

# Mirror Gradient: Towards Robust Multimodal Recommender Systems via Exploring Flat Local Minima

Anonymous Author(s)\*

## ABSTRACT

Multimodal recommender systems utilize various types of information to model user preferences and item features, helping users discover items aligned with their interests. The integration of multimodal information mitigates the inherent challenges in recommender systems, e.g., the data sparsity problem and cold-start issues. However, it simultaneously magnifies certain risks from multimodal information inputs, such as information adjustment risk and inherent noise risk. These risks pose crucial challenges to the robustness of recommendation models. In this paper, we analyze multimodal recommender systems from the novel perspective of *flat local minima* and propose a concise yet effective gradient strategy called Mirror Gradient (MG). This strategy can implicitly enhance the model’s robustness during the optimization process, mitigating instability risks arising from multimodal information inputs. We also provide strong theoretical evidence and conduct extensive empirical experiments to show the superiority of MG across various multimodal recommendation models and benchmarks. Furthermore, we find that the proposed MG can complement existing robust training methods and be easily extended to diverse advanced recommendation models, making it a promising new and fundamental paradigm for training multimodal recommender systems.

## CCS CONCEPTS

• Information systems → Recommender systems; • Computing methodologies → Knowledge representation and reasoning; Artificial intelligence.

## KEYWORDS

Recommender systems, Multimodal, Flat local minima, Robust

## 1 INTRODUCTION

**Relevance to the Web and to the track.** Recommender systems play a crucial role in helping users navigate the wealth of choices on the web and discover suitable items or online services. In fact, the integration of deep learning techniques into recommender systems has become widespread [4, 11, 41, 47]. These techniques leverage historical user-item interactions to model user preferences, thereby facilitating the personalized recommendation of items. In recent years, with the emergence of rich multimodal content information encompassing texts, images, and videos, multimodal recommender systems [33, 51] alleviate challenges [53] such as data sparsity and cold start. However, incorporating multimodal information into recommender systems also increases some inevitable risks about the input distribution shift.

The first risk is **inherent noise risk** which always appears in the training phase. Some previous works [38, 45] show that the performance of recommender systems encounters substantial challenges when confronted with input containing some noise in multimodal

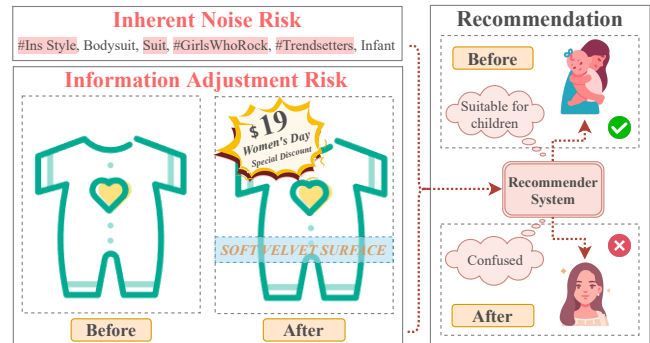


Figure 1: An illustrative example of multimodal risks. Merchants add popular tags (e.g., "ins style") and broad keywords (e.g., "suit") to the text of the bodysuit to increase the likelihood of the item being recommended. At the same time, merchants dynamically change the item’s visual features in real-time due to Women’s Day marketing campaigns and the emphasis on the superiority of the item’s material. These actions make it difficult for the recommender system to accurately determine the target user for the current item, leading to incorrect recommendations for young girls.

information. These noises are intrinsic, such as subpar image quality of items or the presence of a significant number of irrelevant or error information in items’ features. These factors contribute to inherent noise introduced to the model’s input, and the introduction of multimodal data in multimodal recommender systems makes mitigating inherent noise risk more challenging. Another risk is **information adjustment risk**. After the recommender system has been trained based on multimodal data, it is well-known that multimodal data is prone to frequent adjustments. For example, merchants need to keep pace with trends or promotional activities to tailor keywords for items, and the descriptive images of items must be adjusted in line with iterative updates. This implies that in practical scenarios, the multimodal information within recommender systems often undergoes frequent modifications. A straightforward solution to address this risk is to update the model with the latest dataset, but as the volume of data increases, the cost of iterating the model significantly escalates, especially in multimodal recommender systems. In summary, these two risks pose a significant challenge [8, 27, 38, 45] to multimodal recommender systems. Fig. 1 is an **e-commerce case** to illustrate the effect of the risks. On one hand, during the training phase of recommender systems, inherent noise exists in multimodal features such as the text, which is intrinsic and has a potentially negative effect in the inference phase. On the other hand, during the inference phase of recommender systems, the multimodal feature as shown in Fig. 1 of the bodysuit is edited due to promotional activities and the emphasis on the

item’s material. These two risks confuse the recommender system resulting in incorrect recommendations, which will be explicitly and quantitatively assessed in Section 6.

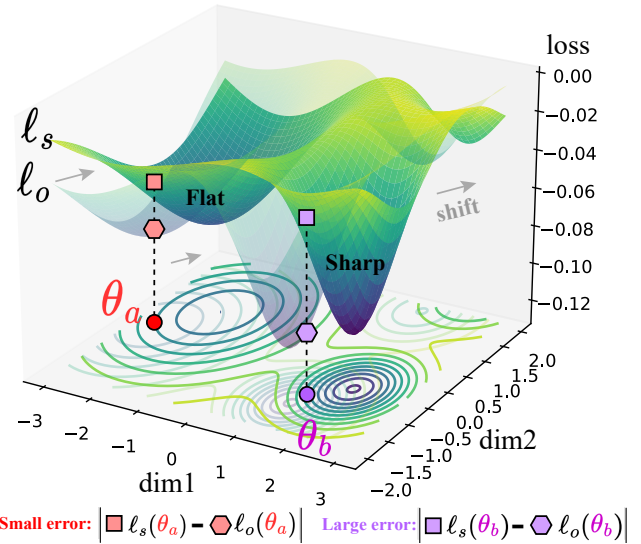
To mitigate the aforementioned risks of information adjustment and inherent noise, the necessity for building a more robust multimodal recommender system becomes apparent. For enhancing the robustness, prior efforts [8, 27] have chiefly employed adversarial training techniques. These methods improve the robustness of multimodal recommender systems by explicitly countering noise at input during the training phase. Different from them, in this paper, we first rethink the robustness of multimodal recommender systems from the loss landscape perspective by considering the *flat local minima* of multimodal recommendation models. Then, we propose a concise yet effective training strategy named Mirror Gradient (MG) for implicitly improving system robustness.

Specifically, we can first present an intuitive insight into why a recommender system should consider the flat local minima, which are located in large weight space regions with very similar low loss values [21]. In Fig. 2,  $\ell_o$  represents the original loss landscape, which is associated with the model’s parameters, architecture, data distribution, etc. When the system’s inputs face shifts in data distribution due to risks like information adjustment or inherent noise,  $\ell_o$  (transparent surface) also shifts to  $\ell_s$  (opaque surface). If the model’s parameters are optimized to a sharp local minimum  $\theta_b$ , the error caused by this shift  $|\ell_s(\theta_b) - \ell_o(\theta_b)|$  may be significantly larger than  $|\ell_s(\theta_a) - \ell_o(\theta_a)|$  of flat local minima  $\theta_a$ . This indicates that the system is not robust while the parameters are in sharp local minima. Therefore, we should strive to guide the learnable parameters of a multimodal recommender system towards flat local minima during training to enhance the model’s robustness against potential risks of input distribution shifts.

