

GeMi: A Graph-based, Multimodal Recommendation System for Narrative Scroll Paintings

Haimonti Dutta^{1,2*}, Pruthvi Moluguri³, Jin Dai¹,
Saurabh Amarnath Mahindre⁴

¹Department of Management Science and Systems, The State University
of New York, Jacobs Management Center, Buffalo, 14260, NY, USA.

²Institute for Artificial Intelligence and Data Science, The State
University of New York, Lockwood Hall, Buffalo, 14260, NY, USA.

³Department of Computer Science and Engineering, The State
University of New York, Davis Hall, Buffalo, 14260, NY, USA.

⁴Technology, Information, Internet, eBay*, Embassy Tech Village,
Bengaluru, 560103, Karnataka, India.

*Corresponding author(s). E-mail(s): haimonti@buffalo.edu;
Contributing authors: pmolugur@buffalo.edu; jindai@buffalo.edu;
smahindr@buffalo.edu;

Abstract

Recommendation Systems are effective in managing the ever-increasing amount of multimodal data available today and help users discover interesting new items. These systems can handle various media types such as images, text, audio, and video data, and this has made it possible to handle content-based recommendation utilizing features extracted from items while also incorporating user preferences. Graph Neural Network (GNN)-based recommendation systems are a special class of recommendation systems that can handle relationships between items and users, making them particularly attractive for content-based recommendations. Their popularity also stems from the fact that they use advanced machine learning techniques, such as deep learning on graph-structured data, to exploit user-to-item interactions. The nodes in the graph can access higher-order neighbor information along with state-of-the-art vision-language models for processing multimodal content, and there are well-designed algorithms for

*This work was done prior to the author joining his current affiliation.

embedding, message passing, and propagation. In this work, we present the design of a GNN-based recommendation system on a novel data set collected from field research. Designed for an endangered *performing* art form, the recommendation system uses multimodal content (text and image data) to suggest similar paintings for viewing and purchase. To the best of our knowledge, there is no recommendation system designed for narrative scroll paintings – our work therefore serves several purposes, including art conservation, a data storage system for endangered art objects, and a state-of-the-art recommendation system that leverages both the novel characteristics of the data and preferences of the user population interested in narrative scroll paintings.

Keywords: Graph Neural Networks, content-based recommendation, multimodal vision language models, folk art recommendation

1 Introduction

Recommendation Systems (RS) (Deldjoo et al. (2020), Zhou et al. (2023a)) help organize a large number of items (such as books, music, appliances) available to users. They find items of interest by predicting the preference for candidate items and recommending the most pertinent ones using historical data. Data-driven RS falls into three main categories: collaborative filtering (CF), Content-based Filtering (CBF), and Hybrid. In CF models, the emphasis is on a community of users, their behavioral patterns and preferences. The CBF methods depend on the user’s behavioral data along with information about the items, such as words in a text document, color in images, and audio content. Some recommendation systems are known to combine collaborative filtering methods with content-based methods for added benefits. All three categories of recommendation systems use multimodal data (images, text, audio, video) extensively.

Conventional methods for integrating data from different modalities in RS include concatenation or summation (He and McAuley (2016); Liu et al. (2017)). However, with the advent of deep learning techniques with non-linear modeling capabilities, each modality can be represented separately and combined subsequently to generate a suitable representation. The recent increase in interest in graph-based recommendation systems has inspired a new line of work that explores higher-order interactions between users and items by leveraging the power of graph neural networks (He et al. (2020); Wu et al. (2023); Zhang et al. (2021a, 2022); Zhou et al. (2023c,b); Wei et al. (2019, 2020); Wang et al. (2023a); Zhou et al. (2023b)). For example, MMGCN (Wei et al. (2019)) exploits user-item interactions to guide representation learning in each modality by constructing graph convolutional networks with nodes having topological features and aggregating features of neighbors. GRCN (Wei et al. (2020)) refines this technique by designing a graph refining layer that is designed to identify noisy edges with high confidence in user preference. However, user and item representations learned by graph aggregations are often inherently noisy, since user-item interactions are sparse. These systems also suffer from cold start problems, where there is not enough data (such as clicks, ratings, or purchases) to learn preferences.

One solution to the user-item sparsity problem is to reinforce node (users and items) features. However, the noise present in the nodes often makes it difficult to extract useful information for recommendations. Noise in item features (such as blurry images, errors in spelling in text, damaged soundtracks, or misleading descriptions) causes uncertainty in representation. The item features may also contain irrelevant information, such as background color or distracting objects in the image. In multimodal settings, complementary information from different modalities is leveraged to deal with sparsity, but noisy features can lead to undesirable results. If certain modalities are missing or absent, the multimodal system must accommodate this. To address these challenges in multimodal learning, Variational Auto Encoders (VAEs) have become popular (Yi and Chen (2022); Liang et al. (2018); Ma et al. (2019); Li and She (2017)) in which modality-specific encoders learn a Gaussian variable parameterized by a mean and variance vector for each node. The mean vector is modeled to represent semantic information, and the variance vector captures the uncertainty of the corresponding modality.

In this paper, we describe a multimodal content-based recommendation system augmented with user preference information, that is designed for a unique, sociocultural application to conserve and study an endangered performing art from eastern India. *Singing painters* are a group of itinerant storytellers who paint scrolls and use a song or poem to explain the artwork. The art is usually inspired by epics, regional folktales, and events that transpire in the day-to-day lives of the painters (McCutchion and Bhowmick (1999)). Songs or poems are typically learned by word-of-mouth. Sporadic, ad hoc efforts have been made to store samples of artwork or song texts in non-digitized formats, but to the best of our knowledge, a meticulous study of this multimodal data from images and text, or the design of a recommendation system to organize it, has never been undertaken. In recent work (Dutta and Tiwari (2024)), we have prepared a custom dataset of digitized images, aligning lyrics of songs with images, and documenting, to a good extent, regional variations in style, theme, color choice, and other relevant details. We use an enhanced version of this dataset¹ to design the recommendation engine, which can suggest artwork similar to what appreciative users like. In doing so, it was found that while there have been many techniques for the construction of multimodal recommendation systems for known benchmark datasets (such as Amazon (Ni et al. (2019)), MovieLens (Harper and Konstan (2015)) or Netflix (Bennett and Lanning (2007))), the application of these to a novel data set can quickly become tedious if careful analysis and principled experimentation are not adhered to. We report our findings on the design of the multimodal recommendation system, GeMi, carefully benchmarking it with state-of-the-art systems found in the literature. The dataset used in this research is available from: <https://github.com/Haimonti/Folk-Art-Recommendation-System>.

The primary contributions of this paper are:

- A unique, multimodal dataset collected from field surveys and curated for research
- The GeMi Recommendation System - A content-based recommendation using state-of-the-art vision-language models with sophisticated multimodal fusion, graph

¹More data is available from subsequent field studies in other districts of the state.

structure learning, and incorporating user preferences through collaborative filtering style plug-and-play modules.

This paper is organized as follows: Section 2 presents Related Work; the data collection procedure along with non-digitized sources are described in Section 3; the recommendation system is presented in Section 4; empirical results comparing the performance of GeMi to state-of-the-art transductive and inductive homogeneous and heterogeneous graphs are presented in Section 5; discussion of enhancements possible is proposed in Section 6, and Section 7 concludes the paper.

2 Related Work

In this section, we present prior work on art recommendation and a brief history of narrative scroll paintings. Appendix A provides a review of the use of multimodal data, Variational Auto Encoders (VAEs), and Large Language Models (LLMs) to augment Graph Neural Networks (GNNs) for building recommendation systems.

2.1 Art Recommendation

Museums have recognized the need to provide visitors with personalized experiences when visiting collections and have adopted recommendation systems to meet this requirement. For example, a visitor’s profile might store the names of their favorite artists or their painting techniques, extracted from short textual descriptions associated with the artworks (Wang et al. (2008, 2009)). The user profile is then matched to the features of the items to provide personalized suggestions. The Cultural Heritage Information Personalization (CHIP) project (Aroyo et al. (2007)) built a recommendation engine for the Rijks museum in Netherlands. In the Folksonomy-based Item Recommender syStem (FIRSt) (Semeraro et al. (2012)), user-generated content from social tagging is integrated into a content-based recommender, allowing users to express item preferences by entering a numerical rating in addition to annotating items with free tags. In the Personal Experience with Active Cultural Heritage (PEACH) project (Stock et al. (2007); Messina et al. (2019); Kuflik et al. (2011)), a PDA-based museum tour application is developed, which is location-aware and whose content is adapted to the needs of the visitor. Approaches that make suggestions to museum visitors based on their history within the physical environment and textual information associated with each item in their history have become fairly popular (Bohnert and Zukerman (2009); Grieser et al. (2011)). Yilma et al. (Yilma and Leiva (2023)) focus on efficiently capturing the latent semantic relationships of visual art for personalized recommendation. They show that textual features are as useful as visual ones, but a fusion of both captures most of the latent semantic information. Kuflik et al. (Kuflik et al. (2011)) showed the benefits of using graph-based recommendation. An anti-recommendation system, entitled “Art I don’t like” (Frost et al. (2019)), suggests artwork dissimilar to the ones users like, thereby suggesting that removing access to opposing viewpoints can lead to *filter bubbles* (Pariser (2011)) and intellectual isolation. However, to the best of our knowledge, an art recommendation system solely designed for narrative scroll paintings does not exist to date.

2.2 A brief history of the dissemination of narrative scroll paintings in South Asia

Narrative scroll paintings are a popular form of *performative* art in Eastern India. It incorporates several modes of communication—visual messages, oral traditions, and music — which blend to describe social transformations, stories of migrations, and socio-political and religious reflections (Korom (2006); Bajpai (2015)). Although the exact date of origin of this profession is not known, ancient lore and traditions indicate that it dates back to the 5th-6th century (possibly even earlier) (Mair (2019)).

These paintings from eastern India have spread to other parts of the country, including the western and southern regions (such as Deccani picture showmen, the “Paithan” paintings of Chitrakathis in Maharashtra, the “Garoda” showmen of Gujarat, the narrative paintings of Telengana and the Bhopas of Rajasthan) (Jain (1998)). Furthermore, confronted by foreign invasions in the east, many scroll paintings have found their way to neighboring countries such as Nepal, Tibet, and other regions along the Silk Road. These traditions were also used to spread religious learning and made their way into China. It is also known that the technique of picture recitation made its way to the western heartland of the Turkish-speaking world by the end of the sixteenth century and possibly much earlier. Evliyā Efendi (ca. 1611–1660) provides a fascinating description of a group known as “Painter Fortunetellers” in his account of life in Istanbul. It must also be noted that in Indonesia, dramatic performances were well known as *wayang* - a transformative realization used for religious instruction. *Wayang* means “shadow,” and thus picture scrolls, puppets, shadows, and dancers are all technically *wayang*. Quoting Mair (Mair (2019)), “Indian inputs were essential for the formation of (Chinese) *pien*, and it is even more obvious that *wayang* owes an enormous debt to Indian storytelling and theater.” Unfortunately, very little of this art form has been documented or digitized.

3 Data

In recent work, the first author undertook extensive field surveys in eastern India² which has allowed digitization of the scrolls, video, and audio recordings of performances, followed by transcription, transliteration, and translation (original songs are in Bengali) of lyrics of songs. The data from this field study was used to design the first version of the content-based recommendation system (Dutta and Tiwari (2024)) for art-lovers and curators. A second field survey was conducted a year later, helping to augment the data in the repository. It must be noted that while the images and text collected from the first survey were well aligned, those from the second were not. The missing data (either the text of the song was not available or there was no painting) led to several challenges in the alignment of multimodal data. The data collected remain the copyright of the artists and song creators, but the Institutional Repository located at the State University of New York at Buffalo has obtained permission and consent to archive and store it.

²This involved surveying villages in the districts of Birbhum, Bankura and Purulia in West Bengal.

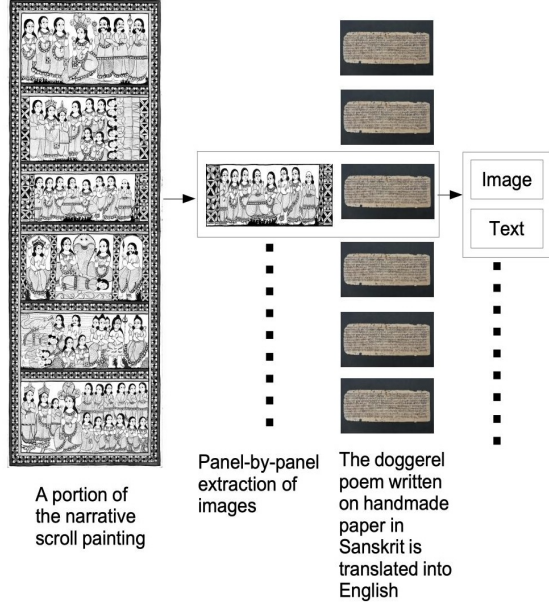


Fig. 1: An illustration of a panel-by-panel extraction of images from a narrative scroll painting and alignment with the text of an ancient doggerel poem written on handmade paper using black ink. Many parts of the original manuscript of the poem are missing or lost.

3.1 Existing non-digitized sources of data

- **Text:** There is no known database of songs that can be leveraged for the design of a recommendation system. Therefore, the process of collecting songs and their lyrics had to be done manually through field studies. However, we are aware of several sources where some preliminary attempts have been made to collect the songs: (1) Books ([McCutchion and Bhowmick \(1999\)](#), [Dutta \(1939\)](#), [Dimock \(1963\)](#)) and theses ([Yazijian \(2007\)](#)) (2) Personal collections, Dr Korom ([Korom \(2006\)](#)) – these are available on tapes and currently not digitized.
- **Images:** Existing non-digitized narrative scroll paintings are found in the collections of the British Museum, the Museum of International Folk Art (Santa Fe, New Mexico), Victoria and Albert Museum (UK), and the Gurusaday Museum (Kolkata, India). However, it must be noted that the process of alignment of images with songs is non-trivial in most cases, since sporadic conservation efforts have varying objectives and may not contain all the relevant data necessary for our study.

4 The Recommendation System

Figure 1 provides an illustration of the multimodal data extraction procedure and Figure 2 represents the overall architecture of the recommendation system. It has the following components: (a) Fine-tuned vision language models for multimodal feature extraction (b) Graph Structure Learning (c) Node Classification and, (d) downstream content-based recommendation. The following sections describe them in more detail.

Preliminaries Our basic premise is that a recommendation of panels and songs should reflect the concept(s) a user is interested in. We study the user-panel preference matrices to develop an algorithm for recommendation. Let \mathcal{U}, \mathcal{P} denote the set of users and panels, respectively. Each user $u_i \in \mathcal{U}$ is associated with a set of panels $p_j \in \mathcal{P}, i \in 1, \dots, |\mathcal{U}|; j \in 1, \dots, |\mathcal{P}|$ which s/he likes. Figure 1 shows a schematic of the panel-by-panel image extraction. The preference score $y_{ij} = 1$, if the user u_i likes panel p_j and 0 otherwise. Each panel p_j has a multimodal representation i.e. $p_j \in \mathbb{R}^{D_m}$ where D_m represents the dimensionality of modality $m \in \mathcal{M}$ and \mathcal{M} is the set of modalities. The multimodal recommendation seeks to accurately find panels a user might like, given his/her known preferences.

In the following subsections, we describe each component of the GeMi recommendation system.

4.1 Fine-tuned vision language models for multimodal feature extraction

We describe a multimodal feature extraction framework that transforms image and text associated with narrative scroll paintings into semantically structured, high-dimensional representations suitable for downstream graph-based recommendation.

Given the symbolic density, narrative continuity, and cultural specificity inherent in them, unimodal representations are insufficient to capture the full semantic spectrum of the data. Instead, we adopt a modular multimodal encoding strategy that integrates complementary inductive biases from language modeling, cross-modal contrastive learning, and probabilistic latent-variable modeling.

Formally, let each panel be represented by a tuple $\mathcal{P}_i = (x_i^{\text{img}}, x_i^{\text{text}})$, where $x_i^{\text{img}} \in \mathcal{X}_{\text{img}}$ denotes the visual depiction of the panel and $x_i^{\text{text}} \in \mathcal{X}_{\text{text}}$ denotes its associated textual description. Our objective is to construct a family of embedding functions

$$\Phi^{(k)} : \mathcal{X}_{\text{img}} \times \mathcal{X}_{\text{text}} \rightarrow \mathbb{R}^{d_k}, \quad (1)$$

each capturing distinct semantic views of the same panel content, which are later integrated through graph-based propagation. To this end, we employ three conceptually distinct yet complementary representation mechanisms:

1. A Large Language Model (LLM) (LLaMA, [Touvron et al. \(2023\)](#)) for semantic abstraction and canonicalization of textual descriptions.
2. A sigmoid-Contrastive Language Image Pretraining model (Sig-CLIP, [Radford et al. \(2021\)](#)) for joint image-text alignment,

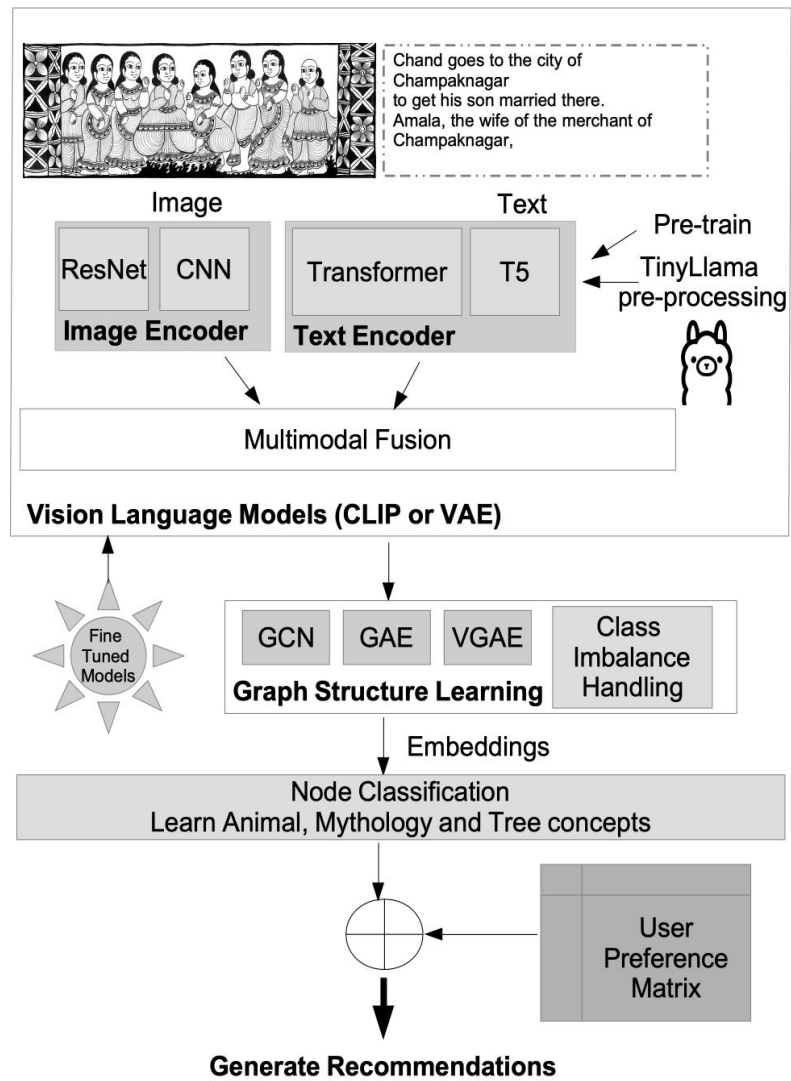


Fig. 2: The GeMi recommendation system. Fine-tuned Vision Language models are used to extract features, learn latent graph structures, and embeddings using graph convolutions. User-Panel preference matrices are used to augment recommendations.

3. A multimodal variational autoencoder (VAE) for probabilistic fusion of visual and linguistic information.

4.1.1 LLM-Based Semantic Text Canonicalization

Textual descriptions associated with narrative scroll paintings are often noisy, stylistically heterogeneous, and semantically under-specified. To reduce lexical variance while preserving high-level semantic intent, we employ a Large Language Model as a deterministic semantic canonicalization operator that rewrites raw text into a compact, concept-centric representation before multimodal fusion.

