

Trading Hard Negatives and True Negatives: A Debiased Contrastive Collaborative Filtering Approach

Chenxiao Yang¹, Qitian Wu¹, Jipeng Jin¹, Xiaofeng Gao^{1*}, Junwei Pan² and Guihai Chen¹

¹MoE Key Lab of Artificial Intelligence, Department of Computer Science and Engineering,
Shanghai Jiao Tong University

²Tencent Inc.

{chr26195, echo740, jinjipeng, gaoxiaofeng}@sjtu.edu.cn, jonaspan@tencent.com, gchen@cs.sjtu.edu.cn

Abstract

Collaborative filtering (CF), as a standard method for recommendation with implicit feedback, tackles a semi-supervised learning problem where most interaction data are unobserved. Such a nature makes existing approaches highly rely on mining negatives for providing correct training signals. However, mining proper negatives is not a free lunch, encountering with a tricky trade-off between mining informative hard negatives and avoiding false ones. We devise a new approach named as Hardness-Aware Debiased Contrastive Collaborative Filtering (HDCCF) to resolve the dilemma. It could sufficiently explore hard negatives from two-fold aspects: 1) adaptively sharpening the gradients of harder instances through a set-wise objective, and 2) implicitly leveraging item/user frequency information with a new sampling strategy. To circumvent false negatives, we develop a principled approach to improve the reliability of negative instances and prove that the objective is an unbiased estimation of sampling from the true negative distribution. Extensive experiments demonstrate the superiority of the proposed model over existing CF models and hard negative mining methods.

1 Introduction

Collaborative Filtering (CF) [Pan *et al.*, 2008] is a standard approach to deal with implicit feedback (e.g., click, watch, purchase, etc.) in recommender systems, wherein observed user-item interactions are assigned with positive labels, and the rest are unlabeled. A common practice in CF methods is to uniformly draw negative instances from the unlabeled portion, a.k.a. *negative sampling* [Chen *et al.*, 2017], and then use both positive and negative instances for training, as has been adopted by existing *point-wise* [Mnih and

Salakhutdinov, 2008] or *pair-wise* [Rendle *et al.*, 2009] approaches. However, this CF paradigm is considered as insufficient to provide informative and reliable training signals. Hence, enormous efforts have been made for improving the quality of negative instances for CF-based recommendation.

Particularly, *hard negative mining* has shown to be an effective approach, which aims to exploit negative user-item pairs whose embeddings are close yet expected to be far apart [Wu *et al.*, 2017; Park and Chang, 2019], as a means to provide informative training signals. A line of works [Rendle and Freudenthaler, 2014; Chen *et al.*, 2017] fallen into this category attempt to replace the uniform negative sampling distribution by some predefined surrogates, based on certain prior knowledge such as that more frequent items constitutes better negatives [Wu *et al.*, 2019a; Chen *et al.*, 2017]. In contrast, another line of works [Rendle and Freudenthaler, 2014; Park and Chang, 2019] seek to adaptively mine negatives by carefully examining relevance score of user-item pairs, which are generally more effective but often require sophisticated training techniques such as generative adversarial network [Park and Chang, 2019], reinforcement learning [Ding *et al.*, 2019] and bi-level optimization [Shu *et al.*, 2019]. This leads to the **first** trade-off between efficiency and effectiveness.

On the other hand, the soundness of these works resides on a problematic assumption that “all unlabeled interactions are true negative instances”, which is against the actual setting where unlabeled user-item pairs may potentially become positive instances, once the item is exposed to the user. These instances are termed as *false negatives*. The incorporation of false negatives would provide erroneous supervised signals for training and seriously degrade the performance [Hernández-Lobato *et al.*, 2014]. While it may sound attractive to identify and remove these instances, it is challenging to distinguish hard negatives and false negatives, given that both of them have large relevance scores in appearance and auxiliary information is often not available. Few works attempt to address this issue, especially in the context of negative mining for CF. This presents the **second** trade-off between informative negatives and reliable negatives.

Towards navigating these trade-offs, we propose a new framework named as Hardness-Aware Debiased Contrastive Collaborative Filtering (HDCCF). Specifically, a contrastive loss function is devised in place of conventional point-wise

*This work was supported by the National Key R&D Program of China [2020YFB1707903]; the National Natural Science Foundation of China [61872238, 61972254], the Shanghai Municipal Science and Technology Major Project [2021SHZDZX0102], and the Tencent Rhino-Bird Renewed Research Program. Xiaofeng Gao is the corresponding author.

and pair-wise objectives, which will be shown by gradient analysis that can automatically and adaptively concentrate on optimizing hard negatives by contrasting with peers, notably, without relying on complex training tricks. We also devise an efficient sampling strategy that implicitly explores negative instances by incorporating item frequency information, without actually conducting negative sampling. On top of the new hardness-aware objective, we further propose a principled method to eliminate the risk of false negatives. Needless to explicitly distinguish hard and false negatives, this is achieved by directly debiasing the objective, such that its expectation is strictly equivalent to the ideal loss function that resembles sampling under true negative distribution.

There are also several additional novel designs in our framework: 1) It considers both negative users and items to avoid the case where all negative items are relatively discriminative for a specific user, and vice versa; 2) Two auxiliary contrastive losses are introduced to model user-user and item-item relationship, which could help to obtain more meaningful user and item representations; 3) A neural modulated mechanism is designed that takes a user's diverse preference on different items into account in the loss function. We validate the effectiveness of HDCCF by comparison experiments and ablation studies. The results demonstrate the superiority of HDCCF as well as the effectiveness of its components.

