

Cross-Modal Content Inference and Feature Enrichment for Cold-Start Recommendation

Haokai Ma¹, Zhuang Qi¹, Xinxin Dong¹, Xiangxian Li¹, Yuze Zheng¹, Xiangxu Meng¹ and Lei Meng^{*1,2}

¹ Shandong University, Jinan, China

² Shandong Research Institute of Industrial Technology, Jinan, China

Email: {mahaokai, z_qi, dongxinxin, xiangxian_lee, zhengyuze}@mail.sdu.edu.cn, {mxx, lmeng}@sdu.edu.cn

Abstract—Multimedia recommendation aims to fuse the multi-modal information of items for feature enrichment to improve the recommendation performance. However, existing methods typically introduce multi-modal information based on collaborative information to improve the overall recommendation precision, while failing to explore its cold-start recommendation performance. Meanwhile, these above methods are only applicable when such multi-modal data is available. To address this problem, this paper proposes a recommendation framework, named **Cross-modal Content Inference and Feature Enrichment Recommendation (CIERec)**, which exploits the multi-modal information to improve its cold-start recommendation performance. Specifically, CIERec first introduces image annotation as the privileged information to help guide the mapping of unified features from the visual space to the semantic space in the training phase. And then CIERec enriches the content representation with the fusion of collaborative, visual, and cross-modal inferred representations, so as to improve its cold-start recommendation performance. Experimental results on two real-world datasets show that the content representations learned by CIERec are able to achieve superior cold-start recommendation performance over existing visually-aware recommendation algorithms. More importantly, CIERec can consistently achieve significant improvements with different conventional visually-aware backbones, which verifies its universality and effectiveness.

Index Terms—Cross-Modal, Content Inference, Feature Enrichment, Cold-Start Recommendation

I. INTRODUCTION

Personalized recommendation aims to capture the users' preference and provide them with the appropriate items [1]–[4]. However, the cold-start problem is a ubiquitous challenge in recommendation, which leads to bias in the training of traditional collaborative filtering recommenders and visually-aware recommenders (e.g., Matrix Factorization (MF) [5] and Visual Bayesian Personalized Ranking (VBPR) [6]). That is, the randomization problem and the popularity bias problem. Content-based recommendation, which leverage the multi-modal information to improve its cold-start performance, has attracted considerable attention [7], [8]. However, the data in web and mobile applications is diverse and unstable [9], [10], and the performance of existing algorithms is generally limited by the learning of heterogeneous multi-modal representations, which is not always available. Therefore, robust cross-modal inference methods for heterogeneous modal cold-start recommendation are urgently needed.

* indicates corresponding author.




Fig. 1. The main framework of CIERec. CIERec designs a novel cross-modal inference strategy to infer the cross-modal semantic content representations by introducing the privileged information into the existing visually-aware cold-start recommendation task.

Existing cold-start recommendation methods can be categorized into two classes, based on the targets they aim at. One class is the information-based methods, which generally incorporate heterogeneous content features from the user's or item's auxiliary information [11]–[13] into the cold-start recommendation. Another class is representation-based methods, which are able to learn the fine-grained representations by dynamically optimizing the recommendation models [8], [14]–[16], thus improving the cold-start recommendation performance. These features are mainly obtained by encoding the items' multi-modal content information, and the accuracy and stability of these algorithms are decreased when these heterogeneous information is not available.

To address these issues, we present **CIERec**, a novel **C**ross-modal **C**ontent **I**nference and **F**eature **E**nrichment **C**old-start **R**ecommendation framework. It introduces image annotation as the privileged information to guide the mapping progress of the content representation from the visual space to the semantic space and enriches it by fusing the collaborative, visual, and inferred semantic features to improve its cold-start performance. Figure 1 shows the main framework of CIERec. Specifically, we first model the collaborative interactions in the collaborative representation learning (CRL) module. We then extract the item's uniform embedding with a traditional visual encoder (e.g. ResNet18) in the source-modal representation learning (SMRL) module. Next, to deal with the difficulty of mapping heterogeneous features, we propose a novel cross-modal inference strategy to map content representations from the visual space to the semantic space with the guidance of prior knowledge at the cross-modal representation learning (CMRL) module, which design follows a learning paradigm called learning using privileged information (LUPI) [17].

Finally, a multi-modal fusion method is used to integrate the user embeddings, the item embeddings, the visual features, and the inferred semantic features for the final recommendation in the multi-modal representation fusion (MRF) module. As observed, CIERec is able to alleviate the absence of heterogeneous modalities in cold-start recommendations through cross-modal inference, thereby improving the stability and accuracy of existing visually-aware cold-start recommendation models.

To validate the effectiveness of the proposed CIERec, we conduct the performance comparison on the pre-processed Allrecipes [18] and Amazon_CDs [6] datasets against previous advanced visually-aware recommendation algorithms. We also conduct extensive experiments, including the ablation study to verify the effectiveness of the different components in the proposed CIERec, and the case study to visualize the distributions of heterogeneous modal representations. The experimental results demonstrate that the advantages of Tri-CDR are: (1) the privileged information can help to model the mapping relationship of content information in the visual and semantic space. (2) The dual-gating module in the CMRL module can optimize the cross-modal inference process of heterogeneous representations. (3) CIERec can further improve its cold-start performance with the fusion of multi-modal representations, allowing the downstream task to focus on the user’s preference information from multiple perspectives. Furthermore, CIERec achieves consistent and significant improvements over the existing visually-aware recommendation methods with different benchmarks (e.g., MF [5] and VBPR [6]), proving its effectiveness and universality. Overall, the main contributions are summarized as follows:

- This paper proposes a novel Cross-modal Content Inference and Feature Enrichment Recommendation framework, CIERec, which can improve the cold-start performance of existing visually-aware recommendation methods through cross-modal semantic inference and multi-modal representation fusion.
- We design a content-enriched cross-modal inference strategy to model the heterogeneous representation inference process and extract the fine-grained multi-modal representations based on the leverage of privilege information. It can be applied as a cross-modal inference module for the general task of multi-modal representation learning.
- We conduct extensive experiments to verify the effectiveness of the proposed CIERec on multiple datasets with different visually-aware recommendation models. The experimental results demonstrate CIERec’s effectiveness, universality, and stability.