We adopt the LLaMA family of autoregressive decoder-only transformers, which model a token sequence $\mathbf{x} = (x_1, \dots, x_T)$ via causal language modeling: $P(\mathbf{x}) = \prod_{t=1}^T P(x_t \mid x_1, \dots, x_{t-1})$, implemented using multi-head causal self-attention. LLaMA follows a pre-normalized transformer architecture with RMSNorm, rotary positional embeddings (RoPE, Su et al. (2024)), and SwiGLU feed-forward blocks (Shazeer (2020)), yielding improved stability and semantic expressivity under constrained inference budgets. We adopt TinyLlama (Zhang et al. (2024)) specifically because its compact parameter count (1.1B) permits inference within the memory constraints of our compute environments without requiring quantization, while its LLaMA-2 architecture retains sufficient semantic expressivity for the canonicalization task, which requires entity recognition and lexical normalization rather than open-ended generation. Let $\mathbf{h}^{(0)}$ denote the input token embeddings augmented with RoPE. Each transformer layer $\ell \in \{1, \dots, L\}$ computes $\mathbf{h}^{(\ell)} = \text{Block}^{(\ell)}(\mathbf{h}^{(\ell-1)})$, and the final hidden states parameterize a categorical distribution over the vocabulary via a linear projection and softmax.

In our pipeline, the language model is not used as a direct feature encoder. Instead, we formalize it as a deterministic rewriting function $\mathcal{R}_{\text{LLaMA}} : \mathcal{X}_{\text{text}} \rightarrow \tilde{\mathcal{X}}_{\text{text}}$, which maps raw panel descriptions to semantically normalized textual forms that emphasize entities, actions, and symbolic motifs while suppressing surface-level linguistic noise. For each panel i , the rewritten text is given by $\tilde{x}_i^{\text{text}} = \mathcal{R}_{\text{LLaMA}}(x_i^{\text{text}})$.

The canonicalized text $\tilde{x}_i^{\text{text}}$ is subsequently passed on to downstream multimodal encoders (SigCLIP and multimodal VAE), where it serves as the linguistic input for contrastive alignment and latent fusion. By enforcing semantic consistency at the textual level, this rewriting stage reduces cross-modal noise and improves the quality of learned image-text representations without altering the downstream encoder architectures.

4.1.2 Contrastive Language Image Pretraining (CLIP) models and their variants

Following semantic canonicalization of textual descriptions, we learn a joint image-text embedding space via contrastive vision-language pretraining. The goal is to embed visual panels and their associated texts into a shared latent space such that semantically aligned image-text pairs are geometrically proximate, while mismatched pairs are separated. Let $\mathcal{D} = \{(I_i, T_i)\}_{i=1}^N$ denote a corpus of paired images $I_i \in \mathcal{I}$ and texts $T_i \in \mathcal{T}$. We learn dual encoders $f_\theta : \mathcal{I} \rightarrow \mathbb{R}^d$, $g_\phi : \mathcal{T} \rightarrow \mathbb{R}^d$, which map each modality into a common d -dimensional embedding space. The architecture follows a dual-tower design: encoders are trained jointly but operate independently at inference time.

Given a minibatch $\mathcal{B} = \{(I_i, T_i)\}_{i=1}^B$, we obtain ℓ_2 -normalized embeddings $\mathbf{x}_i = \frac{f_\theta(I_i)}{\|f_\theta(I_i)\|_2}$, $\mathbf{y}_i = \frac{g_\phi(T_i)}{\|g_\phi(T_i)\|_2}$, $\mathbf{x}_i, \mathbf{y}_i \in \mathbb{R}^d$. Pairwise image–text similarity is measured by cosine similarity, $s_{ij} = \mathbf{x}_i^\top \mathbf{y}_j$, scaled by a temperature parameter $t > 0$, $\ell_{ij} = t s_{ij}$. The CLIP objective treats image-to-text and text-to-image matching as symmetric classification problems. Define $p_{ij}^{I \rightarrow T} = \frac{\exp(\ell_{ij})}{\sum_{k=1}^B \exp(\ell_{ik})}$, $p_{ij}^{T \rightarrow I} = \frac{\exp(\ell_{ij})}{\sum_{k=1}^B \exp(\ell_{kj})}$. The resulting contrastive loss is

$$\mathcal{L}_{\text{CLIP}} = -\frac{1}{2B} \sum_{i=1}^B (\log p_{ii}^{I \rightarrow T} + \log p_{ii}^{T \rightarrow I}),$$

which maximizes agreement between matched image–text pairs while implicitly using all other samples in the batch as negatives.

Multimodal Embedding Construction: Let \mathcal{M} denote a trained CLIP-family encoder with parameters (θ, ϕ) . For a panel (I_i, T_i) , we define a unified multimodal embedding operator $\mathcal{E}_{\mathcal{M}} : (\mathcal{I}, \mathcal{T}) \rightarrow \mathbb{R}^d$, given by simple feature fusion of normalized image and text embeddings: $\mathbf{e}_i^{(\mathcal{M})} = \mathcal{E}_{\mathcal{M}}(I_i, T_i) = \frac{1}{2}(\mathbf{x}_i + \mathbf{y}_i)$.

Textual descriptions vary substantially in length and narrative structure. To obtain representations that are robust to localized phrasing effects, we model text encoding as an expectation over multiple subcontexts. Given a text sequence T_i of length $|T_i|$, we sample C contiguous substrings $\{T_i^{(c)}\}_{c=1}^C$ of fixed maximum length s , drawn uniformly from all valid substrings. The final multimodal representation is defined as

$$\mathbf{e}_i^{(\mathcal{M})} = \frac{1}{C} \sum_{c=1}^C \mathcal{E}_{\mathcal{M}}(I_i, T_i^{(c)}),$$

which can be interpreted as a Monte-Carlo approximation to marginalizing the embedding over textual subcontexts.

CLIP variants: We present three variants of CLIP-style encoders. All of them have an identical architectural backbone and embedding construction procedure, differing only in the training objective and parameterization used to obtain the encoder weights.

Pretrained CLIP The first variant employs a CLIP model trained on large-scale web image-text pairs using the softmax-based contrastive objective. Let (θ_0, ϕ_0) denote the parameters obtained from this pretraining process. The resulting embeddings $\mathbf{e}_i^{\text{CLIP-pre}} = \mathcal{E}_{(\theta_0, \phi_0)}(I_i, T_i)$ provide a strong general-purpose multimodal alignment prior and serve as a baseline for domain-specific adaptation.

Dataset-Specific Softmax-Contrastive Fine-Tuned CLIP This variant corresponds to a CLIP model whose parameters $(\theta_{\text{ft}}, \phi_{\text{ft}})$ are obtained by contrastive training on the image-text pairs drawn exclusively from our dataset. This dataset-specific fine-tuning explicitly aligns the representation space with the stylistic, iconographic, and linguistic distributions characteristic of narrative scroll paintings, while retaining the inductive bias of CLIP’s original cross-modal alignment mechanism. The resulting multimodal representations, $\mathbf{e}_i^{\text{CLIP-ft}} = \mathcal{E}_{(\theta_{\text{ft}}, \phi_{\text{ft}})}(I_i, T_i)$ exhibit enhanced sensitivity to dataset-specific semantics relative to the pretrained baseline,

thereby providing a stronger foundation for downstream graph-based propagation and label-conditioned recommendation.

Dataset-Specific Sigmoid-Contrastive Fine-Tuned CLIP (SigCLIP) This variant adopts the same dual-encoder CLIP architecture as the preceding models but replaces the softmax-based contrastive objective with a sigmoid-based pairwise loss during fine-tuning. Unlike softmax-based contrastive learning, which normalizes similarities across an entire batch, the sigmoid-contrastive formulation treats each image–text pairing as an independent binary classification problem. Let \mathbf{x}_i and \mathbf{y}_j denote the ℓ_2 -normalized image and text embeddings produced by the encoders. The similarity logit for a pair (i, j) is defined as $\ell_{ij} = t \mathbf{x}_i^\top \mathbf{y}_j + b$, where $t > 0$ is a learnable temperature parameter and b is an optional learnable bias term that stabilizes optimization under large class imbalance. Pairwise labels are assigned as

$$z_{ij} = \begin{cases} +1, & \text{if } i = j \quad (\text{matched image–text pair}), \\ -1, & \text{if } i \neq j \quad (\text{mismatched pair}). \end{cases}$$

The sigmoid-contrastive objective optimized during fine-tuning is then

$$\mathcal{L}_{\text{SigCLIP}} = -\frac{1}{B} \sum_{i=1}^B \sum_{j=1}^B \log \sigma(z_{ij} \ell_{ij}),$$

where $\sigma(\cdot)$ denotes the logistic sigmoid function. This loss eliminates batchwise normalization and instead enforces alignment through independent pairwise decisions, thereby decoupling gradient contributions across samples within a batch. As the architecture and embedding operator remain unchanged, the resulting multimodal representations $\mathbf{e}_i^{\text{SigCLIP-ft}} = \mathcal{E}_{(\theta_{\text{sig}}, \phi_{\text{sig}})}(I_i, T_i)$ differ from the pretrained and softmax-fine-tuned variants solely due to the altered contrastive learning signal.

4.1.3 Multimodal Variational Autoencoder

We introduce a multimodal variational autoencoder (VAE) that learns a probabilistic latent representation jointly grounded in visual and textual modalities. Our formulation follows the Auto-Encoding variational Bayes framework (AEVB, [Zhi-Han \(2022\)](#)), extended to heterogeneous multimodal inputs, and is used as a frozen feature extractor during downstream evaluation.

For each narrative panel i , let x_i^{img} denote the image modality and x_i^{text} the associated textual description. We posit a continuous latent variable $\mathbf{z}_i \in \mathbb{R}^d$ that generates both modalities. The joint distribution factorizes as

$$p(\mathbf{z}_i) = \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad p_\theta(x_i^{\text{img}}, x_i^{\text{text}} | \mathbf{z}_i) = p_\theta(x_i^{\text{img}} | \mathbf{z}_i) p_\theta(x_i^{\text{text}} | \mathbf{z}_i),$$

where conditional independence given \mathbf{z}_i enables tractable multimodal modeling while allowing each decoder to specialize to its respective modality. Exact posterior inference $p(\mathbf{z}_i | x_i^{\text{img}}, x_i^{\text{text}})$ is intractable. We therefore introduce modality-specific variational

approximations, each parameterized as a diagonal-covariance Gaussian:

$$q_\phi^{\text{img}}(\mathbf{z}_i | x_i^{\text{img}}) = \mathcal{N}\left(\boldsymbol{\mu}_i^{\text{img}}, \text{diag}(\boldsymbol{\sigma}_i^{\text{img}})^2\right),$$

$$q_\phi^{\text{text}}(\mathbf{z}_i | x_i^{\text{text}}) = \mathcal{N}\left(\boldsymbol{\mu}_i^{\text{text}}, \text{diag}(\boldsymbol{\sigma}_i^{\text{text}})^2\right).$$

These encoders independently extract modality-conditioned evidence about the latent semantic state \mathbf{z}_i .

4.2 Modality Fusion

To combine information from different modalities, we adopt a Product-of-Experts (PoE) formulation. The fused variational posterior is defined as

$$q_\phi(\mathbf{z}_i | x_i^{\text{img}}, x_i^{\text{text}}) \propto q_\phi^{\text{img}}(\mathbf{z}_i | x_i^{\text{img}}) q_\phi^{\text{text}}(\mathbf{z}_i | x_i^{\text{text}}),$$

which remains Gaussian when all experts are Gaussian. Let $\boldsymbol{\Lambda}_i^{(m)} = \text{diag}(\boldsymbol{\sigma}_i^{(m)})^{-2}$ denote the precision matrix for modality $m \in \{\text{img}, \text{text}\}$. The fused posterior parameters admit the closed-form solution

$$\begin{aligned} \boldsymbol{\Lambda}_i &= \boldsymbol{\Lambda}_i^{\text{img}} + \boldsymbol{\Lambda}_i^{\text{text}}, \\ \boldsymbol{\mu}_i &= \boldsymbol{\Lambda}_i^{-1} \left(\boldsymbol{\Lambda}_i^{\text{img}} \boldsymbol{\mu}_i^{\text{img}} + \boldsymbol{\Lambda}_i^{\text{text}} \boldsymbol{\mu}_i^{\text{text}} \right), \\ \boldsymbol{\Sigma}_i &= \boldsymbol{\Lambda}_i^{-1} \end{aligned}$$

This fusion mechanism performs uncertainty-aware integration: modalities with lower posterior variance exert greater influence on the shared latent representation. For each panel, we use the posterior mean of the fused distribution as a deterministic feature representation:

$$\mathbf{e}_i^{\text{VAE}} = \boldsymbol{\mu}_i$$

This choice eliminates sampling variance while retaining the uncertainty-aware fusion learned during training.

4.3 Concept (or Target Label) Learning

Twenty annotators, previously trained, were asked to identify the presence or absence of three different concepts in a panel: (a) animals, (b) trees, and (c) mythological characters. The goal is to find out how good humans are at identifying these concepts. Krippendorff’s alpha (Krippendorff (2004)) is used to measure inter-annotator agreement. The following were the values of alpha for different concepts: (a) Myth. char: 0.69 (b) Tree: 0.85 (c) Animals: 0.69.

4.4 Graph Structure Learning (GSL)

4.4.1 Learning homogeneous graphs

To discover latent structures of graphs from panels, we first use the multimodal features to learn $|\mathcal{P} \times \mathcal{P}|$ similarity matrices. Using the hypothesis that similar panels are likely to be of interest more often to the end-user than dissimilar ones [McPherson et al. \(2003\)](#), the semantic relationship between panels is quantified by using similarity measures such as cosine similarity ([Wang et al. \(2020\)](#), [Zhu et al. \(2021\)](#), [Zhang and Zitnik \(2020\)](#), [Chen et al. \(2020\)](#)) kernel-based functions ([Li et al. \(2018\)](#)), or attention mechanisms ([Chen et al. \(2020\)](#)). In this work, the parameter-free cosine similarity is used so that the similarity $S_m = \frac{(e_i^m)^T e_j^m}{\|e_i^m\| \|e_j^m\|}$, where e_i^m, e_j^m represents multimodal feature vectors. Since the adjacency matrix of a graph has positive values and a fully connected complete graph is often undesirable, we suppress the negative values in S_m . In addition, motivated by the fact that real-world graphs are noisy and can contain many task-unrelated edges, we apply a sparsity constraint. Borrowing an idea from ϵ -proximity thresholding graphs (ϵ -graphs [Bentley \(1975\)](#)), we create an edge between two nodes if their similarity is smaller than a predefined, data-driven, threshold ϵ .

4.4.2 Transductive versus Inductive Graph Learning

Training neural networks on graphs can be done in either the transductive or inductive setting ([Xu et al. \(2020\)](#); [Lachaud et al. \(2023\)](#)). In the transductive setting, the model has access to the test features in the training phase. In the inductive setting, the test data remain unseen. We explore the differences in inductive versus transductive learning primarily because field studies conducted at different times from disparate geographic locations have added data to our corpus at a steady pace, necessitating re-training of models and relying on inductive learning.

4.5 Node Classification

Let $\mathcal{G} = (A, X)$ denote a graph where $A \in R^{N \times N}$ is the adjacency matrix and $X \in R^{N \times F}$ is the node feature matrix with the i -th entry $x_i \in R^F$ denoting the feature of the node v_i . Given \mathcal{G} , the goal of GSL is to simultaneously learn an adjacency matrix A^* and its corresponding node representations for downstream tasks. A GSL model has two trainable components: (1) a GNN encoder f_θ that takes a graph as input and generates an embedding. (2) a structure learning module g_ϕ that contains the edge connectivity information of the graph.

The model parameters (θ, ϕ) are trained with the following learning objective:

$$\mathcal{L} = \mathcal{L}_{tsk}(Z^*, Y) + \lambda \mathcal{L}_{reg}(A^*, Z^*, \mathcal{G}) \quad (2)$$

where the first term on the right-hand side represents the original task, the second – regularizer on the learned graph, and λ is a hyperparameter to be tuned.

Thus, node classification can be used to assign labels (such as the presence/absence of a concept) to nodes in a graph based on the properties of nodes and the relationships

between them. The labels assigned to nodes (panels) are used for relevance estimation in the recommendation system.

The formulation used in Equation 2 relies on the assumption that connected edges in a graph are likely to share the same label. This assumption may restrict modeling capacity as graph edges need not encode similarity but may contain additional information. To address this, Graph Auto Encoders (GAEs Kipf and Welling (2017)) have been proposed, wherein the graph structure is modeled using a neural network $f(X, A)$. We call this GeMi GAE in our design and empirical analysis.

Marrying ideas from deep neural networks and approximate Bayesian inference, it is possible to further enhance the framework for encoding graph structures, using the Variational Graph Auto Encoder (VGAE, Kipf and Welling (2016)). VGAEs are capable of learning interpretable latent representations for undirected graphs. This can be demonstrated by using a Graph Convolutional Network (GCN) encoder and an inner product decoder. Assume stochastic latent variable z_i summarized in an $N \times D$ matrix Z . The inference model parameterized by a two-layer GCN is given by

$$q(z|X, A) = \prod_{i=1}^N q(z_i|X, A)$$

with $q(z_i|X, A) = \mathcal{N}(z_i|\mu_i, \text{diag}(\sigma_i^2))$. Also, $\mu = GCN_\mu(X, A)$ is the matrix of mean vectors of μ_i and $\log\sigma = GCN_\sigma(X, A)$. The two-layer GCN is defined as follows:

$$GCN(X, A) = \tilde{A} \text{ReLU}(\tilde{A}XW_0)W_1$$

with W_i being the weight matrices; GCN_μ and GCN_σ share the first layer of parameters W_0 and \tilde{A} is a symmetrically normalized adjacency matrix. The generative model is given by:

$$p(A|Z) = \prod_{i=1}^N \prod_{j=1}^N p(A_{ij}|z_i, z_j)$$

with $p(A_{ij} = 1|z_i, z_j) = \sigma(z_i^T z_j)$, $\sigma(\cdot)$ is the logistic sigmoid function. Finally, the learning step involves optimizing the variational lower bound \mathcal{L} w.r.t the variational parameters W_i :

$$\mathcal{L} = E_{q(Z|X, A)}[\log p(A|Z)] - KL[q(Z|X, A)||p(Z)]$$

where $KL[q||p]$ is the Kullback-Leibler divergence between $q(\cdot)$ and $p(\cdot)$. This model is referred to as GeMi VGAE in our architecture.

The three graph encoder variants in GeMi — GCN, GAE, and VGAE — are chosen to span a principled spectrum of inductive biases. GCN provides a discriminative baseline that propagates supervised label signals directly through the graph. GAE introduces an unsupervised reconstruction objective, allowing the model to leverage unlabeled structural information without relying solely on sparse concept labels. VGAE extends this with a probabilistic latent space, which is particularly well-suited

to our setting: the scroll painting dataset is small and noisy, and a variational formulation explicitly models uncertainty in node representations. Together, these three architectures allow us to study the trade-off between supervised signal strength, structural self-supervision, and latent uncertainty modeling — questions that are non-trivial to answer a priori for a novel dataset of this kind.

4.6 Strategies for Handling Class Imbalance

The concepts to be learned through node classification may be imbalanced (for example, *Tree*). In a graph-based recommendation system, this skew may cause (i) gradients to be dominated by majority labels, (ii) decision thresholds to shift toward predicting negatives for rare labels, and (iii) graph neighborhoods to reflect dense majority clusters rather than minority structure (Markwald and Demidova (2024); Zhao et al. (2021)). To counter these effects, we employ balancing strategies at the following three levels:

- graph structure
- loss reweighting and hard-example emphasis
- optimization regularization/stabilization

Details about these methods are presented in Section 5.

4.7 Incorporating Collaborative Filtering (CF) Methods

GeMi separates the process of multimodal feature extraction from user-panel interaction information in recommendation tasks. After learning panel representations, the user-panel interactions are incorporated with downstream collaborative filtering tasks. This allows collaborative filtering to be served as a plug-and-play module in the overall architecture.

