



Improving Multi-modal Recommender Systems by Denoising and Aligning Multi-modal Content and User Feedback

Guipeng Xv*
xuguipeng@stu.xmu.edu.cn
School of Informatics, Xiamen
University
Xiamen, China

Xinyu Li
xinyuli@stu.xmu.edu.cn
School of Informatics, Xiamen
University
Xiamen, China

Ruobing Xie
xrbsnowing@163.com
Tencent
Beijing, China

Chen Lin†
chenlin@xmu.edu.cn
School of Informatics, Xiamen
University
Xiamen, China

Chong Liu
nickcliu@tencent.com
Tencent
Beijing, China

Feng Xia
xiafengxia@tencent.com
Tencent
Beijing, China

Zhanhui Kang
kegokang@tencent.com
Tencent
Beijing, China

Leyu Lin
goshawklin@tencent.com
Tencent
Beijing, China

Abstract

Multi-modal recommender systems (MRSs) are pivotal in diverse online web platforms and have garnered considerable attention in recent years. However, previous studies overlook the challenges of (1) noisy multi-modal content, (2) noisy user feedback, and (3) aligning multi-modal content and user feedback. To tackle these challenges, we propose **Denoising and Aligning Multi-modal Recommender System (DA-MRS)**. To mitigate noise in multi-modal content, **DA-MRS** first constructs item-item graphs determined by consistent content similarity across modalities. To denoise user feedback, **DA-MRS** associates the probability of observed feedback with multi-modal content and devises a denoised BPR loss. Furthermore, **DA-MRS** implements Alignment guided by User preference to enhance task-specific item representation and Alignment guided by graded Item relations to provide finer-grained alignment. Extensive experiments verify that **DA-MRS** is a plug-and-play framework and achieves significant and consistent improvements across various datasets, backbone models, and noisy scenarios.

CCS Concepts

• **Information systems** → **Recommender systems; Multimedia and multimodal retrieval.**

*Work done during internship at 2023 Tencent Rhino-Bird Research Elite Program.

†Corresponding author.

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 the author(s) 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.

KDD '24, August 25–29, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0490-1/24/08

<https://doi.org/10.1145/3637528.3671703>

Keywords

Multi-modal Recommender System, Noisy Multi-modal Content, Noisy User Feedback, Aligning Multi-modal Content and User Feedback

ACM Reference Format:

Guipeng Xv, Xinyu Li, Ruobing Xie, Chen Lin, Chong Liu, Feng Xia, Zhanhui Kang, and Leyu Lin. 2024. Improving Multi-modal Recommender Systems by Denoising and Aligning Multi-modal Content and User Feedback. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24), August 25–29, 2024, Barcelona, Spain*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3671703>

1 Introduction

Recommender systems (RSs) are pivotal in diverse online web platforms [1, 2, 5, 18, 23, 33], which typically provide abundant multi-modal content information, i.e., pictures and textual descriptions of items. The multi-modal information can supplement user feedback, alleviate the data sparsity problem in user feedback, and improve recommendation performance [36]. Thus, multi-modal recommender systems (MRSs) have garnered significant attention in recent years [14, 19, 24, 34, 35, 39].

Conventional MRSs are *feature-based methods* that integrate multi-modal features to enhance item representation [9, 19, 25, 26, 39], e.g., merging multi-modal features with item embeddings derived by matrix factorization [9], or refining the propagation on user-item interaction graph [19, 25, 26, 39]. These methods use multi-modal content to establish individual item ground knowledge, while item-item collaborative relations are captured through high-order item-user-item relations [34]. Recently, *structure-based methods* that explicitly extract collaborative item-item relations from multi-modal content have shown superior to feature-based methods [14, 32, 34, 35, 37, 38]. They commonly involve two major steps. The first step is to construct an item-item similarity graph from multi-modal content, while the second step is to fuse the

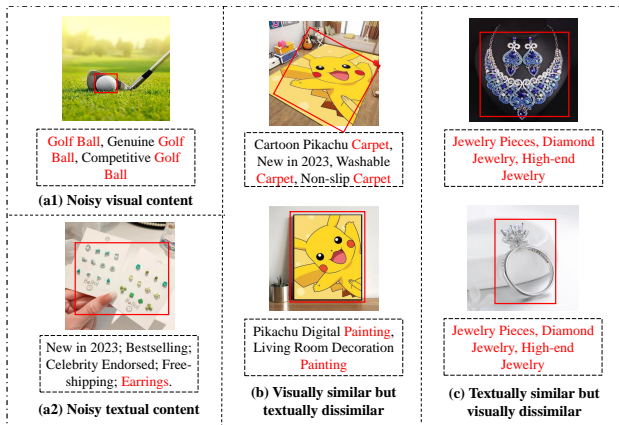


Figure 1: An illustration of multi-modal content in MRSs. (a1) The presence of a golf course background and a golf club in the image of a golf ball potentially distorts the understanding of the item. (a2) The textual description of some earrings contains many irrelevant words. (b) The carpet and the painting are visually similar, but their usages are different. (c) The two jewelry pieces have identical textual descriptions. One has a simple look, and the other has an intricate design. They are usually preferred by different people.

item-item similarities with the user-item interactions. Since similar items are likely to attract users with alike tastes, the item-item graph assists in identifying candidate items and contributes to the collaborative learning process.

Although structure-based methods have shown promising results, three intertwined challenges remain under-explored.

C1: Noisy Multi-modal Content. Multi-modal content usually contains details irrelevant to the item (e.g., Figure 1 shows some examples of noisy multi-modal content). It is also possible that the multi-modal content is inconsistent due to system errors, e.g., the unauthorized use of unrelated pictures and textual descriptions to describe an item to gain attention. Existing structure-based methods tend to ignore the noise present in multi-modal content, leading to the inclusion of false positive links that connect dissimilar items in the constructed item-item graphs [32]. As a result, the inaccurate item-item graph can disrupt item representation learning, which ultimately leads to a decrease in recommendation performance.

C2: Noisy User Feedback. User feedback data is often contaminated with noise, such as various kinds of bias [3, 12, 29] and erroneous clicks [13, 30]. MRSs rely on user feedback as supervision signals and noisy user feedback will hinder MRSs from learning actual user preferences. Although several efforts have been made to address noisy feedback data in the pure collaborative filtering setting [4, 6, 20, 27], none is designed to utilize multi-modal content. Since multi-modal content is perceived by the user before the occurrence of an actual interaction, e.g., a fashion-goer prefers trending elements in the displayed image, it can be naturally used to assess the confidence of an observed interaction. Nonetheless, due to the noise in multi-modal content, it is crucial to carefully design the denoising process for user feedback when leveraging multi-modal content.

C3: Aligning Multi-modal Content and User Feedback. Most MRSs adopt item-level alignment, i.e., they match the multi-modal content of each item by contrastive learning [19, 35, 39]. However, such alignment is insufficient. (1) *It is limited to improving task-specific understanding of the items.* The goal of RS is centered around user preference. Although item-level alignment can improve the general understanding of the items, it pays little attention to how the multi-modal content aligns from the perspective of each user’s preference. (2) *It fails to distinguish items at a finer granularity.* With multi-modal content, the item-item relations demonstrate graded similarities [19, 35], i.e., items that are similar in multiple modalities have stronger correlations than items that are similar in a single modality. Existing methods oversimplify item similarity by assuming items are either similar or different and fail to capture the subtle differences in more similar items.

We propose a framework called Denoising and Aligning Multi-modal Recommender System (DA-MRS). DA-MRS improves over other structure-based methods in addressing the three challenges. In item-item graph construction, to deal with the noisy multi-modal content, DA-MRS first constructs multiple modality-specific Item-item Semantic Graph with more accurate links by considering consistent similarities across modalities. DA-MRS also constructs an Item-item Behavior Graph to compensate for the pure content similarity and provide more comprehensive and reliable item-item collaborative relations. In learning user and item representations, to eliminate the impact of noisy feedback, DA-MRS defines the probabilistic generation of feedback signals. By associating the probability of observed feedback signals with estimated user preference from multi-modal content, DA-MRS derives a denoised version of the commonly employed BPR loss. To achieve task-specific alignment, DA-MRS employs Alignment guided by User preference to minimize the gap between each user’s preference distribution over items inferred from multi-modal content and feedback signals. To achieve finer-grained alignment, DA-MRS further implements Alignment guided by graded Item relations to contrast most similar (i.e., across all modalities), less similar (i.e., on a single modality), and dissimilar items.

In summary, our contributions are four-fold:

- (1) We propose a novel solution to obtain accurate item-item structures based on multi-modal consistency.
- (2) We point out that the problem of noisy feedback can be solved based on multi-modal content and propose a probabilistic generative model with a strong theoretical basis to solve the problem.
- (3) We present more effective multi-modal alignment in MRSs, i.e., Alignment guided by User preference enhances task-specific item representations, and Alignment guided by graded Item relations provides finer-grained alignment.
- (4) DA-MRS is a plug-and-play framework, and extensive experiments have demonstrated its effectiveness in significantly and consistently improving recommendation performance across various datasets, backbone models, and noisy scenarios.

