

Combining LLM Semantic Reasoning with GNN Structural Modeling for Multi-View Multi-Label Feature Selection

Zhiqi Chen^{1,2}, Yuzhou Liu^{1,2*}, Jiarui Liu^{1,2}, Wanfu Gao^{1,2*}

¹College of Computer Science and Technology, Jilin University, China

²Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, Jilin University, China
chenzq24@mails.jlu.edu.cn, liuyuzhou@jlu.edu.cn, liujr24@mails.jlu.edu.cn, gaowf@jlu.edu.cn

Abstract

Multi-view multi-label feature selection aims to identify informative features from heterogeneous views, where each sample is associated with multiple interdependent labels. This problem is particularly important in machine learning involving high-dimensional, multimodal data such as social media, bioinformatics or recommendation systems. Existing Multi-View Multi-Label Feature Selection (MVMLFS) methods mainly focus on analyzing statistical information of data, but seldom consider semantic information. In this paper, we aim to use these two types of information jointly and propose a method that combines Large Language Models (LLMs) semantic reasoning with Graph Neural Networks (GNNs) structural modeling for MVMLFS. Specifically, the method consists of three main components. (1) LLM is first used as an evaluation agent to assess the latent semantic relevance among feature, view, and label descriptions. (2) A semantic-aware heterogeneous graph with two levels is designed to represent relations among features, views, and labels: one is a semantic graph representing semantic relations, and the other is a statistical graph. (3) A lightweight Graph Attention Network (GAT) is applied to learn node embedding in the heterogeneous graph as feature saliency scores for ranking and selection. Experimental results on multiple benchmark datasets demonstrate the superiority of our method over state-of-the-art baselines, and it is still effective when applied to small-scale datasets, showcasing its robustness, flexibility, and generalization ability.

Introduction

With the rapid development of machine learning, people tend to associate multi-modal data sources (typically such as texts and images) to train more powerful models, so multi-view multi-label learning has attracted increasing attention in various domains such as social media, bioinformatics, or recommendation systems (Bharati, Mondal, and Podder 2024; Hao et al. 2024a; Han et al. 2025a). However, the combination of different data sources not only gives useful information, but also introduces noisy, irrelevant, or redundant features (Lin et al. 2023b; Cohen et al. 2023; Han, Hu, and Gao 2024). Thus, Multi-View Multi-Label Feature Selection (MVMLFS), aiming at identifying informative features from

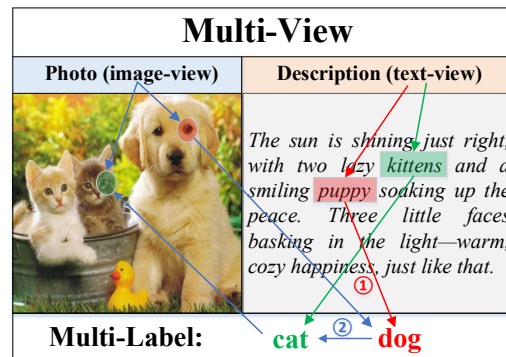


Figure 1: An example of semantic relations in MVMLFS. ① In text view, feature “puppy” has a similar meaning to the label “dog”, which can indicate it is an important feature for “dog”. ② In image view, since labels “dog” and “cat” have high semantic similarity, feature “eye” (which is informative for recognizing “dog”) may also be informative for “cat”.

multi-modal data sources to reduce dimensionality, becomes critical for the learning process to shorten training time, enhance generalization, avoid overfitting, and finally improve performance (Zhang et al. 2020; Hao, Liu, and Gao 2025).

To better complete the MVMLFS task, many methods have been given in recent years, such as MSFS (Zhang et al. 2020), M2LD (Liu et al. 2022), ELSMML (Liu et al. 2023), DHLI (Hao, Liu, and Gao 2024b), and EF2FS (Hao, Gao, and Hu 2025). Although these methods achieve good performance, they usually rely on statistical methods, which means they focus on analyzing statistical information, such as the distributions of data (Lee et al. 2024; Wang, Zhang, and Zhou 2025a). However, semantic information, as a kind of important information in many research fields, has seldom been considered in feature selection (Zhang et al. 2018; Lyu et al. 2024a; Li, Tan, and Liu 2025).

Our key insight is that semantic information can also be useful to construct extra relations for helping feature selection. As illustrated in Figure 1, on one hand, the semantic relations between features and labels can affect the value of features directly, such as: according to the similar meaning of words puppy and dog, the feature “puppy” is important for the label “dog”. On the other hand, the feature impor-

*Corresponding author

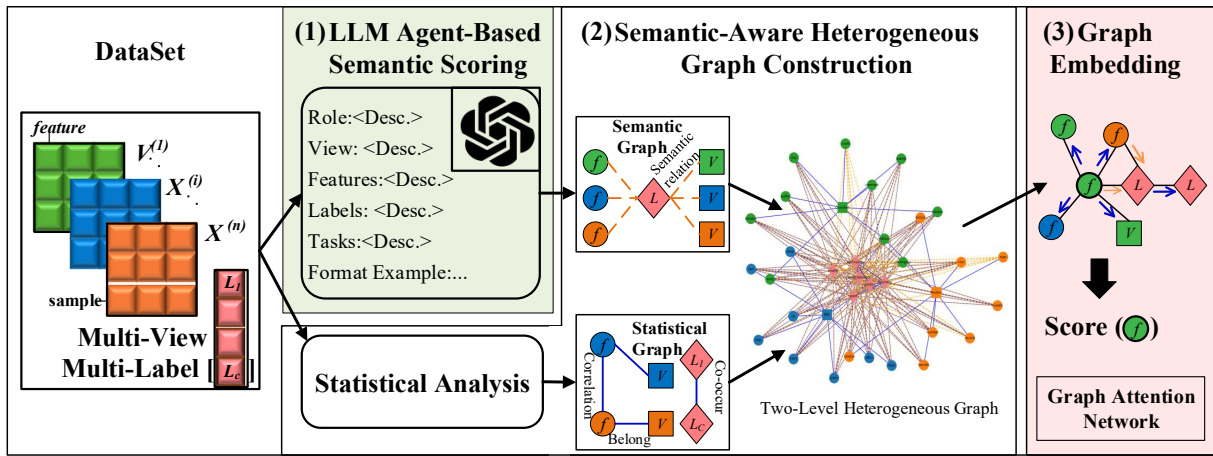


Figure 2: Framework of our method. It consists of three main steps: (1) establishing semantic relations among features, views, and labels by LLM; (2) combining semantic and statistical information, a two-level heterogeneous graph is constructed as the structural modeling of the dataset; (3) employing Graph Attention Network to learn graph embeddings, enabling the estimation of importance scores for feature nodes.

