

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Cognitive-Inspired Hierarchical Attention Fusion With Visual and Textual for Cross-Domain Sequential Recommendation

Permalink

<https://escholarship.org/uc/item/8n60x457>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors

Wu, Wangyu
Chen, Zhenhong
Song, Siqi
[et al.](#)

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Hierarchical Attention Fusion of Visual and Textual Representations for Cross-Domain Sequential Recommendation

Wangyu Wu^{a, b}, Zhenhong Chen^c, Siqi Song^{a, b}
Xianglin Qiu^{a, b}, Xiaowei Huang^b, Fei Ma^{a, 1}, Jimin Xiao^{a, 1}

^a Xi'an Jiaotong-Liverpool University ^b The University of Liverpool

^c Microsoft

Wangyu.wu@liverpool.ac.uk
{fei.ma, jimin.xiao}@xjtlu.edu.cn

Abstract

Cross-Domain Sequential Recommendation (CDSR) predicts user behavior by leveraging historical interactions across multiple domains, focusing on modeling cross-domain preferences through intra- and inter-sequence item relationships. Inspired by human cognitive processes, we propose Hierarchical Attention Fusion of Visual and Textual Representations (HAF-VT), a novel approach integrating visual and textual data to enhance cognitive modeling. Using the frozen CLIP model, we generate image and text embeddings, enriching item representations with multimodal data. A hierarchical attention mechanism jointly learns single-domain and cross-domain preferences, mimicking human information integration. Evaluated on four e-commerce datasets, HAF-VT outperforms existing methods in capturing cross-domain user interests, bridging cognitive principles with computational models and highlighting the role of multimodal data in sequential decision-making.

Keywords: Hierarchical Attention Fusion; Cognitive Modeling; Cross-Domain Sequential Recommendation; CLIP-based Embeddings

Introduction

Sequential Recommendation (SR) has emerged as a prominent approach for modeling dynamic user preferences by analyzing historical interaction sequences (Shani, Heckerman, & Brafman, 2002; He, Kang, & McAuley, 2017; He & McAuley, 2016; J. Li et al., 2017). While SR focuses on predicting the next item of interest within a single domain, it often suffers from data sparsity and domain bias, limiting its ability to capture the full spectrum of user behavior. To address these limitations, Cross-Domain Sequential Recommendation (CDSR) (M. Ma et al., 2019; Zhuang et al., 2020; Guo et al., 2021; Lei et al., 2021; Cao, Cong, Sheng, Liu, & Wang, 2022; C. Li et al., 2022) has been proposed, leveraging information from multiple domains to enhance recommendation accuracy and provide a more comprehensive understanding of user preferences.

Early efforts in CDSR, such as π -Net (M. Ma et al., 2019), model item interaction sequences within a single domain and transfer learned representations to other domains through gated transfer modules. Subsequent work, like MIFN (M. Ma et al., 2022), incorporates external knowledge graphs to strengthen cross-domain connections. However, these approaches primarily focus on intra-sequence relationships within individual domains, often neglecting inter-




Figure 1: (a) In the traditional CDSR framework, the input consists solely of item ID features. (b) In our HAF-VT framework, we incorporate additional image and title information to complement the existing item features. Therefore, our method enriches item representations.

sequence relationships across domains. This oversight leads to persistent domain bias and suboptimal performance.

Moreover, existing CDSR methods fail to fully exploit the rich visual information associated with items. Human decision-making is inherently multimodal, with visual cues playing a critical role in shaping user preferences (Hegarty, 2011). For instance, users often form initial impressions of items based on their visual appearance, even when textual descriptions are similar. This observation aligns with cognitive theories emphasizing the integration of multimodal stimuli in preference formation (Shams & Seitz, 2008). To bridge this gap, we propose **Hierarchical Attention Fusion of Visual and Textual Representations (HAF-VT)**, a novel framework that integrates visual and textual embeddings to enrich item representations and model cross-domain user preferences. By leveraging the frozen CLIP model (Radford et al., 2021), we generate aligned image and text embeddings, cap-

¹Corresponding author

turing both intra- and inter-sequence relationships. Our hierarchical attention mechanism jointly learns single-domain and cross-domain preferences, mimicking the human ability to integrate information across contexts.

Our main contributions are summarized as follows:

- We introduce the CDSR framework that integrates visual, textual, and item embeddings, enriching item representations and enhancing recommendation performance through multimodal cognitive modeling.
- We propose a hierarchical attention mechanism that captures intra-sequence and inter-sequence relationships, enabling more accurate next-item prediction by aligning with human cognitive processes.
- We conduct extensive experiments on four reorganized e-commerce datasets, demonstrating that HAF-VT outperforms state-of-the-art baselines across all evaluation metrics.

