

A Survey on Multi-View Knowledge Graph: Generation, Fusion, Applications and Future Directions

Zihan Yang^{1,2}, Xiaohui Tao¹, Taotao Cai^{1*}, Yifu Tang³, Haoran Xie⁴, Lin Li⁵
Jianxin Li⁶ and Qing Li⁷

¹University of Southern Queensland, Australia

²University of Melbourne, Australia

³Swinburne University of Technology, Australia

⁴Lingnan University, Hong Kong

⁵Wuhan University of Technology, China

⁶Edith Cowan University, Australia

⁷The Hong Kong Polytechnic University, Hong Kong

zihany1@student.unimelb.edu.au, {xiaohui.tao, taotao.cai}@unisq.edu.au,

craigtang@swin.edu.au, hrxie@ln.edu.hk, cathylin@whut.edu.cn,

jianxin.li@ecu.edu.au, csqli@comp.polyu.edu.hk

Abstract

Knowledge Graphs (KGs) have revolutionized structured knowledge representation, yet their capacity to model real-world complexity and heterogeneity remains fundamentally constrained. The emerging paradigm of Multi-View Knowledge Graphs (MVKGs) addresses this gap through multi-view learning, but existing research lacks systematic integration. This survey provides the first systematic consolidation of MVKG methodologies, with four pivotal contributions: 1) The first unified taxonomy of view generation paradigms that rigorously categorizes view into four types: *structure*, *semantic*, *representation*, and *knowledge & modality*; 2) A novel methodological typology for view fusion that systematically classifies techniques by fusion targets (feature, decision, and hybrid); 3) Task-centric application mapping that bridges theoretical MVKG constructs to node/link/graph-level downstream tasks; 4) A forward-looking roadmap identifying underexplored challenges. By unifying fragmented methodologies and formalizing MVKG design principles, this survey serves as a roadmap for advancing KG versatility in complex AI-driven scenarios. In doing so, it paves the way for more efficient knowledge integration, enhanced decision-making, and cross-domain learning in real world.

1 Introduction

KGs [Singhal, 2012] are structured representations of knowledge that model entities as nodes and the relationships between them as edges, which are enriched with attributes that

*Corresponding author.

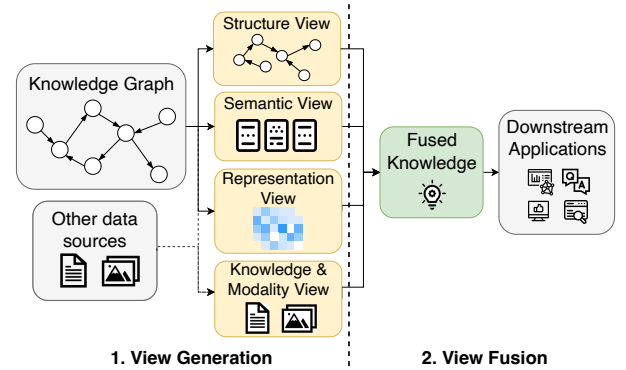


Figure 1: General workflow of multi-view learning in Knowledge Graph (KG). Dashed lines indicate the introduction of external data combined with the original KG to form new views.

capture additional information [Hogan *et al.*, 2021]. This graph-based framework provides a powerful and interpretable way to encode and organize information from various data sources. Over the past decade, KGs have been widely employed in diverse fields, including natural language processing (NLP) [Chen and Luo, 2019; Huang *et al.*, 2019], recommendation system [Guo *et al.*, 2020], search engine [Zou, 2020] and biomedical research [Nicholson and Greene, 2020]. Recent advancements, such as the large language models (LLMs) [Radford, 2018], multi-modal learning [Xu *et al.*, 2023], and graph neural networks (GNNs) [Kipf and Welling, 2016] has opened new avenues for KGs to integrate with cutting-edge technologies.

Motivation of MVKG. As real-world data grows increasingly complex, traditional KGs face significant challenges in modeling and interpreting heterogeneous data effectively. Monolithic frameworks for KG reasoning are often insufficient for capturing the structural and semantic variations in-

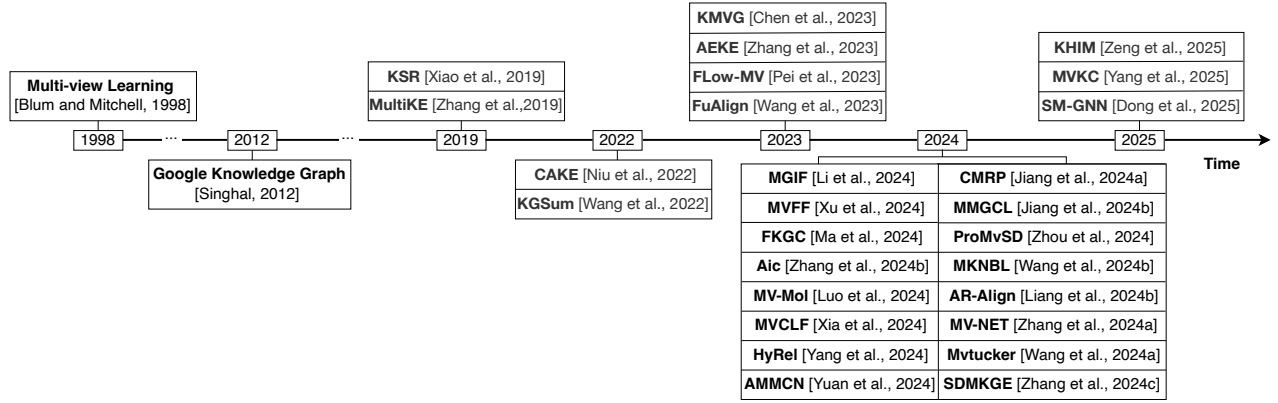


Figure 2: A timeline of key developments in Multi-View Knowledge Graph (MVKG), beginning with foundational work in multi-view learning (1988) and Google Knowledge Graph (2012). The first work on MVKG was proposed in 2019, followed by steady growth, then a significant surge in 2024. Expectations are for continued advancements in the future.

