LSG-Rec: A Unified LLM-Enhanced Framework for Cold-Start and Cross-Domain Recommendation

Zhihao Wang
Trip.com
Shanghai, China
wang_zhihao@trip.com

Ziyi Hu
Trip.com
Shanghai, China
ziyi.hu@trip.com

Pengcheng Zhuo
Trip.com
Shanghai, China
pc.zhuo@trip.com

Wei Ding
Trip.com
Shanghai, China
wding3@trip.com

Saijie Ni*
Trip.com
Shanghai, China
sj.ni@trip.com

***Abstract*—Item cold-start and long-tail recommendation remain critical challenges in travel platforms due to the sparse interaction data. Although Large Language Models (LLMs) possess rich world knowledge and can extract item semantic information, this is interaction-agnostic and disconnected from collaborative signals. This often leads to popularity bias and "blurry" embeddings that fail to capture nuanced user preferences. To bridge this gap, we propose a unified framework that aligns LLM-enhanced multimodal abstractions with interaction-based dynamics. Our approach incorporates: (1) a LLM-powered encoder to distill high-fidelity semantics through commonsense reasoning; (2) a Minimax-driven alignment module that filters popularity-biased noise to reconstruct precise collaborative representations; and (3) a cross-domain heterogeneous GNN that facilitates robust preference transfer across diverse categories, such as UGC (User Generated Content) and PGC (Professionally Generated Content). Extensive experiments conducted on the homepage feed of the Our App demonstrate that our framework effectively warms up cold items and significantly outperforms state-of-the-art baselines in both long-tail and cross-domain scenarios.**

I. Introduction

The explosion of digital content has rendered recommender systems indispensable for mitigating information overload. While conventional ID-based Collaborative Filtering (CF) excels at capturing blockbuster trends, it inherently discriminates against **long-tail products**—the vast majority of items that suffer from severe data sparsity. This challenge escalates into the **item cold-start problem** in fast-moving domains such as e-commerce and short-video feeds, where the item life-cycle is often shorter than the time required to collect sufficient interaction feedback.

To address this, content-based methods have been proposed to bridge the gap between item metadata and CF representations [1]. Recent advancements leverage Multi-modal Modeling and LLMs to distill high-level semantics. However, as noted in Chen et al[2], a significant gap remains between multi-modal embedding extraction and actual user interest modeling. Furthermore, generative models that project content into the embedding space often suffer from "blurry" collaborative signals, failing to capture the distinct co-occurrence patterns present in warm data [3]. These models are also restricted by the topology of the interaction graph; since cold nodes are isolated, they cannot benefit from high-order connectivity unless the graph is explicitly reconstructed through interaction-based high-order connectivity [4].

*Corresponding author

Specifically, in **cross-domain** scenarios, the challenge is twofold: LLMs lack domain-specific interaction awareness, and the semantic resolution of encoders remains too coarse to capture the subtle "shared tastes" spanning across different catalogs. To reconcile LLM-based semantics with collaborative dynamics, we propose a unified framework that synergistically orchestrates Content-to-Collaborative Filtering (CB2CF), Graph Neural Networks (GNNs), and LLM-powered encoders. Our contributions are summarized as follows:

- **LLM-Enhanced Semantic Abstraction**: We utilize LLMs to distill multimodal content, injecting commonsense reasoning to interpret niche long-tail attributes that traditional encoders miss.
- **CB2CF Alignment**: Inspired by the bridge between content and CF [1], we treat LLM vectors as pseudo-IDs and employ a minimax game to filter out popularity-biased noise, ensuring alignment with real user patterns.
- **Cross-Domain Heterogeneous GNN**: We construct a unified graph using meta-path attention to transfer preferences across domains while suppressing domain-specific quirks through gradient reversal, effectively warming up cold items in new categories.