II. RELATED WORK

A. Multi-modal Learning in Recommendation

As a common multimedia analysis method, multi-modal learning has been widely used in the fields of computer vision [19]–[23], data mining [24]–[27], information retrieval [28], [29], and recommendation [12], [30], [31]. In recommendation, multimedia recommendation aims to incorporate the items’

content information to model the fine-grained representation and thereby improve its performance. These methods can be classified as multi-source information embedding methods [12], [32], [33] and heterogeneous information inference methods [30], [34], according to the available information. The multi-source information embedding methods refer to the introduction of multiple sources of heterogeneous information from the heterogeneous information networks [35] and the knowledge graph [33] as the content information to complement the collaborative information, which can decrease the dependence of the recommendation model on the interaction information. Heterogeneous information inference methods aim to learn the mapping function from one modality space to another modality space, thus making the content representation extracted from the same item has similar distributions. These above methods typically perform cross-modal inference based on the user’s interactive information [31] or multi-modal information [30], [34] of items, thus reducing the gap between the user’s interest manifold and the visual semantic manifold. However, existing cross-modal content inference methods are rarely demonstrated against their cold-start recommendation performance, and lack the in-depth analysis of the effectiveness of the source-modal information and the inferred information for recommendation.

B. Visually-aware Recommendation

With the development of image analysis techniques in the field of CV [36]–[40] and NLP [41]–[43], a series of works have validated that the incorporation of visual features of items can improve recommendation performance [6], [7], [18], [30]. Inspired by the significant performance of Convolution Neural Network (CNN) in image classification, most visually-aware recommendation methods incorporate pre-extracted features as the item’s embedding into the recommendation [6], [44], [45]. However, previous visually-aware recommendation methods with pre-extracted features typically utilize the items’ visual features related to their category information while ignoring the users’ personalized interests. Therefore, recent visually-aware recommendation methods extract visual features with an end-to-end approach to jointly optimize image encoders and recommenders. Deepstyle [46] replaces collaborative representations with the real-time extraction of the items’ visual features to capture the user’s multi-dimensional preferences. PiNet [30] learns the visual features that contain both visual, semantic, and collaborative information by jointly optimizing the image encoder. However, these algorithms only take into account the visual information of the items, leading to the suboptimal performance.

C. Cold-Start Recommendation

The cold start problem is common in recommendation due to the imbalance of interactions, where a few users or items dominate the interactions in the dataset. It may lead to the ‘popular bias’ [47], whereby active users perform better than cold users and popular items are more likely to be recommended than cold items. Existing cold-start recommendation methods can

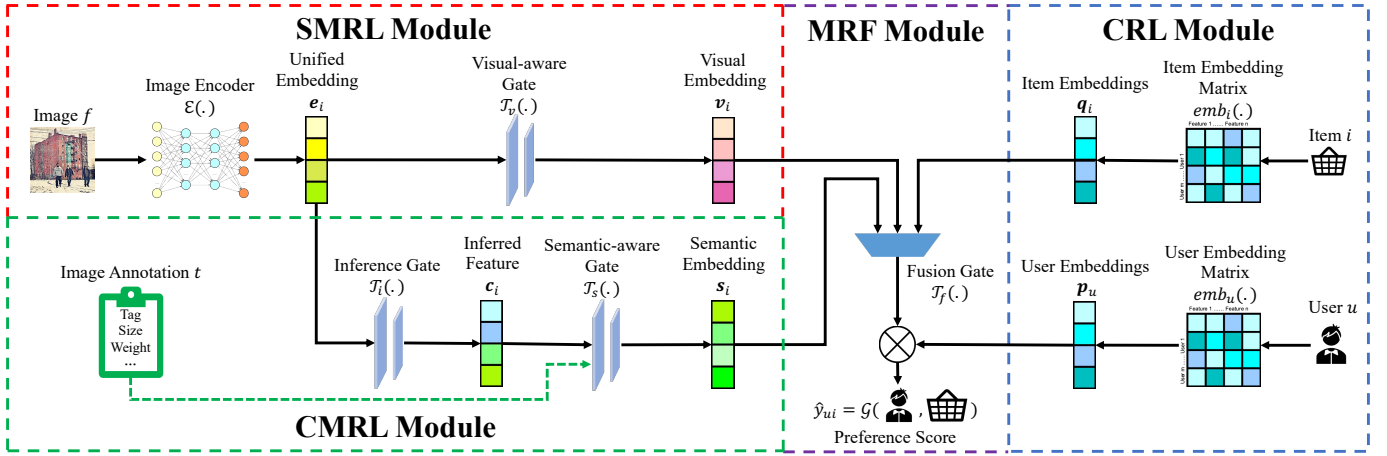


Fig. 2. The overall structure of CIERec. The CRL module is used to learn the user’s collaborative representation p_u and the item’s collaborative representation q_i ; the SMRL module can be used to extract the visual feature v_i ; the CMRL module learns the cross-modal inferred feature c_i and then generate the semantic feature s_i from the unified feature e_i ; the MRF module is able to fuse the aforementioned representations and calculate the final preference score.

be categorized into two categories by their targets, namely information-based methods that introduce auxiliary data (e.g., heterogeneous information [12] and social information [11]) into recommendation, and the representation-based methods that modify recommendation models dynamically through meta-learning [14], transfer learning [16], contrastive learning [8] and other techniques. However, these algorithms have difficulty in achieving idealized recommendations performance when heterogeneous information is not available for certain modalities. To address this problem, existing studies usually deal with interaction information specifically. DropoutNet [48] adopts the dropout mechanism for the observed interactions to improve the model generalization capability. A series of studies [45], [49] attempt to leverage the generalization capabilities of adversarial learning, targeting at augmenting the unobserved interactions.

