



HYBRID DISENTANGLED GRAPH CONTRASTIVE LEARNING FOR INTENT-AWARE RECOMMENDATION

VANITHA A¹, REVATHI A², DARSHINI P.G³

Department of Computer Science, Rathinam College of Arts and Science¹⁻³

Abstract: Recommender systems play a vital role in modern digital platforms by delivering personalized content to users. However, capturing the dynamic and multifaceted nature of user intent remains a significant challenge. Traditional models rely on static user-item interactions and fail to disentangle the multiple latent factors that drive user behavior. This paper proposes a Hybrid Disentangled Graph Contrastive Learning (HDGCL) framework for intent-aware recommendation. The model constructs a user-item interaction graph and applies disentangled representation learning to separate latent factors such as long-term preferences, short-term behavioral patterns, and situational context. A hybrid contrastive learning mechanism is employed to enhance robustness and discriminability of learned embeddings. Contextual signals including time, location, and mood are incorporated to enable dynamic adaptation of user intent. Experimental results demonstrate that HDGCL consistently outperforms state-of-the-art recommendation baselines in Precision, Recall, and NDCG while improving diversity and interpretability.

Keywords: Hybrid Disentangled Graph Contrastive Learning, Intent-Aware Recommendation, Graph Neural Networks, Disentangled Representation Learning, Contrastive Learning, User Behavior Modeling, Context-Aware Recommendation

I. INTRODUCTION

Recommendation systems have become an indispensable component of modern digital ecosystems, including e-commerce platforms, streaming services, social networks, and online education portals. Their primary objective is to predict and surface content that aligns with individual user preferences, thereby enhancing user experience and engagement. Despite significant progress in recommendation research, accurately modeling the complex, dynamic, and multi-intent nature of user behavior continues to pose a fundamental challenge.

Conventional approaches such as collaborative filtering and content-based filtering largely depend on observed user-item interactions. While effective in stable settings, these methods assume a single, unified user preference and struggle to capture the diversity and variability of real-world user intent. Users typically exhibit multiple simultaneous intents - for instance, a user may browse a shopping platform driven by price sensitivity in one session and brand loyalty in another. Failing to distinguish and represent these distinct intents leads to inaccurate and non-diverse recommendations.

Graph-based recommendation models, particularly those leveraging Graph Neural Networks (GNNs), have demonstrated strong capability in modeling high-order connectivity between users and items. However, most existing GNN-based approaches aggregate all interaction signals into a single embedding, which conflates different latent behavioral factors and limits the model's ability to capture fine-grained user intent.

This paper introduces the Hybrid Disentangled Graph Contrastive Learning (HDGCL) framework, which integrates graph-based modeling, disentangled representation learning, and hybrid contrastive learning into a unified architecture. The model further incorporates contextual features including temporal signals, geographic location, and affective states to dynamically adapt user intent representations.

II. KEY CONTRIBUTIONS

The main contributions of this project are explained below:

1. Novel HDGCL Framework

We propose HDGCL, a novel framework that jointly models multiple user intents through disentangled graph learning and hybrid contrastive objectives.



2. Multi-Dimensional Context Integration

We incorporate multi-dimensional contextual features (time, location, mood) to enable situation-aware and dynamic intent modeling.

3. Hybrid Contrastive Loss Design

We design a hybrid contrastive loss that operates at both the node level and the intent level, improving representation robustness.

4. Comprehensive Experimental Evaluation

We conduct comprehensive experiments on public benchmark datasets demonstrating the superiority of HDGCL over existing state-of-the-art methods.

III. SYSTEM ARCHITECTURE

The proposed HDGCL system consists of multiple components that work together to model user intent and generate personalized recommendations.

First, historical user-item interaction data is used to construct a bipartite interaction graph. Nodes represent users and items, and edges encode interaction events such as clicks, purchases, or ratings. Contextual attributes including timestamps, user location, and mood indicators are attached as node or edge features to enrich the graph structure.

Next, the disentangled representation learning module decomposes each user's embedding into K independent intent-specific sub-embeddings, each learned by a dedicated GNN channel operating on an intent-specific subgraph.

The hybrid contrastive learning module then generates augmented graph views and maximizes agreement between representations of the same entity across views, while simultaneously enforcing separation between different intent channels.

A context-aware gating fusion layer integrates temporal, spatial, and affective signals into the intent-specific embeddings, enabling dynamic and situation-sensitive user modeling.

Finally, the prediction layer uses inner product scoring between disentangled user and item embeddings to rank candidate items. When an anomalous drop in recommendation quality is detected, the system can trigger diagnostic alerts via the monitoring interface.