II. Related Work

A. Cold-Start and Cross-Domain Problems in Rec Systems

Content-based filtering (CBF) leverages item features for cold-start recommendations but suffers from sparse or noisy features[5], making hybrid CF-CBF methods a mainstream solution—for instance, DropoutNet[6] simulates cold-start by dropping warm items in training and fuses content features; recent advancements like SimRec[7] and CCFCRec[8] further boost performance via similarity integration, generative modeling and multi-source data fusion. Cross-domain recommendation (CDR) shares core challenges with cold-start while focusing on inter-domain knowledge migration. LeCDSR[9] uses LLMs to generate item semantic embeddings and cross-domain user profiles to supplement missing semantic information and characterize cross-domain user preferences accurately. CL-DTCDR[10] designs cross-domain contrastive learning tasks to optimize user/item embeddings for effective knowledge transfer with few overlapping users. DPM-CDR[11] adopts distributional preference modeling to align domain-invariant latent preferences, minimizing cross-domain distribution gaps without explicit graph structures.

B. Graph Neural Networks for Cold-Start and Cross-Domain Recommendation

For cold-start and cross-domain recommendation, graph neural networks (GNNs) leverage relational information in user-item graphs to compensate for data scarcity and excel in modeling inter-domain relationships and facilitating knowledge transfer. Cao et al. presented GIFT[12], which builds a heterogeneous graph to transfer knowledge from warm to cold videos, significantly improving performance. Guo et al. proposed DA-GCN[13], a domain-aware attentive GCN for shared-account cross-domain sequential recommendation, which constructs domain-specific graphs and uses dual attention mechanisms to refine user representations. These works verify GNNs' effectiveness in cold-start scenarios and cross-domain data via collaborative signal capture and knowledge transfer.

C. Summary

Despite their respective contributions, traditional methods and GNNs both exhibit notable limitations in practical deployment. Traditional methods suffer from the over-reliance of content-based transfer on high-quality shared features, which are frequently unavailable in real-world scenarios. GNNs, while mitigating cold-start issues via knowledge transfer from warm entities and addressing cross-domain challenges through dedicated graph construction and transfer mechanisms, still fail to meet the rigorous industrial robustness demands against extreme data sparsity, severe domain heterogeneity, and stringent real-time efficiency requirements.

III. Methodology

In this section, we present our LSG-Rec framework. The overall framework is illustrated in Fig.1. First LLM is applied to extract semantic information and generate tags for each item. Then, a dual-tower model (CB2CF) fuses semantic information with collaborative filtering signals. Finally, we construct an item-tag bipartite graph and leverage graph neural networks (GNNs) to integrate node information. After joint training of these three components, the final model will generate embedding for every item, greatly improving cold-start performance and cross domain recommendation performance.

A. LLaMA

By integrating product images and textual data into Large Multimodal Models(LMMs) to generate semantic embeddings and tags, we can leverage the model's inherent understanding of multimodal information. This process yields an embedding-based representation capturing an objective summary of product attributes.

Our system supports the seamless integration of various state-of-the-art models. Depending on the computational budget and task-specific requirements, the embedding backbone can be switched from dedicated unimodal encoders (e.g., DINOv2 for images or BGE-M3 for text) to advanced models like GPT-4o or Qwen2.5-VL. This modularity ensures that the framework can leverage the evolving capabilities of LMMs to obtain superior semantic representations.

B. CB2CF

LLMs generate item embeddings rich in semantic information, offering a deep understanding of content and descriptions. In parallel, user historical behaviors provide a direct, collective reflection of item correlations through co-occurrence patterns, a relationship powerfully leveraged by conventional item-to-item (I2I) collaborative filtering [14]. Compared to traditional I2I methods that rely exclusively on these behavioral signals, LLMs introduce a complementary, generalized perspective on item relatedness.To achieve a deep fusion of these two information paradigms—semantic and collaborative—we propose CB2CF, a dual-tower model with a shared structure and shared weights, inspired by [1] and [15].

1) Shared Weights Dual Tower.: Unlike conventional dual-tower architectures that process different inputs in separate, specialized networks, our model employs a single, unified encoder for both information types. This shared tower takes as input a combined representation that integrates LLM-generated embeddings with static product-side information (e.g., category, price). Through this design, the model is compelled to learn a unified representation space where semantic signals, item static information and collaborative signals mutually enhance and disambiguate each other, rather than being processed in isolation. This approach effectively bridges the semantic-CF gap, leading to more robust item representations, especially for items with sparse interactions.