III. TECHNIQUE

A. Framework Overview

CIERec introduces a cross-modal inference and feature enrichment framework to enable the enrichment of content representations through feature-level cross-modal inference. As shown in Figure 2, CIERec can be divided into four main modules, including the Collaborative Representation Learning (CRL) module, the Source-Modal Representation Learning (SMRL) module, the Cross-Modal Representation Learning (CMRL) module, and the Multi-modal Representation Fusion (MRF) module, as illustrated in the following sections.

B. Collaborative Representation Learning (CRL)

As shown in Figure 2, CRL learns the collaborative representation p_u and the item collaborative representation q_i from the embedding matrix based on the randomly sampled user u and item i , which is similar to the traditional collaborative filtering method. The learning process can be expressed as:

$$\begin{aligned} p_u &= \text{emb}_u(u) \\ q_i &= \text{emb}_i(i) \end{aligned} \quad (1)$$

where $\text{emb}_u(\cdot)$ denotes the embedding matrix of users, and $\text{emb}_i(\cdot)$ denotes the embedding matrix of items.

C. Source-Modal Representation Learning (SMRL)

CIERec generates visual feature v_i in the SMRL module to complement the collaborative representation q_i . Specifically, as illustrated in Figure 2, the SMRL module extracts the uniform embedding $\mathcal{E}(f) \mapsto e_i$ from the image f through the visual encoder $\mathcal{E}(\cdot)$. Inspired by the dual-gating mechanism [30], we develop a novel multimodal-gating function to generate the visual feature $\mathcal{T}_v(e_i) \mapsto v_i$ by mapping with the visual-aware gate \mathcal{T}_v , and constrains its optimization with the gradient-regularization gate \mathcal{R} . The overall computational equation can be expressed as:

$$v_i = \mathcal{T}_v(\mathcal{E}(f)) \quad (2)$$

There is significant heterogeneity between the image regions focused by visual feature v_i and semantic feature s_i , which is difficult to be directly mapped from the uniform embedding e_i . Therefore, CIERec proposes a visual-aware gate to control the delivery of visual information. The visual-aware gate $\mathcal{T}_v(\cdot)$ of CIERec introduces a self-learning gate vector g_v and the user representation p_u in addition to the uniform embedding e_i , and maps them to the visual feature space, which is defined as follows:

$$v_i = \mathcal{T}_v(e_i) = \text{MLP}^v \left(\frac{e_i \odot \delta(p_u \| e_i \| g_v)}{\max(\|e_i \odot \delta(p_u \| e_i \| g_v)\|_2, \epsilon)} \right) \quad (3)$$

where $\|$ denotes the concatenate operator, \odot denotes the dot product operator, $\|\cdot\|_2$ denotes the l_2 regularization function, $\epsilon = 1e^{-12}$ denotes a small value to avoid division by zero, $\delta(\cdot)$ denotes a fully connected layer and $\text{MLP}^v(\cdot)$ denotes a fully connected layer followed by the LeakyReLU activation function. We further describe the gradient-regularization gate \mathcal{R} in Sec. III-D3.

D. Cross-modal Representation Learning (CMRL)

In addition to generating the visual feature v_i , we firstly performs cross-modal semantic inference based on the inference gate \mathcal{T}_i to learn the inference feature c_i from the unified embedding e_i ; And then we conducts semantic fusion based on the prior knowledge, that is, fuses the semantic knowledge s of image annotation t with the cross-modal inference feature c_i to generate semantic feature s_i via the semantic-aware gate \mathcal{T}_s , which in turn complements the multimodal representation.

1) *Semantic Inference*: In preliminary experiments, we found that the incorporation of semantic information can significantly improve the recommendation performance of traditional visually-aware recommendation methods. Existing multimedia recommendation methods generally rely on the modal richness of the recommendation datasets, which performs poorly when semantic information is unavailable. To address this issue, the CMRL module generates the cross-modal inference feature $\mathcal{T}_i(e_i) \mapsto c_i$ by mapping from the unified embedding e_i with the inference gate $\mathcal{T}_i(\cdot)$ and uniformly optimizes it with the gradient-regularization gate \mathcal{R} . It is defined as follows:

$$c_i = \frac{e_i \odot \delta(e_i \| g_s)}{\max(\|e_i \odot \delta(e_i \| g_s)\|_2, \epsilon)} \quad (4)$$

where g_s denotes the learnable inference gate vector, $\|$ denotes the concatenate operator, $\|\cdot\|_2$ denotes the ℓ_2 regularization function, $\epsilon = 1e^{-12}$ denotes a small value to avoid division by zero and c_i denotes the inferred feature.

2) *Semantic fusion for privileged information*: The inferred feature c_i derived from the uniform embedding e_i contains a certain amount of redundant information, leading to bias in the subsequent recommendation process. To address this issue, the CMRL module also introduces the prior semantic information to improve the model's ability to represent semantic modalities. The CMRL module uses LSTM [50] as the building block for the semantic element encoder to capture the relationship between the prior semantic representation and the inferred representation. The process of extracting the semantic embedding $\mathcal{T}_s(c_i) \mapsto s_i$ can be represented as:

$$s_i = \text{LSTM}(c_i \| \hat{s}_i) \quad (5)$$

where \hat{s}_i denotes the privileged knowledge encoded from the image annotation, $\text{LSTM}(\cdot)$ denotes the the LSTM operator, s_i denotes the calculated semantic embedding.