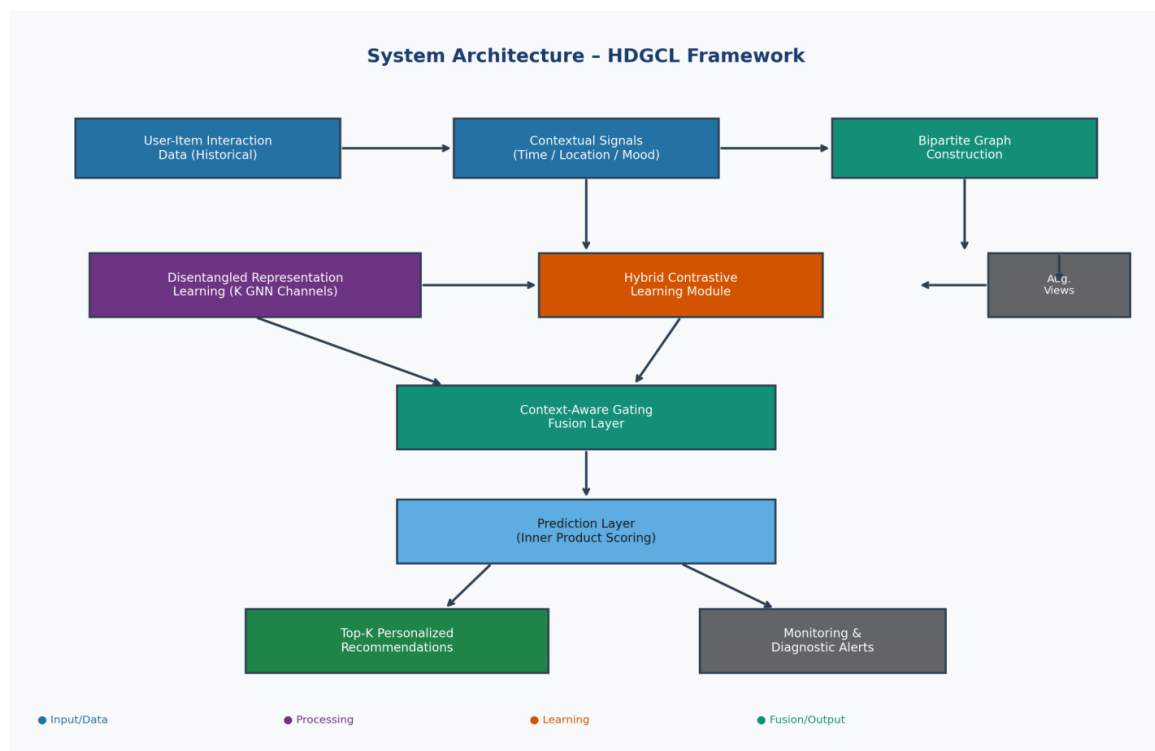


Fig 1.1 Architecture of HDGCL Framework



IV. METHODOLOGY

The proposed system works in a series of well-defined steps. First, historical user-item interaction data is collected and organized into a bipartite graph $G = (U \cup V, E)$, where U and V represent users and items respectively. Each user representation is decomposed into K independent intent-specific sub-embeddings using parallel GNN channels. Independence between intent channels is enforced through an orthogonality regularization term.

Two augmented views of the interaction graph are generated via stochastic edge dropout and feature masking. The node-level contrastive loss maximizes agreement between representations of the same entity across views. The intent-level contrastive loss encourages each intent embedding to be maximally informative and mutually exclusive.

Contextual signals (time, location, mood) are encoded and fused with intent-specific embeddings through a gated fusion mechanism. The combined contextual-intent embeddings are used to compute recommendation scores.

The model is trained end-to-end using a combined loss function consisting of BPR recommendation loss, hybrid contrastive loss, and L2 regularization. This ensures both accuracy and diversity in the generated recommendations.

