and c ($a \in [0, +\infty)$, $b \in (-\infty, +\infty)$, and $c \in (0, +\infty)$) in Hard-BPR loss. In this section, we study the effect of these three coefficients on model performance to provide insights for fine-tuning the model. Experiment results on Tmall under MF are presented in Figure 4(a), Figure 4(b) and Figure 4(c).

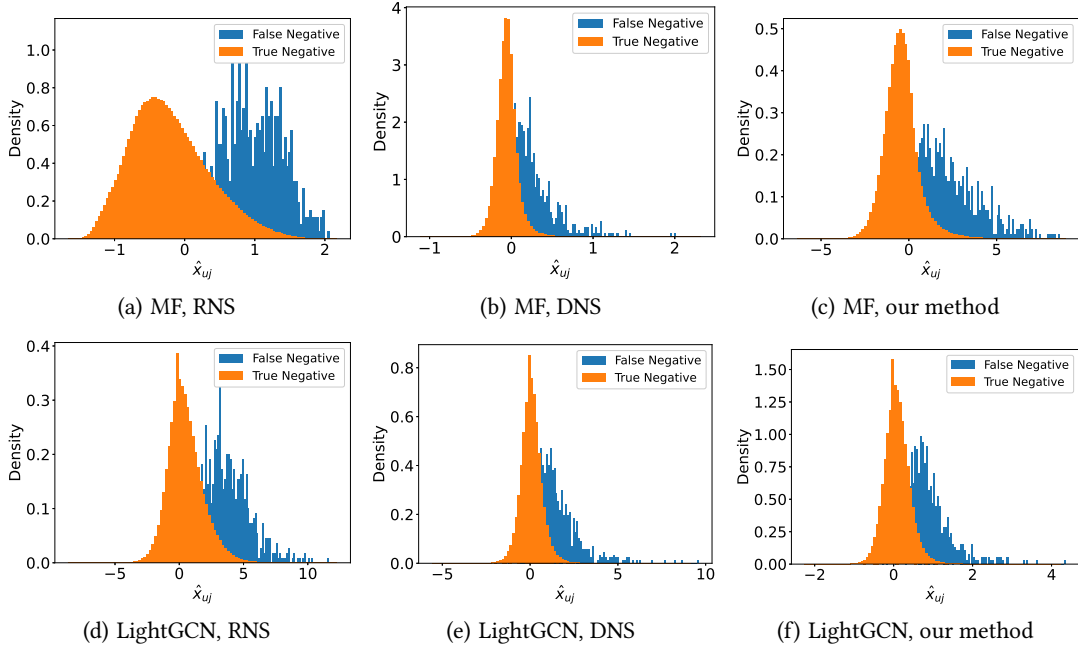


Figure 3: The distributions of predicted true negative scores and predicted false negative scores, employing different base recommendation models and different negative sampling methods (RNS, DNS and the proposed method).

The left column of Figures shows the curves of $\Delta_g(\cdot)/c$ under various settings (for gradient magnitude, dividing by a constant has no influence on model training under Adam).

In Figure 4(a), we set a and c as 1 and vary the b value to investigate the effect of Δ_g curve's position on model performance. We observe that the model performs optimally when b is 0 or 0.9, while suboptimal performance is observed when the Δ_g curve is positioned towards the far left or far right. Hence, ensuring the Δ_g curve concentrated within an appropriate user-item score range is critical for effective model training. In addition, compared to DNS method, our method with b as 0.9 exhibits faster convergence, higher Recall@50 and mitigated overfitting.

Both b and c affect the position of Δ_g curve along x-axis, while c also controls the curve's shape. In Figure 4(b), we fix the coefficient a to 0.1, and vary the curve from a sharp to a broader shape by reducing the coefficient c . Meantime, we correspondingly adjust the b to maintain the maximum point x_{\max} of curves as the same. We observe that a sharp Δ_g curve results in a bad performance in model effectiveness and efficiency, because it places excessive importance on a small range of \hat{x}_{uij} around x_{\max} for model learning. It leads to information loss from other hard negatives. Conversely, when the curve becomes boarder under a smaller c , more hard negatives can participate in parameter updating, thereby enhancing learning speed and model performance. Notably, excessive extension of the curve tails can also lead to performance deterioration. For example, at a c value of 0.4 or 0.24, the training curve initially peaks and then declines.

Figure 4(c) demonstrates that under various a settings, similar Δ_g curves can be obtained by adjusting b and c , leading to comparable performance. The coefficient a is critical to make the Δ_g curve bell-shaped, but it is not hard to tune. In practice, we can manually set a to a small value such as 1, while fine-tuning is primarily required for b and c .