2 Proposed Model

Notations. Let \mathcal{U} and \mathcal{I} denote a set of users and items, \mathcal{I}_u (resp. \mathcal{I}'_u) denote a set of items that user u has (resp. has not) interacted with, \mathcal{U}_i (resp. \mathcal{U}'_i) denote a set of users that item i have (resp. not) interacted with. Observed interaction data are represented by a set of user-item pairs $\mathcal{D} = \{(u, i)\}$. Unobserved user-item pairs are denoted as $\mathcal{D}' = \mathcal{U} \times \mathcal{I} - \mathcal{D}$. For user $u \in \mathcal{U}$, our goal is to recommend a fix-sized set of ordered items $\mathcal{X}_u \subset \mathcal{I}'_u$.

2.1 Hardness-Aware Contrastive CF

For an observed instance $(u, i) \in \mathcal{D}$, we uniformly sample S negative items (resp. users) that have no observed interaction with user u (resp. item i), denoted as $\mathcal{N}_u^- \subset \mathcal{I}'_u$ (resp. $\mathcal{N}_i^- \subset \mathcal{U}'_i$), where S is the *negative sampling number*. Then, the user-item contrastive loss \mathcal{L}_{u-i} is defined as

$$\mathcal{L}_{u-i} = - \sum_{(u,i) \in \mathcal{D}} \log \frac{e^{f(u,i)/\tau}}{F^i(\mathcal{N}_i^-) + F^u(\mathcal{N}_u^-) + e^{f(u,i)/\tau}}, \quad (1)$$

where $f : (u, i) \rightarrow \mathbb{R}$ is a similarity measure which outputs the relevance score of user u and item i , and $\tau \in \mathbb{R}^+$ is a scalar temperature parameter (omitted in the following for brevity). The F^u (resp. F^i) in Eqn. (1) is called *negative score* for user u (resp. item i), and is formulated as

$$F^u(\mathcal{N}_u^-) = \sum_{i^- \in \mathcal{N}_u^-} e^{f(u, i^-)}, \quad F^i(\mathcal{N}_i^-) = \sum_{u^- \in \mathcal{N}_i^-} e^{f(u^-, i)}, \quad (2)$$

where both (u, i^-) and (u^-, i) are unobserved (i.e., unlabeled) instances.

Hardness-Aware Property by Gradient Analysis. To explain the efficacy of the above contrastive loss formulation for hard negative mining, we peer into its gradients with respect to observed and unobserved instances for analysis. Denote the probability of an unobserved instance (u', i') being recognized as positive by

$$P(u', i') = \frac{e^{f(u', i')/\tau}}{F^i(\mathcal{N}_i^-) + F^u(\mathcal{N}_u^-) + e^{f(u, i)/\tau}}, \quad (3)$$

where $(u', i') \in \{(u, i^-)\}_{i^- \in \mathcal{N}_u^-} \cup \{(u^-, i)\}_{u^- \in \mathcal{N}_i^-}$

Then, the gradients with respect to the relevance score of observed and unobserved instances are computed as

$$\frac{\partial \mathcal{L}_{u-i}}{\partial f(u, i)} = -\frac{1}{\tau} \sum_{(u', i')} P(u', i'), \quad \frac{\partial \mathcal{L}_{u-i}}{\partial f(u', i')} = \frac{1}{\tau} P(u', i'), \quad (4)$$

These equations reveal two properties [Wang and Liu, 2021]:

1. For each individual unobserved instance (u', i') , the gradient is proportional to $P(u', i')$ and thus is also proportional to $e^{f(u', i')/\tau}$;
2. The gradient for observed instance is equal to the sum of gradients for all unobserved instances.

These properties have several implications in the context of hard negative mining in CF. **First**, according to the first property, a *harder* negative instance with larger relevance score has larger magnitude of gradients, which indicates the loss function could automatically concentrates on optimizing harder negative instances. The hardness level for each negative instance is adaptively updated for each iteration, and could be controlled by tuning temperature τ . **Second**, the gradient is re-scaled by the sum of relevance scores of peer negative instances, which indicates the hardness for each negative instance is relative to the hardness of peer negatives in the loss function, distinguishing us from pair-wise loss functions. **Third**, according to the second property, the gradients of negative instances, whose sum is determined by $\partial \mathcal{L}_{u-i} / \partial f(u, i)$, are distributed over each negative instance, and thus are not sensitive to label noise, which is known as a limitation of BPR loss. **Fourth**, by considering two types of negative instances for every (u, i) in Eqn. (1) (i.e., negative items and users), we could jointly mine negative instances from two facets and avoid the case when all negative items (resp. users) are easily discriminative for an individual user (resp. item).

