

# Adaptive Hardness Negative Sampling for Collaborative Filtering

Riwei Lai<sup>1,2</sup>, Rui Chen<sup>1\*</sup>, Qilong Han<sup>1\*</sup>, Chi Zhang<sup>1</sup>, Li Chen<sup>2</sup>

<sup>1</sup> College of Computer Science and Technology, Harbin Engineering University

<sup>2</sup> Department of Computer Science, Hong Kong Baptist University

{lai, ruichen, hanqilong, zhangchi20}@hrbeu.edu.cn, lichen@comp.hkbu.edu.hk

## Abstract

Negative sampling is essential for implicit collaborative filtering to provide proper negative training signals so as to achieve desirable performance. We experimentally unveil a common limitation of all existing negative sampling methods that they can only select negative samples of a fixed hardness level, leading to the false positive problem (FPP) and false negative problem (FNP). We then propose a new paradigm called adaptive hardness negative sampling (AHNS) and discuss its three key criteria. By adaptively selecting negative samples with appropriate hardnesses during the training process, AHNS can well mitigate the impacts of FPP and FNP. Next, we present a concrete instantiation of AHNS called  $AHNS_{p<0}$ , and theoretically demonstrate that  $AHNS_{p<0}$  can fit the three criteria of AHNS well and achieve a larger lower bound of normalized discounted cumulative gain. Besides, we note that existing negative sampling methods can be regarded as more relaxed cases of AHNS. Finally, we conduct comprehensive experiments, and the results show that  $AHNS_{p<0}$  can consistently and substantially outperform several state-of-the-art competitors on multiple datasets.

## Introduction

Collaborative filtering (CF), as the most representative technique for recommendation, focuses on modeling user interests from observed user-item interactions (Wang et al. 2019; He et al. 2020). In many cases, it is not always possible to obtain a large amount of high-quality explicit feedback. As a result, implicit feedback, such as clicks or purchases, has become a default choice to train a CF model (Lai et al. 2023). In implicit feedback, each observed interaction normally indicates a user’s interest in an item and corresponds to a positive training sample. As for negative training samples, a widely adopted approach is to randomly select some uninteracted items for users. An implicit CF model is then optimized to give positive samples higher scores than negative ones (Rendle et al. 2009).

Similar to many semi-supervised learning problems, existing implicit CF models highly rely on mining negative samples to provide proper negative training signals. Without auxiliary data describing items, two lines of works have




Figure 1: Average hardness of selected negative items in RNS, DNS, and DENS on two Amazon datasets.

been proposed. The first line consists of *static negative sampling*, which assigns a static probability for each candidate to be sampled. For example, random negative sampling (RNS) (Rendle et al. 2009) chooses uninteracted items with equal probability, and popularity-biased negative sampling (PNS) (Chen et al. 2017; Wu et al. 2019) adopts item-popularity-biased distributions to favor popular items. The other line is *hard negative sampling*, such as dynamic negative sampling (DNS) (Zhang et al. 2013) and disentangled negative sampling (DENS) (Lai et al. 2023), which focuses on selecting hard negative samples that are difficult to be distinguished from the positive samples with dynamic distributions. Such hard negative samples can provide more informative training signals so that user interests can be better characterized (Xu et al. 2022).

Although the above two lines of works on negative sampling have achieved some promising results, we point out that all these methods can only select negative samples of a certain “hardness” level, preventing them from achieving better performance. Without loss of generality, assume that positive samples’ predicted scores are always positive. We can define the *hardness* of a negative sample as its relative predicted score, i.e., the ratio of its predicted score to that of its corresponding positive sample, in order to smooth the influence of the simultaneous increase in the predicted scores of all items during the training process. As illustrated in Fig. 1, throughout the training process, RNS can only select easy negative samples with hardness around 0, while DNS and DENS can only choose hard negative samples with hardness around 0.3 and 0.4, respectively.

Unavoidably, these fixed hardness negative sampling methods may suffer from two significant problems: (1) *false positive problem* (FPP): as shown in the upper part of Fig. 2,

\*Corresponding authors

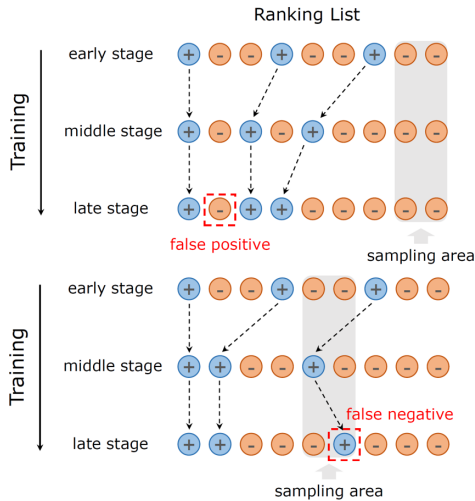


Figure 2: Issues of fixed hardness negative sampling.

when only easy negative samples can be selected during the training process, items of no interest but with initially high predicted scores may not be sufficiently updated and will still be recommended to users, resulting in suboptimal recommendation results; (2) *false negative problem* (FNP): as shown in the lower part of Fig. 2, if only hard negative samples with a fixed hardness level are selected during the training process, items of interests but not interacted yet may be selected as negative and ranked lower in the recommendation list, which worsens recommendation results. We have conducted extensive experiments to verify the existence of FPP and FNP (see RQ2 of Experiments for more details).

