

# Multi-Aspect Cross-modal Quantization for Generative Recommendation

Fuwei Zhang<sup>1</sup>, Xiaoyu Liu<sup>1</sup>, Dongbo Xi<sup>2</sup>, Jishen Yin<sup>2</sup>, Huan Chen<sup>2</sup>, Peng Yan<sup>2</sup>, Fuzhen Zhuang<sup>1,3</sup>,  
Zhao Zhang<sup>3\*</sup>

<sup>1</sup>Institute of Artificial Intelligence, Beihang University

<sup>2</sup>Meituan

<sup>3</sup>SKLCCSE, School of Computer Science and Engineering, Beihang University

{zhangfuwei, liuxiaoyv, zhuangfuzhen}@buaa.edu.cn, {xidongbo,yinjishen,chenhuan15,yanpeng04}@meituan.com, zhangzhao.cs.ai@gmail.com

## Abstract

Generative Recommendation (GR) has emerged as a new paradigm in recommender systems. This approach relies on quantized representations to discretize item features, modeling users' historical interactions as sequences of discrete tokens. Based on these tokenized sequences, GR predicts the next item by employing next-token prediction methods. The challenges of GR lie in constructing high-quality semantic identifiers (IDs) that are hierarchically organized, minimally conflicting, and conducive to effective generative model training. However, current approaches remain limited in their ability to harness multimodal information and to capture the deep and intricate interactions among diverse modalities, both of which are essential for learning high-quality semantic IDs and for effectively training GR models. To address this, we propose Multi-Aspect Cross-modal quantization for generative Recommendation (MACRec), which introduces multimodal information and incorporates it into both semantic ID learning and generative model training from different aspects. Specifically, we first introduce cross-modal quantization during the ID learning process, which effectively reduces conflict rates and thus improves codebook usability through the complementary integration of multimodal information. In addition, to further enhance the generative ability of our GR model, we incorporate multi-aspect cross-modal alignments, including the implicit and explicit alignments. Finally, we conduct extensive experiments on three well-known recommendation datasets to demonstrate the effectiveness of our proposed method.

**Code** — <https://github.com/zhangfw123/MACRec>

## Introduction

Recommendation systems play a crucial role in helping users navigate information overload by providing personalized item suggestions. As a result, they have now been widely applied across various domains, such as e-commerce (Smith and Linden 2017; Chen et al. 2019; Xiaoyu et al. 2025; Zhang et al. 2022a; Xi et al. 2020; Zhang et al. 2024, 2022b; Chen et al. 2024, 2025), social media platforms (Davidson et al. 2010; Covington, Adams, and Sargin 2016), online advertising (Xi et al. 2019, 2021), and

\*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

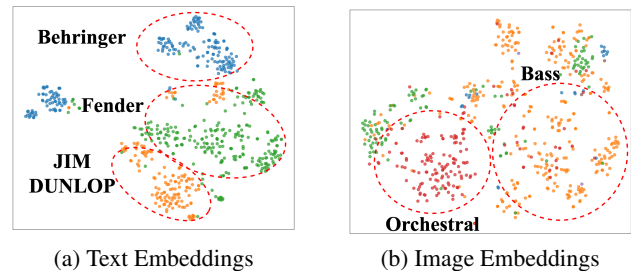


Figure 1: Visualization examples of different modality embeddings on Instrument.

short video platforms (Zhu et al. 2024, 2025; Bin et al. 2025).

With the development of large language models (LLMs), numerous LLM-enhanced recommendation approaches (Du et al. 2024) have been proposed. Among them, a new paradigm called generative recommendation (GR) has recently gained significant attention. In GR, the recommendation task is reformulated as a next-token prediction problem, where the model is provided with a sequence of item semantic identifiers representing a user's interaction history and is tasked with generating the semantic identifier of the next item that aligns with the user's interests. By harnessing the powerful natural language understanding and sequence modeling capabilities of LLMs, GR offers enhanced flexibility and expressiveness, enabling more accurate modeling of user intent and richer contextual integration (Li et al. 2024; Zhang et al. 2025a,b; Liu et al. 2025).

A number of studies have explored the paradigm of GR. These works primarily focus on quantizing item embeddings into unique semantic identifiers (Rajput et al. 2023; Wang et al. 2024; Zhai et al. 2025), and subsequently constructing user interaction sequences via the item semantic ID sequences. In this framework, recommendations are produced by generating the semantic IDs of recommended items. Existing approaches predominantly emphasize the discretization of textual embeddings to form these semantic IDs. However, relying solely on embeddings derived from a single modality might lead to limited semantic discriminability. For example, products from the same brand, such as differ-

ent types of instruments produced by that brand, often have smaller distances between their encoded representations due to the prominence of brand-related textual features. This reduces the distinctiveness of each instrument, making it more difficult to differentiate between products of the same brand with different functions. Such interference may reduce the quality of semantic ID representations, potentially affecting recommendation performance. Moreover, commonly used quantization methods such as Residual Quantization Variational Autoencoder (RQ-VAE) are prone to significant semantic loss in deeper hierarchical structures (Zhou et al. 2025a,b). This semantic loss may cause the model to lack clear semantic guidance when assigning tokens, resulting in an almost random distribution and thereby weakening the semantic hierarchy of the generated item representations.

To address this challenge, we incorporate multimodal information, including images, to overcome the limitations of textual representations. Visual features in images, such as shape and color, are often more distinctive and sensitive than textual descriptions, thereby providing more effective support for understanding and conveying information. In Figure 1, we visualize the embeddings generated from text and images, respectively. The results indicate that text-based embeddings are more effective at clustering items from the same brand, while image-based embeddings excel at distinguishing between different types of musical instruments. These observations suggest that different modalities capture complementary aspects of user preferences. Based on these findings, we argue that relying solely on a single modality is insufficient to comprehensively represent the semantic characteristics of items. Thus, we integrate multimodal item information into the GR to enhance semantic expressiveness and improve the model’s generative capability.

To address these challenges, we propose **Multi-Aspect Cross-modal quantization for generative Recommendation (MACRec)**, a novel generative recommendation framework that effectively integrates multimodal information. By leveraging cross-modal information, MACRec learns hierarchically meaningful semantic IDs for items and enhances the training of GR models. Specifically, during the item quantization stage, we propose a cross-modal quantization method that incorporates cross-modal contrastive learning into each layer of residual quantization to enhance information interaction across different modalities and reduce semantic loss. Meanwhile, we leverage multimodal alignment to optimize the reconstructed representations, thereby further enhancing the representational capacity of the codebook. During the training phase of the GR model, we employ multi-aspect alignment strategies to enhance the model’s understanding of semantic IDs and to enable the learning of shared features across different modalities. These strategies include implicit alignment in the latent space through contrastive methods, as well as explicit alignment within the generative task. Through a series of cross-modal interactions, our model achieves significant improvements in recommendation performance.