tance can be transmitted based on the semantic relations, for example, semantically similar labels such as “dog” and “cat” often have high semantic relevance and share common informative features, reflecting their underlying semantic correlation. Thus, since the feature “eye” is informative for the label “dog”, it may also be informative for the label “cat”. In this paper, we intend to introduce semantic information to the MVMLFS process. To achieve this goal, two main issues need to be solved.

- How to gain proper semantic information of features, views or labels? In some situations, these objects are described by natural language, and they need enough domain knowledge to understand their meanings. While in other situations, these objects (especially features) may be named by just sequence numbers, which limits the usability of semantic information.
- How to effectively use the semantic information for supporting MVMLFS? On one hand, relying only on semantic information may not be enough, and it needs to be combined with other statistical information about the data. On the other hand, the high-dimensional features mean a high cost of semantic analysis, and it is necessary to balance the payback and cost.

To address these issues, we propose a method that combines Large Language Models (LLMs) semantic reasoning with Graph Neural Networks (GNNs) structural modeling for MVMLFS. Specifically, the method consists of three main components, as shown in Figure 2.

Firstly, semantic scoring evaluation based on LLMs. Benefiting from the great power of LLMs on semantic reasoning, we use them as the agent to assess the latent semantic relevance among feature, view, and label definitions. For the feature with a meaningless name, we use “view+number” to gain its “pseudo-semantic” information by considering that the view often has its clear semantics, and features in one view have the same semantic space.

Secondly, semantic-aware heterogeneous graph construction. GNNs have particular advantages in structural modeling for relation analysis. Thus, we design a two-level heterogeneous graph, in which one is semantic graph representing semantic relations among features, views, and labels; the other is statistical graph consisting of three kind of edges to represent different relations, including belong relations between features and views, correlations among features, views, and labels calculated by mutual information, and co-occurrence among labels (Lin et al. 2023a; Han et al. 2025b). In this way, the heterogeneous graph is constructed to integrate the semantic and statistical information of data.

Finally, embedding-based feature importance estimation. A lightweight Graph Attention Network (GAT) is applied to learn node embedding as feature importance scores. The reasons for using GAT are: it can learn different correlation degrees from different types of edges, and it considers the different importance of various neighbors to the central node. According to the importance scores, it ranks and selects the most informative features across all views.

In summary, our main contributions are given as follows:

- To the best of our knowledge, this is the first attempt to leverage LLMs to introduce extra semantic information for supporting multi-view multi-label feature selection.
- A semantic-aware two-level heterogeneous graph is constructed for supporting the joint analysis of new semantic information with traditional statistical information. Moreover, GAT is used to learn the node embedding in a heterogeneous graph for quantifying the importance of features, so that the selection can better consider the global information.
- The extensive experiments on benchmark multi-view multi-label datasets demonstrate the superiority of our method over state-of-the-art baselines. Moreover, it shows the effectiveness of LLMs and GNNs and can also achieve good performance on the small-scale datasets.

Related Work

Multi-View Multi-Label Feature Selection

Multi-label feature selection has been extensively studied to reduce input dimensionality while preserving label dependency structures (Zhang and Zhou 2014; Lin et al. 2025; Nie et al. 2023). In multi-view scenarios, the challenge intensifies due to feature heterogeneity and view complementarity (Wang, Zhang, and Zhou 2025c; Zeng et al. 2024; Dong et al. 2025). View-independent methods treat each view separately and fuse results post hoc, ignoring cross-view label interactions (Zhang et al. 2019; Huang et al. 2025; Han et al. 2023). View-concatenation methods concatenate all views into a single matrix before selection (Fang et al. 2024; Hao, Liu, and Gao 2024a), which may introduce noise and redundancy. Joint optimization frameworks learn feature relevance jointly across views and labels using sparse regularization (Zhao et al. 2017), matrix factorization (Mei et al. 2025), or tensor-based formulations (Zhang et al. 2021b). However, most of the methods are limited by linear assumptions and lack semantic alignment across views and labels.

Graph-Based Feature Selection

Graphs have emerged powerfully for modeling relational structures among features, samples, and labels. Early works constructed feature similarity graphs (Nie, Zhu, and Li 2016), while recent advances explore label-aware or semantic graphs (Jin et al. 2020; Hao et al. 2024b). Graph convolutional networks and graph attention networks further enhance this line by enabling task-driven embedding learning (Veličković et al. 2018; Huang et al. 2025; Wang, Zhang, and Zhou 2025b). Some studies leverage label co-occurrence graphs (Zhang et al. 2023) or sample-feature graphs (Qin et al. 2024) to enhance selection, yet they often overlook cross-modal or cross-view semantics in feature-label relations. Our work extends this by introducing multi-type heterogeneous graphs with semantic edges and LLM guidance.

Large Language Models for Semantic Scoring

With the rapid development of natural language processing, LLMs (e.g., GPT, LLaMA) have shown strong potential in aligning text representations and enabling semantic inference (OpenAI, Achiam, and et al. 2024; Qin et al. 2025; Zhang et al. 2025). Recent research has utilized LLMs for prompt-driven data labeling (Le Mens and Gallego 2025), knowledge graph construction (Zhu et al. 2024), feature recommendation (Lyu et al. 2024b), and few-shot tabular learning (Han et al. 2024). However, LLMs are seldom used for fine-grained semantic edge construction in feature-label graphs, especially in multi-view settings. Unlike prior work that treats LLMs as end-to-end black-box classifiers, our method leverages LLMs as semantic scoring agents, refining features, labels, and views connections via natural language prompts and integrating them into downstream graph learning.