3) *Gradient-regularization Gate*: CIERec fuses the gradients of the two heterogeneous representations with the gradient-regularization gate \mathcal{R} , aiming to allow the visual encoder to trade off the visual feature v_i and the inferred semantic feature s_i from the unified embedding e_i during the backpropagation. The visual gradient-regularization gate is implemented based on the deep Q-network (DQN) [51] approach. It selects the action $s^{(t)}$ via a classifier MLP_{DQN}^v to map the 5D state vector of $s^{(t-1)}$ at batch t , where MLP_{DQN}^v denotes a fully-connected layer followed by a Softmax activation function. And then it penalizes this action by the

TABLE I
STATISTICS OF THE EXPERIMENTED DATASETS.

Datasets	Users	Items	Interactions	Elements	Sparsity
Allrecipes	68,768	45,630	1,093,845	2,736	99.97%
Amazon_CDs	67,282	40,314	752,724	467	99.97%

$J_v^{(t)} = \exp(-\mathcal{L}_v^{(t+1)})$ obtained by the feedback from the recommendation model, with the loss function as following:

$$\mathcal{L}_{dq_n} = -\log \sigma \left(s_{\max}^{(t)} \cdot J_v^{(t)} \right) \quad (6)$$

where $s_{\max}^{(t)}$ is the probability of selecting $s^{(t)}$, and $\sigma(\cdot)$ is the Sigmoid Function. Similarly, we have $J_s^{(t)} = \exp(-\mathcal{L}_s^{(t+1)})$ for the semantic gradient-regularization gate.

E. Multimodal Representation Fusion (MRF) Module

To demonstrate the universality of CIERec, we take the traditional collaborative filtering algorithms MF and VBPR as the backbone models for this module, which represent user and product as an embedding vector, with the core idea of estimating the user's preference as the inner product of their embedding vectors [5]. In addition to the user representation p_u and the collaborative representation q_i , MRF module also receives the visual embedding v_i and the semantic embedding s_i for recommendation. The fusion operation of multi-modal representations can be represented as:

$$f_i = \mathcal{T}_f(q_i, v_i, s_i) = \text{MLP}(q_i \| v_i \| s_i) \quad (7)$$

where \mathcal{T}_f denotes the fusion gate of the multi-modal representations, $\text{MLP}(\cdot)$ denotes a fully connected layer with a LeakyReLU activation function, and f_i denotes the multi-modal fused representation of item i . The process of calculating preference scores for the BPR-MF and VBPR algorithms is defined as following:

$$\begin{aligned} \hat{y}_{MF} &= \alpha + \beta_u + \beta_i + p_u^\top f_i \\ \hat{y}_{VBPR} &= \alpha + \beta_u + \beta_i + \beta_c + p_u^\top f_i + a_u^\top c_i \end{aligned} \quad (8)$$

where α denotes the self-learning global bias, β_u , β_i and β_c denote the self-learning bias of user u , item i and the content c respectively, a_u denotes the implicit representation, \hat{y}_{MF} and \hat{y}_{VBPR} denote the preference score of MF and VBPR.

IV. EXPERIMENTS

A. Experimental Setup

1) *Datasets*: We conduct experiments on two real-world recommendation datasets, where Allrecipes was constructed by Gao [18] and Amazon_CDs was extracted from the original Amazon dataset [6] to meet the needs for this task. To verify the effectiveness of CIERec in alleviating the cold-start problems in recommendation, we divided Allrecipes and Amazon_CDs into two parts with the interaction boundary of 3, i.e., the cold-start set with few interactions and the warm-start set with a relatively large number of interactions. Table I summarizes the statistics of the datasets. Both datasets follow

TABLE II
PERFORMANCE COMPARISON OF CIEREC WITH EXISTING BASELINE ALGORITHMS ON AMAZON_CDs AND ALLRECIPES DATASETS.

Methods	Algorithms	Amazon_CDs Dataset						Allrecipes Dataset					
		Cold		Warm		All		Cold		Warm		All	
		R@10	NDCG@10	R@10	NDCG@10	R@10	NDCG@10	R@10	NDCG@10	R@10	NDCG@10	R@10	NDCG@10
Backbone	MF	0.1719	0.1483	0.1763	0.3678	0.1734	0.2233	0.2427	0.1936	0.1962	0.4555	0.2275	0.2793
	VBPR	0.1724	0.1533	0.1933	0.4087	0.1796	0.2411	0.2539	0.2050	0.2204	0.4948	0.2429	0.2999
Multi-modal Learning	MF(Image)	0.1634	0.1459	0.1930	0.3994	0.1736	0.2331	0.2450	0.1975	0.2085	0.4782	0.2330	0.2894
	MF(Semantics)	0.1843	0.1589	0.2107	0.4086	0.1933	0.2447	0.2376	0.1886	0.2103	0.4683	0.2286	0.2801
	VECF	0.1796	0.1555	0.1989	0.4062	0.1871	0.2426	0.2608	0.2062	0.2284	0.4977	0.2502	0.3016
	HAFR-non-i	0.1872	0.1619	0.2052	0.4174	0.1934	0.2498	0.2600	0.2095	0.2290	0.5064	0.2499	0.3067
	PiNet	0.2100	0.1833	0.2240	0.4367	0.2148	0.2705	0.2770	0.2209	0.2377	0.5106	0.2641	0.3158
Cold-Start Learning	DropoutNet	0.1744	0.1507	0.1876	0.3790	0.1789	0.2292	0.2503	0.2002	0.2043	0.4619	0.2353	0.2859
	AMF	0.1717	0.1484	0.1818	0.3750	0.1752	0.2263	0.2494	0.2006	0.2069	0.4734	0.2355	0.2899
	AMR	0.1826	0.1620	0.2000	0.4160	0.1886	0.2494	0.2619	0.2106	0.2238	0.4977	0.2494	0.3046
	CLCRec	0.1963	0.1723	0.2087	0.4391	0.2005	0.2641	0.2571	0.2056	0.2327	0.5012	0.2491	0.3024
Cross-modal Learning	CIERec(MF)	<u>0.2257</u>	<u>0.1981</u>	<u>0.2461</u>	<u>0.4706</u>	<u>0.2327</u>	<u>0.2918</u>	<u>0.2789</u>	<u>0.2228</u>	<u>0.2427</u>	<u>0.5165</u>	<u>0.2670</u>	<u>0.3190</u>
	CIERec(VBPR)	0.2376	0.2091	0.2602	0.4957	0.2454	0.3076	0.2832	0.2269	0.2447	0.5209	0.2706	0.3232