2) Model Struture Detail.: The architecture of our proposed CB2CF model is designed to process item pairs for relationship prediction, as illustrated in Fig.2. Each item in the pair is processed independently through an identical network structure with shared parameters, following a siamese network paradigm.

For each item i, the input is composed of two components:

- A semantic embedding $\mathbf{e}_{\text{LLM}}^{(i)} \in \mathbb{R}^{d_{\text{LLM}}}$ derived from the item's textual content using a LLM
- The item's static side information processed through an embedding layer for categorical features, yielding $\mathbf{e}_{\text{static}}^{(i)} \in \mathbb{R}^{d_{\text{static}}}$

These components are concatenated to form the combined input representation:

$$\mathbf{x}^{(i)} = \text{Concat}(\mathbf{e}_{\text{LLM}}^{(i)}, \mathbf{e}_{\text{static}}^{(i)}) \tag{1}$$

This combined representation is fed through a series of three fully-connected layers with non-linear activations, ultimately projecting each item into a compact, 16-dimensional embedding space. The forward pass for a single item can be formally described as:

$$\mathbf{e}^{(i)} = W_3\sigma(W_2\sigma(W_1x^{(1)} + b_1) + b_2) + b_3 \tag{2}$$

where $W_1 \in \mathbb{R}^{d_1 \times d_{\text{in}}}, W_2 \in \mathbb{R}^{d_2 \times d_1}$ and $W_3 \in \mathbb{R}^{d_3 \times d_2}$ are weight matrices.

For a pair of items (i, j), the model computes their respective embeddings $\mathbf{e}^{(i)}$ and $\mathbf{e}^{(j)}$ using the same shared network. The dot product between these two embeddings is scaled by

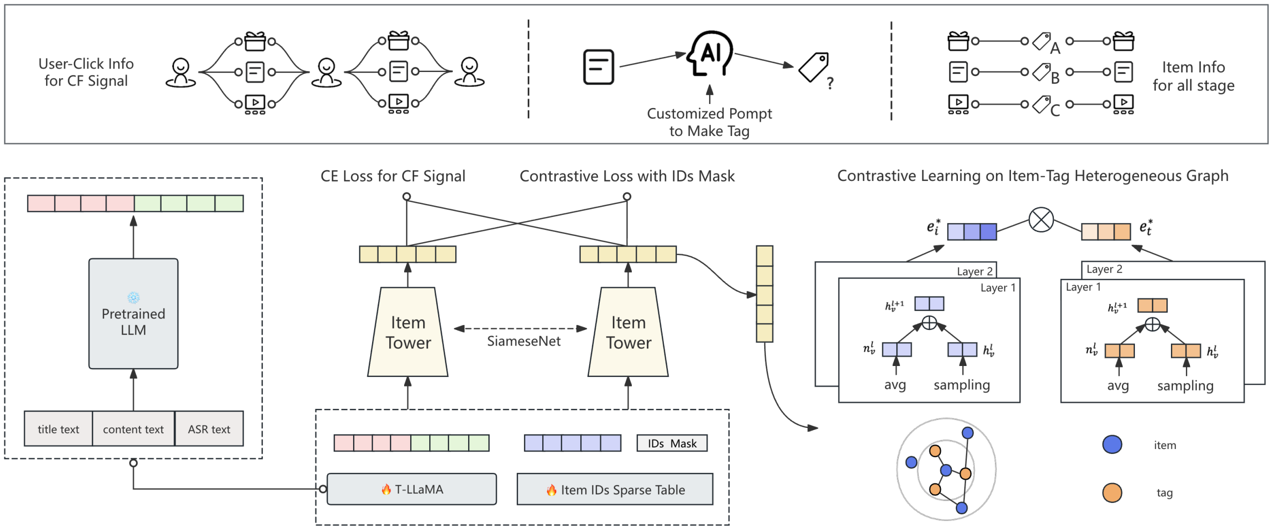

Fig. 1: Overall architecture of the proposed method.

a temperature coefficient τ to produce the output logit for predicting the relationship between the items:

$$\text{logit}_{ij} = \frac{\mathbf{e}^{(i)} \cdot \mathbf{e}^{(j)}}{\tau} \tag{3}$$

where $\tau > 0$ serves as a scaling factor that modulates the sharpness of the similarity distribution, which is inspired by [16].