2 Related Work

Multi-modal Recommender Systems. Conventional multi-modal recommender systems (MRSs) are *feature-based methods*. The multi-modal features are integrated by either (1) *direct fusion*, e.g., VBPR [9] directly concatenates visual embeddings with ID embeddings, (2)

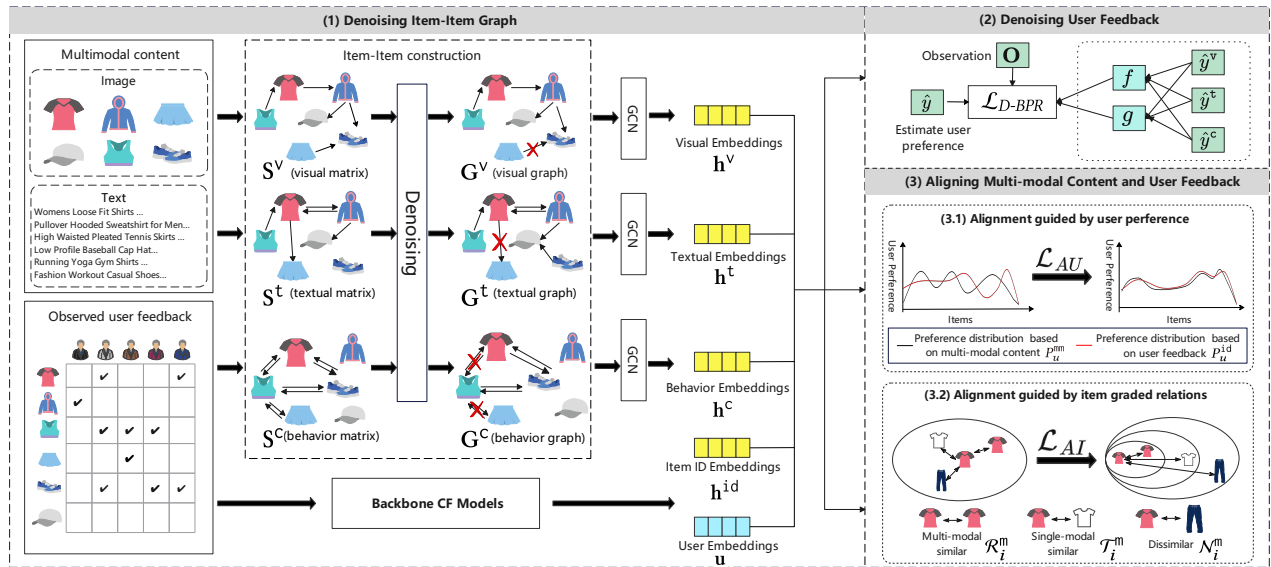


Figure 2: Overall framework of DA-MRS. It consists of three major components, i.e., Denoising Item-item Graph, Denoising User Feedback, and Aligning Multi-modal Content and User Feedback

graph neural network, e.g., MMGCN [26], GRCN [25] and SLM-Rec [19] perform graph convolutions to fuse multi-modal content and ID embeddings, (3) *contrastive learning*, e.g., BM3 [39] fuse multi-modal content and ID embeddings by contrastive learning, or (4) *adversarial learning*, e.g., MMSL [24] applies adversarial learning to fuse multi-modal information.

Recent *structure-based methods* first construct an item-item similarity graph, where the graph construction methods are roughly identical. Then, different methods are used to integrate the item-item graph and user-item interaction. For example, LATTICE [34] incorporates collaborative filtering (CF) approaches with normalized item semantic embeddings; MICRO [35] devises a contrastive framework to fuse multimodal item relationships; FREEDOM [38] freezes the item-item multi-modal graph and denoises the user-item graph by edge pruning; DRAGON [37] uses user-user co-occurrence graph and item-item multi-modal graph to enhance the user-item heterogeneous graph.

Recent studies have acknowledged the benefit of multi-modal alignment in MRSs [19, 35, 39]. Most current alignment methods in MRSs are item-level alignment using contrastive learning. For example, MIRCO [35] and SLMRec [19] align the content of different modalities. BM3 [39] aligns the representations of behavior with multi-modal content. They pull the representations learned from *the different perspectives of the same item* close while pushing the representations of *different items* apart [19, 35, 39].

Denoising Recommender Systems. Recent studies [6, 20, 22] have shown that feedback data inevitably contains noise, such as bias [3, 12, 29] or erroneous user clicks [13, 30]. All existing works are implemented in conventional RSs. We categorize the existing methods into the following categories: (1) *Reweighting methods*. These methods typically reweight each feedback during training. For instance, ADT [20] reweights the feedback according to the loss values. SGDL [6] assign weights for samples based on the similarity of clean samples collected in the initial training stages. Some methods also use user dwell time [28] or item attributes [13]

to denoise implicit feedback. (2) *Ensemble methods*. These methods usually train multiple models and use information from other models to denoise implicit feedback [22]. (3) *Multi-tasking methods*. These methods typically use additional tasks, e.g., multi-view graph contrastive learning in SGL [27].

Remarks. DA-MRS belongs to structure-based MRS. We address the problem of noisy multi-modal content, which is more severe in structure-based MRSs because inaccurate similarities will be propagated and aggregated along false positive links in the item-item graphs. There exists only one study that addresses multi-modal noise by behavior-guided purifier [32]. DA-MRS proposes a different approach, which is based on consistency across modalities. Furthermore, the alignment in DA-MRS is *beyond item-level*. Regarding denoising feedback, DA-MRS overcomes the omission of multi-modal content and presents the first theoretically derived approach to address the feedback noise in MRSs.

3 Methodology

As shown in Figure 2, DA-MRS consists of three modules. Following *structure-based methods* [14, 34, 35], DA-MRS extracts item-item graphs from multi-modal content that reflect static semantic relations among items. Since the multi-modal content is noisy, we propose Denoising Item-item Graph (Section 3.2) to accurately capture item-item semantic relations and build multiple modality-specific item-item semantic graphs. To further utilize dynamic behavior information, we build an item-item behavior graph. Then, the representations of users and items are derived from item-item graphs and the backbone CF model, and the user feedback is utilized as supervised signals to optimize the representations. Due to the noisy nature of user feedback, we propose Denoising User Feedback (Section 3.3) to eliminate the impact of erroneous feedback signals and derive the objective based on the well-known BPR loss. Finally, we align the multi-modal content and user feedback through two aligning methods (Section 3.4) guided by the user preference and the graded item relations.

3.1 Preliminaries

Let \mathcal{U} and \mathcal{I} denote the user and item sets. $|\mathcal{U}|$ and $|\mathcal{I}|$ denote the number of users and items, respectively. The user-item interaction matrix is $\mathbf{O} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$, where $\mathbf{O}_{ui} = 1$ suggests the user u interacts (e.g., clicks, views, etc.) with item i , otherwise $\mathbf{O}_{ui} = 0$. The content of each item i for each modality $m \in \mathcal{M}$ is pre-processed (e.g., by a pre-trained model), and the feature vector is denoted as $\mathbf{e}_i^m \in \mathbb{R}^{d_m}$, where d_m is the embedding dimension, e.g., $\mathcal{M} = \{v, t, a\}$ for visual, textual, and acoustic modalities, respectively. Given \mathbf{O} and \mathbf{e}^m , $m \in \mathcal{M}$, the task of multi-modal recommender systems (MRSs) is to deliver a ranking list of possible recommendations that each user u may prefer, according to the predicted user-item preference score \hat{y}_{ui} .

3.2 Denoising Item-item Graph

Given the user feedback and the multi-modal content, existing *structure-based methods* [14, 34, 35] typically create item-item graphs by connecting each item to its top- k most similar items in each modality. Their construction strategies have two problems. (1) Most studies [14, 34, 38] merge similar items in each modality into one item-item graph. If the similarity is mistakenly amplified based on noisy content, false positive links are introduced into the graph. For example, "carpet" and "painting" in Figure 1(b) can be falsely connected since their noisy visual attributes are similar. (2) They focus on semantic relations extracted from multi-modal content while neglecting behavior relations extracted from user feedback. Thus, they can not fully reveal the collaborative relations among items.

To address these problems, instead of constructing one graph, we construct multiple item-item graphs, namely Item-item Semantic Graph (IIS-Graph) and Item-item Behavior Graph (IIB-Graph). Each IIS-Graph is constructed in one modality to distinguish modality-specific semantic relations, and the construction is based on consistent similarity across modalities to avoid false positive links. The IIB-Graph is constructed from user feedback to represent co-occurrence behaviors. IIS-Graph and IIB-Graph are complementary, i.e., they mitigate the noise and sparsity problems of each other. By using them together, a more comprehensive item-item collaborative relationship can be established.

3.2.1 Item-item Semantic Graph. We initialize a dense matrix \mathbf{S}^m , $m \in \mathcal{M}$, where each element $\mathbf{S}_{i,j}^m$ measures the similarity between the two items i and j in modality m . We employ cosine similarity as the similarity metric due to its parameter independence and lower computational complexity, i.e., $\mathbf{S}_{i,j}^m = ((\mathbf{e}_i^m)^T \mathbf{e}_j^m) / (\|\mathbf{e}_i^m\| \|\mathbf{e}_j^m\|)$.

To prune false positive links, we first discard entries with smaller similarities in the dense matrix \mathbf{S}^m . This step avoids the impact of amplified similarity in certain modalities. For example, in E-commerce platforms, since most retailers use verbose descriptions, the textual similarities tend to be higher than visual similarities. Specifically, let $\bar{\mathbf{S}}^m = (\sum_i \sum_j \mathbf{S}_{i,j}^m) / (|\mathcal{I}|^2)$ represents the average similarity in modality m , if $\mathbf{S}_{i,j}^m < \bar{\mathbf{S}}^m$, we set $\mathbf{S}_{i,j}^m = 0$. Then, we discard entries that exhibit inconsistency across modalities. This step prevents semantic relations from being mistakenly added due to noisy content. For example, if two irrelevant items are assigned identical pictures due to a system error, they may be similar in visual

modality but dissimilar in textual modality. Specifically, entries with small similarities in other modalities are deleted, i.e., $\mathbf{S}_{i,j}^m = 0$ if $\exists m', \mathbf{S}_{i,j}^{m'} = 0$.

Next, we use the k -Nearest Neighbors method to construct the IIS-Graph's adjacency matrix \mathbf{A}^m . For each item $i \in \mathcal{I}$, we retrieve the top K items with the highest similarity and generate a list of elements called top- k ($\mathbf{S}_{i,:}^m$). To enhance computational efficiency, we set the non-zero elements in top- k ($\mathbf{S}_{i,:}^m$) to 1. The adjacency matrix of the IIS-Graph \mathbf{G}^m is defined as

$$\mathbf{A}_{i,j}^m = \begin{cases} 1, & \mathbf{S}_{i,j}^m \in \text{top-}k(\mathbf{S}_{i,:}^m) \ \& \ \mathbf{S}_{i,j}^m > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

3.2.2 Item-item Behavior Graph. We initialize an item-item co-occurrence matrix \mathbf{S}^c , where each element $\mathbf{S}_{i,j}^c$ records the frequency of two items i and j clicked by a same user. The idea is that if two items appear together in the user's clicked lists, they are likely to be semantically relevant. Then, we prune infrequent elements, i.e., $\mathbf{S}_{i,j}^c < \xi_B$, where ξ_B is the pruning threshold. The pruning step avoids the impact of random behaviors, e.g., co-occurrence caused by the user randomly clicking an item. Next, we employ the k -Nearest Neighbors method to process the matrix. For each item $i \in \mathcal{I}$, we retrieve the top K items with the highest similarity and generate top- k ($\mathbf{S}_{i,:}^c$). The adjacency matrix of Item-item Behavior Graph \mathbf{G}^c is defined as

$$\mathbf{A}_{i,j}^c = \begin{cases} \mathbf{S}_{i,j}^c, & \mathbf{S}_{i,j}^c \in \text{top-}k(\mathbf{S}_{i,:}^c) \ \& \ i \neq j \ \& \ \mathbf{S}_{i,j}^c \geq \xi_B, \\ 1, & i = j, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

3.2.3 User and Item Representation. We can treat the co-occurrence as another modality and add IIB-Graph to the modality-specific item-item graphs, i.e., $\mathcal{M} = \mathcal{M} \cup \{c\}$. We perform graph convolutions on each item-item graph. Among the various graph convolution methods, we select LightGCN [10] as the convolution kernel for message propagation and aggregation because of its simplicity in computation and widespread adoption. We stack l layers and obtain the last layer's representations as the embeddings for each modality, i.e., \mathbf{h}_i^v , \mathbf{h}_i^t and \mathbf{h}_i^c are item representations learned on IIS-Graph \mathbf{G}^v , IIS-Graph \mathbf{G}^t and IIB-Graph \mathbf{G}^c , respectively.