herent in many datasets. These limitations hinder the full utilization of data and constrain generalization performance. To address these challenges, the concept of multi-view learning [Blum and Mitchell, 1998] presents a promising solution. Multi-view learning in KGs typically involves two key steps (illustrated in Figure 1): (i) generating multiple views that capture diverse aspects of the data, and (ii) fusing the view-specific knowledge. This process enables a more nuanced understanding of the heterogeneous datasets and overcomes the limitations of traditional KG reasoning approaches in managing data variations. Combining the principles of multi-view learning and KG, a MVKG extends traditional KGs by structuring data into distinct views, then integrates these views to maximize the utility of their combined information.

Scope of the Survey. Although both MVKGs and multi-modal knowledge graphs (MMKGs) [Liu *et al.*, 2019] such as Richpedia [Wang *et al.*, 2020] and MKGAT [Sun *et al.*, 2020] may involve data from different modalities, they differ significantly in their data sources and objectives. MMKGs primarily focus on integrating multi-modal data, such as text, images, and audio, to construct a unified KG. Such KGs have been covered in previous surveys [Zhu *et al.*, 2022; Liang *et al.*, 2024a]. In contrast to MMKGs, MVKGs leverage existing KGs to achieve a more comprehensive understanding of knowledge. Specifically, MVKGs mainly operate on existing KGs and generate multiple views by interpreting the data from various dimensions, including structure [Chen *et al.*, 2023], semantics [Xiao *et al.*, 2019; Zhang *et al.*, 2019] or representation encoders [Pei *et al.*, 2023]. In some cases, data from other modalities can be combined with the original KG and added as an additional view to provide complementary knowledge [Luo *et al.*, 2024]. Once these views are generated and represented, they are fused to enable a more comprehensive and effective utilization of the information. This facilitates better solutions to diverse downstream tasks, such as entity alignment [Wang *et al.*, 2023], recommendation [Wang *et al.*, 2024c] and KG completion [Niu *et al.*, 2022]. To the best of our knowledge, this is the first survey paper on MVKGs that provides a comprehensive overview of view generation and fusion, applications, and fu-

ture research directions.

Contributions. This work aims to present a comprehensive survey on MVKGs, focusing view generation and fusion techniques, downstream applications, and future directions. The key contributions of this survey are as follows:

- We offer the first systematic and comprehensive review of MVKGs, distinguishing them from traditional KGs and MMKGs by providing a formal definition of key concepts underlying MVKGs in Section 2. A brief timeline of key techniques and reviewed works is given in Figure 2.
- We introduce a unified framework for MVKGs, emphasizing two steps: view generation and view fusion. Section 3 categorizes different view types based on the generation process as displayed in Figure 3, while Section 4 provides an in-depth discussion of representative fusion techniques.
- We examine the diverse downstream applications of MVKGs in Section 5, categorized by granularity into node-level, link-level, and graph-level tasks.
- In Section 6, we discuss the key challenges and outline promising research directions by offering insights to guide future advancements in MVKGs.

2 Preliminary

2.1 Knowledge Graph

Definition. KGs are structured representations of knowledge that capture entities, attributes, and relationships. A KG is denoted as $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$, where \mathcal{E} is entity set, \mathcal{R} is relation set, \mathcal{T} is triple set. Each triple $(h, r, t) \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ represents a relational fact, where $h, t \in \mathcal{E}$ are head and tail entities, and $r \in \mathcal{R}$ is the relation connecting them.

2.2 Multi-View Learning in Knowledge Graph

Multi-view learning originally emerged in machine learning as a paradigm to address the limitations of single-view learning, in particular the inability to comprehensively capture the complexity and diversity of data from a single perspective [Zhao *et al.*, 2017]. In KGs, multi-view learning is especially valuable because real-world knowledge is inherently multi-

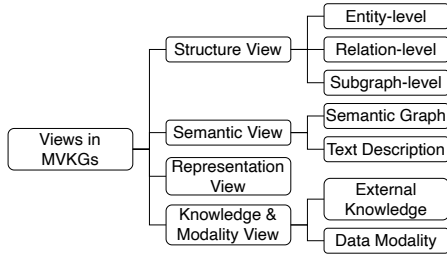


Figure 3: A taxonomy of view types in MVKGs based on view generation process, highlighting *structure*, *semantic*, *representation*, and *knowledge & modality* views. The structure view is further divided into entity-, relation-, and subgraph-level; semantic view can be constructed via semantic graph or text description; knowledge & modality View emphasizes external knowledge and data modality.

faceted, requiring integration from multiple perspectives to fully represent entities, relationships, attributes.

View Definition. A view is a distinct representation or perspective of data that captures specific aspects of its structure, semantics, or context. Within KGs, a view v_i is a subset of the KG \mathcal{G} and optionally an extra data source \mathcal{K} , denoted as $v_i = g_i(\mathcal{G}, \mathcal{K})$, where g_i is a mapping function that extracts or constructs the i -th view.