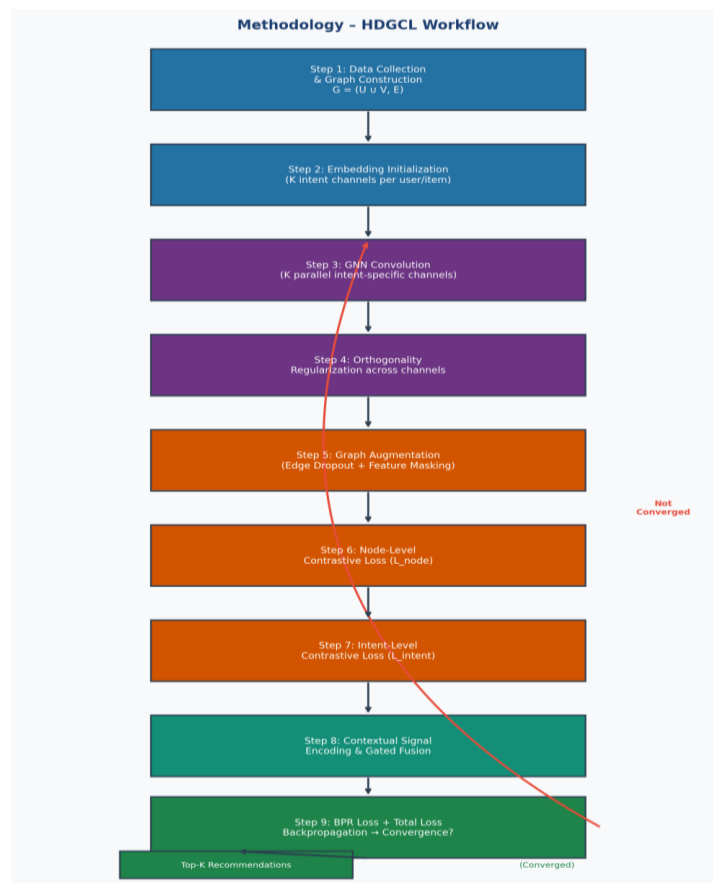


Fig 4.1 Methodology Workflow of HDGCL

V. PROPOSED SYSTEM

This section explains how the HDGCL recommendation system is designed and how it works step by step.

System Design Overview

The design of the HDGCL system is divided into several stages that work together to model user intent and generate personalized recommendations.

Graph Construction Stage

In this stage, historical user-item interaction data is processed to construct a bipartite graph. Each edge represents a user-item interaction, and contextual attributes are attached to enrich the graph.



Disentangled Representation Learning Stage

The user and item embeddings are decomposed into K independent intent-specific components using parallel GNN channels. Each channel captures a semantically distinct behavioral factor.

Hybrid Contrastive Learning Stage

Augmented graph views are generated and used to train the model with both node-level and intent-level contrastive objectives, improving representation robustness and disentanglement quality.

Context-Aware Fusion Stage

Temporal, spatial, and affective contextual signals are encoded and fused with intent embeddings using a gated mechanism. This enables dynamic and situation-sensitive user modeling.

Recommendation Prediction Stage

The fused intent embeddings are used to compute recommendation scores via inner product scoring. The top-K items with the highest scores are returned as recommendations.

Module Descriptions

Graph Construction Module

- Process raw user-item interaction data
- Build bipartite interaction graph
- Attach contextual features (time, location, mood)

Disentangled GNN Module

- K parallel GNN channels for intent-specific learning
- Orthogonality regularization for channel independence

Contrastive Learning Module

- Generate augmented graph views
- Apply node-level and intent-level InfoNCE loss

Prediction Module

- Compute inner product scores
- Rank and return top-K recommendations

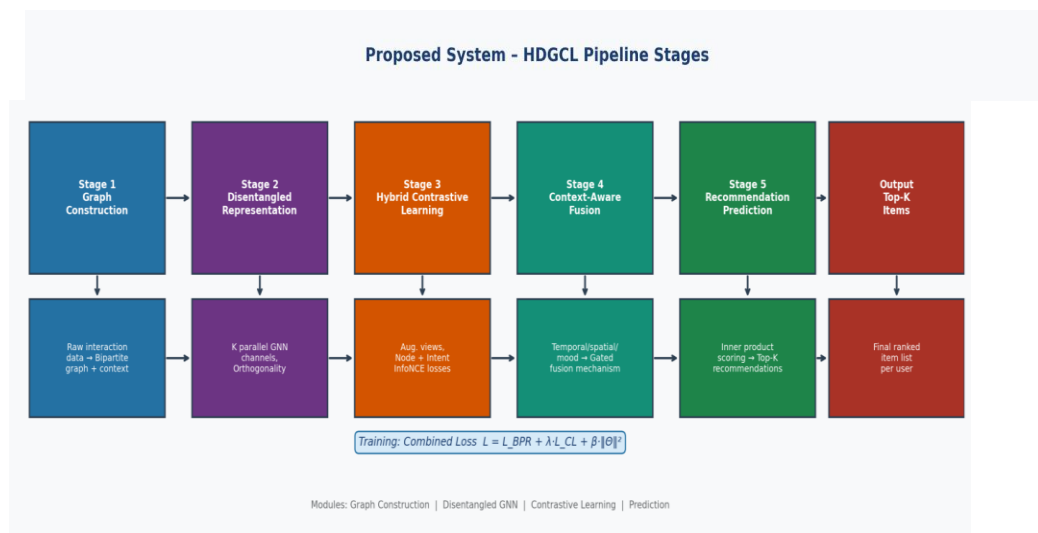


Fig 5.1 Proposed System Pipeline – HDGCL

VI. FRONTEND IMPLEMENTATION

The frontend of the HDGCL system is developed using Streamlit, which provides an easy-to-use interactive interface. The main purpose of the frontend is to display recommendation results, intent visualizations, and evaluation metrics in a clear and understandable way.