To address the above two problems and obtain better recommendation results, we propose to adaptively select negative samples with different hardness levels during the training process. A straightforward attempt is to introduce curriculum learning (Chu et al. 2021) into negative sampling, where a predefined pacing function is utilized to schedule the hardness levels of negative samples in different training epochs. However, such an implementation still selects negative samples with a fixed hardness level within the same epoch, rather than *adaptively select negative samples with different hardness levels for different positive samples*.

In this paper, we introduce a brand new negative sampling paradigm called Advaptive Hardness Negative Sampling (AHNS), and analyze its three key criteria. We then present a concrete instantiation of AHNS called  $AHNS_{p<0}$ , where  $p$  is a predefined smoothing parameter we will explain later. Comprehensive theoretical analyses are performed to confirm that  $AHNS_{p<0}$  satisfies the three criteria of AHNS and prove that implicit CF models with  $AHNS_{p<0}$  can achieve a larger lower bound on normalized discounted cumulative gain (NDCG) than with a fixed hardness negative sampling method. Furthermore, we discuss the relation between AHNS and other negative sampling methods and note that existing negative sampling methods can be considered as more relaxed cases of AHNS.

Our main contributions are summarized as follows:

- We are the first to identify and address FPP and FNP in existing negative sampling methods via adaptively selecting hardnesses of negative samples, which brings a new perspective of negative sampling for implicit CF.
- We propose a new negative sampling paradigm AHNS with three criteria, which generalizes existing negative sampling methods. We present a concrete instantiation  $AHNS_{p<0}$  and theoretically show that it can fit the three criteria well and achieve a larger lower bound on NDCG.
- We conduct extensive experiments to demonstrate that  $AHNS_{p<0}$  can achieve significant improvements over several representative state-of-the-art negative sampling methods.

## Related Work

### Static Negative Sampling

Static negative sampling focuses on identifying good distributions to draw negative samples. For example, as the simplest and most prevalent static negative sampling method, Bayesian personalized ranking (BPR) (Rendle et al. 2009) randomly selects uninteracted items as negative. However, this method makes it hard to guarantee the quality of selected negative samples, and thus some studies (Chen et al. 2017; Wu et al. 2019; Yang et al. 2020) propose to replace the uniform distribution with other distributions. Inspired by the word-frequency-based and node-degree-based negative sampling distributions for network embedding (Mikolov et al. 2013), NNCF (Chen et al. 2017) and NCE-PLRec (Wu et al. 2019) adopt an item-popularity-based sampling distribution to select more popular items as negative, which helps to alleviate the widespread popularity bias issue in recommender systems (Chen et al. 2023).

### Hard Negative Sampling

Hard negative sampling methods emphasize the importance of oversampling hard negative samples to speed up the training process and find more precise delineations of user interests. More specifically, it is achieved by either assigning higher sampling probabilities to items with larger predicted scores (Zhang et al. 2013; Ding et al. 2020; Huang et al. 2021; Zhu et al. 2022; Lai et al. 2023; Shi et al. 2023; Zhao et al. 2023) or leveraging adversarial learning techniques (Wang et al. 2017; Cai and Wang 2018; Park and Chang 2019). For instance, dynamic negative sampling (DNS) (Zhang et al. 2013) selects the item with the highest predicted score in a candidate negative sample set. SRNS (Ding et al. 2020) oversamples items with both high predicted scores and high variances to tackle the false negative problem. DENS (Lai et al. 2023) disentangles relevant and irrelevant factors of items and identifies the best negative samples with a factor-aware sampling strategy. Instead of directly selecting negative samples from uninteracted items, MixGCF (Huang et al. 2021) synthesizes hard negative samples by mixing positive information into negative samples, which further improves the performance.

However, we experimentally find that all the above negative sampling methods can only select negative samples of a fixed hardness level during the training process, leading

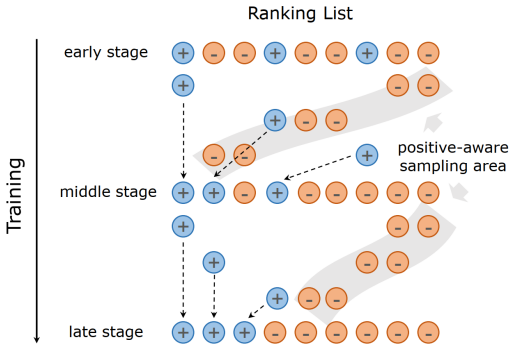


Figure 3: An illustration of adaptive hardness negative sampling.

to the false positive problem and false negative problem. Driven by this limitation, we propose an adaptive hardness negative sampling paradigm, which adaptively selects negative samples with appropriate hardnesses and achieves better recommendation results.

## Proposed Method

### Problem Formulation

In this section, we formulate the problem of negative sampling in implicit CF. Let  $\mathcal{U}$  and  $\mathcal{I}$  be the set of users and the set of items, respectively. We denote the set of observed interactions, i.e., implicit feedback, by  $\mathcal{O}^+ = \{(u, i^+) \mid u \in \mathcal{U}, i^+ \in \mathcal{I}\}$ , where each pair  $(u, i^+)$  indicates an interaction between user  $u$  and item  $i^+$ . Implicit CF aims to characterize user interests from their observed interactions. Interacted items are generally used to form positive pairs, while uninteracted items are considered candidates to generate negative samples. Specifically, given a positive pair  $(u, i^+)$ , a negative sampling strategy identifies an item  $i^-$  that has not been previously interacted by  $u$  as a negative sample. The implicit CF model is then optimized to give positive pairs higher scores than negative pairs by the Bayesian personalized ranking (BPR) loss function (Rendle et al. 2009):

$$\mathcal{L}_{\text{BPR}} = \sum_{(u, i^+, i^-)} -\ln \sigma(\mathbf{e}_u^\top \mathbf{e}_{i^+} - \mathbf{e}_u^\top \mathbf{e}_{i^-}), \quad (1)$$

where  $\mathbf{e}_u$ ,  $\mathbf{e}_{i^+}$ , and  $\mathbf{e}_{i^-}$  are the embeddings of user  $u$ , positive sample  $i^+$ , and negative sample  $i^-$ , respectively, the inner product is used to measure the score of positive and negative pairs, and  $\sigma(\cdot)$  is the sigmoid function.