Multi-View Learning. Multi-view learning is a machine learning paradigm that leverages multiple views of data to improve model performance. In the context of KGs, multi-view learning aims to integrate and reason over diverse and complementary information captured from different views. Formally, given a set of views $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ derived from \mathcal{G} , the goal of multi-view learning is to obtain fused knowledge \mathcal{Z} by integrating the information from all views, such that $\mathcal{Z} = f(\mathcal{V}; \Theta)$ where f is a fusion function and Θ represents the learnable parameters of the model.

As illustrated in Figure 1, the multi-view learning process in KG typically involves two key steps:

1. **View Generation:** Construct a set of views $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$, where each view $v_i = g_i(\mathcal{G}, \mathcal{K})$ is derived from a unique perspective of the KG.
2. **View Fusion:** Combine multiple views to produce fused knowledge \mathcal{Z} using the fusion function f .

Multi-View Knowledge Graph (MVKG) Definition. A MVKG extends a traditional KG by integrating multiple views to capture diverse and complementary aspects of knowledge through the application of multi-view learning. Formally, an MVKG is defined as $(\mathcal{G}, \mathcal{V}, f, \Theta)$.

2.3 Reasoning over MVKG

Building upon the foundation of MVKGs, reasoning over MVKGs plays a crucial role in leveraging fused knowledge from multiple views. As shown in Figure 4, reasoning over MVKGs involves addressing four interconnected challenges: *Construction*, *Integration*, *Alignment*, and *Translation*.

Construction. A critical challenge in reasoning over MVKGs is the construction of multiple views, which are extracted from the underlying KG and may be enriched by external data sources. Each view captures a distinct aspect of the data, such as structural patterns or semantic information.




Figure 4: Relationships among four fundamental challenges in reasoning over MVKG.

This step lays the foundation for the subsequent challenges.

Integration. Building upon the construction of multiple views, another pivotal challenge in MVKG reasoning is integration, which plays a central role in combining the complementary knowledge captured by each view. By fusing diverse views, knowledge integration enables a more comprehensive and robust understanding of data. This process supports subsequent challenges such as alignment and translation.

Alignment. Following integration, a critical challenge is alignment, which aims to align entities or relations that share the same semantics across different views or data sources. Typically, alignment in MVKGs focuses on cross-view alignment, where entities or relations from different views within the same MVKG are matched. When external KGs are incorporated, cross-KG alignment may also be required. This process ensures consistency and coherence across views, further guiding the view construction.

Translation. Alongside alignment, translation addresses the transfer and transformation of knowledge across different views. This process aims to generate or retrieve missing information while maintaining semantic and informational consistency. Supported by knowledge integration, view translation not only enhances the completeness and consistency of MVKGs but also optimizes the view construction process, enabling the creation of more complete and accurate views.

3 View Generation

As a core component of MVKGs, view generation creates diverse views of the given KG, each highlighting distinct aspects of the underlying data. To better understand the unique characteristics captured by views and their applications in downstream tasks, we organize views in existing MVKGs into four main categories based on specific data features to be emphasized: *structure*, *semantic*, *representation*, and *knowledge & modality*, as illustrated in Figure 1.

3.1 Structure View

The structure view focuses on the connectivity and topological relationships between entities and relations in a KG. It captures the relational patterns and structural characteristics of entities within the KG, emphasizing the graph’s inherent connectivity and hierarchy. Based on the level of structural analysis, the structure view can be further categorized into three types: *entity-level*, *relation-level*, and *subgraph-level*.

Entity-level

The entity-level structure view analyzes the topological organization of individual entities in a KG by capturing local




Figure 5: Structure view in MVKGs: illustration at entity-level, relation-level, and subgraph-level.

structural patterns from their immediate relationships. For example, CAKE [Niu *et al.*, 2022] uses a fact view for entity prediction based on first-order relationships, while ProMvSD [Zhou *et al.*, 2024] defines the node view as a unit for evaluating semantic contradictions in relational facts.

Several methods capture entity-level structure in MVKGs. For example, MVCLF [Xia *et al.*, 2024] uses sampling to generate views from first-order connections, while KGSum [Wang *et al.*, 2022] employs Graph Transformer for node representations. MGIF [Li *et al.*, 2024] applies global-aware convolution to process relational and entity features separately, and MultiKE [Zhang *et al.*, 2019] enriches entity embeddings with CNN-based attribute representations.

Entity-level views have been widely applied in KG tasks. CAKE shows effectiveness in entity alignment and KG completion using local embeddings. In recommendation systems, KHIM [Zeng *et al.*, 2025] integrates entity-level knowledge into a hierarchical intent modeling module to improve user and item representations. Despite these benefits, the entity-level view faces challenges such as semantic inconsistencies [Zhou *et al.*, 2024] and computational inefficiencies [Dong *et al.*, 2025] in large-scale KGs.

Relation-level

The relation-level structure view models interactions between entities via relational connections in MVKGs. Instead of entities, it highlights how relations shape knowledge representation and graph structure. For example, TransE in MultiKE treats relations as translation vectors from head to tail entities, providing a basic approach for relational semantics.

Various methods exploit relational structures in MVKGs. For instance, FuAlign [Wang *et al.*, 2023] defines an entity’s neighborhood as its neighbors and their relations by using context embeddings to capture local patterns. Additionally, AEKE [Zhang *et al.*, 2023] builds a relational hypergraph by treating triples as nodes linked through shared entities, while ProMvSD extends this idea with a triple view where triples act as hypernodes.

Dynamic approaches can selectively extract relations. CMRP [Jiang *et al.*, 2024a] employs dynamic edge selec-

tion to focus on relevant relations, and MGIF introduces interaction-aware convolution for encoding relation-specific features. Moreover, SM-GNN [Dong *et al.*, 2025] aggregates relation-aware information using attention mechanisms to differentiate contributions. In heterogeneous graphs, KG-Sum constructs an entity-sentence graph to update sentence representations based on entity interactions. MVKC [Yang *et al.*, 2025] divides user-item-entity graph into a user-item view and an item-entity view.