While existing works contribute to either multi-view feature selection, graph modeling, or LLM-guided tasks, the proposed method uniquely combines all three. A semantic-aware heterogeneous graph is constructed with typed nodes

and edges, integrating LLM-derived semantic priors with GAT-based structural learning. This novel integration enables a principled yet practical method for multi-view multi-label feature selection in semantically complex domains.

Methodology

Problem Formulation

Since our method introduces extra semantic information into MVMLFS, we reformulate the problem by incorporating natural language descriptions of features, views, and labels.

Given a multi-view multi-label dataset $D = \{F^{(i)} \in \mathbb{R}^{n \times d(i)}\}_{i=1}^V$, with V different views, and a binary label matrix is denoted by $L \in \{0, 1\}^{n \times d(L)}$, where n is the number of instances, $d(L)$ is the number of labels, and $F^{(i)}$ is the i -th view with $d(i)$ -dimensional features. To simplify the description, we denote the set of features in the i -th view as $F^{(i)}$, then $f_m^{(i)} \in F^{(i)}$ denotes the m -th feature in the i -th view, and all the features in the dataset is denoted as $F = \bigcup_{i=1}^V F^{(i)}$. Similarly, $l_j \in L$ denotes the j -th label.

Moreover, to enhance semantic understanding, we associate each feature, view, and label with a corresponding textual description in natural language, we have three kinds of text sets: (1) $T^F = \{T_i^f\}_{i=1}^V$, where $T_i^f = \{t_1, t_2, \dots, t_{d(i)}\}$ is a set of texts describing the $d(i)$ features in the i -th view. (2) $T^V = \{t_1, t_2, \dots, t_V\}$ gives text descriptions of the V views. (3) $T^L = \{t_1, t_2, \dots, t_{d(L)}\}$ gives texts describing the meaning of $d(L)$ labels.

Based on the above definitions, the objective is to select a subset of discriminative features from each view and finally aggregate them into a global set of features.

Semantic Score Computation Based on LLM

To incorporate semantic understanding into the feature selection process, we leverage an LLM to quantify semantic scores that capture the relevance among features, views, and labels. This semantic scoring mechanism complements statistical structures by introducing language-level conceptual associations, enabling more informed selection decisions across heterogeneous views.

To achieve the goal, an LLM-based agent is designed to perform semantic analysis on the description texts of features, views, and labels. In this process, we give structured prompts tailored for semantic scoring. Although the detailed contents are adjusted for different analyzing objects (such as feature-label, view-label, or label-label), the prompt mainly consists of five key components, that is $P = \{Role + (Views, Features, Labels) + Task\}$:

- *Role*, defines the thinking perspective of LLM, “*You are a data scientist working on multi-view multi-label feature selection. Your goal is dedicated to exploring the impact of semantic relationships among features, labels, and views on feature selection results.*”
- *Views*, the names and the description texts of the views.
- *Features*, the list of their names and descriptions. Note that if the feature name is without explicit semantics, the pseudo-semantic “view+number” is adopted instead.

- *Labels*, the list of label names representing semantic categories.
- *Tasks*, it needs to identify semantically relevant or redundant features and assign semantic scores in the range [0, 1] between the given objects.

According to the prompt, when given any two objects (feature, view, or label), the LLM-agent returns a semantic score between them, denoted by $LLMScore$. For example, given a feature $f_m^{(i)}$ and a label l_j , semantic score is given as:

$$LLMScore(f_m^{(i)}, l_j) =$$

$$P(\text{Role} + (t_i \in T^v, t_m \in T_i^f, t_j \in T^L) + \text{Task})$$

In this way, we can gain semantic scores to establish extra relations among features, views, and labels, providing additional guidance for the subsequent feature selection process.

Semantic-Aware Heterogeneous Graph Construction

To comprehensively capture both data-driven and knowledge-driven relationships, we construct a unified heterogeneous graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ by integrating statistical dependencies and semantic correlations. This graph comprises three types of nodes *views*, *features*, and *labels* as well as multiple types of edges, each reflecting a distinct relational perspective.

Semantic Graph Construction Beyond capturing statistical dependencies, we construct a semantic graph to model implicit semantic associations among features, labels, and views. This graph complements the statistical structure by introducing edges derived from a language model-based semantic scoring mechanism. Three types of semantic relations are considered: *feature-label*, *view-label*, and *label-label* semantic edges. Edges are established between nodes with scores exceeding a predefined threshold δ , and their weights are set proportionally to their semantic relevance. For example, the semantic feature-label edge is defined as:

$$E_{fls} = \{(f_m^{(i)}, l_j) \mid LLMScore_{mj}^{(i)} > \delta\}$$

Similarly, semantic edges between views and labels, and between labels themselves, are generated in the same manner. These semantic relations are fused into the heterogeneous graph structure. And unified semantic graph captures multi-level dependencies among features, views, and labels, providing enriched structural signals for guiding downstream feature selection.

Statistical Graph Construction To model the statistical relationships among views, features, and labels, we construct a graph \mathcal{G}_s , where edges encode different types of statistical dependencies. Specifically, we define the following edge types and their associated weights:

(1) Feature-View Belonging Edges. Each feature f_i is inherently associated with a specific view v . We connect them with an unweighted edge:

$$E_{fv} = \{(f_i, v) \mid f_i \in F_v\}$$

(2) Feature-Label Correlation Edges. We compute the mutual information (MI) between each feature $f_i \in F_v$ and label $l_j \in L$. An edge is added between f_i and l_j if $MI(f_i, l_j) > \tau_1$. The edge weight is defined as the normalized MI value:

$$E_{fl} = \{(f_i, l_j) \mid MI(f_i, l_j) > \tau_1\}, w_{fl} = \frac{MI(f_i, l_j)}{\max_{i,j} MI(f_i, l_j)}$$

(3) Feature-Feature Correlation Edges. To capture relevance among features, we calculate MI for each feature pair (f_i, f_j) . An edge is added if $MI(f_i, f_j) > \tau_2$. The edge weight reflects the normalized correlation strength:

$$E_{ff} = \{(f_i, f_j) \mid MI(f_i, f_j) > \tau_2\}, w_{ff} = \frac{MI(f_i, f_j)}{\max_{i,j} MI(f_i, f_j)}$$

