Expert Systems With Applications 279 (2025) 127356

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Unified link prediction modeling for enhanced knowledge graph completion task☆

Tri D.T. Nguyen [a,b,*], Ubaid Ur Rehman [c], Musarrat Hussain [d], Rao Faizan [b], Jamil Hussain [e], Sung-Ho Bae [b], Jung Uk Kim [b], Seong Tae Kim [b,**], Sungyoung Lee [b,**]

[a] Sunchon National University, Suncheon-si, 57922, Republic of Korea
[b] Kyung Hee University (Global Campus), Yongin-si, 17104, Republic of Korea
[c] School of Computer Science and Engineering, Constructor University, Bremen, 28759, Germany
[d] UiT The Arctic University of Norway, Tromsø, 9019, Norway
[e] Department of Artificial Intelligence Data Science, Sejong University, Seoul, 209, Republic of Korea

ARTICLE INFO

Keywords:
Knowledge graph completion
Graph neural network
Link prediction

ABSTRACT

Link prediction in knowledge graphs (KGs) aims to identify missing links between entities. Existing studies primarily focus on specific scenarios, such as transductive (seen-to-seen) or inductive (seen-to-unseen, unseen-to-seen, or unseen-to-unseen) settings, individually. However, real-world challenges arise when both unseen entities and unseen relations appear simultaneously during testing, creating a more complex and realistic scenario. To address this gap, we propose a unified method with three key components designed to enhance adaptability and accuracy across diverse link prediction settings: (1) entity-independent modeling with triple-view graph, which employs graph neural networks (GNNs) to learn relation patterns independently of entities, enabling more effective inductive knowledge graph completion; (2) contrastive learning-based relation-context modeling, which mitigates the lack of connectivity information in KGs by allowing GNNs to propagate and learn meaningful representations for unseen entities; and (3) unseen relations modeling with augmented schema, which leverages the KG's ontological schema to flexibly model unseen relations and predict missing links without requiring extensive retraining. Extensive experiments on multiple benchmark datasets demonstrate the effectiveness of our framework in terms of accuracy and adaptability across various scenarios, outperforming state-of-the-art baselines and underscoring its potential for addressing real-world knowledge graph completion challenges.

1. Introduction

Knowledge graphs (KGs) organize factual information in triples (entity, relation, entity). Real-world KGs such as FreeBase (Bollacker et al., 2008) and DBpedia (Becker & Bizer, 2008) often suffer from incompleteness (Färber et al., 2018). Accurately predicting missing links is critical, as incomplete KGs can lead to inaccurate insights or decisions in applications like recommendation systems, medical diagnosis, and knowledge-based search engines. To address this issue, knowledge graph completion (KGC) methods, including link prediction (LP), focus on identifying missing connections between entities. Many studies have adopted KG embedding techniques, which represent entities and relations as vectors that capture their semantic properties. This vector representation enables computational prediction of missing links, as demonstrated in studies such as Wang et al. (2017). However, these methods typically operate in a transductive manner, meaning they can only predict triples involving entities and relations present in the training data (referred to as *seen entities* and *relations*). When new entities

☆ This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (IITP-2022-0-00078, Explainable Logical Reasoning for Medical Knowledge Generation). It was also supported by (a) the ITRC (Information Technology Research Center) support program (RS-2023-00259004) supervised by the IITP (Institute for Information & Communications Technology Planning & Evaluation), and (b) the MSIT (Ministry of Science and ICT), Korea, under the Grand Information Technology Research Center support program (IITP-2024-2020-0-01489). This work was also supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00252391).

* Corresponding author at: Sunchon National University, Suncheon-si, 57922, Republic of Korea.
** Corresponding authors.
E-mail addresses: tringuyendt@khu.ac.kr (T.D.T. Nguyen), urehman@constructor.university (U.U. Rehman), musarrat.hussain@uit.no (M. Hussain), rao.faizan@oslab.khu.ac.kr (R. Faizan), jamil@sejong.ac.kr (J. Hussain), shbae@khu.ac.kr (S.-H. Bae), ju.kim@khu.ac.kr (J.U. Kim), st.kim@khu.ac.kr (S.T. Kim), sylee@oslab.khu.ac.kr (S. Lee).

https://doi.org/10.1016/j.eswa.2025.127356
Received 31 July 2024; Received in revised form 10 March 2025; Accepted 19 March 2025
Available online 3 April 2025

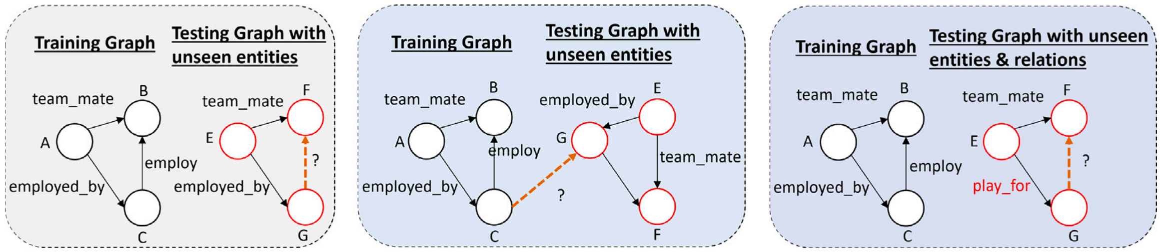

Fig. 1. Consideration of LP scenarios: (a) a testing graph with unseen entities; (b) a testing graph with an empty enclosing subgraph; (c) a testing graph featuring both new entities and new relations.

or relations (termed *unseen entities or relations*) appear during testing, retraining the entire KG to generate embeddings becomes impractical due to the significant computational resources required for large-scale, dynamically evolving KGs (P1).

A widely studied and notable approach for handling LP tasks involving previously unseen entities is the inductive method, which aims to infer valid missing links for these entities. Inspired by the ability of graph neural networks (GNNs) to aggregate local information, several inductive models have been proposed. The core idea is to derive entity-independent logical rules that apply to relationships within the KG in an end-to-end manner, as discussed in Meilicke et al. (2018), Sun et al. (2019). For example, a rule such as $employed\_by(A, C) \wedge team\_mate(A, B) \rightarrow employ(C, B)$ can be derived from the training graph in Fig. 1 and then used to infer a testing triple involving unseen entities, such as $(G, employ, F)$. GraIL (Teru et al., 2020) is a pioneering study that enables inductive link inference by reasoning over an enclosing subgraph structure extracted from a target entity's k-hop neighbors. GraIL and its subsequent extensions (Chen et al., 2021; Lin et al., 2022) have demonstrated the viability of using GNN-based message passing on subgraphs for inductive KGC. However, these methods face a critical *topological limitation*: when no connected subgraph exists around the target link, the model lacks essential connectivity information between unseen entities and the established KG. This limitation weakens GNNs' ability to learn effectively, as they rely heavily on graph structure for meaningful representation learning (P2).

In practical applications, real-world KGs are rarely static, as they continuously evolve with new concepts and relations added regularly, particularly in fields such as healthcare, e-commerce, and social media. Without a strategy to manage these unseen relations, link prediction models quickly become outdated, potentially compromising the quality and relevance of predictions in dynamically growing knowledge domains. To address this issue, several studies have proposed methods for handling LP tasks involving unseen relations without retraining KGs from scratch. The objective of these approaches is to derive embeddings for previously unseen relations using external resources, such as textual descriptions (Geng et al., 2021; Wang et al., 2021) or auxiliary triples. While effective, these methods pose practical challenges related to resource availability and computational cost. Notably, approaches that rely on external textual data for relation representation require significant processing power for text embedding generation and integration, making them impractical for dynamically evolving graphs. Moreover, dependence on external sources can lead to inconsistencies, particularly when such resources are sparse or domain-specific, resulting in limited adaptability and potential data alignment issues. These challenges highlight the need for an internally driven, efficient method to handle unseen relations directly within the KG's ontological framework (P3).