Related Work

Models for Cognitive Science

Recent advancements in computer vision have demonstrated significant potential for addressing key research questions in cognitive science (L. Zhou et al., 2025; L. Zhou, Chen, et al., 2024). A growing body of work has focused on image similarity and categorization (Yao, Xu, Hou, & Chuah, 2024; Yao & Chuah, 2024; L. Zhou, Xia, et al., 2024; Chen, Wu, & Li, 2024; Qiu et al., 2024; P. Wu et al., 2025), which, while related to the phenomenon of naming, does not directly address it. These studies collectively reveal that deep visual representations, particularly those derived from the final layers of neural networks, align closely with human perceptual judgments. This alignment is often measured by comparing human similarity ratings with model-generated similarity scores, highlighting the ability of deep learning models to capture human-like visual representations. More directly relevant to the present work are studies by (W. Wu, Dai, et al., 2025; Gualdoni, Brochhagen, Mädebach, & Boleda, 2023; Günther, Marelli, Tureski, & Petilli, 2023; W. Wu et al., 2024), which leverage deep learning-based computer vision models to construct visual semantic representations for extensive vocabularies. These approaches bridge the gap between visual and linguistic domains, offering insights into how visual information can be mapped to semantic structures. Recent research has also begun to explore cognitive modeling using the CLIP (Contrastive Language–Image Pretraining) framework. The integration of self-supervised learning with vision-language pretraining has spurred the development of numerous CLIP-based models, each aiming to enhance the alignment between visual and textual data (Cha, Mun, & Roh, 2023). These models extend the original CLIP architecture by introducing novel techniques to improve representation quality and strengthen the connection between images and text.

Despite these advancements, existing methods often fail to fully integrate visual and textual information to enhance model perception. In contrast, our approach employs a Hierarchical Attention mechanism to effectively capture and combine information from multiple feature sources, thereby improving the overall performance and alignment with human cognitive processes.

Sequential Recommendation

Sequential recommendation (Wang et al., 2019) serves as the core framework of our approach, where user behavior is modeled as a sequence of time-sensitive items with the goal of predicting the likelihood of the next item in the sequence. Initially, Markov chains were utilized to capture the sequential dependencies within the item sequences, as seen in FPMC (Rendle, Freudenthaler, & Schmidt-Thieme, 2010). To capture more complex interactions between historical items, deep learning techniques such as recurrent neural networks (Chung, Gulcehre, Cho, & Bengio, 2014; Hochreiter & Schmidhuber, 1997), convolutional neural networks (LeCun, Bengio, et al., 1995), and attention mechanisms (Vaswani et al., 2017) have been integrated into recommender systems (Hidasi, Karatzoglou, Baltrunas, & Tikk, 2015; G. Zhou et al., 2019; Tang & Wang, 2018; Kang & McAuley, 2018a; G. Zhou et al., 2018). In addition, SURGE (Chang et al., 2021) uses metric learning to reduce the complexity of the item sequences. More recently, works like DFAR (Lin et al., 2023) and DCN (Lin et al., 2022) have explored capturing deeper, more intricate relationships within sequential recommendation data. In this work, we extend these ideas by applying cross-domain learning to sequential recommendation models, facilitating knowledge transfer between multiple domains.

Cross-Domain Recommendation

Cross-domain recommender systems (Cantador, Fernández-Tobías, Berkovsky, & Cremonesi, 2015) have emerged as a promising approach to address the challenges of sparse data and cold-start issues typically faced in sequential recommendation tasks (W. Wu, Song, et al., 2025). Early models in this area were grounded in single-domain recommendation frameworks, under the assumption that auxiliary user behaviors from different domains could contribute to improved user modeling in the target domain (Singh & Gordon, 2008). Many popular methods rely on transfer learning (S. J. Pan & Yang, 2009), which involves transferring user or item embeddings from the source domain to the target domain to enhance the modeling process. Notable examples of such approaches include MiNet (Ouyang et al., 2020), CoNet (Hu, Zhang, & Yang, 2018), and itemCST (W. Pan, Xiang, Liu, & Yang, 2010), among others. However, in practice, industrial platforms typically aim to improve all product domains simultaneously, rather than exclusively enhancing the target domain without considering the source domain. This has led to the development of dual learning (Qiu, Li, et al., 2025; Qiu, Wu, et al., 2025), which enables mutual improvements across both




Figure 2: The overview of our proposed HAF-VT is as follows. Firstly, the Feature Preparation Process embeds all items from domains X and Y using a learnable item matrix for ID-based embeddings E_{id} , a frozen CLIP image encoder for image-based embeddings E_{img} , and a text encoder for text-based embeddings E_{tex} . Then, the input sequence S is forwarded to the Image Fusion Prediction Process. S consists of sub-sequences S^X , S^Y , and S^{X+Y} , which are processed through an embedding mapping layer to generate item ID-based, image-based, and text-based sequence embeddings. These embeddings are passed through multiple attention layers to capture both intra- and inter-sequence relationships, producing attention-aggregated embeddings. Finally, cosine similarity is applied between these aggregated embeddings and E_{id} , E_{img} , and E_{tex} to predict the next item.

the source and target domains, and has already been successfully applied in cross-domain recommender systems. Despite these advancements, many existing methods fail to account for the role of visual perceptual features and do not integrate multimodal fusion with attention mechanisms, which could significantly enhance performance.