### Method Design

**Paradigm.** To achieve the adaptive selection in the hardnesses of negative samples and alleviate the false positive problem (FPP) and false negative problem (FNP), we propose the adaptive hardness negative sampling (AHNS) paradigm. As shown in Fig. 3, unlike fixed hardness negative sampling, AHNS simultaneously satisfies the following three key criteria:

- **C1: The hardness of a selected negative sample should be positive-aware.** Instead of setting a specific hardness

### Algorithm 1: AHNS<sub>p<0</sub>

---

```

1: Input: Set of observed interactions  $\mathcal{O}^+ = \{(u, i^+) \mid u \in \mathcal{U}, i^+ \in \mathcal{I}\}$ , number of candidate negatives  $M$ , predefined hyperparameters  $\alpha, \beta$  and  $p$ 
2: Output: Set of training triples  $\mathcal{T}$ 
3:  $\mathcal{T} \leftarrow \{\}$   $\triangleright$  Initialize an empty set for training triples
4: for each positive pair  $(u, i^+)$  in  $\mathcal{O}^+$  do
5:    $\mathcal{C} \leftarrow \{\}$   $\triangleright$  Initialize an empty set for candidate negative samples
6:   for  $m = 1$  to  $M$  do
7:      $i_m \leftarrow$  Randomly sample an uninteracted item
8:     Add  $i_m$  to  $\mathcal{C}$ 
9:   end for
10:   $\mathcal{R} \leftarrow \{\}$   $\triangleright$  Initialize an empty set for ratings of candidate negative samples
11:  for each candidate negative sample  $i_m$  in  $\mathcal{C}$  do
12:     $r_m \leftarrow |\mathbf{e}_u^\top \mathbf{e}_{i_m} - \beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)^{p+1}|$ 
13:    Add  $r_m$  to  $\mathcal{R}$ 
14:  end for
15:   $i^- \leftarrow$  Select  $i_m$  with the smallest  $r_m$  in  $\mathcal{R}$ 
16:  Add  $(u, i^+, i^-)$  to  $\mathcal{T}$ 
17: end for
18: return  $\mathcal{T}$ 

```

---

level of negative samples for each training epoch like curriculum learning (Chu et al. 2021), AHNS is expected to identify the appropriate hardness of a negative sample according to its corresponding positive sample.

- **C2: The hardness of a selected negative sample should be negatively correlated with the predicted score of its corresponding positive sample.** On the one hand, for positive samples with higher predicted scores, AHNS should select items with lower hardnesses as negative, which can effectively avoid the FNP. On the other hand, for positive samples with lower predicted scores, AHNS should select items with higher hardnesses as negative, which can accelerate the optimization of positives and enable negatives with higher hardnesses to be sufficiently updated, thus alleviating the FPP.
- **C3: The hardness of selected negative samples should be adjustable.** To cover a variety of practical recommendation scenarios, e.g., different datasets or evaluation metrics (Shi et al. 2023), AHNS should be able to adjust the hardness of selected negative samples.

**Instantiation.** Next, we give a concrete instantiation of AHNS called AHNS<sub>p<0</sub>, whose entire procedure is detailed in Algo. 1. Specifically, for a positive pair  $(u, i^+)$ , we follow conventional methods (Chen et al. 2022; Lai et al. 2023) and adapt the two-pass sampling idea, which first randomly samples a fixed size of uninteracted items to form a candidate set, and then selects a negative sample from the candidate set according to predefined rating functions and sampling rules. For the first pass, the size  $M$  of the candidate set  $\mathcal{C}$  is usually much smaller than the total number of items  $|\mathcal{I}|$ , which can boost the sampling efficiency. For the second pass, the rating function and sampling rule play a critical role in iden-

tifying the final negative sample and are the focus of all negative sampling methods. Therefore, we introduce three hyperparameters and carefully design the rating function in  $\text{AHNS}_{p<0}$ . For each candidate negative item  $i_m \in \mathcal{C}$ , the rating function is formulated as:

$$r_m = \left| \mathbf{e}_u^\top \mathbf{e}_{i_m} - \beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)^{p+1} \right|, \quad (2)$$

where  $\alpha > 0$ ,  $\beta > 0$  and  $p < 0$  are predefined hyperparameters, whose effects will be given in the subsequent Thm. 3. After calculating the ratings of all candidate negative items, we obtain a rating set  $\mathcal{R}$ , and then the final negative sample is identified by selecting  $i_m$  with the smallest  $r_m$  in  $\mathcal{R}$ :

$$i^- = i_{\arg \min_m r_m}. \quad (3)$$

### Theoretical Analysis

In this section, we conduct in-depth analyses on  $\text{AHNS}_{p<0}$ . We first show that  $\text{AHNS}_{p<0}$  satisfies the three criteria of AHNS, and then establish that implicit CF models with  $\text{AHNS}_{p<0}$  can achieve a larger lower bound on normalized discounted cumulative gain than with a fixed hardness negative sampling method as training progresses.