Following other structure-based methods [14, 34, 35], DA-MRS can plug in various collaborative filtering (CF) methods that model user-item interactions. We feed the user ID embeddings, item ID embeddings, and user observed feedback \mathbf{O} to the backbone CF method and obtain the user embeddings \mathbf{u}_u for user u and item embeddings \mathbf{h}_i^{id} for item i . Note that the user representation is obtained solely by the backbone CF method. We use the item embeddings learned from IIS-Graph, IIB-Graph, and the backbone CF method to obtain the item representation:

$$\mathbf{t}_i = \text{Meanpooling}(\mathbf{h}_i^{\text{id}}, \text{Meanpooling}(\mathbf{h}_i^v, \mathbf{h}_i^t, \mathbf{h}_i^c)). \quad (3)$$

3.3 Denoising User Feedback

After obtaining the user embeddings \mathbf{u} and item embeddings \mathbf{t} , conventional RSs usually use Bayesian Personalized Ranking [17] (BPR) loss. Let y denotes the true user behavior; the probability a user u prefers item i over item j (i.e., $y_{ui} > y_{uj}$) is determined by the model parameters Θ .

$$p(y_{ui} > y_{uj} | \Theta) = \sigma(\hat{y}_{ui} - \hat{y}_{uj}), \quad (4)$$

where \widehat{y}_{ui} is the predicted user-item preference score. It is commonly defined as $\widehat{y}_{ui} = \mathbf{u}_u^T \mathbf{t}_i$, where the user and item representations are part of the model parameters $\mathbf{u} \in \Theta$, $\mathbf{t} \in \Theta$, $\sigma(\cdot)$ represents the sigmoid function.

A training set $\mathcal{D} = \{\langle u, i, j \rangle\}$ is constructed from the observations \mathbf{O} [17]. Each triple in the training set $\langle u, i, j \rangle$ contains a positive observation $\mathbf{O}_{ui} = 1$, and a randomly sampled negative observation $\mathbf{O}_{uj} = 0$. The BPR loss assumes a triple $\langle u, i, j \rangle$ implies $y_{ui} > y_{uj}$. The model parameters are optimized via the BPR loss:

$$\begin{aligned} \mathcal{L}_{BPR} &= \ln(p(\Theta|\mathcal{D})) \\ &\propto \ln(p(\mathcal{D}|\Theta) \times p(\Theta)) \\ &= \sum_{\langle u, i, j \rangle \in \mathcal{D}} \ln(p(y_{ui} > y_{uj}|\Theta)) + \lambda_{\Theta} \|\Theta\|^2, \end{aligned} \quad (5)$$

where λ_{Θ} is the regularization coefficient.

However, due to the presence of noisy feedback [3, 13, 29, 30], **the observation triple $\langle u, i, j \rangle$ is not equivalent to true user behavior $y_{ui} > y_{uj}$** . Our basic idea is to treat the observation triple as a random variable, and the probability of the observation triple is conditioned on the true user behavior. Inspired by the BPR loss [17], we assume **the true user behavior is a ranked list, and the pair-wise rank is either $y_{ui} > y_{uj}$ or $y_{ui} < y_{uj}$** ¹.

When the true user behavior $y_{ui} > y_{uj}$, we assume that the observation triple is drawn from a Bernoulli distribution² parameterized by $f(u, i)$, i.e., $p(\langle u, i, j \rangle \in \mathcal{D} | y_{ui} > y_{uj}, \Theta) = \text{Bernoulli}(f(u, i))$. If the observation is correct and reliable, then $f(u, i) = 1$. In other words, the value of $f(u, i)$ measures the reliability of the observation $\langle u, i, j \rangle \in \mathcal{D}$. Intuitively, since users respond to their preferred multi-modal content, e.g., a fashion-goer prefers trending elements in the displayed image, we can use the estimated preference score on multi-modal content to define $f(u, i)$. That is, the stronger a user is attracted to the item's multi-modal content, the more reliable the observation triple is. Furthermore, the more consistent the user is attracted across different modalities, the more possible the triple can be observed. Specifically, we estimate the user-item preference score on each modality and calculate the mean and variance across different modalities.

$$\begin{aligned} f(u, i) &= (\mu_{ui})^{\alpha} \times (e^{-s_{ui}^2})^{\beta}, \\ \mu_{ui} &= \frac{\sum_{m \in \mathcal{M}} \sigma(\widehat{y}_{ui}^m)}{|\mathcal{M}|}, \quad s_{ui}^2 = \frac{\sum_{m \in \mathcal{M}} (\mu_{ui} - \sigma(\widehat{y}_{ui}^m))^2}{|\mathcal{M}|}, \\ \widehat{y}_{ui}^m &= (\mathbf{u}_u)^T \mathbf{h}_i^m, \end{aligned} \quad (6)$$

where μ_{ui} denotes the average estimated preference score of user u on item i across all modalities. We employ the sigmoid function $\sigma(\cdot)$ to ensure that $\mu \in (0, 1)$. s_{ui}^2 denotes the variance of the estimated preference scores of user u for item i across all modalities. Since $s_{ui}^2 > 0$, $e^{-s_{ui}^2} \in (0, 1)$. $\alpha, \beta > 0$ are two hyper-parameters. Therefore, the function $f(u, i) \in (0, 1)$ ensures a Bernoulli probability. In Equation 6, the more the user prefers the multi-modal content (i.e., larger μ_{ui}) and the more consistent the user preference is across modalities (i.e., smaller s_{ui}^2), the more reliable the observation triple is (i.e., larger $f(u, i)$).

If the true user behavior $y_{ui} < y_{uj}$, we define another Bernoulli distribution $p(\langle u, i, j \rangle \in \mathcal{D} | y_{ui} < y_{uj}, \Theta) = \text{Bernoulli}(g(u, i))$ to be

parameterized by $g(u, i)$. The value of $g(u, i)$ quantifies the probability of incidents that the observation contradicts the user's true behavior. Intuitively, one possible cause of such incidents is when a user is influenced by a certain modality of an item and interacts with it. For example, a user is attracted by an item's high-quality product image and clicks it, although afterward, he finds the item is not what he seeks. Another possible cause is when there are no comparable products on the market. For example, a user does not like the item, but since there is no substitute, he clicks it and attempts to evaluate it with an open mind. To model such phenomena, we can define $g(u, i)$ based on the maximal estimated preference score on the target item i in any single modality and the estimated preference score on competitors.

$$\begin{aligned} g(u, i) &= \begin{cases} \sigma(\max_m(\widehat{y}_{ui}^m) - \overline{\widehat{y}_n})^{\gamma}, & \overline{\widehat{y}_n} > \mu_{ui}, \\ 0, & \text{otherwise,} \end{cases} \\ \overline{\widehat{y}_n} &= \frac{\sum_{\langle u, j \rangle \in \mathcal{B}} \sigma(\widehat{y}_{uj})}{|\mathcal{B}|}, \end{aligned} \quad (7)$$

where $\max_m(\widehat{y}_{ui}^m)$ represents the maximal estimated preference score of user u on item i in any modality. $\overline{\widehat{y}_n}$ represents the average user preference score of negative samples within the mini-batch \mathcal{B} . $\gamma > 0$ is a hyper-parameter. We employ the sigmoid function $\sigma(\cdot)$ to ensure $g(u, i) \in (0, 1)$. Thus, when the user does not like the item (i.e., $\overline{\widehat{y}_n} > \mu_{ui}$), the more the user is attracted by a certain modality (i.e., larger $\max_m(\widehat{y}_{ui}^m)$) and the less the user prefers other products (i.e., smaller $\overline{\widehat{y}_n}$), the more possible the observation triple appears (i.e., larger $g(u, i)$).

Based on the above reasoning, we derive the objective for denoised user feedback as³,

$$\begin{aligned} \mathcal{L}_{D-BPR} &= \sum_{\langle u, i, j \rangle \in \mathcal{D}} \ln \left(f(u, i) \sigma(\mathbf{u}_u^T (\mathbf{t}_i - \mathbf{t}_j)) + g(u, i) (1 - \sigma(\mathbf{u}_u^T (\mathbf{t}_i - \mathbf{t}_j))) \right) \\ &\quad + \lambda_{\Theta} \|\Theta\|^2. \end{aligned} \quad (8)$$

3.4 Aligning Multi-modal Content and User Feedback

Alignment between multi-modal content and user feedback is under-explored in current MRSs. The recommendation performance is damaged because the embeddings learned from multi-modal content and user feedback usually reside in different regions of the feature space. To better integrate multi-modal content and user feedback, we align them. The recommender system consists of items and users. Naturally, our alignment is split into two parts.

3.4.1 Alignment guided by user preference. The item-level alignment can be seen as instance-level alignment on parallel corpora in traditional multi-modal systems. We believe that only instance-level alignment is insufficient in MRSs because the goal of recommender systems is essentially to predict user preference instead of understanding multi-modal content. Our motivation is to use the user preference to orient multi-modal content. For example, if a user prefers a pen over a pencil, then the estimated preference from multi-modal content for the pen should be larger than the estimated preference for the pencil.

Specifically, we extract distinct users from the mini-batch \mathcal{B} and form \mathcal{B}_u . For each user $u \in \mathcal{B}_u$, we compute preference scores on

¹The original BPR paper [17] also only consider $y_{ui} > y_{uj}$ or $y_{ui} < y_{uj}$.

²We consider implicit feedback so the observation matrix is binary.