Based on these observations, we propose a unified link prediction method that not only effectively infers missing links in partially inductive scenarios (e.g., considering only unseen entities in a testing graph) but also seamlessly operates in fully inductive LP settings (e.g., when both unseen entities and relations occur simultaneously in testing). Specifically, the proposed method addresses problem (P1) by introducing the entity-independent modeling using triple-view graph (EITripVG) component. Unlike previous work (Liu et al., 2021), our approach begins by extracting an enclosing subgraph for each triple rather than relying on the entire KG. We then construct a triple-view transformed graph to explicitly capture reasoning rules for relationships between entities through GNN-based message passing. To tackle problem (P2), we introduce the contrastive learning-based relation-context modeling (CLRCM) component. CLRCM encodes the embeddings of entities in the target triple based on their associated relations and then uses these embeddings to infer missing links without relying on GNNs' aggregation mechanism. Finally, inspired by Geng, Chen, Pan et al. (2023), we leverage the KG's ontological schema, which is commonly available and contains complementary relational semantics, to develop the unseen relations modeling with augmented schema (URMAS) component for addressing problem (P3). URMAS encodes features of unseen relations based on their associated seen relations in the ontological schema and utilizes them for link prediction through the CLRCM component.

To summarize, our primary contributions include: (1) Developing a comprehensive method that effectively addresses various challenging link prediction (LP) scenarios, including both partially and fully inductive settings. (2) Training our method in an end-to-end manner to minimize computational costs by reasoning over an enclosing graph and incorporating pruning techniques in the GNN-based message-passing process. (3) Demonstrating the effectiveness of our approach through experimental results, which show superior performance compared to existing methods.

The remainder of this paper is structured as follows. Section 2 reviews related work on knowledge graph completion and inductive link prediction methods. Section 3 defines the problem and highlights the key challenges addressed in this study. Section 4 details the methodology, including the design and implementation of the proposed framework. Section 5 describes the experimental setup, evaluation metrics, and results. Section 6 discusses the findings, implications, and limitations. Finally, Section 7 concludes the paper and outlines potential directions for future research.

2. Related work

2.1. Transductive methods

Transductive methods learn embeddings for each entity and relation while preserving the inherent semantics of KGs. These methods are categorized into three main approaches: (a) translation-based methods, such as TransE and RotatE (Bordes et al., 2013; Sun et al., 2019); (b) factorization-based methods, such as TuckER (Balažević et al., 2019); and (c) GNN-based methods, including R-GCN and CompGCN (Schlichtkrull et al., 2018; Vashishth et al., 2019). However, these methods assume that knowledge graphs remain static, limiting their ability to predict missing links involving previously unseen entities.

2.2. Inductive methods

2.2.1. Partially inductive

Partially inductive methods are designed to generalize and reason about previously unseen entities. Current approaches fall into two categories: rule-based methods and graph-based methods. Rule-based methods explicitly learn logical rules for reasoning, which are independent of specific entities, allowing them to infer connections between unseen nodes. Some studies, such as RUGE and RLvLR (Guo et al., 2018; Omran et al., 2019), have introduced frameworks that use end-to-end training to jointly learn logical rules and their confidence levels. However, these methods ignore the topological information in KGs, resulting in low expressive ability and poor scalability for large-scale KGs. On the other hand, graph-based methods have recently emerged as promising approaches for the inductive LP task. GraIL, introduced by Teru et al. (2020), was the first method designed to model the structures of enclosing subgraphs around target triples for inferring missing links. This technique effectively addresses the scalability issues typically associated with rule-based methods. Later, ConGLR (Lin et al., 2022) incorporated relational paths and enclosing subgraphs to improve relation-context entity embeddings, thereby enhancing the ability to reason about missing links in an entity-independent manner.

Recent studies have further advanced graph-based methods. TACO (Wang et al., 2023) models correlations between relations by integrating topology-aware patterns, effectively capturing graph-level and edge-level interactions through a relational correlation network. This approach enhances generalization for inductive link prediction by leveraging distinct topological correlations. Similarly, SE4LP (Si et al., 2024) employs subgraph-centric embedding learning to scale inductive LP tasks. By applying contrastive discrimination on induced subgraphs, SE4LP captures both intrinsic and structural link features, improving model performance and scalability. InGram (Lee et al., 2023) introduces a relation graph-based inductive KG embedding approach, using auxiliary relation graphs to capture hierarchical patterns and refine embeddings, thereby enhancing KG models' adaptability to unseen entities. Additionally, SiaILP (Zhang & Liu, 2024), a path-based inductive approach utilizing Siamese networks, emphasizes path and relation embeddings to generalize across unseen entities without retraining, achieving state-of-the-art benchmarks on inductive KG datasets.

Nevertheless, these methods face several challenges: (1) When the enclosing subgraph is sparse or nonexistent, they struggle to infer missing links inductively, as GNNs lack the necessary structural information to generate meaningful embeddings. (2) They often assume that the relationships present during testing are identical to those in the training phase, which does not accurately reflect real-world KGs, particularly those constructed using open information extraction systems like NELL or publicly editable platforms like Wikidata.

2.2.2. Fully inductive

Fully inductive approaches refer to methods capable of inferring missing links involving both new entities and new relations during the testing phase. Currently, only a few methods have been proposed to handle this challenging LP scenario. MaKEr (Chen et al., 2022) is one such method that focuses on fully inductive LP. The authors apply meta-learning techniques to construct a set of unseen entities and relations for LP tasks during training. They then use GNNs to encode features and generate embeddings for both unseen entities and relations based on structural information in the KG. Through this approach, embeddings for unseen entities and relations can be obtained via knowledge extrapolation. Additionally, the authors in Geng, Chen, Pan et al. (2023) proposed a relational message-passing network for fully inductive KGC. Their method leverages the KG's ontological schema to encode embeddings for unseen relations based on associated seen relations. Although these approaches effectively handle fully inductive LP tasks, they are designed to work with specific LP scenarios, resulting in narrow solutions that lack generalizability across diverse LP settings.

3. Problem formulation

A KG is defined as $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$, where $\mathcal{E}$ and $\mathcal{R}$ denote the set of entities and relations, respectively. $\mathcal{T} = \{(h,r,t) | h,t \in \mathcal{E}; r \in \mathcal{R}\}$ represents a set of triples. The link prediction task aims to predict a missing relation $r_{pred} \in \mathcal{R}$, given entities $h, t, \in \mathcal{E}$, denoted as (h, r_{pred}, t), to enhance the completeness of $\mathcal{G}$. Notably, this formulation does not explicitly indicate that a relation between h and t exists; rather, it queries the system to determine if a plausible relation can connect these two entities. During the testing phase, a set of new entities, $\mathcal{E}'$, is introduced, where $\mathcal{E}' \cap \mathcal{E} = \emptyset$, to predict a triple (h,r,t) with $r \in \mathcal{R}$ and $h,t \in \mathcal{E}'$. Additionally, we examine a testing scenario that incorporates a set of unseen relations, $\mathcal{R}'$, which are distinct from the existing relations, $\mathcal{R}$, such that $\mathcal{R}' \cap \mathcal{R} = \emptyset$.

This study examines three specific test scenarios to explore link prediction adaptability in knowledge graphs. The first scenario, a test graph with unknown entities and known relations, evaluates the model's ability to generalize to new entities while leveraging familiar relational structures. The second scenario, a test graph with an empty surrounding subgraph, assesses the model's robustness in cases with minimal structural information, where conventional GNN-based methods may struggle. The third scenario, involving both novel entities and relations, measures the model's effectiveness in fully inductive settings, where it encounters previously unseen entity and relation types. This structured approach enables a comprehensive evaluation of the proposed method's adaptability to partially and fully inductive link prediction scenarios, as well as sparse graph structures, addressing key challenges in knowledge graph completion.

4. Methodology

This section details our proposed methodology. As illustrated in Fig. 2, our approach dynamically adapts to each scenario by leveraging a relation index mapping created during training to invoke the appropriate components during testing. If the system encounters a new relation (i.e., one with no existing mapping), it activates the unseen relations modeling with augmented schema component, which utilizes the KG's ontological schema to generate embeddings for the unseen relation. Conversely, if no new relation is detected, the method triggers the entity-independent modeling using triple-view graph component. This component applies triple-view graph transformation and a GNN-based message-passing strategy to perform partially inductive link prediction. The results are then integrated with the output of the contrastive learning-based relation-context modeling component, which encodes entity embeddings based on their relational context, to predict the missing link for a given query. The following sections provide a detailed explanation of each step in our methodology.