User-User and Item-Item Relations. Besides modeling user-item interactions, we further extend the advantage of negative mining to neighbored users (i.e. a pair of users (u, u') that have interactions with the same item) and neighbored items [Sarwar *et al.*, 2001; Kabbur *et al.*, 2013] by proposing two auxiliary contrastive losses. We uniformly sample P positive items (resp. users) for user u (resp. item i), denoted as \mathcal{N}_u^+ (resp. \mathcal{N}_i^+), where P is the *positive neighbor sampling number*. The auxiliary loss function \mathcal{L}_{u-u} is formulated as:

$$\mathcal{L}_{u-u} = - \sum_{(u,i) \in \mathcal{D}} \sum_{u^+ \in \mathcal{N}_i^+} \log \frac{e^{f(u, u^+)}}{e^{f(u, u^+)} + \sum_{u^- \in \mathcal{N}_i^-} e^{f(u, u^-)}}, \quad (5)$$

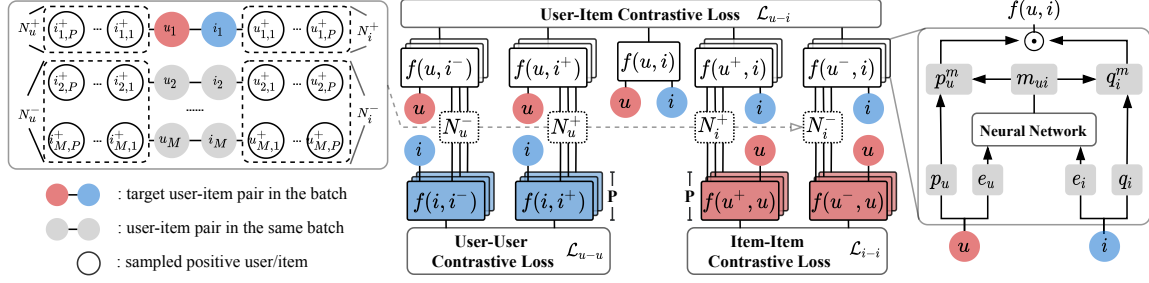


Figure 1: Overview of Hardness-Aware Debiased Contrastive Collaborative Filtering (HDCCF) framework.

where \mathcal{N}_u^- and \mathcal{N}_i^- are the same as those used in Eqn.(1), and \mathcal{L}_{i-i} could be defined in the same way. Then, the final loss function \mathcal{L} is the weighted sum of three terms:

$$\mathcal{L} = \mathcal{L}_{u-i} + \lambda_u \mathcal{L}_{u-u} + \lambda_i \mathcal{L}_{i-i}, \quad (6)$$

where λ_u and λ_i are weights to balance the importance for each type of relation.

Neural Modulated Similarity Model. For the similarity function $f(u, i)$, one could use dot-product [Koren *et al.*, 2009], Euclidean distance [Hsieh *et al.*, 2017] or parameterize it with neural network [He *et al.*, 2017]. Particularly in this paper, we propose to use the following neural modulated similarity model:

$$f(u, i) = (\mathbf{m}_{ui} \odot \mathbf{p}_u)^\top \cdot (\mathbf{m}_{ui} \odot \mathbf{q}_i), \quad (7)$$

where $\mathbf{p}_u, \mathbf{q}_i \in \mathbb{R}^d$ are user and item embeddings, and $\mathbf{m}_{ui} \in \mathbb{R}_+^d$ is a modulating vector for element-wise scaling, which is computed by

$$\mathbf{m}_{ui} = \sigma(g(\mathbf{e}_u \parallel \mathbf{e}_u \| (\mathbf{e}_i \odot \mathbf{e}_u))), \quad (8)$$

where $\mathbf{e}_u, \mathbf{e}_i \in \mathbb{R}^d$ are user and item latent factors, \odot denotes Hadamard product, \parallel denotes vector concatenation, and $g: \mathbb{R}^{3d} \rightarrow \mathbb{R}^d$ is a neural network. The key insight is that using a fixed user embedding may fail to represent one's diverse preference on distinct items in the loss function (items may also have multiple attributes that could attract a user), especially in our case where the objective incorporate more candidates of items. To mitigate this issue, our design could capture a user's varying preferences by allowing more flexible representations, which could improve the discrimination ability of the model, empirically verified by ablation studies.

2.2 Sampling Strategy

The sampling approach mentioned in last subsection samples S negative users/items and P positive users/items for each target observed user-item pair. Suppose the batch size is M , there are additional $2M \times (S + P)$ (where $M \approx S, S \gg P$) users/items to be sampled for each iteration besides target user-item pairs, which is impractical when scaling the training. Alternatively, we adopt a sampling strategy that uses the positive users (resp. items) from other observed instances in the same mini-batch as the negative users (resp. items) for the target instance. Formally,

$$\mathcal{N}_{u_k}^- = \bigcup_{m \in \{1, \dots, M\} \setminus k} \mathcal{N}_{u_m}^+, \quad \mathcal{N}_{i_k}^- = \bigcup_{m \in \{1, \dots, M\} \setminus k} \mathcal{N}_{i_m}^+, \quad (9)$$

where $\mathcal{N}_{u_k}^-$ and $\mathcal{N}_{i_k}^-$ are multi-sets, which allow multiple appearances of the same item or user in the set. In this way, the negative instance number is enlarged from S to $P \times (M - 1)$ with lower sampling overhead (from $2M \times (S + P)$ to $2M \times P$). We also highlight that such sampling strategy is free of explicit negative sampling.

Frequency-Aware Properties. To further shed lights on HDCCF's hard negative mining capability and its relation with frequency-based sampling methods, we investigate on: for a specific positive user-item interaction (u, i) in a mini-batch, 1) the number of times a negative item i' (resp. u') appears in user-item contrastive loss, denoted as $n_{i'}^{u-i}$ (resp. $n_{u'}^{u-i}$) and 2) the number of times a negative item i' (resp. u') appears in auxiliary contrastive losses, denoted as $n_{i'}^{i-i}$ (resp. $n_{u'}^{u-u}$).