To this end, we propose a concise gradient strategy MG that inverts the gradient signs appropriately during training to make the multimodal recommendation models approach flat local minima easier compared to models with normal training. Additionally, we conduct extensive experiments and analysis to validate the effectiveness of MG across various datasets and systems empirically. To elaborate our MG strategy, we first model it formally and then analyze how it improves the model’s robustness by driving the parameters towards flat local minima implicitly from a theoretical perspective. The visualization method from Li et al. [25] supports our theoretical analysis and shows that MG indeed can help the model achieve flatter minima. Besides, we also empirically verify that MG is versatile, as it is compatible with most optimizers and other adversarial training-based robust recommendation methods.

In summary, our contributions are threefold:

- We analyze the robustness of multimodal recommender systems from the perspective of flat local minima.
- From the perspective of flat local minima, we propose Mirror Gradient (MG), a concise yet effective gradient strategy that guides recommender system models toward flat local minima, enhancing model robustness. We also provide theoretical evidence for its effectiveness.
- Extensive experiments demonstrate the efficacy and versatility of MG. We also discuss the limitations of MG.



**Figure 2: Illustration of flat local minima.** When the distribution of the inputs shifts, for example, facing the risks of inherent noise and information adjustment, the loss landscape  $\ell_o$  of the recommender system also shifts (to  $\ell_s$ ) accordingly. The parameters  $\theta_a$  located in flat local minima are more robust compared to  $\theta_b$  in sharp local minima.

## 2 RELATED WORKS

**Multimodal Recommender Systems.** Traditional recommender systems [16, 48, 54] model the interaction between users and items, relying on extensive user-item interaction data to ensure accurate recommendations. In the presence of diverse multimodal information, multimodal recommender systems [1, 28, 31] leverage supplementary multimodal information to complement historical user-item interactions, mitigating challenges like data sparsity [53] and cold start [24, 34] in the recommendation. Early researchers often employed collaborative filter [41, 47] or matrix factorization [15] for multimodal recommendation modeling. Recently, many works [42, 52] employ graph neural networks for multimodal recommender systems, with self-supervised learning [40, 56] also gaining traction in this domain.

**Robust Recommender Systems.** Recent studies [5, 45] have shed light on the vulnerability of recommender systems, highlighting how disturbances introduced by noise can significantly undermine the accuracy of recommendations. In pursuit of bolstering the robustness of recommender systems, a multitude of efforts [3, 8, 27, 38] have focused on adversarial training. This approach, which operates under the premise that each instance may serve as a potential target for attacks [5], introduces controlled perturbations to either the input data or model parameters to enhance robustness. However, most existing studies have overlooked the potential risks arising from information adjustment in multimodal recommender systems.

**Flat Local Minima.** Flat local minima have been consistently linked to the favorable generalization capabilities of deep neural networks [14, 17–19, 35]. In the wake of this insight, numerous researchers have surfaced, exemplified by works such as Du

et al. [7], Zhao et al. [49], Zhuang et al. [57], which strive to enhance model performance through exploring flat local minima. Specifically, Foret et al. [10] introduces a novel approach named Sharpness-Aware Optimization (SAM), wherein the optimization process hinges on addressing a mini-max problem to achieve an optimal sharpness value. Kwon et al. [23], on the other hand, proposes a scale-invariant variant of SAM, named ASAM, bolstered by an adaptive radius mechanism aimed at augmenting training stability. Moreover, Mi et al. [30] innovatively delve into the realm of sparse perturbation with their SSAM (Sparse Sharpness-Aware Minima) approach, strategically focusing on perturbations that encapsulate the most critical yet sparse dimensions of the problem space.

### 3 PRELIMINARY

**Multimodal Recommender Systems.** Considering a set of users  $\mathcal{U} = \{u_1, u_2, \dots, u_{|\mathcal{U}|}\}$ , and a set of items  $\mathcal{I} = \{i_1, i_2, \dots, i_{|\mathcal{I}|}\}$ , each user  $u \in \mathcal{U}$  is associated with an item set  $\mathcal{I}_u \subseteq \mathcal{I}$  about which  $u$  has expressed explicit positive feedback. Besides, each item  $i \in \mathcal{I}$  has multimodal information denoted by visual features  $v_i \in \mathcal{V}$  and textual features  $t_i \in \mathcal{T}$  in this paper. Then given a multimodal recommendation model denoted as  $\mathbf{R}(\cdot)$ ,

$$y_{u,i} = \mathbf{R}(u, i, v_i, t_i, \mathcal{I}_u \mid \Theta), \quad (1)$$

where  $\Theta$  represents the model parameters of  $\mathbf{R}(\cdot)$ , and score  $y_{u,i}$  signifies the preference of user  $u$  towards item  $i$ . A higher score suggests that item  $i$  is more suitable to be recommended to user  $u$ . **Loss Function of Recommender Systems.** Most works [15, 50] optimize the model parameters  $\Theta$  of multimodal recommender systems using Bayesian personalized ranking loss [32]. This optimization seeks to ensure that  $y_{u,i}$ , where  $i \in \mathcal{I}_u$ , is greater than  $y_{u,i'}$  when  $i' \notin \mathcal{I}_u$ , thus promoting positive interactions while discouraging negatives ones. Additionally, some recommender systems introduce supplementary losses [40, 56] to enhance their performance. We adopt the unified notation  $L_{\mathbf{R}}(\cdot)$  to represent those losses.

## 4 METHODOLOGY

In this section, we first elaborate on the algorithm of the proposed MG. Then, we introduce the theoretical insight of MG.

### 4.1 Mirror Gradient

MG is a concise and easily implementable approach that enhances the gradient of the model during the optimization process of recommender systems. This enhancement is equivalent to adding a regularization term to improve the system's robustness on input. The proposed MG consists of two phases in each training epoch: Normal Training and Mirror Training.

During Normal Training, the conventional gradient descent is applied to the loss function  $L_{\mathbf{R}}(\cdot)$  with the current learnable parameters  $\Theta_{t-1}$ , as follows:

$$\Theta_t = \Theta_{t-1} - \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta_{t-1}), \quad (2)$$

where  $\eta$  represents the learning rate. As shown in the Algorithm 1, we use an interval  $\beta$  to control the effect of MG on each training epoch. After updating per  $\beta - 1$  iterations using Eq. (2), we employ the Mirror Training strategy to update the parameter  $\Theta_{t-1}$  as

---

#### Algorithm 1 The training algorithm of Mirror Gradient

---

**Input:** The recommendation model  $\mathbf{R}(\cdot)$ ; learning rate  $\eta$ ; the number of iteration  $T$ ; The scaling coefficients  $\alpha_1, \alpha_2 \in \mathbb{R}^+$  and  $\alpha_1 > \alpha_2$ . The interval  $\beta \in \mathbb{N}^+$  of mirror training.

**Output:** Model parameters  $\Theta$ .

```

1: for  $t$  from 1 to  $T$  do
2:   if  $t \bmod \beta == 0$  do                                ▶ Mirror Training
3:      $\Theta'_t = \Theta_{t-1} - \alpha_1 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta_{t-1})$ ;
4:      $\Theta_t = \Theta'_t + \alpha_2 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta'_t)$ ;
5:   else do                                             ▶ Normal Training
6:      $\Theta_t = \Theta_{t-1} - \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta_{t-1})$ ;
7:   end if
8: end for
9: return  $\Theta$ 

```

---

follows:

$$\begin{cases} \Theta'_t = \Theta_{t-1} - \alpha_1 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta_{t-1}), \\ \Theta_t = \Theta'_t + \alpha_2 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta'_t). \end{cases} \quad (3)$$

Here, in order to control the relative size of updates introduced by mirror training, we introduce two positive scaling coefficients,  $\alpha_1$  and  $\alpha_2$ , with  $\alpha_1 > \alpha_2$ .

Although the MG we proposed is highly simple, it possesses a strong theoretical insight and remarkable versatility. This enables consistent performance improvements across a wide array of experimental scenarios. Furthermore, in Section 6, we demonstrate the compatibility of MG with various optimizers and existing robust recommender system techniques. Moreover, it can achieve superior performance compared to some conventional optimization strategies about flat local minima.