Here, we summarize our contributions:

- To capture richer and more discriminative semantics, we propose a novel cross-modal quantization method that

integrates contrastive learning into residual quantization and reconstruction, yielding hierarchically meaningful semantic IDs for items.

- To enable the model to learn common features from different modalities, we employ multi-aspect alignment strategies, including both implicit alignment in the latent space and explicit alignment in the generative task.
- We conduct experiments on three widely used recommendation datasets, and our approach significantly outperforms state-of-the-art GR models.

## Related Work

In the related work section, we mainly introduce traditional sequential recommendation methods and generative recommendation methods under different modalities.

### Sequential Recommendation

Single-modal sequential recommendation focuses on modeling user behavior sequences based solely on interaction data, aiming to capture users’ dynamic preferences over time. Early approaches employ neural networks such as GRU4Rec (Hidasi et al. 2015), STAMP (Liu et al. 2018), and NARM (Li et al. 2017) to learn sequential patterns and temporal dependencies. The introduction of attention-based models like SASRec (Kang and McAuley 2018) further improves the ability to model long-range dependencies within sequences. More recently, pretrained language models (PLMs) such as BERT4Rec (Sun et al. 2019) have significantly advanced performance by leveraging large-scale self-supervised pretraining. In addition, prompt-based methods like P5 (Geng et al. 2022) and M6-Rec (Cui et al. 2022) reformulate recommendation tasks as language modeling problems, which further enhances model generalization and flexibility.

Multi-modal sequential recommendation enriches sequential representations by incorporating various item modalities (e.g., text, images), thereby improving recommendation quality (Liu et al. 2024b). Recent approaches commonly employ deep and graph neural networks, such as MMGCN (Wei et al. 2019) and GRCN (Wei et al. 2020), to integrate heterogeneous features for enhanced user-item interaction modeling. Additionally, contrastive learning and multimodal pretraining methods such as MMGCL (Yi et al. 2022) and MISSRec (Wang et al. 2023) further strengthen user interest modeling. The VIP5 (Geng et al. 2023) framework extends prompt-based techniques to multimodal settings and has advanced the field’s performance.

### Generative Recommendation

With the rapid development of large language models (LLMs), the potential of generative recommendation (GR) has attracted increasing attention. Early works such as TIGER (Rajput et al. 2023) discretize item sequences into tokens, enabling generative recommendation paradigms. LC-Rec (Zheng et al. 2024) utilizes the natural language understanding abilities of LLMs to support diverse task-specific fine-tuning in recommendation. LETTER (Wang

et al. 2024) extends TIGER by introducing collaborative filtering embeddings and an additional loss function to improve codebook utilization. In multimodal generative recommendation, MMGRec (Liu et al. 2024a) uses a Graph RQ-VAE to generate item representations by integrating multimodal features with collaborative signals. MQL4GRec (Zhai et al. 2025) further advances the field by encoding multimodal and cross-domain item information into a unified quantized language, facilitating knowledge transfer and achieving better performance than previous methods.

However, existing multimodal GR models typically encode each modality separately to obtain semantic IDs for different modalities. They do not consider cross-modal interactions during the quantization process, which makes them more prone to hierarchical semantic loss. In this paper, we are the first to introduce cross-modal learning during quantization, enabling IDs to capture the advantageous features of different modalities. In addition, we also incorporate both implicit and explicit alignment methods during GR training, allowing features from different modalities to complement each other more effectively.

## Methodology

In this section, we organize our approach into two main modules. Figure 2 illustrates the overall architecture, which includes a cross-modal item quantization module for generating discrete semantic IDs, and the training phase of the GR model with multi-aspect alignment.

### Cross-modal Item Quantization

Original RQ-VAE approaches primarily focus on quantizing a single embedding into discrete tokens. However, when dealing with multimodal information such as images and text, simply concatenating their embeddings and quantizing them as a unified token set, or independently quantizing each modality into separate tokens, both have some limitations. First, the dimensions of image and text embeddings often differ substantially, and straightforward concatenation tends to bias the quantization process toward the modality with higher dimensionality. Second, independently quantizing each modality and subsequently training the GR model fails to fully exploit the complementary nature of cross-modal information. In response to these challenges, we propose the Cross-modal Item Quantization framework. Our method effectively utilizes multimodal contrastive learning in the learning of quantization and reconstruction, enabling the generation of semantically discriminative item identifiers while simultaneously reducing identifier collision rates.

**Dual-modality Pseudo-label Generation** To perform quantization across different modalities using contrastive learning, we first generate pseudo-labels by clustering items according to their text and visual embeddings, which are then used to construct positive samples. Specifically, for a given item  $i$ , its text information is encoded into an embedding  $\mathbf{t}_i$  using an open-source large language model such as LLaMA (Touvron et al. 2023). Simultaneously, the visual

content of the item’s image is encoded into an embedding  $\mathbf{v}_i$  using a Vision Transformer (ViT) (Dosovitskiy et al. 2020).

Subsequently, we perform K-means clustering independently on the textual and visual embeddings, partitioning them into  $K$  clusters. The clustering process can be formulated as follows:

$$\mathcal{C}_{\text{text}} = \text{KMeans}(\{\mathbf{t}_i\}_{i=1}^N), \quad \mathcal{C}_{\text{vision}} = \text{KMeans}(\{\mathbf{v}_i\}_{i=1}^N), \quad (1)$$

where  $\mathcal{C}_{\text{text}}$  and  $\mathcal{C}_{\text{vision}}$  denote the resulting cluster assignments (pseudo-labels) for the text and vision modalities, respectively, and  $N$  is the total number of items.