5 RELATED WORK

Negative sampling approaches in implicit collaborative filtering are developed into several branches in recent years. The simplest way is to follow a static distribution such as uniform distribution [29] or popularity-based distribution [5], however, which may encounter challenges such as gradient vanishing or fail to adapt to dynamic changes in predicted user preferences. To address these limitations and enhance the quality of sampled negatives, various hard negative sampling strategies, including GAN-based samplers (e.g., IR-GAN [34], AdvIR [27]), DNS and its extensions, and softmax-based sampling approaches, have been proposed. GAN-based samplers leverage GANs to iteratively generate informative negative instances. DNS [38], a simple yet effective framework, involves random sampling of negative candidates for each user and subsequent selection of the highest-scored candidate for model updating. Extensions of DNS, such as MixGCF [15], SRNS [6], GDNS [40], and DNS(M, N) [32], have further contributed to the evolution of hard negative sampling strategies. Softmax-based samplers, such as AdaSIR [4] or Softmax-v(ρ, N) [32], assign resampling weights to each negative candidate. This strategy aims to align the negative sampling probability with the ideal softmax distribution, where greater weights




Figure 4: Parameter study of the proposed method. Experiments on Tmall under MF. The left column: curves of $\Delta_g(\cdot)/c$. The right column: training curves of our method.

are typically allocated to negative samples that are more informative. Moreover, Shi et al. [32] proves that the BPR loss combined with DNS is an exact estimator of OPAUC, and they extend DNS to $\text{DNS}(M, N)$ to suit various Top-K metrics. In our study, without applying important resampling, we directly modify the loss function to ensure the contribution of hard negatives on parameter updating.

For hard negative sampling, the inadvertent inclusion of false negatives poses a challenge on effective model learning. Cai et al. [2] formulates the false negative problem as learning from labeled data with inherent bias, and adopts a Coupled Estimation Technique to correct the bias. To avoid false negatives, SRNS develops a scored-based and variance-based sampling strategy, and GDNS proposes a gain-aware sampler to distinguish true hard negatives.

However, both SRNS and GDNS center their efforts on crafting sophisticated sampling progress to avoid the inclusion of false negatives, lacking a desired level of elegance. The investigation of overfitting related to false negatives in the study [31] provides valuable insights into our research.

6 CONCLUSION

This paper aims to address the false negative problem of hard negative sampling in implicit collaborative filtering. Inspired by the equivalence theory, our approach diverges from focusing on the sampling process design. Instead, we put efforts on loss function crafting, proposing an enhanced BPR objective, Hard-BPR, specifically tailored for dynamic hard negative sampling. Employing DNS under Hard-BPR is simple and efficient enough for real-world deployments. This approach not only demonstrates resilience to false negatives during hard sampling, but also exhibits a notable advantage in terms of time complexity. Moreover, equipping DNS with Hard-BPR serves as an exact estimator of OPAUC. Extensive experiments are conducted to assess the proposed method's robustness and effectiveness, and the parameter study is conducted to provide guidance for optimal tuning.

A APPENDIX

A.1 Derivation of $\Delta_g(\cdot)$ properties.

The first derivative of $\Delta_g(x)$ is calculated as:

$$\begin{aligned}
 \Delta'_g(x) &= \frac{c^2 \left(1 - \frac{1}{e^{-b-cx}+1}\right) e^{-b-cx}}{(e^{-b-cx}+1)^2 \left(a + \frac{1}{e^{-b-cx}+1}\right)} - \frac{c^2 e^{-b-cx}}{(e^{-b-cx}+1)^3 \left(a + \frac{1}{e^{-b-cx}+1}\right)} \\
 &\quad - \frac{c^2 e^{-b-cx} \left(1 - \frac{1}{e^{-b-cx}+1}\right)}{(e^{-b-cx}+1)^3 \left(a + \frac{1}{e^{-b-cx}+1}\right)^2} \\
 &= - \frac{c^2 e^{b+cx} \left(a \left(e^{2(b+cx)} - 1\right) + e^{2(b+cx)}\right)}{(e^{b+cx}+1)^2 (ae^{b+cx} + a + e^{b+cx})^2}
 \end{aligned} \tag{14}$$

By solving $\Delta'_g(x) = 0$, the critical point of $\Delta_g(x)$ is calculated as $x_c = \left(-b + \ln\left(\frac{\sqrt{a}}{\sqrt{1+a}}\right)\right)/c$. Since the graph of $\Delta_g(x) = 0$ ($a > 0$) is unimodal, x_c is the maximum point, termed as x_{\max} .