Proposition 1. *The expectation of $n_{i'}^{u-i}$ is proportional to the number of times this item appears in interaction dataset \mathcal{D} , i.e., $n_{i'}^{u-i} \propto |\mathcal{U}_{i'}|$. This property holds true for user u' in \mathcal{L}_{u-i} . Formally we have*

$$\begin{aligned} \mathbb{E}_{\substack{(u,i) \sim p^o \\ i^+ \sim p_u^o}} [n_{i'}^{u-i}] &= \frac{M-1}{|\mathcal{D}|-1} \cdot P \cdot |\mathcal{U}_{i'}|, \\ \mathbb{E}_{\substack{(u,i) \sim p^o \\ u^+ \sim p_i^o}} [n_{u'}^{u-i}] &= \frac{M-1}{|\mathcal{D}|-1} \cdot P \cdot |\mathcal{I}_{u'}|, \end{aligned} \quad (10)$$

where the observed interaction (u, i) is sampled from \mathcal{D} with distribution $p^o(u, i)$, and the neighbored item i^+ is sampled from \mathcal{I}_u with distribution $p_u^o(i^+)$.

Proposition 2. *The expectation of $n_{i'}^{i-i}$ is also proportional to $|\mathcal{I}_{u'}|$. This property holds true for user u' in user-user contrastive loss \mathcal{L}_{u-u} . Formally we have*

$$\begin{aligned} \mathbb{E}_{\substack{(u,i) \sim p^o \\ i^+ \sim p_u^o}} [n_{i'}^{i-i}] &= \frac{M-1}{|\mathcal{D}|-1} \cdot P \cdot |\mathcal{U}_{i'}|, \\ \mathbb{E}_{\substack{(u,i) \sim p^o \\ u^+ \sim p_i^o}} [n_{u'}^{u-u}] &= \frac{M-1}{|\mathcal{D}|-1} \cdot P \cdot |\mathcal{I}_{u'}|, \end{aligned} \quad (11)$$

The proof of these propositions is shown in the appendix. As an observation, both $\mathbb{E}[n_{i'}^{u-i}]$ and $\mathbb{E}[n_{i'}^{i-i}]$ are proportional to $|\mathcal{U}_{i'}|$, which indicates that such sampling strategy is essentially frequency-aware, which enforces the loss function to concentrate on more frequent (popular) items. Since popular items are treated as harder negative instances [Chen *et al.*, 2017; Wu *et al.*, 2019a], such sampling strategy implicitly agrees with the negative mining efficacy of HDCCF.

2.3 Debiased Contrastive Loss

As mentioned before, a user-item pair (u, i^-) that we regard as a negative instance is potentially a positive interaction (i.e., *false negative instance*). The existence of false negative instance could introduce bias in the training signals, and hence may cause sub-optimal results. Particularly, in our case, there are two types of false negative instances: 1) The user-item pair is an observed instance, i.e., $(u, i^-) \in \mathcal{D}$; 2) Though (u, i^-) is unobserved, the interaction will occur once i^- is exposed to u . In a similar spirit with [Robinson *et al.*, 2021] that considers a simpler case (without the first type of false negatives and user-user/item-item losses) in the general contrastive learning setting, we propose to eliminate the effects of false negatives by first formulating the expected loss function and then devising unbiased versions of Eqn. (1) and Eqn. (5) without violating their hardness-aware properties.

Formulation of Expected Loss Function. Our analysis mainly focus on the user side for brevity, while the same also applies to the item side. Given an item i , suppose a negative user u^- is drawn from \mathcal{U}_i^- with a *negative sampling distribution* $p_i(u^-)$, i.e., a uniform distribution. Drawing from this distribution may either yield a false negative instance or a real negative instance. Suppose their probabilities are ω_u^+ and ω_u^- (i.e., $1 - \omega_u^+$) respectively. To investigate on the formulation of expected loss function, we denote $p_i^+(u^-)$ (resp. $p_i^-(u^-)$) as the sampling distribution for false (resp. real) negative instance, which are unknown for us. The marginalization of the negative sampling distribution induces a decomposition form $p_i(u^-) = \omega_u^+ \cdot p_i^+(u^-) + \omega_u^- \cdot p_i^-(u^-)$. Reorganizing it yields the following expression for real negative sampling distribution

$$p_i^-(u') = \begin{cases} 0 & , (u', i) \in \mathcal{D}, \\ \frac{p_i(u')}{\omega_u^-} - \frac{\omega_u^+ \cdot p_i^+(u')}{\omega_u^-} & , (u', i) \notin \mathcal{D}. \end{cases} \quad (12)$$

Equipped with these notations, we can formulate the ideal optimization objective for \mathcal{L}_{u-i} as

$$\mathcal{L}_{u-i}^{ideal} = - \mathbb{E}_{(u,i) \sim p^o} \left[\frac{e^{f(u,i)}}{e^{f(u,i)} + Q \mathbb{E}_{u^- \sim p_i^-} [e^{f(u^-,i)}] + Q \mathbb{E}_{i^- \sim p_u^-} [e^{f(u,i^-)}]} \right], \quad (13)$$

where Q is constant to facilitate the analysis. By comparison between the original formulation of optimization objective in Eqn. (1) and the expected objective in Eqn. (13), we immediately notice the bias essentially stems from the underlying negative sampling distribution. Eliminating the effects of false negatives boils down to approximating the ideal optimization objective using biased observations in datasets. This is challenging due to the existence of two types of false negatives as stated before, and the intractability of the real negative sampling distribution.