5 Empirical Results

5.1 Aims

Our aim is to answer the following research questions:

- **RQ1: (Baseline Comparison)** How does our system (GeMi) perform when compared to the state-of-the-art graph-based recommendation systems built using homogeneous item-item graphs and other heterogeneous user-item graphs used in collaborative filtering settings?
- **RQ2: (Feature Generation)** Do LLM-enhanced vision-language features contribute to better model performance of the recommendation system?
- **RQ3: (Graph Structure Learning - Inductive versus Transductive training)** Does an inductive or transductive training procedure provide better results for recommendation?




Fig. 3: Distribution of the number of panels per scroll in the two phases of data collection from field surveys.

5.2 Materials

5.2.1 Data

We conducted a series of experiments on a meticulously curated dataset of panels and associated songs collected during two phases of field studies. There are a total of 106 panels from 13 scrolls, paired with 13 distinct songs collected in Phase 1. In Phase 2, 15 distinct scrolls with a total of 83 panels were collected. Figures 3a and Figure 3b show the distribution of the number of panels per scroll collected in these two phases. Although the data from phase 1 are well aligned (every panel illustrates an excerpt of a song), only 26.67% of the scrolls in phase 2 have associated songs. Each panel in the data set is annotated with the presence or absence of three key concepts (labels = 0 or 1): tree, animal and mythology.

5.2.2 Machines

All inductive experiments were conducted on a university-managed High Performance Computing (HPC) compute node under a single-GPU configuration. Unless otherwise specified, each run reserved 8 CPU cores, 64 GB RAM, and 1 GPU per job. The system supported a software stack with Jupyter, PyTorch 3, and CUDA, and computations were run on AVX-512 capable CPUs. All transductive experiments were conducted using Google Colab as the execution environment and were run with Python 3 and a single NVIDIA Tesla T4 GPU as the hardware accelerator, providing 16 GB of GPU memory.

5.2.3 Baseline Algorithms

We compare the performance of GeMi to the state-of-the-art transductive and inductive graph-based recommendation systems. These algorithms are briefly described in this section.

- Transductive homogenous baselines
 1. LATTICE (Zhang et al. (2021b)) constructs KNN item graphs for each modality and refines them into a skip-connected latent adjacency that is softmax-fused across modalities. A LightGCN-style propagation is then applied on the fused graph to strengthen BPR-optimized CF backbones.
 2. PMGT (Liu et al. (2021b)) performs pretraining on an item multimodal graph through a Transformer encoder, optimized with graph structure reconstruction and masked node feature reconstruction.
 3. HUIGN (Wei et al. (2022)) models multi-level user intents through hierarchical graph convolution learned from a co-interacted item graph, forming a fine-to-coarse hierarchy of latent intents that jointly refine user and item embeddings.
 4. DualGNN (Wang et al. (2023b)) applies simplified GNN propagation on modality-specific user-item bipartite graphs to learn per-modality user preferences, which are then attention-fused and propagated over a user co-occurrence graph to handle modality missing in user representations.
 5. FREEDOM (Zhou and Shen (2023)) improves multimodal recommendation by freezing a KNN-based item-item graph constructed from raw multimodal features while denoising the user-item graph through degree-sensitive edge pruning to suppress noisy interaction.
 6. MICRO (Zhang et al. (2023)) learns modality-specific item-item graphs from multimodal features. It refines embeddings with graph convolutions and aligns modalities through a contrastive fusion objective.
 7. DGVAE (Zhou and Miao (2024)) uses a disentangled graph variational auto-encoder to learn interpretable multimodal user preferences, by encoding both user-item interactions and user-word preferences on a fixed item-item graph.

Since all of the above algorithms construct homogeneous item-item graphs, they are closest in spirit to the GeMi system.

- Transductive heterogeneous baselines
 1. MambaRec (Ren et al. (2025)) integrates local alignment through a Dilated Refinement Attention Module (DREAM) with multi-scale convolutions and dual attention, and global modality regularization through Maximum Mean Discrepancy (MMD) and contrastive objectives to enforce cross-modal semantic consistency.
 2. PGL (Yu et al. (2025)) extracts a principal user-item subgraph (via global-aware truncation or local-aware sampling) to preserve local structural information, with InfoNCE-based feature masking, and also leveraging multimodal item-item graphs.
 3. HPMRec (Chen et al. (2025a)) encodes user-item representations with multi-component hypercomplex embeddings. It performs nonlinear cross-modal fusion through hypercomplex multiplication and uses prompt-aware compensation to dynamically align components.
 4. COHESION (Xu et al. (2025a)) is a composite GCN that integrates graphs user-item, user-user, and item-item, refines embeddings via behavior modality,

and then carefully fuses representations of the modality, with an adaptive BPR objective to balance learning.

5. SMORE (Ong and Khong (2025)) fuses visual and textual features in the frequency domain, with adaptive spectral filtering to suppress modality-specific noise before fusion. It models both uni-modal and fused preferences through multi-view graph learning and a modality-aware preference module.
 6. CMDL (Lin et al. (2025b)) disentangles multimodal representations into modality-invariant and modality-specific components using a novel mutual-information-based disentanglement regularization.
- Inductive baselines: All the inductive baselines are studied both in the homogeneous (item-item) and heterogeneous (user-item) settings.
 1. GraphSAGE (Hamilton et al. (2017)): In the homogeneous setting, this algorithm first constructs an item-item graph using cosine similarity on item embeddings and a fixed percentile threshold; it then trains a two-layer GraphSAGE encoder to learn the aggregation function over node features. In the heterogeneous setting, it constructs a user-item bipartite graph through neighborhood aggregation.
 2. PinSAGE³ (Cao et al. (2025); Ying et al. (2018)) It replaces purely feature-based neighborhoods for node aggregation with random-walk-based neighborhoods by precomputing co-visitation counts and rewiring the item-item graph with top-k most frequently visited neighbors.
 3. GATNE-I (Cen et al. (2019)) learns item embeddings on a single-modality KNN graph using a skip-gram objective with negative sampling.

5.2.4 User-Preference Matrix Generation

We simulate user preferences by constructing a synthetic user-preference matrix derived from the training corpus through Monte-Carlo sampling. Let the training set consist of N_{tr} panels. Each panel i is annotated with a binary multi-label indicator vector \mathbf{y}_i , where $\mathbf{y}_i = (y_{i1}, \dots, y_{i|\mathcal{L}|}) \in \{0, 1\}^{|\mathcal{L}|}$, over a fixed semantic label set \mathcal{L} , where $|\mathcal{L}| = 3$ corresponds to the categories *Animal*, *Mythology*, and *Tree*.

Synthetic User Construction: We generate a population of U synthetic users, each intended to represent a plausible latent preference profile over the label space. For each synthetic user $u \in \{1, \dots, U\}$, we sample a subset of K distinct training items, $\mathcal{I}_u \subset \{1, \dots, N_{\text{tr}}\}$, $|\mathcal{I}_u| = K$, $1 \leq K \leq N_{\text{tr}}$. The interaction sets are sampled independently according to the uniform distribution over all K -subsets of the training index set:

$$\mathcal{I}_u \sim \text{Unif}(\{\mathcal{S} \subseteq \{1, \dots, N_{\text{tr}}\} \mid |\mathcal{S}| = K\})$$

This sampling scheme treats each training item as an equally likely candidate interaction and enforces diversity within a user profile by prohibiting repeated items. When the training set size is smaller than K , the effective interaction count is implicitly capped at N_{tr} to ensure feasibility.

³Our implementation of the PinSAGE algorithm is available from https://github.com/Haimonti/Folk-Art-Recommendation-System/tree/Inductive_Jin/Inductive_settings and is an improved version of Cao et al. (2025).

Preference Estimation: Given the sampled interaction set \mathcal{I}_u , we estimate the latent interests of user u by aggregating the semantic labels associated with the selected items. Specifically, we compute the empirical label frequency vector $\hat{\mathbf{f}}_u = \frac{1}{K} \sum_{i \in \mathcal{I}_u} \mathbf{y}_i \in [0, 1]^{|\mathcal{L}|}$, where each component \hat{f}_{ul} represents the fraction of sampled interactions that exhibit label ℓ . To obtain a discrete and interpretable preference representation, we apply a fixed threshold $\tau \in (0, 1)$ to the empirical frequencies, yielding a binary preference vector $\mathbf{p}_u = \mathbb{I}[\hat{\mathbf{f}}_u \geq \tau] \in \{0, 1\}^{|\mathcal{L}|}$, where the indicator function is applied elementwise. Equivalently, a label is deemed preferred by user u if it appears in at least a τ -fraction of the sampled interactions. Equivalently, for each label $\ell \in \{1, \dots, |\mathcal{L}|\}$,

$$p_{ul} = \begin{cases} 1, & \text{if } \frac{1}{K} \sum_{i \in \mathcal{I}_u} y_{il} \geq \tau, \\ 0, & \text{otherwise} \end{cases}$$

This formulation naturally adapts to the interaction budget K , ensuring that preference activation reflects repeated evidence rather than isolated occurrences.

User-Preference Matrix: Aggregating the preference vectors across all synthetic

users yields the user-preference matrix $\mathbf{P} = \begin{bmatrix} \mathbf{p}_1^\top \\ \mathbf{p}_2^\top \\ \vdots \\ \mathbf{p}_U^\top \end{bmatrix} \in \{0, 1\}^{U \times |\mathcal{L}|}$. Each row of \mathbf{P}

encodes a distinct synthetic user’s inferred interests over the semantic label space, while each column corresponds to a specific label category. The total number of training item appearances across all profiles is UK , and assuming approximately uniform sampling, each training item participates in roughly $\frac{UK}{N_{\text{tr}}}$ synthetic profiles on average, promoting broad coverage of the training distribution.

5.3 Methods

The following details are relevant:

- **Parameters of GeMi variants:** Table 1 summarizes the architectural and optimization configurations of the three GeMi variants, namely GeMi-GCN, GeMi-GAE, and GeMi-VGAE, under the transductive protocol. In all cases, graph construction is based on cosine-similarity k-nearest-neighbor connectivity, supervised learning employs focal binary cross-entropy with logits, optimization is performed using Adam, and regularization includes dropout, edge dropout, and gradient clipping. Evaluation is conducted using fixed synthetic user profiles with $k = 5$ interactions and a preference threshold of 0.2.

To ensure comparability across variants, several core hyperparameters are intentionally aligned. These include dropout (0.20), optimizer (Adam), gradient clipping with maximum norm 2.0, kNN-based graph construction (with variant-specific k values of 30, 30, and 25), edge dropout (0.10), ReLU activation in graph and head layers, and focal loss parameters $\gamma = 2.0$ and $\alpha = 0.25$. Furthermore, the same recommendation evaluation setup ($k = 5$ interactions and preference threshold as 0.2) is maintained across all models.

GeMi Variant	Parameters
GeMi-GCN	Graph encoder: 2-layer Graph Convolutional Network; GCN hidden dimension: 128; GCN dropout: 0.20; Training epochs: 450; Optimizer: Adam; Activation: ReLU; Learning rate: 3×10^{-4} ; Weight decay: 2×10^{-3} ; Gradient clipping: max norm 2.0; Loss (balanced): focal binary cross-entropy with logits; Focal parameters: $\gamma = 2.0$, $\alpha = 0.25$; kNN neighbors (k): 30; Tree-edge augmentation neighbors (k_{tree}): 25; Tree augmentation max nodes: 2500; Edge dropout: 0.10; $k = 5$ interactions; Preference threshold: 0.2
GeMi-GAE	Graph encoder: Graph Auto-Encoder (GAE) with a 2-layer GCN encoder; Hidden dimension: 128; Latent dimension (d_z): 64; Dropout: 0.20; Training epochs: 400; Optimizer: Adam; Activation: ReLU; Learning rate: 2×10^{-3} ; Weight decay: 2×10^{-4} ; Gradient clipping: max norm 2.0; kNN neighbors (k): 30; Edge dropout: 0.10; Self-supervised loss: GAE reconstruction loss \mathcal{L}_{rec} ; Supervised loss: focal binary cross-entropy with logits; Focal parameters: $\gamma = 2.0$, $\alpha = 0.25$; Joint objective: $\mathcal{L} = \mathcal{L}_{rec} + \lambda_{sup}\mathcal{L}_{sup}$ with $\lambda_{sup} = 0.60$; Augmentation neighbors: $k_{animal} = 18$, $k_{myth} = 18$, $k_{tree} = 35$; Augmentation max nodes per label: 2500; $k = 5$ interactions; Preference threshold: 0.2
GeMi-VGAE	Graph encoder: Variational Graph Auto-Encoder (VGAE) with GCN-based parameterization; Hidden dimension: 256; Latent dimension (d_z): 128; Dropout: 0.20; Log-standard-deviation clamp: $\log \sigma \in [-10, 10]$; Training epochs: 500; Optimizer: Adam; Activation: ReLU; Learning rate: 3×10^{-3} ; Weight decay: 5×10^{-4} ; Gradient clipping: max norm 2.0; kNN neighbors (k): 25; Tree augmentation neighbors: $k_{tree} = 25$; Tree augmentation max nodes: 2500; Edge dropout: 0.10; Reconstruction loss: VGAE reconstruction loss \mathcal{L}_{rec} ; KL regularization: VGAE KL loss \mathcal{L}_{KL} ; KL annealing: $\beta_{KL}(t)$ linearly increased to $\beta_{max} = 1.0$; Self-supervised VGAE objective: $\mathcal{L}_{vgae} = \mathcal{L}_{rec} + \beta_{KL}\mathcal{L}_{KL}$; Supervised loss: focal binary cross-entropy with logits; Focal parameters: $\gamma = 2.0$, $\alpha = 0.25$; Joint SSL objective: $\mathcal{L} = \mathcal{L}_{vgae} + \lambda_{ssl}\mathcal{L}_{sup}$ with $\lambda_{ssl} = 0.60$; $k = 5$ interactions; Preference threshold: 0.2

Table 1: Parameters of GeMi Variants

- Balanced Graph Structure Learning using Symmetric KNN:** Instead of constructing edges using a global similarity threshold (which tends to over-connect dense majority regions and under-connect rare regions), we construct a *symmetric k-nearest-neighbor (kNN) graph* over all nodes. Let $\mathbf{X} \in \mathbb{R}^{N \times d}$ be node features. We first ℓ_2 -normalize: $\tilde{\mathbf{x}}_i = \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|_2 + \varepsilon}$. We compute cosine similarities $s_{ij} = \cos(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$ and, for each node i , select its top- k neighbors: $\mathcal{N}_k(i) = \arg \text{topk}_{j \neq i} s_{ij}$. We then add *bidirectional* edges (symmetric closure): $A_{ij} = \mathbb{I}[j \in \mathcal{N}_k(i)] \vee \mathbb{I}[i \in \mathcal{N}_k(j)]$. Optionally, we apply a similarity floor $s_{ij} \leftarrow \max(s_{ij}, 0)$ to avoid negative edges. Symmetric kNN enforces roughly uniform local neighborhood size, improving representation learning for sparse concepts by ensuring every node participates in comparable amounts of graph propagation.

- **Class Imbalance Handling Strategies:** We employ several different strategies for handling class imbalance. These include:

1. **Minority-Aware Edge Augmentation for the *Tree* Label:** To explicitly address extreme sparsity in the *Tree* class, we add extra edges among *Tree-positive* training nodes. Let $\mathcal{V}_{\text{tree}} = \{i \mid y_{i,\text{tree}} = 1\}$ be indices of nodes labeled positive for *Tree* (restricted to train nodes in practice). We compute cosine similarities within $\mathcal{V}_{\text{tree}}$ and connect each such node to its top- k_{tree} neighbors: $\mathcal{N}_{k_{\text{tree}}}^{\text{tree}}(i) = \arg \text{topk}_{j \in \mathcal{V}_{\text{tree}}, j \neq i} \cos(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$. We then union these edges with the base kNN graph: $\mathcal{E} \leftarrow \mathcal{E} \cup \mathcal{E}_{\text{tree}}$. For computational feasibility and to prevent over-connection, we cap the number of *Tree* nodes used in augmentation (random subsampling if needed). Adding intra-minority edges increases within-class information flow, stabilizes embeddings for minority positives, and reduces the chance that these nodes are overwhelmed by messages dominated by majority neighborhoods.
2. **Edge Dropout to Reduce Majority-Connectivity Overfitting:** During each training epoch, we apply *edge dropout* by randomly removing edges with probability p_e . For an edge set \mathcal{E} , we sample Bernoulli masks $\delta_e \sim \text{Bernoulli}(1-p_e)$ and keep: $\mathcal{E}' = \{e \in \mathcal{E} \mid \delta_e = 1\}$. Edge dropout acts as structural regularization, encouraging robustness to missing edges and preventing reliance on a small subset of dominant connections, which benefits minority generalization.
3. **Class-Balanced Reweighting via Positive Weights:** We compute per-label positive weights from the *training* labels to counteract skewed prevalence. For label ℓ , let $P_\ell = \sum_{i \in \text{train}} y_{i\ell}$ be positives and $N_\ell = n_{\text{train}} - P_\ell$ negatives. We set: $w_\ell = \frac{N_\ell}{P_\ell}$, (with clamping to avoid division by zero), and use weighted BCE-with-logits:

$$\mathcal{L}_{\text{WBCE}} = \frac{1}{n_{\text{train}}} \sum_{i \in \text{train}} \sum_{\ell} \left(-w_\ell y_{i\ell} \log \sigma(z_{i\ell}) - (1 - y_{i\ell}) \log(1 - \sigma(z_{i\ell})) \right),$$

where $z_{i\ell}$ are logits and $\sigma(\cdot)$ is the sigmoid. Without reweighting, the empirical risk is dominated by abundant negatives and frequent labels, causing the optimizer to favor conservative predictions (especially for rare labels). w_ℓ increases the gradient contribution of minority positives, effectively equalizing the cost of misclassifying rare positives versus frequent ones.

4. **Focal Loss to Emphasize Hard and Minority Examples:** We further incorporate *focal modulation* on top of BCE-with-logits. Let $p_{i\ell} = \sigma(z_{i\ell})$. Define

$$p_t = \begin{cases} p_{i\ell}, & y_{i\ell} = 1, \\ 1 - p_{i\ell}, & y_{i\ell} = 0. \end{cases}$$

We apply the focusing factor $(1 - p_t)^\gamma$ and an α -balance term: $\alpha_t = \alpha y_{i\ell} + (1 - \alpha)(1 - y_{i\ell})$, yielding the focal-BCE objective:

$$\mathcal{L}_{\text{focal}} = \frac{1}{n_{\text{train}}} \sum_{i \in \text{train}} \sum_{\ell} \alpha_t (1 - p_t)^\gamma \text{BCEWithLogits}(z_{i\ell}, y_{i\ell}; w_\ell).$$

Focal loss reduces weights on well-classified examples and concentrates learning on hard cases—often boundary samples and minority positives—improving minority recall without excessively inflating false positives.

5. Training Stabilization Under Reweighted Objectives: Reweighting (pos-weights) and focal modulation can increase gradient variance. We therefore adopt stabilization measures:
 - (a) Gradient clipping: after backpropagation, we clip the global norm

$$\|\nabla_{\theta}\mathcal{L}\|_2 \leftarrow \min(\|\nabla_{\theta}\mathcal{L}\|_2, \tau), \quad (3)$$

preventing rare-label spikes from destabilizing updates.

- (b) Regularization and capacity tuning: we use a larger hidden dimension and dropout in the GCN and employ weight decay, improving generalization under class imbalance.
- (c) Edge dropout (above) as a structural regularizer complements loss balancing by preventing topology overfitting.

Balancing strategies intentionally amplify minority gradients; clipping and regularization ensure these amplified gradients do not lead to exploding updates, oscillations, or overfitting, especially in multi-label regimes where some samples contribute to multiple losses.