**Cross-modal Quantization with Contrastive Learning** Residual-Quantized Variational AutoEncoder (RQ-VAE) (Lee et al. 2022; Rajput et al. 2023) is an embedding quantization method that builds upon multi-layer vector quantization (VQ). In this framework, VQ discretizes continuous embeddings by mapping them to the closest entries in a learnable codebook, thereby effectively compressing the representation space. Residual quantization (RQ) further enhances this process by sequentially applying multiple VQ layers, where each layer performs VQ on the residuals from the previous layer. In multimodal scenarios, it is necessary to quantize information from different modalities. For a given item, both the text and visual embeddings are first processed by an encoder composed of a multi-layer perceptron (MLP) to obtain latent representations, which are denoted as  $\mathbf{z}^t = \text{T-Encoder}(\mathbf{t})$  and  $\mathbf{z}^v = \text{V-Encoder}(\mathbf{v})$ , respectively. Here,  $\mathbf{z}^t$  and  $\mathbf{z}^v$  represent the latent representation for text and visual information, respectively, which are used as the residuals ( $\mathbf{r}_0^t = \mathbf{z}^t, \mathbf{r}_0^v = \mathbf{z}^v$ ) for quantization in the first VQ layer.

At the  $l$ -th layer, there is a learnable codebook  $\mathcal{C}_l^{v/t} = \{\mathbf{e}_{l,k}^{v/t}\}_{k=0}^M$  for each modality, where  $M$  denotes the codebook size. For each modality, the residual at the  $l$ -th layer is quantized by finding the closest codebook vector for the current residual, as shown below.

$$c_l^t = \arg \min_k \|\mathbf{r}_l^t - \mathbf{e}_{l,k}^t\|_2, \quad c_l^v = \arg \min_k \|\mathbf{r}_l^v - \mathbf{e}_{l,k}^v\|_2, \quad (2)$$

$$\mathbf{r}_{l+1}^t = \mathbf{r}_l^t - \mathbf{e}_{l,c_l^t}^t, \quad \mathbf{r}_{l+1}^v = \mathbf{r}_l^v - \mathbf{e}_{l,c_l^v}^v, \quad (3)$$

where  $c_l^t$  and  $c_l^v$  represent the codewords for text and visual information at the  $l$ -th layer of residual quantization.  $\mathbf{r}_{l+1}^t$  and  $\mathbf{r}_{l+1}^v$  are the residual vectors for the next layer.

However, the quantization of text and visual embeddings described above is performed in a fully independent manner, without any interaction between the two modalities. This leads to several drawbacks: (1) Due to the similarity between certain embeddings in the text and visual domains, codebook collapse may occur, resulting in low utilization; (2) The complementary strengths of text and visual information in representing different aspects of the data are not fully leveraged. To address these issues, we introduce cross-modal contrastive learning by leveraging the multimodal pseudo-labels constructed in the previous section. Specifically, we optimize the residual representations in each layer of the

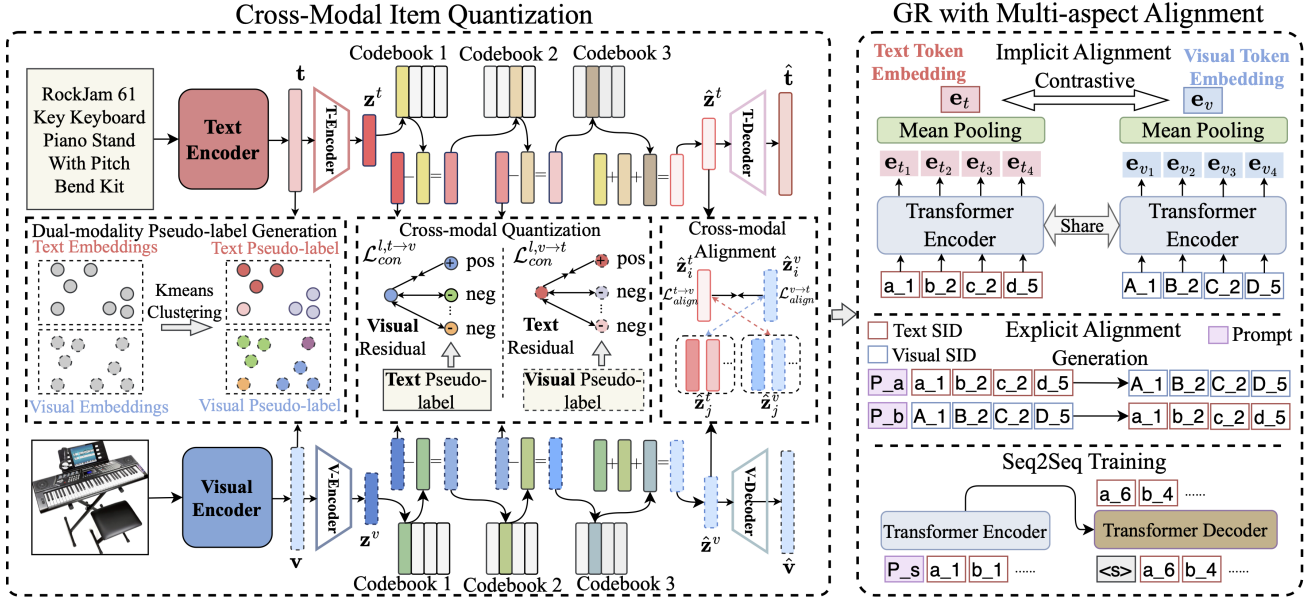


Figure 2: Overall architecture of MACRec. Left: Cross-modal Item Quantization, including Dual-modality Pseudo-label Generation, Cross-modal Quantization, and Cross-modal Reconstruction Alignment, aiming to generate high-quality semantic IDs across different modalities. Right: Generative Recommendation with Multi-aspect Alignment, aligning features from different modalities via implicit alignment, explicit alignment, and training the GR model with Seq2Seq task.

codebooks. More concretely, we use visual pseudo-labels to enhance the residual representations of the text modality, and conversely, use textual pseudo-labels to optimize those of the visual modality. The detailed InfoNCE (Oord, Li, and Vinyals 2018) loss for the  $l$ -th layer are as follows:

$$\mathcal{L}_{\text{con}}^{l,v \rightarrow t} = -\frac{1}{B} \sum_{i=1}^B \log \left( \frac{\exp(\langle \mathbf{r}_i^t, \mathbf{r}_{i,\text{pos}}^t \rangle / \tau)}{\sum_{j=1}^B \exp(\langle \mathbf{r}_i^t, \mathbf{r}_j^t \rangle / \tau)} \right), \quad (4)$$

$$\mathcal{L}_{\text{con}}^{l,t \rightarrow v} = -\frac{1}{B} \sum_{i=1}^B \log \left( \frac{\exp(\langle \mathbf{r}_i^v, \mathbf{r}_{i,\text{pos}}^v \rangle / \tau)}{\sum_{j=1}^B \exp(\langle \mathbf{r}_i^v, \mathbf{r}_j^v \rangle / \tau)} \right), \quad (5)$$

$$\mathcal{L}_{\text{con}}^l = \mathcal{L}_{\text{con}}^{l,t \rightarrow v} + \mathcal{L}_{\text{con}}^{l,v \rightarrow t}, \quad (6)$$

where  $\mathbf{r}_{i,\text{pos}}^t$  and  $\mathbf{r}_{i,\text{pos}}^v$  represent the positive samples for the  $i$ -th item in the batch that share the same vision pseudo-label  $\mathcal{C}_{\text{vision}}$  and text pseudo-label  $\mathcal{C}_{\text{text}}$ , respectively. Here,  $\langle \cdot, \cdot \rangle$  denotes the inner product,  $B$  is the batch size, and  $\tau$  is the temperature parameter used in the contrastive loss.