Methodology

Problem Formulation

In the CDSR task, user interaction sequences occur in domain X and domain Y , with their item sets denoted as $\mathcal{X}$ and $\mathcal{Y}$, respectively. Let S represent the overall interaction sequence of the user in chronological order, which consists of three sub-sequences $(S^X, S^Y, S^{X+Y}) \in S$. Specifically, $S^X = [x_1, x_2, \dots, x_{|S^X|}]$, $x \in \mathcal{X}$ and $S^Y = [y_1, y_2, \dots, y_{|S^Y|}]$, $y \in \mathcal{Y}$ are the interaction sequences within each domain, where $|\cdot|$ denotes the total item number. Additionally, $S^{X+Y} = [x_1, y_1, x_2, \dots, x_{|S^X|}, \dots, y_{|S^Y|}]$ represents the merged sequence by combining S^X and S^Y . In general, the goal of the CDSR task is to predict the probabilities of candidate items across both domains and select the item with the highest probability as the next recommendation.

Overall framework

As illustrated in Fig. 2, the proposed HAF-VT framework is designed to model cross-domain sequential preferences through a hierarchical cognitive architecture. The framework begins with Item Representation Encoding, where items from domains X and Y are processed through parallel pathways to construct distinct representations. A learnable semantic embedding matrix E_{id} is initialized to capture identity-specific

features, while a frozen CLIP encoder generates perceptual embeddings E_{img} and conceptual embeddings E_{tex} from item images and textual descriptions, respectively. This multimodal encoding stage simulates the human capacity for integrating visual, semantic, and linguistic information.

The framework then proceeds to the Preference Simulation stage, where behavioral sequences S are processed. Each sequence S consists of three sub-sequences: S^X , S^Y , and S^{X+Y} , representing within-domain and cross-domain interactions. The embedding layer retrieves and restructures the multimodal embeddings (E_{id} , E_{img} , and E_{tex}) according to the sequence order, generating dynamic cognitive states for each sub-sequence. These states are further processed through hierarchical attention layers to model both intra-sequence and inter-sequence relationships, simulating the interplay between local working memory and global contextual integration.

Finally, the framework performs Decision Generation by comparing the attention-aggregated sequence embeddings with the original multimodal embeddings (E_{id} , E_{img} , and E_{tex}) using cosine similarity. This process mimics human decision-making by combining evidence from multiple sensory modalities to predict the next item in the sequence.

Visual and Textual Feature Integration

We enhance user preferences by integrating multimodal features (image and textual) into the CDSR framework. This subsection provides a detailed explanation of feature preparation and sequence representation combined with image and textual features within a single domain, and applies this approach to domain X , domain Y , and domain $X + Y$.

1) *Item Presentations Preparation with Image Feature:* First, we prepare the feature representations for all items. As

shown on the left in Fig. 2, we construct a learnable item matrix based on item IDs for all items in domain X and Y , denoted as $E_{id} \in \mathbb{R}^{(|\mathcal{X}|+|\mathcal{Y}|) \times q}$. Here, $|\mathcal{X}| + |\mathcal{Y}|$ represents the total number of items from domains X and Y , and q is the learnable item embedding dimension. Simultaneously, we employ a pre-trained and frozen CLIP model to generate image embeddings for each item. These embeddings are used to form image matrix $E_{img} \in \mathbb{R}^{(|\mathcal{X}|+|\mathcal{Y}|) \times e}$, where e is the image embedding dimension. These two matrices serve as the base for generating embeddings for sequences.

2) *Item Presentations Preparation with Textual Feature:* To further enrich the item representations, we leverage textual information from item titles using the same frozen CLIP model. For each item, the CLIP text encoder generates textual embeddings based on the item's title, forming the textual matrix $E_{tex} \in \mathbb{R}^{(|\mathcal{X}|+|\mathcal{Y}|) \times e}$. This textual representation captures semantic information that complements the visual and ID-based features, enabling a more comprehensive understanding of user preferences.

3) *Enhanced Similarity Score with Multimodal Features:* We first generate a sequence representation to represent the user interaction sequence, which is then used to calculate the similarity with items for predicting the next item in the Prediction Process. As shown on the right in Fig. 2, we propose an embedding layer to process the input sequences $\mathcal{S}$, which consists of three sub-sequences $\mathcal{S}^X$, $\mathcal{S}^Y$, and $\mathcal{S}^{X+Y}$. These sequences include item IDs, images, and textual data. Accordingly, the embedding layer produces item ID-based, image-based, and textual-based embeddings for each input sequence.

Once we obtain E_{id} , E_{img} , and E_{tex} , the item ID-based sequence embeddings $F_{id} \in \mathbb{R}^{|\mathcal{S}| \times q}$ can be generated by placing the appropriate embedding from E_{id} in the order of the items in the sequence $\mathcal{S}$. Here, $|\mathcal{S}|$ denotes the total item number in $\mathcal{S}$. Similarly, the image-based sequence embeddings $F_{img} \in \mathbb{R}^{|\mathcal{S}| \times e}$ and textual-based sequence embeddings $F_{tex} \in \mathbb{R}^{|\mathcal{S}| \times e}$ are produced from E_{img} and E_{tex} , respectively, based on the same sequence.