- **Transductive and Inductive Graph Learning:** The GeMi system is studied in both the transductive and inductive settings – for each setting, state-of-the-art baseline algorithms for homogeneous graphs are used for comparison. In addition, in the transductive setting, we also study collaborative filtering style algorithms built on heterogeneous graphs. While the baseline algorithms have been discussed in Section 5.2.3, Table A1 and Table A2 elucidate the parameters used to run them.
- **Relevance Estimation and Recommendation:** We formalize relevance estimation as a label-conditioned semantic consistency problem between inferred user preferences and recommended items. Let \mathcal{D}_{tr} and \mathcal{D}_{te} denote the training and test splits, respectively, with cardinalities $|\mathcal{D}_{\text{tr}}| = N_{\text{tr}}$ and $|\mathcal{D}_{\text{te}}| = N_{\text{te}}$. Each item (panel) $i \in \mathcal{D}_{\text{tr}} \cup \mathcal{D}_{\text{te}}$ is associated with a binary multi-label annotation

$$\mathbf{y}_i = (y_{i1}, \dots, y_{i|\mathcal{L}|}) \in \{0, 1\}^{|\mathcal{L}|}, \quad (4)$$

where \mathcal{L} denotes the fixed semantic label set (with $|\mathcal{L}| = 3$ in our setting). The label indicators are treated as ground-truth semantic descriptors of item content. We consider a family of embedding models indexed by $m \in \mathcal{M}$. For each model m , items in the training and test splits are embedded into a d_m -dimensional latent space, $\mathbf{e}_i^{(m)} \in \mathbb{R}^{d_m}$, $i \in \mathcal{D}_{\text{tr}} \cup \mathcal{D}_{\text{te}}$

1. Synthetic User Preference Model: Let $u \in \{1, \dots, U\}$ denote a user. Each user is associated with an inferred binary preference vector $\mathbf{p}_u = (p_{u1}, \dots, p_{u|\mathcal{L}|}) \in \{0, 1\}^{|\mathcal{L}|}$, where $p_{u\ell} = 1$ indicates that user u is inferred to prefer semantic label $\ell \in \mathcal{L}$. The construction of \mathbf{p}_u follows the Monte Carlo sampling and thresholding procedure described in Section 5.2.4, and is treated as fixed for the remainder of the evaluation.

2. **User Representation in Latent Space:** Given a synthetic user u , we associate a latent user representation $\mathbf{e}_u^{(m)} \in \mathbb{R}^{d_m}$ by aggregating the embeddings of the items sampled to generate the user profile. This representation is interpreted as a point estimate of the user’s latent semantic interest under model m .
3. **Scoring and Ranking Function:** For a given model m and synthetic user u , each test item $t \in \mathcal{D}_{\text{te}}$ is assigned a relevance score $s^{(m)}(u, t) = \text{sim}(\mathbf{e}_u^{(m)}, \mathbf{e}_t^{(m)})$, where $\text{sim}(\cdot, \cdot)$ denotes cosine similarity, $\text{sim}(\mathbf{x}, \mathbf{z}) = \frac{\mathbf{x}^\top \mathbf{z}}{\|\mathbf{x}\|_2 \|\mathbf{z}\|_2}$. The recommended list of size K_{rec} is defined as $\mathcal{R}_u^{(m)} = \arg \max_{\substack{S \subset \mathcal{D}_{\text{te}} \\ |S|=K_{\text{rec}}}} \sum_{t \in S} s^{(m)}(u, t)$, i.e., the set of test items attaining the top- K_{rec} similarity scores.
4. **Label-Conditioned Relevance Definition** Relevance is evaluated independently for each semantic label. For user u , item t , and label ℓ , we define the label-wise relevance indicator

$$r_{u,t,\ell} = \mathbb{I}[p_{u\ell} = 1 \wedge y_{t\ell} = 1]$$

Thus, a recommended item is deemed relevant for label ℓ if and only if (i) the user is inferred to prefer label ℓ , and (ii) the item expresses label ℓ in its ground-truth annotation. Equivalently, for each user u and label ℓ , we define the label-conditioned candidate set

$$\mathcal{C}_{u,\ell} = \{t \in \mathcal{D}_{\text{te}} \mid p_{u\ell} = 1 \wedge y_{t\ell} = 1\},$$

and the set of relevant recommendations as

$$\mathcal{RR}_{u,\ell}^{(m)} = \mathcal{R}_u^{(m)} \cap \mathcal{C}_{u,\ell}$$

5. **Evaluation Metric (Precision@K):** For each synthetic user u , label ℓ , and model m , we compute label-specific Precision@ K_{rec} as

$$\text{P@}K_{\text{rec}}^{(m)}(u, \ell) = \frac{|\mathcal{RR}_{u,\ell}^{(m)}|}{K_{\text{rec}}} = \frac{1}{K_{\text{rec}}} \sum_{t \in \mathcal{R}_u^{(m)}} r_{u,t,\ell} \quad (5)$$

This yields, for each user, a vector of label-wise precision scores

$$\boldsymbol{\pi}_u^{(m)} = (\text{P@}K_{\text{rec}}^{(m)}(u, 1), \dots, \text{P@}K_{\text{rec}}^{(m)}(u, |\mathcal{L}|)) \in \mathbb{R}^{|\mathcal{L}|} \quad (6)$$

Aggregating over the population of synthetic users, we report the empirical mean and standard deviation

$$\mu_\ell^{(m)} = \frac{1}{U} \sum_{u=1}^U \text{P@}K_{\text{rec}}^{(m)}(u, \ell), \quad \sigma_\ell^{(m)} = \sqrt{\frac{1}{U} \sum_{u=1}^U (\text{P@}K_{\text{rec}}^{(m)}(u, \ell) - \mu_\ell^{(m)})^2} \quad (7)$$

We report Precision@K as the primary metric because the recommendation task in GeMi is explicitly framed as a precision-oriented retrieval problem: given a user’s

known concept preferences, the system surfaces a short list of K panels likely to match those preferences. In this setting, returning a small number of highly relevant items is more valuable than exhaustive retrieval — a user browsing an art collection benefits more from a curated short list than from high recall across hundreds of panels. We acknowledge, however, that Precision@ K alone does not capture ranking quality; future work incorporating NDCG or Mean Average Precision would provide a more complete evaluation, particularly as the dataset grows.

5.4 Results

In this section, we present empirical analysis in the transductive and inductive settings and answer the research questions proposed in Section 5.1.

5.4.1 Performance Comparison (RQ1)

Transductive experiments

Tables 2–5 compare GEMi against state-of-the-art, graph-based baselines under transductive setting across homogeneous item–item graphs and heterogeneous user–item graphs. Overall, GEMi achieves the strongest and most consistent performance across concepts, particularly on Mythology and Tree, while remaining competitive on Animal. The results demonstrate that combining LLM-enhanced multimodal features with graph-based representation learning yields clear improvements over prior homogeneous and heterogeneous graph recommenders.

On homogeneous item–item graphs with Llama-SigCLIP features, baselines such as LATTICE, PMGT, HUIGN, DualGNN, FREEDOM, MICRO, and DGVAE achieve moderate performance, with MICRO reaching 0.61 on Animal and HUIGN reaching 0.55 on Mythology. However, GEMi-GCN substantially improves Tree to 0.62 and Mythology to 0.61, indicating strong relational propagation when label signals are sparse. GEMi-GAE further improves Animal to 0.59 while maintaining balanced results across labels. Most notably, GEMi-VGAE achieves 0.69 on Mythology, the highest value in this configuration, showing that variational graph encoding better captures uncertainty and latent structure in item–item relations. Thus, even when strong homogeneous baselines are considered, all three GEMi variants either match or surpass the best-performing competitors in at least one label dimension.

Under homogeneous graphs with Llama-VAE features, the improvements become more pronounced. While baselines remain below 0.60 on most dimensions, GEMi-GCN reaches 0.60 on Animal and 0.65 on Mythology. GEMi-GAE achieves the overall best Mythology performance at 0.78, significantly exceeding prior methods. GEMi-VGAE reaches 0.62 on Animal, again outperforming baselines. These results show that when richer generative text representations are used, the graph encoder variants in GEMi consistently dominate homogeneous graph baselines.

A similar trend holds for heterogeneous user–item graphs. In the Llama-SigCLIP configuration, competitive multimodal CF baselines such as MambaRec, PGL, HPM-Rec, COHESION, SMORE, and CMDL achieve values around 0.60–0.67 on Mythology but remain weaker on Tree. GEMi-GCN improves Tree to 0.36 while maintaining 0.65 on Mythology. GEMi-GAE achieves 0.78 on Mythology, the strongest result in this

Model	Animal	Mythology	Tree
LATTICE	0.57 ± 0.21	0.47 ± 0.25	0.20 ± 0.18
PMGT	0.52 ± 0.21	0.46 ± 0.23	0.33 ± 0.22
HUIGN	0.48 ± 0.16	0.55 ± 0.25	0.10 ± 0.11
DualGNN	0.55 ± 0.21	0.47 ± 0.26	0.25 ± 0.21
FREEDOM	0.58 ± 0.22	0.41 ± 0.24	0.22 ± 0.19
MICRO	0.61 ± 0.20	0.47 ± 0.26	0.26 ± 0.21
DGVAE	0.55 ± 0.21	0.45 ± 0.26	0.24 ± 0.22
GeMi GCN	0.47 ± 0.27	0.61 ± 0.16	0.62 ± 0.30
GeMi GAE	0.59 ± 0.30	0.55 ± 0.13	0.49 ± 0.29
GeMi VGAE	0.51 ± 0.22	0.69 ± 0.16	0.48 ± 0.21

Table 2: Transductive Llama-SigCLIP features, homogeneous graph

Model	Animal	Mythology	Tree
LATTICE	0.53 ± 0.22	0.54 ± 0.13	0.21 ± 0.25
PMGT	0.47 ± 0.19	0.52 ± 0.17	0.25 ± 0.20
HUIGN	0.49 ± 0.20	0.51 ± 0.12	0.23 ± 0.22
DualGNN	0.41 ± 0.18	0.56 ± 0.32	0.19 ± 0.27
FREEDOM	0.51 ± 0.17	0.55 ± 0.24	0.24 ± 0.16
MICRO	0.45 ± 0.25	0.58 ± 0.23	0.26 ± 0.19
DGVAE	0.44 ± 0.23	0.60 ± 0.21	0.27 ± 0.21
GeMi GCN	0.60 ± 0.24	0.65 ± 0.20	0.36 ± 0.26
GeMi GAE	0.48 ± 0.21	0.78 ± 0.23	0.27 ± 0.15
GeMi VGAE	0.62 ± 0.25	0.56 ± 0.19	0.27 ± 0.28

Table 3: Transductive Llama-VAE features, homogeneous graph

setting. GEMi-VGAE reaches 0.62 on Animal, again outperforming heterogeneous competitors. This indicates that GEMi effectively integrates collaborative user-item structure with multimodal feature signals.

In the Llama-VAE heterogeneous setting, the pattern persists. Although CMDL and HPMRec perform competitively on Animal and Mythology, none of the baselines reach the 0.78 Mythology score achieved by GEMi-GAE. Similarly, GEMi-GCN and GEMi-VGAE maintain strong performance on Animal and Tree relative to prior methods. Across all heterogeneous experiments, the Tree label particularly benefits from the explicit graph propagation and label-aware augmentation in GEMi, where prior CF models struggle due to sparsity.

Collectively, these findings answer RQ1 by demonstrating that GEMi, through its integration of LLM/VLM-derived features and graph-based relational learning, outperforms both homogeneous item-item GNN recommenders and heterogeneous collaborative filtering baselines under transductive training.

Model	Animal	Mythology	Tree
MambaRec	0.50 ± 0.24	0.60 ± 0.17	0.12 ± 0.13
PGL	0.43 ± 0.26	0.60 ± 0.20	0.25 ± 0.18
HPMRec	0.47 ± 0.26	0.63 ± 0.21	0.34 ± 0.21
COHESION	0.41 ± 0.23	0.59 ± 0.20	0.27 ± 0.17
SMORE	0.40 ± 0.20	0.66 ± 0.24	0.24 ± 0.16
CMDL	0.41 ± 0.25	0.67 ± 0.22	0.23 ± 0.18
GeMi GCN	0.60 ± 0.24	0.65 ± 0.20	0.36 ± 0.26
GeMi GAE	0.48 ± 0.21	0.78 ± 0.23	0.27 ± 0.15
GeMi VGAE	0.62 ± 0.25	0.56 ± 0.19	0.27 ± 0.28

Table 4: Transductive Llama-SigCLIP features (user-item)

Model	Animal	Mythology	Tree
MambaRec	0.52 ± 0.19	0.57 ± 0.26	0.15 ± 0.28
PGL	0.44 ± 0.23	0.60 ± 0.21	0.27 ± 0.21
HPMRec	0.46 ± 0.18	0.64 ± 0.22	0.23 ± 0.14
COHESION	0.50 ± 0.24	0.55 ± 0.11	0.25 ± 0.26
SMORE	0.41 ± 0.54	0.59 ± 0.13	0.17 ± 0.31
CMDL	0.57 ± 0.28	0.61 ± 0.15	0.16 ± 0.25
GeMi GCN	0.60 ± 0.24	0.65 ± 0.20	0.36 ± 0.26
GeMi GAE	0.48 ± 0.21	0.78 ± 0.23	0.27 ± 0.15
GeMi VGAE	0.62 ± 0.25	0.56 ± 0.19	0.27 ± 0.28

Table 5: Transductive Llama-VAE features heterogeneous graph

Model	Animal	Mythology	Tree
GraphSAGE	0.48 ± 0.14	0.56 ± 0.23	0.34 ± 0.28
PinSAGE	0.52 ± 0.21	0.57 ± 0.22	0.49 ± 0.14
GATNE-I	0.52 ± 0.20	0.60 ± 0.21	0.37 ± 0.19
GeMi GCN	0.56 ± 0.24	0.66 ± 0.20	0.44 ± 0.20
GeMi GAE	0.50 ± 0.22	0.63 ± 0.15	0.34 ± 0.23
GeMi VGAE	0.52 ± 0.22	0.70 ± 0.21	0.44 ± 0.20

Table 6: Inductive Llama-SigCLIP features (item-item)

Inductive experiments

Tables 6–7 summarize inductive performance under homogeneous item–item graphs. Overall, the best-performing GEMi variants are label-dependent.

- Animal: The strongest inductive item–item result is achieved by GeMi-GCN with LLM-SigCLIP features (0.56), which is similar to the best inductive competitors with LLM-VAE.

Model	Animal	Mythology	Tree
GraphSAGE	0.56 ± 0.25	0.58 ± 0.12	0.33 ± 0.16
PinSAGE	0.56 ± 0.25	0.57 ± 0.12	0.32 ± 0.17
GATNE-I	0.56 ± 0.25	0.60 ± 0.19	0.34 ± 0.22
GeMi GCN	0.46 ± 0.22	0.32 ± 0.19	0.56 ± 0.34
GeMi GAE	0.51 ± 0.29	0.61 ± 0.18	0.29 ± 0.23
GeMi VGAE	0.50 ± 0.27	0.61 ± 0.23	0.29 ± 0.23

Table 7: Inductive Llama-VAE features (item-item)

Model	Animal	Mythology	Tree
GraphSAGE	0.63 ± 0.18	0.59 ± 0.19	0.26 ± 0.18
PinSAGE	0.50 ± 0.21	0.65 ± 0.17	0.21 ± 0.15
GATNE-I	0.51 ± 0.20	0.60 ± 0.21	0.36 ± 0.19
GeMi GCN	0.71 ± 0.25	0.39 ± 0.05	0.35 ± 0.13
GeMi GAE	0.47 ± 0.20	0.71 ± 0.20	0.20 ± 0.22
GeMi VGAE	0.52 ± 0.23	0.77 ± 0.11	0.53 ± 0.19

Table 8: Inductive Llama-SigCLIP features heterogeneous graph

Model	Animal	Mythology	Tree
GraphSAGE	0.58 ± 0.24	0.52 ± 0.15	0.26 ± 0.12
PinSAGE	0.46 ± 0.20	0.61 ± 0.11	0.39 ± 0.22
GATNE-I	0.56 ± 0.24	0.61 ± 0.20	0.33 ± 0.22
GeMi GCN	0.36 ± 0.12	0.59 ± 0.07	0.35 ± 0.13
GeMi GAE	0.69 ± 0.25	0.77 ± 0.11	0.34 ± 0.15
GeMi VGAE	0.53 ± 0.19	0.78 ± 0.10	0.53 ± 0.19

Table 9: Inductive Llama-VAE features heterogeneous

- Mythology: The best inductive result is GeMi-VGAE with LLM-SigCLIP features (0.70). The strongest competing baseline is GATNE-I (0.60), while GraphSAGE and PinSAGE are lower.
- Tree: Tree is the hardest label overall, but GEMi still achieves the best inductive item-item performance. Under LLM+SigCLIP, GeMi-GCN reaches 0.56 ± 0.34 , substantially above GraphSAGE/PinSAGE/GATNE-I. A possible explanation is that “tree” is often a secondary or background element in many panels: it occupies a small region, is visually less salient than the main subjects, and may be weakly described (if at all) in the accompanying text. As a result, both the multimodal features and the induced neighborhood structure provide a limited discriminative signal for this label compared to more semantically dominant categories such as Mythological characters or Animals.

In the user–item inductive setting (Tables 8–9), GeMi variants remain competitive and sometimes improve further for specific labels (e.g., Animal: GeMi-GCN with LLM-SigCLIP (0.71); Mythology: GeMi-VGAE with LLM-VAE (0.78); Tree: GeMi-VGAE with LLM-VAE (0.53)), which is expected since the user–item bipartite graph provides additional interaction structure beyond purely feature-based item neighborhoods.

5.4.2 Feature Generation (RQ2)

Transductive experiments We first compare base feature representations (SigCLIP + Graph and VAE + Graph in Tables 10 and 11) against their LLM-enhanced counterparts within the same GeMi backbone (GCN, GAE, VGAE). We then analyze the relative behavior of LLM + SigCLIP versus LLM + VAE across label dimensions.

LLM-enhanced vs Base features within GeMi Comparing tables 10 and 11 with tables 2–5 results shows a consistent and substantial improvement when LLM-generated features are used. For example, under homogeneous item–item graphs, Base-SigCLIP + GCN achieves 0.48 on Animal and 0.70 on Mythology, whereas LLM-SigCLIP + GCN improves Animal to 0.60 while maintaining strong Mythology performance. Similarly, Base-VAE + GCN reaches 0.49 on Animal and 0.38 on Tree, while LLM-VAE + GCN improves Animal to 0.60 and maintains competitive Tree performance at 0.36.

The gains are even more pronounced for the autoencoder variants. Base-VAE + GAE achieves only 0.33 on Animal and 0.47 on Mythology, whereas LLM-VAE + GAE improves these to 0.48 and 0.78, respectively. This large margin on Mythology indicates that LLM-enhanced representations better encode semantic structure that benefits reconstruction-based graph objectives. Likewise, LLM-SigCLIP + VGAE improves Animal from 0.51 (Base) to 0.62 and stabilizes performance across other labels.

Across heterogeneous user–item graphs, the same pattern holds. While competitive baselines such as MambaRec, HPMRec, SMORE, and CMDL remain within the 0.57–0.67 range on Mythology, the LLM-enhanced GeMi variants reach up to 0.78. These results confirm that feature quality, rather than solely graph architecture, is a critical factor in performance.

LLM + SigCLIP vs LLM + VAE Comparing the two LLM-enhanced feature pipelines reveals complementary behavior across labels. On Animal, LLM-SigCLIP + VGAE achieves 0.62, slightly outperforming LLM-VAE + GCN at 0.60, suggesting that discriminative contrastive alignment benefits visually grounded categories. On Mythology, LLM-VAE + GAE achieves the strongest result at 0.78, exceeding LLM-SigCLIP variants, indicating that generative multimodal alignment captures abstract narrative semantics more effectively. For Tree, LLM-SigCLIP + GCN reaches 0.36, while LLM-VAE + GCN attains a similar 0.36 in heterogeneous settings, showing that both pipelines can capture structural category cues when paired with graph propagation.

Overall, the results demonstrate that LLM-enhanced vision-language representations consistently outperform their base counterparts across all GeMi backbones. The improvements are particularly significant for semantically rich labels such as Mythology, where LLM conditioning provides clearer contextual signals.

Model	Animal	Mythology	Tree
SigCLIP + GCN	0.48 ± 0.25	0.70 ± 0.21	0.13 ± 0.15
SigCLIP + GAE	0.54 ± 0.28	0.52 ± 0.17	0.20 ± 0.17
SigCLIP + VGAE	0.51 ± 0.25	0.62 ± 0.19	0.18 ± 0.18

Table 10: Transductive Base-SigCLIP

Model	Animal	Mythology	Tree
VAE + GCN	0.49 ± 0.23	0.58 ± 0.20	0.38 ± 0.24
VAE + GAE	0.33 ± 0.20	0.47 ± 0.21	0.23 ± 0.20
VAE + VGAE	0.34 ± 0.20	0.53 ± 0.13	0.20 ± 0.18

Table 11: Transductive Base-VAE features

Inductive experiments Following the transductive analysis, we examine whether LLM-enhanced features improve inductive performance within GEMi. Tables 12–13 report inductive results using base SigCLIP and base VAE features, while Tables 6–9 report inductive results using LLM-enhanced pipelines (LLM-SigCLIP / LLM-VAE), under the same inductive protocol. In general, LLM-enhanced features outperformed base features, with the magnitude depending on the labels.