Subgraph-level

The subgraph-level structure view extracts meaningful subgraphs from MVKGs to model complex relationships or temporal dynamics. Unlike entity or relation level, it captures dependencies across multiple entities and relations, which is essential for temporal reasoning and subgraph-based inference.

Temporal dynamics are modeled explicitly using subgraph-level views. For example, Mvtucker [Wang *et al.*, 2024a] introduces a temporal view to explore interactions between entities, relations and time, while MVFF [Xu *et al.*, 2024] captures evolving patterns via temporal structure view. MV-NET [Zhang *et al.*, 2024a] extends this idea by proposing multiple temporal views at each timestamp. Additionally, MMGCL [Jiang *et al.*, 2024b] employs a dynamic sliding window to extract meta-knowledge from a manageable temporal graph.

Beyond temporal analysis, subgraph-level views also model pathway-based semantics and triple-level interactions. ProMvSD introduces a pathway view to trace information flow along sub-paths and computes semantic gaps, while SM-GNN uses a triple view to aggregate features of relations and tail entities. Moreover, Aic [Zhang *et al.*, 2024b] explores multi-granularity subgraph extraction, highlighting the importance of diverse subgraph-level analysis.

3.2 Semantic View

The semantic view in MVKGs focuses on capturing the semantic information of entities, relations, and their interactions. Unlike structure view that emphasizes graph topology, semantic view usually leverages *semantic graph* or *text description* to enhance knowledge representation and reasoning.

Semantic Graph

A semantic graph is a derived structure that captures implicit semantic information in the original KG by generating new relationships or structures. It is built by extracting, transforming, or generating relational structures from entity attributes, relation semantics, or contextual patterns. An example is HyRel [Yang *et al.*, 2024], which reorganizes relations based on their relative positions and shared entity frequency.

Various methods construct semantic graphs in MVKGs. AEKE uses triple-level hypergraphs to model attribute semantics, while HyRel and MKNBL [Wang *et al.*, 2024b] fuse entity and relation features to produce enriched representations. For dynamic data, KMVG [Chen *et al.*, 2023] builds directed session graphs to capture sequential patterns. Hierarchical methods such as KSR [Xiao *et al.*, 2019] organize knowledge into multi-level views and clusters. Additionally, AMMCN [Yuan *et al.*, 2024] combines user-item graph with KG to construct item-item graph for recommendation tasks. Despite their interpretability and flexibility, semantic graphs

face challenges in computational complexity and scalability [Zhang *et al.*, 2023; Chen *et al.*, 2023], and their performance heavily depends on the quality of a KG [Wang *et al.*, 2024b].

Text Description

The text-based semantic view leverages textual information, such as names or descriptions, to enrich MVKG representations. Using NLP techniques as shown in [Wang *et al.*, 2023], this approach extracts and encodes textual semantics to improve the understanding of entities and their relationships.

A common method is to employ pretrained language models (PLMs) for encoding text. For instance, SDMKGE [Zhang *et al.*, 2024c] uses two Siamese networks based on PLMs to separately encode entity and structure semantics. Similarly, AR-Align [Liang *et al.*, 2024b] initializes entity names and attributes using LaBSE to capture contextual semantics. FKGC [Ma *et al.*, 2024] also utilizes text descriptions to achieve richer semantic representations. Entity names often serve as a universal textual feature for semantic enrichment. For instance, FuAlign emphasizes their importance as a text source, while MultiKE uses literal embeddings to capture name semantics. Another approach focuses on domain-scalable prompts that integrate textual inputs with graph data. CMRP converts text into prompts that preserve semantic and logical content while emphasizing graph information. But while text-based views enhance reasoning, challenges like computational overhead [Jiang *et al.*, 2024a] and ambiguity in descriptions [Wang *et al.*, 2023] must be carefully addressed.

3.3 Representation View

The representation view focuses on generating diverse representations of a MVKG by applying different encoding techniques or perturbations to the same data. This approach uses multiple encoders or augmentation strategies to create distinct views of the KG, enabling richer feature extraction and improved downstream performance. Unlike structure or semantic views that rely on explicit graph topology or semantic information, the representation view derives alternative embeddings by altering feature space properties.

One common strategy is to use multiple encoders to generate different views. FLOW-MV [Pei *et al.*, 2023] employs three encoders: a few-shot learner for relation representation, a relation knowledge distiller for another different view, and a perturbed few-shot learner that creates a new view via perturbation. Similarly, MVFF uses two encoders, one based on tensor decomposition and another based on a relational GNN. Another approach is to directly perturb and augment embeddings to create new views such as AR-Align. MVKC further extends this idea with a stochastic scheme to generate two augmented views for contrastive learning.

Representation views significantly enhance contrastive learning and few-shot generalization, as demonstrated in FLOW-MV. However, challenges such as computational overhead [Xu *et al.*, 2024] and representation inconsistency [Yang *et al.*, 2025] must be addressed for scalable deployment in real-world knowledge graphs.

3.4 Knowledge & Modality View

The knowledge & modality view enhances MVKGs by incorporating *external knowledge* or additional *data modality*.




Figure 6: Feature fusion in MVKGs illustrated with two views.

This approach enriches KG representations by integrating complementary information from diverse sources, enabling more comprehensive and robust reasoning.