³The detailed derivation of this loss objective is provided in Appendix A.1

multi-modal content with respect to all $i \in \mathcal{I}$ by $\hat{y}_{ui}^{\text{mm}} = (\mathbf{u}_u)^T \mathbf{h}_i^{\text{mm}}$, where $\mathbf{h}^{\text{mm}} = \text{MeanPooling}(\mathbf{h}^{\text{v}}, \mathbf{h}^{\text{t}}, \mathbf{h}^{\text{c}})$. Then, we compute the preference distribution of the user u over items based on multi-modal content,

$$P_u^{\text{mm}} = \text{softmax}([\hat{y}_{u1}^{\text{mm}}, \dots, \hat{y}_{u|\mathcal{I}|}^{\text{mm}}]), \quad (9)$$

where $|\mathcal{I}|$ represents the number of items, $\text{softmax}()$ represents the softmax function.

Similarly, we can compute the preference distribution P_u^{id} of u over the itemset based on user feedback by $\hat{y}_{ui}^{\text{id}} = (\mathbf{u}_u)^T \mathbf{h}_i^{\text{id}}$. \mathcal{L}_{AU} aligns the two preference distributions,

$$\mathcal{L}_{AU} = \sum_{u \in \mathcal{B}_u} KL[P_u^{\text{mm}} || P_u^{\text{id}}] + KL[P_u^{\text{id}} || P_u^{\text{mm}}], \quad (10)$$

where $KL[]$ represents the KL divergence.

3.4.2 Alignment guided by item graded relations. Contrastive learning is an efficient alignment method that aligns the positive samples and makes the negative samples more distinguishable. In defining positive and negative samples, current MRSs [35, 39] simplify the multi-modal relation as a binary relation, which is sub-optimal. (1) Considering similar items within a single modality as *positive samples* can result in false positives, impeding item representation learning. (2) Treating dissimilar items within a single modality as *negative samples* may overlook some potentially useful samples.

We believe the relations extracted from multi-modal content are graded in nature, i.e., items can be similar in multiple modalities, similar in a single modality, and dissimilar. Exploiting the graded relation is beneficial for finer-grained alignment. For example, when a user wants to purchase a jacket with a similar style to a previously bought shirt, similar shirts (visually similar to the jacket but textually dissimilar) and similar jackets (visually and textually similar) can improve our understanding of the preferred jackets.

To represent the similarity grades, we construct two types of positive samples for each modality, i.e., multi-modal similar items and single-modal similar items. *Note that the multi-modal similar items and single-modal similar items differ in each modality m.* For each modality m , we first compute a similarity matrix, $\mathbf{T}_{i,j}^m = ((\mathbf{h}_i^m)^T \mathbf{h}_j^m) / (\|\mathbf{h}_i^m\| \|\mathbf{h}_j^m\|)$. Then, for each item $i \in \mathcal{I}$, we use the softmax function $\text{softmax}()$ to normalize the similarity scores, i.e., $\tilde{\mathbf{T}}_{i,:}^m = \text{softmax}(\mathbf{T}_{i,:}^m)$. Next, we modify the modal-aware multi-modal similarity by adding the aggregated multi-modal similarity, $\mathbf{R}_{i,:}^m = \tilde{\mathbf{T}}_{i,:}^m + \sum_m \tilde{\mathbf{T}}_{i,:}^m$, i.e., incorporating the multi-modal similarity while highlighting the current modality. Thus, the *multi-modal similar items* are defined as the top k similar items with largest $\mathbf{R}_{i,:}^m$, denoted as \mathcal{R}_i^m . To obtain *single-modal similar items*, we first remove the items in \mathcal{R}_i^m , i.e., $\tilde{\mathbf{T}}_{i,j}^m = 0, \forall j \in \mathcal{R}_i^m$. Then we retrieve top k similar items with largest $\tilde{\mathbf{T}}_{i,j}^m$, and build the single-modal similar itemset \mathcal{T}_i^m . Consequently, we construct the *dissimilar itemset* \mathcal{N}_i^m . For each $j \in \mathcal{T}_i^m$, we set $\tilde{\mathbf{T}}_{i,j}^m = 0$. Then, the dissimilar itemset contains item in mini-batch \mathcal{B} , i.e., $\mathcal{N}_i^m = \{j | j \in \mathcal{B} \& \mathbf{T}_{i,j}^m > 0\}$.

We believe multi-modal similar items should be closer in the representation space than single-modal similar items, and single-modal similar items should be closer than dissimilar items. Accordingly, we propose the contrastive learning loss,

$$\mathcal{L}_{AI-MM} = \sum_m -\log \frac{\sum_{j \in \mathcal{R}_i^m} \varphi(\mathbf{h}_i^m, \mathbf{h}_j^m)}{\sum_{j \in \mathcal{R}_i^m} \varphi(\mathbf{h}_i^m, \mathbf{h}_j^m) + \sum_{k \in \mathcal{T}_i^m} \varphi(\mathbf{h}_i^m, \mathbf{h}_k^m) + \sum_{l \in \mathcal{N}_i^m} \varphi(\mathbf{h}_i^m, \mathbf{h}_l^m)},$$

Table 1: Statistics of the datasets. $|\mathcal{D}|$, $|\mathcal{U}|$, $|\mathcal{I}|$ represents the number of observations, users, and items. \bar{S}^{v} , \bar{S}^{t} and \bar{S}^{a} represents the average visual, textual and acoustic similarity.

Datasets	$ \mathcal{D} $	$ \mathcal{U} $	$ \mathcal{I} $	Sparsity	\bar{S}^{v}	\bar{S}^{t}	\bar{S}^{a}
Baby	160,792	19,445	7,050	0.9988	0.2240	0.2627	-
Sports	296,337	35,598	18,357	0.9995	0.2085	0.2184	-
Clothing	278,677	39,387	23,033	0.9997	0.2239	0.3880	-
TikTok	68,722	9,308	6,710	0.9989	0.8556	0.7113	0.1245

$$\mathcal{L}_{AI-S} = \sum_m -\log \frac{\sum_{k \in \mathcal{T}_i^m} \varphi(\mathbf{h}_i^m, \mathbf{h}_k^m)}{\sum_{k \in \mathcal{T}_i^m} \varphi(\mathbf{h}_i^m, \mathbf{h}_k^m) + \sum_{l \in \mathcal{N}_i^m} \varphi(\mathbf{h}_i^m, \mathbf{h}_l^m)}, \quad (11)$$

where $\varphi(\mathbf{h}_i^m, \mathbf{h}_j^m) = \exp(\text{sim}(\mathbf{h}_i^m, \mathbf{h}_j^m)/\tau)$, τ is the temperature, $\text{sim}()$ is the cosine similarity.

The final loss consists of the denoised BPR loss, the aligning user preference loss, and the aligning graded item relations loss.

$$\mathcal{L} = \mathcal{L}_{D-BPR} + \lambda_1 \mathcal{L}_{AU} + \lambda_2 (\mathcal{L}_{AI-MM} + \mathcal{L}_{AI-S}), \quad (12)$$

where λ_1, λ_2 are two hyper-parameters.

4 Experiment

In this section, we study the following research questions:

- (1) **RQ1:** How does **DA-MRS** perform, compared with various conventional recommender systems (RSs), multi-modal recommender systems (MRSs), and denoising RSs? (Section 4.2)
- (2) **RQ2:** How does **DA-MRS** perform in noisy multi-modal content and noisy feedback scenarios? (Section 4.3)
- (3) **RQ3:** How does each component in **DA-MRS** perform? (Section 4.4)

4.1 Experiments Settings

Datasets. We conduct experiments on three **commonly used** Amazon review dataset ⁴ following previous works [14, 24, 32, 34, 35, 38, 39]: Baby, Sports, and Clothing. To test the effectiveness of **DA-MRS** with various modalities and its usability in various situations, we also conduct experiments on the TikTok dataset following [24]. We process the datasets following previous works [24, 32, 34, 35, 38]. Additional dataset processing details are in Appendix A.2.1. The statistics of datasets are summarized in Table 1.

Evaluation metrics. We use three widely-used evaluation metrics [14, 32, 34, 35, 38]: *Recall@K*, *Precision@K*, and *NDCG@K*, simplified as R , P , and N . Higher values of $R@K$, $P@K$, and $N@K$ indicate more accurate recommendation results.

Implementation. We make our code available online to ease reproducibility ⁵. Other implementation details are in the **Appendix A.2.2**.

4.2 Performance Comparison

4.2.1 Comparison with RSs and MRSs. First, we compare **DA-MRS** with conventional recommender systems and multi-modal recommender systems. The competitors are (1) Conventional RSs, including conventional matrix factorization method (MF-BPR [17]), graph neural networks methods (NGCF [21] and LightGCN [10]), and graph contrastive learning methods (SGL [27] and SimGCL [31]). (2) MRSs include *feature-based methods* (VBPR [9], MMGCN [26],

⁴<http://jmcauley.ucsd.edu/data/amazon/links.html>

⁵<https://github.com/XMUDM/DA-MRS>

Table 2: Performance comparison with conventional RSs, Multi-modal RSs, and Denoising RSs. The best performance is highlighted in bold, and the second best is highlighted by underlines. *vs. Vanilla* represents the relative improvements over the vanilla. *vs. Best* represents the relative improvements over the best baseline in percentage. We report the results of each Denoising RS on its best backbone. * denotes results are copied from its original paper or from MMRc framework.