the data partitioning method used in Allrecipes, where the train set includes the earliest 60% of the interaction data for each user, the test set includes the latest 30% of the interaction data, and the remaining 10% is used as the valid set.

2) *Evaluation measures*: Following the classical cold-start recommendation works [8], [52], two widely-used metrics are adopted to evaluate the performance of the cold start recommendation, including Recall (R) and Normalized discounted cumulative gain (NDCG) [8]. Following [30], we randomly select one negative sample for each positive sample in training, while in testing five hundred items are randomly selected as the negative samples (have no interaction with the user) from the dataset along with all positive items (have interactions with the user) to form the ranking candidate for each user. R@K and NDCG@K calculate the performance of positive samples in the Top-k ranking items for all sampled items. To alleviate the problem of randomness, we repeat the evaluation process five times and report the average value.

3) *Implementation details*: Based on the efficiency and performance of ResNet18 in recommendation [53] and prediction [54], CIERec used it as the visual encoder to extract the uniform representations with the dimension of 512. The cross-modal recommendation model was optimized by Adagrad with the learning rate from 0.0001 to 0.5, and the DQN model was optimized by Adam with the learning rate from 0.00001 to 0.005. Both the batch size and the dimension of the CIERec were selected in 32, 64, 28, 256, and these optimizers are decayed proportionally for every four epochs, with the decay rate chosen from 0.1, 0.5.

4) *Baselines*: We compare CIERec with both multi-modal learning models and cold-start learning models:

- MF [5] is a classical recommendation model to utilize the implicit feedback information with Bayesian Personalized Ranking.
- VBPR [6] is a factorization model to incorporate pre-extracted visual features into recommendation.

- MF(Image/Semantic) [5] is a variant of MF where we replace the collaborative embedding with visual or semantic feature. For fair comparisons, we use the same pre-extracted visual features and semantic feature for training, noted as MF(Image) and MF(Semantic).
- VECF [55] capture the visual and textual features by a multi-modal attention network seamlessly.
- HAFR-non-i [18] learns the user’s preference via the visual images, ingredients and the collaborative information of the interacted recipes.
- PiNet [30] is a heterogeneous multi-task learning framework that learns visual features containing with semantic and collaborative information, which is also the baseline method for the proposed CIERec.
- DropoutNet [48] learn a DNN-based latent model via the dropout mechanism based on the idea that cold start is equivalent to the missing data problem.
- AMF [49] learns the effective collaborative feature via adding gradient-based perturbations to item embedding.
- AMR [45] adds random-based perturbations and gradient-based perturbations to the pre-extracted visual features to model the effective visual information.
- CLCRec [8] is a SOTA model in cold-start recommendation that optimizes the dependence between items’ embedding and content information via contrastive learning.

B. Performance Comparison

In this section, we compare CIERec with the algorithm mentioned in section IV-A4 for performance comparison. For each algorithm, we have fine-tuned its hyper-parameters to obtain its best performance in the experiments. It can be observed from Table II that:

- MF [5] with content features achieves consistent improvements over MF on most datasets and metrics, whereas replacing the collaborative information only with content information from a single modality can lead to a

TABLE III
RECALL@10 ON ABLATION STUDY OF CIEREC(MF) AND CIEREC(VBPR).

Algorithms	CIERec(MF)						CIERec(VBPR)					
	Amazon_CDs			Allrecipes			Amazon_CDs			Allrecipes		
	Cold	Warm	All	Cold	Warm	All	Cold	Warm	All	Cold	Warm	All
Base	0.1634	0.1930	0.1736	0.2450	0.2085	0.2330	0.1724	0.1933	0.1796	0.2539	0.2204	0.2429
Base+CI	0.1688	0.1766	0.1715	0.2405	0.2071	0.2296	0.1685	0.1850	0.1742	0.2445	0.2130	0.2342
Base+CI+TA	0.1906	0.2043	0.1953	0.2536	0.2177	0.2419	0.1843	0.2135	0.1943	0.2739	0.2368	0.2617
Base+CI+TA+GR	0.2006	0.2145	0.2054	0.2718	0.2325	0.2589	0.2059	0.2184	0.2102	0.2777	0.2407	0.2656
Base+CI+TA+GR+PI	0.2257	0.2461	0.2327	0.2789	0.2427	0.2670	0.2376	0.2602	0.2454	0.2832	0.2447	0.2706