External Knowledge

Integrating external knowledge bases enriches MVKGs by adding extra semantic constraints and relational details. For example, FKGC introduces a commonsense view that uses high-quality negative sampling based on complex relations. This improves sample quality and applies semantic constraints to reduce overfitting. Similarly, CAKE leverages commonsense knowledge to enhance the original KG, capturing implicit relationships and constraints. Such external knowledge provides critical contextual information absent from the original KG that leads to a more accurate domain representation.

Data Modality

Incorporating additional data modalities further enriches MVKGs by integrating non-relational data like text, images, or molecular structures. For instance, MV-Mol [Luo *et al.*, 2024] combines chemical structure data with unstructured biomedical texts to create a unified representation that captures diverse aspects of the data. By fusing structured KG knowledge with unstructured modalities, MV-Mol improves the accuracy and interpretability of predictions in biomedical applications. This broadens the scope and adaptability of MVKGs for complex real-world tasks.

4 View Fusion

View fusion integrates information from multiple views in a MVKG to enhance representation and reasoning capabilities. By combining diverse views, fusion methods provide a more comprehensive understanding of underlying data. Depending on the fusion targets, fusion can occur at *feature* level (combining embeddings or raw features) or *decision* level (merging outputs or predictions from different views).

4.1 Feature Fusion

Feature fusion combines embeddings or raw features from different views to create a unified and enriched representation. This occurs before the final representation is generated, enabling models to capture complementary information for more robust and expressive embeddings. Figure 6 shows that embeddings from separate encoders are fused to integrate knowledge across views. Techniques used in feature fusion include direct combination [Xu *et al.*, 2024; Wang *et al.*, 2023], contrastive learning [Jiang *et al.*, 2024b], attention mechanisms [Wang *et al.*, 2022; Li *et al.*, 2024], and vector transformation [Wang *et al.*, 2024a].




Figure 7: Decision fusion in MVKGs illustrated with two views.

Combination-based. A straightforward approach is to directly combine feature vectors using concatenation, element-wise addition, or pooling, often enhanced by weighting mechanisms to emphasize important views. For example, MVFF and FuAlign concatenate features, while MKNBL and KMVG explore pooling and concatenation. SM-GNN selects features from more similar views for fusion based on feature consistency. Weighted methods, such as the weighted view averaging in MultiKE, assign importance scores to different views to improve fusion quality.

Contrastive Learning. Contrastive learning methods align features from different views by maximizing similarity for positive pairs and minimizing for negative pairs [Zeng *et al.*, 2025; Yang *et al.*, 2024]. Some models use dual contrastive learning across multiple views [Yang *et al.*, 2024], while others refine instance-level representations [Zhang *et al.*, 2023]. MVCLF compares features across views, and MV-Mol applies cross-modal contrastive learning for modality alignment. Additionally, MMGCL proposes a meta-knowledge transfer contrastive approach, while FLOW-MV uses contrastive knowledge distillation for few-shot learning.

Attention Mechanism. Attention-based fusion techniques dynamically weight features from different views to capture their importance. KGSum uses cross-attention to combine textual and graph context, while MGIF adopts self-attention to correct multi-perspective feature maps. KMVG employs soft-attention mechanism to fuse item representations from session-view graphs. Reinforcement learning techniques, such as collaborative policy learning scheme in CMRP, can also be seen as a form of attention, where the model learns to optimize edge selection strategies based on mutual rewards.

Vector-based. Vector-based methods transform features into a shared space or decomposes them into unified representations. MultiKE induces an orthogonal mapping matrix to project view-specific embeddings into shared latent space. MVFF utilizes tensor decomposition to represent entities and relations with time features. MvTucker models MVKGs as n th-order binary tensors and uses tensor products to capture interactions between views.

4.2 Decision Fusion

Decision fusion in MVKGs combines outputs from multiple views to make a unified decision. Unlike feature fusion, decision fusion usually merges predictions, scores, or losses at a higher level. As shown in Figure 7, each view processes its own embedding, and their decisions are combined to generate the final fused knowledge. Decision fusion can be done through simple combination or weight-based methods.

Combination-based. Simple combination techniques merge the outputs of different views without explicit weighting. For

instance, MV-Mol combines the losses from two subtasks - KG embedding and KG completion - to incorporate multi-view knowledge. Similarly, SDMKGE sums the entity semantic similarity matrix and the structure semantic similarity matrix to obtain the final semantic similarity matrix. While computationally efficient this method lacks adaptability, as it usually assumes equal contributions from all views.

Weight-based. Weight-based methods assign importance scores to the outputs of different views, either statically or dynamically, to balance their contributions. For example, ProMVSD uses a static trade-off coefficient λ to balance the influence of two hyperviews when assessing the suspicious score of a triple. In contrast, MV-NET designs an adaptive scoring module that customizes scores for different queries, utilizing ConvTransE to calculate the score of a quadruple for each view. This dynamic approach ensures that the fusion process adapts to the specific context of the query.

4.3 Hybrid Fusion

Hybrid fusion in MVKGs refers to the integration of different fusion strategies, with the approach typically combining feature fusion and decision fusion at different stages of the learning process. Unlike purely feature- or decision-based methods, hybrid fusion leverages both early and late fusion techniques.

Early Fusion Methods. Early fusion techniques integrate features from different views at an initial stage, often using contrastive learning or attention mechanisms to align representations. For instance, AR-Align uses contrastive learning to align entity names and attributes early on, while AMMCN aligns embeddings from different views before combining them. MVKC uses contrastive learning for structure views and augmented views, enabling effective information exchange. FKGC introduces adaptive attention to adjust the interaction between the structure and text views, ensuring focus on the most relevant features during early fusion.