3) Sample Selection Strategy: Our sampling strategy is designed to ensure high-quality training signals while maintaining computational efficiency:

- **Positive Samples**: We construct positive pairs from the top-5 item pairs generated by item-based collaborative filtering (itemCF). Each positive pair (i, j) is required to have at least one item with sufficient exposure history, ensuring reliable collaborative signals.
- **Negative Samples**: Negative samples are obtained through in-batch negative sampling. For each positive pair (i, j) in a training batch, we utilize all other items in the same batch as negative candidates, excluding any items that would form actual positive pairs with the anchor item to prevent false negatives.

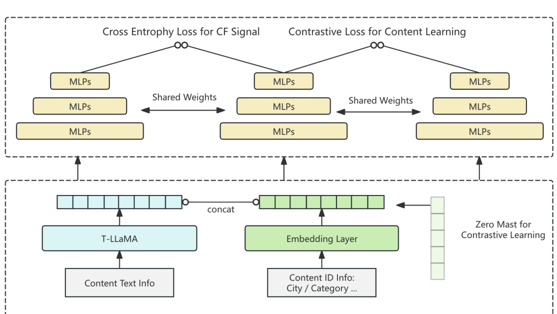

Fig. 2: Overall architecture of the proposed method.

This strategy ensures that we leverage high-confidence collaborative relationships as positive signals while efficiently utilizing batch composition for negative sampling without introducing label conflicts.

C. Graph Neural Networks

Leveraging LLMs, we automatically annotate every item with a diverse set of open-domain tags that capture its multifaceted characteristics. Since several item may share the same tag, their interplay naturally forms an item–tag bipartite graph whose structure is most effectively exploited by Graph Neural Networks(GNNs).

1) Homogeneous Neighborhood Augmentation via Contrastive Learning.: To combat the extreme sparsity of the item-tag bipartite graph we replace the conventional "interaction = positive" paradigm with homogeneous-neighbourhood contrastive learning: after a two-step random walk that lands on nodes of the same type (item→item or tag→tag), we treat the reached homogeneous neighbours as positive pairs and all remaining nodes as negatives. This strategy expands the expected number of positives from the scarce n direct neighbours to n^2 homogeneous neighbours, yielding an abundant, balanced contrastive signal for every vertex. The resulting InfoNCE-style[17] objective pulls each node closer to its type-consistent neighbours and pushes it away from others, dramatically densifying the supervisory signal and amplifying the benefit of contrastive augmentation without any extra data.

Homogeneous neighborhood augmentation via contrastive learning can increase the number of positive pairs in the case of sparse data, making the effect of contrastive learning more obvious. Simultaneously, it is more rational to use contrastive learning to encourage homogeneous neighbors to approach each other than to minimize the distance between item and tag nodes with different node types.

Lightweight Parameters According to previous works [18], [19], feature transformation and nonlinear activation, which frequently appear in the design of GNNs, degrade collaborative filtering. Traditional GNN is proposed for node classification on the attribute graph. In collaborative filtering scenarios with inadequate node features, the implementation of multi-layer nonlinear feature transformation will not bring any benefits but will increase the difficulty of model training. Therefore, we also avoid these two designs and remove all the trainable parameters in GNN module. After determining the sampled neighbors, we add the averaged features of neighbors to the node's features with weight.

D. Loss Function

Our overall training objective combines three loss components that respectively address the recommendation task, feature debiasing, and graph-based representation learning.

1) Cross Entropy Loss in CB2CF: The primary cross-entropy loss for the recommendation task is defined as:

$$\mathcal{L}_{CE}^{(\text{CB2CF})} = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log(\hat{y}_i) + (1-y_i)\log(1-\hat{y}_i)\right] \quad (4)$$

where $\hat{y}_i$ is the predicted probability for sample i, and $y_i \in \{0,1\}$ is the corresponding ground truth label.