4.1. Entity-independent modeling using triple-view graph

Subgraph extraction. Building upon the method presented in GraIL (Teru et al., 2020), our first step involves extracting an enclosing subgraph $\mathcal{G}_s(u, r_t, v)$ centered around the target triple (u, r_t, v). The process follows these steps: First, we extract two k-hop neighborhood sets, $\mathcal{N}_k(u)$ and $\mathcal{N}_k(v)$, for the two target nodes u and v, respectively. Next, the enclosing subgraph is constructed by intersecting $\mathcal{N}_k(u) \cap \mathcal{N}_k(v)$ to identify the common nodes within a k-hop distance.

Triple-view graph transformation. Inductive LP requires performing the LP task in an entity-independent manner. However, conventional GNNs do not meet this requirement, as they encode entity embeddings and update features in an entity-dependent way. To address this limitation, we transform the enclosing subgraph $\mathcal{G}_s$ into a new graph structure, allowing the GNN to infer missing links without relying on specific entity embeddings.

$$\mathcal{G}_s(u, r, v) \xrightarrow{\text{transform}} \mathcal{G}_s^{\text{triple\_view}}(u_t, r_t, v_t). \tag{1}$$

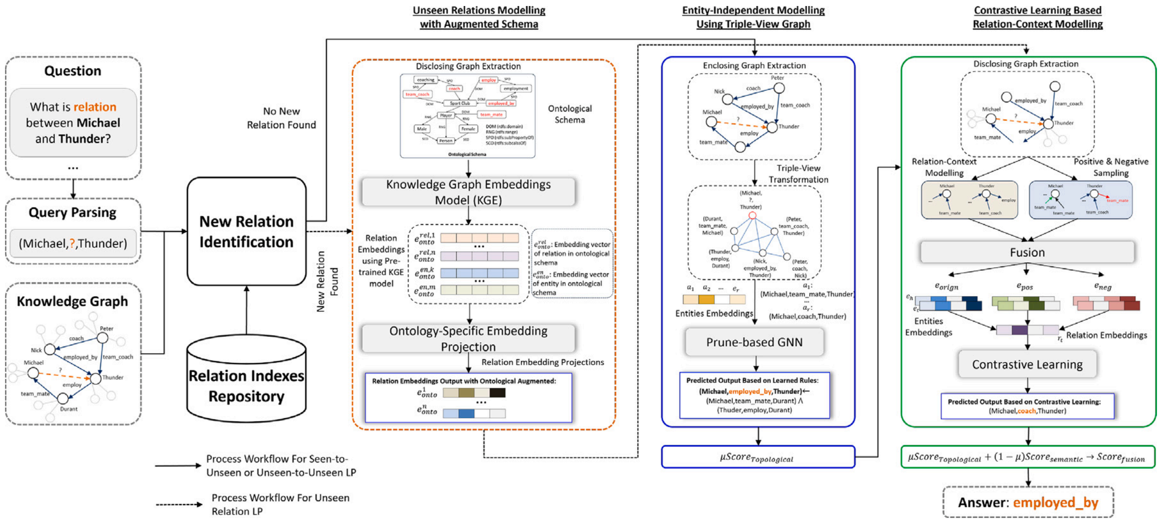

Fig. 2. Overview of our proposed method: (1) constructing a framework to predict missing links in partially inductive settings; (2) developing an approach to infer missing links when the enclosing subgraph is empty; (3) utilizing a strategy to handle the LP task with unseen relations.

where u_t, v_t, and r_t represent new entities and relations in the transformed graph, respectively. The details of the graph transformation are as follows: Each triple $(u, r, v) \in \mathcal{G}_s$ is converted into a pair of entities, represented as a new node in $\mathcal{G}_s^{triple\_view}$. Specifically, the transformation follows the rule $(u, r, v) \rightarrow u_t \equiv (u, v)$ if $idx(u) > idx(v)$, and vice versa, where $idx(.)$ denotes the *index* of an entity stored in the graph database. The relation r_t in the transformed graph is established by connecting any two entities u_t and v_t if they share a common instance. For example, $(u_t, r_t, v_t) \iff u_t \equiv (u, \underline{v}),\ v_t \equiv (\underline{v}, w); \forall u, v, w \in \mathcal{G}_s$. Additionally, individual entities u or $v \in \mathcal{G}_s$ are also transformed into new entities (u, u) or (v, v) in $\mathcal{G}_s^{triple\_view}$, respectively.

After completing the graph transformation, we initialize the node feature $\mathbf{h}_i^{(0)}$ for each entity $i \in \mathcal{G}_s^{triple\_view}$, $i \equiv (u, v)$, and $u, v \in \mathcal{G}_s$:

$$\mathbf{h}_i^{(0)} = \left[\text{one-hot}(l(u, v))\right]. \tag{2}$$

The function $l(.)$ outputs 1 if a plausible triple (u, r, v) exists in the KG $\mathcal{G}_s$ for any $r \in \mathcal{R}$ and 0 otherwise. Additionally, we encode the reverse direction of each node i, for example, (v, r^{-1}, u), using a similar representation. Consequently, the dimension of each node feature in the triple-view graph is $h_i \in \mathbb{R}^{2\times|\mathcal{R}|}$.

Prune-based GNN. We adapt a GNN to iteratively update node features in $\mathcal{G}_s^{triple\_view}$ by aggregating features from neighboring nodes. The updated features of entity i in the kth GNN layer are defined as follows:

$$\mathbf{h}_i^{(k)} = \text{ReLU}\left(\sum_{j\in\mathcal{N}_i} \alpha_{ij}^{(k)} \mathbf{W}^{(k)} \mathbf{h}_j^{(k-1)} + \mathbf{W}_{\text{self}}^{(k)} \mathbf{h}_i^{(k-1)}\right), \tag{3}$$

$$\alpha_{ij}^{(k)} = \sigma\left(\mathbf{A}_1^{(k)} \mathbf{s}_{ij}^{(k)} + \mathbf{b}_1^{(k)}\right), \tag{4}$$

$$\mathbf{s}_{ij}^{(k)} = \text{ReLU}\left(\mathbf{A}_2^{(k)} \left[\mathbf{h}_i^{(k-1)} \oplus \mathbf{h}_j^{(k-1)}\right] + b_2^{(k)}\right). \tag{5}$$

In this context, $\mathcal{N}_i$ represents all neighbors of entity i. $\mathbf{W}^{(k)}$ is the transformation matrix learned at the kth layer, while $\mathbf{W}_{self}^{(k)}$ is the matrix used for aggregating features from the node itself. The attention weight for the edge connecting nodes i and j at the kth layer is represented by $\alpha_{ij}^{(k)}$, computed using latent embeddings from the preceding layer. The attention weight matrix at layer k is denoted as $\mathbf{A}_1^{(k)}$, and the weight matrix for intermediate features at the same layer is $\mathbf{A}_2^{(k)}$. Finally, $b_1^{(k)}$ and $b_2^{(k)}$ are the bias terms for layer k.

It is important to note that the node feature update process requires all nodes to be updated recursively. To reduce computational overhead during message passing, we employ an efficient pruning strategy, as described by Li et al. (2020). In this strategy, the features of node i at the kth layer are updated using only a subset of $|N - k|$ nodes, where N represents the total number of nodes in the triple-view graph $\mathcal{G}_s^{triple\_view}$. Consequently, the feature updated for node i in Eq. (3) is computed using only a subset of nodes at each kth layer. The node feature at the final layer K, obtained through the above update process, is then used to infer missing links for a pair of entities (u, v) in the original graph $\mathcal{G}_s$ since entity $i \in \mathcal{G}_s^{triple\_view}$ corresponds to a triple $(u, r, v) \in \mathcal{G}_s$. The final score of the target triple (u, r_t, v) is denoted as $score_{topo}(u, r_t, v)$ and is computed as follows:

$$\text{score}_{\text{topo}}(u, r_t, v) = \max_p \left(\mathbf{h}_i^{(K)}\right), \tag{6}$$

where $\mathbf{h}_i^{(K)}$ represents the embedding of the pair (u, v) in the triple-view graph $\mathcal{G}_s^{triple\_view}$ at the final layer K of GNN, and $\max_p$ denotes the element p with the highest value in $\mathbf{h}_i^{(K)}$. During training, we use a margin-based loss function that ensures a balanced number of negative triples by replacing either the head or tail entity:

$$\mathcal{L}_t = \sum_{(u,r_t,v)\in\mathcal{G}_s} \max\Bigg(0, \text{score}_{\text{topo}}(u', r_t, v') - \text{score}_{\text{topo}}(u, r_t, v) + \epsilon\Bigg). \tag{7}$$

In this equation, ϵ is the margin hyperparameter, and (u, r_t, v) and (u', r_t, v') represent the positive and negative samples, respectively.