Modification of Eqn.(1) Toward eliminating the effects of false negatives, we proceed to modify both formulations of Eqn. (1) and Eqn. (5) such that they agree with the ideal opti-

mization objective. Specifically, the debiased user-item contrastive loss can be formulated as:

$$\tilde{\mathcal{L}}_{u-i} = - \sum_{(u,i) \in \mathcal{D}} \log \frac{e^{f(u,i)}}{\tilde{F}^u(\mathcal{N}_i^-, \mathcal{N}_i^+) + \tilde{F}^i(\mathcal{N}_u^-, \mathcal{N}_u^+) + e^{f(u,i)}}, \quad (14)$$

where $\tilde{F}^u(\mathcal{N}_i^-, \mathcal{N}_i^+)$ and $\tilde{F}^i(\mathcal{N}_u^-, \mathcal{N}_u^+)$ are debiased negative scores for negative users and negative items respectively, and the former one is defined as:

$$\sum_{u^- \in \mathcal{N}_i^-} \pi_0^u(u^-, i) \cdot e^{f(u^-, i)} - \sum_{u^+ \in \mathcal{N}_i^+ \cup \{u\}} \pi_1^u(u^+, i) \cdot e^{f(u^+, i)}, \quad (15)$$

where $\pi_0^u(u^-, i), \pi_1^u(u^+, i) \in \mathbb{R}$ are constants w.r.t. $|\mathcal{N}_i^-|, |\mathcal{N}_i^+|$ and ω_u^+ . Their exact formulations will be given in the appendix. By replacing u by i , we can get $\tilde{F}^i(u, i, \mathcal{N}_i^-, \mathcal{N}_i^+)$ in the same way.

Theorem 1. Equation. (14) is an unbiased estimation of the ideal user-item contrastive loss where negative instances are drawn from the real negative distributions p_i^- and p_u^- .

The proof is shown in appendix. To prevent negative values in the logarithm, we can constrain the negative scores to be greater than its theoretical lower bound in practice

$$\begin{aligned} \tilde{F}^u(\mathcal{N}_i^-, \mathcal{N}_i^+) &\leftarrow \max \{ \tilde{F}^u(\mathcal{N}_i^-, \mathcal{N}_i^+), |\mathcal{N}_i^-| e^{1/\tau} \}, \\ \tilde{F}^i(\mathcal{N}_u^-, \mathcal{N}_u^+) &\leftarrow \max \{ \tilde{F}^i(\mathcal{N}_u^-, \mathcal{N}_u^+), |\mathcal{N}_u^-| e^{1/\tau} \}. \end{aligned} \quad (16)$$

Modification of Eqn.(5) In the similar spirit, we develop unbiased formulations for user-user and item-item contrastive losses \mathcal{L}_{u-u} and \mathcal{L}_{i-i} in the following forms

$$\tilde{\mathcal{L}}_{u-u} = - \sum_{(u,i) \in \mathcal{D}} \sum_{u^+ \in \mathcal{N}_i^+} \log \frac{e^{f(u,u^+)}}{e^{f(u,u^+)} + \sum_{u^- \in \mathcal{N}_i^-} \pi^{u-u}(u, u^-) \cdot e^{f(u,u^-)}}, \quad (17)$$

\mathcal{D}^{u-u} is a set of observed neighbored users. By replacing u by i , we can get $\tilde{\mathcal{L}}_{i-i}$ in the same way.

3 Related Works

Negative mining plays an important role in CF approaches, with significant influences on the recommendation performance [Chen *et al.*, 2017; Rendle *et al.*, 2009]. A high-quality negative instance should satisfy: 1) It should be hard for model to discriminate, so as to provide useful information for training [Wu *et al.*, 2017]; 2) It should be a reliable negative instance sampled from the distribution of true negatives, rather than those erroneously recognized as negatives [Hernández-Lobato *et al.*, 2014]. Most existing works on hard negative mining [Ding *et al.*, 2019; Ding *et al.*, 2018; Park and Chang, 2019] either learns the sampling distribution with a separate model, or generates negative instances with adversarial training. Despite of the promising results, they often require complex designs, architectures, or side information such as user’s “view” behavior which is not always available [Ding *et al.*, 2018]. Besides, the risk of false negative instances are overlooked in these works. Deviating from

Datasets	Metrics	PMF	BPR	SVD++	NeuMF	ENMF	IRGAN	SD-GAR	HDCCF	Imp.
ML-1M	HR@10	0.7109	0.7162	0.7230	0.6991	0.7273	0.7205	<u>0.7323</u>	0.7596	3.72%
	NDCG@10	0.4396	0.4435	0.4496	0.4283	0.5193	0.4705	<u>0.5320</u>	0.5788	8.80%
	NDCG@50	0.5048	0.5069	0.5132	0.4936	0.5731	0.5310	<u>0.5842</u>	0.6233	6.69%
Yelp	HR@10	0.3048	0.3093	0.3268	0.3230	<u>0.3710</u>	0.3197	0.3459	0.3911	5.41%
	NDCG@10	0.1609	0.1644	0.1830	0.1835	<u>0.2212</u>	0.1776	0.1929	0.2413	9.09%
	NDCG@50	0.2181	0.2209	0.2413	0.2453	<u>0.2763</u>	0.2470	0.2684	0.2983	7.96%
Gowalla	HR@10	0.7341	0.7402	0.7303	0.7456	0.7895	0.7721	<u>0.8060</u>	0.8311	3.11%
	NDCG@10	0.5353	0.5411	0.5292	0.5498	0.6124	0.5894	<u>0.6321</u>	0.6633	4.94%
	NDCG@50	0.5693	0.5736	0.5672	0.5777	0.6390	0.6102	<u>0.6631</u>	0.6895	3.98%

Table 1: Experiment results of HDCCF and competitors. The bold value marks the best one in one row, while the underlined value corresponds to the best one among all the baselines. Improvements are statistically significant with $p < 0.01$.