Late Fusion Methods. Late fusion techniques combine the outputs of different views at a later stage, often using filtering, summing, or concatenation to refine the final predictions. For instance, AMMCN sums and concatenates embeddings in late fusion. CAKE filters candidate entities with commonsense knowledge before predicting final answer entities. FKGC also uses commonsense knowledge to filter out irrelevant candidates, narrowing the range for link predictions. MVKC sums and concatenates the representations from different views to generate the final output, ensuring that both structure and augmented views contribute to the decision.

5 Downstream Applications

Having explored the generation and fusion of multiple views in MVKGs, we turn our focus to the diverse downstream applications that leverage these enriched knowledge representations. By integrating diverse views, MVKGs provide a structured knowledge representation, making them valuable for various applications at different levels of graph analysis. These applications can be categorized into *node-level*, *link-level*, and *graph-level*. The following section examines the impact of MVKGs across these levels of analysis.

5.1 Node-level

Node-level applications in MVKGs focus on analyzing individual entities by leveraging multi-view representations to enhance entity alignment, retrieval, enrichment, and prediction tasks.

Entity Alignment & Retrieval. Entity alignment aims to match equivalent entities across different KGs, addressing inconsistencies in naming conventions, languages, and structural representations. AR-Align, FuAlign, and SDMKGE highlight how multi-view fusion improves cross-lingual entity alignment by incorporating graph structures and auxiliary semantic information. Additionally, MultiKE emphasizes the importance of treating entity features equally across views to ensure robust alignment. In entity retrieval, multi-view representations refine entity search results by integrating semantic and structural knowledge [Xiao *et al.*, 2019].

Entity Enrichment & Prediction. Entity enrichment integrates additional attributes into entity representations to enhance their utility for downstream tasks. For example, AEKE enhances entity enrichment through a multi-view framework by modeling KG topology and attributes as hypergraphs. MV-Mol demonstrates how enriched representations aid molecular property prediction by refining entity attributes with diverse knowledge sources. Entity prediction involves inferring missing attributes or relationships. MV-NET evaluates multi-view entity prediction models across diverse datasets, showing clear effectiveness in predicting missing entities using structural and semantic information.

5.2 Link-level

Link-level applications in MVKGs focus on tasks related to entity-entity interactions, including link prediction and recommendation systems.

Link Prediction. Link prediction aims to infer missing or potential edges in KG by analyzing multi-view structural relationships. MVFF and HyRel demonstrate how multi-view representations improve prediction by capturing hidden dependencies and semantic similarities between entities.

Recommendation. MVKGs enhance recommendation systems by leveraging semantic information for refined user-item interaction predictions. KMVG integrates session-based and pairwise views via a global item-item graph to improve accuracy. AR-Align and KHIM capture user preferences through multi-view interactions and semantic graphs. Contrastive learning propagates neighboring node information while distinguishing importance levels [Xia *et al.*, 2024]. To tackle data sparsity and cold-start issues, MKNBL employs a multi-channel knowledge-aware network using KGs as side information. MMGCL and MVKC further optimize recommendations via meta-knowledge-enhanced contrastive learning and GNN-based multi-view fusion.

5.3 Graph-level

Graph-level applications in MVKGs extend beyond individual entities and links, focusing on holistic graph-wide tasks such as KG completion [Pei *et al.*, 2023; Niu *et al.*, 2022] or multi-domain applications [Jiang *et al.*, 2024a]. These applications leverage multi-view structures, semantic represen-

tations, and cross-domain integration to improve reasoning, retrieval, and interpretability.

KG Completion. KG completion (KGC) infers missing triples by leveraging multi-view structural and semantic relationships. FLOW-MV addresses long-tail entity issues in low-resource KGC, while CAKE refines predictions with a coarse-to-fine link prediction module using commonsense and fact-based views. MGIF enables cross-view knowledge sharing to resolve multi-domain semantic inconsistencies. FKGC demonstrates that integrating multiple perspectives enhances completion accuracy. In multilingual settings, SM-GNN advocates simplified GNN to handle large-scale graphs with heterogeneous linguistic structures.

Multi-Domain Tasks. MVKGs facilitate knowledge transfer across different domains, allowing for integration between graph learning and NLP tasks. For example, CMRP demonstrates the versatility of MVKGs by evaluating them on thirteen graph learning datasets and ten NLP datasets, showing that graph-based reasoning enhances language-based tasks.

6 Conclusion and Future Directions

In conclusion, this paper provides a comprehensive survey of MVKGs, covering definitions, methodologies, and applications. We now highlight the current research landscape and suggest future directions to address existing challenges and further expand the field.

Benchmarking and Evaluation. Despite numerous MVKG models, standardized evaluation methodologies are lacking. Current benchmarks focus on single-task performance without assessing view quality, informativeness, or fusion generalization. Future research should establish comprehensive frameworks that systematically compare MVKG models across view generation and fusion techniques.

Advanced Representation Learning. Enhancing representation learning techniques is another key direction. This includes exploring multi-modal and cross-lingual embeddings to capture richer semantic information across diverse data sources and languages. Additionally, development of embedding models that leverage higher-order neighborhood information, logic paths, and global KG structures can provide more comprehensive representations.

Few-shot Learning and Noise Reduction. Few-shot learning and noise reduction are crucial for low-resource and noisy environments. Soft sampling for k-shot learning enhances performance with limited labeled data, while knowledge distillation reduces noise and improves robustness, especially in recommendation systems with incomplete or noisy data.

Integration with LLMs. Integrating MVKGs with LLMs enables better reasoning and inference by aligning textual knowledge with structured KG data. Developing prompt templates can enhance this synergy. Future research may explore using LLMs to generate or refine views, bridging unstructured text and structured knowledge.