2) Contrastive Learning Loss for Feature Debiasing in CB2CF: To mitigate feature dominance through representation consistency in the CB2CF module, we employ a contrastive learning loss:

$$\mathcal{L}_{CL}^{(\text{CB2CF})} = -\frac{1}{|B|}\sum_{i\in B}\log\left(\frac{\exp\left(\text{sim}(\mathbf{e}_i, \mathbf{e}_i^m)/\tau_1\right)}{\sum_{\substack{j\in B\\ j\neq i}}\exp\left(\text{sim}(\mathbf{e}_i, \mathbf{e}_j)/\tau_1\right)}\right) \quad (5)$$

where $\mathbf{e}_i$ and $\mathbf{e}_i^m$ are the original and masked embeddings of item i, $\tau_1 > 0$ is the temperature hyperparameter, and $\text{sim}(\mathbf{u}, \mathbf{v}) = \frac{\mathbf{u}\cdot\mathbf{v}}{\|\mathbf{u}\|\|\mathbf{v}\|}$ denotes cosine similarity. The masking operation encourages robust representations that are invariant to partial feature removal.

3) Graph Contrastive Learning Loss in GNNs: For the GNNs component, we employ a homogeneous-neighborhood contrastive learning strategy to combat the sparsity of the item-tag bipartite graph. After a two-step random walk that lands on nodes of the same type (item→item or tag→tag), we treat the reached homogeneous neighbors as positive pairs. This expands the expected number of positives from the scarce direct neighbors to quadratic homogeneous neighbors, yielding abundant contrastive signals.

The graph contrastive loss combines item and tag components:

$$\mathcal{L}_{CL}^{(\text{GNNs})} = \mathcal{L}_C^{\text{I}} + \alpha\mathcal{L}_C^{\text{T}} \quad (6)$$

where α is a balancing hyperparameter, and $\mathcal{L}_C^{\text{I}}$ and $\mathcal{L}_C^{\text{T}}$ represent the InfoNCE-style contrastive losses for item and tag nodes respectively, with temperature τ_2.

4) Total Training Objective: The overall training objective combines all three loss components with weighting factors:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{CE}^{(\text{CB2CF})} + \lambda_1\cdot\mathcal{L}_{CL}^{(\text{CB2CF})} + \lambda_2\cdot\mathcal{L}_{CL}^{(\text{GNNs})} \quad (7)$$

where $\lambda_1, \lambda_2 \in [0,1]$ control the emphasis on the debiasing and graph contrastive objectives respectively.

This unified loss formulation enables joint optimization of recommendation accuracy, representation robustness, and graph-structured knowledge propagation, ensuring that all components of the LSG-Rec framework work synergistically towards the common goal of improving recommendation quality across diverse scenarios.

IV. Experiments

To comprehensively evaluate the effectiveness of our proposed LSG-Rec framework, we conduct experiments from two distinct perspectives: cold-start recommendation and cross-domain recommendation. This dual-scenario evaluation framework allows us to assess the model's capabilities in addressing fundamental challenges in real-world recommendation systems.

We evaluate LSG-Rec using both offline and online experiments. The offline experiments utilize a real-world dataset comprising 5 million user interaction logs sampled from a major online travel platform's information feed. This dataset provides realistic user behavior patterns and diverse item distributions, offering a robust testbed for evaluating our method.

A. Offline Evaluation

1) User Intent Satisfaction: In cold-start scenarios, accurately matching user intent is crucial. We compare LSG-Rec against other I2I recall methods by measuring their *Intent Satisfaction Rate* at different recall depths. Satisfaction is determined by whether the recalled item aligns with the user's trigger (e.g., city, province, or specific intent entity).

TABLE I: Intent Satisfaction Rate (%)

Method	City Match	Province Match	Entity Match
N2V	42.83	80.31	49.09
BCE	19.76	31.72	62.89
LSG-Rec (Ours)	**47.66**	**83.91**	**54.22**

As shown in Table I, LSG-Rec achieves superior intent satisfaction across all match types in cold-start scenarios. The integration of semantic understanding enabled by our framework's LLM and CB2CF components leads to more precise item alignment compared to methods relying solely on co-occurrence statistics.