### 4.2 Theoretical Insight of MG

In this part, we introduce Lemma 4.1 and Theorem 4.2 which can help us analyze how MG helps the model's parameters tend towards flat local minima from a theoretical perspective, thereby enhancing the input robustness of the recommender system.

**LEMMA 4.1.** [6, 20] Consider a neural network  $f(x)$  with  $L$  layers and learnable parameters  $\theta$ .  $h_i, 1 \leq i \leq L$ , denotes the feature map from  $i$  th layer. For any scalar function  $g$  of  $h_L$ , we have

$$\|\nabla_x g_{\theta}(x)\|_2^2 \cdot \sum_{j=1}^L \mathcal{O}\left(\frac{1 + \|h_j\|_2^2}{\|\nabla_x h_j\|_2^2}\right) \leq \|\nabla_{\theta} g_{\theta}(x)\|_2^2. \quad (4)$$

**THEOREM 4.2.** Mirror Training in Eq. (3) is equal to introducing an implicit regularization term  $\|\nabla_{\Theta} L_{\mathbf{R}}(\Theta)\|_2^2$  to the original optimization objective  $L_{\mathbf{R}}(\Theta)$ .

**PROOF.** In Mirror Training, we have

$$\begin{aligned} \Theta_t &= \Theta'_t + \alpha_2 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta'_t) \\ &= \Theta_{t-1} - \alpha_1 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta_{t-1}) \\ &\quad + \alpha_2 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta_{t-1} - \alpha_1 \eta \nabla_{\Theta} L_{\mathbf{R}}(\Theta_{t-1})). \end{aligned} \quad (5)$$

Next, since  $\eta$  is small, we can use Taylor expansion for estimating  $\nabla_{\Theta} L_R(\Theta_{t-1} - \alpha_1 \eta \nabla_{\Theta} L_R(\Theta_{t-1}))$ , and we have

$$\begin{aligned} \Theta_t &\approx \Theta_{t-1} - \alpha_1 \eta \nabla_{\Theta} L_R(\Theta_{t-1}) + \alpha_2 \eta \nabla_{\Theta} L_R(\Theta_{t-1}) \\ &\quad - \alpha_1 \alpha_2 \eta^2 \nabla_{\Theta}^2 L_R(\Theta_{t-1})^\top \nabla_{\Theta} L_R(\Theta_{t-1}) \\ &= \Theta_{t-1} - (\alpha_1 - \alpha_2) \eta \nabla_{\Theta} L_R(\Theta_{t-1}) \\ &\quad - \frac{1}{2} \cdot \alpha_1 \alpha_2 \eta^2 \nabla_{\Theta} \|\nabla_{\Theta} L_R(\Theta_{t-1})\|_2^2. \end{aligned} \quad (6)$$

Therefore, from Eq. (6), we can find that the equivalent objective function for Mirror Training  $L_M$  is

$$L_M = (\alpha_1 - \alpha_2) \underbrace{L_R(\Theta)}_{\text{main term}} + \alpha_1 \alpha_2 \eta \cdot \underbrace{\|\nabla_{\Theta} L_R(\Theta_{t-1})\|_2^2}_{\text{regularization term}}, \quad (7)$$

where  $\alpha_1 \alpha_2 \eta > 0$  and  $\alpha_1 - \alpha_2 > 0$ .  $\square$

The essence of MG, as revealed in Theorem 4.2, lies in the addition of a regularization term concerning gradient magnitude to the original objective function  $L_R(\Theta)$  implicitly. It's worth noting that the magnitude of gradients near flat local minima is quite small. And since  $\alpha_1 \alpha_2 \eta > 0$ , Eq. (6) requires that the norm of gradient  $\|\nabla_{\Theta} L_R(\Theta)\|_2^2$  should be sufficiently small, i.e., MG will lead the model's parameters towards flatter minima. Furthermore, from Lemma 4.1 and Eq.(7), let the scalar function is  $L_R$  in recommender system, we have

$$\begin{aligned} L_M &\geq \alpha_1 \alpha_2 \eta \cdot \|\nabla_{\Theta} L_R(\Theta_{t-1})\|_2^2 \\ &\geq \alpha_1 \alpha_2 \eta \cdot \sum_{j=1}^L \mathcal{O}\left(\frac{1 + \|h_i\|_2^2}{\|\nabla_x h_i\|_2^2}\right) \cdot \underbrace{\|\nabla_x L_R\|_2^2}_{\text{robustness term}}. \end{aligned} \quad (8)$$

In general,  $(1 + \|h_i\|_2^2)/(\|\nabla_x h_i\|_2^2)$  is bounded and positive. Taking BM3 on the Baby as an example, its value is around 96.16 with the well-trained system. Therefore, from Eq. (8) and Eq. (7), we can infer that MG also aims to minimize the impact of inputs on the loss,  $\|\nabla_x L_R\|_2^2$ , while enhancing the model's robustness against input perturbations.

Furthermore, although Eq. (7) reveals that our proposed MG is equivalent to adding a regularization term  $\|\nabla_{\Theta} L_R(\Theta_{t-1})\|_2^2$  during the optimization process of recommender systems, we do not recommend directly including this term in the loss function. On one hand, computing this term requires the prior calculation of the gradient  $\nabla_{\Theta} L_R(\Theta_{t-1})$ , implying additional computational overhead during inference and not easy to implement. On the other hand, this kind of explicit loss term is not conducive to optimization and generally results in relatively poor performance [10]. In fact, our proposed MG is an implicit optimization of the additional regularization term shown in Eq. (7). It is also straightforward to implement. Hence, MG possesses greater practical applicability and potential.

## 5 EXPERIMENTS

In this section, we present our experimental setup and empirical results.

### 5.1 Experimental Settings

**Datasets.** The dataset statistics have been summarized in Table 1. We primarily employ four multimodal datasets from Amazon [29],

Dataset	# Users	# Items	# Interactions	Sparsity
Pinterest	3,226	4,998	9,844	99.94%
Baby	19,445	7,050	160,792	99.88%
Sports	35,598	18,357	296,337	99.95%
Clothing	39,387	23,033	237,488	99.97%
Electronics	192,403	63,001	1,689,188	99.99%

**Table 1: Statistics of datasets. These datasets comprise textual and visual features in the form of item descriptions and images.**

Model	REC	PREC	MAP	NDCG
VBPR	0.0182	0.0042	0.0098	0.0122
VBPR + MG	0.0203	0.0046	0.0110	0.0136
<b>Improv.</b>	<b>11.54%</b>	<b>9.52%</b>	<b>12.24%</b>	<b>11.48%</b>
MMGCN	0.0140	0.0033	0.0075	0.0094
MMGCN + MG	0.0157	0.0036	0.0084	0.0106
<b>Improv.</b>	<b>12.14%</b>	<b>9.09%</b>	<b>12.00%</b>	<b>12.77%</b>
GRCN	0.0226	0.0051	0.0126	0.0155
GRCN + MG	0.0250	0.0057	0.0139	0.0171
<b>Improv.</b>	<b>10.62%</b>	<b>11.76%</b>	<b>10.32%</b>	<b>10.32%</b>
DualGNN	0.0238	0.0054	0.0132	0.0162
DualGNN + MG	0.0249	0.0056	0.0139	0.0170
<b>Improv.</b>	<b>4.62%</b>	<b>3.70%</b>	<b>5.30%</b>	<b>4.94%</b>
BM3	0.0280	0.0062	0.0157	0.0192
BM3 + MG	0.0285	0.0063	0.0159	0.0195
<b>Improv.</b>	<b>1.79%</b>	<b>1.61%</b>	<b>1.27%</b>	<b>1.56%</b>
FREEDOM	0.0252	0.0056	0.0139	0.0171
FREEDOM + MG	0.0260	0.0058	0.0144	0.0176
<b>Improv.</b>	<b>3.17%</b>	<b>3.57%</b>	<b>3.60%</b>	<b>2.92%</b>
<b>Avg. Improv.</b>	<b>7.31%</b>	<b>6.54%</b>	<b>7.46%</b>	<b>7.33%</b>

**Table 2: Top-5 recommendation performance of baselines with or without MG on Electronics. "Improv." indicates the relative enhancement of MG compared to the baseline. "Avg. Improv." represents the average improvement.**

including Baby, Sports, Clothing, and Electronics. These datasets comprise textual and visual features in the form of item descriptions and images. Our data preprocessing methodology follows the approach outlined in Zhou et al. [51]. Furthermore, we utilize the dataset Pinterest to assess the compatibility of Mirror Gradient with adversarial training methods in alignment with the official implementation of the adversarial training baseline AMR [38].