**Cross-modal Reconstruction Alignment** At the same time, given the  $L$  layers of codebooks, the quantized representation can be obtained by summing the corresponding codebook vectors of each layer for the item, denoted as  $\hat{\mathbf{z}}^t = \sum_{l=0}^{L-1} \mathbf{e}_{l,c_k^t}^t$ ,  $\hat{\mathbf{z}}^v = \sum_{l=0}^{L-1} \mathbf{e}_{l,c_k^v}^v$ . In order to further utilize the quantized representations from different modalities to refine the codebook representations and balance codebook utilization, we introduce an alignment loss based on contrastive learning. This loss encourages bidirectional alignment between the quantized representations of different modalities

for the same item, as formulated below:

$$\mathcal{L}_{\text{align}}^{t \rightarrow v} = -\frac{1}{B} \sum_{i=1}^B \log \left( \frac{\exp(\langle \hat{\mathbf{z}}_i^t, \hat{\mathbf{z}}_i^v \rangle / \tau)}{\sum_{j=1}^B \exp(\langle \hat{\mathbf{z}}_i^t, \hat{\mathbf{z}}_j^v \rangle / \tau)} \right), \quad (7)$$

$$\mathcal{L}_{\text{align}}^{v \rightarrow t} = -\frac{1}{B} \sum_{i=1}^B \log \left( \frac{\exp(\langle \hat{\mathbf{z}}_i^v, \hat{\mathbf{z}}_i^t \rangle / \tau)}{\sum_{j=1}^B \exp(\langle \hat{\mathbf{z}}_i^v, \hat{\mathbf{z}}_j^t \rangle / \tau)} \right), \quad (8)$$

$$\mathcal{L}_{\text{align}} = \mathcal{L}_{\text{align}}^{t \rightarrow v} + \mathcal{L}_{\text{align}}^{v \rightarrow t}, \quad (9)$$

where  $\hat{\mathbf{z}}_i^t, \hat{\mathbf{z}}_i^v$  are the quantization embeddings for text and vision modality of the same item.

Similar to the RQ-VAE architecture, we decode and reconstruct the quantized representations of different modalities separately. The decoded textual and visual embeddings can be represented as  $\hat{\mathbf{t}} = \text{T-Decoder}(\hat{\mathbf{z}}^t)$  and  $\hat{\mathbf{v}} = \text{V-Decoder}(\hat{\mathbf{z}}^v)$ , respectively. The training of RQ-VAE can be achieved by optimizing both the reconstruction loss and the residual quantization loss, as follows:

$$\mathcal{L}_{\text{recon}}^t = \|\mathbf{t} - \hat{\mathbf{t}}\|_2^2, \mathcal{L}_{\text{recon}}^v = \|\mathbf{v} - \hat{\mathbf{v}}\|_2^2, \quad (10)$$

$$\mathcal{L}_{\text{rq}}^m = \sum_{l=0}^{L-1} (\|\text{sg}[\mathbf{r}_l^m] - \mathbf{e}_{l,c_k^m}^m\|_2^2 + \alpha \|\mathbf{r}_l^m - \text{sg}[\mathbf{e}_{l,c_k^m}^m]\|_2^2), \quad (11)$$

$$\mathcal{L}_{\text{RQ-VAE}} = \mathcal{L}_{\text{recon}}^t + \mathcal{L}_{\text{recon}}^v + \mathcal{L}_{\text{rq}}^t + \mathcal{L}_{\text{rq}}^v, \quad (12)$$

where  $m$  denotes different modalities.  $\text{sg}[\cdot]$  represents the stop-gradient operation.  $\alpha$  is a loss coefficient. The superscripts  $t$  and  $v$  denote the textual and visual modalities, respectively.

Finally, the overall training objective for learning the semantic identifier, denoted as  $\mathcal{L}_{ID}$ , is defined as follows:

$$\mathcal{L}_{ID} = \mathcal{L}_{RQ-VAE} + \lambda_{con}^l \sum_{l=0}^{L-1} \mathcal{L}_{con}^l + \lambda_{align} \mathcal{L}_{align}, \quad (13)$$

where  $\mathcal{L}_{RQ-VAE}$  represents the reconstruction and codebook learning losses for both text and visual modalities, and  $\lambda_{con}^l$  and  $\lambda_{align}$  are trade-off hyperparameters that balance the contribution of the contrastive loss  $\mathcal{L}_{con}$  and the alignment loss  $\mathcal{L}_{align}$ , respectively.

For cases where conflicts occur among certain item IDs, we adopt the same conflict resolution strategy as proposed by Zhai et al., where codewords are reassigned according to the distance between items and the codebook.

### Generative Recommendation with Multi-aspect Alignment

Through the aforementioned trained RQ-VAE model, we obtain discrete semantic IDs for both texts and images, which can be represented as “<a.1><b.2><c.3>” for text and “<A.1><B.2><C.3>” for images. For the training of GR models, it is necessary to construct Seq2Seq training data based on these discrete semantic IDs of items. The recommendation model is then trained in a next-token prediction manner. To further optimize the sharing and interaction of information across different modalities, we design both implicit alignment and explicit alignment mechanisms.