2) Consistency with Downstream Ranking Model: We evaluate the alignment between our cold-start recall and the final ranking stage by measuring consistency metrics over the top-10 ranked items.

Table II demonstrates that LSG-Rec achieves the best overall consistency with the downstream ranking model in cold-start scenarios. Our framework significantly outperforms baseline approaches in Recall@10, indicating that the LLM-enhanced embeddings from our CB2CF module provide more relevant cold-start items that align with final ranking criteria.

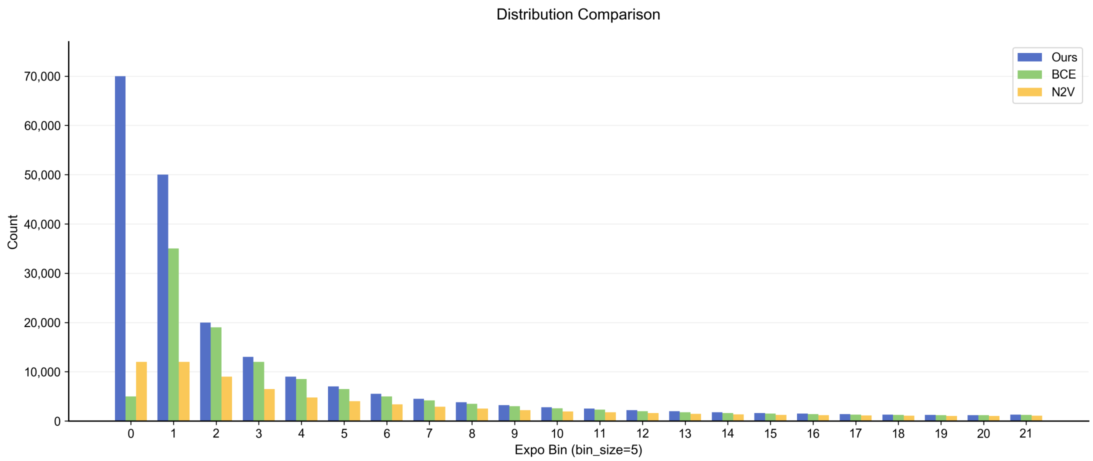

Fig. 3: Distribution of historical exposure counts for items recalled by different methods in cold-start scenarios. LSG-Rec recalls a significantly higher proportion of items with low past exposure, demonstrating its ability to address the cold-start problem.

TABLE II: Consistency with ranking model

Method	Recall@10	HitRate@10
N2V	0.006	0.1
BCE	0.0049	0.0846
LSG-Rec (Ours)	**0.0147**	**0.2032**

3) Cold-start and Long-tail Item Recall: A key advantage of LSG-Rec is its effectiveness in recommending cold-start and long-tail items. We analyze the historical exposure distribution of items recalled by different methods. As illustrated in Figure 3, LSG-Rec successfully recalls a substantially larger proportion of items with minimal historical exposure compared to traditional collaborative filtering methods. This capability stems from the framework's ability to leverage LLM-based semantic content understanding, bypassing the dependency on extensive interaction history.

B. Online AB-Testing

1) Cold-start Scenarios: We deploy LSG-Rec in a production environment and conduct large-scale A/B tests over a four-week period, focusing specifically on cold-start scenarios.

TABLE III: Online A/B testing results for cold-start scenarios (relative improvement over baseline)

Group	**CTR**	**clickPV**	**Concentration**
Cold-start (G2)	+3.77%	+4.43%	-0.11%
General Rec (G3)	+0.73%	+0.87%	-1.22%

Key business metrics are evaluated: Click-Through Rate (CTR), Click Per View (clickPV), and Recommendation Concentration (a diversity metric where lower values indicate less concentration on popular items).