Datasets	Baby			Sports			Clothing			TikTok		
	R@20	P@20	N@20	R@20	P@20	N@20	R@20	P@20	N@20	R@20	P@20	N@20
MF	0.0570	0.0033	0.0251	0.0681	0.0039	0.0319	0.0318	0.0017	0.0152	0.0558	0.0028	0.0220
NGCF	0.0592	0.0032	0.0233	0.0724	0.0041	0.0315	0.0425	0.0022	0.0184	0.0752	0.0038	0.0319
LightGCN	0.0738	0.0040	0.0323	0.0851	0.0048	0.0383	0.0510	0.0027	0.0229	0.0916	0.0046	0.0406
SGL	0.0808	0.0045	0.0357	0.0939	0.0053	0.0430	0.0595	0.0031	0.0272	0.0972	0.0049	0.0411
SimGCL	0.0778	0.0044	0.0354	0.0896	0.0050	0.0406	0.0525	0.0027	0.0234	0.0972	0.0049	0.0410
VBPR	0.0697	0.0039	0.0295	0.0856*	0.0048*	0.0384*	0.0385	0.0020	0.0165	0.0420	0.0021	0.0164
MMGCN	0.0603	0.0034	0.0255	0.0630	0.0035	0.0260	0.0350	0.0018	0.0147	0.0870	0.0043	0.0283
GRCN	0.0844	0.0047	0.0360	0.0878	0.0050	0.0396	0.0669	0.0035	0.0289	0.0624	0.0031	0.0251
SLMRec	0.0861	0.0047	0.0384	0.1033	0.0057	0.0463	0.0707	0.0037	0.0315	0.1008	0.0050	0.0417
BM3	0.0847	0.0047	0.0369	0.0971	0.0054	0.0437	0.0641	0.0034	0.0294	0.1064	0.0053	0.0454
MMSSL	0.0918	0.0051	0.0409	0.1010	0.0057	0.0455	0.0752	0.0039	0.0340	0.0921*	0.0046*	0.0392*
LATTICE	0.0845	0.0047	0.0366	0.0941	0.0052	0.0414	0.0710*	0.0036*	0.0316*	0.0939	0.0047	0.0433
MICRO	0.0865	0.0045	0.0389	0.0988*	0.0052*	0.0457*	0.0782*	0.0040*	0.0351*	0.0936	0.0047	0.0432
T-CE	0.0730	0.0041	0.032	0.0582	0.0033	0.0261	0.0499	0.0028	0.0224	0.0699	0.0035	0.0286
R-CE	0.0729	0.0041	0.032	0.0697	0.0039	0.0307	0.0414	0.0022	0.0185	0.0706	0.0035	0.0262
DeCA	0.0613	0.0035	0.0264	0.0488	0.0028	0.0214	0.0359	0.0019	0.0151	0.0627	0.0031	0.0245
DA-MRS +MF	0.0881	0.0049	0.0385	0.0998	0.0056	0.0430	0.0913	0.0047	0.0409	0.0643	0.0032	0.0213
<i>vs. Vanilla</i>	54.56%	48.48%	53.39%	46.55%	43.59%	34.80%	187.11%	176.47%	169.08%	15.23%	14.29%	3.29%
DA-MRS +VBPR	0.0749	0.0042	0.0324	0.0923	0.0051	0.0415	0.0746	0.0039	0.0347	0.0525	0.0026	0.0219
<i>vs. Vanilla</i>	7.46%	7.69%	9.83%	7.70%	6.25%	8.07%	93.77%	95.00%	110.30%	25.00%	23.81%	33.54%
DA-MRS +LightGCN	0.0994	0.0055	0.0435	0.1125	0.0063	0.0498	0.0963	0.0050	0.0433	0.1100	0.0055	0.0493
<i>vs. Vanilla</i>	34.69%	37.50%	34.67%	32.20%	31.25%	30.03%	88.82%	85.19%	89.08%	20.09%	19.57%	21.43%
<i>vs. Best</i>	8.28%	7.84%	6.36%	7.81%	10.53%	7.56%	23.14%	25.00%	23.36%	3.38%	3.77%	8.59%

GRCN [25], SLMRec [19], BM3 [39], and MMSSL [24]) and *structure-based methods* (LATTICE [34] and MICRO [35]). We use either their original implementations or the implementations in MMRc⁶ with default parameters.

We have three observations from Table 2: (1) **DA-MRS** significantly outperforms both conventional RSs and MRSs. Specifically, **DA-MRS** improves over the strongest baselines averagely by 7.49%, 8.63%, 23.83%, and 5.25% on Baby, Sports, Clothing, and TikTok, respectively. This indicates that **DA-MRS** can *effectively capture user preferences and get more accurate recommendation results*. (2) The performance of MRSs is generally better than conventional RSs. This indicates that *multi-modal content can supplement user feedback and reflect user interests*. At the same time, we observed that simply fusing the multi-modal content (i.e., VBPR, MMGCN) is not as effective as conventional RSs, indicating that *how to utilize multi-modal information to improve recommendations is a challenging task*. (3) **DA-MRS** can be applied to various modalities. The results of **DA-MRS** on the TikTok dataset are consistent with those obtained on the Amazon dataset. This validates that **DA-MRS** can generalize well to handle visual, textual, and acoustic content.

4.2.2 Comparison with Denoising RSs. Next, we compare **DA-MRS** with various denoising RSs. Since denoising frameworks can be applied to different backbones, we employ three representative backbone CF models: conventional CF method MF [17], graph CF method LightGCN [10], and MRS VBPR [9]. We choose these three backbones because they are widely applicable in many MRSs [14, 32,

34, 35, 38, 39] and denoising frameworks [20, 22, 27, 31]. The competitors of denoising RSs include reweighting methods (T-CE [20] and R-CE [20]), ensemble method (DeCA [22]), and multi-tasking method (SGL [27]). We do not compare with SGDL [6] since we encounter a "CUDA out of memory" error when implementing SGDL on the small Baby dataset on an NVIDIA GeForce RTX 3090. We conduct experiments using publicly available original code from research papers. We carefully tune their hyper-parameters and the results of each denoising RS on its best backbone. The results of denoising RSs on all backbones are shown in Table 5 in the Appendix.

From Table 2, we observe that **DA-MRS** consistently improves the performance of different backbone models. Averagely, **DA-MRS** boosts the MF backbone by 70.57%, the VBPR backbone by 35.70%, and the LightGCN backbone by 43.71% on the four datasets. On the contrary, most existing denoising methods (i.e., T-CE, R-CE, and DeCA) do not obtain satisfactory results. Even after carefully tuning their hyper-parameters, their performance drops by at most 48.65% compared with the vanilla backbone model. Possible reasons for the superiority of **DA-MRS** can be attributed to the following factors. (1) **DA-MRS** uses multi-modal content to denoise the feedback data, along with alignment between the multi-modal content and user feedback. Thus, **DA-MRS** is more robust and accurate in correcting the noisy feedback. (2) T-CE, R-CE, and DeCA require explicit user-item ratings and are optimized by the CE loss, which limits their applicability. When only implicit feedback is available, their performance is severely damaged.

The improvement of **DA-MRS** on VBPR is less significant when compared with MF and LightGCN. We believe that VBPR *directly utilizes* noisy multi-modal content as ground knowledge of each

⁶<https://github.com/enoch/MMRc>

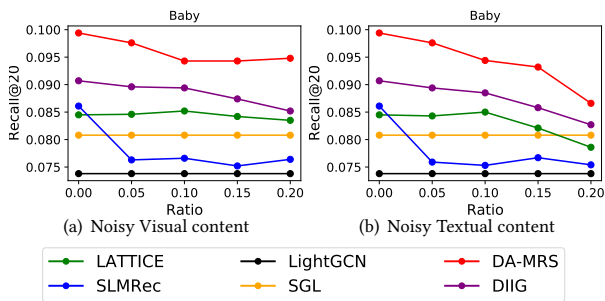


Figure 3: Performance in various noisy multi-modal content scenarios on Baby dataset

item, offsetting some of the improvements made by DA-MRS. However, DA-MRS still significantly improves over VBPR by at least 7.46% on all metrics.

DA-MRS is effective on different datasets. We observe stable improvements of at least 7.46%, 6.25%, 85.19%, and 3.29% compared with backbones on the Baby, Sports, Clothing, and TikTok datasets. This shows the effectiveness of DA-MRS in encompassing diverse dataset sizes, sparsity levels, distributions of multi-modal content similarity, and application scenarios. The improvement on the Clothing dataset is significantly higher than those on the Baby, Sports, and TikTok. From Table 1, we find that the Clothing dataset exhibits the sparsest user feedback, and the average similarity among different modalities is the highest. This finding validates our assumption that leveraging multi-modal content can more effectively address the issue of sparse user feedback.

We analyze the computational complexity of the models, and DA-MRS has significantly lower complexity than MMSSL and the same complexity magnitude as LATTICE. The theoretical analysis and more results on the training time and GPU memory cost of each model are provided in Appendix A.3.

4.3 Performance on Various Noisy Scenarios

4.3.1 Performance on noisy multi-modal content. To testify DA-MRS’s ability to denoise multi-modal content, we deliberately introduce noise to the multi-modal content of each dataset and evaluate the RS’s performance. To construct a dataset with (more) noisy multi-modal content, we first randomly sample items from the training set. For each sampled item i , we randomly sample another item j , $j \neq i$ from the training set and then replace the raw modal features of item i with the features of item j . We randomly replace 5%, 10%, 15%, and 20% of the item modal features in the dataset. The replacement ratio is limited to 20% to prevent excessive noise that could potentially cause all models to fail. We conduct replacements on either visual or textual modality. Each time, we only modify the features of one modality while keeping the feedback data unchanged.

Competitors. We compare DA-MRS with (1) the backbone LightGCN, (2) the best denoising framework SGL, and (3) two well-performing MRSs LATTICE and SLMRec. Since Denoising Item-item Graph is the major component in DA-MRS to deal with noisy multi-modal content, we additionally report the performance of Denoising Item-item Graph (DIIG). We use LightGCN as the backbone and obtain the item embeddings through Equation 3, the

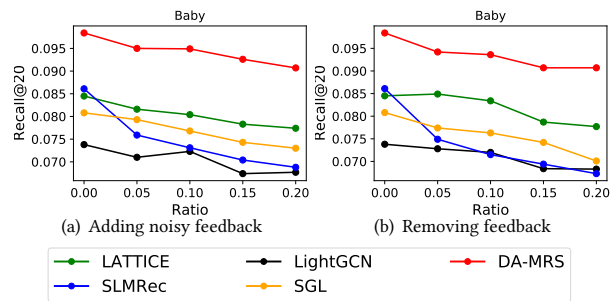


Figure 4: Performance in various noisy user feedback scenarios on Baby dataset

model DIIG is optimized by the vanilla BPR loss in Equation 5. We use $R@20$ as the evaluation metrics.

Figure 3 reports results in Baby. We can observe that (1) *With different noise ratios, DA-MRS stably outperforms all competitors.* DA-MRS achieves significantly better results; its results on 20% noise surpass the best results of other methods on none noise. Denoising Item-item Graph and aligning multi-modal content and user feedback in DA-MRS can alleviate the performance declines caused by noisy multi-modal. (2) Denoising the multi-modal content is crucial. LightGCN and SGL do not utilize multi-modal content, so their performance remains unchanged. When the multi-modal content is too noisy, i.e., replacing 20% item modal features, MRSs (i.e., SLMRec and LATTICE) perform worse than SGL. The noisy multi-modal content will *mislead the item representation modeling and disrupt the performance of recommendations, demonstrating the importance of denoising multi-modal content.* (3) DIIG frequently performs second best in different noise ratios, which shows the effectiveness of this component in denoising multi-modal content. Compared with DIIG, DA-MRS is optimized by D-BPR and has two alignment tasks AU and AI, achieving better results. This demonstrates the effectiveness of denoising user feedback and aligning tasks to improve recommendation performance.