Datasets	#Users	#Items	#Interactions	Density
MovieLens	6039	3415	999611	4.847%
Yelp	23056	15575	648687	0.181%
Gowalla	72454	56173	1360493	0.033%

Table 2: Statistics of three datasets.

those works, we devise a hardness-aware loss function that can automatically detect hard and reliable negatives.

Existing optimization objectives in CF approaches could be roughly categorized into point-wise [Mnih and Salakhutdinov, 2007], pair-wise [Rendle *et al.*, 2009] and list-wise [Wu *et al.*, 2018]. The most relevant approaches are based on list-wise, which also consider multiple instances in the loss function. However, state-of-the-art list-wise approach [Wu *et al.*, 2018] based on a permutation probability only optimize the upper bound rather than the original negative log-likelihood. Another related work in CF [Chen *et al.*, 2020] that is also free of negative sampling essentially relies on a pair-wise objective with predefined frequency-based weight. Collaborative filtering has been extensively studied in other recommendation situations, e.g., social recommendation [Ma *et al.*, 2008; Wu *et al.*, 2019b], sequential recommendation [Kang and McAuley, 2018; Wu *et al.*, 2021], multi-task learning [Ma *et al.*, 2018; Yang *et al.*, 2022], etc. While our paper mainly focuses on the general setting where only the user-item interactions are assumed as input, the proposed methodology can be trivially extended to other cases to incorporate more information.

4 Experiments

Datasets. We perform extensive experiments on three publicly accessible datasets from various domains with different sparsities. MovieLens [Harper and Konstan, 2015] is a widely adopted benchmark dataset for collaborative filtering. We use two versions, namely MovieLens(ML)-100K and MovieLens(ML)-1M. Yelp is a dataset of user ratings on businesses, and we use the filtered subset created by [He *et al.*, 2016] for evaluation. Gowalla [Cho *et al.*, 2011] is collected from a popular location-based social network, which contains users’ check-in history with time spanning from February 2009 to October 2010. Each user’s interactions are sorted by the timestamps ascendingly. Then the testing data (resp. validation data) comprise the last (resp. second to last) interacted item of each user, while the remaining are used as training data. The statistics of filtered datasets are given in Table 2.

Evaluation Protocol and Metrics. Following [He *et al.*, 2017; Tay *et al.*, 2018; Rendle *et al.*, 2009], we adopt the *leave one out* protocol for model evaluation. We evaluate the ranking performance of the proposed model based on two widely used metrics: *Hit Ratio at K* (HR@K), and *Normalized Discounted Cumulative Gain at K* (NDCG@K).

Competitors. We consider nine baseline models for collaborative filtering, including four classic collaborative filtering models (PMF [Koren *et al.*, 2009], BPR [Rendle *et al.*, 2009], SVD++ [Koren, 2008], NeuMF [He *et al.*, 2017]), a non-sampling approach ENMF [Chen *et al.*, 2020] and two hard negative mining methods (IRGAN [Wang *et al.*, 2017], SD-GAR [Jin *et al.*, 2020]).

4.1 Performance Comparison

We report experiment results of HDCCF and other comparative models in Tab. 1. As we can see, the proposed HDCCF consistently outperforms other comparative methods and achieve state-of-the-art results w.r.t. different metrics throughout three dataset. Specifically, HDCCF on average achieves 4.08% improvement for HR@10, 7.61% improvement for NDCG@10 and 6.21% improvement for NDCG@50. The results demonstrate that HDCCF is a powerful approach for recommendation in implicit feedback. There are some other findings. First, four classic methods (i.e., PMF, BPR, SVD++, NeuMF) have the worst performance, which implies the important role of negative mining on promising recommendation performance. Second, SVD++ has better performance than PMF on average since it considers user neighbor information while PMF fails to do so, which also justifies the design of our auxiliary losses. Third, HDCCF outperforms adversarial hard negative mining methods, which demonstrate the effectiveness of HDCCF for mining reliable and hard negative instances with more simple and flexible designs.

4.2 Ablation Study

We conduct a series of ablation studies to investigate the necessities of some key components in our model and how these components contribute to the overall results.

Debiased Contrastive Losses. We compare the performance of HDCCF and five variants of it with simplified loss functions by removing or replacing the contrastive losses. Specifically, we have two findings in Fig. 2. First, our HDCCF consistently outperforms the variants of HDCCF which




Figure 2: Ablation studies on MovieLens for HDCCF.