The results, presented as relative improvements over the baseline, are summarized in Table III. LSG-Rec demonstrates the most substantial gains in the cold-start scenario, achieving a CTR increase of 3.77% and a clickPV increase of 4.43%. This significant increase confirms that the LLM-derived semantic embeddings integrated within our framework effectively compensate for sparse interaction data, thus enabling high-quality recommendations for new or rarely-exposed items. In the general recommendation scenario, LSG-Rec delivers more moderate but consistent improvements, with CTR and clickPV increasing by 0.73% and 0.87%, respectively, while also producing a pronounced reduction in concentration (–1.22%), indicating increased diversity of recommendations.

2) Cross-domain Recall Scenario: We evaluate LSG-Rec in a real-world cross-domain recall scenario: hotspot-triggered content recommendation. In this scenario, trending topics or popular events (hotspots) serve as triggers to retrieve relevant content across different domains.

TABLE IV: A/B testing results for hotspot-triggered cross-domain recall (relative lift over baseline)

Date	**Click PV Lift (%)**	**Click UV Lift (%)**
day1	+11.93%	+6.38%
day2	+25.98%	+14.07%
day3	+4.73%	+6.78%
day4	+14.56%	+13.60%
day5	+1.83%	-0.39%
day6	+6.54%	+11.91%
day7	+19.24%	+16.77%
Average	**+12.12%**	**+9.87%**

We conduct a seven-day A/B test to evaluate LSG-Rec's performance in the hotspot-triggered cross-domain recall scenario. The A/B testing results, summarized in Table IV, demonstrate LSG-Rec's effectiveness in hotspot-triggered cross-domain recall. Over the seven-day testing period, our framework achieved an average lift of +12.12% in Click PV and +9.87% in Click UV compared to the baseline system.

The cross-domain recall capability demonstrated by LSG-Rec is particularly valuable for platforms with diverse content offerings, as it helps users discover relevant items outside their typical interaction patterns, leading to increased user engagement and satisfaction. The experimental results confirm that our framework effectively addresses the domain boundary challenge through its integrated semantic and collaborative approach.

C. Ablation Study on Key Modules

TABLE V: Ablation Study Results on Downstream Ranking Consistency (Top-10)

Variant	Recall@10	HitRate@10
w/o CB2CF module	0.0049	0.0846
w/o GNN block	0.0114	0.1608
Full LSG-Rec	**0.0147**	**0.2032**

To analyze the contributions of the core modules within the LSG-Rec framework, we conduct ablation studies focusing on the CB2CF and GNN components. The experiments are performed on our offline dataset which is mentioned above, with **HitRate@10 and Recall@10 of the downstream ranking model** serving as the primary evaluation metric.

The results of the ablation study are summarized in Table V. Removing the CB2CF module leads to a significant drop in both Recall@10 (0.0049) and HitRate@10 (0.0846), indicating that this module plays a central role in capturing user-item collaborative signals and enhancing alignment with the ranking stage. Removing the GNN block also results in notable performance degradation, with Recall@10 and HitRate@10 falling to 0.0114 and 0.1608, respectively, suggesting that graph-structure modeling further refines the representations and strengthens the recall-ranking consistency. The full LSG-Rec framework achieves the highest Recall@10 (0.0147) and HitRate@10 (0.2032), validating the complementary and synergistic effects of the CB2CF and GNN modules.

V. Conclusion

We present LSG-Rec, a unified recommendation framework integrating LLMs, dual-tower CB2CF module, and GNN. Extensive experiments and production A/B tests demonstrate its performance in cold-start and cross-domain scenarios. Our key contributions include:

1) (1) Significant cold-start improvement with 3.77% CTR and 4.43% clickPV gains through LLM-derived semantic embeddings;
2) (2) Effective cross-domain recall achieving 12.12% Click PV and 9.87% Click UV lifts over seven-day tests;
3) (3) Verified modular synergy where CB2CF aligns semantic-collaborative signals while GNN enhances representation consistency.

LSG-Rec bridges semantic understanding and collaborative filtering, providing a production-ready solution. Future work will extend to multi-modal inputs and sequential recommendation tasks.