References

- [Blum and Mitchell, 1998] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100, 1998.
- [Chen and Luo, 2019] Hainan Chen and Xiaowei Luo. An automatic literature knowledge graph and reasoning network modeling framework based on ontology and natural language processing. *Advanced Engineering Informatics*, 42:100959, 2019.
- [Chen et al., 2023] Qian Chen, Zhiqiang Guo, Jianjun Li, and Guohui Li. Knowledge-enhanced multi-view graph neural networks for session-based recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 352–361, 2023.
- [Dong et al., 2025] Bingbing Dong, Chenyang Bu, Yi Zhu, Shengwei Ji, and Xindong Wu. Simplified multi-view graph neural network for multilingual knowledge graph completion. *Frontiers of Computer Science*, 19(7):197324, 2025.
- [Guo et al., 2020] Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering*, 34(8):3549–3568, 2020.
- [Hogan et al., 2021] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. Knowledge graphs. *ACM Computing Surveys (Csur)*, 54(4):1–37, 2021.
- [Huang et al., 2019] Xiao Huang, Jingyuan Zhang, Dingcheng Li, and Ping Li. Knowledge graph embedding based question answering. In *Proceedings of the twelfth ACM international conference on web search and data mining*, pages 105–113, 2019.
- [Jiang et al., 2024a] Wenyuan Jiang, Wenwei Wu, Le Zhang, Zixuan Yuan, Jian Xiang, Jingbo Zhou, and Hui Xiong. Killing two birds with one stone: Cross-modal reinforced prompting for graph and language tasks. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1301–1312, 2024.
- [Jiang et al., 2024b] Yuezihan Jiang, Changyu Li, Gaode Chen, Peiyi Li, Qi Zhang, Jingjian Lin, Peng Jiang, Fei Sun, and Wentao Zhang. Mmgcl: Meta knowledge-enhanced multi-view graph contrastive learning for recommendations. In *Proceedings of the 18th ACM Conference on Recommender Systems*, pages 538–548, 2024.
- [Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [Li et al., 2024] Duantengchuan Li, Fobo Shi, Xiaoguang Wang, Chao Zheng, Yuefeng Cai, and Bing Li. Multi-perspective knowledge graph completion with global and interaction features. *Information Sciences*, 666:120438, 2024.
- [Liang et al., 2024a] Wanying Liang, Pasquale De Meo, Yong Tang, and Jia Zhu. A survey of multi-modal knowledge graphs: Technologies and trends. *ACM Computing Surveys*, 56(11):1–41, 2024.
- [Liang et al., 2024b] Yan Liang, Weishan Cai, Minghao Yang, and Yuncheng Jiang. An unsupervised multi-view contrastive learning framework with attention-based reranking strategy for entity alignment. *Neural Networks*, 179:106583, 2024.
- [Liu et al., 2019] Ye Liu, Hui Li, Alberto Garcia-Duran, Mathias Niepert, Daniel Onoro-Rubio, and David S Rosenblum. Mmkg: multi-modal knowledge graphs. In *The Semantic Web: 16th International Conference, ESWC 2019, Portorož, Slovenia, June 2–6, 2019, Proceedings 16*, pages 459–474. Springer, 2019.
- [Luo et al., 2024] Yizhen Luo, Kai Yang, Massimo Hong, Xing Yi Liu, Zikun Nie, Hao Zhou, and Zaiqing Nie. Learning multi-view molecular representations with structured and unstructured knowledge. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2082–2093, 2024.
- [Ma et al., 2024] Ruixin Ma, Hao Wu, Xiaoru Wang, Weihe Wang, Yunlong Ma, and Liang Zhao. Multi-view semantic enhancement model for few-shot knowledge graph completion. *Expert Systems with Applications*, 238:122086, 2024.
- [Nicholson and Greene, 2020] David N Nicholson and Casey S Greene. Constructing knowledge graphs and their biomedical applications. *Computational and structural biotechnology journal*, 18:1414–1428, 2020.
- [Niu et al., 2022] Guanglin Niu, Bo Li, Yongfei Zhang, and Shiliang Pu. Cake: A scalable commonsense-aware framework for multi-view knowledge graph completion. *arXiv preprint arXiv:2202.13785*, 2022.
- [Pei et al., 2023] Shichao Pei, Ziyi Kou, Qiannan Zhang, and Xiangliang Zhang. Few-shot low-resource knowledge graph completion with multi-view task representation generation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1862–1871, 2023.
- [Radford, 2018] Alec Radford. Improving language understanding by generative pre-training. 2018.
- [Singhal, 2012] A. Singhal. Introducing the knowledge graph: Things, not strings. <https://www.blog.google/products/search/introducing-knowledge-graph-things-not/>, 2012. Accessed: 2025-05-24.
- [Sun et al., 2020] Rui Sun, Xuezhi Cao, Yan Zhao, Junchen Wan, Kun Zhou, Fuzheng Zhang, Zhongyuan Wang, and Kai Zheng. Multi-modal knowledge graphs for recommender systems. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pages 1405–1414, 2020.