4.3.2 Performance on Noisy feedback. To testify DA-MRS’s ability to denoise user feedback, we contaminate the feedback data in two manners. (1) Randomly generate interactions that do not exist in the dataset and add them to the training set. (2) Randomly remove interactions from the training set. We only modified the training set and left the validation set and test set unchanged. Again, to generate noisy scenarios and avoid excessive noise, we change 5%, 10%, 15%, and 20% of the user feedback in the training set.

Figure 4 reports results in Baby. We can observe that (1) **DA-MRS** consistently achieves the best results, i.e., the highest $R@20$. Even our worst results on 20% noisy feedback surpass the best results of other methods on none noisy feedback, providing strong evidence for *the high performance of DA-MRS on noisy user feedback.* (2) As the ratio of adding noisy feedback increases, the results of all methods generally decline. These scenarios simulate situations where the observed user behavior and user true preference are inconsistent. **DA-MRS** exhibits the smallest decrease rate 7.8% compared with other competitors, demonstrating that *DA-MRS can handle the issue of noisy user feedback and generate stable recommendations.* (3) As the feedback data is removed, the user feedback becomes sparser. Compared with other baselines, **DA-MRS** exhibits the smallest

Table 3: Ablation results of DA-MRS on Baby dataset

Datasets	Baby		
	R@20	P@20	N@20
LightGCN	0.0738	0.0040	0.0323
IIG	0.0897	0.0050	0.0397
DIIG	0.0915	0.0051	0.0400
DIIG + D-BPR	0.0920	0.0051	0.0400
DIIG + AU	0.0935	0.0052	0.0411
DIIG + AI	0.0956	0.0053	0.0419
DIIG + AUI	0.0977	0.0054	0.0427
DA-MRS - $f(u, i)$	0.0983	0.0054	0.0429
DA-MRS - $g(u, i)$	0.0992	0.0055	0.0430
DA-MRS	0.0994	0.0055	0.0435

decrease rate 9.0%. This shows that **DA-MRS** is more effective in solving the sparsity problem in user feedback caused by noisy missing interaction.

4.4 Ablation Study

4.4.1 Impact of each component in DA-MRS. In this section, we aim to evaluate the effectiveness of each component in **DA-MRS**, i.e., DIIG, D-BPR, AI, and AU. We use LightGCN as the backbone and build several variants of **DA-MRS** upon it by adding IIG (the item-item graphs without pruning the false positive links); DIIG (the denoised item-item graphs); DIIG + AU (DIIG with additional alignment guided by user preference); DIIG + AI (DIIG with additional alignment guided by graded item relations); DIIG + AUI (DIIG with both alignments); These variants are optimized by the vanilla BPR loss in Equation 5. We also set a variant DIIG + D-BPR that DIIG is optimized by D-BPR without alignment.

We have following observations from Table 3: (1) Denoising Item-item Graph is effective. IIG and DIIG achieve higher results than LightGCN, indicating that *multi-modal content can supplement user feedback and generate a superior-quality recommendation*. DIIG gets better recommendations than IIG, demonstrating the necessity of denoising the item-item graphs. (2) Aligning Multi-modal Content and User Feedback is crucial. Compared with DIIG, using either AU or AI leads to improvements. Furthermore, when AU and AI are combined, greater improvements are achieved. This shows *the proposed alignment methods AU and AI can be combined to achieve larger improvements*. (3) Denoising User Feedback is necessary. The only difference between **DA-MRS** and DIIG + AUI is that **DA-MRS** employs D-BPR for model optimization instead of vanilla BPR. While DIIG + D-BPR performs slightly better than DIIG alone on the Baby dataset, **DA-MRS** achieves the best recommendations, demonstrating that D-BPR can generate more accurate feedback signals and enhance recommendation accuracy, especially when alignment is applied. This further demonstrates that the components of **DA-MRS** can be combined to achieve the most significant improvements.

4.4.2 Impact of Bernoulli distribution in Denoising User Feedback. To verify the effectiveness of two Bernoulli distribution in Denoising User Feedback, we have two variations: (1) **DA-MRS - $f(u, i)$** : using only $Bernoulli(g(u, i))$ in Equation 8 by setting $f(u, i) = 1$; (2) **DA-MRS - $g(u, i)$** : using only $Bernoulli(f(u, i))$ in Equation 8 by setting $g(u, i) = 0$.

Table 4: Performance of different strategies to select positive and negative samples in Alignment guided by graded Item relations on Baby.

Datasets	Baby		
	R@20	P@20	N@20
DIIG+AI	0.0956	0.0053	0.0419
DIIG+SP	0.0912	0.0050	0.0398
DIIG+MP	0.0940	0.0052	0.0407

We can observe from Table 3 that using one Bernoulli distribution leads to improvements. Specifically, using only $Bernoulli(f(u, i))$ leads to more improvement. This is because feedback consistent with the true user preference is more frequent than feedback contradicting the true user preference. Combining the two Bernoulli distributions yields the best results, suggesting that considering both scenarios achieves the best denoising efficiency.

4.4.3 Impact of Positive and Negative Sample Strategy in Alignment guided by graded Item relations. To investigate the impact of different positive and negative sample strategies in Alignment guided by graded Item relations, we conduct an experiment on the Baby dataset. Specifically, we implement DIIG with three different strategies to select positive and negative samples: (1) AI: It is our method that considers multi-modal similar items and single-modal similar items as graded positive examples. (2) SP: Considering multi-modal similar and single-modal similar items as positive examples, dissimilar items as negative examples. (3) MP: Considering multi-modal similar items as positive examples, single-modal similar and dissimilar items as negative examples.

From Table 4, we can observe that (1) AI achieves the best recommendation performance, indicating that our method *effectively utilizes graded item relations* and assists in the recommendation task. (2) SP performs the worst because it treats multi-modal similarities and single-modal similarities equally, which *introduces some false positive relationships*. (3) MP performs worse than AI because it ignores the single-modal similarities, which *overlooks some potentially useful samples*.

5 Conclusion

DA-MRS is a plug-and-play multi-modal recommendation framework to deal with noisy multi-modal content and noisy user feedback simultaneously for the first time. **DA-MRS** sheds insight into several perspectives. (1) The content noise can be mitigated effectively based on similarity consistency across modalities. (2) Multi-modal content can be utilized in estimating the confidence of noisy user feedback with a probabilistic generative model, which opens opportunities for modifying BPR loss to fit related denoising scenarios. (3) Coarse-grained item-level alignment is sub-optimal in MRSs. The proposed Alignment guided by graded Item relations has the potential to be adapted to other multi-modal tasks.

Acknowledgments

Chen Lin is supported by the Natural Science Foundation of China (No.62372390). Ruobing Xie is supported by the Young Elite Scientists Sponsorship Program by CAST (2023QNRC001).