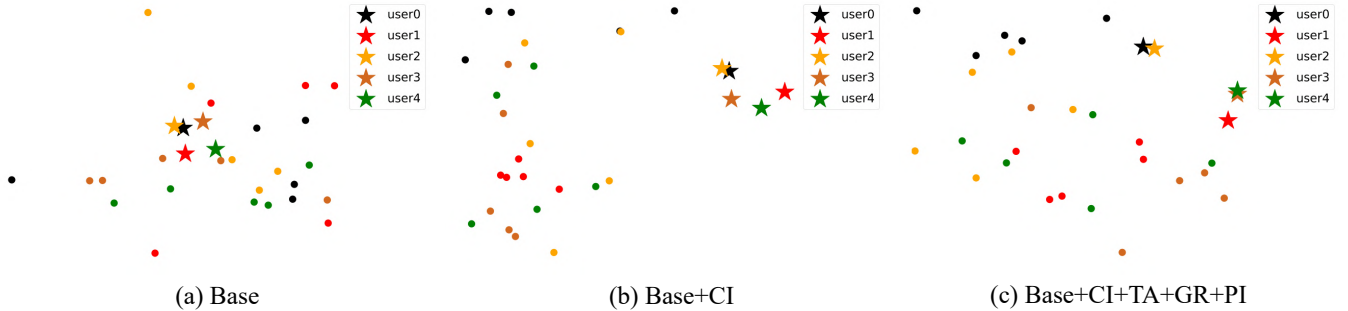


Fig. 3. Visualization of the learned t-SNE transformed representations derived from (a) Base, (b) Base+CI, and (c) Base+CI+TA+GR+PI mentioned in Sec.IV-C. Each star denotes a user and the points with the same color indicate the interacted item.

decrease in cold-start performance (MF(Image) achieved a 4.9% decrease on Amazon_CDs and MF(Semantics) obtained a 1.9% decrease on Allrecipes), demonstrating the necessity of fine-grained processing of multi-modal information in cold-start recommendations.

- PiNet [30] outperforms the traditional visually-aware recommendation methods (e.g., VECF [55] and HAFR-non-i [18]). This is mainly owing to the fact that it introduces semantic information in addition to visual information to regularize the learning process of the content representation.
- CIERec outperforms existing multi-modal learning methods in all performance metrics, which validates that CIERec can bring significant and consistent improvement over existing recommendation methods by fusing the user representation, the collaborative representation, and the multi-modal content information obtained from cross-modal inference.
- CIERec achieves higher cold-start improvements on the Amazon_CDs, which is due to its fewer classes of semantic elements. That is, the same semantic element category corresponds to more items, thus making the cross-modal inferred semantic information more representative. It also proves the importance of semantic information in cold-start recommendation.

C. Ablation Study

In addition to the overall performance comparison, we further explore the effectiveness of the combinations of components in CIERec. Specifically, we use Base, CI, TA, GR and PI to represent using the components of plain recommendation model, cross-modal inference, task aware gate, gradient regularization and privileged information respectively. The ablation results are shown in Table III and we have the following findings:

- For 'CI', that is, cross-modal inference without any constraints, the application of naive inference in recommendations leads to a significant degradation of its performance. This may be explained by the fact that a large number of visual information is lost in the cross-modal transformation.
- For 'TA' and 'GR', 'TA' facilitates the modeling of mapping relationships between heterogeneous modalities through task-aware gates in each modality; while 'GR' can help to learn the optimization direction of heterogeneous representations through gradient-regularization gates within the constraints of reinforcement learning. CIERec can further improve its performance by combining these components.
- The 'PI' component is able to introduce the prior knowledge (e.g., the textual annotations of the image of the item) as the privileged information based on the existing components, which helps to enhance the model's ability to mine visual information. As such, it allows the

original visual features to be fully utilized to achieve the best cold-start recommendation results.

D. Case Study

In this section, we attempt to investigate how the CIERec facilitates the learning process of the multi-modal representation in the embedding space. To this end, we randomly selected five users and items they interacted with to explore how these embeddings varied in different methods.

As shown in Figure 3, the relevance of users and items is well reflected in the t-SNE space, namely, the more relevant representations are embedded in the more similar positions. We found that naive cross-modal inference leads to a collapse of its training procedure, whereas the embedding representations learned by CIERec show a significant clustering effect, that is, points with the same color tend to form clusters. These observations demonstrate that CIERec is able to effectively facilitate the learning process of cross-modal representations by augmenting content information, so that the representation of users and the items they interact with tend to be close to each other, which may be one of the reasons for CIERec's superior performance.

V. CONCLUSION

This paper proposes a novel cross-modal content inference and feature enrichment recommendation framework, CIERec, which conducts the cross-modal inference from the visual space to the semantic space based on the items' prior knowledge, and combines a multi-modal representation fusion method to trade-off the heterogeneous representation modeling process from the multi-modal information. Experimental results demonstrate that the introduction of cross-modal inferred information is able to improve the items' representation from multiple perspectives, which makes CIERec superior to existing methods in cold-start recommendation.

Future work of this study will focus on two main directions. First, the heterogeneous alignment techniques may help to model the mapping relationships between visual features and cross-modal inferred features at a multi-granularity level, leading to an increased information gain. Second, we will further improve the representational capability of CIERec by filtering noisy information in cross-modal inference with the graph convolutional network.

ACKNOWLEDGMENT

This work is supported in part by the National Natural Science Foundation of China (Grant no. 62006141), by the Excellent Youth Scholars Program of Shandong Province (Grant no. 2022HWYQ-048), and the TaiShan Scholars Program (Grant no. tsqn202211289)