4.2. Contrastive learning based relation-context modeling

To address the scenario of an empty enclosing subgraph, we extract a disclosing subgraph $\mathcal{DG}_{(u,r_t,v)}$ for each triple (u, r_t, v) as input to this component. The extraction process follows a similar workflow to that of the enclosing subgraph, but instead of finding the intersection, we compute the difference $\mathcal{N}_k(u) \setminus \mathcal{N}_k(v)$ to obtain the disclosing subgraph.

Relation-context modeling. In KGs, the semantic interpretation attributed to an entity is inherently influenced by the relations it is associated with. For example, *Michael* may be recognized as a *Sport player* due to the relation *teammate*, while *Thunder* is considered a *sports employer* because of its associated relations *employ* and *employed_by*. Consequently, an appropriate relation between a *sports player* and a *sports employer* should be *employed_by*. Following this intuition, we encode entity features by fusing features of their associated relations. We denote a set of features specific to relations $\mathcal{R}$ as $\mathcal{F}$:

$$\mathcal{F} = \left\{\mathbf{f}_k \mid r_k \in \mathcal{R}\right\}, \tag{8}$$

where $\mathbf{f}_k \in \mathbb{R}^d$ is a learned embedding representing the semantics of each relation $r_k \in \mathcal{R}$. The semantic information of an entity $u_i \in \mathcal{E}$ is then modeled as a relation-contextual vector, denoted as:

$$\mathcal{M}_i = \left\{ m_i^k \mid u_i \in \mathcal{E}, r_k \in \mathcal{R} \right\}, \tag{9}$$

where $m^k i$ represents the count of triples involving the entity u_i with the relation r_k. If u_i is not associated with r_k, then $m_i^k = 0$. It is important to note that the relation-contextual vector for each entity depends solely on the entity's associated relations. As a result, Eq. (9) applies to both known entities in $\mathcal{G}$ and unknown entities in $\mathcal{G}'$.

Positive and negative sampling. Motivated by the advantages of contrastive learning in graph representation learning (Sun et al., 2021), we adopt a similar approach to train the LP task through a contrastive learning process. To achieve this, we sample both positive and negative triples for training. Positive triples are generated by randomly increasing the number of existing associated relations r_k for each entity u_i. This follows the intuition that an entity's inherent semantic characteristics remain unchanged when no new relations are introduced. In contrast, negative triples are sampled by randomly adding new relations to entity u_i. For example, the relation-specific context of the entity *Michael* would change significantly if new relations such as *employ* or *father_of* were introduced, as the new semantics of *Employer* and *Father* would become associated with him.

Fusion. Once positive and negative triples are obtained, each entity is represented by integrating its associated relation features. The semantic representation of an entity $\mathbf{u}_i$ for entity i in the original graph $\mathcal{G}_s$ is defined as follows:

$$\mathbf{u}_i = \frac{\sum_{k=0}^{|\mathcal{R}|-1} m_i^k \cdot \mathbf{f}^k}{\sqrt{\sum k = 0^{|\mathcal{R}|-1} (m_i^k)^2}}. \tag{10}$$

This formulation ensures that $\mathbf{u}_i$ resides in Euclidean space. Similarly, we can obtain the semantic representation of entity $\mathbf{u}_i^{pos}$ in positive triples and entity $\mathbf{u}_i^{neg}$ in negative triples using their corresponding relation-contextual vectors $\mathcal{M}_i^{pos}$ and $\mathcal{M}_i^{neg}$.

Contrastive learning. We design a contrastive learning loss using a triplet loss function, which aims to increase the similarity between the positive pair $(\mathbf{u}^{pos} i, \mathbf{u} i)$ and decrease the similarity between the negative pair $(\mathbf{u}^{neg} i, \mathbf{u} i)$. Formally, the loss function is defined as:

$$\mathcal{L}_c = \max\left(\mathrm{sim}(\mathbf{u}_i^{\mathrm{pos}}, \mathbf{u}_i) - \mathrm{sim}(\mathbf{u}_i^{\mathrm{neg}}, \mathbf{u}_i) + \gamma, 0\right), \tag{11}$$

where γ is the hyperparameter that defines the margin, and $sim(.)$ measures the similarity between two features using Euclidean distance. Ultimately, the score function that determines the semantic likelihood of the triple (u, r_t, v) is defined as follows:

$$\mathrm{score}_{\mathrm{sem}}(u, r_t, v) = \left\langle \mathbf{u}, \mathbf{r}_t^{\mathrm{sem}}, \mathbf{v} \right\rangle, \tag{12}$$

where $\mathbf{r}_t^{sem} \in \mathbb{R}^d$ is the learned embedding of relation r_k, and $\langle . \rangle$ denotes the element-wise product.

4.3. Unseen relations modeling with augmented schema

To address the challenge of unseen relations in the link prediction task, our proposed method incorporates the ontological schema of the KG in an a priori manner (Geng, Chen, Pan et al., 2023). A KG is typically accompanied by an ontology that serves as its schema, enhancing semantic richness and improving overall quality. Ontologies contain structured vocabularies that define complex semantic relationships between both known and unknown relations. By leveraging information from the ontology associated with a given KG, we can predict triples involving previously unseen relations more effectively.

We begin by pretraining the schema graph using KG embedding techniques, such as TransE, to generate embedding vectors for both familiar and unfamiliar relations (Bordes et al., 2013). These learned vectors are then projected into the same embedding space as relations $r_t \in \mathbb{R}^d$ to facilitate their participation in the link prediction task. The mapping function is implemented using two fully connected linear layers, as follows

$$\mathbf{h}_{r_t}^0 = \mathbf{W}_1(\mathbf{W}_2 \mathbf{h}_{r_t}^{\mathrm{onto}}), \tag{13}$$

where $\mathbf{h}_{r_t}^{onto}$ is the embedding vector of r_i learned by a KG embedding method on the schema graph. The projected vector $\mathbf{h}_{r_t}^0$ is then used as the initial representation of relations in the CLRM component to perform the LP process. It is important to note that EITripVG does not participate in the LP process when unseen relations occur, as it is limited to modeling only seen relations during training.

4.4. Joint training strategy

The final objective of the learning process integrates the supervised loss from Eq. (7) with the contrastive learning loss from Eq. (11):

$$\mathcal{L} = \mathcal{L}_p + \varrho \mathcal{L}_c, \tag{14}$$

where ϱ controls the influence of the contrastive learning mechanism. By employing this joint training strategy, the proposed method effectively models various LP scenarios, enabling the inference of missing links in both partially and fully inductive settings.

5. Experiments

5.1. Experimental configurations

Dataset. To evaluate partially inductive link prediction, we focus on entities that have not been previously encountered, using benchmarks identified in GraIL (Teru et al., 2020). These benchmarks are derived from three well-known transductive knowledge graph completion datasets: FB15k-237, NELL-995, and WN18RR (Dettmers et al., 2018; Toutanova et al., 2015; Xiong et al., 2017). To ensure a comprehensive evaluation, four inductive benchmark variants of varying sizes are created for each dataset. Each benchmark consists of a distinct training graph and testing graph, with non-overlapping entity sets. During training, 80% of the triples are used for model training, while an additional 10% are set aside for validation (see Table 1).

To evaluate link prediction performance in scenarios where enclosing subgraphs are absent, we follow the strategy outlined in Zhang et al. (2023) to develop specialized benchmark datasets. Specifically, we introduce three additional datasets: FB15k-237-emsg, NELL-995-emsg, and WN18RR-emsg. These datasets consist of testing graphs where the target triples (h, r_{pred}, t) do not have an enclosing subgraph between h and t. To ensure balanced evaluation, the ratio of test cases with subgraph information to those without is maintained at 1:1 within the testing graphs. Table 2 provides detailed statistics on the number of relations, entities, and triples for these benchmark datasets.