For simplicity, we take the image-based sequence embeddings F_{img} as an example to illustrate how HAF-VT obtains the next item prediction. We employ an Attention Layer to get the enhanced sequence image-based embeddings $H_{img} \in \mathbb{R}^{|\mathcal{S}| \times e}$ as follows:

$$H_{img} = \text{Attention}(F_{img}), \quad (1)$$

Afterward, we extract the last embedding vector $h_{img} \in \mathbb{R}^{1 \times e}$ from H_{img} as the sequence representation, which serves as the attention-aggregated embedding of the sequence. This vector captures the user preferences, specifically focusing on the most recent interaction within the sequence. Then, h_{img} is compared against E_{img} using cosine similarity. In this way, the alignment between user preferences and the embedding of each item across all domains is assessed. The similarity score is computed as follows:




Figure 3: The process of transforming a user sequence into a representation sequence. Initially, the user's sequence of interactions is converted into an embedded representation. This sequence is then passed through an attention layer to capture both intra- and inter-sequence relationships. Finally, the attention-aggregated sequence representations are compared with item embeddings, and the item with the highest similarity score is selected as the predicted next item.

$$\text{Sim}(h_{img}, E_{img}) = \frac{h_{img} \cdot E_{img}^T}{\|h_{img}\| \|E_{img}^T\|}, \quad (2)$$

where T denotes the transpose function. Here, a higher similarity score indicates that the sequence preference is more aligned with this item. In the same manner, we can obtain the item ID-based sequence representative vector $h_{id} \in \mathbb{R}^{1 \times q}$ and textual-based sequence representative vector $h_{tex} \in \mathbb{R}^{1 \times e}$, along with their respective similarity scores. In the next section, we will introduce how the multiple attention layer is used to fuse ID-based, image-based, and textual-based predictions for CDSR.

Hierarchical Attention in Preference Formation

In traditional sequential recommendation methods (Hidasi et al., 2015; Kang & McAuley, 2018b), behavioral sequences from domains X and Y are often processed indiscriminately, leading to potential dominance by the domain with higher behavioral frequency. This imbalance can skew the representation of user preferences. To address this issue, we propose a hierarchical attentional framework that separately processes sequences $\mathcal{S}^X$, $\mathcal{S}^Y$, and $\mathcal{S}^{X+Y}$, ensuring balanced representation across domains. Our approach integrates item identity (ID), visual features, and textual embeddings to model user preferences more comprehensively. As illustrated in Fig. 3, to model the interactions within each sequence more effectively, we incorporate a self-attention mechanism with query, key, and value (QKV) pairs. The attention mechanism allows the model to focus on relevant parts of the sequence when making predictions, capturing both local and global dependencies within the sequence.

The attention mechanism computes the output as a weighted sum of the values, where the weights are determined by the similarity between the query and the keys:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V, \quad (3)$$

where Q , K , and V are the query, key, and value matrices, respectively, and d_k is the dimension of the key vectors. This

mechanism is applied within each domain’s sequence processing to capture both intra-domain and cross-domain relationships more effectively.

The framework generates nine sequence representations through multiple attention layers: $h_{id}^X, h_{img}^X, h_{tex}^X, h_{id}^Y, h_{img}^Y, h_{tex}^Y, h_{id}^{X+Y}, h_{img}^{X+Y},$ and h_{tex}^{X+Y} . These representations capture both intra-domain and cross-domain user preferences. For domain X , the prediction probability for the next item x_{t+1} in sequence S^X is computed as follows:

$$P_{id}^X(x_{t+1} | S) = \text{softmax} \left(\text{Sim}(h_{id}^X, E_{id}^X) \right), \quad x_t \in S^X, \quad (4)$$

where x_{t+1} denotes the predicted next item, h_{id}^X is the ID-based sequence representation, and E_{id}^X is the ID embedding matrix for domain X . Similarly, the prediction probabilities based on visual and textual features are calculated as:

$$P_{img}^X(x_{t+1} | S) = \text{softmax} \left(\text{Sim}(h_{img}^X, E_{img}^X) \right), \quad x_t \in S^X, \quad (5)$$

$$P_{tex}^X(x_{t+1} | S) = \text{softmax} \left(\text{Sim}(h_{tex}^X, E_{tex}^X) \right), \quad x_t \in S^X, \quad (6)$$

where $h_{img}^X, E_{img}^X, h_{tex}^X,$ and E_{tex}^X represent the image-based sequence representation, image embedding matrix, text-based sequence representation, and text embedding matrix, respectively. The final prediction probability for the next item combines ID, image, and text-based predictions:

$$\begin{aligned} P^X(x_{t+1} | S) &= \alpha P_{id}^X(x_{t+1} | S) \\ &\quad + \beta P_{img}^X(x_{t+1} | S) \\ &\quad + (1 - \alpha - \beta) P_{tex}^X(x_{t+1} | S), \end{aligned} \quad (7)$$

where α and β are weighting factors that balance the contributions of ID, visual, and textual features. This formulation captures structural, perceptual, and semantic aspects of user preferences.