REFERENCES

- [1] Y. Zhang, Y. Liu, Y. Xu, H. Xiong, C. Lei, W. He, L. Cui, and C. Miao, "Enhancing sequential recommendation with graph contrastive learning," *arXiv preprint arXiv:2205.14837*, 2022.
- [2] X. Du, X. Wang, X. He, Z. Li, J. Tang, and T.-S. Chua, "How to learn item representation for cold-start multimedia recommendation?" in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 3469–3477.
- [3] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "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*, 2020, pp. 639–648.
- [4] F. Sun, J. Liu, J. Wu, C. Pei, X. Lin, W. Ou, and P. Jiang, "Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer," in *Proceedings of the 28th ACM international conference on information and knowledge management*, 2019, pp. 1441–1450.
- [5] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "Bpr: Bayesian personalized ranking from implicit feedback," *arXiv preprint arXiv:1205.2618*, 2012.
- [6] R. He and J. McAuley, "Vbpr: visual bayesian personalized ranking from implicit feedback," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, 2016.
- [7] H. Ma, X. Li, L. Meng, and X. Meng, "Comparative study of adversarial training methods for cold-start recommendation," in *Proceedings of the 1st International Workshop on Adversarial Learning for Multimedia*, ser. ADVM '21, 2021, p. 28–34.
- [8] Y. Wei, X. Wang, Q. Li, L. Nie, Y. Li, X. Li, and T.-S. Chua, "Contrastive learning for cold-start recommendation," in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 5382–5390.
- [9] X. Yu, L. Meng, X. Tian, S. Fauvel, B. Huang, Y. Guan, Z. Shen, C. Miao, and C. Leung, "Usability analysis of the novel functions to assist the senior customers in online shopping," in *Social Computing and Social Media. User Experience and Behavior*, G. Meiselwitz, Ed., 2018.
- [10] H. Wu, X. Chen, X. Li, H. Ma, Y. Zheng, X. Li, X. Meng, and L. Meng, "A visually-aware food analysis system for diet management," in *2022 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*. IEEE, 2022, pp. 1–1.
- [11] J. Yu, M. Gao, J. Li, H. Yin, and H. Liu, "Adaptive implicit friends identification over heterogeneous network for social recommendation," in *Proceedings of the 27th ACM international conference on information and knowledge management*, 2018, pp. 357–366.
- [12] C. Shi, B. Hu, W. X. Zhao, and S. Y. Philip, "Heterogeneous information network embedding for recommendation," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 2, pp. 357–370, 2018.
- [13] D.-K. Chae, J.-S. Kang, S.-W. Kim, and J. Choi, "Rating augmentation with generative adversarial networks towards accurate collaborative filtering," in *The World Wide Web Conference*, 2019, pp. 2616–2622.
- [14] H. Lee, J. Im, S. Jang, H. Cho, and S. Chung, "Melu: Meta-learned user preference estimator for cold-start recommendation," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 1073–1082.
- [15] O. Barkan, N. Koenigstein, E. Yogev, and O. Katz, "Cb2cf: a neural multiview content-to-collaborative filtering model for completely cold item recommendations," in *Proceedings of the 13th ACM Conference on Recommender Systems*, 2019, pp. 228–236.
- [16] S. Kang, J. Hwang, D. Lee, and H. Yu, "Semi-supervised learning for cross-domain recommendation to cold-start users," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019, pp. 1563–1572.
- [17] L. Meng, L. Chen, X. Yang, D. Tao, H. Zhang, C. Miao, and T.-S. Chua, "Learning using privileged information for food recognition," in *Proceedings of the 27th ACM International Conference on Multimedia*, 2019, pp. 557–565.
- [18] X. Gao, F. Feng, X. He, H. Huang, X. Guan, C. Feng, Z. Ming, and T.-S. Chua, "Hierarchical attention network for visually-aware food recommendation," *IEEE Transactions on Multimedia*, vol. 22, no. 6, pp. 1647–1659, 2019.
- [19] H. Chen, Y. Deng, Y. Li, T.-Y. Hung, and G. Lin, "Rgb salient object detection via disentangled cross-modal fusion," *IEEE Transactions on Image Processing*, vol. 29, pp. 8407–8416, 2020.
- [20] J. Li, H. Ma, X. Li, Z. Qi, L. Meng, and X. Meng, "Unsupervised contrastive masking for visual haze classification," in *Proceedings of the 2022 International Conference on Multimedia Retrieval*, 2022.
- [21] X. Li, H. Ma, L. Meng, and X. Meng, "Comparative study of adversarial training methods for long-tailed classification," in *Proceedings of the 1st International Workshop on Adversarial Learning for Multimedia*, 2021.
- [22] X. Chen, L. Wu, M. He, L. Meng, and X. Meng, "Mlfont: Few-shot chinese font generation via deep meta-learning," in *Proceedings of the 2021 International Conference on Multimedia Retrieval*, 2021, pp. 37–45.