**Theorem 1.**  *$\text{AHNS}_{p<0}$  satisfies C2 of AHNS.*

*Proof.* Consider a positive pair  $(u, i^+)$ . Let  $i_*^-$  be the ideal negative sample selected by  $\text{AHNS}_{p<0}$ . According to Eq. (2) and Eq. (3), we have:

$$\mathbf{e}_u^\top \mathbf{e}_{i_*^-} = \beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)^{p+1}. \quad (4)$$

To simplify the calculation process, we substitute  $\mathbf{e}_u^\top \mathbf{e}_{i^+}$  with  $(\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)$  to calculate the hardness of  $i_*^-$ :

$$\begin{aligned} \text{Hardness}(i_*^-) &= \frac{\mathbf{e}_u^\top \mathbf{e}_{i_*^-}}{\mathbf{e}_u^\top \mathbf{e}_{i^+}} \approx \frac{\mathbf{e}_u^\top \mathbf{e}_{i_*^-}}{\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha} \\ &= \frac{\beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)^{p+1}}{\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha} \\ &= \beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)^p. \end{aligned} \quad (5)$$

Based on the chain rule, we have:

$$\begin{aligned} \frac{d\text{Hardness}(i_*^-)}{d(\mathbf{e}_u^\top \mathbf{e}_{i^+})} &= \frac{d(\beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)^p)}{d(\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)} \cdot \frac{d(\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)}{d(\mathbf{e}_u^\top \mathbf{e}_{i^+})} \\ &= p \cdot \beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i^+} + \alpha)^{p-1}. \end{aligned} \quad (6)$$

Clearly,  $d\text{Hardness}(i_*^-)/d(\mathbf{e}_u^\top \mathbf{e}_{i^+}) < 0$  always holds when  $\mathbf{e}_u^\top \mathbf{e}_{i^+} > 0$ ,  $\alpha > 0$ ,  $\beta > 0$  and  $p < 0$ , which means that the hardness of  $i_*^-$  is always negatively correlated with the predicted score of  $i^+$  – the above completes the proof.  $\square$

**Theorem 2.**  *$\text{AHNS}_{p<0}$  satisfies C1 of AHNS.*

*Proof.* Consider two different positive pairs  $(u, i_1^+)$  and  $(u, i_2^+)$ . Let  $i_{1*}^-$  and  $i_{2*}^-$  be the ideal negative samples selected by  $\text{AHNS}_{p<0}$  for  $(u, i_1^+)$  and  $(u, i_2^+)$ , respectively. According to Eq. (5), we have:

$$\begin{aligned} \text{Hardness}(i_{1*}^-) &= \beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i_1^+} + \alpha)^p, \\ \text{Hardness}(i_{2*}^-) &= \beta \cdot (\mathbf{e}_u^\top \mathbf{e}_{i_2^+} + \alpha)^p. \end{aligned} \quad (7)$$




Figure 4: Hardness of ideal negative sample  $i_*^-$  w.r.t. different  $p$ .

It has been proved in Thm. 1 that  $\text{Hardness}(i_*^-)$  monotonically decreases as  $\mathbf{e}_u^\top \mathbf{e}_{i^+}$  increases. Thus when  $\mathbf{e}_u^\top \mathbf{e}_{i_1^+} \neq \mathbf{e}_u^\top \mathbf{e}_{i_2^+}$ ,  $\text{Hardness}(i_{1*}^-) \neq \text{Hardness}(i_{2*}^-)$  – the above completes the proof.  $\square$

**Theorem 3.**  *$\text{AHNS}_{p<0}$  satisfies C3 of AHNS.*

*Proof.* According to Eq. (5), we plot the curves of the hardness of  $i_*^-$  under different values of the predicted score of  $i^+$  in Fig. 4. It is clear that  $p$  affects the magnitude of the curves, smaller  $p$  leads to larger magnitudes. In addition, all curves pass through the point  $(1 - \alpha, \beta)$ , indicating the effect of  $\alpha$  and  $\beta$  in adjusting the hardness of selected negative samples – the above completes the proof.  $\square$

**Theorem 4.** *As training progresses, implicit CF models with  $\text{AHNS}_{p<0}$  can achieve a larger lower bound on normalized discounted cumulative gain (NDCG) than with a fixed hardness negative sampling method.*

*Proof.* Given a user  $u$ , let  $\pi_{f_u}$  be the ranking function induced by recommender system  $f$  for user  $u$ , and  $\pi_{f_u}(i)$  the rank of item  $i$ . Let  $y$  be a binary indicator:  $y_i = 1$  if item  $i$  has been interacted by  $u$ , otherwise  $y_i = 0$ . Let  $\mathcal{I}(u) = \{i \mid y_i = 1\}$  be the set of items interacted by  $u$  and  $\mathbb{I}$  be the indicator function.

First, we consider discounted cumulative gain (DCG). With  $1 + z \leq 2^z$  when  $z \geq 1$ , we have the following:

$$\begin{aligned} \text{DCG}(u) &= \sum_{i=1}^{|\mathcal{I}|} \frac{2^{y_i} - 1}{\log_2(1 + \pi_{f_u}(i))} \\ &= \sum_{i=1}^{|\mathcal{I}(u)|} \frac{1}{\log_2(1 + \pi_{f_u}(i))} \geq \sum_{i=1}^{|\mathcal{I}(u)|} \frac{1}{\pi_{f_u}(i)} \\ &= \sum_{i=1}^{|\mathcal{I}(u)|} \frac{1}{1 + \sum_{j \in |\mathcal{I}| \setminus \{i\}} \mathbb{I}(\mathbf{e}_u^\top \mathbf{e}_j - \mathbf{e}_u^\top \mathbf{e}_i > 0)} \\ &\geq \sum_{i=1}^{|\mathcal{I}(u)|} \frac{1}{1 + \sum_{j \in |\mathcal{I}| \setminus \{i\}} \exp(\mathbf{e}_u^\top \mathbf{e}_j - \mathbf{e}_u^\top \mathbf{e}_i)}. \end{aligned} \quad (8)$$