To optimize the framework, we minimize the following loss function for domain X :

$$\mathcal{L}^X = \sum_{x_t \in S^X} -\log P^X(x_{t+1} | S). \quad (8)$$

The same method is applied to domains Y and $X + Y$, yielding loss terms $\mathcal{L}^Y$ and $\mathcal{L}^{X+Y}$. The final loss is a weighted combination of these terms:

$$\mathcal{L} = \mathcal{L}^X + \lambda_1 \mathcal{L}^Y + \lambda_2 \mathcal{L}^{X+Y}, \quad (9)$$

where λ_1 and λ_2 control the relative importance of each domain.

During evaluation, we aggregate predictions from all domains to recommend the next item. Let X be the target domain and Y the source domain. The combined prediction score for item x_i is computed as:

$$P(x_i | S) = \mathcal{P}^X(x_i | S) + \lambda_1 \mathcal{P}^Y(x_i | S) + \lambda_2 \mathcal{P}^{X+Y}(x_i | S). \quad (10)$$

The recommended item is selected by maximizing the prediction score within the target domain X :

$$\text{argmax}_{x_i \in \mathcal{X}} P(x_i | S). \quad (11)$$

This process allows the model to incorporate cross-domain information effectively, leveraging the predictions from both the target domain and the source domain. By weighting the contributions from each domain, the model can refine its recommendation, taking into account the richer information that might come from the source domain while focusing on the target domain for the final prediction.

The weighting factors λ_1 and λ_2 allow the model to control how much influence the source domain (Y) and the combined domain ($X + Y$) have on the recommendation for the target domain. This flexibility in adjusting the weights is crucial for handling different types of datasets or domain-specific characteristics, ensuring that the model adapts to varying degrees of relevance between domains.

Finally, by selecting the item with the highest prediction score from the target domain, the model aims to provide the most contextually appropriate recommendation, which is critical for real-world recommendation systems that need to balance both accuracy and relevance across multiple domains.

Experiments

In this section, we introduce the experimental setup, providing detailed descriptions of the dataset, evaluation metrics, and implementation specifics. We then compare our method with state-of-the-art approaches on the Amazon dataset (Wei et al., 2021). Finally, we conduct ablation studies to validate the effectiveness of the proposed method.

Dataset and Evaluated Metric. Our experiments are conducted on the Amazon dataset (Wei et al., 2021) to construct the CDSR scenarios. Following previous works (M. Ma et al., 2019, 2022), we select four domains to generate two CDSR scenarios for our experiments: “Food-Kitchen” and “Movie-Book”. We first extract users who have interactions in both domains. Then we filter out users and items with fewer than 10 interactions. Additionally, to meet the sequential constraints, we retain cross-domain interaction sequences that contain at least three items from each domain. In the training/validation/test split, the latest interaction sequences are equally divided into the validation and test sets, while the remaining interaction sequences are used for the training set. The statistics of our refined datasets for the CDSR scenarios are summarized in Table 1. We adopt Mean Reciprocal Rank (MRR) (Voorhees, 1999) and Normalized Discounted Cumulative Gain (NDCG@){5, 10} (Järvelin & Kekäläinen, 2002) as evaluation metrics to assess the performance of our model. These metrics are widely used in recommendation systems.

Implementation Details. For fair comparisons, we adopt the same hyperparameter settings as in previous works (M. Ma et al., 2022). We set the learnable embedding size q to 256, the CLIP-based image and text embedding size e to 512, and the mini-batch size to 256, with a dropout

Table 1: Statistics of two CDSR scenarios.

Scenarios	#Items	#Train	#Valid	#Test	Avg.length
Food	29,207	34,117	2,722	2,747	9.91
Kitchen	34,886		5,451	5,659	
Movie	36,845	58,515	2,032	1,978	11.98
Book	63,937		5,612	5,730	

Table 2: Experimental results (%) on the Food-Kitchen scenario.

Model (Food-Kitchen)	MRR (Food)	NDCG		MRR (Kitchen)	NDCG	
		@5	@10		@5	@10
GRU4Rec (Hidasi et al., 2015)	5.79	5.48	6.13	3.06	2.55	3.10
SASRec (Kang & McAuley, 2018b)	7.30	6.90	7.79	3.79	3.35	3.93
SR-GNN (S. Wu et al., 2019)	7.84	7.58	8.35	4.01	3.47	4.13
MIFN (M. Ma et al., 2022)	8.55	8.28	9.01	4.09	3.57	4.29
Tri-CDR (H. Ma et al., 2024)	8.35	8.18	8.89	4.29	3.63	4.33
Ours HAF-VT	9.16	8.95	9.98	5.03	4.64	5.35

rate of 0.3. The L_2 regularization coefficient is selected from $\{0.0001, 0.00005, 0.00001\}$, and the learning rate is chosen from $\{0.001, 0.0005, 0.0001\}$. Training is conducted for 100 epochs using an NVIDIA 4090 GPU, with the Adam optimizer (Kingma & Ba, 2015) to update all parameters.