#### Implicit Alignment for Cross-modal Semantic IDs

First, we aim for the model to better recognize the commonality between the semantic IDs of different modalities belonging to the same item. To achieve this goal, we align them at the latent space level after encoding. Specifically, based on the encoder-decoder architecture as the GR model, we encode both the textual and visual semantic IDs into latent representations using the encoder of GR model. We then align the representations of different modalities for the same item in the latent space through contrastive learning. Suppose the textual semantic ID of an item is denoted as  $t-sid$  and the visual semantic ID as  $v-sid$ . The implementation is as follows,

$$\mathbf{e}^t = \text{MeanPool}(\text{T5-Encoder}(t-sid)), \quad (14)$$

$$\mathbf{e}^v = \text{MeanPool}(\text{T5-Encoder}(v-sid)), \quad (15)$$

$$\mathcal{L}_{implicit}^{t \rightarrow v} = -\frac{1}{B} \sum_{i=1}^B \log \left( \frac{\exp(\langle \mathbf{e}_i^t, \mathbf{e}_i^v \rangle / \tau)}{\sum_{j=1}^B \exp(\langle \mathbf{e}_i^t, \mathbf{e}_j^v \rangle / \tau)} \right), \quad (16)$$

$$\mathcal{L}_{implicit}^{v \rightarrow t} = -\frac{1}{B} \sum_{i=1}^B \log \left( \frac{\exp(\langle \mathbf{e}_i^v, \mathbf{e}_i^t \rangle / \tau)}{\sum_{j=1}^B \exp(\langle \mathbf{e}_i^v, \mathbf{e}_j^t \rangle / \tau)} \right), \quad (17)$$

$$\mathcal{L}_{implicit} = \mathcal{L}_{implicit}^{t \rightarrow v} + \mathcal{L}_{implicit}^{v \rightarrow t}. \quad (18)$$

#### Explicit Alignment with Different Generation Tasks

Furthermore, for most of the generative models, we can align the representations of images and items by designing different training tasks. Inspired by Zhai et al., we propose both item-level and sequence-level cross-modal alignment

strategies. For item-level alignment, we use the textual semantic ID of an item as input to generate its visual semantic ID, and vice versa, using the visual semantic ID as input to generate the textual semantic ID. For sequence-level alignment, we construct a prediction task where a historical sequence of textual semantic IDs is used to predict the visual semantic ID of the next recommended item. Similarly, a sequence of visual semantic IDs is used to predict the textual semantic ID of the next item. These additional explicit alignment tasks are incorporated into the sequential recommendation training.

**Training Objections and Inference** For multimodal GR, there are two main tasks for recommendation. The first is to predict the textual semantic ID of the next item based on the historical sequence of item textual semantic IDs. The second is to predict the visual semantic ID of the next item using the historical sequence of item visual semantic IDs. By integrating the aforementioned alignment strategies, the final training objective is formulated as follows:

$$\mathcal{L}_{rec} = -\sum_{t=1}^{|y|} \log P_{\theta}(y_t | y < t, x) + \lambda_{implicit} \mathcal{L}_{implicit}. \quad (19)$$

During the inference stage, we generate multiple candidate semantic IDs for different modalities using constrained beam search (Rajput et al. 2023). Finally, the ensemble of the results (Zhai et al. 2025) from two modalities is performed by averaging the scores of both modalities to obtain the final inference result.

## Experiment

We analyze our model’s effectiveness and address: 1) **RQ1**: How does MACRec compare to state-of-the-art baselines? 2) **RQ2**: How do different modules impact MACRec’s performance? 3) **RQ3**: What is the impact of our method on the item collision rate? 4) **RQ4**: What is the impact of our method on the distribution of code allocation? 5) **RQ5**: How do different hyperparameters affect MACRec?

### Experimental Setup

**Datasets.** We employ three real-world recommendation datasets, all constructed from the Amazon Product Reviews dataset, which contains user reviews and item metadata collected between May 1996 and October 2018. Specifically, we conducted experiments on datasets from three different categories: “Musical Instruments”, “Arts, Crafts and Sewing”, and “Video Games”. The detailed statistics of these datasets are summarized in Table 1.

Datasets	#Users	#Items	#Interactions	Sparsity	Avg. len
Instruments	17112	6250	136226	99.87%	7.96
Arts	22171	9416	174079	99.92%	7.85
Games	42259	13839	373514	99.94%	8.84

Table 1: Statistics of the datasets. **Avg. len** represents the average length of item sequences.

Dataset	Metrics	BERT4Rec	SASRec	FDSA	S <sup>3</sup> -Rec	MISSRec	P5-CID	VIP5	TIGER	MQL4GRec	MACRec
Instruments	HR@1	0.0450	0.0318	0.0530	0.0339	0.0723	0.0512	0.0737	0.0754	<u>0.0763</u>	<b>0.0819*</b>
	HR@5	0.0856	0.0946	0.0987	0.0937	<u>0.1089</u>	0.0839	0.0892	0.1007	0.1058	<b>0.1110*</b>
	HR@10	0.1081	0.1233	0.1249	0.1123	<u>0.1361</u>	0.1119	0.1071	0.1221	0.1291	<b>0.1363</b>
	NDCG@5	0.0667	0.0654	0.0775	0.0693	0.0797	0.0678	0.0815	0.0882	0.0902	<b>0.0965*</b>
	NDCG@10	0.0739	0.0746	0.0859	0.0743	0.0880	0.0704	0.0872	0.0950	0.0997	<b>0.1046*</b>
Arts	HR@1	0.0289	0.0212	0.0380	0.0172	0.0479	0.0421	0.0474	0.0532	<u>0.0626</u>	<b>0.0685*</b>
	HR@5	0.0697	0.0951	0.0832	0.0739	0.1021	0.0713	0.0704	0.0894	<u>0.0997</u>	<b>0.1046*</b>
	HR@10	0.0922	0.1250	0.1190	0.1030	<u>0.1321</u>	0.0994	0.0959	0.1167	0.1254	<b>0.1329*</b>
	NDCG@5	0.0502	0.0610	0.0583	0.0511	0.0699	0.0607	0.0586	0.0718	<u>0.0816</u>	<b>0.0868*</b>
	NDCG@10	0.0575	0.0706	0.0695	0.0630	0.0815	0.0662	0.0635	0.0806	<u>0.0898</u>	<b>0.0953*</b>
Games	HR@1	0.0115	0.0069	0.0163	0.0136	<u>0.0201</u>	0.0169	0.0173	0.0166	0.0200	<b>0.0208*</b>
	HR@5	0.0426	0.0587	0.0614	0.0527	<b>0.0674</b>	0.0532	0.0480	0.0523	0.0645	<u>0.0671</u>
	HR@10	0.0725	0.0985	0.0988	0.0903	<u>0.1048</u>	0.0824	0.0758	0.0857	0.1007	<b>0.1078*</b>
	NDCG@5	0.0270	0.0333	0.0389	0.0351	0.0385	0.0331	0.0328	0.0345	<u>0.0421</u>	<b>0.0435*</b>
	NDCG@10	0.0366	0.0461	0.0509	0.0468	0.0499	0.0454	0.0418	0.0453	<u>0.0538</u>	<b>0.0565*</b>

Table 2: Performance comparison of methods on three datasets. The best and second-best results are in bold and underlined, respectively. \* denotes statistical significance ( $p$ -value  $< 0.05$ ) against the best baseline.