To evaluate fully inductive scenarios, where the testing graph contains previously unseen relations, we use benchmarks established by Geng, Chen, Pan et al. (2023). Specifically, we modify the original GraIL benchmarks by replacing the testing graphs with new ones that include unseen relations, while keeping the training graphs unchanged. For example, in the second benchmark of NELL-995 (version 2), which contains 88 relations, we replace the original testing graph with one from NELL-995 (version 3), which includes 122 relations, 51 of which are not present in version 2. This adjustment creates a new benchmark for the fully inductive setting using the NELL-995 dataset. The detailed configuration of the benchmark datasets used for evaluating fully inductive scenarios is presented in Table 3.

In our approach, we enhance relation semantics by incorporating knowledge graph ontologies. To ensure broad applicability, we develop schema graphs that capture relation-aware semantics by extracting and restructuring information from public ontologies of specific datasets, such as NELL-995. We specifically use the schema described

Table 1
Partially inductive LP benchmark datasets with unseen entities.

		FB15k-237			NELL-995			WN18RR		
		#R	#E	#T	#R	#E	#T	#R	#E	#T
v1	TR	180	1594	5226	14	3103	5540	9	2746	6678
	TE	142	1093	2404	14	225	1034	8	922	1991
v2	TR	200	2608	12 085	88	2564	10 109	10	6954	18 968
	TE	172	1660	5092	79	2086	5521	10	2757	4863
v3	TR	215	3668	22 394	142	4647	20 117	11	12 078	32 150
	TE	183	2501	9137	122	3566	9668	11	5084	7470
v4	TR	219	4707	33 916	76	2092	9289	9	3861	9842
	TE	200	3051	14 554	61	2795	8520	9	7084	15 157

Table 2
LP benchmarks with empty subgraphs.

	FB15k-237-emsg		NELL-995-emsg		WN18RR-emsg	
	#R	#E	#T	#R	#E	#T
TR	180	1594	14	3103	9	2746
TE	142	1093	14	225	8	922

Table 3
Fully inductive LP benchmarks with unseen relations.

	NELL-995.fully.v1			NELL-995.fully.v2		
	#R	#E	#T	#R	#E	#T
TR	14	3103	5540	88	2564	10 109
TE	106	2271	5550	116	2803	6749
	NELL-995.fully.v3					
	#R	#E	#T			
TR	76	2092	9289			
TE	110	3140	8308			

in Geng, Chen, Zhuang et al. (2023), which consists of 1,186 nodes and 3,055 triples, covering all relations and their associated semantics. Future research will focus on evaluating ontology-enhanced settings for WN18RR and FB15k-237, as these datasets currently lack publicly available ontologies.

Evaluation metric. To ensure a comprehensive assessment, we evaluate performance using triple classification and entity prediction tasks. The goal of triple classification is to determine the validity of a given triple using a binary classification metric—the area under the precision–recall curve (AUC-PR). For each positive triple in the test set, a corresponding negative triple is sampled for comparison. For incomplete triples (h, r_{pred}, t), link prediction involves ranking candidate relations based on their predicted likelihood of being the missing relation. Better performance is indicated when the true relation ranks higher. The key performance metrics include mean reciprocal rank (MRR) and Hits@n. MRR is the average ranking position of the correct relation, while Hits@n is the percentage of test triples where the correct relation appears within the top-n predictions. To assess semi-inductive knowledge graph completion (KGC), we employ standard metrics like Hits@10 and AUC-PR. For a comprehensive evaluation of fully inductive KGC, we use AUC-PR, MRR, and Hits@10. In cases where subgraphs are empty, link prediction performance is evaluated using mean reciprocal rank (MRR) and Hits@1, Hits@5, and Hits@10. For both triple classification and entity prediction tasks, negative triples are generated by randomly replacing either the head or the tail entity. Each experiment is conducted five times, and the average results are reported to ensure robust comparison.

Hyperparameter settings. For subgraph extraction, a *3*-hop enclosing subgraph is used. During training, the hyperparameters are set as follows: a learning rate of 0.001, a dropout rate of 0.5, and an embedding dimension of 32. These configurations have been determined to be optimal based on the validation set. We finetune the margin parameter

Table 4
Link prediction with only unseen entities (Hits@10).

Datasets	Methods	v1	v2	v3	v4
FB15k-237	GraIL	62.19	80.67	81.53	87.38
	INDIGO	62.82	82.21	82.93	89.45
	Ingram	35.77	40.76	40.92	36.21
	SiaILP	69.31	80.98	80.78	79.42
	Ours	**65.49**	**82.75**	**83.33**	**89.57**
WN18RR	GraIL	82.45	78.68	58.43	73.41
	INDIGO	82.43	78.75	58.60	73.70
	Ingram	42.63	45.01	40.53	44.06
	SiaILP	83.41	83.55	74.18	80.79
	Ours	**83.55**	**79.81**	**59.71**	**73.94**
NELL-995	GraIL	59.50	93.25	91.41	73.19
	INDIGO	59.61	93.50	92.00	80.82
	Ingram	37.70	42.92	44.79	41.29
	SiaILP	83.30	68.97	64.08	67.43
	Ours	**60.00**	**93.86**	**93.11**	**88.45**

Table 5
Triple classification with only unseen entities (AUC-PR).

Datasets	Methods	v1	v2	v3	v4
FB15k-237	GraIL	83.16	89.91	90.13	93.24
	INDIGO	85.34	91.12	92.08	95.23
	Ingram	61.83	78.20	67.70	74.45
	SiaILP	67.37	70.98	70.54	70.83
	Ours	**85.75**	**91.47**	**93.08**	**95.73**
WN18RR	GraIL	93.12	93.01	84.40	91.52
	INDIGO	94.40	94.25	86.23	93.03
	Ingram	89.81	90.71	76.34	82.41
	SiaILP	70.76	69.59	65.06	66.34
	Ours	**95.62**	**94.71**	**86.54**	**93.89**
NELL-995	GraIL	85.14	91.56	92.42	86.83
	INDIGO	86.51	93.02	93.12	88.20
	Ingram	58.53	73.73	57.74	79.32
	SiaILP	58.03	67.03	68.31	61.58
	Ours	**86.80**	**93.60**	**93.02**	**91.38**

ϵ in Eq. (7) over the set $\{0.0005, 1e-7, 5e-8, 1e-9\}$, the margin loss γ in Eq. (11) over $\{0.5, 1, 5, 10\}$, and the joint margin loss hyperparameter ϱ in Eq. (14) over $\{0.01, 0.1, 0.5, 1\}$. The optimal configuration is found to be $\epsilon = 5e-8$, $\gamma = 1$, and $\varrho = 0.1$. The model is trained using the Adam optimizer. All experiments are conducted using PyTorch and executed on an NVIDIA RTX 3080.

Baselines. For partially inductive link prediction (LP), we compare our model to GraIL and INDIGO, two pioneering methods (Liu et al., 2021; Teru et al., 2020). GraIL is a graph-based model that performs inductive reasoning for missing link prediction by leveraging enclosing subgraphs. INDIGO is also a graph-based approach but infers missing links by reasoning within a triple-view transformation graph. Additionally, we incorporate two recent methods: Ingram (Lee et al., 2023), which employs relation-level aggregation to generate embeddings and has demonstrated effectiveness in partially inductive LP tasks, and SiaILP (Zhang & Liu, 2024), a path-based Siamese neural network

Table 6
Link prediction with empty subgraph (MRR and Hits@n).

Datasets	Models	MRR	Hits@1	Hits@5	Hits@10
FB15k-237-emsg	GraIL	27.34	21.46	32.45	34.61
	TACT	22.98	13.17	31.82	40.43
	Ours	**29.91**	**31.21**	**34.90**	**40.49**
NELL-995-emsg	GraIL	19.63	11.25	**23.87**	39.02
	TACT	15.47	7.24	22.46	32.98
	Ours	**20.30**	**18.25**	23.00	**40.11**
WN18RR-emsg	GraIL	40.58	31.72	46.95	61.07
	TACT	43.97	32.56	**57.24**	59.18
	Ours	**49.08**	**38.11**	55.31	**62.45**

model designed to generalize without fine-tuning, achieving robust performance across diverse inductive scenarios. For empty subgraph link prediction, we compare our model to GraIL and TACT (Chen et al., 2021), as both were specifically designed to address this scenario. For fully inductive LP, we use TACT and RMPI (Geng, Chen, Pan et al., 2023) as baseline methods, as both were developed to handle fully inductive link prediction tasks.