Figure 3: Positive and negative user-item relevance score distributions on ML-100K dataset.

remove $\tilde{\mathcal{L}}_{i-i}$ (w/o-I), $\tilde{\mathcal{L}}_{u-u}$ (w/o-U) and both auxiliary contrastive losses (w/o-U&I), respectively. This is because $\tilde{\mathcal{L}}_{i-i}$ and $\tilde{\mathcal{L}}_{u-u}$ can take the advantage of hard negative mining to fully exploit the relations of neighbored users and items. Second, the variant of HDCCF that only preserves the user-item contrastive loss (w/o-U&I) also consistently outperforms the variant that replaces the contrastive loss by BPR loss (BPR+). This result validates that the good recommendation attributes to the proposed loss functions rather than other model designs. Third, HDCCF outperforms the biased variant by a large margin, which demonstrate the importance of removing false negative instances especially in our case where false negatives could be either observed or unobserved.

Modulated Similarity Model. We compare HDCCF with a variant with a fixed modulating vector, i.e., an all-ones modulating vector. Specifically, we visualize the distributions of user-item relevance scores (which are normalized to range 0 to 100) for two models on ML-100K dataset in Fig. 3. As shown in the figure, the modulated similarity model pushes the upper bound of negative instances’ relevance scores to the left, and the lower bound of positive instances’ relevance scores to the right (which is more significant than the former one). Consequently, the intersecting interval (which is filled with red shade in the figure) of our modulated similarity model is considerably smaller than the unmodulated variant. This empirical result conforms to the conjecture that the modulation mechanism is helpful to distinguish positive interactions against negative instances by using a more flexible user and item representations.

4.3 Hyper-parameter Analysis

Impact of Bias Correction Probabilities. Two bias correction probabilities ω_i^+ and ω_u^+ are introduced in Sec.2.3 and denote the probabilities that items and users sampled from

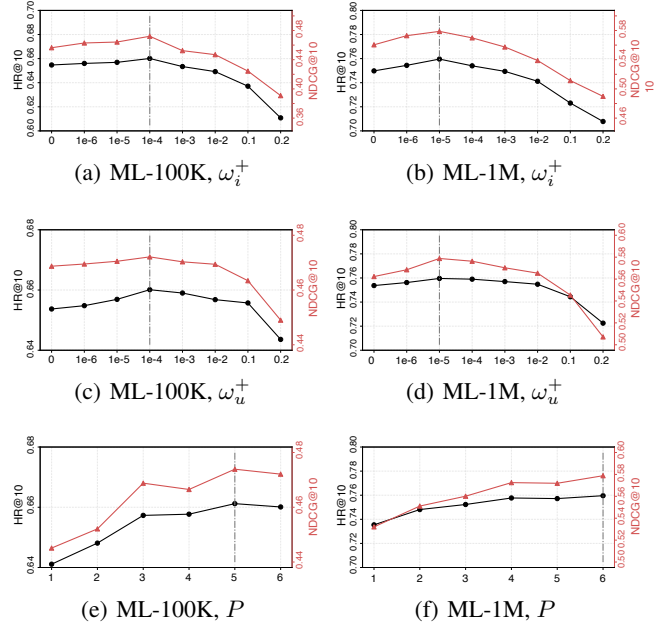


Figure 4: Performance of HDCCF on MovieLens w.r.t. different hyper-parameters. The vertical dotted lines mark the best value.

unobserved data are true negative samples. It is intractable to derive the precise values of ω_i^+ and ω_u^+ , and they are different across datasets. Therefore, we treat them as hyper-parameters in implementation and here study how these parameters affect the recommendation performance. Figure 4(a)-(d) show HR@10 and NDCG@10 of HDCCF with ω_i^+ and ω_u^+ varying from 0 to 0.2 and other hyper-parameters unchanged. As we can see, a certain range of values for ω_i^+ and ω_u^+ can correct the sampling bias, while out-of-range ω_i^+ and ω_u^+ will harm the recommendation performance.

Impact of Positive Neighbor Sampling Number. For our neighbor sampling strategy, P is the hyper-parameter denoting the number of positive user (or item) neighbors for each interaction in the batch. Intuitively, larger P brings higher time and space costs, but improves the stability of training process. To validate this conjecture, we study how P affects the recommendation performance. Figure 4(e)-(f) show the performance of HDCCF with P varying from 1 to 6 and other hyper-parameters unchanged. As we can see, larger P can bring up improvements for recommendation performance, since it can reduce the variance brought by sampling over observed interactions (which helps to stabilize the training process and alleviate over-fitting) and improve the hard negative mining ability by contrasting with more instances.

5 Conclusion

In this paper, we propose hardness-aware debiased contrastive collaborative filtering framework, which resolves the dilemma of hard negative mining and the reliability of negative instances. Comprehensive experiments and ablation studies on three real-world datasets demonstrate the effectiveness of our framework.