(4) Label-Label Co-occurrence Edges. We compute co-occurrence statistics across the label matrix L . For labels l_i and l_j , let $c_{ij} = \sum_{n=1}^N \mathbb{I}[L_{ni} = 1 \wedge L_{nj} = 1]$. An edge is added if $c_{ij} > 0$, with weight normalized across all label pairs:

$$E_{ll} = \{(l_i, l_j) \mid c_{ij} > 0\}, w_{ll} = \frac{c_{ij}}{\max_{i,j} c_{ij}}$$

(5) View-Label Correlation Edges. To capture the overall relationship between views and labels, we compute the average mutual information between all features in a view v and label l_j :

$$\bar{M}_{vj} = \frac{1}{|F_v|} \sum_{f_i \in F_v} MI(f_i, l_j)$$

If $\bar{M}_{vj} > \tau_3$, we add an edge, with the weight defined as:

$$E_{vl} = \{(v, l_j) \mid \bar{M}_{vj} > \tau_3\}, w_{vl} = \frac{\bar{M}_{vj}}{\max_{v,j} \bar{M}_{vj}}$$

Through this construction, a weighted statistical graph is obtained that captures rich multi-type dependencies across views, features, and labels. These weighted edges serve as informative priors for subsequent feature selection.

Type-Aware Heterogeneous Graph Attention Network

To effectively model the intricate semantic interactions and deliver messages among views, features, and labels, a type-aware Heterogeneous Graph Attention Network is adopted that extends the standard GAT to heterogeneous graphs. This is achieved via a relation-specific convolution framework that captures type-aware message passing across distinct node and edge types.

Specifically, for each edge relation $r = (\tau_s, \tau_r, \tau_t)$, which represents a connection from a source node of type τ_s to a target node of type τ_t through relation type τ_r , a dedicated GAT layer is instantiated. The embedding of a target node v of type τ_t is computed as:

$$H_v^{(\tau_t)} = \sum_{(\tau_s, \tau_r, \tau_t)} GAT^{(\tau_r)}(h_u^{(\tau_s)}, h_v^{(\tau_t)})$$

Where $h_u^{(\tau_s)}$ and $h_v^{(\tau_t)}$ denote the input features of neighboring source and target nodes, and $\text{GAT}^{(\tau_r)}$ is a relation-specific attention mechanism. For computational efficiency, we adopt single-head attention and omit self-loops.

The outputs from different relation types targeting the same node type are aggregated by summation:

$$H^{(\tau)} = \sum_{r \in \mathcal{R}_{\rightarrow \tau}} H_r^{(\tau)}$$

Where $\mathcal{R}_{\rightarrow \tau}$ is the set of all relations pointing to node type τ .

Once node representations are computed, we derive feature importance scores from the learned embeddings of feature nodes. For each feature node f , its relevance score is defined as:

$$\text{score}(f) = h_f$$

Where $h_f \in \mathbb{R}$ is the scalar output embedding of node f .

We rank all feature nodes based on their scores and select the top- k features to form the final feature subset:

$$S_k = \text{TopK}(\{\text{score}(f) \mid f \in \mathcal{V}_{\text{feature}}\}, k)$$

Experiments

Experimental Setup

Datasets We evaluate the proposed method on six publicly available multi-view multi-label datasets: SCENE (Chua et al. 2009), VOC07 (Everingham et al. 2010), MIRFlickr (Huiskes and Lew 2008), OBJECT (Chua et al. 2009), 3Sources (Greene and Cunningham 2009), and Yeast (Elisseeff and Weston 2001), which cover various domains including vision, news, and genes with different available semantic information. The key characteristics of these datasets are summarized in Table 1. Each view is identified by its corresponding individual name, with the dimensionality shown in braces. Missing views are denoted by ‘-’. The number of samples, features, and labels is denoted by n , d , and c , respectively.

Comparing Methods We compare our method with three representative multi-label feature selection methods based on information theory, including ENM (Gonzalez-Lopez, Ventura, and Cano 2020), MLSMFS (Zhang et al. 2021a), and STFS (Gao et al. 2023). In addition, we evaluate three multi-view multi-label feature selection methods, including MSFS (Zhang et al. 2020), DHLI (Hao, Liu, and Gao 2024b), and EF2FS (Hao, Gao, and Hu 2025). Notably, the adopted learning framework also enables the extraction of a feature importance matrix, which provides insights into the relative contribution of each feature to the final prediction.

Evaluation Metrics For all datasets, MLKNN with $k = 10$ is employed as the base classifier. Feature selection is performed with selection ratios ranging from 2% to 20% of the total number of features, increasing in 2% increments. In each experiment, 70% of the samples are randomly selected for training and the remaining 30% for testing. To ensure robustness, each experiment is repeated 10 times,

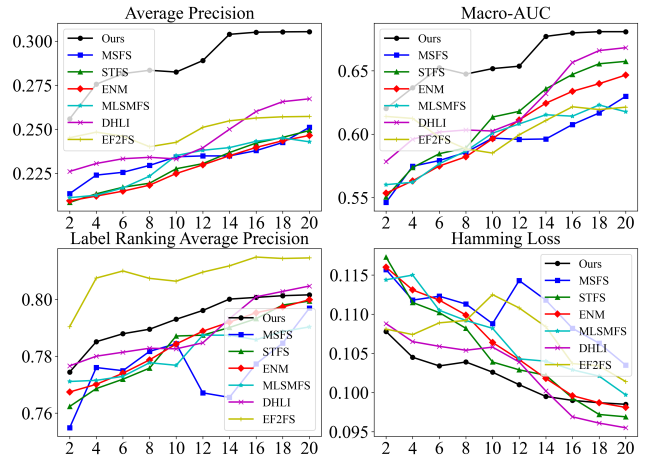


Figure 3: Seven methods on SCENE in terms of Average Precision, Macro-AUC, Label Ranking Average Precision, and Hamming loss.