To prove the symmetry about the maximum point x_{\max} :

$$\begin{aligned}
 \Delta_g(x + x_{\max}) &= \frac{c \left(1 - \frac{1}{e^{\left(-c \left(\frac{\log\left(\frac{\sqrt{a}}{\sqrt{a+1}}\right) - b\right)}{c} + x \right) - b}} \right)}{e^{\left(-c \left(\frac{\log\left(\frac{\sqrt{a}}{\sqrt{a+1}}\right) - b\right)}{c} + x \right) - b} + 1} \right) \left(\frac{1}{e^{\left(-c \left(\frac{\log\left(\frac{\sqrt{a}}{\sqrt{a+1}}\right) - b\right)}{c} + x \right) - b}} + a \right)}{\left(\sqrt{\frac{a}{a+1}} e^{cx} + 1 \right) \left(\sqrt{\frac{a}{a+1}} a e^{cx} + \sqrt{\frac{a}{a+1}} e^{cx} + a \right)} \\
 &= \frac{\sqrt{\frac{a}{a+1}} c e^{cx}}{\left(\sqrt{\frac{a}{a+1}} e^{cx} + 1 \right) \left(\sqrt{\frac{a}{a+1}} a e^{cx} + \sqrt{\frac{a}{a+1}} e^{cx} + a \right)} \quad (15)
 \end{aligned}$$

$$\begin{aligned}
 \Delta_g(-x + x_{\max}) &= \frac{c \left(1 - \frac{1}{e^{\left(-c \left(\frac{\log\left(\frac{\sqrt{a}}{\sqrt{a+1}}\right) - b\right)}{c} - x \right) - b}} \right)}{e^{\left(-c \left(\frac{\log\left(\frac{\sqrt{a}}{\sqrt{a+1}}\right) - b\right)}{c} - x \right) - b} + 1} \right) \left(\frac{1}{e^{\left(-c \left(\frac{\log\left(\frac{\sqrt{a}}{\sqrt{a+1}}\right) - b\right)}{c} - x \right) - b}} + a \right)}{\left(\sqrt{\frac{a}{a+1}} + e^{cx} \right) \left(a \left(\sqrt{\frac{a}{a+1}} + e^{cx} \right) + \sqrt{\frac{a}{a+1}} \right)} \\
 &= \frac{\sqrt{\frac{a}{a+1}} c e^{cx}}{\left(\sqrt{\frac{a}{a+1}} + e^{cx} \right) \left(a \left(\sqrt{\frac{a}{a+1}} + e^{cx} \right) + \sqrt{\frac{a}{a+1}} \right)} \quad (16)
 \end{aligned}$$

Hence, we have $\Delta_g(x + x_{\max}) = \Delta_g(-x + x_{\max})$.