Next, we consider the ideal DCG (IDCG). Let  $\pi_{f_u}^*$  be the ideal ranking function, which can sort the items in the ground truth order:

$$\begin{aligned} IDCG(u) &= \sum_{i=1}^{|\mathcal{I}|} \frac{2^{y_i} - 1}{\log_2(1 + \pi_{f_u}^*(i))} \\ &= \sum_{i=1}^{|\mathcal{I}(u)|} \frac{1}{\log_2(1 + i)} \leq \sum_{i=1}^{|\mathcal{I}(u)|} 1 = |\mathcal{I}(u)|. \end{aligned} \quad (9)$$

Clearly, we have:

$$\frac{1}{IDCG(u)} \geq \frac{1}{|\mathcal{I}(u)|}. \quad (10)$$

Finally, we consider NDCG:

$$\begin{aligned} NDCG(u) &= \frac{DCG(u)}{IDCG(u)} \geq \frac{1}{|\mathcal{I}(u)|} DCG(u) \\ &\geq \frac{1}{|\mathcal{I}(u)|} \sum_{i=1}^{|\mathcal{I}(u)|} \frac{1}{1 + \sum_{j \in |\mathcal{I}| \setminus \{i\}} \exp(\mathbf{e}_u^\top \mathbf{e}_j - \mathbf{e}_u^\top \mathbf{e}_i)} \\ &\approx \frac{1}{|\mathcal{I}(u)|} \sum_{i=1}^{|\mathcal{I}(u)|} \frac{1}{1 + \exp(\mathbf{e}_u^\top \mathbf{e}_{i^-} - \mathbf{e}_u^\top \mathbf{e}_i)}. \end{aligned} \quad (11)$$

As illustrated in Fig. 4, it is not difficult to derive that as the predicted score of  $i^+$  increases, the hardness of  $i_-^*$  sampled by  $\text{AHNS}_{p < 0}$  (the solid lines) is lower compared to that of a fixed hardness negative sampling method (the dotted line), leading to a lower value of  $\exp(\mathbf{e}_u^\top \mathbf{e}_{i^-} - \mathbf{e}_u^\top \mathbf{e}_i)$ . Thus implicit CF models with  $\text{AHNS}_{p < 0}$  can achieve a larger lower bound of NDCG. This completes the proof.  $\square$

## Discussion

In this section, we first discuss the relation between AHNS and other negative sampling methods. We point out that existing negative sampling methods can be considered as more relaxed cases that satisfy part of the three criteria of AHNS. For example, DENS (Lai et al. 2023) proposes a positive gating layer to disentangle items' factors for negative sampling. Thus the hardness of selected negative samples becomes positive-aware and satisfies C1 of AHNS. By using an anti-curriculum pacing function to schedule the hardnesses of negative samples for different training epochs, CuCo (Chu et al. 2021) partially satisfies C2 of AHNS. To adapt to different datasets and top- $K$  metrics,  $\text{DNS}(M, N)$  (Shi et al. 2023) adjusts the hardnesses of selected negative samples via predefined hyperparameters, which satisfies C3 of AHNS.

In addition, we note that the main idea of AHNS in negative sampling is consistent with that of focal loss (Lin et al. 2017) in object detection, i.e., **putting more focus on lower-ranked positives and higher-ranked negatives (hard, misclassified examples)**, which may bring some new insights into negative sampling for implicit CF.

## Experiments

In this section, we perform extensive experiments to evaluate  $\text{AHNS}_{p < 0}$  and answer the following research questions:

Dataset	#user ( $ \mathcal{U} $ )	#item ( $ \mathcal{I} $ )	#inter. ( $ \mathcal{R} $ )	avg. inter. per user	density
ML-1M	6.0k	3.7k	1000.2k	165.6	4.47%
Phones	27.9k	10.4k	194.4k	7.0	0.07%
Sports	35.6k	18.4k	296.3k	8.3	0.05%
Tools	16.6k	10.2k	134.5k	8.1	0.08%

Table 1: The statistics of four datasets.

- **RQ1:** How does  $\text{AHNS}_{p < 0}$  perform compared with previous negative sampling methods?
- **RQ2:** Does  $\text{AHNS}_{p < 0}$  achieve adaptive selection in the hardnesses of negative samples and alleviate the false positive problem (FPP) and false negative problem (FNP)?
- **RQ3:** What are the impacts of the hyperparameters (e.g.,  $\alpha, \beta$ ) on  $\text{AHNS}_{p < 0}$ ?
- **RQ4:** Does  $\text{AHNS}_{p < 0}$  have an advantage in terms of sampling efficiency?

## Experimental Setup

**Datasets and Evaluation Metrics.** We consider four widely used public benchmark datasets in experiments: MovieLens-1M<sup>1</sup> (ML-1M), Amazon-Phones<sup>2</sup> (Phones), Amazon-Sports<sup>2</sup> (Sports) and Amazon-Tools<sup>2</sup> (Tools). Following (He et al. 2020; Shi et al. 2023), we randomly split each user's interactions into training/test sets with a ratio of 80%/20%, and build the validation set by randomly sampling 10% interactions of the training set. Tab. 1 summarizes the statistics of the four datasets. We report the recommendation performances in terms of  $Recall@20$  ( $R@20$ ) and  $NDCG@\{20, 50\}$  ( $N@\{20, 50\}$ ), where higher values indicate better performances.