- [23] L. Wu, X. Chen, L. Meng, and X. Meng, "Multitask adversarial learning for chinese font style transfer," in *2020 international joint conference on neural networks (IJCNN)*. IEEE, 2020, pp. 1–8.
- [24] L. Meng, A.-H. Tan, and D. Wunsch, *Adaptive resonance theory in social media data clustering*. Springer, 2019.
- [25] L. Meng, A.-H. Tan, and D. Xu, "Semi-supervised heterogeneous fusion for multimedia data co-clustering," *IEEE Transactions on Knowledge and Data Engineering*, vol. 26, no. 9, pp. 2293–2306, 2013.
- [26] L. Meng, A.-H. Tan, and D. C. Wunsch, "Adaptive scaling of cluster boundaries for large-scale social media data clustering," *IEEE transactions on neural networks and learning systems*, vol. 27, no. 12, pp. 2656–2669, 2015.
- [27] L. Meng, A.-H. Tan, and C. Miao, "Salience-aware adaptive resonance theory for large-scale sparse data clustering," *Neural Networks*, vol. 120, pp. 143–157, 2019.
- [28] H. Fu, R. Wu, C. Liu, and J. Sun, "Mcen: Bridging cross-modal gap between cooking recipes and dish images with latent variable model," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 14 570–14 580.
- [29] L. Meng, A.-H. Tan, D. C. Wunsch II, L. Meng, A.-H. Tan, and D. C. Wunsch II, "Online multimodal co-indexing and retrieval of social media data," *Adaptive Resonance Theory in Social Media Data Clustering: Roles, Methodologies, and Applications*, pp. 155–174, 2019.
- [30] L. Meng, F. Feng, X. He, X. Gao, and T.-S. Chua, "Heterogeneous fusion of semantic and collaborative information for visually-aware food recommendation," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 3460–3468.
- [31] X. Wu, S. Cetintas, D. Kong, M. Lu, J. Yang, and N. Chawla, "Learning from cross-modal behavior dynamics with graph-regularized neural contextual bandit," in *Proceedings of The Web Conference 2020*, 2020, pp. 995–1005.
- [32] Y. Zhu, J. Lin, S. He, B. Wang, Z. Guan, H. Liu, and D. Cai, "Addressing the item cold-start problem by attribute-driven active learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 4, pp. 631–644, 2019.
- [33] X. Wang, X. He, Y. Cao, M. Liu, and T.-S. Chua, "Kgat: Knowledge graph attention network for recommendation," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 950–958.
- [34] M. Jian, T. Jia, X. Yang, L. Wu, and L. Huo, "Cross-modal collaborative manifold propagation for image recommendation," in *Proceedings of the 2019 on International Conference on Multimedia Retrieval*, 2019, pp. 344–348.
- [35] Y. Lu, Y. Fang, and C. Shi, "Meta-learning on heterogeneous information networks for cold-start recommendation," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1563–1573.
- [36] J. Liu, J. Xiao, H. Ma, X. Li, Z. Qi, X. Meng, and L. Meng, "Prompt learning with cross-modal feature alignment for visual domain adaptation," in *Artificial Intelligence: Second CAAI International Conference, CICA 2022, Beijing, China, August 27–28, 2022, Revised Selected Papers, Part I*, 2022.
- [37] Y. Wang, X. Li, H. Ma, Z. Qi, X. Meng, and L. Meng, "Causal inference with sample balancing for out-of-distribution detection in visual classification," in *Artificial Intelligence: Second CAAI International Conference, CICA 2022, Beijing, China, August 27–28, 2022, Revised Selected Papers, Part I*, 2022.
- [38] C. Lin, S. Zhao, L. Meng, and T. Chua, "Multi-source domain adaptation for visual sentiment classification," *CoRR*, vol. abs/2001.03886, 2020. [Online]. Available: <https://arxiv.org/abs/2001.03886>
- [39] X. Li, L. Wu, X. Chen, L. Meng, and X. Meng, "Dse-net: Artistic font image synthesis via disentangled style encoding," in *2022 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2022, pp. 1–6.
- [40] Y. Wang, X. Li, Z. Qi, J. Li, X. Li, X. Meng, and L. Meng, "Meta-causal feature learning for out-of-distribution generalization," in *European Conference on Computer Vision*. Springer, 2023, pp. 530–545.
- [41] W. Guo, Y. Zhang, X. Cai, L. Meng, J. Yang, and X. Yuan, "Ld-man: Layout-driven multimodal attention network for online news sentiment recognition," *Trans. Multi.*, vol. 23, p. 1785–1798, jan 2021. [Online]. Available: <https://doi.org/10.1109/TMM.2020.3003648>
- [42] P. Dong, L. Wu, L. Meng, and X. Meng, "Hr-prgan: High-resolution story visualization with progressive generative adversarial networks," *Inf. Sci.*, vol. 614, no. C, p. 548–562, oct 2022. [Online]. Available: <https://doi.org/10.1016/j.ins.2022.10.083>
- [43] —, "Disentangled representations and hierarchical refinement of multi-granularity features for text-to-image synthesis," in *Proceedings of the 2022 International Conference on Multimedia Retrieval*, 2022, pp. 268–276.
- [44] J. Chen, H. Zhang, X. He, L. Nie, W. Liu, and T.-S. Chua, "Attentive collaborative filtering: Multimedia recommendation with item-and component-level attention," in *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*, 2017, pp. 335–344.
- [45] J. Tang, X. Du, X. He, F. Yuan, Q. Tian, and T.-S. Chua, "Adversarial training towards robust multimedia recommender system," *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 5, pp. 855–867, 2019.
- [46] Q. Liu, S. Wu, and L. Wang, "Deepstyle: Learning user preferences for visual recommendation," in *Proceedings of the 40th international acm sigir conference on research and development in information retrieval*, 2017, pp. 841–844.
- [47] J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, and X. He, "Bias and debias in recommender system: A survey and future directions," *arXiv preprint arXiv:2010.03240*, 2020.
- [48] M. Volkovs, G. Yu, and T. Poutanen, "Dropoutnet: Addressing cold start in recommender systems," *Advances in neural information processing systems*, vol. 30, 2017.
- [49] X. He, Z. He, X. Du, and T.-S. Chua, "Adversarial personalized ranking for recommendation," in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018, pp. 355–364.
- [50] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, 1997.
- [51] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [52] Z. Zhou, L. Zhang, and N. Yang, "Contrastive collaborative filtering for cold-start item recommendation," *arXiv preprint arXiv:2302.02151*, 2023.
- [53] H. Ma, X. Li, L. Meng, and X. Meng, "Comparative study of adversarial training methods for cold-start recommendation," in *Proceedings of the 1st International Workshop on Adversarial Learning for Multimedia*, 2021, pp. 28–34.
- [54] Q. Li, L. Shen, S. Guo, and Z. Lai, "Wavelet integrated cnns for noise-robust image classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7245–7254.
- [55] X. Chen, H. Chen, H. Xu, Y. Zhang, Y. Cao, Z. Qin, and H. Zha, "Personalized fashion recommendation with visual explanations based on multimodal attention network: Towards visually explainable recommendation," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019, pp. 765–774.