To reproduce the results of the comparison methods, we used the publicly available source code provided by their authors. All experiments were conducted on the same datasets generated in Section 5.1 to ensure a fair comparison. Hyperparameters and experimental settings were kept consistent with those reported in the respective publications to maintain reproducibility. This approach ensures that our evaluation framework adheres to standardized protocols, allowing for accurate performance benchmarking.

5.2. Main results

Partially inductive LP. Our evaluation results demonstrate a significant improvement in the partially inductive LP task, particularly in terms of the Hits@10 metric across various benchmarks in Table 4. For the FB15k-237 dataset, our proposed method achieves performance gains over previous approaches. Our method demonstrates notable advancements across all FB15k-237 versions, consistently outperforming GraIL and INDIGO, with particularly substantial improvements in FB15k-237 v4. Ingram and SiaILP show lower performance relative to our method, with a notable gap in FB15k-237 v4, where our method surpasses SiaILP by over 10%.

For the NELL-995 dataset, our method achieves superior performance compared to GraIL and INDIGO in three out of four variations. In NELL-995 v1, our approach demonstrates higher average performance than both GraIL and INDIGO. Similarly, in NELL-995 v3 and NELL-995 v4, our method maintains higher Hits@10 performance, with significant improvements in NELL-995 v4. In NELL-995 v2, the Hits@10 difference between our method and both GraIL and INDIGO is minimal. Ingram and SiaILP lag behind our method on NELL-995, with SiaILP performing moderately well but still trailing in NELL-995 v3 and NELL-995 v4. Ingram, however, consistently underperforms, highlighting its limitations in handling complex scenarios.

In the WN18RR dataset, our proposed method consistently achieves better performance compared to both GraIL and INDIGO across all dataset variants. For example, in WN18RR v1, our method exhibits a notable increase in performance over GraIL while maintaining comparable results to INDIGO. Additionally, in WN18RR v3, the improvement in Hits@10 is evident. For the remaining variants, our method shows significant enhancements in WN18RR v2 and maintains comparable performance in WN18RR v4 relative to both GraIL and INDIGO. Ingram performs significantly worse across all variants, reflecting its limited adaptability in this dataset, while SiaILP achieves competitive results in some cases but falls short overall compared to our method.

In the triple classification task, evaluated using the AUC-PR metric (Table 5), our method consistently outperforms existing baselines across all datasets and their respective versions. On the FB15k-237 dataset, our method achieves the highest AUC-PR scores, particularly in FB15k-237 v4, where it reaches 95.73%, significantly surpassing GraIL and INDIGO. Similarly, on the NELL-995 dataset, our method demonstrates superior performance across all variations, with notable improvements in NELL-995 v4. While SiaILP performs moderately well in some versions, it consistently lags behind our method. Meanwhile, Ingram exhibits significantly lower scores, highlighting its limitations in handling the complexities of these datasets. On the WN18RR dataset, our method maintains strong results across all versions, achieving 93.89% in WN18RR v4, again outperforming GraIL, INDIGO, and other baselines. SiaILP and Ingram exhibit weaker performance across all versions, with SiaILP showing occasional competitiveness but failing to match our method's effectiveness, while Ingram struggles to adapt to the dataset's complexities.

The primary reason for these performance improvements lies in our method's ability to leverage both the structural information of the knowledge graph (KG) for predicting missing links and the semantic context of entities based on their associated relations. As a result, the obtained entity embeddings are more expressive compared to those generated by GraIL and INDIGO, which encode entity embeddings solely based on the KG's topological structure. This advantage is particularly critical for inductive link prediction, where models must generalize to unseen entities and relations. In contrast, Ingram and SiaILP exhibit lower overall performance as their evaluations were conducted using parameter settings recommended in their respective papers, which were not specifically optimized for the datasets used in our experiments. This lack of parameter tuning may have resulted in suboptimal embeddings and degraded performance, particularly in scenarios requiring fine-grained generalization. Due to certain constraints, the current experiments focus on addressing the partially inductive scenario, excluding cases where the enclosing subgraph is empty. Although our proposed method considers this aspect, GraIL and INDIGO do not account for empty subgraphs.

Empty subgraph. The results in Table 6 present the performance of various models on link prediction tasks with empty subgraphs, evaluated using different metrics: mean reciprocal rank (MRR), Hits@1, Hits@5, and Hits@10. The datasets used for evaluation include FB15k-237-emsg, NELL-995-emsg, and WN18RR-emsg. Three models are compared: GraIL, TACT, and our proposed model. Several key observations can be drawn from these results.

First, the proposed model consistently outperforms both GraIL and TACT across all datasets and evaluation metrics. In the FB15k-237-emsg dataset, our method achieves the highest performance in all metrics: MRR (29.91%), Hits@1 (31.21%), Hits@5 (34.90%), and Hits@10 (40.49%). This suggests that the proposed model is particularly effective at ranking correct entities higher and retrieving a greater number of correct entities within the top 10 predictions compared to other models.

For the NELL-995-emsg dataset, the proposed model also demonstrates superior performance, achieving an MRR of 20.30%, Hits@1 of 18.25%, and Hits@10 of 40.11%. Although its Hits@5 score is slightly lower than that of GraIL, the model's overall performance suggests a balanced and effective approach for link prediction in this dataset.

In the WN18RR-emsg dataset, the proposed model's superiority is most pronounced. It achieves an MRR of 49.08%, Hits@1 of 38.11%, Hits@5 of 55.31%, and Hits@10 of 62.45%, significantly outperforming both GraIL and TACT across all metrics. This highlights the robustness of the proposed model in handling more complex datasets with empty subgraphs.

Comparing the other models, while GraIL performs competitively, particularly in Hits@5 for the NELL-995-emsg dataset and Hits@1 for the WN18RR-emsg dataset, it generally falls short compared to the proposed model. TACT, although a strong baseline, underperforms relative to both GraIL and the proposed model across most metrics

Table 7
Link prediction with unseen relations in testing (AUC-PR, MRR, and Hits@10).

Methods	NELL-995.fully.v1			NELL-995.fully.v2			NELL-995.fully.v3		
	AUC-PR	MRR	Hits@10	AUC-PR	MRR	Hits@10	AUC-PR	MRR	Hits@10
TACT	91.47	69.55	88.76	90.02	75.13	91.02	90.21	73.88	**91.33**
RMPI	91.12	71.25	89.14	**91.36**	79.45	**91.24**	**92.83**	**78.55**	88.99
Ours	**91.53**	**72.03**	**89.48**	91.01	**79.51**	90.15	92.10	78.12	88.50

and datasets. These observations suggest that the proposed model has stronger generalization capability and is more effective in link prediction tasks, particularly in scenarios involving empty subgraphs.

Overall, the proposed model's consistently superior performance across different datasets and evaluation metrics highlights its effectiveness and robustness. Its ability to achieve higher scores in MRR, Hits@1, Hits@5, and Hits@10 demonstrates that it not only ranks correct entities higher but also retrieves more correct entities within the top predictions. This makes the proposed model a promising approach for further research and practical applications in knowledge graph completion and link prediction tasks involving empty subgraphs.

Fully inductive link prediction. Table 7 presents the performance comparison of three models, TACT, RMPI, and our proposed method, on link prediction tasks with unseen relations. These models were evaluated across three different versions of the NELL-995 dataset: NELL-995.fully.v1, NELL-995.fully.v2, and NELL-995.fully.v3. The evaluation metrics used to assess performance include AUC-PR, MRR, and Hits@10.

In the NELL-995.fully.v1 dataset, the proposed method outperforms both TACT and RMPI across all metrics. Specifically, TACT achieves an AUC-PR of 91.47%, an MRR of 69.55%, and a Hits@10 of 88.76%. RMPI, while slightly lower in AUC-PR at 91.12%, performs better in MRR (71.25%) and Hits@10 (89.14%) compared to TACT. The proposed method, however, excels by achieving an AUC-PR of 91.53%, an MRR of 72.03%, and a Hits@10 of 89.48%, demonstrating its superior ability to handle unseen relations in this dataset.