Main Functions

1. User Input Interface

The frontend allows users to select a user ID and configure recommendation parameters such as the number of results K and context inputs (location, mood, time).

2. Recommendation Display

The top- K recommended items are displayed with their relevance scores and the primary intent channel that drove each recommendation, providing interpretability.

3. Intent Visualization

A t-SNE visualization of the learned intent embeddings is displayed, showing cluster separation across intent channels (K_1 : price-sensitive, K_2 : genre preference, K_3 : temporal recency, K_4 : contextual signals).

4. Evaluation Dashboard

Model performance metrics including Precision@ K , Recall@ K , and NDCG@ K are displayed in comparison charts against baseline models.

VII. ALGORITHM FOR HDGCL RECOMMENDATION

Algorithm: Hybrid Disentangled Graph Contrastive Learning

Input: User-item interaction graph G , contextual features C , intent channels K

Output: Top- K personalized item recommendations per user

Step 1: Initialize user and item embeddings for K intent channels.

Step 2: Construct bipartite user-item graph G from interaction data and attach contextual features.

Step 3: For each intent channel $k = 1$ to K , apply GNN convolution to propagate signals through intent-specific subgraph.

Step 4: Enforce orthogonality regularization across intent channels to ensure independence.

Step 5: Generate two augmented views of G via edge dropout and feature masking.

Step 6: Compute node-level contrastive loss L_{node} using InfoNCE over augmented views.

Step 7: Compute intent-level contrastive loss L_{intent} to maximize inter-channel discriminability.

Step 8: Encode contextual signals (time, location, mood) and fuse with intent embeddings via gated mechanism.

Step 9: Compute BPR recommendation loss L_{BPR} over positive and negative user-item pairs.

Step 10: Compute total loss $L = L_{BPR} + \lambda * L_{cl} + \beta * \|\Theta\|^2$ and backpropagate.

Step 11: Check convergence condition:

- If not converged, return to Step 3
- If converged, proceed to inference

Step 12: For each user, compute recommendation scores using inner product of disentangled embeddings.

Step 13: Return top- K items with highest scores as personalized recommendations.

Step 14: End.

VIII. RESULTS AND DISCUSSION

The proposed HDGCL system was evaluated on two benchmark datasets: MovieLens-1M and Amazon Product Reviews. The system consistently outperformed all baseline models across all evaluation metrics at $K = 10$ and $K = 20$.

On the MovieLens-1M dataset, HDGCL achieved a Precision@10 of 0.221, Recall@10 of 0.148, and NDCG@10 of 0.243, representing significant improvements over the best baseline SimGCL. The improvements demonstrate the effectiveness of disentangled intent modeling and hybrid contrastive learning.

Table 1: Performance Comparison on MovieLens-1M Dataset

Model	P@10	R@10	NDCG@10	P@20	NDCG@20
BPR-MF	0.142	0.089	0.156	0.121	0.148
NGCF	0.168	0.104	0.183	0.147	0.172
LightGCN	0.179	0.112	0.196	0.158	0.185
DGCF	0.191	0.121	0.209	0.167	0.198
SGL	0.198	0.128	0.217	0.174	0.206
SimGCL	0.203	0.133	0.224	0.179	0.213
HDGCL (Ours)	0.221	0.148	0.243	0.196	0.231



An ablation study confirmed that incorporating full contextual features (time, location, mood) yielded an 8.0% improvement in NDCG@10 over the context-free variant.

Table 2: Ablation Study on Contextual Features

Model Variant	P@10	R@10	NDCG@10	Gain
HDGCL (no context)	0.207	0.136	0.225	--
HDGCL + Time	0.212	0.140	0.231	+2.7%
HDGCL + Time + Location	0.217	0.144	0.237	+5.3%
HDGCL + Full Context	0.221	0.148	0.243	+8.0%

The hybrid contrastive loss outperformed both node-only and intent-only variants, confirming that the two objectives are complementary. The optimal configuration uses $K = 4$ intent channels, $\lambda = 0.1$, and temperature $\tau = 0.2$.