Table 3: Experimental results (%) on the Movie-Book scenario.

Model (Movie-Book)	MRR (Movie)	NDCG		MRR (Book)	NDCG	
		@5	@10		@5	@10
GRU4Rec (Hidasi et al., 2015)	3.83	3.14	3.73	1.68	1.34	1.52
SASRec (Kang & McAuley, 2018b)	3.79	3.23	3.69	1.81	1.41	1.71
SR-GNN (S. Wu et al., 2019)	3.85	3.27	3.78	1.78	1.40	1.66
PSJNet (Sun et al., 2022)	4.63	4.06	4.76	2.44	2.07	2.35
MIFN (M. Ma et al., 2022)	5.05	4.21	5.20	2.51	2.12	2.31
Tri-CDR (H. Ma et al., 2024)	5.15	4.62	5.05	2.32	2.08	2.22
Ours HAF-VT	6.27	5.21	5.98	2.84	2.45	2.76

Performance Comparisons

We evaluate our method against existing state-of-the-art (SOTA) approaches on the ‘‘Food-Kitchen’’ and ‘‘Movie-Book’’ CDSR scenarios, as summarized in Tab. 2 and Tab. 3, respectively. Our approach achieves superior performance, surpassing the current SOTA methods in terms of final evaluation metrics, thereby demonstrating its effectiveness in cross-domain recommendation tasks.

Ablation Studies

We conducted an ablation study to systematically evaluate the contributions of our proposed components: *Textual Fusion*, *Visual Fusion*, and *Hierarchical Attention*. As illustrated in Tab. 4, we first established a baseline using the original framework without any additional modules, achieving an MRR of 5.03. The introduction of *Visual Fusion* alone resulted in a performance gain of 0.6 points, highlighting the importance of incorporating visual signals for cross-domain recommendation tasks. Similarly, adding *Textual Fusion*

Table 4: Ablation study on the movie dataset

original-framework	Textual Fusion	Visual Fusion	Hierarchical Attention	MRR
✓				5.03
✓		✓		5.63
✓	✓			5.48
✓	✓		✓	5.88
✓		✓	✓	6.08
✓	✓	✓	✓	6.27

without hierarchical attention improved the MRR to 5.48, demonstrating the complementary role of textual information in enhancing the model’s understanding of user preferences.

When both *Textual Fusion* and *Hierarchical Attention* were integrated, the MRR further increased to 5.88, underscoring the effectiveness of hierarchical structures in capturing domain-specific features. Notably, the combination of *Visual Fusion* and *Hierarchical Attention* achieved an MRR of 6.08, indicating that visual signals and hierarchical modeling synergistically contribute to performance improvements.

Finally, the full framework, incorporating *Textual Fusion*, *Visual Fusion*, and *Hierarchical Attention*, achieved the highest MRR of 6.27. This result demonstrates the cumulative impact of our proposed modules, with each component playing a distinct yet complementary role in enhancing the model’s ability to capture cross-domain user preferences. Overall, our ablation study confirms the effectiveness of the HAF-VT method in advancing CDSR performance.

Conclusion

In this work, we propose a novel framework, *Hierarchical Attention with Visual and Textual Fusion* (HAF-VT), for Cross-Domain Sequential Recommendation. Departing from conventional methods that focus exclusively on temporal patterns in user interaction sequences, our approach integrates multimodal representations by harnessing the robust visual and textual understanding capabilities of the CLIP model. Specifically, we generate frozen image and text embeddings, which are then fused with a learnable item matrix to enrich item representations. To further enhance the model’s ability to capture cross-domain user preferences, we introduce a hierarchical attention mechanism. This mechanism is designed to process sequential information from multiple domains and modalities, enabling the model to discern and leverage distinct features from each source. By combining these innovations, our HAF-VT framework achieves state-of-the-art performance in CDSR, demonstrating the effectiveness of integrating multimodal signals and hierarchical modeling for sequential recommendation tasks.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62471405, 62331003, 62301451), Suzhou Basic Research Program (SYG202316), XJTLU REF-22-01-010, and the SIP AI Innovation Platform (YZCXPT2022103).