**Baselines.** To evaluate our approach, we compare it with representative recent methods, including BERT4Rec (Sun et al. 2019), SASRec (Kang and McAuley 2018), FDSA (Zhang et al. 2019), S<sup>3</sup>-Rec (Zhou et al. 2020), MISSRec (Wang et al. 2023), P5 (Geng et al. 2022), VIP5 (Geng et al. 2023), TIGER (Rajput et al. 2023), and MQL4GRec (Zhai et al. 2025).

**Metrics.** To assess recommendation effectiveness, we adopt top-k hit rate (HR@K) and normalized Discounted Cumulative Gain (NDCG@K), where K is set to 1, 5, and 10. Consistent with prior studies (Geng et al. 2022; Hua et al. 2023), we utilize a leave-one-out evaluation protocol. Rather than sampling, we conduct full ranking assessments across the entire item collection.

**Implementation Details.** For the multimodal generative baseline model MQL4GRec, we did not utilize pre-training on millions of additional-category datasets to ensure a fair comparison. Text and image features are obtained using LLaMA and ViT-L/14, respectively. For RQ-VAE, the codebook size  $M$  is 256 with 4 levels. We adopt the AdamW optimizer (batch size 1024, learning rate 0.001). The number of clusters  $K$  is 512. Following Rajput et al. (2023); Zhai et al. (2025), we use T5 as the backbone, whose encoder and decoder each have 4 transformer layers with 6 attention heads (dimension 64). The layer-wise contrastive weight  $\lambda_{con}^l$  is applied from the third layer onward:  $\lambda_{con}^{0,1} = 0$ ,  $\lambda_{con}^{2,3} = 0.1$ . We set  $\lambda_{align} = 0.001$ ,  $\lambda_{implicit} = 0.01$ , and temperature  $\tau = 0.1$ . Results are averaged over five random seeds.

### Performance Analysis (RQ1)

Table 2 presents the experimental results of MACRec on three datasets. From the table, we can draw the following conclusions: (1) MACRec achieves the best performance across all three datasets, demonstrating the effectiveness of our proposed approach; (2) MACRec significantly outperforms the state-of-the-art multimodal generative recommen-

dation model MQL4GRec, indicating that the constructed semantic IDs and the cross-modal alignment training strategy can effectively enhance the recommendation performance; (3) Compared with traditional multimodal sequential recommendation models, our model achieves a remarkable improvement in NDCG, suggesting that our multimodal generative recommendation framework can more accurately recommend items that users are interested in.

Model	Instruments	Arts	Games
MACRec	<b>0.1363</b>	<b>0.1329</b>	<b>0.1078</b>
w/o $\mathcal{L}_{con}^l$	0.1289	0.1283	0.1018
w/o $\mathcal{L}_{align}$	0.1310	0.1301	0.1026
w/o $\mathcal{L}_{implicit}$	0.1312	0.1296	0.1042
w/o Explicit Alignment	0.1296	0.1299	0.1037

Table 3: Ablation study (HR@10) on three datasets.

### Ablation Study (RQ2)

Table 3 presents the ablation study results for different loss alignment strategies. From the table, we observe the following: (1) Removing any of the proposed modules leads to a certain degree of performance degradation in recommendation, which demonstrates the effectiveness of our proposed components. (2) Excluding  $\mathcal{L}_{con}$  results in the largest performance drop across all three datasets, highlighting the effectiveness of our contrastive learning-based cross-modal quantization approach. (3) Both implicit and explicit alignment play important roles in the training of generative models.

### Item Collision Analysis (RQ3)

Table 4 compares the item collision rates during the quantization process between MACRec and MQL4GRec. The results show that our model can effectively reduce the item

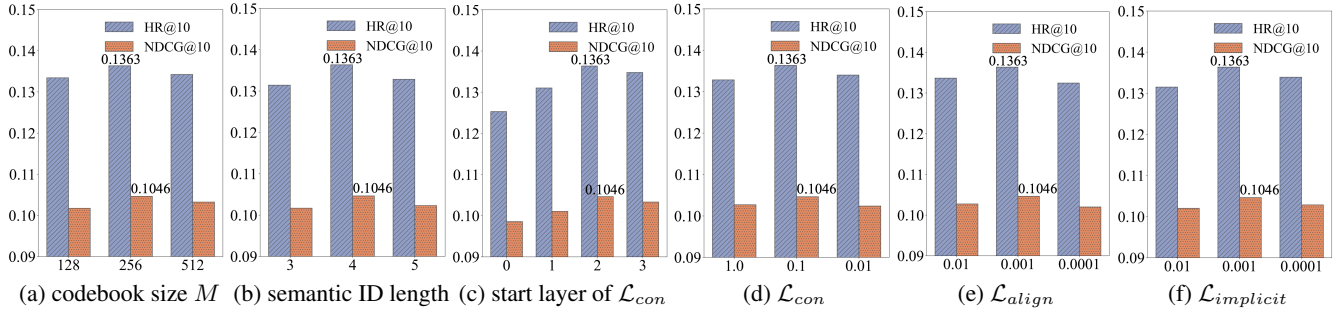


Figure 3: Performance of MACRec over different hyper-parameters on Instruments.




Figure 4: Code assignment distribution on the 2-th RQ layer.

collision rate in the quantization process by fully leveraging the complementarity between different modalities. The decrease in collision rate further demonstrates that our method enables a more balanced distribution of items in the codebook, thereby reducing the probability that similar items are encoded with the same semantic ID.

Dataset	Text		Image	
	MQL4GRec	MACRec	MQL4GRec	MACRec
Instruments	2.76	<b>2.38</b>	3.71	<b>3.23</b>
Arts	5.15	<b>4.24</b>	5.71	<b>5.29</b>
Games	3.51	<b>2.91</b>	26.10	<b>25.24</b>

Table 4: Item ID Collision Rate (%) comparison between MQL4GRec and MACRec on three datasets.