**Baseline Methods.** We compare  $\text{AHNS}_{p < 0}$  with a wide range of representative negative sampling methods:

- **RNS** (Rendle et al. 2009) randomly selects uninteracted items as negative.
- **SSM** (Wu et al. 2022) achieves better performances by sampling more items as negative.
- **DNS** (Zhang et al. 2013) chooses the item with the highest predicted score in a candidate set as negative.
- **MixGCF** (Huang et al. 2021) synthesizes harder negative samples by injecting information from positive samples.
- **DENS** (Lai et al. 2023) identifies better negative samples by disentangling factors of items.
- **$\text{DNS}(M, N)$**  (Shi et al. 2023) controls the sampling hardness via predefined hyperparameters.
- **GuCo** (Chu et al. 2021) proposes a negative sampling method adopting curriculum learning in graph representation learning. We transfer this method to CF.

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

<sup>2</sup><https://jmcauley.ucsd.edu/data/amazon/>

Method	ML-1M			Phones			Sports			Tools		
	R@20	N@20	N@50	R@20	N@20	N@50	R@20	N@20	N@50	R@20	N@20	N@50
RNS	22.86	35.46	37.41	11.06	5.98	7.35	6.73	3.60	4.68	5.53	2.99	3.75
SSM	24.87	37.74	39.71	11.37	6.13	7.48	7.08	3.80	4.87	5.72	3.10	3.88
DNS	24.66	36.64	38.31	12.08	6.64	7.99	7.74	4.25	5.32	6.66	3.78	4.52
MixGCF	24.75	37.54	38.95	12.20	6.73	8.13	7.68	4.32	5.36	6.82	3.88	4.59
DENS	25.07	37.67	39.11	12.16	6.68	8.13	7.90	4.35	5.50	6.66	3.76	4.55
DNS( $M, N$ )	25.09	37.58	39.22	12.27	6.75	8.15	7.84	4.31	5.35	6.86	3.76	4.61
CuCo	25.12	37.53	39.20	12.19	6.68	8.11	7.68	4.25	5.36	6.76	3.82	4.59
AHNS $_{p=-1}$	<u>25.17</u>	37.72	39.31	<u>13.02</u>	<u>7.08</u>	<u>8.71</u>	<u>8.42</u>	<u>4.58</u>	<b>5.82</b>	<u>7.27</u>	<u>4.02</u>	<b>4.95</b>
AHNS $_{p=-2}$	<b>25.51</b>	<b>38.77</b>	<b>40.57</b>	<b>13.03</b>	<b>7.14</b>	<b>8.74</b>	<b>8.52</b>	<b>4.61</b>	<u>5.81</u>	<b>7.42</b>	<b>4.05</b>	<u>4.92</u>
Improv.	1.6%	2.7%	2.2%	6.2%	5.8%	7.2%	7.8%	6.0%	5.8%	8.2%	4.4%	7.4%

Table 2: Performances (%) of AHNS $_{p=-1}$ , AHNS $_{p=-2}$ , and baseline methods. The best results are in bold, and the second best are underlined. Improvements are calculated over the best baseline method and statistically significant with  $p$ -value  $< 0.01$ .

**Implementation Details.** We strictly follow the experimental setting in DENS (Lai et al. 2023). We utilize matrix factorization (MF) as the implicit CF model. The embedding dimension is fixed to 64, and the embedding parameters are initialized with the Xavier method. We optimize all parameters with Adam (Kingma and Ba 2015) and use the default learning rate of 0.001 and default mini-batch size of 2,048. The number of training epochs is set to 100. For AHNS $_{p<0}$ , the candidate negative size  $M$  is searched in the range of  $\{4, 8, 16, 32, 64\}$ . The hyperparameters  $\alpha$  and  $\beta$  are tuned over  $\{0.1, 0.2, \dots, 0.9, 1.0\}$  independently. The hyperparameters of all baseline methods are carefully tuned by grid search. Our code is publically available at <https://github.com/Riwei-HEU/AHNS>.

### RQ1: Performance Comparison

Tab. 2 shows the performances of AHNS $_{p=-1}$ , AHNS $_{p=-2}$ , and baseline methods. We can observe the following:

- Compared to randomly selecting uninteracted items as negative (RNS), increasing the number (SSM) or the hardness (DNS, MixGCF, DENS, etc.) of negative samples leads to a substantial performance improvement.
- By introducing curriculum learning into negative sampling, CuCo draws negative samples with different hardnesses in different training epochs, achieving comparable performances to hard negative sampling methods.
- Benefiting from positive-aware adaptive selection in the hardnesses of negative samples, AHNS $_{p=-1}$  and AHNS $_{p=-2}$  significantly outperform RNS by on average 20%. Meanwhile, the two methods also show a huge performance boost over other hard negative sampling methods and the curriculum-learning-based method.