- **Animal:** Under the SigCLIP pipeline, LLM-enhancement improves the best GEMi score from 0.52 (Base-SigCLIP + GCN, Table 12) to 0.56 (LLM-SigCLIP + GCN, Table 6), suggesting that the normalized text input can strengthen the vision–language embedding used for inductive neighborhood propagation.
- **Mythology:** The best base-SigCLIP result is 0.57 (Table 12), whereas LLM-SigCLIP reaches 0.70 (Table 6). This indicates that LLM-enhanced textual message are helpful for mythology, where semantic disambiguation is often text-sensitive.
- **Tree:** Improvements are more mixed. Base-SigCLIP peaks at 0.44 (Table 12) and LLM-SigCLIP stays in a similar range (0.44, Table 6), while the VAE pipeline shows a larger jump from 0.51 (Base-VAE + VGAE, Table 13) to 0.56 (LLM-VAE + GCN, Table 7). This is consistent with Tree being a visually subtle category.

We further compare the two LLM-enhanced feature pipelines under the inductive setting: LLM-SigCLIP (Table 6 and Table 8) versus LLM-VAE (Table 7 and Table 9). Overall, the better pipeline is label-dependent, and the pattern differs slightly between homogeneous item–item and heterogeneous user–item graphs. In homogeneous graphs, we observe the following:

- **Animal:** LLM-SigCLIP has the strongest GEMi configuration for Animal (0.56), while LLM-VAE provides comparable baselines but does not surpass the best GEMi score in this label.
- **Mythology:** LLM-SigCLIP is clearly stronger: GEMi-VGAE achieves 0.70, whereas the best LLM-VAE GEMi result is 0.61.
- **Tree:** The trend reverses: LLM-VAE performs better for Tree under item–item inductive graphs, with GEMi-GCN reaching 0.56 compared to the best LLM-SigCLIP

Model	Animal	Mythology	Tree
SigCLIP + GCN	0.52 \pm 0.26	0.57 \pm 0.14	0.39 \pm 0.27
SigCLIP + GAE	0.50 \pm 0.25	0.54 \pm 0.23	0.37 \pm 0.27
SigCLIP + VGAE	0.51 \pm 0.35	0.50 \pm 0.27	0.44 \pm 0.21

Table 12: Inductive Base-SigCLIP

Model	Animal	Mythology	Tree
VAE + GCN	0.47 \pm 0.27	0.51 \pm 0.28	0.45 \pm 0.23
VAE + GAE	0.39 \pm 0.23	0.61 \pm 0.22	0.30 \pm 0.25
VAE + VGAE	0.51 \pm 0.25	0.33 \pm 0.18	0.51 \pm 0.30

Table 13: Inductive Base-VAE features

GEM_I around 0.44. This indicates that the VAE fusion maybe able to better preserve the background semantics for Tree in the inductive regime.

On the other hand, in heterogeneous settings, the two LLM-enhanced pipelines have closer performance. For the animal concept, LLM-SigCLIP attains 0.71 versus 0.69 for LLM-VAE. For the mythology concept, LLM-VAE is slightly higher (0.78) than LLM-SigCLIP (0.77), and for Tree, the best results are essentially tied (0.53 in both user-item Tables).

Overall, inductive results show that LLM enhancement is consistently beneficial over base features, while the relative advantage between LLM-SigCLIP and LLM-VAE is label- and graph-dependent: SigCLIP tends to favor Mythology in homogeneous graphs, whereas VAE is more helpful for visually subtle labels such as Tree

5.4.3 Graph Structure Learning (RQ3)

We examine whether transductive or inductive graph training yields better recommendation performance. We analyze this by comparing GCN, GAE, and VGAE across (i) homogeneous item-item graphs and (ii) heterogeneous user-item graphs, under both training paradigms.

Homogeneous item-item graphs (Llama-SigCLIP) Comparing Tables 2 (transductive) and the corresponding inductive table shows label-dependent behavior. For the Animal concept, transductive GAE (0.59) and inductive GCN (0.56) are comparable, but transductive GAE remains slightly stronger. For the Mythology concept, inductive VGAE (0.70) marginally surpasses transductive VGAE (0.69), indicating that for semantically rich categories, inductive learning can generalize well when graph structure is stable. However, for the Tree concept, transductive GCN (0.62) clearly outperforms inductive GCN (0.44), suggesting that dense propagation over the full graph during training benefits structurally sparse labels.

Homogeneous item-item graphs (Llama-VAE) Under VAE features (Table 3), transductive training is consistently stronger. For the Animal concept, transductive VGAE (0.62) exceeds inductive variants (around 0.50). For Mythology, transductive

GAE achieves 0.78, substantially higher than inductive GAE/VGAE (around 0.61). For Tree, transductive GCN (0.36) and inductive GCN (0.56) show complementary strengths, but the inductive advantage is concentrated in GCN only, while GAE and VGAE favor transductive settings.

Heterogeneous user–item graphs For Llama-SigCLIP features (Table 4), transductive GeMi variants consistently outperform inductive baselines. Transductive GCN reaches 0.60 (Animal), 0.65 (Mythology), and 0.36 (Tree), whereas inductive GraphSAGE, PinSAGE, and GATNE-I remain below these values on at least two labels. The same pattern holds for Llama-VAE features (Table 5), where transductive GAE achieves 0.78 on Mythology, clearly exceeding inductive counterparts (maximum around 0.61).

Model-wise trends (GCN vs GAE vs VGAE) Across settings, GCN benefits most from transductive propagation on structurally oriented labels such as Tree. GAE and VGAE show strong advantages on Mythology under transductive training, particularly with LLM-enhanced features, where reconstruction-based objectives can exploit full-graph connectivity. Inductive training occasionally matches or slightly exceeds transductive performance on specific labels (e.g., Mythology with VGAE in homogeneous SigCLIP), but these gains are not consistently observed across all labels or graph types.

Overall, transductive training provides more reliable and consistently superior performance, especially in heterogeneous user–item graphs and for semantically complex labels such as Mythology. Inductive training can generalize competitively in certain homogeneous settings, but it does not consistently outperform transductive learning. Thus, for the recommendation scenario studied here, transductive graph structure learning yields stronger and more stable results across graph formulations and GeMi variants.

6 Discussion

- In Section 4.1.1, we presented an LLM-based semantic text canonicalization procedure. This has been implemented using TinyLlama (Zhang et al. (2024)) in the GeMi system. Other LLM variants which can replace TinyLlama include BLOOM (Le Scao and Fan (2023)), Gemma2 (Team et al. (2024)), Qwen (Hui and Yang (2024)) and other open source variants. Experimentation with these pretrained models to study their impact on performance is an active area of research.
- Instead of adapting pretrained language models to downstream tasks, a trend in recent years has been to use the *pretrain, prompt and inference* paradigm, reformulating downstream recommendation tasks with hard/soft prompts (Liu et al. (2023c)). In this paradigm, fine-tuning can be avoided, and the pretrained model itself can be directly employed to predict panel ratings, or even output sub-tasks related to recommendation targets such as explanations. For example, Geng et al. (Geng et al. (2023)) have utilized a pretrained CLIP component to convert images into tokens. These tokens are added to textual tokens of an item to create a personalized multimodal soft prompt. Different prompting strategies such as Fixed-pretrained

Models, Tuning-free prompting or Prompt and pretrained Tuning can be explored in the GeMi system.

- In the current version of the GeMi system, we have not explored the possibility of content creation using diffusion models (Wei and Fang (2025); Lin et al. (2025a)). However, given that the multimodal data is not always well aligned (for example, the text of the song may exist but there is no illustration available) it may be possible to use generative AI techniques for content creation with the possibility of enhanced personalization of the recommendation system.
- Our current online system – UPFAR (described in Appendix E) can be used to generate data about the users of the system and also learn their preferences. We have incorporated this real-world data into GeMi.
- Investigation of algorithms for dynamic pricing mechanisms for narrative scrolls is likely to benefit the marginalized artists who often suffer from severe economic hardships. This will be addressed in a future version of GeMi.

7 Conclusion

Singing painters are a group of itinerant storytellers who describe moral and mythological stories in long, narrative, scroll paintings. This performing art is now endangered, with only several hundred painters still practicing the art form. An effort is being made to conserve the artwork and lyrics of their songs for future generations, first by extensive field surveys and data collection, followed by careful annotation and processing of the data leading to the development of a recommendation system – GeMi. This multimodal, graph-based recommendation system has comparable performance with state-of-the-art transductive and inductive graph learning algorithms and incorporates user preferences as in collaborative filtering. The GeMi system has great potential to be extended with generative AI capabilities, or prompt-based recommendations, and with the incorporation of pricing models to enable the purchase of narrative scrolls from an online platform.

Acknowledgments. The first author acknowledges support from the following sources: Gnam Award from the Asia Research Institute, a faculty grant from the Office of International Education at SUNY, Buffalo, and a fellowship from the National Endowment for the Humanities and the American Institute of Indian Studies. The authors would like to thank Daricha Foundation for its help with data collection during COVID, Jishnu Basak for editing images, and students in a graduate class on “Optimization Methods in Machine Learning” and “Distributed Computing and Big Data Technologies” for help with processing data collected from field research and annotations. Finally, this work derives inspiration from many discussions with Elise Auerbach, James Nye, and colleagues from the Digital India Learning Initiatives at the University of Chicago.

Appendix A GNN-based Recommendation Systems

GNNs have been extensively used to analyze graph-structured data for node classification (Kipf and Welling (2017); Zhu et al. (2021)), link prediction (Zhang and Chen (2018)), and information retrieval tasks (Zhang et al. (2021c)). In recent years, their use in recommendation systems (Wu et al. (2022b); Gao et al. (2023); Anand and Maurya (2024); Sharma et al. (2024)) has become prevalent. A large-scale recommendation system typically has three stages: (a) Matching - filtering candidate items from a very large (usually in millions) set of items, reducing the scale and complexity of the problem; (b) Ranking - Scoring candidate items by a ranking model, and (c) Re-ranking - Updating the ranked list to satisfy additional criteria or business needs. GNNs⁴ are known to have been used for all three stages.

A.0.1 Use of GNNs for Matching

In the matching stage, GNN-based models are often used to develop sophisticated embedding representations of architectures (Zhang et al. (2021c); Gao et al. (2023); Wang et al. (2019); He et al. (2020); Wu et al. (2021); Sun et al. (2020a); Ying et al. (2018); Liu et al. (2021a, 2023b); Yu and Qin (2020); Mu et al. (2022); Liu et al. (2023a)). These models can capture higher-order similarity among users and items and also capture structural connectivity. A guiding principle for the development of these models is that users with similar interactions tend to have similar preferences, and these can be elucidated through multiple rounds of information propagation. van den Berg et al. (2018) proposes a graph auto-encoder framework based on differentiable message passing on a bipartite, user-item interaction graph representing movie ratings. In Neural Graph Collaborative Filtering (NGCF)(Wang et al. (2019)), the authors argue that embeddings obtained from pre-defined user-item features do not sufficiently capture the collaborative filtering effect; instead they exploit the user-item graph structure by propagating embeddings on it. Sun et al. (2020a) proposes Neighbor Interaction Aware (NIA)-GCN, which can explicitly model the relational information between neighbor nodes and exploit the heterogeneous nature of the user-item bipartite graph. Wang et al. (2020) claim that user-item interactions are often modeled in a uniform manner, which fails to model diverse relationships and disentangle user intents in representations. Their proposed model, Disentangled Graph Collaborative Filtering (DGCF), solves this problem by modeling a distribution over intents for each user-item interaction, and iteratively refining the intent-aware interaction graphs and representations. EdGe-wise mOdulation (EGO, Chen et al. (2022)) fusion enhances user-item interactions by adaptively distilling edge-wise multimodal information and learning to modulate unimodal nodes under the supervision of other modalities. Adaptive Anti-Bottleneck MultiModal Graph Learning Network (Cai et al. (2022a)) designs a collaborative representation learning module and a semantic module for learning representations of users and micro-videos, respectively. Spurious correlations are learned in Du et al. (2022).

⁴A comprehensive review of GNNs and their use is outside the scope of this work, but we refer the interested reader to the following surveys: (Wu et al. (2022b); Gao et al. (2023); Anand and Maurya (2024); Sharma et al. (2024))

Liu et al. (2021a) notice that GCN-based recommendation models suffer from over-smoothing, and high-order neighboring users with no common interests of a user can be involved in the user’s embedding learning in the graph convolution operation. Consequently, to rectify this, a novel Interest-aware Message-Passing GCN (IMP-GCN) recommendation model, which performs high-order graph convolution inside sub-graphs, is proposed. HUIGN (Wei et al. (2022)) exhibits user intents in a hierarchical graph structure.

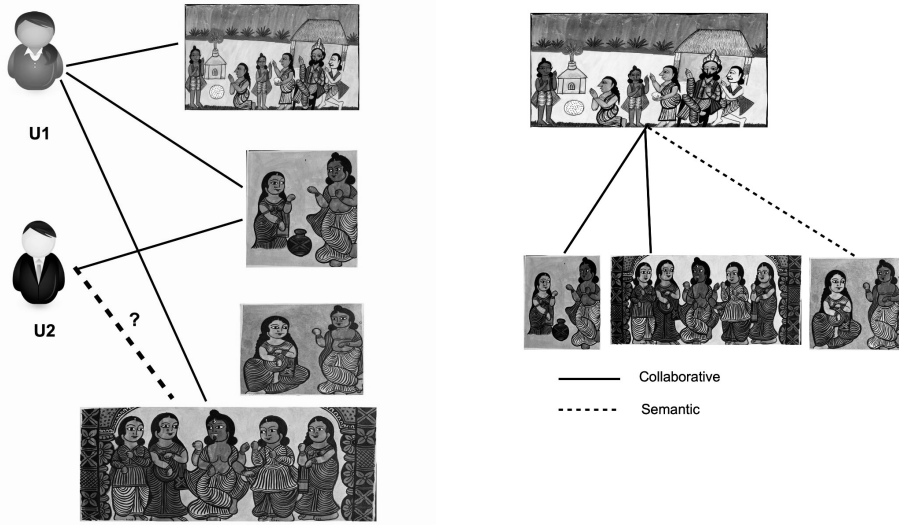
Behavior-driven Feature Adapter (BeFA Fan et al. (2025)) reduces information drift and omission among extracted features. Multi-level sElf-supervised learNing for mulTimOdal Recommendation (MENTOR, Xu et al. (2025b)) addresses the label sparsity problem and the modality alignment problem by enhancing features and the GCN. EliMRec (Liu et al. (2022c)) was designed to eliminate single-modal bias. Dual-GNN (Wang et al. (2023a)) addresses multimodal fusion problems arising from missing modalities. Modality alignment for matching is explored in PretrAining Modality-Disentangled Representations Model (PAMD Han et al. (2022)). In Self-Supervised Graph Learning (SGL) (Wu et al. (2021)), multiple views of a node are generated to maximize the agreement between different views of the same node compared to that of other nodes. This is assumed to rectify problems arising from higher degree nodes exerting a larger influence on recommendations and noisy interactions. MMGCL (Yi et al. (2022)) design modality edge dropout and modality masking to generate multiple views of users and items.

DGMRec (Kim et al. (2025)) disentangles modality features from an information-based perspective and generates missing modality features by integrating aligned features from other modalities and leveraging user preferences for modalities. Modality alignment is further explored in Tran and Lauw (2022). COHESION (Xu et al. (2025c)) promotes a dual-stage fusion strategy to reduce the impact of irrelevant information, refining all modalities. GUIDER (Li et al. (2025a)) denoises user feedback. MoDiCF (Li et al. (2025b)) rectifies the difficulty in accurately generating missing data due to the limited ability to capture modality distributions.

Hamming Spatial GCN (HS-GCN) (Liu et al. (2023b)) develops a hash framework to model both the first- and the high-order similarities in the Hamming space through the user-item bipartite graphs. Yu and Qin (2020) observe that when dealing with very large graphs, the convolution operation is over-simplified to make it computationally efficient, but this over-simplification impairs its normal function. Therefore, a Low Pass Collaborative Filter (LCF) is designed to remove the noise caused by quantization in the observed data, which also reduces the complexity of graph convolution in an unscathed way. To deal with the large size of modality encoders and the complexity of fusing high dimensional features, Zhuang et al. (2025) proposes the use of a light-weight Wasserstein Knowledge Distillation framework.

While GNN models can be used efficiently to inject higher-order connections into embeddings in user-item bipartite graphs, they are known to have several limitations:

- **Modeling semantic relationships between items:** Figure A1 illustrates an example from GeMi wherein item-item semantic relationships are as important in understanding latent graph structures as collaborative signals. In the example, panels {1, 2, 4} are of interest to the first user, and {2} has only been liked by the



(a) Collaborative item relations

(b) Semantic relationships between items

Fig. A1: A toy example from GeMi illustrating two types of item relations – collaborative and semantic. Both can be used to learn the latent graph structures.

second; the collaborative signal helps to learn whether panel {4} is of interest. In addition, since the same character is being discussed in all four panels, there is a semantic relation between panel {1} and {3} which is not inherently captured by this collaborative signal, and the multimodal representation must be able to learn. We are aware of only a few works that capture this semantic information (such as Meng et al. (2025a)).

- **Handling state-of-the-art vision-language models and multimodal fusion:** Many recommendation systems use multimodal information about users and items to model preferences (Xu et al. (2025c) and references there-in), and we are aware of at least two open-source repositories/toolboxes for implementing multimodal system models (MMRec (Zhou et al. (2023a,c)) and Cornac (Salah et al. (2020); Truong et al. (2021b,a))). For example, LGMRec (Guo et al. (2024)) first constructs a local graph embedding module to independently learn collaborative and modality-related embeddings of users and items with local topological relations. Then, a global hypergraph embedding module is designed to capture global user and item embeddings by modeling insightful global dependency relations. Liu et al. (2021b) proposes a pretraining strategy to learn item representations by considering both item side information and their relationships. Multimodal Graph Attention Networks (MGAT) (Tao et al. (2020)) tries to disentangle personal interests at the level of modality. Consequently, multimodal interaction graphs are first built for

each modality; following this a gated attention mechanism is used to identify varying importance scores of different modalities to user preference. In the Multimodal Knowledge Graph Attention Network (MKGAT) (Sun et al. (2020b)), a multimodal graph attention technique is used to propagate information on knowledge graphs. In PGMT (Liu et al. (2021b)), the pretrained Inception-v4 network is used to extract image features, and the pretrained BERT model is used to extract sentence features. LATTICE (Zhang et al. (2021a)) learns item-item structures for each modality⁵ and aggregates multiple modalities to obtain latent item graphs. To explore highly complicated relationships amongst users a modality-aware Heterogeneous Information Graph (M-HIG) is used to generate high-quality user and micro-video embeddings (Cai et al. (2022b)). The model has two key components: (a) a heterogeneous graph that utilizes a random walk-based sampling strategy to sample neighbors for users and micro-videos, and (b) a hierarchical intra- and inter-feature aggregation network to better capture the complex structure and rich semantic information in the heterogeneous graph.

Along with GNNs, several algorithms exist which use Graph Auto-Encoders or Variational Graph Auto-encoders (VGAE) for feature extraction. In GraphCAR (Xu et al. (2018)), two GCNs are used to auto-encode the latent vector space of users and items, and preference scores are obtained by taking the inner product of the two latent factor vectors. In MVGAE (Yi and Chen (2022)), modality-specific variational encoders learn a Gaussian variable for each node – the mean vector represents semantic information, and the variance vector denotes the noise level of the corresponding modality. In addition, modality-specific Gaussian node embeddings are fused according to the product-of-experts principle, where the semantic information in each modality is weighted based on the estimated uncertainty level. We explored the use of MVGAE for the purpose of ranking in the GeMi system. We use one discriminative (Contrastive Language Image Pretraining (CLIP)) and one generative (VAE) vision language model in the GeMi system (described in detail in Section 4.1). Finally, we are aware of several non-GNN based techniques (such as PAMD) which are not directly relevant to the discussion here.