- [Wang *et al.*, 2020] Meng Wang, Haofen Wang, Guilin Qi, and Qiushuo Zheng. Richpedia: a large-scale, comprehensive multi-modal knowledge graph. *Big Data Research*, 22:100159, 2020.
- [Wang *et al.*, 2022] Pancheng Wang, Shasha Li, Kunyuan Pang, Liangliang He, Dong Li, Jintao Tang, and Ting Wang. Multi-document scientific summarization from a knowledge graph-centric view. *arXiv preprint arXiv:2209.04319*, 2022.
- [Wang *et al.*, 2023] Chenxu Wang, Zhenhao Huang, Yue Wan, Junyu Wei, Junzhou Zhao, and Pinghui Wang. Fualign: Cross-lingual entity alignment via multi-view representation learning of fused knowledge graphs. *Information Fusion*, 89:41–52, 2023.
- [Wang *et al.*, 2024a] Hao Wang, Jing Yang, Laurence T Yang, Yuan Gao, Jihong Ding, Xiaokang Zhou, and Huazhong Liu. Mvtucker: Multi-view knowledge graphs representation learning based on tensor tucker model. *Information Fusion*, 106:102249, 2024.
- [Wang *et al.*, 2024b] Li-e Wang, Yuelan Qi, Zhigang Sun, and Xianxian Li. Mknbl: Joint multi-channel knowledge-aware network and broad learning for sparse knowledge graph-based recommendation. *Neurocomputing*, 575:127277, 2024.
- [Wang *et al.*, 2024c] Xiaofeng Wang, Zengjie Zhang, Guodong Shen, Shuaiming Lai, Yuntao Cheng, and Shuailei Zhu. Multi-view knowledge graph convolutional networks for recommendation. *Applied Soft Computing*, page 112633, 2024.
- [Xia *et al.*, 2024] Boshen Xia, Jiwei Qin, Lu Han, Aohua Gao, and Chao Ma. Knowledge filter contrastive learning for recommendation. *Knowledge and Information Systems*, 66(11):6697–6716, 2024.
- [Xiao *et al.*, 2019] Han Xiao, Yidong Chen, and Xiaodong Shi. Knowledge graph embedding based on multi-view clustering framework. *IEEE Transactions on Knowledge and Data Engineering*, 33(2):585–596, 2019.
- [Xu *et al.*, 2023] Peng Xu, Xiatian Zhu, and David A Clifton. Multimodal learning with transformers: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10):12113–12132, 2023.
- [Xu *et al.*, 2024] Hongcai Xu, Junpeng Bao, Hui Li, Chao He, and Feng Chen. A multi-view temporal knowledge graph reasoning framework with interpretable logic rules and feature fusion. *Electronics*, 13(4):742, 2024.
- [Yang *et al.*, 2024] Jing Yang, Xiaowen Jiang, Yuan Gao, Laurence T Yang, and Jieming Yang. Generalize to fully unseen graphs: Learn transferable hyper-relation structures for inductive link prediction. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 1274–1282, 2024.
- [Yang *et al.*, 2025] Xiao-Jun Yang, Yang-Hui Wu, Zhi-Hao Zhang, Jing Wang, and Fei-Ping Nie. Multitype view of knowledge contrastive learning for recommendation. *Neural Networks*, 181:106690, 2025.
- [Yuan *et al.*, 2024] Xu Yuan, Huinan Wu, Longfei Wang, Xiya Bu, Zhengnan Gao, and Ruixin Ma. Attribute mining multi-view contrastive learning network for recommendation. *Expert Systems with Applications*, 253:124224, 2024.
- [Zeng *et al.*, 2025] Jin Zeng, Nan Wang, and Jinbao Li. Knowledge-driven hierarchical intents modeling for recommendation. *Expert Systems with Applications*, 259:125361, 2025.
- [Zhang *et al.*, 2019] Qingheng Zhang, Zequn Sun, Wei Hu, Muhao Chen, Lingbing Guo, and Yuzhong Qu. Multi-view knowledge graph embedding for entity alignment. *arXiv preprint arXiv:1906.02390*, 2019.
- [Zhang *et al.*, 2023] Qinggang Zhang, Junnan Dong, Qiaoyu Tan, and Xiao Huang. Integrating entity attributes for error-aware knowledge graph embedding. *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [Zhang *et al.*, 2024a] Fuwei Zhang, Zhao Zhang, Fuzhen Zhuang, Zhiqiang Zhang, Jun Zhou, and Deqing Wang. Multi-view temporal knowledge graph reasoning. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 4263–4267, 2024.
- [Zhang *et al.*, 2024b] Kai Zhang, Zhiying Tu, Dianhui Chu, Xiaoping Lu, and Lucheng Chen. Aic: an industrial knowledge graph with abstraction-instance-capability reasoning abilities for personalized customization. *Journal of Intelligent Manufacturing*, 35(7):3419–3440, 2024.
- [Zhang *et al.*, 2024c] Xin Zhang, Yu Liu, and Zhehuan Zhao. Semantics driven multi-view knowledge graph embedding for cross-lingual entity alignment. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 11811–11815. IEEE, 2024.
- [Zhao *et al.*, 2017] Jing Zhao, Xijiong Xie, Xin Xu, and Shiliang Sun. Multi-view learning overview: Recent progress and new challenges. *Information Fusion*, 38:43–54, 2017.
- [Zhou *et al.*, 2024] Yunfeng Zhou, Cui Zhu, and Wenjun Zhu. Promvsd: Towards unsupervised knowledge graph anomaly detection via prior knowledge integration and multi-view semantic-driven estimation. *Information Processing & Management*, 61(4):103705, 2024.
- [Zhu *et al.*, 2022] Xiangru Zhu, Zhixu Li, Xiaodan Wang, Xueyao Jiang, Pinglei Sun, Xuwu Wang, Yanghua Xiao, and Nicholas Jing Yuan. Multi-modal knowledge graph construction and application: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 36(2):715–735, 2022.
- [Zou, 2020] Xiaohan Zou. A survey on application of knowledge graph. In *Journal of Physics: Conference Series*, volume 1487, page 012016. IOP Publishing, 2020.