For the NELL-995.fully.v2 dataset, the performance trends show slight variations. TACT achieves an AUC-PR of 90.02%, an MRR of 75.13%, and a Hits@10 of 91.02%. RMPI demonstrates stronger performance, with an AUC-PR of 91.36% and a Hits@10 of 91.24%, but a slightly lower MRR at 79.45%. The proposed method maintains a competitive edge, achieving an AUC-PR of 91.01%, an MRR of 79.51%, and a Hits@10 of 90.15%. These results suggest that while RMPI performs slightly better in AUC-PR and Hits@10, the proposed method still demonstrates robust performance, particularly in MRR.

In the NELL-995.fully.v3 dataset, TACT achieves an AUC-PR of 90.21%, an MRR of 73.88%, and a Hits@10 of 91.33%. RMPI excels with the highest AUC-PR of 92.83% and an MRR of 78.55%, although it falls slightly short in Hits@10 at 88.99%. The proposed method also demonstrates strong performance, achieving an AUC-PR of 92.10%, an MRR of 78.12%, and a Hits@10 of 88.50%. These results highlight the efficacy of the proposed method in maintaining high precision and ranking quality, even though RMPI slightly outperforms it in some metrics.

In general, the table highlights the effectiveness and robustness of the proposed method across different versions of the NELL-995 dataset. The proposed method consistently achieves high performance, often outperforming TACT and delivering competitive results compared to RMPI. This underscores its capability in handling link prediction tasks with unseen relations, effectively balancing precision, ranking quality, and recall. The consistently high performance across different datasets further affirms the generalizability of the proposed approach.

Case studies. In Fig. 3, we illustrate three different link prediction scenarios drawn from NELL-995.fully.v1, NELL-995-emsg, and FB15k-237 v4, demonstrating how our proposed method handles unseen relations, empty subgraphs, and general inductive settings. Each example presents a target triple, its 2-hop enclosing subgraph, and the associated schema graph where applicable. We compare the predicted scores of various models, highlighting the advantages of our unified approach in diverse inductive settings.

The first example focuses on the triple *(coach:vince lombardi, coach won trophy, super bowl)*, where the relation *coach won trophy* is unseen during testing. Compared to TACT and RMPI, our model assigns a higher confidence score to this triple, highlighting its superior capability in reasoning with unseen relations. The reasoning process benefits from ontological schemas, by identifying semantic relationships between unseen and seen relations. For example, if *coach won trophy* shares a domain-range relationship with a known relation *team won trophy* through the intermediate concept *award trophy tournament*, this additional semantic knowledge strengthens our model's inference. Additionally, one-hop neighboring relations, such as *team won trophy*, along with two-hop neighboring relations, such as *team plays sport*, help establish logical patterns that support the plausibility of the target relation. As a result, our model outperforms TACT and RMPI by effectively integrating structural information and schema-based reasoning.

The second case demonstrates a scenario where the enclosing subgraph is empty, making traditional GNN-based models ineffective. The target triple *(mammal:rabbits, mammal type, animal:small_mammals)* lacks direct connectivity, posing a challenge for structure-dependent models like GraIL and TACT. Our model, however, effectively compensates for missing structural information by utilizing relation-context encoding, which derives entity embeddings from their associated relations rather than direct connections.

Particularly, relations such as *predator prey* and *animal category* provide indirect contextual information that help determine the missing link. By encoding relation-specific features, our model can infer that *rabbits*, classified under a *general mammal category*, should also be categorized as *small mammals*. The resulting improvement in prediction score highlights the effectiveness of our contrastive learning-based relation-context modeling in cases where enclosing subgraphs are sparse or nonexistent.

The third example involves (/m/0j0k, countries within, /m/0jdd), a case where both entities are unseen. Note that FB-237 dataaset encodes entities information for privacy concerns). This scenario poses a challenge for traditional inductive models, which typically rely on seen entity embeddings. While methods like SiaILP and InGram perform relatively well by learning from seen relation structures, our approach surpasses them in generalization capability.

By combining the triple-view graph transformation and informative neighboring relations surrounding the target triple, our model assigns a more accurate relation embedding to *countries within*. Unlike models such as INDIGO that transforms the whole KG, our approach focuses on selecting and utilizing informative subgraphs for the partially inductive link prediction scenario. It ensures computational efficiency and enables low-complexity predictions, making our method scalable for large-scale knowledge graphs.

6. Discussion

Our proposed method, particularly the URMAS component, is broadly applicable to a variety of KG domains that incorporate ontological schemas or structured taxonomies. This approach is especially well-suited for domains where ontological resources define semantic and hierarchical relationships between entities, enabling the effective handling of unseen relations and entities in evolving KGs. For example,

Predicted Score	Relations Ids	2-hop Enclosing Subgraph	Schema Graph
	Target Triple: (coach:vince lombardi, coach won trophy, super bowl)		
TACT: 0.152 RMPI: 0.29 **Ours: 0.301**	**coach won trophy (213)** award trophy tournament (47) team plays sport (3) team won trophy (16)	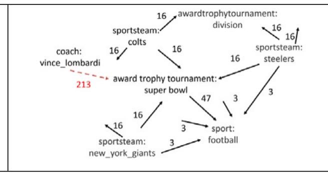	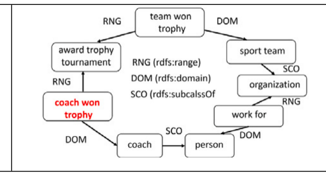
	Target Triple: (mammal:rabbits, mammal type, animal:small_mammals)		
GraIL: 0.145 TACT: 0.209 **Ours: 0.235**	**predator prey (24)** animal category (2) mammal type (48)	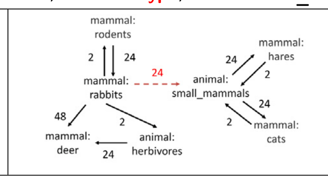	/
	Target Triple: (/m/0j0k, countries within, /m/0jdd)		
GraIL: 0.40 INDIGO: 0.512 InGram: 0.578 SialLP: 0.531 **Ours: 0.658**	**countries within (67)** contains (41) adjoins (94) exported To (164) jurisdiction (101)	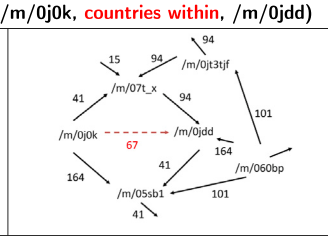	/

Fig. 3. Link prediction use cases with three possible scenarios from NELL-995.fully.v1, NELL-995-emsg, and FB15k-237 v4. Unseen relations are highlighted in red, and the numbers represent the encoded IDs of each relation.

in biomedical and healthcare domains, which rely on ontologies such as SNOMED CT, UMLS, or the Human Disease Ontology, URMAS can predict emerging links related to new diseases, treatments, or symptoms. Similarly, in scientific research (e.g., gene and protein databases) and e-commerce (e.g., product knowledge graphs), where ontological structures such as the Gene Ontology or Google's product schema are available, the method supports link prediction in dynamically growing KGs without requiring retraining on vast datasets.

However, the method has certain limitations. First, it relies on the availability and quality of ontological schemas. Domains that lack structured schemas or have sparse taxonomies may experience reduced performance. In such cases, integrating additional inductive techniques, such as text-based embeddings or rule-based reasoning, could provide complementary support. Second, for domains with extremely large or complex ontologies (e.g., SNOMED CT in healthcare), the computational load for schema-based processing could become significant. However, it remains more efficient than embedding-based models that rely on external text resources. Future work will explore computational optimizations, such as ontology pruning or approximate embeddings, to further reduce this overhead. Finally, our method may face challenges in highly dynamic KGs, where ontological structures change frequently. In such environments, periodic model updates or a hybrid approach incorporating text- or rule-based methods could improve adaptability.

7. Conclusion

In this paper, we address link prediction tasks in both partially inductive and fully inductive settings during the testing stage, reflecting real-world knowledge graph (KG) practices. Our approach comprises three key components: (1) entity-independent modeling using triple-view graph, (2) contrastive learning-based relation-context modeling, and (3) unseen relations modeling with augmented schema. Each component is tailored to tackle a challenging LP scenario, with subsequent components complementing the previous ones to provide a comprehensive solution. Specifically, the first component handles the partially inductive scenario when only unseen entities occur during testing, using graph neural networks (GNNs). The second component is designed for scenarios where the enclosing subgraph of the target triple is empty. The third component addresses unseen relations by leveraging additional resources, such as the KG's ontological schema. Future work may include conducting more ablation experiments and extending the methodology to different domains, such as healthcare.