VI. Acknowledgment

The authors would like to express their sincere gratitude to Xiaoru Dai and Yiting Wu for their visionary leadership and continuous support. Their insightful suggestions and invaluable guidance have played a crucial role in shaping this research. We are also deeply grateful for the generous allocation of high-performance computing resources provided by the Xiaoru Dai, which has been instrumental in conducting the large-scale experiments presented in this work.

References

[1] 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. 1, 2

[2] G. Chen, J. Cao, H. Li, R. Sun, Y. Jiang, Q. Zhang, J. Lin, and K. Gai, "A multi-modal modeling framework for cold-start short-video recommendation," in *Kuaishou Technology*, 2023. 1

[3] Z. Zhou, L. Zhang, and N. Yang, "Contrastive collaborative filtering for cold-start item recommendation," in *Proceedings of the ACM Web Conference (WWW)*, 2023. 1

[4] J. Kim, E. Kim, C. Kim, K. Yeo, and S. Lee, "Content-based graph reconstruction for cold-start item recommendation," in *Seoul National University*, 2023. 1

[5] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *Proceedings of the 10th international conference on World Wide Web*, 2001, pp. 285–295. 1

[6] M. Volkovs, G. Yu, and T. Poutanen, "Dropoutnet: Addressing cold start in recommender systems," *Advances in neural information processing systems*, vol. 30, 2017. 1

[7] S. Brody and S. Lagziel, "Simrec: Mitigating the cold-start problem in sequential recommendation by integrating item similarity," *arXiv preprint arXiv:2410.22136*, 2024. 1

[8] Z. Zhou, L. Zhang, and N. Yang, "Contrastive collaborative filtering for cold-start item recommendation," in *Proceedings of the ACM web conference 2023*, 2023, pp. 928–937. 1

[9] S. Wang, J. Zhu, K. Wang, and S. Ruan, "Lecdsr: Large language model enhanced cross-domain sequential recommendation," *Information Fusion*, p. 103762, 2025. 1

[10] J. Lu, G. Sun, X. Fang, J. Yang, and W. He, "A contrastive learning framework for dual-target cross-domain recommendation," in *Proceedings of the 31st ACM international conference on multimedia*, 2023, pp. 6332–6339. 1

[11] J. Du, Z. Ye, B. Guo, Z. Yu, and L. Yao, "Distributional domain-invariant preference matching for cross-domain recommendation," in *2023 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2023, pp. 81–90. 1

[12] Y. Cao, S. Hu, Y. Gong, Z. Li, Y. Yang, Q. Liu, and S. Ji, "Gift: Graph-guided feature transfer for cold-start video click-through rate prediction," in *Proceedings of the 31st ACM international conference on information & knowledge management*, 2022, pp. 2964–2973. 2

[13] L. Guo, L. Tang, T. Chen, L. Zhu, Q. V. H. Nguyen, and H. Yin, "Da-gcn: A domain-aware attentive graph convolution network for shared-account cross-domain sequential recommendation," *arXiv preprint arXiv:2105.03300*, 2021. 2

[14] G. Linden, B. Smith, and J. York, "Amazon. com recommendations: Item-to-item collaborative filtering," *IEEE Internet computing*, vol. 7, no. 1, pp. 76–80, 2003. 2

[15] Y. Huang, W. Wang, L. Zhang, and R. Xu, "Sliding spectrum decomposition for diversified recommendation," in *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 2021, pp. 3041–3049. 2

[16] X. Yi, J. Yang, L. Hong, D. Z. Cheng, L. Heldt, A. Kumthekar, Z. Zhao, L. Wei, and E. Chi, "Sampling-bias-corrected neural modeling for large corpus item recommendations," in *Proceedings of the 13th ACM conference on recommender systems*, 2019, pp. 269–277. 3

[17] M. Gutmann and A. Hyvärinen, "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models," in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2010, pp. 297–304. 3

[18] 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*, 2020, pp. 639–648. 4

[19] F. Wu, A. Souza, T. Zhang, C. Fifty, T. Yu, and K. Weinberger, "Simplifying graph convolutional networks," in *International conference on machine learning*. Pmlr, 2019, pp. 6861–6871. 4