- **Exploring enhanced linguistic (such as sentence- or discourse-level) properties and controllability (such as prompting techniques) of embeddings using pretrained Large Language Models:** LLMs⁶ have been used extensively as support tools for recommendation (Wei et al. (2024); Chen et al. (2025b); Ji et al. (2025); Meng et al. (2025b)), such as to develop prompts using user-item data, item metadata, and the candidate items generated by other multimodal recommendation systems. The use of LLMs as supportive components within existing recommendation systems is less explored (Fioretti et al. (2025)). The GeMi framework explores this supporting role of an LLM to explore sentence and discourse-level linguistic properties of pretrained models (discussed in Section 4.1.1).

⁵Visual features are obtained from published literature, and textual features embeddings are obtained by concatenating the title, descriptions, categories, and brand of each item utilizing sentence-transformers to obtain a 1,024-dimensional representation.

⁶We point the reader to several surveys (Lin et al. (2024); Huang et al. (2025); Wu et al. (2024); Shehmir (2025)) on where and how LLMs can be adapted for use in recommendation systems. This is outside the scope of the current work.

A.0.2 Use of GNNs for ranking

In the ranking phase of recommendation, a small set of candidate items are identified and a ranking score generated. GNN-based ranking models (Li et al. (2019); Su et al. (2020); Guo et al. (2021); Li et al. (2021); Wu et al. (2022a); Liu et al. (2022a)) usually consist of two components - an encoder that captures the desired feature interactions and a decoder (predictor) where the ranking score is estimated by integrating different feature embeddings. Li et al. (2019) show that simple unstructured combination of feature fields often limits the ability to model sophisticated interactions between features and thus propose to represent the multi-field features in a graph structure, where each node corresponds to a feature field and different fields can interact through edges. Zheng et al. (2023) study the effect of price in recommendation models by recommending Price-aware user preference modeling (PUP) wherein a heterogeneous graph comprising of four different kinds of nodes – users, items, prices, and categories is first constructed and embeddings are propagated on them. Finally, a pairwise interaction based decoder is built, leveraging ideas from Factorization Machines (Rendle (2010)). Liu et al. (2022b) design an Attribute-aware Attentive Graph Convolutional Network (A^2 -GCN) which accounts for missing attributes and allows users to have different preferences for items. Su et al. (2022) proposes a hypergraph neural network based model, which generates feature interactions of arbitrary order that help learn recommendations. Guo et al. (2021) design a dual graph enhanced embedding that exploits divide-and-conquer and curriculum-learning-inspired organized learning to address feature sparsity and user-behavior-induced sparsity. Li et al. (2021) develops a Sequence-aware Heterogeneous graph neural Collaborative Filtering (SHCF) model which incorporates high-order heterogeneous collaborative signals and sequential information to eliminate the data sparsity problem and incorporate dynamic user preferences. The GeMi framework explores three different graph architectures – Graph Convolutional Networks (GCN), Graph Auto Encoders (GAE) and Variational Graph Auto Encoders (VGAE), which are discussed in detail in Section 4.5.

A.0.3 Use of GNNs for Re-ranking

Re-ranking has been widely accepted as improving recommendation performance on large-scale commercial platforms such as Amazon and Google (Pei et al. (2019); Liu et al. (2020)). Once candidate items are generated and ranked, a re-ranking algorithm is usually run on top candidates to further refine recommendation results. In **Item Relationship Graph Neural Networks for Personalized Recommendation** (IRGPR Liu et al. (2020)) two graphs - an item relationship graph and an user-item scoring graph are fused to obtain diverse semantic information from heterogeneous node representations. A fully connected GNN is then used to re-rank candidate items. While neural models can be used in the re-ranking phase, it is not uncommon to have simple rule-based functions developed from user preferences or item relationships. In GeMi, we use a label-consistency rule as a re-ranking function.

Model	Parameters
LATTICE	Graph convolution layers: 2; Graph mixing coefficient (λ): 0.6; Loss function: Bayesian Personalized Ranking (BPR); Optimizer: Adam; Learning rate: 5×10^{-4} ; Weight decay: 1×10^{-4} ; Activation: linear graph propagation; Embedding normalization: ℓ_2 normalization;
PMGT	Neighbors per sequence: 32; Hidden size: 128; Transformer layers: 4; Attention heads: 4; Feed-forward dimension: 256; Activation: GELU; Dropout: 0.1; Loss functions: BCE-with-logits + MSE; Masking ratio: 0.16; Random replacement ratio: 0.02; Optimizer: AdamW; Learning rate: 8×10^{-4} ; Weight decay: 1×10^{-4}
HUIGN	Hierarchical intent layers ($K^{(1)}, K^{(2)}, K^{(3)}$): [32, 8, 4]; Co-interaction threshold: ≥ 2 ; Loss terms: BPR + entropy regularization + independence regularization + ℓ_2 weight decay; Entropy weight (λ_1): 0.1; Independence weight (λ_2): 0.1; Weight decay (λ_3): 1×10^{-5} ; Optimizer: Adam; Learning rate: 1×10^{-4}
DualGNN	Graph propagation depth: 2; Graph mixing coefficient (α): 0.6; Contrastive temperature (τ): 0.2; Node feature dropout: 0.2; Loss terms: InfoNCE alignment + orthogonality regularization + graph smoothness regularization; Loss weights: $\lambda_{\text{align}}=1.0$, $\lambda_{\text{ortho}}=0.05$, $\lambda_{\text{smooth}}=0.05$; Optimizer: Adam; Learning rate: 8×10^{-4} ; Weight decay: 1×10^{-4}
FREEDOM	Multimodal graph mixing: $\alpha_v=0.1$ (image) and $\alpha_t=0.9$ (text); Item-Item propagation layers (L_{ii}): 1; User-Item propagation layers (L_{ui}): 2; Degree-sensitive edge pruning ratio (ρ): 0.8; Training loss: BPR + multimodal auxiliary BPR regularizers; Regularization weight (λ): 1×10^{-3} ; Optimizer: Adam; Learning rate: 1×10^{-3} ; Weight decay: 1×10^{-5}
MICRO	Item-view layers: 1 (per modality); Graph mixing coefficient (α): 0.7; Contrastive objective: InfoNCE (temperature $\tau=0.5$); Orthogonality regularizer: cross-view residual decorrelation; Loss weights: $\beta_c=0.1$ (contrastive), $\beta_s=0.1$; Optimizer: Adam; Learning rate: 5×10^{-4} ; Weight decay: 1×10^{-4}
DGVAE	Modality fusion weight (α_v): 0.1 (visual), $1 - \alpha_v$ (text); Graph convolution layers: 2; Number of latent factors (K): 4; Soft assignment temperature (τ): 0.1; Prior scale (σ_0): 0.1; Objective: BCE-with-logits + KL regularization + mutual-information proxy regularization; Regularization weights: $\beta=0.1$ (KL), $\lambda_{\text{MI}}=0.2$ (MI); Optimizer: Adam; Learning rate: 5×10^{-4} ; Weight decay: 1×10^{-4}

Table A1: Transductive Homogeneous Item-Item Baseline Parameters

Appendix B Parameters of Baseline Transductive Algorithms

Tables A1 and A2 present the parameters of the homogeneous and heterogeneous transductive baseline algorithms described in Section 5.2.3.

Model	Parameters
MambaRec	UI graph layers (L_{ui}): 2; Item-Item layers (L_{mm}): 2; Dimensionality reduction factor: $r = 2$; Normalization + dropout: LayerNorm(d) + Dropout(0.2); Loss: BPR pairwise ranking loss + ℓ_2 regularization + contrastive InfoNCE + optional MMD alignment; Contrastive weight: 1.0; InfoNCE temperature: 0.2; MMD weight: 1×10^{-2} ; MMD kernel: RBF with bandwidth $\sigma = 1.0$; Regularization weight: 1×10^{-3} ; Optimizer: Adam; Learning rate: 1×10^{-3} ; Weight decay: 1×10^{-4}
PGL	UI graph layers (L_{ui}): 2; Item-Item layers (L_{mm}): 2; Dropout: 0.2; Loss: BPR pairwise ranking + contrastive InfoNCE loss; Contrastive temperature (τ): 0.2; Regularization weight (λ): 1×10^{-3} on contrastive term; Optimizer: Adam; Learning rate: 1×10^{-3} ; Weight decay: 1×10^{-4}
HPMRec	Hypercomplex algebra dimension (A): 4; UI graph layers (L_{ui}): 2; Item-Item multimodal layers (L_{mm}): 2; Hypercomplex interaction: Cayley-Dickson multiplication with scale $\alpha = 0.1$; Dropout: 0.2; Loss: BPR pairwise ranking + ℓ_2 regularization + self-supervised hypercomplex loss; Reg weight: 1×10^{-3} ; SSL weight: 0.05; Optimizer: Adam; Learning rate: 1×10^{-3} ; Weight decay: 1×10^{-4}
COHESION	UI graph layers (L_{ui}): 2; Multimodal item-item layers (L_{mm}): 2; UI edge dropout: 0.2; Post-fusion stabilizer: LayerNorm + Linear + LeakyReLU; User-User neighborhood aggregation: $k_{uu} = 40$; Loss: BPR pairwise ranking + ℓ_2 regularization; Reg weight: 1×10^{-3} ; Optimizer: Adam; Learning rate: 1×10^{-3} ; Weight decay: 1×10^{-4}
SMORE	UI graph layers (L_{ui}): 2; Item-Item layers (L_{ii}): 2; Spectrum (FFT) convolution: rFFT over feature dimension with learnable complex filters (3 branches: w_a, w_b, w_f); Nonlinear fusion branch: frequency-domain ($\text{FFT}(z) \odot \text{FFT}(z) \odot w_f$) then iFFT; Dropout: 0.2; Auxiliary contrastive loss: InfoNCE; Contrastive weight: 0.15; Temperature (τ): 0.2; Regularization: light ℓ_2 ; Reg weight: 1×10^{-4} ; Optimizer: Adam; Learning rate: 2×10^{-3} ; Weight decay: 1×10^{-4} ; Gradient clipping: max norm 5.0
CMDL	UI graph layers (L_{ui}): 2; Content-view augmentation: feature dropout $p = 0.25$ + Gaussian noise $\sigma = 0.05$; Item content gating; Loss: BPR pairwise ranking + ℓ_2 regularization + contrastive alignment; Contrastive weight: 0.1; Contrastive temperatures: $\tau_{pos} = 0.2$ (items), $\tau_{neg} = 0.2$; Regularization weight: 1×10^{-3} ; Optimizer: AdamW; Learning rate: 1×10^{-3} ; Weight decay: 1×10^{-4}

Table A2: Transductive Heterogenous User-Item Baseline Parameters

Appendix C Parameters of Baseline Inductive Algorithms

Tables B3 and B4 present the parameters of the homogeneous and heterogeneous inductive baseline algorithms described in Section 5.2.3.

Model	Parameters
GraphSAGE	Base percentile: 90; GNN layers: 2; SAGEConv aggregator: max; SAGEConv settings: normalize: False; Hidden channels: 12; Output channels: 2; Activation: ReLU (between layers); Loss: cross-entropy on item labels; Optimizer: Adam; Learning rate: 0.0001;
PinSAGE	Percentile threshold: 90; Random-walk neighborhood: 50; walk length: 5; new neighbor per node(topm): 20; GNN encoder: same as GraphSAGE above (2 layers, hidden 12, out 2, Adam lr 0.0001; cross-entropy).
GATNE-I	kNN on training items ($k = 30$); Similarity floor: 0.0; Feature pre-processing: L2-normalize features before cosine; Embedding dim: 128; Negative samples: 10; Optimizer: Adam; Learning rate: 0.001; Weight decay: 0.000001.

Table B3: Inductive Homogeneous Item–Item Baseline Parameters

Model	Parameters
GraphSAGE	GNN layers: 2; SAGEConv aggregator: max; SAGEConv settings: project=True, normalize=False; Hidden channels: 12; Output channels: 2; Activation: ReLU; Loss: cross-entropy on item labels; Optimizer: Adam; Learning rate: 1e-4; Weight decay: default;
PinSAGE	Percentile threshold: 90; Random-walk neighborhood: 50; walk length: 5; new neighbor per node(topm): 20; GNN encoder: same as GraphSAGE above (2 layers; hidden 12; out 2; Adam lr 1e4; cross-entropy)
GATNE-I	kNN on training items ($k = 30$); Similarity floor: 0.0; Feature preprocessing: L2-normalize features before cosine; Embedding dim: 128; Negative samples: 10; Optimizer: Adam; Learning rate: 1e-3; Weight decay: 1e-6. No learning on test interactions.

Table B4: Inductive Heterogenous User–Item Baseline Parameters

Appendix D Ablation Studies

This section presents ablation studies on various parameters of the GeMi system in both transductive and inductive settings.

Transductive experiments

We conduct a sensitivity analysis focusing specifically on the shared hyperparameters presented in Table 1. The objective is to quantify the robustness of GeMi under perturbations of training regularization, graph sparsity, supervised balancing strength, and evaluation configuration. By varying top- k recommendation cutoff, edge dropout rate, focal loss parameters, while keeping architecture-specific components fixed, we isolate the contribution of these global design choices.

We analyze the sensitivity of GeMi with respect to the Top- k recommendation cutoff in the transductive setting (Fig. D2). Across all variants, moderate values of k yield the most stable and reliable performance trends, while very small or large k lead to suboptimal evaluation behavior.

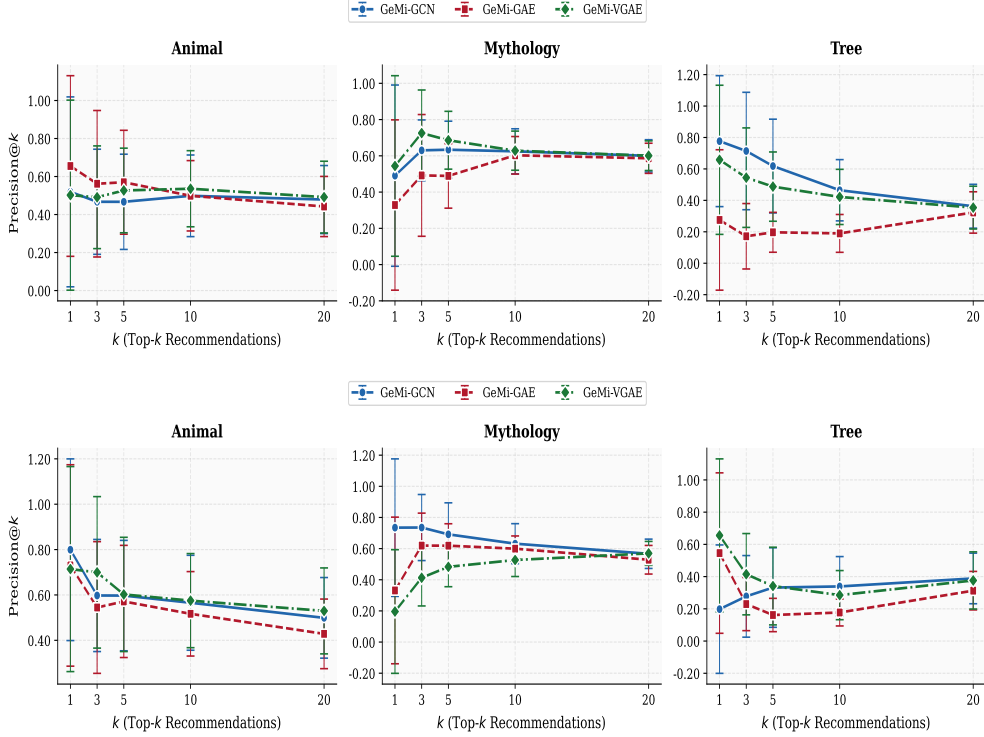


Fig. D2: Effect of the Top- k recommendation cutoff on Precision@5 across Animal, Mythology, and Tree categories for GeMi variants under different feature backbones (LlamaSigCLIP and LlamaVAE). Results are reported as mean \pm standard deviation over multiple runs.

For Animal, Precision@5 is highest at lower k (approximately $k = 1-5$) and gradually declines as k increases, indicating that the most relevant items are ranked early, while larger recommendation lists introduce less relevant candidates, reducing precision. GeMi-GCN achieves the strongest peak performance, followed by GeMi-VGAE, whereas GeMi-GAE shows comparatively larger degradation as k increases.

For Mythology, performance remains relatively stable across increasing k , with slight improvements at moderate values ($k = 5-10$) before marginal decline. This suggests that semantically abstract categories benefit from a slightly broader recommendation list, though larger k introduces diminishing returns due to inclusion of lower-relevance items. GeMi-VGAE demonstrates greater robustness across the entire range of k .

For Tree, which is comparatively sparse and challenging, Precision@5 is lower at very small k and improves with moderate values ($k = 5-10$), after which it stabilizes. This indicates that a slightly larger candidate pool helps retrieve relevant items, but further increases introduce noise. GeMi-VGAE again exhibits more stable behavior, while GeMi-GAE is the most sensitive.

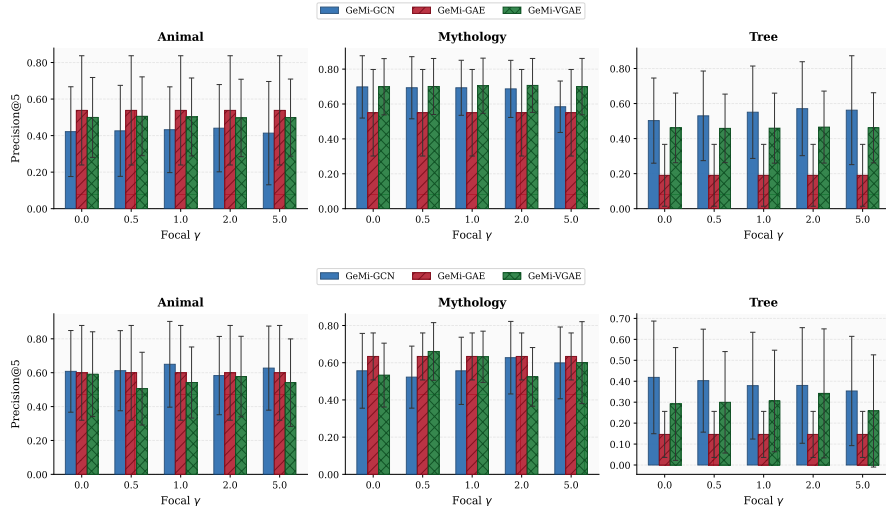


Fig. D3: Effect of the focal loss focusing parameter γ on Precision@5 across Animal, Mythology, and Tree categories for GeMi variants under different feature backbones (LlamaSigCLIP and LlamaVAE). Results are reported as mean \pm standard deviation over multiple runs.

Overall, the results highlight a trade-off governed by the recommendation cutoff: smaller k emphasizes high-confidence predictions, while larger k increases coverage at the expense of precision. GeMi exhibits controlled sensitivity to this parameter, with consistent performance trends across variants. GeMi-GCN achieves the strongest peak performance, whereas GeMi-VGAE provides greater robustness across varying k , supporting the use of moderate Top- k values for balanced and effective recommendation performance.

We further analyze the sensitivity of GeMi with respect to the focal loss focusing parameter γ , using both LlamaSigCLIP and LlamaVAE feature backbones, as illustrated in Fig. D3. Overall, the models exhibit relatively stable behavior across a broad range of γ values, indicating that GeMi is not overly sensitive to moderate variations in the focal modulation strength.

For the LlamaSigCLIP setting, the Animal category shows marginal variation across $\gamma \in [0, 5]$, with GeMi-GAE and GeMi-VGAE maintaining stable performance, while GeMi-GCN remains slightly lower but consistent. This suggests that for visually grounded and relatively well-separated categories, the benefit of increasing γ is limited, as class imbalance is already sufficiently handled by the graph structure and feature quality. In contrast, for Mythology, a slight improvement is observed as γ increases from 0 to around 1–2, after which performance saturates. This indicates that moderate focusing helps emphasize harder samples in semantically abstract categories, but excessive focusing yields diminishing returns. For Tree, performance remains relatively stable with small gains at moderate γ , suggesting that focal reweighting provides mild benefits for sparse categories but does not drastically alter performance trends.