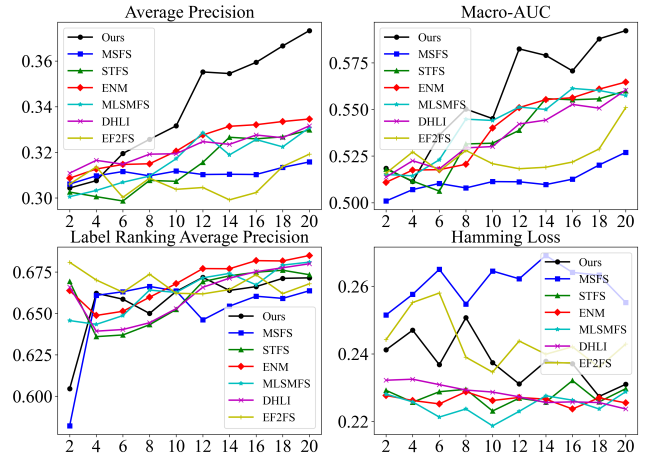


Figure 4: Seven methods on Yeast in terms of Average Precision, Macro-AUC, Label Ranking Average Precision, and Hamming loss.

and the mean performance along with the standard deviation is reported. To evaluate the effectiveness of feature selection, four widely used multi-label classification metrics are adopted (Zhang and Zhou 2014; Gibaja and Ventura 2015): Average Precision (AP), Macro-Area Under the Curve (AUC), Label Ranking Average Precision (LRAP), and Hamming Loss (HL). For AP, AUC, and LRAP, higher values indicate better performance, whereas lower HL values reflect better outcomes. These metrics collectively assess how well the selected features preserve the predictive power with respect to the label space.

Results and Discussion

Table 2 and Table 3 present the performance of our method implemented with GPT4omini, compared with other representative methods. As can be seen, our method achieves

Feature views	SCENE	VOC07	MIRFlickr	OBJECT	3Sources	Yeast
View1(d_1)	CH(64)	DH(100)	DH(100)	CH(64)	BBC(1000)	GE(79)
View2(d_2)	CM(225)	GIST(512)	GIST(512)	CM(225)	Reuters(1000)	PP(24)
View3(d_3)	CORR(144)	HH(100)	HH(100)	CORR(144)	Guardian(1000)	–
View4(d_4)	EDH(73)	–	–	EDH(73)	–	–
View5(d_5)	WT(128)	–	–	WT(128)	–	–
Samples(n)	4400	3817	4053	6047	169	2417
Features(d)	634	712	712	634	3000	103
Labels(c)	33	20	38	31	6	14

Table 1: Characteristics of the datasets in our experiments.

Datasets	Ours	ENM	MLSMFS	STFS	MSFS	DHLI	EF2FS
$AP \uparrow$							
SCENE	0.2888±0.016	0.2274±0.014	0.2308±0.013	0.2290±0.014	0.2328±0.010	0.2440±0.016	0.2499±0.006
VOC07	0.1801±0.021	0.1314±0.005	0.1155±0.002	0.1358±0.007	0.1114±0.001	0.1280±0.002	0.1135±0.003
MIRFlickr	0.3131±0.027	<u>0.2931±0.011</u>	0.2599±0.001	0.2879±0.012	0.2434±0.001	0.2545±0.003	0.2692±0.004
OBJECT	0.1746±0.013	0.1022±0.013	0.0981±0.012	0.1040±0.011	0.1067±0.010	0.1174±0.019	<u>0.1252±0.010</u>
3Sources	0.2086±0.021	<u>0.2044±0.003</u>	0.1880±0.009	0.1946±0.014	0.1925±0.001	0.2033±0.004	0.1928±0.002
Yeast	0.3398±0.025	<u>0.3231±0.010</u>	0.3164±0.011	0.3141±0.012	0.3109±0.003	0.3215±0.006	0.3072±0.007
$AUC \uparrow$							
SCENE	0.6581±0.021	0.6026±0.033	0.5965±0.024	0.6125±0.037	0.5929±0.023	<u>0.6215±0.032</u>	0.6070±0.014
VOC07	0.6410±0.033	0.5947±0.020	0.5363±0.006	<u>0.6087±0.021</u>	0.5001±0.001	0.5754±0.008	0.5066±0.008
MIRFlickr	0.6151±0.034	0.6132±0.020	0.5501±0.001	0.6029±0.015	0.5091±0.001	0.5457±0.005	0.5642±0.009
OBJECT	0.6557±0.016	0.6193±0.030	0.6108±0.023	0.6188±0.028	0.5947±0.019	<u>0.6339±0.033</u>	0.6279±0.011
3Sources	0.5116±0.041	<u>0.5036±0.011</u>	0.4828±0.026	0.5028±0.011	0.5004±0.002	0.4963±0.011	0.5004±0.003
Yeast	0.5574±0.029	0.5395±0.021	<u>0.5422±0.018</u>	0.5365±0.020	0.5118±0.007	0.5364±0.016	0.5248±0.010

Table 2: Experimental results (mean ± std) for AP and AUC, where the 1st/2nd best results are shown in boldface/underline.

most of the superior performance compared to other baseline methods on six datasets under the evaluation metrics AP and AUC. Figure 3 and Figure 4 clearly illustrate our performance metrics on SCENE and Yeast datasets, respectively.

Our method consistently achieves the best performance in terms of AP and AUC across all datasets. On SCENE, VOC07, MIRFlickr, OBJECT, 3Sources, and Yeast, it outperforms the second-best method in AP by margins of +3.9%, +4.4%, +2.0%, +4.9%, +0.4%, and +1.7%, respectively, highlighting its ability to capture discriminative features in diverse, high-dimensional multi-label data. Regarding AUC, our method achieves 0.6410 on VOC07 and 0.6151 on MIRFlickr, demonstrating strong robustness in ranking relevant labels. On the OBJECT dataset, it obtains an AUC of 0.6557, surpassing the closest competitor by more than 2%. In terms of LRAP, our method consistently ranks among the top performers on VOC07, OBJECT, and 3Sources, and second on SCENE and MIRFlickr, indicating enhanced feature selection in complex multi-view settings. For HL, our method attains the lowest values on SCENE, VOC07, and 3Sources, and ranks second on OBJECT.