CRediT authorship contribution statement

Tri D.T. Nguyen: Writing – original draft. **Ubaid Ur Rehman:** Writing – review & editing. **Musarrat Hussain:** Writing – review & editing. **Rao Faizan:** Writing – review & editing. **Jamil Hussain:** Writing – review & editing. **Sung-Ho Bae:** Writing – review & editing. **Jung Uk Kim:** Writing – review & editing. **Seong Tae Kim:** Writing – review & editing. **Sungyoung Lee:** Supervision, Funding acquisition.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Tri D.T. Nguyen, Seong Tae Kim, and Sungyoung Lee reports financial support was provided by National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2023-00252391). Tri D.T. Nguyen, Seong Tae Kim, and Sungyoung Lee reports financial support was provided by MSIT(Ministry of Science and ICT), Korea, under the Grand Information Technology Research Center support program(IITP-2024-2020-0-01489). Tri D.T. Nguyen, Seong Tae Kim, and Sungyoung Lee reports financial support was provided by ITRC(Information Technology Research Center) support program(RS-2023-00259004) supervised by the IITP(Institute for Information & communications Technology Planning & Evaluation). Tri D.T. Nguyen, Seong Tae Kim, and Sungyoung Lee reports financial support was provided by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (IITP-2022-0-00078, Explainable Logical Reasoning for Medical Knowledge Generation). If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

Balažević, I., Allen, C., & Hospedales, T. M. (2019). Tucker: Tensor factorization for knowledge graph completion. ArXiv preprint arXiv:1901.09590.

Becker, C., & Bizer, C. (2008). DBpedia mobile: A location-enabled linked data browser. *Ldow*, *369*, 2008.

Bollacker, K., Evans, C., Paritosh, P., Sturge, T., & Taylor, J. (2008). Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on management of data* (pp. 1247–1250).

Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., & Yakhnenko, O. (2013). Translating embeddings for modeling multi-relational data. *Advances in Neural Information Processing Systems*, *26*.

Chen, J., He, H., Wu, F., & Wang, J. (2021). Topology-aware correlations between relations for inductive link prediction in knowledge graphs. In *Proceedings of the AAAI conference on artificial intelligence (Vol. 35, No. 7)* (pp. 6271–6278).

Chen, M., Zhang, W., Yao, Z., Chen, X., Ding, M., Huang, F., & Chen, H. (2022). Meta-learning based knowledge extrapolation for knowledge graphs in the federated setting. ArXiv preprint arXiv:2205.04692.

Dettmers, T., Minervini, P., Stenetorp, P., & Riedel, S. (2018). Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence (Vol. 32, No. 1)*.

Färber, M., Bartscherer, F., Menne, C., & Rettinger, A. (2018). Linked data quality of dbpedia, freebase, opencyc, wikidata, and yago. *Semantic Web*, *9*(1), 77–129.

Geng, Y., Chen, J., Chen, Z., Pan, J. Z., Ye, Z., Yuan, Z., Jia, Y., & Chen, H. (2021). Ontozsl: Ontology-enhanced zero-shot learning. In *Proceedings of the web conference 2021* (pp. 3325–3336).

Geng, Y., Chen, J., Pan, J. Z., Chen, M., Jiang, S., Zhang, W., & Chen, H. (2023). Relational message passing for fully inductive knowledge graph completion. In *2023 IEEE 39th international conference on data engineering* (pp. 1221–1233). IEEE.

Geng, Y., Chen, J., Zhuang, X., Chen, Z., Pan, J. Z., Li, J., Yuan, Z., & Chen, H. (2023). Benchmarking knowledge-driven zero-shot learning. *Journal of Web Semantics*, *75*, Article 100757.

Guo, S., Wang, Q., Wang, L., Wang, B., & Guo, L. (2018). Knowledge graph embedding with iterative guidance from soft rules. In *Proceedings of the AAAI conference on artificial intelligence (Vol. 32, No. 1)*.

Lee, J., Chung, C., & Whang, J. J. (2023). InGram: Inductive knowledge graph embedding via relation graphs. In *International conference on machine learning* (pp. 18796–18809). PMLR.

Li, J., Zhang, T., Tian, H., Jin, S., Fardad, M., & Zafarani, R. (2020). Sgcn: A graph sparsifier based on graph convolutional networks. In *Advances in knowledge discovery and data mining: 24th Pacific-Asia conference, PAKDD 2020, Singapore, May 11–14, 2020, proceedings, Part I 24* (pp. 275–287). Springer.

Lin, Q., Liu, J., Xu, F., Pan, Y., Zhu, Y., Zhang, L., & Zhao, T. (2022). Incorporating context graph with logical reasoning for inductive relation prediction. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval* (pp. 893–903).

Liu, S., Grau, B., Horrocks, I., & Kostylev, E. (2021). Indigo: Gnn-based inductive knowledge graph completion using pair-wise encoding. *Advances in Neural Information Processing Systems*, *34*, 2034–2045.

Meilicke, C., Fink, M., Wang, Y., Ruffinelli, D., Gemulla, R., & Stuckenschmidt, H. (2018). Fine-grained evaluation of rule-and embedding-based systems for knowledge graph completion. In *The semantic web–ISWC 2018: 17th international semantic web conference, Monterey, CA, USA, October 8–12, 2018, proceedings, Part I 17* (pp. 3–20). Springer.

Omran, P. G., Wang, K., & Wang, Z. (2019). An embedding-based approach to rule learning in knowledge graphs. *IEEE Transactions on Knowledge and Data Engineering*, *33*(4), 1348–1359.

Schlichtkrull, M., Kipf, T. N., Bloem, P., Van Den Berg, R., Titov, I., & Welling, M. (2018). Modeling relational data with graph convolutional networks. In *The semantic web: 15th international conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, proceedings 15* (pp. 593–607). Springer.

Si, J., Xie, C., Zhou, J., Yu, S., Chen, L., Xuan, Q., & Miao, C. (2024). Inductive subgraph embedding for link prediction. *Mobile Networks and Applications*, 1–12.

Sun, Z., Deng, Z.-H., Nie, J.-Y., & Tang, J. (2019). Rotate: Knowledge graph embedding by relational rotation in complex space. ArXiv preprint arXiv:1902.10197.

Sun, M., Xing, J., Wang, H., Chen, B., & Zhou, J. (2021). Mocl: data-driven molecular fingerprint via knowledge-aware contrastive learning from molecular graph. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining* (pp. 3585–3594).

Teru, K., Denis, E., & Hamilton, W. (2020). Inductive relation prediction by subgraph reasoning. In *International conference on machine learning* (pp. 9448–9457). PMLR.

Toutanova, K., Chen, D., Pantel, P., Poon, H., Choudhury, P., & Gamon, M. (2015). Representing text for joint embedding of text and knowledge bases. In *Proceedings of the 2015 conference on empirical methods in natural language processing* (pp. 1499–1509).

Vashishth, S., Sanyal, S., Nitin, V., & Talukdar, P. (2019). Composition-based multi-relational graph convolutional networks. ArXiv preprint arXiv:1911.03082.

Wang, J., Chen, H., Lv, Q., Shi, Z., Chen, J., He, H., Xie, H., Lian, D., Chen, E., & Wu, F. (2023). Learning complete topology-aware correlations between relations for inductive link prediction. ArXiv preprint arXiv:2309.11528.

Wang, Q., Mao, Z., Wang, B., & Guo, L. (2017). Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, *29*(12), 2724–2743.

Wang, B., Shen, T., Long, G., Zhou, T., Wang, Y., & Chang, Y. (2021). Structure-augmented text representation learning for efficient knowledge graph completion. In *Proceedings of the web conference 2021* (pp. 1737–1748).

Xiong, W., Hoang, T., & Wang, W. Y. (2017). Deeppath: A reinforcement learning method for knowledge graph reasoning. ArXiv preprint arXiv:1707.06690.

Zhang, C., & Liu, X. (2024). Inductive link prediction in knowledge graphs using path-based neural networks. In *2024 international joint conference on neural networks* (pp. 1–9). IEEE.

Zhang, Y., Wang, W., Yin, H., Zhao, P., Chen, W., & Zhao, L. (2023). Disconnected emerging knowledge graph oriented inductive link prediction. In *2023 IEEE 39th international conference on data engineering* (pp. 381–393). IEEE.