For the LlamaVAE setting, a similar but slightly more pronounced trend is observed. In the Animal category, GeMi-GCN benefits from moderate γ (around 1), achieving peak performance, while higher values introduce slight instability, likely due to over-emphasis on difficult samples. GeMi-GAE and GeMi-VGAE remain relatively stable, though GeMi-VGAE shows minor degradation at extreme γ values. In the Mythology category, both GeMi-GAE and GeMi-VGAE achieve improved performance at moderate γ (around 0.5–1), reflecting better handling of hard and ambiguous samples, while very high γ leads to marginal decline due to over-focusing. For the Tree category, which is the most imbalanced, moderate γ improves performance for GeMi-VGAE, but GeMi-GCN shows slight degradation as γ increases, indicating that excessive focusing may suppress useful majority-class signals.

Across both feature settings, the results highlight that γ controls a trade-off between emphasizing hard samples and maintaining stable gradient contributions from easier examples. Lower values ($\gamma \approx 0$) reduce focal loss to standard BCE, leading to more uniform learning, while higher values ($\gamma > 2$) over-amplify difficult samples, potentially introducing optimization instability. Empirically, moderate values ($\gamma \in [0.5, 2]$) consistently yield the best or near-optimal performance across categories and model variants.

In summary, GeMi demonstrates robustness to the focal loss parameter, with performance varying smoothly across γ . Moderate focusing provides consistent benefits, particularly for semantically complex and imbalanced categories, while extreme values are unnecessary and can slightly degrade performance. These findings support the choice of $\gamma = 2.0$ as a balanced and effective default across different feature backbones and GeMi variants.

We finally analyze the sensitivity of GeMi to edge dropout, which induces stochastic sparsification of the interaction graph during training, under both LlamaSigCLIP and LlamaVAE feature backbones (Fig. D4). Overall, GeMi demonstrates strong robustness, with only minor variations in Precision@5 across dropout rates from 0.0 to 0.4, indicating stable representation learning under structural perturbations.

In the LlamaSigCLIP setting, performance remains largely consistent across all dropout levels. For Animal, GeMi-GCN shows a slight improvement at moderate dropout (~ 0.1), reflecting mild regularization benefits, while higher dropout slightly degrades performance due to excessive edge removal. GeMi-GAE and GeMi-VGAE remain nearly invariant. In Mythology, GeMi-VGAE consistently achieves the best performance (~ 0.69) with negligible variation, while GeMi-GCN exhibits mild gains at moderate dropout. In Tree, GeMi-GCN and GeMi-VGAE remain stable, indicating that moderate sparsification does not hinder information propagation even in sparse regimes.

In the LlamaVAE setting, slightly higher sensitivity is observed. For Animal, GeMi-GCN benefits from moderate dropout (peak around 0.1), suggesting improved generalization via removal of redundant edges, while higher dropout reduces performance. In Mythology, GeMi-GCN peaks at moderate dropout (~ 0.2), indicating effective noise filtering, whereas excessive sparsification degrades performance. GeMi-GAE remains stable, and GeMi-VGAE shows consistent but lower performance. In

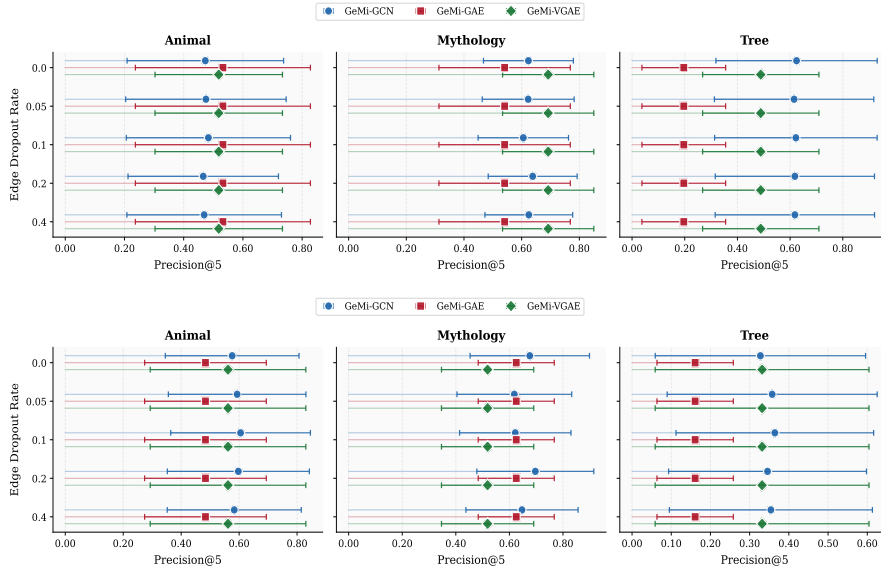


Fig. D4: Effect of the edge dropout parameter on Precision@5 across Animal, Mythology, and Tree categories for GeMi variants under different feature backbones (LlamaSigCLIP and LlamaVAE). Results are reported as mean \pm standard deviation over multiple runs.

Tree, moderate dropout again benefits GeMi-GCN, while higher dropout reduces connectivity and slightly harms performance.

Overall, the results highlight a trade-off between regularization and information loss. Low dropout preserves full connectivity but retains noise, moderate dropout (0.05–0.2) improves robustness and generalization, and high dropout (≥ 0.4) weakens message passing. GeMi-VGAE exhibits the highest robustness, while GeMi-GCN benefits most from moderate regularization. These findings support the use of moderate edge dropout as an effective and stable design choice.

Inductive experiments

We study the sensitivity of inductive performance to the kNN graph parameter k , which controls (i) the density of the training item-item subgraph and (ii) the connection of each test item to the training set. Across labels and backbones, the effect of k is non-monotonic: increasing k is not uniformly beneficial, and in some cases smaller k will generate better performance. This result implies a bias-variance tradeoff where larger k would introduces more weakly similar neighbors and may causing the result over smoothing, while too small k can lead to an overly sparse graph where information passing become weaker.

Non-monotonicity: The impact of k can be strongly non-monotonic even when holding the model and feature pipeline fixed. For instance, under a SigCLIP + GAE setting, performance improves markedly when decreasing k from 30 to 10, but drops




Fig. D5: RQ4 (kNN ablation): Precision@5 vs. kNN graph construction parameter k for Animal. Error bars denote mean \pm 1 std across runs.




Fig. D6: RQ4 (kNN ablation): Precision@5 vs. kNN graph construction parameter k for Mythology. Error bars denote mean \pm 1 std across runs.

again at $k = 5$. This implies that the best fit knn for graph density lies in a middle range rather than at extremes.

Label-dependent sensitivity: The optimal k is label-dependent, but the sensitivity pattern is not uniform across backbones and feature pipelines. In many settings, Tree label tends to be more k -sensitive, while Mythology can range from relatively stable to strongly k -sensitive. A reasonable explanation is that Tree is a weaker signal (often visually subtle and not always dominant in text), so the choice of neighborhood size more directly affects whether sufficient informative signal can be aggregated, additionally, this also motivates our Tree-specific edge augmentation.

Feature dependence: The preferred k is also feature dependent. Different embedding leads to different nearest-neighbor quality. A smaller k can preserve higher-quality neighbors in one embedding space but may overly sparsify the graph in another. Therefore, k acts as a smoothing hyperparameter for message passing, and its preferred range depends on the base feature geometry.

Backbone dependence: Compared with GCN and GAE, VGAE tends to show higher sensitivity to k , with performance sometimes shows sharper rises and drops as k varies. This aligns with VGAE’s generative objective, where changes in graph density can reshape the latent structure and lead to local optima.




Fig. D7: RQ4 (kNN ablation): Precision@5 vs. kNN graph construction parameter k for Tree. Error bars denote mean \pm 1 std across runs.



Fig. E8: Characteristics of users working with UPFAR

Appendix E An online tool for folk art recommendation

In ongoing work ([Hotwani and Dutta \(2025\)](#)), the authors have developed a system – User Preferences for Folk Art Recommendation (UPFAR) – to collect information (such as cultural exposure, museum visits, art and music familiarity, experience with prior performances of narrative scroll paintings) and ratings (like(= 1)/not(= 0)) for panels. Figure [E8](#) shows the categories of data collected from the users, Figure [E9](#) the architecture diagram, and the tool deployed online (Figure [E10](#)). Data from this system was ingested into the GeMi pipeline. The following details are relevant:

- Data pre-processing: The user-panel interaction data was cleaned by removing invalid entries, removing duplicates, and keeping only those interactions that belong to the training panels, to avoid any leakage of test information into user profiles.
- Computing preferences: To convert raw ratings into user preference over labels, the ratings were normalized using min-max normalization in $[0, 1]$. Then,
 - For each user, a baseline rating (their average normalized rating) was computed.
 - For each label (animal, myth, tree), the average rating of items that contain that label was computed.

Using this, relative lift is defined as follows:

$$lift = \text{average rating for that label} - \text{user baseline}$$

which implies that if a user rates a concept higher than their average, it gets a positive lift, otherwise negative.

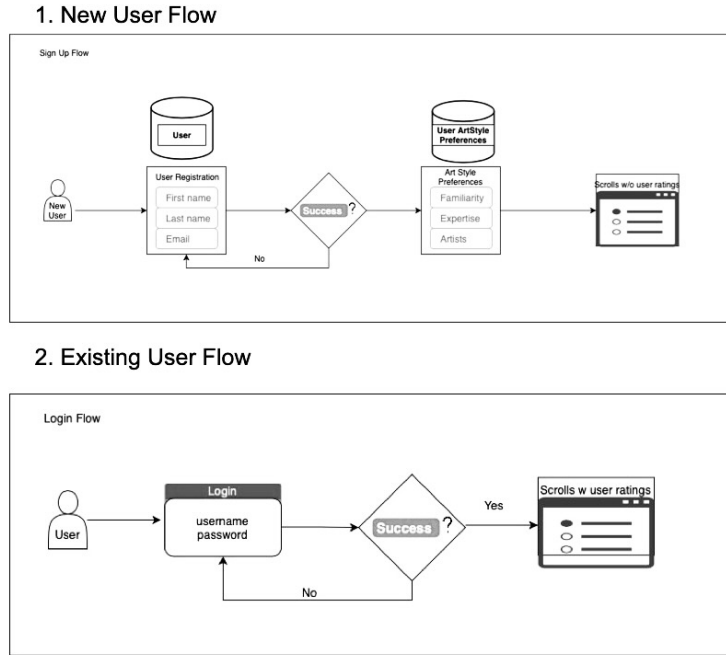


Fig. E9: Architecture diagram of UPFAR

- Dealing with sparsity: Many users have very few interactions for a label. So, a Bayesian-style smoothing was applied that combined the user’s lift with a global prior lift (computed across all users). This is weighted by how many examples the user has, thus preventing unstable estimates for sparse users.
- Developing binary preferences: The lift for each user is mapped into a probability-like preference using a sigmoid function as follows:

$$\text{preference} = \sigma(\text{gain} \times \text{lift}).$$

The gain controls how sharply the preferences move away from 0.5. This gives us a clean preference matrix, where each user has a value between 0 and 1 for animal, myth, and tree concepts. At the same time, for each user, we also store their top-k interacted panels ($k = 5$) based on highest ratings. This is used later to build user embeddings during evaluation.

- Augmentation of user profiles: The real-world data collected from UPFAR is very small. To circumvent this problem, we have generated a large, realistic dataset through bootstrapped augmentation. For each synthetic user (target = 10,000 users), we randomly pick a real user as a base, sample 5 interactions from that user’s top-k items (with replacement), and with some probability ($p = 0.3$), we randomly replace a few items with other observed items to introduce diversity. For preferences, the base user’s base preference vector is used to apply gain scaling (to make some

1. Login

UP-FAR: User Preferences for Folk Art Recommendation

Username Password

[New User Registration](#)

2. New User Registration

User Registration Form

Full Name:

Username:

Password:

Email Address:

Address:

City:

District:

3. Art Style Preferences

Narrative Scroll Painting Preferences

Art Style Information

Names of artists who inspire you:

Degree of interest in Rural/Folk Art Origins:

Familiarity with Mythology:

South Asian

South American

Greek

Roman

Australian

African

Familiarity with Narrative Scroll Painting Techniques from Indian States:

4. Scroll Panel Preferences

Panel 4



Image 7 in Panel 4 in Scroll 1

Overall Rating:

Write Your Review:

Fig. E10: Deployed screens from UPFAR

users more extreme or more neutral), small bias shift (to simulate user tendency) and gaussian noise per label (to diversify), and then clip back values to the desired range $[0, 1]$.

Table E5 presents results from the transductive setting. Our results indicate that graph-structure learning using VGAE has significant benefits for this data, but either SIGCLIP or VAE can be used to learn the concepts. Interestingly, only little benefit is obtained from enhancements via LLM canonicalization (for example to learn mythology concepts using SIGCLIP features). Table E6 presents results from the inductive setting. SIGCLIP with GCN or GAE representations are able to provide good recommendations, while the Tree concept remains hard to recommend for real-world data as well.

References

Anand V, Maurya AK (2024) A survey on recommender systems using graph neural network. ACM Trans Inf Syst 43(1)

Model	Animal	Mythology	Tree
SIGCLIP + GCN	0.52 ± 0.22	0.68 ± 0.17	0.40 ± 0.20
SIGCLIP + GAE	0.55 ± 0.22	0.49 ± 0.14	0.33 ± 0.22
SIGCLIP + VGAE	0.44 ± 0.16	0.74 ± 0.18	0.43 ± 0.22
VAE + GCN	0.27 ± 0.09	0.59 ± 0.04	0.01 ± 0.04
VAE + GAE	0.50 ± 0.22	0.69 ± 0.11	0.22 ± 0.16
VAE + VGAE	0.60 ± 0.11	0.65 ± 0.09	0.60 ± 0.21
Llama SIGCLIP + GCN	0.56 ± 0.19	0.71 ± 0.16	0.44 ± 0.25
Llama SIGCLIP + GAE	0.43 ± 0.20	0.58 ± 0.15	0.40 ± 0.17
Llama SIGCLIP + VGAE	0.59 ± 0.16	0.69 ± 0.15	0.44 ± 0.13
Llama VAE + GCN	0.45 ± 0.19	0.53 ± 0.19	0.39 ± 0.24
Llama VAE + GAE	0.43 ± 0.20	0.43 ± 0.23	0.33 ± 0.16
Llama VAE + VGAE	0.48 ± 0.17	0.39 ± 0.20	0.31 ± 0.22

Table E5: Results from **Transductive** model training using real-world data to guide user-preference matrix generation.

Model	Animal	Mythology	Tree
SIGCLIP + GCN	0.45 ± 0.14	0.79 ± 0.25	0.23 ± 0.23
SIGCLIP + GAE	0.39 ± 0.20	0.58 ± 0.20	0.24 ± 0.26
SIGCLIP + VGAE	0.65 ± 0.27	0.71 ± 0.13	0.30 ± 0.21
VAE + GCN	0.64 ± 0.18	0.61 ± 0.09	0.24 ± 0.30
VAE + GAE	0.53 ± 0.09	0.54 ± 0.13	0.19 ± 0.16
VAE + VGAE	0.60 ± 0.20	0.64 ± 0.19	0.25 ± 0.19
Llama SIGCLIP + GCN	0.56 ± 0.27	0.52 ± 0.19	0.25 ± 0.20
Llama SIGCLIP + GAE	0.49 ± 0.22	0.54 ± 0.13	0.36 ± 0.23
Llama SIGCLIP + VGAE	0.47 ± 0.21	0.59 ± 0.17	0.26 ± 0.14
Llama VAE + GCN	0.31 ± 0.23	0.31 ± 0.15	0.33 ± 0.17
Llama VAE + GAE	0.51 ± 0.18	0.56 ± 0.12	0.38 ± 0.12
Llama VAE + VGAE	0.53 ± 0.19	0.48 ± 0.13	0.37 ± 0.17

Table E6: Results from **Inductive** model training using real-world data to guide user-preference matrix generation.

Aroyo L, Brussee R, Rutledge L, et al (2007) Personalized museum experience: The rijksmuseum use case. In: Museums and the Web 2007

Bajpai LP (2015) Myths and folktales in the patachitra art of bengal: Tradition and modernity. Chitralekha International Journal on Art and Design 5(1):124–131

Bennett J, Lanning S (2007) The netflix prize. In: Proceedings of the KDD Cup Workshop 2007. ACM, pp 3–6, URL <http://www.cs.uic.edu/~liub/KDD-cup-2007/NetflixPrize-description.pdf>

Bentley JL (1975) The complexity of finding fixed-radius near neighbors. Communications of the ACM 18(8):421–425

- van den Berg R, Kipf T, Welling M (2018) Graph convolutional matrix completion. KDD Deep Learning Day 1-7
- Bohnert F, Zukerman I (2009) Using keyword-based approaches to adaptively predict interest in museum exhibits. In: Proceedings of the Australasian Joint Conference on Advances in Artificial Intelligence, p 656–665
- Cai D, Qian S, Fang Q, et al (2022a) Adaptive anti-bottleneck multi-modal graph learning network for personalized micro-video recommendation. In: Proceedings of the ACM International Conference on Multimedia, p 581–590
- Cai D, Qian S, Fang Q, et al (2022b) Heterogeneous hierarchical feature aggregation network for personalized micro-video recommendation. IEEE Transactions on Multimedia 24:805–818
- Cao M, Yang FF, Jin Y, et al (2025) Graph neural network for product recommendation on the amazon co-purchase graph. URL <https://arxiv.org/abs/2508.14059>
- Cen Y, Zou X, Zhang J, et al (2019) Representation learning for attributed multiplex heterogeneous network. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, pp 1358–1368
- Chen F, Wang J, Wei Y, et al (2022) Breaking isolation: Multimodal graph fusion for multimedia recommendation by edge-wise modulation. In: Proceedings of the ACM International Conference on Multimedia, p 385–394
- Chen Y, Wu L, Zaki M (2020) Iterative deep graph learning for graph neural networks: Better and robust node embeddings. In: Advances in Neural Information Processing Systems, vol 33. Curran Associates, Inc., pp 19314–19326
- Chen Z, Xu J, Wang H, et al (2025a) Hypercomplex prompt-aware multimodal recommendation. URL <https://arxiv.org/abs/2508.10753>, [arXiv:2508.10753](https://arxiv.org/abs/2508.10753)
- Chen Z, Xu J, Wang H, et al (2025b) Hypercomplex prompt-aware multimodal recommendation. In: Proceedings of the ACM International Conference on Information and Knowledge Management, p 403–414
- Deldjoo Y, Schedl M, Cremonesi P, et al (2020) Recommender systems leveraging multimedia content. ACM Comput Surv 53(5). <https://doi.org/10.1145/3407190>
- Dimock EC (ed) (1963) The Thief of Love: Bengali Tales from Court and Village. University of Chicago Press
- Du X, Wu Z, Feng F, et al (2022) Invariant representation learning for multimedia recommendation. In: Proceedings of the ACM International Conference on Multimedia, p 619–628

- Dutta G (1939) Patua Sangeet. Kolkata University
- Dutta H, Tiwari A (2024) Content-based art recommendation using multimodal graph neural networks. In: IEEE International Conference on Knowledge Graph (ICKG), pp 67–74
- Fan Q, Yu P, Tan Z, et al (2025) Befa: A general behavior-driven feature adapter for multimedia recommendation. In: Proceedings of the AAAI Conference on Artificial Intelligence
- Fioretti ML, Laterza N, Preziosa A, et al (2025) How powerful are llms to support multimodal recommendation? a reproducibility study of llmrec. In: Proceedings of the ACM Conference on Recommender Systems, p 774–782
- Frost S, Thomas MM, Forbes AG (2019) “art i don’t like”: An anti-recommender system for visual art. *Museums and the Web*
- Gao C, Zheng Y, Li N, et al (2023) A survey of graph neural networks for recommender systems: Challenges, methods, and directions. *ACM Trans Recomm Syst* 1(1)
- Geng S, Tan J, Liu S, et al (2023) Vip5: Towards multimodal foundation models for recommendation. In: Proceedings of the ACM Conference on Recommender Systems
- Grieser K, Baldwin T, Bohnert F, et al (2011) Using ontological and document similarity to estimate museum exhibit relatedness. *J Comput Cult Herit* 3(3)
- Guo W, Su R, Tan R, et al (2021) Dual graph enhanced embedding neural network for ctr prediction. In: Proceedings of the ACM SIGKDD Conference on Knowledge Discovery & Data Mining, p 496–504
- Guo Z, Li J, Li G, et al (2024) Lgmrec: local and global graph learning for multimodal recommendation. In: Proceedings of the AAAI Conference on Artificial Intelligence
- Hamilton WL, Ying R, Leskovec J (2017) Inductive representation learning on large graphs. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, pp 1025–1035
- Han T, Wang P, Niu S, et al (2022) Modality matches modality: Pretraining modality-disentangled item representations for recommendation. In: Proceedings of the ACM Web Conference, p 2058–2066
- Harper FM, Konstan JA (2015) The movielens datasets: History and context. *ACM Trans Interact Intell Syst* 5(4)
- He R, McAuley J (2016) Vbpr: Visual bayesian personalized ranking from implicit feedback. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. AAAI Press, p 144–150