### RQ2: Hardness Visualization

To justify the motivation of AHNS, i.e., adaptively selecting hardnesses of negative samples to alleviate FPP and FNP, we plot the curves of average negative hardness and  $NDCG@20$  of RNS, DNS, DENS, and AHNS $_{p=-2}$  in

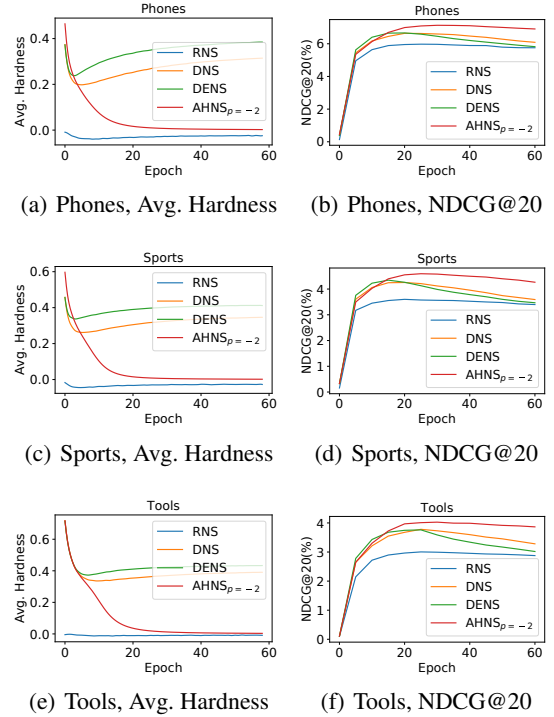


Figure 5: Average negative hardness and  $NDCG@20$  of RNS, DNS, DENS, and AHNS $_{p=-2}$ .

Fig. 5. Due to the space limitation, we only report the results on the three Amazon datasets. From these figures, we have the following key findings:

- As shown in Fig. 5(a), 5(c) and 5(e), compared to fixed hardness negative sampling methods RNS, DNS, and DENS, AHNS $_{p=-2}$  can adaptively adjust the hardnesses of negative samples as training progresses. Specifically, in the early stages of training, AHNS $_{p=-2}$  favors negative samples with higher hardnesses, while in the later stages of training, AHNS $_{p=-2}$  prefers negative samples with lower hardnesses.

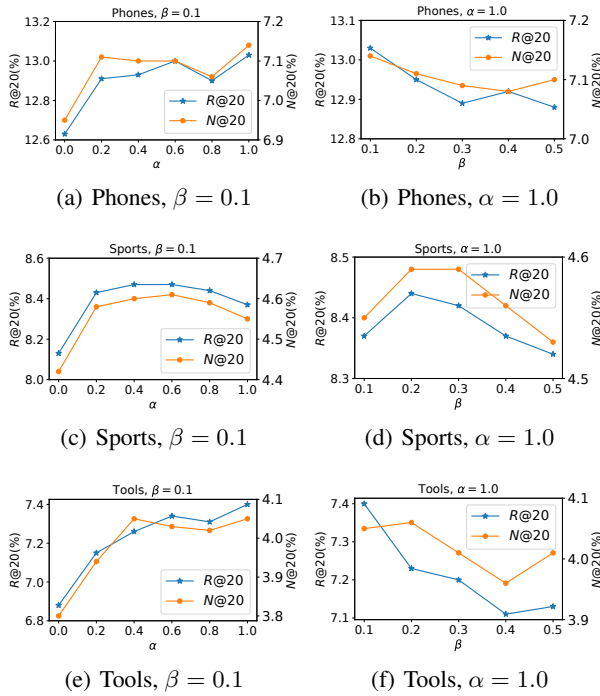


Figure 6: Performance of  $AHNS_{p=-2}$  w.r.t. different hyperparameters.

- As shown in Fig. 5(b), 5(d) and 5(f), the performance of RNS peaks in the early stages of training and remains stable thereafter, DNS and DENS present better performance than RNS but suffer a significant performance drop in the later stages of training, and  $AHNS_{p=-2}$  achieves the best performance while maintaining similar stability as RNS.
- The average negative hardness and  $NDCG@20$  of RNS (the blue line) well verify the existence of FPP, i.e., when only easy negative samples can be selected during the training process, items of no interest but with initially high predicted scores may not be sufficiently updated and will still be recommended to users, leading to the suboptimal performance of RNS.
- The average negative hardness and  $NDCG@20$  of DNS (the orange line) and DENS (the green line) well verify the existence of FNP, i.e., when only hard negative samples can be selected during the training process, items of interests may be selected as negative and ranked lower in the recommendation list, resulting in the performance drop of DNS and DENS.
- The average negative hardness and  $NDCG@20$  of  $AHNS_{p=-2}$  (the red line) well justify our motivation. For positives with lower predicted scores, by selecting items with higher hardnesses as negative,  $AHNS_{p=-2}$  well alleviates FPP and achieves a higher peak; for positives with higher predicted scores, by selecting items with lower hardnesses as negative,  $AHNS_{p=-2}$  well avoids FNP and thus prevents the performance drop.




Figure 7: Time (second) for training per epoch on ML-1M w.r.t. different methods. Best viewed in color.

### RQ3: Hyperparameter Study

As discussed in Thm. 3, hyperparameters  $\alpha$  and  $\beta$  affect the hardnesses of selected negative samples. Here we study how these hyperparameters affect the recommendation performance. Fig. 6 shows  $Recall@20$  and  $NDCG@20$  of  $AHNS_{p=-2}$  under different  $\alpha$  or  $\beta$  values with other hyperparameters unchanged on the three Amazon datasets. We can see that it is intractable to identify the optimal values of  $\alpha$  and  $\beta$  since they are different across datasets and evaluation metrics. However, in practice, we can achieve desirable performance in a relatively wide range of  $\alpha$  or  $\beta$  values, which relieves the overhead of hyperparameter tuning.