References

- [1] Desheng Cai, Shengsheng Qian, Quan Fang, Jun Hu, Wenkui Ding, and Changsheng Xu. 2023. Heterogeneous Graph Contrastive Learning Network for Personalized Micro-Video Recommendation. *IEEE Trans. Multim.* 25 (2023), 2761–2773.
- [2] Xumin Chen, Ruobing Xie, Zhijie Qiu, Peng Cui, Ziwei Zhang, Shukai Liu, Shiqiang Yang, Bo Zhang, and Leyu Lin. 2023. Group-based social diffusion in recommendation. *World Wide Web (WWW)* 26, 4 (2023), 1775–1792.
- [3] Yewen Fan, Nian Si, and Kun Zhang. 2023. Calibration Matters: Tackling Maximization Bias in Large-scale Advertising Recommendation Systems. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
- [4] Zeno Gantner, Lucas Drumond, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2012. Personalized Ranking for Non-Uniformly Sampled Items. In *Proceedings of KDD Cup 2011 competition, San Diego, CA, USA, 2011 (JMLR Proceedings, Vol. 18)*. 231–247.
- [5] Chen Gao, Tzu-Heng Lin, Nian Li, Depeng Jin, and Yong Li. 2023. Cross-Platform Item Recommendation for Online Social E-Commerce. *IEEE Trans. Knowl. Data Eng.* 35, 2 (2023), 1351–1364.
- [6] Yunjun Gao, Yuntao Du, Yujia Hu, Lu Chen, Xinjun Zhu, Ziquan Fang, and Baihua Zheng. 2022. Self-guided learning to denoise for robust recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1412–1422.
- [7] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010 (JMLR Proceedings, Vol. 9)*. 249–256.
- [8] Ruining He and Julian J. McAuley. 2016. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*. ACM, 507–517.
- [9] Ruining He and Julian J. McAuley. 2016. VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*. 144–150.
- [10] 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.
- [11] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- [12] Chen Lin, Xinyi Liu, Guipeng Xu, and Hui Li. 2021. Mitigating Sentiment Bias for Recommender Systems. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*. 31–40.
- [13] Hongyu Lu, Min Zhang, and Shaoping Ma. 2018. Between Clicks and Satisfaction: Study on Multi-Phase User Preferences and Satisfaction for Online News Reading. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*. 435–444.
- [14] Zongshen Mu, Yueting Zhuang, Jie Tan, Jun Xiao, and Siliang Tang. 2022. Learning Hybrid Behavior Patterns for Multimedia Recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia (Lisboa, Portugal) (MM '22)*. 376–384.
- [15] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. 8024–8035.
- [16] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*. 3980–3990.
- [17] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18-21, 2009*. 452–461.
- [18] Yu Shang, Chen Gao, Jiansheng Chen, Depeng Jin, Meng Wang, and Yong Li. 2023. Learning Fine-grained User Interests for Micro-video Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*. 433–442.
- [19] Zhulin Tao, Xiaohao Liu, Yewei Xia, Xiang Wang, Lifang Yang, Xianglin Huang, and Tat-Seng Chua. 2022. Self-supervised learning for multimedia recommendation. *IEEE Transactions on Multimedia* (2022).
- [20] Wenjie Wang, Fuli Feng, Xiangnan He, Liqiang Nie, and Tat-Seng Chua. 2021. Denoising implicit feedback for recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*. 373–381.
- [21] 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.
- [22] Yu Wang, Xin Xin, Zaiqiao Meng, Joemon M Jose, Fuli Feng, and Xiangnan He. 2022. Learning robust recommenders through cross-model agreement. In *Proceedings of the ACM Web Conference 2022*. 2015–2025.
- [23] Zongyi Wang, Yanyan Zou, Anyu Dai, Linfang Hou, Nan Qiao, Luobao Zou, Mian Ma, Zhuoye Ding, and Sulong Xu. 2023. An Industrial Framework for Personalized Serendipitous Recommendation in E-commerce. In *Proceedings of the 17th ACM Conference on Recommender Systems, RecSys 2023, Singapore, Singapore, September 18-22, 2023*. 1015–1018.
- [24] Wei Wei, Chao Huang, Lianghao Xia, and Chuxu Zhang. 2023. Multi-Modal Self-Supervised Learning for Recommendation. In *Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*. 790–800.
- [25] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, and Tat-Seng Chua. 2020. Graph-refined convolutional network for multimedia recommendation with implicit feedback. In *Proceedings of the 28th ACM international conference on multimedia*. 3541–3549.
- [26] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In *Proceedings of the 27th ACM international conference on multimedia*. 1437–1445.
- [27] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2021. Self-supervised graph learning for recommendation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 726–735.
- [28] Ruobing Xie, Lin Ma, Shaoliang Zhang, Feng Xia, and Leyu Lin. 2023. Reweighting Clicks with Dwell Time in Recommendation. In *Companion Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*. 341–345.
- [29] Guipeng Xu, Chen Lin, Hui Li, Jinsong Su, Weiyao Ye, and Yewang Chen. 2022. Neutralizing Popularity Bias in Recommendation Models. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*. 2623–2628.
- [30] Tao Yang, Chen Luo, Hanqing Lu, Parth Gupta, Bing Yin, and Qingyao Ai. 2022. Can Clicks Be Both Labels and Features? Unbiased Behavior Feature Collection and Uncertainty-aware Learning to Rank. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*. 6–17.
- [31] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. 2022. Are Graph Augmentations Necessary? Simple Graph Contrastive Learning for Recommendation. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1294–1303.
- [32] Penghang Yu, Zhiyi Tan, Guanming Lu, and Bing-Kun Bao. 2023. Multi-View Graph Convolutional Network for Multimedia Recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23)*. 6576–6585.
- [33] Fanjin Zhang, Jie Tang, Xueyi Liu, Zhenyu Hou, Yuxiao Dong, Jing Zhang, Xiao Liu, Ruobing Xie, Kai Zhuang, Xu Zhang, Leyu Lin, and Philip S. Yu. 2022. Understanding WeChat User Preferences and “Wow” Diffusion. *IEEE Trans. Knowl. Data Eng.* 34, 12 (2022), 6033–6046.
- [34] Jinghao Zhang, Yanqiao Zhu, Qiang Liu, Shu Wu, Shuhui Wang, and Liang Wang. 2021. Mining Latent Structures for Multimedia Recommendation. In *MM '21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*. 3872–3880.
- [35] Jinghao Zhang, Yanqiao Zhu, Qiang Liu, Mengqi Zhang, Shu Wu, and Liang Wang. 2022. Latent Structure Mining With Contrastive Modality Fusion for Multimedia Recommendation. *IEEE Transactions on Knowledge and Data Engineering* (2022), 1–14.
- [36] Hongyu Zhou, Xin Zhou, Zhiwei Zeng, Lingzi Zhang, and Zhiqi Shen. 2023. A Comprehensive Survey on Multimodal Recommender Systems: Taxonomy, Evaluation, and Future Directions. *CoRR abs/2302.04473* (2023). arXiv:2302.04473
- [37] Hongyu Zhou, Xin Zhou, Lingzi Zhang, and Zhiqi Shen. 2023. Enhancing Dyadic Relations with Homogeneous Graphs for Multimodal Recommendation. In *ECAI 2023 - 26th European Conference on Artificial Intelligence, September 30 - October 4, 2023, Kraków, Poland - Including 12th Conference on Prestigious Applications of Intelligent Systems (PAIS 2023) (Frontiers in Artificial Intelligence and Applications, Vol. 372)*. IOS Press, 3123–3130.
- [38] Xin Zhou and Zhiqi Shen. 2023. A Tale of Two Graphs: Freezing and Denoising Graph Structures for Multimodal Recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23)*. 935–943.
- [39] Xin Zhou, Hongyu Zhou, Yong Liu, Zhiwei Zeng, Chunyan Miao, Pengwei Wang, Yuan You, and Feijun Jiang. 2023. Bootstrapped Latent Representations for Multimodal Recommendation. In *Proceedings of the ACM Web Conference 2023 (Austin, TX, USA) (WWW '23)*. 845–854.

A Appendix

A.1 Mathematical Formulations

The detailed formulation of the denoised BPR Loss in Equation 8 is shown as:

$$\begin{aligned}
 \mathcal{L}_{D-BPR} &= \sum_{(u,i,j) \in \mathcal{D}} \ln p(\langle u, i, j \rangle \in \mathcal{D} | \Theta) + \lambda_{\Theta} \|\Theta\|^2 \\
 &= \sum_{(u,i,j) \in \mathcal{D}} \ln \left(p(\langle u, i, j \rangle \in \mathcal{D} | y_{ui} > y_{uj}, \Theta) p(y_{ui} > y_{uj} | \Theta) \right. \\
 &\quad \left. + p(\langle u, i, j \rangle \in \mathcal{D} | y_{ui} < y_{uj}, \Theta) p(y_{ui} < y_{uj} | \Theta) \right) + \lambda_{\Theta} \|\Theta\|^2 \\
 &= \sum_{(u,i,j) \in \mathcal{D}} \ln \left(f(u, i) \sigma(\widehat{y}_{ui} - \widehat{y}_{uj}) + g(u, i) (1 - \sigma(\widehat{y}_{ui} - \widehat{y}_{uj})) \right) \\
 &\quad + \lambda_{\Theta} \|\Theta\|^2 \\
 &= \sum_{(u,i,j) \in \mathcal{D}} \ln \left(f(u, i) \sigma(\mathbf{u}_u^T (\mathbf{t}_i - \mathbf{t}_j)) + g(u, i) (1 - \sigma(\mathbf{u}_u^T (\mathbf{t}_i - \mathbf{t}_j))) \right) \\
 &\quad + \lambda_{\Theta} \|\Theta\|^2. \tag{13}
 \end{aligned}$$

Table 5: Performance comparison with denoising RSs. Vanilla denotes the backbone model. The best performance is highlighted in bold. vs. Vanilla indicates the improvements over the backbone in percentage. Note that SGL can not be used on MF and VBPR.

Datasets		Baby			Sports			Clothing			TikTok		
Backbone	Method	R@20	P@20	N@20	R@20	P@20	N@20	R@20	P@20	N@20	R@20	P@20	N@20
MF	vanilla	0.0570	0.0033	0.0251	0.0681	0.0039	0.0319	0.0318	0.0017	0.0152	0.0558	0.0028	0.022
	R-CE	0.0632	0.0036	0.0266	0.0619	0.0035	0.0281	0.0315	0.0017	0.0147	0.0332	0.0017	0.0125
	T-CE	0.0630	0.0036	0.0264	0.0499	0.0028	0.0224	0.0499	0.0028	0.0224	0.0302	0.0015	0.0097
	DeCA	0.0464	0.0026	0.0206	0.0488	0.0028	0.0214	0.0222	0.0012	0.0064	0.0466	0.0023	0.0175
	DA-MRS	0.0881	0.0049	0.0385	0.0939	0.0056	0.0430	0.0913	0.0047	0.0409	0.0643	0.0032	0.0213
vs. Vanilla	54.56%	48.48%	53.39%	46.55%	43.59%	34.80%	187.11%	176.47%	169.08%	15.23%	14.29%	-3.18%	
LightGCN	Vanilla	0.0738	0.0040	0.0323	0.0851	0.0048	0.0383	0.0510	0.0027	0.0229	0.0916	0.0046	0.0406
	R-CE	0.0729	0.0041	0.0320	0.0585	0.0033	0.0260	0.0367	0.0019	0.0153	0.0706	0.0035	0.0262
	T-CE	0.0730	0.0041	0.032	0.0582	0.0033	0.0261	0.0366	0.0019	0.0155	0.0647	0.0032	0.0262
	SGL	0.0808	0.0045	0.0357	0.0939	0.0053	0.0430	0.0595	0.0031	0.0272	0.0972	0.0049	0.0411
	DeCA	0.0613	0.0035	0.0264	0.0367	0.0081	0.0246	0.0359	0.0019	0.0151	0.063	0.0032	0.0242
DA-MRS	0.0994	0.0055	0.0435	0.1125	0.0063	0.0498	0.0963	0.0050	0.0433	0.11	0.0055	0.0493	
vs. Vanilla	34.69%	37.50%	34.67%	32.20%	31.25%	30.03%	88.82%	85.19%	89.08%	20.09%	19.57%	21.43%	
VBPR	Vanilla	0.0697	0.0039	0.0295	0.0857	0.0048	0.0384	0.0385	0.0020	0.0165	0.042	0.0021	0.0164
	R-CE	0.0697	0.0039	0.0299	0.0697	0.0039	0.0307	0.0414	0.0022	0.0185	0.0699	0.0035	0.0267
	T-CE	0.0496	0.0027	0.0204	0.0277	0.0016	0.0127	0.0279	0.0016	0.0127	0.0699	0.0035	0.0286
	DeCA	0.0352	0.0020	0.0153	0.0326	0.0134	0.0019	0.0229	0.0012	0.0088	0.0627	0.0031	0.0245
	DA-MRS	0.0749	0.0042	0.0324	0.0923	0.0051	0.0415	0.0746	0.0039	0.0347	0.0525	0.0026	0.0219
vs. Vanilla	7.46%	7.69%	9.83%	7.70%	6.25%	7.46%	93.77%	95.00%	110.30%	25.00%	3.81%	33.54%	

A.2 Experiments Settings

A.2.1 Datasets. We conduct experiments on three categories of the Amazon review dataset ⁷. The Amazon review dataset provides both image and text information about the items and varies in the number of items under different categories. We choose the commonly used Baby, Sports, and Clothing datasets. We process the dataset and modal content following previous works [14, 32, 34, 35, 38]. We apply a 5-core setting on both items and users and ensure each item contains visual and textual modality. The open datasets are pre-split into training/validation/test by 8:1:1. We directly use 4,096-dimensional visual features extracted by a pre-trained CNN model [8] and 384-dimensional textual features extracted by sentence-transformers [16]. We calculate the cosine similarity between each item and then calculate the average visual similarity \overline{S}^v and textual similarity \overline{S}^t .