**Metrics.** For the evaluation of recommendation performance, we pay attention to top-5 accuracy as recommendations in the top positions of rank lists are more important [39], and adopt four widely used metrics [13, 36, 46, 51] including recall (REC), precision (PREC), mean average precision (MAP), and normalized discounted cumulative gain (NDCG). These four evaluation metrics are chosen because they each focus on different crucial aspects. REC measures whether the system captures a user's potential areas of interest, PREC gauges the accuracy of recommendations, MAP focuses on the average accuracy of rankings, and NDCG emphasizes the quality



Model	Baby				Sports				Clothing			
	REC	PREC	MAP	NDCG	REC	PREC	MAP	NDCG	REC	PREC	MAP	NDCG
VBPR	0.0265	0.0059	0.0134	0.0170	0.0353	0.0079	0.0189	0.0235	0.0186	0.0039	0.0103	0.0124
VBPR + MG	0.0273	0.0061	0.0149	0.0184	0.0375	0.0084	0.0203	0.0251	0.0230	0.0048	0.0129	0.0155
Improv.	3.02%	3.39%	11.19%	8.24%	6.23%	6.33%	7.41%	6.81%	23.66%	23.08%	25.24%	25.00%
MMGCN	0.0240	0.0053	0.0130	0.0160	0.0216	0.0049	0.0114	0.0143	0.0130	0.0028	0.0073	0.0088
MMGCN + MG	0.0269	0.0060	0.0139	0.0175	0.0241	0.0054	0.0126	0.0158	0.0153	0.0032	0.0081	0.0100
Improv.	12.08%	13.21%	6.92%	9.38%	11.57%	10.20%	10.53%	10.49%	17.69%	14.29%	10.96%	13.64%
GRCN	0.0336	0.0074	0.0182	0.0225	0.0360	0.0080	0.0196	0.0241	0.0269	0.0056	0.0140	0.0173
GRCN + MG	0.0354	0.0078	0.0186	0.0232	0.0383	0.0086	0.0207	0.0256	0.0276	0.0058	0.0146	0.0179
Improv.	5.36%	5.41%	2.20%	3.11%	6.39%	7.50%	5.61%	6.22%	2.60%	3.57%	4.29%	3.47%
DualGNN	0.0322	0.0071	0.0175	0.0216	0.0374	0.0084	0.0206	0.0253	0.0277	0.0058	0.0153	0.0185
DualGNN + MG	0.0329	0.0073	0.0176	0.0219	0.0387	0.0086	0.0212	0.0261	0.0305	0.0063	0.0165	0.0201
Improv.	2.17%	2.82%	0.57%	1.39%	3.48%	2.38%	2.91%	3.16%	10.11%	8.62%	7.84%	8.65%
SLMRec	0.0343	0.0075	0.0182	0.0226	0.0429	0.0095	0.0233	0.0288	0.0292	0.0061	0.0163	0.0196
SLMRec + MG	0.0381	0.0085	0.0204	0.0253	0.0449	0.0099	0.0242	0.0299	0.0323	0.0067	0.0175	0.0213
Improv.	11.08%	13.33%	12.09%	11.95%	4.66%	4.21%	3.86%	3.82%	10.62%	9.84%	7.36%	8.67%
BM3	0.0327	0.0072	0.0174	0.0216	0.0353	0.0078	0.0194	0.0238	0.0246	0.0051	0.0135	0.0164
BM3 + MG	0.0345	0.0077	0.0183	0.0228	0.0386	0.0086	0.0210	0.0259	0.0259	0.0054	0.0145	0.0174
Improv.	5.50%	6.94%	5.17%	5.56%	9.35%	10.26%	8.25%	8.82%	5.28%	5.88%	7.41%	6.10%
FREEDOM	0.0374	0.0083	0.0194	0.0243	0.0446	0.0098	0.0232	0.0291	0.0388	0.0080	0.0211	0.0257
FREEDOM + MG	0.0397	0.0088	0.0209	0.0261	0.0466	0.0102	0.0242	0.0303	0.0405	0.0084	0.0223	0.0270
Improv.	6.15%	6.02%	7.73%	7.41%	4.48%	4.08%	4.31%	4.12%	4.38%	5.00%	5.69%	5.06%
DRAGON	0.0374	0.0082	0.0202	0.0249	0.0449	0.0098	0.0239	0.0296	0.0401	0.0083	0.0225	0.0270
DRAGON + MG	0.0419	0.0092	0.0219	0.0273	0.0465	0.0102	0.0248	0.0307	0.0437	0.0091	0.0239	0.0290
Improv.	12.03%	12.20%	8.42%	9.64%	3.56%	4.08%	3.77%	3.72%	8.98%	9.64%	6.22%	7.41%
<b>Avg. Improv. †</b>	<b>7.17%</b>	<b>7.91%</b>	<b>6.79%</b>	<b>7.08%</b>	<b>6.22%</b>	<b>6.13%</b>	<b>5.83%</b>	<b>5.90%</b>	<b>10.41%</b>	<b>9.99%</b>	<b>9.38%</b>	<b>9.75%</b>

**Table 3: Top-5 recommendation performance of baselines with or without MG on Baby, Sport, and Clothing. "Improv." indicates the relative enhancement of MG compared to the baseline. "Avg. Improv." represents the average improvement across each dataset. † The Top-5 performance improvement achieved by MG is substantial and significant. See the appendix for further discussion.**

of rankings. These metrics complement each other and collectively provide a comprehensive evaluation, aiding in a holistic understanding of the recommender system’s performance. Additionally, when comparing against the adversarial training method AMR [38], we apply the hits ratio (HR) in alignment with the evaluation approach used in the original AMR paper. The choice of HR serves the purpose of maintaining consistency and comparability with established benchmarks, allowing for a direct comparison of our results with those reported in the paper. All the above-cited metrics range from 0 to 1, the closer to 1 the better.

**Baselines.** We extensively examine MG’s performance across a variety of multimodal recommendation models, encompassing matrix factorization (VBPR [15]), graph neural networks (MMGCN [44], GRCN [43], DualGNN [42], FREEDOM [52], DRAGON [50]), self-supervised learning (SLMRec [40], BM3 [56]), as well as non-multimodal models (LayerGCN [54], SelfCF [55]). We utilize AMR [38] as our foundational adversarial training method, given its widespread popularity in the field. Additionally, we compare our

MG with flat local minima approach SSAM [30], as it represents a leading-edge, versatile solution for addressing flat local minima.

**Implementation Details.** We retain the standard settings for all baselines. Following the settings of some current works in multimodal recommender systems [50, 52, 56], we perform a grid search on hyperparameters  $\alpha_1$  and  $\alpha_2$ , and set  $\beta$  to 3. Unless otherwise specified, Adam [22] serves as the chosen optimizer. The training and evaluation of all models is conducted using the RTX3090 GPU.