References

- Cantador, I., Fernández-Tobías, I., Berkovsky, S., & Cremonesi, P. (2015). Cross-domain recommender systems. In *Recommender systems handbook* (pp. 919–959). Springer.
- Cao, J., Cong, X., Sheng, J., Liu, T., & Wang, B. (2022). Contrastive cross-domain sequential recommendation. In *Proceedings of the 31st acm international conference on information & knowledge management* (pp. 138–147).
- Cha, J., Mun, J., & Roh, B. (2023). Learning to generate text-grounded mask for open-world semantic segmentation from only image-text pairs. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition* (pp. 11165–11174).
- Chang, J., Gao, C., Zheng, Y., Hui, Y., Niu, Y., Song, Y., ... Li, Y. (2021). Sequential recommendation with graph neural networks. In *Proceedings of the 44th international acm sigir conference on research and development in information retrieval* (pp. 378–387).
- Chen, Y., Wu, W., & Li, J. (2024). Adaptive attention-enhanced yolo for wall crack detection. *Applied Sciences* (2076-3417), 14(17).
- Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Nips 2014 workshop on deep learning, december 2014*.
- Gualdoni, E., Brochhagen, T., Mädebach, A., & Boleda, G. (2023). What's in a name? a large-scale computational study on how competition between names affects naming variation. *Journal of Memory and Language*, 133, 104459.
- Günther, F., Marelli, M., Tureski, S., & Petilli, M. A. (2023). Vispa (vision spaces): A computer-vision-based representation system for individual images and concept prototypes, with large-scale evaluation. *Psychological Review*, 130(4), 896.
- Guo, L., Tang, L., Chen, T., Zhu, L., Nguyen, Q. V. H., & Yin, H. (2021). Da-gcn: A domain-aware attentive graph convolution network for shared-account cross-domain sequential recommendation. In *International joint conference on artificial intelligence (ijcai)*.
- He, R., Kang, W.-C., & McAuley, J. (2017). Translation-based recommendation.
- He, R., & McAuley, J. (2016). Fusing similarity models with markov chains for sparse sequential recommendation. In *Ieee international conference on data mining (icdm)*.
- Hegarty, M. (2011). The cognitive science of visual-spatial displays: Implications for design. *Topics in cognitive science*, 3(3), 446–474.
- Hidasi, B., Karatzoglou, A., Baltrunas, L., & Tikk, D. (2015). Session-based recommendations with recurrent neural networks. In *International conference on learning representations (iclr)*.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Hu, G., Zhang, Y., & Yang, Q. (2018). Conet: Collaborative cross networks for cross-domain recommendation. In *Acm international conference on information and knowledge management (cikm)*.
- Järvelin, K., & Kekäläinen, J. (2002). Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*.
- Kang, W.-C., & McAuley, J. (2018a). Self-attentive sequential recommendation. In *2018 ieee international conference on data mining (icdm)* (pp. 197–206).
- Kang, W.-C., & McAuley, J. (2018b). Self-attentive sequential recommendation. In *Ieee international conference on data mining (icdm)*.
- Kingma, P. D., & Ba, L. J. (2015). Adam: A method for stochastic optimization. In *International conference on learning representations (iclr)*.
- LeCun, Y., Bengio, Y., et al. (1995). Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*.
- Lei, C., Liu, Y., Zhang, L., Wang, G., Tang, H., Li, H., & Miao, C. (2021). Semi: A sequential multi-modal information transfer network for e-commerce micro-video recommendations. In *Acm knowledge discovery and data mining (kdd)*.
- Li, C., Zhao, M., Zhang, H., Yu, C., Cheng, L., Shu, G., ... Niu, D. (2022). Recguru: Adversarial learning of generalized user representations for cross-domain recommendation. In *Acm international conference on web search and data mining (wsdm)*.
- Li, J., Ren, P., Chen, Z., Ren, Z., Lian, T., & Ma, J. (2017). Neural attentive session-based recommendation. In *Acm international conference on information and knowledge management (cikm)*.
- Lin, G., Gao, C., Li, Y., Zheng, Y., Li, Z., Jin, D., & Li, Y. (2022). Dual contrastive network for sequential recommendation. In *Proceedings of the 45th international acm sigir conference on research and development in information retrieval* (pp. 2686–2691).
- Lin, G., Gao, C., Zheng, Y., Chang, J., Niu, Y., Song, Y., ... Li, Y. (2023). Dual-interest factorization-heads attention for sequential recommendation. In *Proceedings of the acm web conference 2023* (pp. 917–927).
- Ma, H., Xie, R., Meng, L., Chen, X., Zhang, X., Lin, L., & Zhou, J. (2024). Triple sequence learning for cross-domain recommendation. *ACM Transactions on Information Systems*, 42(4), 1–29.
- Ma, M., Ren, P., Chen, Z., Ren, Z., Zhao, L., Ma, J., & de Rijke, M. (2022). Mixed information flow for cross-domain sequential recommendations. *ACM Transactions on Knowledge Discovery from Data (TKDD)*.
- Ma, M., Ren, P., Lin, Y., Chen, Z., Ma, J., & Rijke, M. d. (2019). π -net: A parallel information-sharing network for shared-account cross-domain sequential recommendations. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval* (pp. 685–694).
- Ouyang, W., Zhang, X., Zhao, L., Luo, J., Zhang, Y., Zou,