To validate the effectiveness of **DA-MRS** using different modal information and its applicability to different scenarios, we conduct

⁷<http://jmcauley.ucsd.edu/data/amazon/links.html>

experiments on the TikTok dataset following [24]. The TikTok dataset, which contains visual, textual, and acoustic modalities, is collected from a streaming media platform, TikTok ⁸, while the Amazon dataset is collected from an E-commerce site. We believe the TikTok dataset is noisier than Amazon datasets because bloggers are generally less motivated than merchants to produce high-quality media. We calculate the cosine similarity between each item and then calculate the average visual similarity \overline{S}^v , textual similarity \overline{S}^t and acoustic similarity \overline{S}^a .

A.2.2 Implementation details. We implement our method in PyTorch [15]. The embedding dimension d is fixed to 64 for all models to ensure fair comparison. We optimize all models with the Adam [11] optimizer, where the batch size is fixed at 4,096. We use the Xavier initializer [7] to initialize the model parameters. We set $k = 10$ for the k -Nearest Neighbors method. We set the pruning threshold $\xi_B = 2$ for constructing Item-item Behavior Graph. As for α and β in Equation 6, we set the $\alpha = \beta = 1.5$ on Baby and Clothing dataset, $\alpha = \beta = 3$ on Sports dataset. The optimal hyper-parameters are determined via grid search on the validation set: the learning rate is tuned amongst $\{1e-4, 1e-3, 1e-2\}$, the γ in Equation 7 is tuned amongst $\{2.0, 1.0\}$, the λ_1 in Equation 12 is tuned amongst $\{10, 1, 0.1, 0.01\}$, the λ_2 in Equation 12 is tuned amongst $\{1, 0.1, 0.01, 1e-3, 1e-4\}$. For convergence consideration, the early stopping and total epochs are fixed at 25 and 1,000, respectively.

A.2.3 Baselines. We compare with five conventional RSs, including three paradigms: (1) conventional CF method, MF-BPR [17], which uses Bayesian personalized ranking (BPR) loss to optimize matrix factorization; (2) graph CF methods, NGCF [21] and LightGCN [10], which encode collaborative relations into the embedding through various graph neural networks; (3) graph contrastive learning methods, SGL ⁹ [27] and SimGCL ¹⁰ [31], which apply various graph contrastive learning to improve the recommendations.

The MRSs we compared have two paradigms, the *feature-based methods* and *structure-based methods*. The feature-based methods use multi-modal content to establish ground knowledge of individual items, including (1) direct fusion method, VBPR [9], which fuses multi-modal content through conventional CF; (2) graph neural network methods, MMGCN [26], GRCN [25] and SLMRec [19], which fuses multi-modal content through GNNs; (3) other fusing methods, BM3 [39] and MMSSL ¹¹ [24], which fuse modal content through contrastive learning and adversarial learning, respectively. The *structure-based methods* construct the semantic graph through similar k -NN and incorporate the multi-modal semantic graph in various methods, including LATTICE [34] and MICRO [35].

As for denoising recommendation frameworks, we have four competitors: (1) Reweighting methods, T-CE ¹² [20] discards the large-loss samples with a dynamic threshold, while R-CE ¹³ [20] adaptively lowers the weights of large-loss samples. (2) Ensemble Method, DeCA ¹⁴ [22] which minimizes the KL-divergence between

⁸<https://www.tiktok.com/>

⁹<https://github.com/wujcan/SGL-Torch>

¹⁰<https://github.com/Coder-Yu/QRec>

¹¹<https://github.com/HKUDS/MMSSL>

¹²<https://github.com/WenjieWWJ/DenoisingRec>

¹³<https://github.com/WenjieWWJ/DenoisingRec>

¹⁴<https://github.com/wangyu-ustc/DeCA>

Table 6: Computational complexity

Model	Graph Convolution	Multi-modal Feature Mapping	Loss
LATTICE+ LightGCN	$O(2L E d/B)$	$O(I 2dm + I ^3 + K I \log(I))$	$O(2dB)$
MMSSL	$O(M 2L E d/B)$	$O(\sum_{m \in M} I d_m d)$	$O((2 + M U I + 2 M)dB + M U I d_m B)$
DA-MRS + LightGCN	$O(2L E d/B)$	$O((M + 1)2LK I d/B)$	$O((4 + 2 M (K + 1) + 2 I)dB)$
DA-MRS + MF	0	$O((M + 1)2LK I d/B)$	$O((4 + 2 M (K + 1) + 2 I)dB)$

Table 7: Performance and training time on three datasets. Time denotes the training time. The unit of training time is "s/epoch". Mem. denotes the GPU Memory. The unit of GPU Memory is "MB". * denotes results are copied from its original paper. - indicates the model cannot be fitted into an NVIDIA GeForce RTX 3090 with 24 GB memory.

Dataset	Baby					Sports					Clothing				
	R@20	P@20	N@20	Time	Mem.	R@20	P@20	N@20	Time	Mem.	R@20	P@20	N@20	Time	Mem.
LATTICE	0.0845	0.0047	0.0366	7.24	4,291	0.0941	0.0052	0.0414	26.79	18,541	0.0710*	0.0036*	0.0316*	-	-
MMSSL	0.0918	0.0051	0.0409	49.68	6,669	0.1010	0.0057	0.0455	183.59	18,181	0.0752	0.0039	0.0340	198.78	23,885
DA-MRS	0.0994	0.0055	0.0435	1.71	5,767	0.1125	0.0063	0.0498	6.51	12,285	0.0963	0.0050	0.0433	7.11	18,061

multiple models while maximizing the likelihood of data observation. (3) Multi-tasking method, SGL¹⁵ [27] uses graph contrastive learning to denoise the model.

A.3 Complexity Analysis

To investigate the complexity of **DA-MRS** compared with other state-of-the-art MRSs, we have the following analysis.

A.3.1 Theory analysis. The computational complexity of a multi-modal recommendation model can be divided into three major components: the graph convolution module, the multi-modal feature mapping module, and the loss computation module. (1) In graph convolution, the computational complexity of **DA-MRS** depends on the backbone model. Taking LightGCN as the backbone, the complexity is $O(\frac{2L|E|d}{B})$, where L is the number of layers in LightGCN, B is the batch size, d is the dimension of embeddings, and $|E|$ is the number of edges in the graph. If MF is used as the backbone, this part has no computational complexity. (2) In multi-modal feature mapping, **DA-MRS** obtains multi-modal features by performing graph convolutions on $|M| + 1$ item-item graphs. **Each graph is constructed before training, and its structure remains fixed during training.** We use LightGCN as the graph convolution kernel, so the complexity during convolution is $O(\frac{2LK|I|d}{B})$, where K is the number of neighbors for each item, and $|I|$ represents the number of items. Therefore, the computational complexity of this part is $O((|M| + 1)\frac{2LK|I|d}{B})$. (3) Moreover, we compute the complexity of the loss computation module of **DA-MRS**. **DA-MRS** includes L_{D-BPR} (which costs $O((2+2(|M|+1))dB)$), L_{AU} (which costs $O(2|I|dB)$), and L_{AI} (which costs $O((2|M|KdB)$). So the complexity of **DA-MRS** + LightGCN is $O(\frac{2Ld(|E|+(|M|+1)K|I|)}{B} + (4 + 2|M|(K + 1) + 2|I|)dB)$.

We also provide the computational complexity analysis for the competitors. (1) LATTICE: In the graph convolution module, it is consistent with **DA-MRS**. *During training, it updates the item-item graph, while the DA-MRS does not require such updates.* It costs $O(|I|2dm)$ to build the similarity matrix between items, $O(|I|^3)$ to normalize the matrix, and $O(K|I|\log(|I|))$ to retrieve the k most

similar items for each item. It is trained using BPR loss, which costs $O(2dB)$. (2) MMSSL: In the graph convolution module, it performs graph convolutions for each modality, resulting in a complexity of $O(|M| * \frac{2L|E|d}{B})$. In the feature transformation part, it uses multiple layers of MLP, which costs $O(\sum_{m \in M} |I|d_m d)$. The loss part includes BPR loss (which costs $O(2dB)$), generator loss (which costs $O(|M||U||I|dB)$), discriminator loss (which costs $O(|M||U||I|d_m B)$), and contrastive learning loss (which costs $O(2|M|dB)$).

We summarize the computational complexity of **DA-MRS** and other models in Table 6. We can observe that **DA-MRS has significantly lower algorithm complexity** than the state-of-the-art multi-modal recommender system LATTICE and MMSSL.

A.3.2 Experimental analysis. We record the average training time and GPU Memory Cost for each model. All models are trained on NVIDIA GeForce RTX 3090 with a batch size 4,096. We use LightGCN as the backbone. We encounter a "CUDA out of memory" error when implementing LATTICE on the clothing dataset, so we copy the results from the original paper [34].

From the Table 7, we can observe that: (1) The training time of **DA-MRS** is significantly shorter than MMSSL, i.e., average 1/27 training time. Note that, compared with MMSSL, **DA-MRS** achieves improvements of 8.27%, 11.39%, and 28.06% in terms of $R@20$ on the Baby, Sports, and Clothing datasets, respectively. (2) The training time of **DA-MRS** is shorter than LATTICE (average 6/25 training time) since the item-item graphs remain fixed during training in **DA-MRS**. Moreover, it significantly improves the recommendation performance. **DA-MRS** improves over LATTICE regarding $R@20$ by 17.63%, 19.55%, and 34.69% on Baby, Sports, and Clothing, respectively. (3) **DA-MRS** costs lower GPU memory than comparative MRSs when handling larger datasets (Sports and Clothing). The growth rate of **DA-MRS** is lower. For example, from the Baby dataset to the Clothing dataset, GPU memory cost for **DA-MRS** increased by 3.13x, and LATTICE increased by 6.58x. **DA-MRS** exhibits superior scalability in GPU memory usage than the comparative MRSs.

Overall, **DA-MRS** has comparable or lower algorithm complexity than the state-of-the-art multi-modal recommender system, and the performance of **DA-MRS** is significantly better than all competitors.

¹⁵<https://github.com/wujcan/SGL-Torch>