## 5.2 Overall Performance

**Observation #1: MG can enhance the performance of diverse multimodal recommender systems consistently.** As shown in Table 3, we conduct an extensive evaluation of MG across eight baseline models using three distinct datasets. Our experimental findings unequivocally illustrate that it is difficult to overlook the enhancements in the performance of multimodal recommender systems with MG across all evaluation metrics. Remarkably, the most substantial improvement is observed in the case of VBPR training on the dataset Clothing, resulting in an impressive performance enhancement of over 20%. To summarize, MG we propose excels at




Figure 3: Visualization of local minima. Training loss landscapes of FREEDOM and BM3 on Baby trained with or without MG.

mitigating inherent noise risks, enabling multimodal recommender systems to capture user preferences within more favorable flat local minima. This capability contributes to the success of our method. **Observation #2: MG exhibits remarkable efficacy on more challenging datasets.** Furthermore, as indicated in Table 2, we extend the evaluation of our proposed MG to the dataset Electronics, which presents a more demanding scenario with a cumulative user and item count surpassing the sum of other datasets presented in Table 3. Remarkably, the increased challenge posed by this dataset does not diminish MG’s performance. On Electronics, the average improvement achieved by MG remains consistent with the figures presented in Table 3.

## 6 ANALYSIS

Our analysis is designed to answer the following research questions (RQs), where RQ1-5 focus on the performance of MG, while RQ6-8 examine its versatility:

- RQ1: How does MG perform in mitigating inherent noise risk?
- RQ2: How does MG perform in mitigating information adjustment risk?
- RQ3: Does MG enable multimodal recommender systems to approach flatter local minima?
- RQ4: Does MG increase model training costs?
- RQ5: Is MG superior to sharpness-aware minimization methods?
- RQ6: Can MG be compatible with various optimizers?
- RQ7: Can MG be compatible with robust recommender systems?
- RQ8: Can MG enhance the performance of non-multimodal recommender systems?

**Further Evaluation for MG’s Robustness (RQ1 & RQ2).** In this part, we directly validate the robustness of the proposed MG by explicitly simulating the two risks mentioned in Section 1 in the inference phase.

(1) Information noise. For models with learnable item embedding layers, e.g., BM3 and FREEDOM, we inject Gaussian noise  $\epsilon \sim \mathcal{N}(0, 10^{-6})$  into their embedding layers. This step aims to simulate the presence of information noise in multimodal recommender systems. The results are presented in Table 4.

(2) Information adjustment. To simulate the adjustment in multimodal information within a recommender system, we insert image captions generated using BLIP-2 [26] into the text of the dataset with a one percent chance. As models with learnable item embedding layers do not rely on the original multimodal information

Model	Org	Noise	Decr. ↓
BM3	0.0327	0.0272	16.92%
BM3 + MG	0.0345	0.0314	<b>8.99%</b>
FREEDOM	0.0374	0.0357	4.60%
FREEDOM + MG	0.0397	0.0390	<b>1.76%</b>

Table 4: REC of models under information noise on Baby. The term "Noise" refers to the injection with noise  $\epsilon \sim \mathcal{N}(0, 10^{-6})$  into the embedding layers of the model. The experimental results of the noise are calculated as the mean value obtained from conducting the noise experiment 10 times repetitively. "Decr." denotes the relative decrease compared to the origin result (Org) after injecting noise.

Model	Org	Adjustment	Decr. ↓
MMGCN	0.0240	0.0223	7.08%
MMGCN + MG	0.0269	0.0260	<b>3.35%</b>
GRCN	0.0336	0.0321	4.46%
GRCN + MG	0.0354	0.0346	<b>2.26%</b>

Table 5: REC of models under information adjustment on Baby. The term "Adjustment" refers to the action of inserting image captions generated using BLIP-2 [26] into the text of items with a 1% chance. This insertion process simulates changes in multimodal information within recommender systems. "Decr." denotes the relative decrease compared to the origin result (Org) after information adjustment.

during the inference phase, we select MMGCN and GRCN, which use deep neural networks to extract multimodal information, as the baselines. The results are outlined in Table 5.

The results of experiments indicate that multimodal recommendation models trained with MG exhibit greater robustness in handling multimodal noise and adjustment.

**Visualization of Local Minima (RQ3).** In order to affirm that MG can help the model approach flat local minima, we visualize the training loss landscapes with and without MG of latest models BM3 and FREEDOM which is more challenge to improve than previous models. Following the methodology outlined by Li et al. [25], we record the trained model parameters as  $p$ . Subsequently, we sample 20 values each for  $m$  and  $n$  at equal intervals from the

Optimizer	Model	REC	PREC	MAP	NDCG	Model	REC	PREC	MAP	NDCG
Adam	GRCN	0.0336	0.0074	0.0182	0.0225	DRAGON	0.0374	0.0082	0.0202	0.0249
	GRCN + MG	0.0354	0.0078	0.0186	0.0232	DRAGON + MG	0.0419	0.0092	0.0219	0.0273
	Improv.	5.36%	5.41%	2.20%	3.11%	Improv.	12.03%	12.20%	8.42%	9.64%
SGD	GRCN	0.0013	0.0003	0.0005	0.0007	DRAGON	0.0182	0.0040	0.0093	0.0118
	GRCN + MG	0.0014	0.0003	0.0006	0.0008	DRAGON + MG	0.0188	0.0041	0.0097	0.0122
	Improv.	7.69%	0.00%	20.00%	14.29%	Improv.	3.30%	2.50%	4.30%	3.39%
RMSprop	GRCN	0.0338	0.0074	0.0184	0.0227	DRAGON	0.0367	0.0081	0.0198	0.0245
	GRCN + MG	0.0345	0.0076	0.0187	0.0231	DRAGON + MG	0.0391	0.0087	0.0201	0.0253
	Improv.	2.03%	2.63%	1.60%	1.73%	Improv.	6.54%	7.41%	1.52%	3.27%
Adagrad	GRCN	0.0283	0.0063	0.0152	0.0189	DRAGON	0.0393	0.0086	0.0215	0.0264
	GRCN + MG	0.0286	0.0064	0.0154	0.0191	DRAGON + MG	0.0408	0.0090	0.0216	0.0269
	Improv.	1.06%	1.59%	1.32%	1.06%	Improv.	3.82%	4.65%	0.47%	1.89%

Table 6: Top-5 recommendation performance of GRCN and DRAGON with or without MG on Baby. "Improv." indicates the relative enhancement of MG compared to the baselines.




Figure 4: Convergence of MG on the dataset Baby.

range  $[-100, 100]$ , and init noise  $n_1, n_2$  from the standard normal distribution.  $n_1, n_2$  have the same shape as  $p$ . We then update the model's parameters to  $(p + mn_1 + nn_2)$  and calculate the corresponding loss values, simulating the shift of the loss landscape as depicted in Fig. 2. By employing this methodology, we can generate training loss landscapes as illustrated in Fig. 3. The flatter the training loss landscape, the flatter the local minima to which the current model converges. Notably, the landscapes associated with MG demonstrate a flatter topography when compared to those of the baseline approaches and prominently encompass an area characterized by low loss (depicted in blue). These visualization results show that MG can make multimodal recommender system approach flatter minima. Moreover, from the visualization results, the loss landscape associated with FREEDOM appears flatter compared to that of BM3. This indicates that FREEDOM may be more robust to noise, aligning with the observations in Table 4.

**Convergence Speed of MG (RQ4).** In this part, we visualize the training loss of both widely used and leading-stage models to confirm MG's superior convergence. Following the training configuration outlined by Zhou et al. [51], we set the maximum number of epochs to 1000 while implementing an early stopping strategy. Subsequently, we visualize the progression of the loss value for multiple multimodal recommendation models before and after the application of MG, as shown in Fig. 4. Noticeably, MG often leads the models to meet the termination criteria with fewer training iterations and achieve a smaller final training loss value.