### Code Assignment Distribution (RQ4)

Figure 4 clearly shows the number of items assigned to each codeword in the second codebook layer for both MACRec and MQL4GRec. The red bars represent the number of items assigned to text semantic IDs, while the blue bars correspond to visual semantic IDs. For clearer visualization, the codewords are sorted in descending order by the number of assigned items, and every 16 codewords are grouped together in a single bar. From the figure, we can observe that MACRec distributes items more evenly across the codewords, indicating a better utilization of codebook capacity

and superior semantic representation.

### Parameter Analysis (RQ5)

We analyze the effects of key hyperparameters as follows: 1) For codebook size, both very small and very large sizes degrade performance. Small codebooks limit the quantization space and semantic associations, while large ones dilute token exposure, hindering robust representation learning. 2) For semantic ID length, very short IDs fail to capture comprehensive semantics, and overly long IDs complicate learning by expanding the generation space, both reducing performance. 3) When choosing the starting layer for cross-modal contrastive loss  $\mathcal{L}_{con}^l$ , performance rises then falls; applying it from the first layer erases modality-specific information, while starting from the third layer achieves the best results by letting later VQ layers leverage cross-modal signals to compensate for semantic loss. 4) Each of the other three contrastive losses has an optimal weight: higher weights strengthen modality fusion, while lower weights lead to insufficient cross-modal interaction.

### Conclusion

In this paper, we address the insufficient cross-modal alignment and interaction in current GR methods during semantic ID learning and generative model training. We propose a novel model, **Multi-Aspect Cross-modal quantization for generative Recommendation (MACRec)**. MACRec introduces cross-modal interactions at two stages: semantic ID training and generative model training. Specifically, we incorporate cross-modal quantization based on contrastive learning to facilitate the construction of semantic IDs that are both hierarchically semantic and independent. Moreover, we incorporate both implicit and explicit cross-modal alignment in the training process of the generative model, further enhancing the model’s understanding of sequential information across different modalities. Extensive experiments demonstrate the superior performance of MACRec in recommendation systems. Finally, we further conduct additional analysis to demonstrate the advantage of our method in codebook utilization.

## Acknowledgments

This work was supported by the National Key Research and Development Program of China under Grant No. 2024YFF0729003, the National Natural Science Foundation of China under Grant Nos. 62176014, 62206266, the Fundamental Research Funds for the Central Universities.

## References

- Bin, X.; Cui, J.; Yan, W.; Zhao, Z.; Han, X.; Yan, C.; Zhang, F.; Zhou, X.; Wu, Q.; and Liu, Z. 2025. Real-time Indexing for Large-scale Recommendation by Streaming Vector Quantization Retriever. *arXiv preprint arXiv:2501.08695*.
- Chen, Q.; Zhao, H.; Li, W.; Huang, P.; and Ou, W. 2019. Behavior sequence transformer for e-commerce recommendation in alibaba. In *Proceedings of the 1st international workshop on deep learning practice for high-dimensional sparse data*, 1–4.
- Chen, W.; Wu, Y.; Zhang, Z.; Zhuang, F.; He, Z.; Xie, R.; and Xia, F. 2024. FairGap: Fairness-aware recommendation via generating counterfactual graph. *ACM Transactions on Information Systems*, 42(4): 1–25.
- Chen, W.; Yuan, M.; Zhang, Z.; Xie, R.; Zhuang, F.; Wang, D.; and Liu, R. 2025. FairDgcl: Fairness-aware recommendation with dynamic graph contrastive learning. *IEEE Transactions on Knowledge and Data Engineering*.
- Covington, P.; Adams, J.; and Sargin, E. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, 191–198.
- Cui, Z.; Ma, J.; Zhou, C.; Zhou, J.; and Yang, H. 2022. M6-rec: Generative pretrained language models are open-ended recommender systems. *arXiv preprint arXiv:2205.08084*.
- Davidson, J.; Liebal, B.; Liu, J.; Nandy, P.; Van Vleet, T.; Gargi, U.; Gupta, S.; He, Y.; Lambert, M.; Livingston, B.; et al. 2010. The YouTube video recommendation system. In *Proceedings of the fourth ACM conference on Recommender systems*, 293–296.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Du, Y.; Luo, D.; Yan, R.; Wang, X.; Liu, H.; Zhu, H.; Song, Y.; and Zhang, J. 2024. Enhancing job recommendation through llm-based generative adversarial networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 8363–8371.
- Geng, S.; Liu, S.; Fu, Z.; Ge, Y.; and Zhang, Y. 2022. Recommendation as language processing (rlp): A unified pre-train, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM conference on recommender systems*, 299–315.
- Geng, S.; Tan, J.; Liu, S.; Fu, Z.; and Zhang, Y. 2023. VIP5: Towards Multimodal Foundation Models for Recommendation. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Hidasi, B.; Karatzoglou, A.; Baltrunas, L.; and Tikk, D. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939*.
- Hua, W.; Xu, S.; Ge, Y.; and Zhang, Y. 2023. How to index item ids for recommendation foundation models. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, 195–204.
- Kang, W.-C.; and McAuley, J. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*, 197–206. IEEE.
- Lee, D.; Kim, C.; Kim, S.; Cho, M.; and Han, W.-S. 2022. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11523–11532.
- Li, J.; Ren, P.; Chen, Z.; Ren, Z.; Lian, T.; and Ma, J. 2017. Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 1419–1428.
- Li, L.; Zhang, Y.; Liu, D.; and Chen, L. 2024. Large Language Models for Generative Recommendation: A Survey and Visionary Discussions. In *Joint 30th International Conference on Computational Linguistics and 14th International Conference on Language Resources and Evaluation, LREC-COLING 2024*, 10146–10159. European Language Resources Association (ELRA).
- Liu, H.; Wei, Y.; Song, X.; Guan, W.; Li, Y.-F.; and Nie, L. 2024a. Mmgrec: Multimodal generative recommendation with transformer model. *arXiv preprint arXiv:2404.16555*.
- Liu, Q.; Hu, J.; Xiao, Y.; Zhao, X.; Gao, J.; Wang, W.; Li, Q.; and Tang, J. 2024b. Multimodal recommender systems: A survey. *ACM Computing Surveys*, 57(2): 1–17.
- Liu, Q.; Zeng, Y.; Mokhosi, R.; and Zhang, H. 2018. STAMP: short-term attention/memory priority model for session-based recommendation. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1831–1839.
- Liu, X.; Zhang, F.; Wu, Y.; Jia, X.; Xia, Z.; Zhuang, F.; Zhang, Z.; Jiang, F.; and Lin, W. 2025. CAT-ID<sup>2</sup>: Category-Tree Integrated Document Identifier Learning for Generative Retrieval In E-commerce. *arXiv preprint arXiv:2511.01461*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Rajput, S.; Mehta, N.; Singh, A.; Hulikal Keshavan, R.; Vu, T.; Heldt, L.; Hong, L.; Tay, Y.; Tran, V.; Samost, J.; et al. 2023. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems*, 36: 10299–10315.
- Smith, B.; and Linden, G. 2017. Two decades of recommender systems at Amazon. com. *Ieee internet computing*, 21(3): 12–18.
- Sun, F.; Liu, J.; Wu, J.; Pei, C.; Lin, X.; Ou, W.; and Jiang, P. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*, 1441–1450.

- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, J.; Zeng, Z.; Wang, Y.; Wang, Y.; Lu, X.; Li, T.; Yuan, J.; Zhang, R.; Zheng, H.-T.; and Xia, S.-T. 2023. Missrec: Pre-training and transferring multi-modal interest-aware sequence representation for recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia*, 6548–6557.
- Wang, W.; Bao, H.; Lin, X.; Zhang, J.; Li, Y.; Feng, F.; Ng, S.-K.; and Chua, T.-S. 2024. Learnable item tokenization for generative recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2400–2409.
- Wei, Y.; Wang, X.; Nie, L.; He, X.; and Chua, T.-S. 2020. Graph-refined convolutional network for multimedia recommendation with implicit feedback. In *Proceedings of the 28th ACM international conference on multimedia*, 3541–3549.
- Wei, Y.; Wang, X.; Nie, L.; He, X.; Hong, R.; and Chua, T.-S. 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.
- Xi, D.; Chen, Z.; Yan, P.; Zhang, Y.; Zhu, Y.; Zhuang, F.; and Chen, Y. 2021. Modeling the sequential dependence among audience multi-step conversions with multi-task learning in targeted display advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 3745–3755.
- Xi, D.; Zhuang, F.; Liu, Y.; Gu, J.; Xiong, H.; and He, Q. 2019. Modelling of bi-directional spatio-temporal dependence and users’ dynamic preferences for missing poi check-in identification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 5458–5465.
- Xi, D.; Zhuang, F.; Zhu, Y.; Zhao, P.; Zhang, X.; and He, Q. 2020. Graph factorization machines for cross-domain recommendation. *arXiv preprint arXiv:2007.05911*.
- Xiaoyu, L.; Wu, Y.; Han, R.; Zhuang, F.; Li, X.; and Lin, W. 2025. A Soft-partitioned Semi-supervised Collaborative Transfer Learning Approach for Multi-Domain Recommendation. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, 5366–5370. Association for Computing Machinery. ISBN 9798400720406.
- Yi, Z.; Wang, X.; Ounis, I.; and Macdonald, C. 2022. Multi-modal graph contrastive learning for micro-video recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 1807–1811.
- Zhai, J.; Mai, Z.-F.; Wang, C.-D.; Yang, F.; Zheng, X.; Li, H.; and Tian, Y. 2025. Multimodal Quantitative Language for Generative Recommendation. In *The Thirteenth International Conference on Learning Representations*.
- Zhang, F.; Liu, X.; Jia, X.; Zhang, Y.; Xia, Z.; Jiang, F.; Zhuang, F.; Lin, W.; and Zhang, Z. 2025a. HierGR: Hierarchical Semantic Representation Enhancement for Generative Retrieval in Food Delivery Search. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, 444–455.
- Zhang, F.; Liu, X.; Jia, X.; Zhang, Y.; Zhang, S.; Li, X.; Zhuang, F.; Lin, W.; and Zhang, Z. 2025b. Multi-level Relevance Document Identifier Learning for Generative Retrieval. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 10066–10080.
- Zhang, F.; Zhang, Z.; Ao, X.; Gao, D.; Zhuang, F.; Wei, Y.; and He, Q. 2022a. Mind the gap: Cross-lingual information retrieval with hierarchical knowledge enhancement. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 4345–4353.
- Zhang, F.; Zhang, Z.; Zhuang, F.; Zhao, Y.; Wang, D.; and Zheng, H. 2024. Temporal knowledge graph reasoning with dynamic memory enhancement. *IEEE Transactions on Knowledge and Data Engineering*, 36(11): 7115–7128.
- Zhang, J.; Zhu, Y.; Liu, Q.; Zhang, M.; Wu, S.; and Wang, L. 2022b. Latent structure mining with contrastive modality fusion for multimedia recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 35(9): 9154–9167.
- Zhang, T.; Zhao, P.; Liu, Y.; Sheng, V. S.; Xu, J.; Wang, D.; Liu, G.; Zhou, X.; et al. 2019. Feature-level deeper self-attention network for sequential recommendation. In *IJCAI*, 4320–4326.
- Zheng, B.; Hou, Y.; Lu, H.; Chen, Y.; Zhao, W. X.; Chen, M.; and Wen, J.-R. 2024. Adapting large language models by integrating collaborative semantics for recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 1435–1448. IEEE.
- Zhou, G.; Deng, J.; Zhang, J.; Cai, K.; Ren, L.; Luo, Q.; Wang, Q.; Hu, Q.; Huang, R.; Wang, S.; et al. 2025a. OneRec Technical Report. *arXiv preprint arXiv:2506.13695*.
- Zhou, G.; Hu, H.; Cheng, H.; Wang, H.; Deng, J.; Zhang, J.; Cai, K.; Ren, L.; Ren, L.; Yu, L.; et al. 2025b. Onerec-v2 technical report. *arXiv preprint arXiv:2508.20900*.
- Zhou, K.; Wang, H.; Zhao, W. X.; Zhu, Y.; Wang, S.; Zhang, F.; Wang, Z.; and Wen, J.-R. 2020. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *Proceedings of the 29th ACM international conference on information & knowledge management*, 1893–1902.
- Zhu, Y.; Chen, J.; Chen, L.; Li, Y.; Zhang, F.; and Liu, Z. 2024. Interest clock: Time perception in real-time streaming recommendation system. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2915–2919.
- Zhu, Y.; Jiang, G.; Chen, J.; Zhang, F.; Wu, Q.; and Liu, Z. 2025. Long-Term Interest Clock: Fine-Grained Time Perception in Streaming Recommendation System. In *Companion Proceedings of the ACM on Web Conference 2025*, 1554–1557.