- H., ... Du, Y. (2020). Minet: Mixed interest network for cross-domain click-through rate prediction. In *Proceedings of the 29th acm international conference on information & knowledge management* (pp. 2669–2676).
- Pan, S. J., & Yang, Q. (2009). A survey on transfer learning. , 22(10), 1345–1359. (Publisher: IEEE)
- Pan, W., Xiang, E., Liu, N., & Yang, Q. (2010). Transfer learning in collaborative filtering for sparsity reduction. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 24). (Issue: 1)
- Qiu, X., Hu, J., Zhou, L., Wu, X., Du, J., Zhang, B., ... Yang, B. (2024). Tfb: Towards comprehensive and fair benchmarking of time series forecasting methods. In *Proc. VLDB endow.* (pp. 2363–2377).
- Qiu, X., Li, X., Pang, R., Pan, Z., Wu, X., Yang, L., ... Yang, B. (2025). Easytime: Time series forecasting made easy. In *Icde*.
- Qiu, X., Wu, X., Lin, Y., Guo, C., Hu, J., & Yang, B. (2025). Duet: Dual clustering enhanced multivariate time series forecasting. In *Sigkdd* (p. 1185–1196).
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... others (2021). Learning transferable visual models from natural language supervision. In *International conference on machine learning* (pp. 8748–8763).
- Rendle, S., Freudenthaler, C., & Schmidt-Thieme, L. (2010). Factorizing personalized markov chains for next-basket recommendation. In *Www* (pp. 811–820).
- Shams, L., & Seitz, A. R. (2008). Benefits of multisensory learning. *Trends in cognitive sciences*, 12(11), 411–417.
- Shani, G., Heckerman, D., & Brafman, R. I. (2002). An mdp-based recommender system. *Journal of Machine Learning Research (JMLR)*.
- Singh, A. P., & Gordon, G. J. (2008). Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 650–658).
- Sun, W., Ren, P., Lin, Y., Ma, M., Chen, Z., Ren, Z., ... de Rijke, M. (2022). Parallel split-join networks for shared account cross-domain sequential recommendations. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*.
- Tang, J., & Wang, K. (2018). Personalized top-n sequential recommendation via convolutional sequence embedding. In *Acm international conference on web search and data mining (wsdm)*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. In *Neurips* (pp. 5998–6008).
- Voorhees, E. M. (1999). The trec-8 question answering track report. In *Text retrieval conference (trec)*.
- Wang, S., Hu, L., Wang, Y., Cao, L., Sheng, Q. Z., & Orgun, M. (2019). Sequential recommender systems: challenges, progress and prospects.
- Wei, Y., Wang, X., Li, Q., Nie, L., Li, Y., Li, X., & Chua, T.-S. (2021). Contrastive learning for cold-start recommendation. In *Proceedings of the 29th acm international conference on multimedia* (pp. 5382–5390).
- Wu, P., Chai, B., Li, H., Zheng, M., Peng, Y., Wang, Z., ... Sun, X. (2025). Spiking point transformer for point cloud classification. *arXiv preprint arXiv:2502.15811*.
- Wu, S., Tang, Y., Zhu, Y., Wang, L., Xie, X., & Tan, T. (2019). Session-based recommendation with graph neural networks. In *Aaai conference on artificial intelligence (aaai)*.
- Wu, W., Dai, T., Chen, Z., Huang, X., Ma, F., & Xiao, J. (2025). Generative prompt controlled diffusion for weakly supervised semantic segmentation. *Neurocomputing*, 130103.
- Wu, W., Qiu, X., Song, S., Huang, X., Ma, F., & Xiao, J. (2024). Prompt categories cluster for weakly supervised semantic segmentation. *arXiv preprint arXiv:2412.13823*.
- Wu, W., Song, S., Qiu, X., Huang, X., Ma, F., & Xiao, J. (2025). Image fusion for cross-domain sequential recommendation. In *Companion proceedings of the acm web conference 2025*.
- Yao, Z., & Chuah, M. C. (2024). Event-guided low-light video semantic segmentation. *arXiv preprint arXiv:2411.00639*.
- Yao, Z., Xu, J., Hou, S., & Chuah, M. C. (2024). Crack-nex: a few-shot low-light crack segmentation model based on retinex theory for uav inspections. In *Icra* (pp. 11155–11162).
- Zhou, G., Mou, N., Fan, Y., Pi, Q., Bian, W., Zhou, C., ... Gai, K. (2019). Deep interest evolution network for click-through rate prediction. In *Aaai* (pp. 5941–5948).
- Zhou, G., Zhu, X., Song, C., Fan, Y., Zhu, H., Ma, X., ... Gai, K. (2018). Deep interest network for click-through rate prediction. In *Kdd* (pp. 1059–1068).
- Zhou, L., Chen, H., Zhou, X., Yuan, Y., Zhu, W., & Zhou, M. (2024). A smart agriculture image protection scheme based on annealing algorithm and affine transformation is optimized for s-box generated by chaos. *Nonlinear Dynamics*, 1–25.
- Zhou, L., Xia, H., Lin, Q., Yang, X., Zhang, X., & Zhou, M. (2024). Two-dimensional hyperchaos-based encryption and compression algorithm for agricultural uav-captured planar images. *Scientific Reports*.
- Zhou, L., Xia, H., Lyu, W., Huang, G., Chen, H., Zhou, H., ... man, z. (2025). Image protection scheme for bridge management systems based on quantum coupling function. *Physica Scripta*.
- Zhuang, F., Zhou, Y., Ying, H., Zhang, F., Ao, X., Xie, X., ... Xiong, H. (2020). Sequential recommendation via cross-domain novelty seeking trait mining. *Journal of Computer Science and Technology (JCST)*.