Method	REC	PREC	MAP	NDCG
-	0.0374	0.0082	0.0202	0.0249
SSAM-F	0.0397	0.0088	0.0217	0.0267
SSAM-D	0.0382	0.0084	0.0197	0.0248
MG	0.0419	0.0092	0.0219	0.0273

Table 7: Top-5 recommendation performance of DRAGON with flat local minima method SSAM and MG on Baby. SSAM-F draws upon fisher information and SSAM-D capitalizes on the principles of dynamic sparse training mask.

**Comparing with Sharpness-aware Minimization (RQ5).** Recently, several general smoothing methods [7, 14, 35, 49, 57] about flat local minima have been proposed. In this section, we compare these methods with MG in multimodal recommender systems. Specifically, we apply the recent method SSAM [30] to the multimodal recommendation model DRAGON on the dataset Baby, and the results are laid out in Table 7. The experimental results show that MG outperforms both SSAM-F, which draws upon fisher information, and SSAM-D, which capitalizes on the principles of dynamic sparse training mask [30]. These experiments suggest the MG's inherent advantage of identifying flat local minima compared with other minimization methods in the scenario of multimodal recommender systems.

**Compatibility with Various Optimizers (RQ6).** As a gradient method, it is necessary for MG to adapt to various optimizers. Therefore, in this part, we evaluate the performance of MG on the dataset Baby under various optimizers and baselines as illustrated in Table 6. The optimizers include Adam [22], SGD [2, 37], RMSprop [12], and Adagrad [9]. Here, we choose the classic multimodal recommender system GRCN and the latest state-of-the-art multimodal recommender system DRAGON as baselines. Despite the considerable performance fluctuations observed among baselines under different optimizers, MG consistently delivers a noticeable enhancement in recommendation accuracy. This consistency highlights the stable performance of MG across a range of optimizers.

Model	HR	NDCG
VBPR	0.1352	0.1005
VBPR + AMR	0.1395	0.1027
VBPR + AMR + MG	0.1457	0.1047

**Table 8: Top-5 recommendation performance of VBPR with adversarial training method AMR and MG on Pinterest.**

Dataset	Metric	LayerGCN	LayerGCN + MG	SelfCF	SelfCF + MG
Baby	REC	0.0314	0.0337	0.0329	0.0348
	PREC	0.0070	0.0075	0.0073	0.0078
	MAP	0.0168	0.0179	0.0175	0.0182
	NDCG	0.0209	0.0223	0.0217	0.0228
Clothing	REC	0.0242	0.0258	0.0246	0.0284
	PREC	0.0051	0.0054	0.0052	0.0059
	MAP	0.0136	0.0141	0.0136	0.0155
	NDCG	0.0163	0.0171	0.0164	0.0188

**Table 9: Top-5 recommendation performance of general models with or without MG.**

**Compatibility with Other Robust Recommendation Methods (RQ7).** Researchers always enhance the robustness of multimodal recommender systems in the way of adversarial training [3, 8, 27, 38]. Therefore, in this part, we perform experiments using the official code of adversarial training method AMR [38] to explore the performance of multimodal recommender systems training with both MG and the robust recommendation method. As presented in Table 8, the experimental results show that the performance of VBPR with AMR can be improved by MG, indicating that MG is compatible with adversarial training aiming at enhancing the robustness of multimodal recommender systems.

**MG for Non-multimodal Systems (RQ8).** The non-multimodal system is a special case of multimodal settings. In this part, we validate the effectiveness of MG across various classical non-multimodal recommender systems, like LightGCN [54] and SelfCF [55]. Although they are subjected to a relatively lower risk of input distribution shift, as indicated in Table 9, our proposed MG also brings a noticeable improvement in the performance of these general models, showing that even in a single-modal context, MG remains effective in identifying superior flat local minima and promoting model robustness.

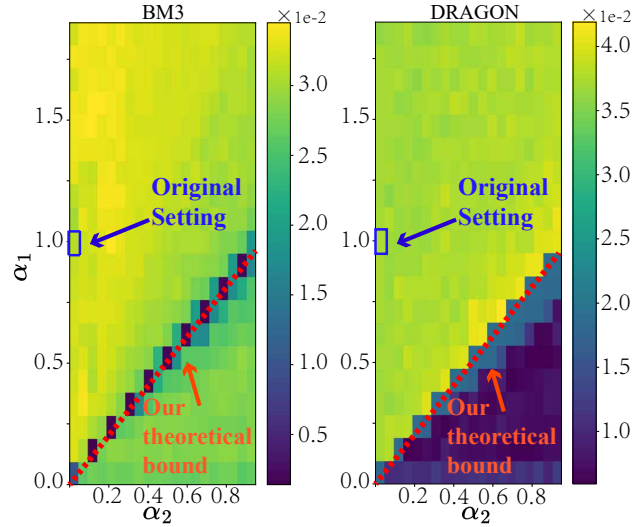
## 7 ABLATION STUDY

**The Interval  $\beta$  in Algorithm 1.**  $\beta$  denotes the interval of mirror training, where MG affects the training of multimodal recommendation models every  $\beta$  iterations. As depicted in Table 10, we set  $\beta$  to 1, 3, 5, and 7, and observe the performance of various multimodal recommendation models. Interestingly, we notice that changes in  $\beta$  do not lead to a significant impact on model performance, with 3 being a suitable choice for the value of  $\beta$ .

**The  $(\alpha_1, \alpha_2)$  in Algorithm 1.**  $\alpha_1$  and  $\alpha_2$  are the scaling coefficients while using Mirror Training. As shown in Fig. 5, when  $\alpha_1 > \alpha_2$ , MG can achieve significant performance improvement by setting appropriate  $(\alpha_1, \alpha_2)$ . From our theoretical results in Eq. (7), when  $\alpha_1 = \alpha_2$  (red dashed line in Fig. 5), the objective function degenerates to

Model	Metric	$\beta = 1$	$\beta = 3$	$\beta = 5$	$\beta = 7$
DualGNN	REC	0.0327	0.0329	<b>0.0331</b>	0.0327
	NDCG	0.0218	<b>0.0219</b>	0.0217	0.0216
BM3	REC	0.0339	<b>0.0345</b>	0.0337	0.0342
	NDCG	0.0223	<b>0.0228</b>	0.0223	0.0225
DRAGON	REC	0.0418	<b>0.0419</b>	0.0410	0.0416
	NDCG	<b>0.0274</b>	0.0273	0.0273	<b>0.0274</b>

**Table 10: Top-5 recommendation performance of models with different MG interval ( $\beta$ ) on Baby.**



**Figure 5: Recall for different  $(\alpha_1, \alpha_2)$  on Baby.**

only the regularization term, leading to optimization difficulties. On the other hand, when  $\alpha_1 < \alpha_2$  (below the red dashed line), the main term becomes a pure gradient ascent term, which might not necessarily be advantageous for optimization and potentially impact the model’s performance. These observations align well with our theoretical results.

## 8 LIMITATIONS

The selection of  $(\alpha_1, \alpha_2)$  lacks a direct and efficient method, often requiring grid search, similar to other advanced recommender system approaches [50, 52, 56]. Moreover, as per Algorithm 1, MG theoretically requires more time for iterations. Fortunately, as indicated by the results in Fig. 4, MG exhibits rapid convergence speed, making the additional computational cost acceptable. In the future, further development of MG is needed to address these limitations.