On Yeast dataset, which involves more complex label semantics, our method does not achieve the best performance in LRAP and HL, with differences of 1.11% and 0.0131

from the top results, respectively. This can be attributed to the challenges in semantic perception posed by intricate label semantics, which limit further gains from the semantic-guided feature selection. As shown in Table 4, the comparable performance of the statistics-based variant without semantic enhancement further supports this observation. Overall, the results demonstrate the consistent superiority of our method over existing baselines.

Ablation Study

To investigate the effectiveness of different components in our method and the impact of different abilities of LLMs, we conduct an ablation study across six datasets, the experiments are shown in Table 4 and Figure 5.

Firstly, the impact of generating semantic scores of different LLMs is evaluated. Two popular models of different scales, GPT4omini and Deepseek R1-14B, are selected for our experiments. Among the variants, the model equipped with GPT4omini achieves the best overall performance, demonstrating that lightweight models can still provide reliable semantic understanding for feature selection. Secondly, we assess the contribution of statistical structure by removing the semantic component. This configuration consistently underperforms compared to the full model, especially on the

Datasets	Ours	ENM	MLSMFS	STFS	MSFS	DHLI	EF2FS
<i>LRAP</i> ↑							
SCENE	<u>0.7930±0.009</u>	0.7848±0.012	0.7810±0.008	0.7834±0.013	0.7763±0.012	0.7890±0.011	0.8087±0.007
VOC07	0.5944±0.017	0.5883±0.008	0.5584±0.002	<u>0.5929±0.011</u>	0.5538±0.014	0.5834±0.003	0.5537±0.019
MIRFlickr	0.6795±0.024	0.6877±0.010	0.6514±0.001	0.6760±0.012	0.5704±0.053	0.6556±0.005	0.6513±0.013
OBJECT	0.4910±0.019	0.4287±0.026	0.4088±0.011	0.4288±0.022	0.4098±0.018	0.4246±0.031	<u>0.4395±0.012</u>
3Sources	0.4636±0.031	<u>0.4557±0.010</u>	0.4275±0.050	0.4528±0.037	0.4041±0.023	0.4524±0.011	<u>0.4442±0.017</u>
Yeast	0.6583±0.020	0.6694±0.013	0.6638±0.014	0.6604±0.016	0.6520±0.025	0.6613±0.016	<u>0.6679±0.007</u>
<i>HL</i> ↓							
SCENE	0.1019±0.003	0.1060±0.007	0.1070±0.005	0.1050±0.007	0.1104±0.004	<u>0.1025±0.005</u>	0.1074±0.004
VOC07	0.0831±0.002	0.0850±0.001	0.0862±0.001	0.0859±0.002	<u>0.0847±0.002</u>	0.0852±0.000	0.0963±0.014
MIRFlickr	0.1799±0.009	0.1710±0.005	0.1919±0.001	<u>0.1756±0.006</u>	0.2288±0.022	0.1896±0.003	0.1934±0.004
OBJECT	<u>0.0587±0.001</u>	0.0611±0.003	0.0590±0.002	0.0616±0.003	0.0654±0.003	0.0575±0.003	0.0594±0.002
3Sources	0.1850±0.009	0.2119±0.007	0.2647±0.043	0.2033±0.010	0.2112±0.017	0.2199±0.007	<u>0.1954±0.006</u>
Yeast	0.2378±0.007	<u>0.2264±0.001</u>	0.2247±0.003	0.2277±0.003	0.2607±0.006	0.2282±0.003	0.2436±0.008

Table 3: Experimental results (mean ± std) for LRAP and HL, where the 1st/2nd best results are shown in boldface/underline.

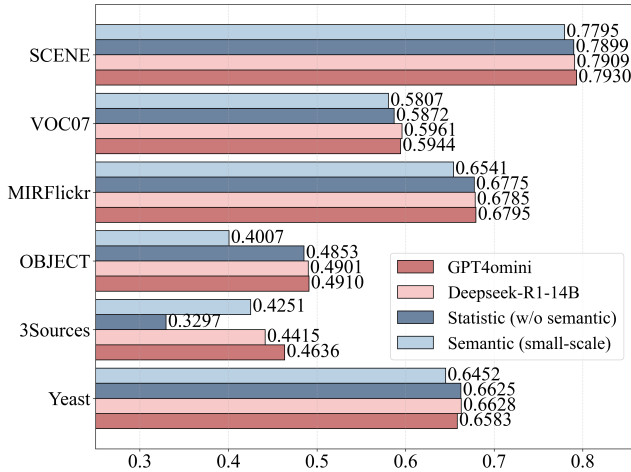


Figure 5: Ablation experimental results in terms of LRAP on six datasets with different settings of the method.

3Sources, which is a text-rich dataset where semantics play a crucial role. This highlights the importance of integrating semantic information, particularly in language-centric tasks.

Finally, to evaluate the performance of the proposed method on small-scale datasets, we investigate whether semantic information is sufficient for effective feature selection. However, relying solely on semantics without incorporating statistical structure results in failed selection due to limited data support. To address this, we retain the semantic component with GPT4omini while introducing partial statistical information. Specifically, when constructing the statistical graph, only 50% of the available data is used, simulating a small-scale dataset scenario. Interestingly, this variant achieves competitive results on small-scale dataset scenarios SCENE and Yeast, suggesting that semantic guidance is effective in small-scale dataset scenarios where sta-

Datasets	Metric	GPT 4omini	Deepseek R1-14B	Statistical (w/o semantic)	Semantic (small-scale)
SCENE	AP	0.2888	<u>0.2875</u>	0.2842	0.2310
	AUC	0.6581	0.6512	<u>0.6523</u>	0.6048
	LRAP	0.7930	<u>0.7909</u>	0.7899	0.7795
	HL	0.1019	<u>0.1030</u>	0.1033	0.1059
Yeast	AP	<u>0.3398</u>	0.3425	0.3025	0.3001
	AUC	0.5574	<u>0.5598</u>	0.5628	0.5042
	LRAP	0.6583	0.6628	<u>0.6625</u>	0.6452
	HL	0.2378	0.2353	<u>0.2335</u>	0.2334

Table 4: Ablation experimental results on SCENE and Yeast.

tistical estimation is less reliable. These findings confirm the complementary benefits of combining statistical and semantic information in our feature selection method.