References

- [Chen *et al.*, 2017] Ting Chen, Yizhou Sun, Yue Shi, and Liangjie Hong. On sampling strategies for neural network-based collaborative filtering. In *SIGKDD*, 2017.
- [Chen *et al.*, 2020] Chong Chen, Min Zhang, Yongfeng Zhang, Yiqun Liu, and Shaoping Ma. Efficient neural matrix factorization without sampling for recommendation. *TOIS*, 38(2), 2020.
- [Cho *et al.*, 2011] Eunjoon Cho, Seth A Myers, and Jure Leskovec. Friendship and mobility: user movement in location-based social networks. In *SIGKDD*, 2011.
- [Ding *et al.*, 2018] Jingtao Ding, Fuli Feng, Xiangnan He, Guanghui Yu, Yong Li, and Depeng Jin. An improved sampler for bayesian personalized ranking by leveraging view data. In *WWW Companion*, 2018.
- [Ding *et al.*, 2019] Jingtao Ding, Yuhan Quan, Xiangnan He, Yong Li, and Depeng Jin. Reinforced negative sampling for recommendation with exposure data. In *IJCAI*, 2019.
- [Harper and Konstan, 2015] F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. In *ACM TIS*, 2015.
- [He *et al.*, 2016] Xiangnan He, Hanwang Zhang, Min-Yen Kan, and Tat-Seng Chua. Fast matrix factorization for online recommendation with implicit feedback. In *SIGIR*, 2016.
- [He *et al.*, 2017] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *WWW*, pages 173–182, 2017.
- [Hernández-Lobato *et al.*, 2014] José Miguel Hernández-Lobato, Neil Houlsby, and Zoubin Ghahramani. Probabilistic matrix factorization with non-random missing data. In *ICML*, 2014.
- [Hsieh *et al.*, 2017] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. Collaborative metric learning. In *WWW*, pages 193–201, 2017.
- [Jin *et al.*, 2020] Binbin Jin, Defu Lian, Zheng Liu, Qi Liu, Jianhui Ma, Xing Xie, and Enhong Chen. Sampling-decomposable generative adversarial recommender. In *NeurIPS*, 2020.
- [Kabbur *et al.*, 2013] Santosh Kabbur, Xia Ning, and George Karypis. Fism: factored item similarity models for top-n recommender systems. In *SIGKDD*, pages 659–667, 2013.
- [Kang and McAuley, 2018] Wang-Cheng Kang and Julian McAuley. Self-attentive sequential recommendation. In *ICDM*, 2018.
- [Koren *et al.*, 2009] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. In *Computer*, pages 30–37, 2009.
- [Koren, 2008] Yehuda Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *SIGKDD*, 2008.
- [Ma *et al.*, 2008] Hao Ma, Haixuan Yang, Michael R Lyu, and Irwin King. Sorec: social recommendation using probabilistic matrix factorization. In *CIKM*, 2008.
- [Ma *et al.*, 2018] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *SIGKDD*, pages 1930–1939, 2018.
- [Mnih and Salakhutdinov, 2007] Andriy Mnih and Russ R Salakhutdinov. Probabilistic matrix factorization. In *NeurIPS*, 2007.
- [Mnih and Salakhutdinov, 2008] Andriy Mnih and Russ R Salakhutdinov. Probabilistic matrix factorization. In *NeurIPS*, pages 1257–1264, 2008.
- [Pan *et al.*, 2008] Rong Pan, Yunhong Zhou, Bin Cao, Nathan N Liu, Rajan Lukose, Martin Scholz, and Qiang Yang. One-class collaborative filtering. In *ICDM*, 2008.
- [Park and Chang, 2019] Dae Hoon Park and Yi Chang. Adversarial sampling and training for semi-supervised information retrieval. In *WWW*, 2019.
- [Rendle and Freudenthaler, 2014] Steffen Rendle and Christoph Freudenthaler. Improving pairwise learning for item recommendation from implicit feedback. In *WSDM*, 2014.
- [Rendle *et al.*, 2009] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, pages 452–461, 2009.
- [Robinson *et al.*, 2021] Joshua Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. *ICLR*, 2021.
- [Sarwar *et al.*, 2001] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *WWW*, pages 285–295, 2001.
- [Shu *et al.*, 2019] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. *NeurIPS*, 2019.
- [Tay *et al.*, 2018] Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. Latent relational metric learning via memory-based attention for collaborative ranking. In *WWW*, pages 729–739, 2018.
- [Wang and Liu, 2021] Feng Wang and Huaping Liu. Understanding the behaviour of contrastive loss. In *CVPR*, pages 2495–2504, 2021.
- [Wang *et al.*, 2017] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *SIGIR*, pages 515–524, 2017.
- [Wu *et al.*, 2017] Chao-Yuan Wu, R Manmatha, Alexander J Smola, and Philipp Krahenbuhl. Sampling matters in deep embedding learning. In *ICCV*, 2017.
- [Wu *et al.*, 2018] Liwei Wu, Cho-Jui Hsieh, and James Sharpnack. Sql-rank: A listwise approach to collaborative ranking. In *ICML*, 2018.
- [Wu *et al.*, 2019a] Ga Wu, Maksims Volkovs, Chee Loong Soon, Scott Sanner, and Himanshu Rai. Noise contrastive estimation for one-class collaborative filtering. In *SIGIR*, 2019.
- [Wu *et al.*, 2019b] Qitian Wu, Hengrui Zhang, Xiaofeng Gao, Peng He, Paul Weng, Han Gao, and Guihai Chen. Dual graph attention networks for deep latent representation of multifaceted social effects in recommender systems. In *WWW*, 2019.
- [Wu *et al.*, 2021] Qitian Wu, Chenxiao Yang, Shuodan Yu, Xiaofeng Gao, and Guihai Chen. Seq2bubbles: Region-based embedding learning for user behaviors in sequential recommenders. In *CIKM*, 2021.
- [Yang *et al.*, 2022] Chenxiao Yang, Junwei Pan, Xiaofeng Gao, Tingyu Jiang, Dapeng Liu, and Guihai Chen. Cross-task knowledge distillation in multi-task recommendation. *AAAI*, 2022.