## 9 CONCLUSION

In this paper, we rethink the robustness challenges in multimodal recommender systems through the perspective of flat local minima. We propose a novel gradient strategy called Mirror Gradient (MG) to address the prevalent risk of input distribution shift in the systems. Extensive experiments across diverse multimodal recommendation models and benchmark datasets and strong theoretical evidence demonstrate the effectiveness, compatibility, and versatility of MG.



## REFERENCES

- [1] Kevin M Boehm, Pegah Khosravi, Rami Vanguri, Jianjiong Gao, and Sohrab P Shah. 2022. Harnessing multimodal data integration to advance precision oncology. *Nature Reviews Cancer* 22, 2 (2022), 114–126.
- [2] Léon Bottou. 2012. Stochastic gradient descent tricks. In *Neural Networks: Tricks of the Trade: Second Edition*. Springer, 421–436.
- [3] Huiyuan Chen and Jing Li. 2019. Adversarial tensor factorization for context-aware recommendation. In *Proceedings of the 13th ACM Conference on Recommender Systems*, 363–367.
- [4] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems* 41, 3 (2023), 1–39.
- [5] Yashar Deldjoo, Tommaso Di Noia, and Felice Antonio Merra. 2021. A survey on adversarial recommender systems: from attack/defense strategies to generative adversarial networks. *ACM Computing Surveys (CSUR)* 54, 2 (2021), 1–38.
- [6] Benoit et al Dherin. 2021. The Geometric Occam’s Razor Implicit in Deep Learning. *arXiv preprint arXiv:2111.15090* (2021).
- [7] Jiawei Du, Hanshu Yan, Jiashi Feng, Joey Tianyi Zhou, Liangli Zhen, Rick Siow Mong Goh, and Vincent YF Tan. 2021. Efficient sharpness-aware minimization for improved training of neural networks. *arXiv preprint arXiv:2110.03141* (2021).
- [8] Yali Du, Meng Fang, Jinfeng Yi, Chang Xu, Jun Cheng, and Dacheng Tao. 2018. Enhancing the robustness of neural collaborative filtering systems under malicious attacks. *IEEE Transactions on Multimedia* 21, 3 (2018), 555–565.
- [9] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research* 12, 7 (2011).
- [10] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. 2020. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412* (2020).
- [11] Chen Gao, Xiang Wang, Xiangnan He, and Yong Li. 2022. Graph neural networks for recommender system. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 1623–1625.
- [12] Alex Graves. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850* (2013).
- [13] Bowei He, Xu He, Yingxue Zhang, Ruiming Tang, and Chen Ma. 2023. Dynamically Expandable Graph Convolution for Streaming Recommendation. *arXiv preprint arXiv:2303.11700* (2023).
- [14] Haowei He, Gao Huang, and Yang Yuan. 2019. Asymmetric valleys: Beyond sharp and flat local minima. *Advances in neural information processing systems* 32 (2019).
- [15] Ruining He and Julian McAuley. 2016. VBPR: visual bayesian personalized ranking from implicit feedback. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 30.
- [16] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, 639–648.
- [17] Sepp Hochreiter and Jürgen Schmidhuber. 1994. Simplifying neural nets by discovering flat minima. *Advances in neural information processing systems* 7 (1994).
- [18] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Flat minima. *Neural computation* 9, 1 (1997), 1–42.
- [19] Zhongzhan Huang, Mingfu Liang, Senwei Liang, and Wei He. 2021. AlterSGD: Finding Flat Minima for Continual Learning by Alternative Training. *arXiv preprint arXiv:2107.05804* (2021).
- [20] Zhongzhan Huang, Wenqi Shao, Xinjiang Wang, Liang Lin, and Ping Luo. 2021. Rethinking the pruning criteria for convolutional neural network. *Advances in Neural Information Processing Systems* 34 (2021), 16305–16318.
- [21] Jean Kaddour, Linqing Liu, Ricardo Silva, and Matt J Kusner. 2022. When do flat minima optimizers work? *Advances in Neural Information Processing Systems* 35 (2022), 16577–16595.
- [22] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [23] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. 2021. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*. PMLR, 5905–5914.
- [24] Xuan Nhat Lam, Thuc Vu, Trong Duc Le, and Anh Duc Duong. 2008. Addressing cold-start problem in recommendation systems. In *Proceedings of the 2nd international conference on Ubiquitous information management and communication*, 208–211.
- [25] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. 2018. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems* 31 (2018).
- [26] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. BliP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597* (2023).
- [27] Ruirui Li, Xian Wu, and Wei Wang. 2020. Adversarial learning to compare: Self-attentive prospective customer recommendation in location based social networks. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, 349–357.
- [28] Daniele Malitesta, Giandomenico Cornacchia, Claudio Pomo, and Tommaso Di Noia. 2023. Disentangling the Performance Puzzle of Multimodal-aware Recommender Systems. In *EvalRS@ KDD (CEUR Workshop Proceedings, Vol. 3450)*. CEUR-WS.org.
- [29] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, 43–52.
- [30] Peng Mi, Li Shen, Tianhe Ren, Yiyi Zhou, Xiaoshuai Sun, Rongrong Ji, and Dacheng Tao. 2022. Make sharpness-aware minimization stronger: A sparsified perturbation approach. *Advances in Neural Information Processing Systems* 35 (2022), 30950–30962.
- [31] Juan Ni, Zhenhua Huang, Yang Hu, and Chen Lin. 2022. A two-stage embedding model for recommendation with multimodal auxiliary information. *Information Sciences* 582 (2022), 22–37.
- [32] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [33] Aghiles Salah, Quoc-Tuan Truong, and Hady W Lauw. 2020. Cornac: A comparative framework for multimodal recommender systems. *The Journal of Machine Learning Research* 21, 1 (2020), 3803–3807.
- [34] Andrew I Schein, Alexandrin Popescul, Lyle H Ungar, and David M Pennock. 2002. Methods and metrics for cold-start recommendations. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, 253–260.
- [35] Guangyuan Shi, Jiaxin Chen, Wenlong Zhang, Li-Ming Zhan, and Xiao-Ming Wu. 2021. Overcoming catastrophic forgetting in incremental few-shot learning by finding flat minima. *Advances in neural information processing systems* 34 (2021), 6747–6761.
- [36] Jiajie Su, Chaochao Chen, Weiming Liu, Fei Wu, Xiaolin Zheng, and Haoming Lyu. 2023. Enhancing Hierarchy-Aware Graph Networks with Deep Dual Clustering for Session-based Recommendation. In *Proceedings of the ACM Web Conference 2023*, 165–176.
- [37] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. 2013. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*. PMLR, 1139–1147.
- [38] Jinhui Tang, Xiaoyu Du, Xiangnan He, Fajie Yuan, Qi Tian, and Tat-Seng Chua. 2019. Adversarial training towards robust multimedia recommender system. *IEEE Transactions on Knowledge and Data Engineering* 32, 5 (2019), 855–867.
- [39] Jiayi Tang and Ke Wang. 2018. Ranking distillation: Learning compact ranking models with high performance for recommender system. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2289–2298.
- [40] 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).
- [41] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative deep learning for recommender systems. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 1235–1244.
- [42] Qifan Wang, Yinwei Wei, Jianhua Yin, Jianlong Wu, Xuemeng Song, and Liqiang Nie. 2021. Dualgnn: Dual graph neural network for multimedia recommendation. *IEEE Transactions on Multimedia* (2021).
- [43] 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.
- [44] 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.
- [45] Chenwang Wu, Defu Lian, Yong Ge, Zhihao Zhu, Enhong Chen, and Senchao Yuan. 2021. Fight fire with fire: towards robust recommender systems via adversarial poisoning training. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1074–1083.
- [46] Xixi Wu, Yun Xiong, Yao Zhang, Yizhu Jiao, Jiawei Zhang, Yangyong Zhu, and Philip S Yu. 2023. ConsRec: Learning Consensus Behind Interactions for Group Recommendation. In *Proceedings of the ACM Web Conference 2023*, 240–250.
- [47] Fuzheng Zhang, Nicholas Jing Yuan, Defu Lian, Xing Xie, and Wei-Ying Ma. 2016. Collaborative knowledge base embedding for recommender systems. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 353–362.
- [48] Lingzi Zhang, Yong Liu, Xin Zhou, Chunyan Miao, Guoxin Wang, and Haihong Tang. 2022. Diffusion-based graph contrastive learning for recommendation with