Conclusion

We propose a novel method for multi-view multi-label feature selection that integrates large language model priors with graph attention mechanisms. By constructing a heterogeneous graph that encodes both LLM-derived semantic information and statistical information based on mutual information as well as label co-occurrence, and applying a lightweight Graph Attention Network, the method produces informative feature embeddings for effective ranking and selection. Extensive experiments on benchmark datasets show consistent improvements over state-of-the-art methods, demonstrating superior performance, robustness, and generalization even on small-scale datasets. In future work, we will explore end-to-end optimization with LLM feedback for closed-loop refinement, integration of learned embeddings into downstream tasks, and self-supervised pretraining on heterogeneous graphs to enhance representation quality.

Acknowledgments

The work is funded by Jilin Provincial Science and Technology Development Plan Project No.20240302084GX, and Changchun Science and Technology Bureau Project 23YQ05.

References

- Bharati, S.; Mondal, M. R. H.; and Podder, P. 2024. A Review on Explainable Artificial Intelligence for Healthcare: Why, How, and When? *IEEE Transactions on Artificial Intelligence*, 5(4): 1429–1442.
- Chua, T.-S.; Tang, J.; Hong, R.; Li, H.; Luo, Z.; and Zheng, Y. 2009. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*, 1–9.
- Cohen, D.; Shnitzer, T.; Kluger, Y.; and Talmon, R. 2023. Few-sample feature selection via feature manifold learning. In *International Conference on Machine Learning*, 6296–6319.
- Dong, Z.; Liu, M.; Wang, S.; Liang, K.; Zhang, Y.; Liu, S.; Jin, J.; Liu, X.; and Zhu, E. 2025. Enhanced then Progressive Fusion with View Graph for Multi-View Clustering. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15518–15527.
- Elisseeff, A.; and Weston, J. 2001. A kernel method for multi-labelled classification. *Advances in neural information processing systems*, 14.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2): 303–338.
- Fang, S.-G.; Huang, D.; Wang, C.-D.; and Tang, Y. 2024. Joint Multi-View Unsupervised Feature Selection and Graph Learning. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 8(1): 16–31.
- Gao, W.; Hao, P.; Wu, Y.; and Zhang, P. 2023. A unified low-order information-theoretic feature selection framework for multi-label learning. *Pattern Recognition*, 134: 109111.
- Gibaja, E.; and Ventura, S. 2015. A Tutorial on Multilabel Learning. *ACM Comput. Surv.*, 47(3).
- Gonzalez-Lopez, J.; Ventura, S.; and Cano, A. 2020. Distributed multi-label feature selection using individual mutual information measures. *Knowledge-Based Systems*, 188: 105052.
- Greene, D.; and Cunningham, P. 2009. A matrix factorization approach for integrating multiple data views. In *Joint European conference on machine learning and knowledge discovery in databases*, 423–438. Springer.
- Han, Q.; Hu, L.; and Gao, W. 2024. Feature relevance and redundancy coefficients for multi-view multi-label feature selection. *Information Sciences*, 652: 119747.
- Han, Q.; Shi, R.; Hu, L.; and Gao, W. 2025a. Multi-label feature selection based on positive sample information weighting. *Knowledge-Based Systems*, 329: 114373.
- Han, Q.; Zhao, Z.; Hu, L.; and Gao, W. 2025b. Enhanced multi-label feature selection considering label-specific relevant information. *Expert Systems with Applications*, 264: 125819.
- Han, S.; Yoon, J.; Arik, S. O.; and Pfister, T. 2024. Large Language Models Can Automatically Engineer Features for Few-Shot Tabular Learning. In *Forty-first International Conference on Machine Learning*.
- Han, Z.; Zhang, C.; Fu, H.; and Zhou, J. T. 2023. Trusted Multi-View Classification With Dynamic Evidential Fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2551–2566.
- Hao, P.; Ding, W.; Gao, W.; and He, J. 2024a. Exploring view-specific label relationships for multi-view multi-label feature selection. *Information Sciences*, 681: 121215.
- Hao, P.; Gao, W.; and Hu, L. 2025. Embedded feature fusion for multi-view multi-label feature selection. *Pattern Recognition*, 157: 110888.
- Hao, P.; Liu, K.; and Gao, W. 2024a. Anchor-guided global view reconstruction for multi-view multi-label feature selection. *Information Sciences*, 679: 121124.
- Hao, P.; Liu, K.; and Gao, W. 2024b. Double-Layer Hybrid-Label Identification Feature Selection for Multi-View Multi-Label Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(11): 12295–12303.
- Hao, P.; Liu, K.; and Gao, W. 2025. Uncertainty-Aware Global-View Reconstruction for Multi-View Multi-Label Feature Selection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(16): 17068–17076.
- Hao, P.; Zhang, P.; Feng, Q.; and Gao, W. 2024b. Label generation with consistency on the graph for multi-label feature selection. *Information Sciences*, 677: 120890.
- Huang, Y.; Lu, M.; Huang, W.; Yi, X.; and Li, T. 2025. TIME-FS: Joint Learning of Tensorial Incomplete Multi-View Unsupervised Feature Selection and Missing-View Imputation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(16): 17503–17510.
- Huiskes, M. J.; and Lew, M. S. 2008. The mir flickr retrieval evaluation. In *Proceedings of the 1st ACM international conference on Multimedia information retrieval*, 39–43.
- Jin, W.; Ma, Y.; Liu, X.; Tang, X.; Wang, S.; and Tang, J. 2020. Graph Structure Learning for Robust Graph Neural Networks. arXiv:2005.10203.
- Le Mens, G.; and Gallego, A. 2025. Positioning Political Texts with Large Language Models by Asking and Averaging. *Political Analysis*, 33(3): 274–282.
- Lee, S. Y.; Kim, S.; Bu, F.; Yoo, J.; Tang, J.; and Shin, K. 2024. Feature distribution on graph topology mediates the effect of graph convolution: homophily perspective. In *Proceedings of the 41st International Conference on Machine Learning*.
- Li, D.; Tan, Z.; and Liu, H. 2025. Exploring Large Language Models for Feature Selection: A Data-centric Perspective. *SIGKDD Explor. Newsl.*, 26(2): 44–53.
- Lin, J.; Chen, Q.; Xue, B.; and Zhang, M. 2025. Evolutionary Sequential Transfer Learning for Multi-Objective