- He X, Deng K, Wang X, et al (2020) 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, p 639–648
- Hotwani S, Dutta H (2025) An online tool to collect user preferences for a folk art recommendation system. MSS Experiential Learning Expo, The State University of New York at Buffalo
- Huang C, Huang H, Yu T, et al (2025) A survey of foundation model-powered recommender systems: From feature-based, generative to agentic paradigms. URL <https://arxiv.org/abs/2504.16420>
- Hui B, Yang Jea (2024) Qwen2.5-coder technical report. URL <https://arxiv.org/abs/2409.12186>, arXiv:2409.12186
- Jain J (ed) (1998) Picture Showmen: Insights into the narrative tradition in Indian Art. Marg Publications, India
- Ji Y, Luo D, Liu C, et al (2025) Generating negative samples for multi-modal recommendation. In: Proceedings of the ACM International Conference on Multimedia, p 6007–6016
- Kim J, Kang H, Kim S, et al (2025) Disentangling and generating modalities for recommendation in missing modality scenarios. In: Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval, p 1820–1829
- Kipf TN, Welling M (2016) Variational graph auto-encoders. URL <https://arxiv.org/abs/1611.07308>
- Kipf TN, Welling M (2017) Semi-Supervised Classification with Graph Convolutional Networks. In: Proceedings of the International Conference on Learning Representations, ICLR '17
- Korom FJ (2006) Village of painters: narrative scrolls from West Bengal. Museum of New Mexico Press, Santa Fe
- Krippendorff K (ed) (2004) Content Analysis, An Introduction to Its Methodology. Sage Publications
- Kuffik T, Stock O, Zancanaro M, et al (2011) A visitor’s guide in an active museum: Presentations, communications, and reflection. J Comput Cult Herit 3
- Lachaud G, Conde-Céspedes P, Trocan M (2023) Comparison between inductive and transductive learning in a real citation network using graph neural networks. In: Proceedings of the 2022 IEEE/ACM International Conference on Advances in Social

Networks Analysis and Mining, p 534–540

- Le Scao T, Fan Aea (2023) Bloom: A 176b-parameter open-access multilingual language model. URL <https://arxiv.org/abs/2211.05100>, [arXiv:2211.05100](https://arxiv.org/abs/2211.05100)
- Li C, Hu L, Shi C, et al (2021) Sequence-aware Heterogeneous Graph Neural Collaborative Filtering, pp 64–72
- Li H, Du H, Li Y, et al (2025a) Teach me how to denoise: A universal framework for denoising multi-modal recommender systems via guided calibration. In: Proceedings of the ACM International Conference on Web Search and Data Mining, p 782–791
- Li J, Wang S, Zhang Q, et al (2025b) Generating with fairness: A modality-diffused counterfactual framework for incomplete multimodal recommendations. In: Proceedings of the ACM on Web Conference 2025, p 2787–2798
- Li R, Wang S, Zhu F, et al (2018) Adaptive graph convolutional neural networks. In: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence. AAAI Press
- Li X, She J (2017) Collaborative variational autoencoder for recommender systems. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, p 305–314
- Li Z, Cui Z, Wu S, et al (2019) Fi-gnn: Modeling feature interactions via graph neural networks for ctr prediction. In: Proceedings of the ACM International Conference on Information and Knowledge Management, p 539–548
- Liang D, Krishnan RG, Hoffman MD, et al (2018) Variational autoencoders for collaborative filtering. In: Proceedings of the World Wide Web Conference, p 689–698
- Lin J, Dai X, Xi Y, et al (2024) How can recommender systems benefit from large language models: A survey. *ACM Trans Inf Syst*
- Lin J, Cao Y, Yu Y, et al (2025a) Diffusion models for recommender systems: From content distribution to content creation. In: Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining, p 6074–6085
- Lin X, Liu R, Cao Y, et al (2025b) Contrastive modality-disentangled learning for multimodal recommendation. *ACM Transactions on Information Systems* 43(3):1–31
- Liu F, Cheng Z, Zhu L, et al (2021a) Interest-aware message-passing gcn for recommendation. In: Proceedings of the Web Conference, p 1296–1305
- Liu F, Cheng Z, Zhu L, et al (2022a) An attribute-aware attentive gcn model for attribute missing in recommendation. *IEEE Transactions on Knowledge and Data*

Engineering 34(9):4077–4088

- Liu F, Cheng Z, Zhu L, et al (2022b) An attribute-aware attentive gcn model for attribute missing in recommendation. *IEEE Transactions on Knowledge and Data Engineering* 34(9):4077–4088
- Liu F, Chen H, Cheng Z, et al (2023a) Disentangled multimodal representation learning for recommendation. *Trans Multi* 25:7149–7159
- Liu H, Wei Y, Yin J, et al (2023b) Hs-gcn: Hamming spatial graph convolutional networks for recommendation. *IEEE Trans on Knowl and Data Eng* 35(6):5977–5990
- Liu P, Zhang L, Gulla JA (2023c) Pre-train, prompt, and recommendation: A comprehensive survey of language modeling paradigm adaptations in recommender systems. *Transactions of the Association for Computational Linguistics* 11:1553–1571
- Liu Q, Wu S, Wang L (2017) Deepstyle: Learning user preferences for visual recommendation. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '17*, p 841–844
- Liu W, Liu Q, Tang R, et al (2020) Personalized re-ranking with item relationships for e-commerce. In: *Proceedings of the ACM International Conference on Information & Knowledge Management*, p 925–934
- Liu X, Tao Z, Shao J, et al (2022c) Elimrec: Eliminating single-modal bias in multimedia recommendation. In: *Proceedings of the ACM International Conference on Multimedia*, p 687–695
- Liu Y, Yang S, Lei C, et al (2021b) Pre-training graph transformer with multimodal side information for recommendation. In: *Proceedings of the ACM International Conference on Multimedia*, p 2853–2861
- Ma J, Zhou C, Cui P, et al (2019) Learning disentangled representations for recommendation. In: *Proceedings of the International Conference on Neural Information Processing Systems*.
- Mair VH (ed) (2019) *Painting and Performance: Chinese Picture Recitation and Its Indian Genesis*. University of Hawaii Press
- Markwald M, Demidova E (2024) Refuel: rule extraction for imbalanced neural node classification. *Machine Learning* 113(9):6227–6246
- McCutchion DJ, Bhowmick SK (1999) *Patuas and Patua Art in Bengal*. Firma KLM
- McPherson M, Smith-Lovin L, Cook JM (2003) Birds of a feather: Homophily in social networks. *Annu Rev Sociol* 27:415–444

- Meng F, Meng Z, Jin R, et al (2025a) Tamer: Interest tree augmented modality graph recommender for multimodal recommendation. In: Proceedings of the ACM International Conference on Multimedia, p 5998–6006
- Meng F, Meng Z, Jin R, et al (2025b) Doge: Llms-enhanced hyper-knowledge graph recommender for multimodal recommendation. Proceedings of the AAAI Conference on Artificial Intelligence 39:12399–12407
- Messina P, Dominguez V, Parra D, et al (2019) Content-based artwork recommendation: integrating painting metadata with neural and manually-engineered visual features. *User Modeling and User-Adapted Interaction* 29(2):251–290
- Mu Z, Zhuang Y, Tan J, et al (2022) Learning hybrid behavior patterns for multimedia recommendation. In: Proceedings of the ACM International Conference on Multimedia, p 376–384
- Ni J, Li J, McAuley J (2019) Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp 188–197
- Ong RK, Khong AWH (2025) Spectrum-based modality representation fusion graph convolutional network for multimodal recommendation. In: Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining. ACM, WSDM '25, p 773–781, <https://doi.org/10.1145/3701551.3703561>, URL <http://dx.doi.org/10.1145/3701551.3703561>
- Pariser E (2011) ‘the filter bubble: What the internet is hiding from you. The Penguin Group
- Pei C, Zhang Y, Zhang Y, et al (2019) Personalized re-ranking for recommendation. In: Proceedings of the ACM Conference on Recommender Systems, p 3–11
- Radford A, Kim JW, Hallacy C, et al (2021) Learning transferable visual models from natural language supervision. In: Proceedings of the 38th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol 139. PMLR, pp 8748–8763
- Ren K, Ju CY, Lee DH (2025) Modality alignment with multi-scale bilateral attention for multimodal recommendation. In: Proceedings of the 34th ACM International Conference on Information and Knowledge Management. ACM, CIKM '25, p 2493–2502, <https://doi.org/10.1145/3746252.3761222>, URL <http://dx.doi.org/10.1145/3746252.3761222>
- Rendle S (2010) Factorization machines. In: Proceedings of the 2010 IEEE International Conference on Data Mining, p 995–1000

- Salah A, Truong Q, Lauw HW (2020) Cornac: A comparative framework for multimodal recommender systems. *Journal of Machine Learning Research* 21(95):1–5. URL <http://jmlr.org/papers/v21/19-805.html>
- Semeraro G, Lops P, De Gemmis M, et al (2012) A folksonomy-based recommender system for personalized access to digital artworks. *J Comput Cult Herit* 5(3)
- Sharma K, Lee Y, Nambi S, et al (2024) A survey of graph neural networks for social recommender systems. *ACM Comput Surv* 56(10)
- Shazeer N (2020) GLU variants improve transformer. *CoRR* abs/2002.05202
- Shehmir KRS. (2025) Llm4rec: A comprehensive survey on the integration of large language models in recommender systems—approaches, applications and challenges. *Future Internet* 17(6)
- Stock O, Zancanaro M, Busetta P, et al (2007) Adaptive, intelligent presentation of information for the museum visitor in peach. *User Modeling and User-Adapted Interaction* 17(3):257–304
- Su J, Ahmed M, Lu Y, et al (2024) Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* 568
- Su Y, Zhang R, Erfani SM, et al (2020) Detecting beneficial feature interactions for recommender systems. In: *AAAI Conference on Artificial Intelligence*, URL <https://api.semanticscholar.org/CorpusID:234771712>
- Su Y, Zhao Y, Erfani S, et al (2022) Detecting arbitrary order beneficial feature interactions for recommender systems. In: *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, p 1676–1686
- Sun J, Zhang Y, Guo W, et al (2020a) Neighbor interaction aware graph convolution networks for recommendation. In: *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, p 1289–1298
- Sun R, Cao X, Zhao Y, et al (2020b) Multi-modal knowledge graphs for recommender systems. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, p 1405–1414
- Tao Z, Wei Y, Wang X, et al (2020) Mgat: Multimodal graph attention network for recommendation. *Information Processing Management* 57(5):102277
- Team G, Riviere M, Pathak Sea (2024) Gemma 2: Improving open language models at a practical size. URL <https://arxiv.org/abs/2408.00118>, [arXiv:2408.00118](https://arxiv.org/abs/2408.00118)
- Touvron H, Lavril T, Izacard G, et al (2023) Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*

- Tran N, Lauw HW (2022) Aligning dual disentangled user representations from ratings and textual content. In: Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining, p 1798–1806
- Truong Q, Salah A, Lauw H (2021a) Multi-modal recommender systems: Hands-on exploration. In: Fifteenth ACM Conference on Recommender Systems, pp 834–837
- Truong Q, Salah A, Tran T, et al (2021b) Exploring cross-modality utilization in recommender systems. *IEEE Internet Computing* 25(4):50–57
- Wang Q, Wei Y, Yin J, et al (2023a) Dualgnn: Dual graph neural network for multimedia recommendation. *Trans Multi* 25:1074–1084
- Wang Q, Wei Y, Yin J, et al (2023b) Dualgnn: Dual graph neural network for multimedia recommendation. *IEEE Transactions on Multimedia* 25:1074–1084
- Wang X, He X, Wang M, et al (2019) Neural graph collaborative filtering. In: Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR’19, p 165–174
- Wang X, Jin H, Zhang A, et al (2020) Disentangled graph collaborative filtering. In: Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval, p 1001–1010
- Wang Y, Stash N, Aroyo L, et al (2008) Recommendations based on semantically enriched museum collections. *Web Semant* 6(4):283–290
- Wang Y, Aroyo LM, Stash N, et al (2009) Cultivating personalized museum tours online and on-site. *Interdisciplinary Science Reviews* 34(2):141–156
- Wei TR, Fang Y (2025) Diffusion models in recommendation systems: A survey. [arXiv:2501.10548](https://arxiv.org/abs/2501.10548)
- Wei W, Ren X, Tang J, et al (2024) Llmrec: Large language models with graph augmentation for recommendation. In: Proceedings of the ACM International Conference on Web Search and Data Mining, p 806–815
- Wei Y, Wang X, Nie L, et al (2019) Mmgcn: Multi-modal graph convolution network for personalized recommendation of micro-video. In: Proceedings of the 27th ACM International Conference on Multimedia, p 1437–1445
- Wei Y, Wang X, Nie L, et al (2020) Graph-refined convolutional network for multimedia recommendation with implicit feedback. In: Proceedings of the 28th ACM International Conference on Multimedia, p 3541–3549
- Wei Y, Wang X, He X, et al (2022) Hierarchical user intent graph network for multimedia recommendation. *IEEE Transactions on Multimedia* 24:2701–2712

- Wu J, Wang X, Feng F, et al (2021) Self-supervised graph learning for recommendation. In: Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval, p 726–735
- Wu J, He X, Wang X, et al (2022a) Graph convolution machine for context-aware recommender system. *Front Comput Sci* 16(6)
- Wu L, Zheng Z, Qiu Z, et al (2024) A survey on large language models for recommendation. *World Wide Web* 27(60)
- Wu S, Sun F, Zhang W, et al (2022b) Graph neural networks in recommender systems: A survey. *ACM Comput Surv* 55(5)
- Wu S, Sun F, Zhang W, et al (2023) Graph neural networks in recommender systems: A survey 55(5)
- Xu D, Ruan C, Korpeoglu E, et al (2020) Inductive representation learning on temporal graphs. URL <https://arxiv.org/abs/2002.07962>, arXiv:2002.07962
- Xu J, Chen Z, Wang W, et al (2025a) Cohesion: Composite graph convolutional network with dual-stage fusion for multimodal recommendation. URL <https://arxiv.org/abs/2504.04452>, arXiv:2504.04452
- Xu J, Chen Z, Yang S, et al (2025b) Mentor: multi-level self-supervised learning for multimodal recommendation. In: Proceedings of the AAAI Conference on Artificial Intelligence
- Xu J, Chen Z, Yang S, et al (2025c) A survey on multimodal recommender systems: Recent advances and future directions. arXiv preprint arXiv:250215711
- Xu Q, Shen F, Liu L, et al (2018) Graphcar: Content-aware multimedia recommendation with graph autoencoder. In: The International ACM SIGIR Conference on Research & Development in Information Retrieval, p 981–984
- Yazijian E (2007) From popular performance to literature: The candimangala of kavikankana mukundarama cakravarti. Ph.D. diss., University of Chicago
- Yi J, Chen Z (2022) Multi-Modal Variational Graph Auto-Encoder for Recommendation Systems. *IEEE Transactions on Multimedia* 24:1067–1079
- Yi Z, Wang X, Ounis I, et al (2022) Multi-modal graph contrastive learning for micro-video recommendation. In: Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval, p 1807–1811
- Yilma BA, Leiva LA (2023) The elements of visual art recommendation: Learning latent semantic representations of paintings. In: Proceedings of the CHI Conference on Human Factors in Computing Systems

- Ying R, He R, Chen K, et al (2018) Graph convolutional neural networks for web-scale recommender systems. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, p 974–983
- Yu P, Tan Z, Lu G, et al (2025) Mind individual information! principal graph learning for multimedia recommendation. In: Proceedings of the AAAI conference on artificial intelligence, pp 13096–13105
- Yu W, Qin Z (2020) Graph convolutional network for recommendation with low-pass collaborative filters. In: Proceedings of the International Conference on Machine Learning. JMLR.org
- Zhang J, Zhu Y, Liu Q, et al (2021a) Mining latent structures for multimedia recommendation. In: Proceedings of the 29th ACM International Conference on Multimedia, p 3872–3880
- Zhang J, Zhu Y, Liu Q, et al (2021b) Mining latent structures for multimedia recommendation. In: Proceedings of the 29th ACM International Conference on Multimedia. ACM, MM '21, p 3872–3880, <https://doi.org/10.1145/3474085.3475259>, URL <http://dx.doi.org/10.1145/3474085.3475259>
- Zhang J, Zhu Y, Liu Q, et al (2023) Latent structure mining with contrastive modality fusion for multimedia recommendation. IEEE Transactions on Knowledge and Data Engineering 35(9):9154–9167. <https://doi.org/10.1109/TKDE.2022.3221949>
- Zhang L, Liu Y, Zhou X, et al (2022) Diffusion-based graph contrastive learning for recommendation with implicit feedback. In: International Conference on Database Systems for Advanced Application, p 232–247
- Zhang M, Chen Y (2018) Link prediction based on graph neural networks. In: Proceedings of the International Conference on Neural Information Processing Systems, p 5171–5181
- Zhang P, Zeng G, Wang T, et al (2024) Tinyllama: An open-source small language model. [arXiv:2401.02385](https://arxiv.org/abs/2401.02385)
- Zhang X, Zitnik M (2020) Gnn-guard: Defending graph neural networks against adversarial attacks. In: Advances in Neural Information Processing Systems, vol 33. Curran Associates, Inc., pp 9263–9275
- Zhang Y, Zhang J, Cui Z, et al (2021c) A graph-based relevance matching model for ad-hoc retrieval. In: AAAI Conference on Artificial Intelligence
- Zhao T, Zhang X, Wang S (2021) Graphsmote: Imbalanced node classification on graphs with graph neural networks. In: Proceedings of the ACM International Conference on Web Search and Data Mining, p 833–841

- Zheng Y, Gao C, He X, et al (2023) Incorporating price into recommendation with graph convolutional networks. *IEEE Transactions on Knowledge and Data Engineering* 35(2):1609–1622
- Zhi-Han Y (2022) Training latent variable models with auto-encoding variational bayes: A tutorial
- Zhou H, Zhou X, Zeng Z, et al (2023a) A comprehensive survey on multi-modal recommender systems: Taxonomy, evaluation, and future directions. *ArXiv abs/2302.04473*
- Zhou H, Zhou X, Zhang L, et al (2023b) Enhancing dyadic relations with homogeneous graphs for multimodal recommendation. In: *ECAI*, pp 3123–3130
- Zhou X, Miao C (2024) Disentangled graph variational auto-encoder for multi-modal recommendation with interpretability. URL <https://arxiv.org/abs/2402.16110>, [arXiv:2402.16110](https://arxiv.org/abs/2402.16110)
- Zhou X, Shen Z (2023) A tale of two graphs: Freezing and denoising graph structures for multimodal recommendation. In: *Proceedings of the 31st ACM International Conference on Multimedia*. ACM, MM '23, p 935–943, <https://doi.org/10.1145/3581783.3611943>, URL <http://dx.doi.org/10.1145/3581783.3611943>
- Zhou X, Lin D, Liu Y, et al (2023c) Layer-refined graph convolutional networks for recommendation. In: *Proceedings of the International Conference on Data Engineering, ICDE 2023*, pp 1247–1259
- Zhu Y, Xu Y, Yu F, et al (2021) Graph contrastive learning with adaptive augmentation. In: *Proceedings of the Web Conference*, pp 2069–2080
- Zhuang Z, Du H, Han H, et al (2025) Bridging the gap: Teacher-assisted wasserstein knowledge distillation for efficient multi-modal recommendation. In: *Proceedings of the ACM on Web Conference*, p 2464–2475