REFERENCES

- [1] Jessa Bekker and Jesse Davis. 2020. Learning from positive and unlabeled data: A survey. *Machine Learning* 109 (2020), 719–760.
- [2] Yinqiong Cai, Jiafeng Guo, Yixing Fan, Qingyao Ai, Ruqing Zhang, and Xueqi Cheng. 2022. Hard Negatives or False Negatives: Correcting Pooling Bias in Training Neural Ranking Models. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 118–127.
- [3] Chong Chen, Weizhi Ma, Min Zhang, Chenyang Wang, Yiqun Liu, and Shaoping Ma. 2023. Revisiting negative sampling vs. non-sampling in implicit recommendation. *ACM Transactions on Information Systems* 41, 1 (2023), 1–25.
- [4] Jin Chen, Defu Lian, Binbin Jin, Kai Zheng, and Enhong Chen. 2022. Learning Recommenders for Implicit Feedback with Importance Resampling. In *Proceedings of the ACM Web Conference 2022*. 1997–2005.
- [5] Ting Chen, Yizhou Sun, Yue Shi, and Liangjie Hong. 2017. On sampling strategies for neural network-based collaborative filtering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 767–776.
- [6] Jingtao Ding, Yuhuan Quan, Quanming Yao, Yong Li, and Depeng Jin. 2020. Simplify and robustify negative sampling for implicit collaborative filtering. *Advances in Neural Information Processing Systems* 33 (2020), 1094–1105.
- [7] Gintare Karolina Dziugaite and Daniel M. Roy. 2015. Neural Network Matrix Factorization. *CoRR* (2015).
- [8] Yongfu Fan, Jin Chen, Yongquan Jiang, Defu Lian, Fangda Guo, and Kai Zheng. 2023. Batch-Mix Negative Sampling for Learning Recommendation Retrievers. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 494–503.
- [9] Scott Fujimoto, David Meger, and Doina Precup. 2020. An equivalence between loss functions and non-uniform sampling in experience replay. *Advances in neural information processing systems* 33 (2020), 14219–14230.
- [10] Wei Gao and Zhi-Hua Zhou. 2012. On the consistency of AUC pairwise optimization. *arXiv preprint arXiv:1208.0645* (2012).
- [11] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yong-Dong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25–30, 2020*. 639–648.
- [12] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [13] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative filtering for implicit feedback datasets. In *2008 Eighth IEEE international conference on data mining*. Ieee, 263–272.
- [14] Tinglin Huang, Yuxiao Dong, Ming Ding, Zhen Yang, Wenzheng Feng, Xinyu Wang, and Jie Tang. 2021. MixGCF: An Improved Training Method for Graph Neural Network-based Recommender Systems. In *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14–18, 2021*, Feida Zhu, Beng Chin Ooi, and Chunyan Miao (Eds.). ACM, 665–674. <https://doi.org/10.1145/3447548.3467408>
- [15] Tinglin Huang, Yuxiao Dong, Ming Ding, Zhen Yang, Wenzheng Feng, Xinyu Wang, and Jie Tang. 2021. Mixgcf: An improved training method for graph neural network-based recommender systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 665–674.
- [16] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [17] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24–27, 2008*. 426–434.
- [18] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 426–434.
- [19] Yehuda Koren, Robert M. Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (2009), 30–37.
- [20] Daniel Lee and H Sebastian Seung. 2000. Algorithms for non-negative matrix factorization. *Advances in neural information processing systems* 13 (2000).
- [21] Seok Kee Lee, Yoon Ho Cho, and Soung Hie Kim. 2010. Collaborative filtering with ordinal scale-based implicit ratings for mobile music recommendations. *Information Sciences* 180, 11 (2010), 2142–2155.
- [22] Defu Lian, Qi Liu, and Enhong Chen. 2020. Personalized ranking with importance sampling. In *Proceedings of The Web Conference 2020*. 1093–1103.
- [23] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *Proceedings of the 2018 world wide web conference*. 689–698.
- [24] Andriy Mnih and Russ R Salakhutdinov. 2007. Probabilistic matrix factorization. *Advances in neural information processing systems* 20 (2007).
- [25] Maryam Khanian Najafabadi and Mohd Naz'ri Mahrin. 2016. A systematic literature review on the state of research and practice of collaborative filtering technique and implicit feedback. *Artificial intelligence review* 45, 2 (2016), 167–201.
- [26] Rong Pan, Yunhong Zhou, Bin Cao, Nathan N Liu, Rajan Lukose, Martin Scholz, and Qiang Yang. 2008. One-class collaborative filtering. In *2008 Eighth IEEE international conference on data mining*. IEEE, 502–511.
- [27] Dae Hoon Park and Yi Chang. 2019. Adversarial sampling and training for semi-supervised information retrieval. In *The World Wide Web Conference*. 1443–1453.
- [28] Steffen Rendle and Christoph Freudenthaler. 2014. Improving pairwise learning for item recommendation from implicit feedback. In *Proceedings of the 7th ACM international conference on Web search and data mining*. 273–282.
- [29] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [30] Ruslan Salakhutdinov and Andriy Mnih. 2007. Probabilistic Matrix Factorization. In *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3–6, 2007*. 1257–1264.
- [31] Kexin Shi, Yun Zhang, Bingyi Jing, and Wenjia Wang. 2022. Soft BPR Loss for Dynamic Hard Negative Sampling in Recommender Systems. *arXiv preprint arXiv:2211.13912* (2022).
- [32] Wentao Shi, Jiawei Chen, Fuli Feng, Jizhi Zhang, Junkang Wu, Chongming Gao, and Xiangnan He. 2023. On the Theories Behind Hard Negative Sampling for Recommendation. In *Proceedings of the ACM Web Conference 2023*. 812–822.
- [33] Rianne van den Berg, Thomas N. Kipf, and Max Welling. 2017. Graph Convolutional Matrix Completion. *CoRR* abs/1706.02263 (2017).
- [34] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. 2017. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*. 515–524.
- [35] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. KGAT: Knowledge Graph Attention Network for Recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4–8, 2019*. 950–958.
- [36] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21–25, 2019*. 165–174.
- [37] Xi Wu, Liangwei Yang, Jibing Gong, Chao Zhou, Tianyu Lin, Xiaolong Liu, and Philip S Yu. 2023. Dimension Independent Mixup for Hard Negative Sample in Collaborative Filtering. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 2785–2794.
- [38] Weinan Zhang, Tianqi Chen, Jun Wang, and Yong Yu. 2013. Optimizing top-n collaborative filtering via dynamic negative item sampling. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. 785–788.
- [39] Yuhuan Zhao, Rui Chen, Riwei Lai, Qilong Han, Hongtao Song, and Li Chen. 2023. Augmented Negative Sampling for Collaborative Filtering. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 256–266.
- [40] Qiannan Zhu, Haobo Zhang, Qing He, and Zhicheng Dou. 2022. A Gain-Tuning Dynamic Negative Sampler for Recommendation. In *Proceedings of the ACM Web Conference 2022*. 277–285.