- Feature Selection in Classification. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 9(1): 1019–1033.
- Lin, Y.; He, Z.; Guo, L.; and Ding, W. 2023a. Multi-label feature selection via positive or negative correlation. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 8(1): 401–415.
- Lin, Y.; Liu, H.; Zhao, H.; Hu, Q.; Zhu, X.; and Wu, X. 2023b. Hierarchical Feature Selection Based on Label Distribution Learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(6): 5964–5976.
- Liu, B.; Li, W.; Xiao, Y.; Chen, X.; Liu, L.; Liu, C.; Wang, K.; and Sun, P. 2023. Multi-view multi-label learning with high-order label correlation. *Information Sciences*, 624: 165–184.
- Liu, W.; Yuan, J.; Lyu, G.; and Feng, S. 2022. Label driven latent subspace learning for multi-view multi-label classification. *Applied Intelligence*, 53.
- Lyu, G.; Kang, W.; Wang, H.; Li, Z.; Yang, Z.; and Feng, S. 2024a. Common-individual semantic fusion for multi-view multi-label learning. In *International Joint Conference on Artificial Intelligence*, 4715–4723.
- Lyu, H.; Jiang, S.; Zeng, H.; Xia, Y.; Wang, Q.; Zhang, S.; Chen, R.; Leung, C.; Tang, J.; and Luo, J. 2024b. LLM-Rec: Personalized Recommendation via Prompting Large Language Models. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Findings of the Association for Computational Linguistics: NAACL 2024*, 583–612.
- Mei, S.; Wang, Q.; Gao, Q.; and Yang, M. 2025. Multi-view clustering based on feature selection and semi-non-negative anchor graph factorization. *Neural Networks*, 184: 107111.
- Nie, F.; Tian, L.; Wang, R.; and Li, X. 2023. Learning Feature-Sparse Principal Subspace. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4858–4869.
- Nie, F.; Zhu, W.; and Li, X. 2016. Unsupervised feature selection with structured graph optimization. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI 16, 1302–1308.
- OpenAI; Achiam, J.; and et al., S. A. 2024. GPT-4 Technical Report. arXiv:2303.08774.
- Qin, L.; Chen, Q.; Zhou, Y.; Chen, Z.; Li, Y.; Liao, L.; Li, M.; Che, W.; and Yu, P. S. 2025. A survey of multilingual large language models. *Patterns*, 6(1).
- Qin, Z.; Chen, H.; Mi, Y.; Luo, C.; Horng, S.-J.; and Li, T. 2024. Multi-label Feature selection with adaptive graph learning and label information enhancement. *Knowledge-Based Systems*, 285: 111363.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. arXiv:1710.10903.
- Wang, X.; Zhang, Y.; and Zhou, Y. 2025a. Bidirectional Probabilistic Multi-Graph Learning and Decomposition for Multi-View Clustering. *IEEE Transactions on Image Processing*, 34: 3609–3621.
- Wang, X.; Zhang, Y.; and Zhou, Y. 2025b. Multimodal Remote Sensing Image Clustering With Multiscale Spectral-Spatial Anchor Graphs. *IEEE Transactions on Geoscience and Remote Sensing*, 63: 1–12.
- Wang, X.; Zhang, Y.; and Zhou, Y. 2025c. Pseudo-Supervision Affinity Propagation for Efficient and Scalable Multiview Clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 36(8): 15282–15293.
- Zeng, P.; Yang, M.; Lu, Y.; Zhang, C.; Hu, P.; and Peng, X. 2024. Semantic Invariant Multi-View Clustering With Fully Incomplete Information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4): 2139–2150.
- Zhang, C.; Yu, Z.; Hu, Q.; Zhu, P.; Liu, X.; and Wang, X. 2018. Latent semantic aware multi-view multi-label classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Zhang, K.; Liang, W.; Cao, P.; Yang, J.; Li, W.; and Zaiane, O. R. 2023. Label Correlation Guided Feature Selection for Multi-label Learning. In *Advanced Data Mining and Applications*, 387–402.
- Zhang, M.-L.; and Zhou, Z.-H. 2014. A Review on Multi-Label Learning Algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 26(8): 1819–1837.
- Zhang, P.; Liu, G.; Gao, W.; and Song, J. 2021a. Multi-label feature selection considering label supplementation. *Pattern Recognition*, 120: 108137.
- Zhang, R.; Nie, F.; Li, X.; and Wei, X. 2019. Feature selection with multi-view data: A survey. *Information Fusion*, 50: 158–167.
- Zhang, X.; Wang, D.; Dou, L.; Zhu, Q.; and Che, W. 2025. A survey of table reasoning with large language models. *Frontiers of Computer Science*, 19(9): 199348.
- Zhang, Y.; Wang, X.; Cai, Z.; Zhou, Y.; and Yu, P. S. 2021b. Tensor-Based Unsupervised Multi-View Feature Selection for Image Recognition. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6.
- Zhang, Y.; Wu, J.; Cai, Z.; and Yu, P. S. 2020. Multi-View Multi-Label Learning With Sparse Feature Selection for Image Annotation. *IEEE Transactions on Multimedia*, 22(11): 2844–2857.
- Zhao, J.; Xie, X.; Xu, X.; and Sun, S. 2017. Multi-view learning overview: Recent progress and new challenges. *Information Fusion*, 38: 43–54.
- Zhu, Y.; Wang, X.; Chen, J.; Qiao, S.; Ou, Y.; Yao, Y.; Deng, S.; Chen, H.; and Zhang, N. 2024. LLMs for Knowledge Graph Construction and Reasoning: Recent Capabilities and Future Opportunities. arXiv:2305.13168.