### RQ4: Efficiency Analysis

As presented in Algo. 1,  $AHNS_{p<0}$  does not introduce additional time cost compared to the simplest hard negative sampling method DNS. Here we empirically compare the time for training each epoch of  $AHNS_{p=-2}$  and other baseline methods on the ML-1M dataset. All the methods are implemented under the same framework and with optimal hyperparameters to ensure fairness. The results are shown in Fig. 7.  $DNS(M, N)$  takes the longest time as it requires an extremely large candidate negative set to adjust the hardness of negative samples. SSM costs the second longest time because multiple negative samples are selected to participate in the training of the CF model. The time difference between  $AHNS_{p=-2}$  and other hard negative sampling methods is marginal, and RNS undoubtedly takes the least time. Considering the performance improvements in Tab. 2 that  $AHNS_{p=-2}$  can bring, we believe that  $AHNS_{p=-2}$  is the best negative sampling method in terms of both efficiency and performance.

## Conclusion

In this paper, we propose a new negative sampling paradigm  $AHNS$  with three key criteria, which enables adaptive selection of hardnesses of negative samples to alleviate FPP and FNP. We devise a concrete instantiation  $AHNS_{p<0}$  and theoretically demonstrate that it can well fit the three criteria of  $AHNS$  and achieve a larger lower bound of  $NDCG$ . Comprehensive experiments confirm that  $AHNS_{p<0}$  provides a promising new research direction for negative sampling to further boost implicit CF models' performance.

## Acknowledgments

This work was supported by the Heilongjiang Key R&D Program of China under Grant No. GA23A915 and the National Natural Science Foundation of China under Grant No. 62072136. It was also partially supported by Hong Kong Baptist University IG-FNRA project under Grant No. RC-FNRA-IG/21-22/SCI/01.

## References

- Cai, L.; and Wang, W. Y. 2018. KBGAN: Adversarial Learning for Knowledge Graph Embeddings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1470–1480.
- Chen, J.; Dong, H.; Wang, X.; Feng, F.; Wang, M.; and He, X. 2023. Bias and Debias in Recommender System: A Survey and Future Directions. *ACM Transactions on Information Systems*, 41(3): 1–39.
- Chen, J.; Lian, D.; Jin, B.; Zheng, K.; and Chen, E. 2022. Learning Recommenders for Implicit Feedback with Importance Resampling. In *Proceedings of the ACM Web Conference 2022*, 1997–2005.
- Chen, T.; Sun, Y.; Shi, Y.; and Hong, L. 2017. On Sampling Strategies for Neural Network-based Collaborative Filtering. In *Proceedings of the 23rd ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 767–776.
- Chu, G.; Wang, X.; Shi, C.; and Jiang, X. 2021. CuCo: Graph Representation with Curriculum Contrastive Learning. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, 2300–2306.
- Ding, J.; Quan, Y.; Yao, Q.; Li, Y.; and Jin, D. 2020. Simplify and Robustify Negative Sampling for Implicit Collaborative Filtering. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*.
- He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 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.
- Huang, T.; Dong, Y.; Ding, M.; Yang, Z.; Feng, W.; Wang, X.; and Tang, J. 2021. MixGCF: An Improved Training Method for Graph Neural Network-based Recommender Systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 665–674.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations*.
- Lai, R.; Chen, L.; Zhao, Y.; Chen, R.; and Han, Q. 2023. Disentangled Negative Sampling for Collaborative Filtering. In *Proceedings of the 16th International Conference on Web Search And Data Mining*, 96–104.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 2980–2988.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, 3111–3119.
- Park, D. H.; and Chang, Y. 2019. Adversarial Sampling and Training for Semi-Supervised Information Retrieval. In *Proceedings of the 28th International Conference on World Wide Web*, 1443–1453.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 452–461.
- Shi, W.; Chen, J.; Feng, F.; Zhang, J.; Wu, J.; Gao, C.; and He, X. 2023. On the Theories Behind Hard Negative Sampling for Recommendation. In *Proceedings of the ACM Web Conference 2023*, 812–822.
- Wang, J.; Yu, L.; Zhang, W.; Gong, Y.; Xu, Y.; Wang, B.; Zhang, P.; and Zhang, D. 2017. IRGAN: A Minimax Game for Unifying Generative and Discriminative Information Retrieval Models. In *Proceedings of the 40rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 515–524.
- Wang, X.; He, X.; Wang, M.; Feng, F.; and Chua, T.-S. 2019. Neural Graph Collaborative Filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 165–174.
- Wu, G.; Volkovs, M.; Soon, C. L.; Sanner, S.; and Rai, H. 2019. Noise Contrastive Estimation for One-Class Collaborative Filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 135–144.
- Wu, J.; Wang, X.; Gao, X.; Chen, J.; Fu, H.; Qiu, T.; and He, X. 2022. On the Effectiveness of Sampled Softmax Loss for Item Recommendation. *arXiv preprint arXiv:2201.02327*.
- Xu, L.; Lian, J.; Zhao, W. X.; Gong, M.; Shou, L.; Jiang, D.; Xie, X.; and Wen, J.-R. 2022. Negative Sampling for Contrastive Representation Learning: A Review. *arXiv preprint arXiv:2206.00212*.
- Yang, Z.; Ding, M.; Zhou, C.; Yang, H.; Zhou, J.; and Tang, J. 2020. Understanding Negative Sampling in Graph Representation Learning. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1666–1676.
- Zhang, W.; Chen, T.; Wang, J.; and Yu, Y. 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.
- Zhao, Y.; Chen, R.; Lai, R.; Han, Q.; Song, H.; and Chen, L. 2023. Augmented Negative Sampling for Collaborative Filtering. In *Proceedings of the 17th ACM Conference on Recommender Systems*.
- Zhu, Q.; Zhang, H.; He, Q.; and Dou, Z. 2022. A Gain-Tuning Dynamic Negative Sampler for Recommendation. In *Proceedings of the Web Conference 2022*, 277–285.