1045	implicit feedback. In <i>International Conference on Database Systems for Advanced Applications</i> . Springer, 232–247.	1103
1046	[49] Yang Zhao, Hao Zhang, and Xiuyuan Hu. 2022. Penalizing gradient norm for efficiently improving generalization in deep learning. In <i>International Conference on Machine Learning</i> . PMLR, 26982–26992.	1104
1047	[50] Hongyu Zhou, Xin Zhou, and Zhiqi Shen. 2023. Enhancing Dyadic Relations with Homogeneous Graphs for Multimodal Recommendation. <i>arXiv preprint arXiv:2301.12097</i> (2023).	1105
1048	[51] Hongyu Zhou, Xin Zhou, Zhiwei Zeng, Lingzi Zhang, and Zhiqi Shen. 2023. A Comprehensive Survey on Multimodal Recommender Systems: Taxonomy, Evaluation, and Future Directions. <i>arXiv preprint arXiv:2302.04473</i> (2023).	1106
1049	[52] Xin Zhou. 2022. A tale of two graphs: Freezing and denoising graph structures for multimodal recommendation. <i>arXiv preprint arXiv:2211.06924</i> (2022).	1107
1050	[53] Xin Zhou, Donghui Lin, and Toru Ishida. 2016. Evaluating reputation of web services under rating scarcity. In <i>2016 IEEE International Conference on Services Computing (SCC)</i> . IEEE, 211–218.	1108
1051	[54] Xin Zhou, Donghui Lin, Yong Liu, and Chunyan Miao. 2023. Layer-refined graph convolutional networks for recommendation. In <i>2023 IEEE 39th International Conference on Data Engineering (ICDE)</i> . IEEE, 1247–1259.	1109
1052	[55] Xin Zhou, Aixin Sun, Yong Liu, Jie Zhang, and Chunyan Miao. 2023. Selfcf: A simple framework for self-supervised collaborative filtering. <i>ACM Transactions on Recommender Systems</i> 1, 2 (2023), 1–25.	1110
1053	[56] Xin Zhou, Hongyu Zhou, Yong Liu, Zhiwei Zeng, Chunyan Miao, Pengwei Wang, Yuan You, and Feijun Jiang. 2023. Bootstrap latent representations for multi-modal recommendation. In <i>Proceedings of the ACM Web Conference 2023</i> . 845–854.	1111
1054	[57] Juntang Zhuang, Boqing Gong, Liangzhe Yuan, Yin Cui, Hartwig Adam, Nicha Dvornek, Sekhar Tatikonda, James Duncan, and Ting Liu. 2022. Surrogate gap minimization improves sharpness-aware training. <i>arXiv preprint arXiv:2203.08065</i> (2022).	1112
1055		1113
1056		1114
1057		1115
1058		1116
1059		1117
1060		1118
1061		1119
1062		1120
1063		1121
1064		1122
1065		1123
1066		1124
1067		1125
1068		1126
1069		1127
1070		1128
1071		1129
1072		1130
1073		1131
1074		1132
1075		1133
1076		1134
1077		1135
1078		1136
1079		1137
1080		1138
1081		1139
1082		1140
1083		1141
1084		1142
1085		1143
1086		1144
1087		1145
1088		1146
1089		1147
1090		1148
1091		1149
1092		1150
1093		1151
1094		1152
1095		1153
1096		1154
1097		1155
1098		1156
1099		1157
1100		1158
1101		1159
1102		1160

## A SUPPLEMENTARY EXPERIMENTS

Throughout all the experiments in this paper, we adopted the Top-5 recommendation performance as the metric due to its capacity to reflect user preferences better, given the heightened significance of recommendations in top positions of rank lists [39]. However, enhancing Top-5 performance is a challenging job on Amazon datasets. The current advanced multimodal recommender system methods, including matrix factorization (VBPR [15]), self-supervised learning (SLMRec [40], BM3 [56]), graph neural networks (MMGCN [44], GRCN [43], DualGNN [42], FREEDOM [52], DRAGON [50]), still achieve relatively low value on Top-5 recommendation performance.

In fact, despite the challenge in enhancing this metric, we believe that the improvement brought about by the proposed MG in this paper is substantial and significant. Specifically, in Table 11 and Table 12, we present the performance evaluation of the BM3 and FREEDOM methods on the dataset Baby under various random seeds. Compared to the performance of the original method under different random seeds, the improvement achieved by MG is notable. This implies that the enhancements from MG are substantial and not obtained through random perturbations or simple parameter-tuning.

BM3	REC	Improv.	PREC	Improv.	MAP	Improv.	NDCG	Improv.
Origin	0.0327	0.00%	0.0072	0.00%	0.0174	0.00%	0.0216	0.00%
Random seed	0.0329	0.61%	0.0073	1.39%	0.0176	1.15%	0.0218	0.93%
	0.0329	0.61%	0.0073	1.39%	0.0176	1.15%	0.0218	0.93%
	0.0329	0.61%	0.0073	1.39%	0.0173	-0.57%	0.0216	0.00%
	0.0323	-1.22%	0.0072	0.00%	0.0175	0.57%	0.0216	0.00%
	0.0323	-1.22%	0.0072	0.00%	0.0174	0.00%	0.0216	0.00%
	0.0322	-1.53%	0.0072	0.00%	0.0174	0.00%	0.0215	-0.46%
	0.0322	-1.53%	0.0072	0.00%	0.0174	0.00%	0.0215	-0.46%
	0.0321	-1.83%	0.0072	0.00%	0.0174	0.00%	0.0215	-0.46%
	0.0321	-1.83%	0.0072	0.00%	0.0174	0.00%	0.0215	-0.46%
MG	0.0345	5.50%	0.0077	6.94%	0.0183	5.17%	0.0228	5.56%

**Table 11: Top-5 recommendation performance of BM3 on the dataset Baby. "Random seed" denotes that the model is trained with a random seed.**

FREEDOM	REC	Improv.	PREC	Improv.	MAP	Improv.	NDCG	Improv.
Origin	0.0374	0.00%	0.0083	0.00%	0.0194	0.00%	0.0243	0.00%
Random seed	0.0376	0.53%	0.0083	0.00%	0.0193	-0.52%	0.0243	0.00%
	0.0376	0.53%	0.0082	-1.20%	0.0195	0.52%	0.0246	1.23%
	0.0374	0.00%	0.0082	-1.20%	0.0194	0.00%	0.0245	0.82%
	0.0374	0.00%	0.0083	0.00%	0.0195	0.52%	0.0246	1.23%
	0.0373	-0.27%	0.0082	-1.20%	0.0196	1.03%	0.0244	0.41%
	0.0373	-0.27%	0.0082	-1.20%	0.0196	1.03%	0.0245	0.82%
	0.0373	-0.27%	0.0082	-1.20%	0.0196	1.03%	0.0245	0.82%
	0.0373	-0.27%	0.0082	-1.20%	0.0196	1.03%	0.0245	0.82%
	0.0372	-0.53%	0.0082	-1.20%	0.0196	1.03%	0.0245	0.82%
MG	0.0397	6.15%	0.0088	6.02%	0.0209	7.73%	0.0261	7.41%

**Table 12: Top-5 recommendation performance of FREEDOM on the dataset Baby. "Random seed" denotes that the model is trained with a random seed.**