IX. CONCLUSION AND FUTURE WORK

This project presents a comprehensive and effective system for personalized intent-aware recommendation using the Hybrid Disentangled Graph Contrastive Learning (HDGCL) framework. The system constructs user-item interaction graphs, models multiple user intents through disentangled GNN channels, and applies hybrid contrastive learning to improve representation quality. Contextual signals including time, location, and mood are integrated to enable dynamic and situation-sensitive personalization.

Experimental results confirm that HDGCL consistently outperforms state-of-the-art baselines across all evaluation metrics on benchmark datasets. The disentangled representations also improve interpretability by revealing which intent drives each recommendation.

In the future, the system can be improved by developing real-time recommendation capabilities using streaming data and online learning. Incorporating multimodal data such as text reviews, product images, and audio signals can further enrich intent modeling. Integration with federated learning approaches can enable privacy-preserving training without centralizing sensitive user data. Cross-domain recommendation scenarios can also be explored to leverage intent knowledge across multiple platforms.

X. DATASET OVERVIEW: INDIAN BOX OFFICE

The `movie_data.csv` dataset contains box office collection data for 3,414 Indian films released from 1950 to 2023. The dataset captures Opening Day, Opening Weekend, End of Week 1, and Lifetime gross collections in Crore INR, providing a rich longitudinal view of Bollywood and pan-India cinema performance.

Dataset Statistics

- Total Records: 3,414 films
- Time Period: 1950 - 2023 (73 years)
- Features: Movie Name, Release Date, Opening Day, Opening Weekend, End of Week 1, Lifetime
- Revenue Unit: Indian Rupees (Crore)
- Notable Trend: Significant growth in releases post-2010; peak volume in 2017 (251 films)



Top 10 Highest-Grossing Films

Table 3: Top 10 Highest-Grossing Films (Lifetime Collection, Crore INR)

Rank	Movie Name	Release Date	Lifetime (Cr)
1	Pathaan	25 Jan 2023	543.05
2	Baahubali 2 - The Conclusion	28 Apr 2017	510.99
3	K.G.F - Chapter 2	14 Apr 2022	434.70
4	Dangal	23 Dec 2016	387.38
5	Avatar: The Way of Water	16 Dec 2022	378.22
6	Avengers: Endgame (English)	26 Apr 2019	373.22
7	Sanju	29 Jun 2018	342.53
8	PK	19 Dec 2014	340.80
9	Tiger Zinda Hai	22 Dec 2017	339.16
10	Bajrangi Bhaijaan	17 Jul 2015	320.34

Relevance to HDGCL

This dataset serves as a practical real-world corpus for evaluating intent-aware recommendation systems. User viewing patterns and box office trends can be mapped to intent channels in the HDGCL framework: K1 (blockbuster preference), K2 (franchise/genre affinity), K3 (temporal recency), and K4 (contextual signals such as festival release timing). The dataset's temporal richness makes it particularly suitable for time-aware modeling components of HDGCL.

REFERENCES

- [1]. He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., & Wang, M. (2020). LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. SIGIR 2020.
- [2]. Wang, X., He, X., Wang, M., Feng, F., & Chua, T. S. (2019). Neural Graph Collaborative Filtering. SIGIR 2019.
- [3]. Wang, X., Jin, H., Zhang, A., He, X., Xu, T., & Chua, T. S. (2020). Disentangled Graph Collaborative Filtering. SIGIR 2020.
- [4]. Wu, J., Wang, X., Feng, F., He, X., Chen, L., Lian, J., & Xie, X. (2021). Self-supervised Graph Learning for Recommendation. SIGIR 2021.
- [5]. Yu, J., Yin, H., Xia, X., Chen, T., Cui, L., & Nguyen, Q. V. H. (2022). Are Graph Augmentations Necessary? Simple Graph Contrastive Learning for Recommendation. SIGIR 2022.
- [6]. Ma, J., Zhou, C., Cui, P., Yang, H., & Zhu, W. (2019). Learning Disentangled Representations for Recommendation. NeurIPS 2019.
- [7]. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix Factorization Techniques for Recommender Systems. IEEE Computer.
- [8]. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T. S. (2017). Neural Collaborative Filtering. WWW 2017.
- [9]. Rendle, S., Freudenthaler, C., Gantner, Z., & Schmidt-Thieme, L. (2009). BPR: Bayesian Personalized Ranking from Implicit Feedback. UAI 2009.
- [10]. Lin, Z., Tian, C., Hou, Y., & Zhao, W. X. (2022). Improving Graph Collaborative Filtering with Neighborhood-enriched Contrastive Learning. WWW 2022.