ASGEA: Exploiting Logic Rules from Align-Subgraphs for Entity Alignment

Anonymous ACL submission

Abstract

Entity alignment (EA) aims to identify entities across different knowledge graphs that represent the same real-world objects. Recent embedding-based EA methods have achieved state-of-the-art performance in EA yet faced interpretability challenges as they purely rely on the embedding distance and neglect the logic rules behind a pair of aligned entities. In this paper, we propose the Align-Subgraph Entity Alignment (ASGEA) framework to exploit logic rules from Align-Subgraphs. ASGEA uses anchor links as bridges to construct Align-Subgraphs and spreads along the paths across KGs, which distinguishes it from the embedding-based methods. Furthermore, we design an interpretable Path-based Graph Neural Network, ASGNN, to effectively identify and integrate the logic rules across KGs. We also introduce a node-level multi-modal attention mechanism coupled with multi-modal enriched anchors to augment the Align-Subgraph. Our experimental results demonstrate the superior performance of ASGEA over the existing embedding-based methods in both EA and Multi-Modal EA (MMEA) tasks. Our code is available at https://anonymous.4open.science/r/ASGEA-B2FE.

1 Introduction

Knowledge Graphs (KGs) provide structured knowledge from real-world facts but suffer from inherent incompleteness. Entity Alignment (EA) addresses this problem by aligning entities across various KGs, facilitating knowledge integration.

Recent advancements in EA have pivoted towards Embedding-Based EA methods (Sun et al., 2018; Cao et al., 2019), utilizing vector embeddings of entities and relations to assess similarity. Notably, approaches employing Graph Neural Networks (GNNs) (Liu et al., 2020; Gao et al., 2022) for encoding focus on graph structural representation by neighbor aggregation, positing that entities with analogous neighborhoods are more likely aligned. However, these methods struggle in scenarios with similar neighborhood structures but different alignment relevance, as shown in Figure 1. This reveals a critical limitation: the uniform representation across entities inadequately distinguishes between relevant and irrelevant neighbors for precise alignment, pointing to the inefficiency of current methods in detailed entity distinction.

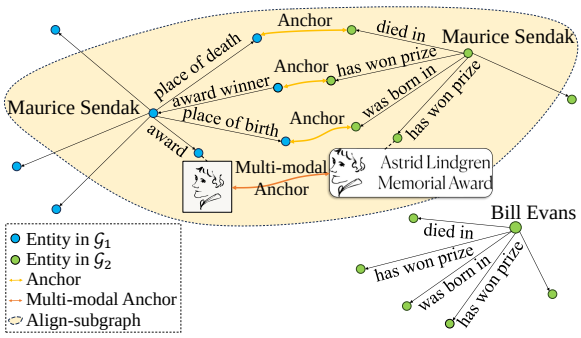

Figure 1: ***(i)*** The entities "*Maurice Sendak*" in $\mathcal{G}_1$ and "*Bill Evans*" in $\mathcal{G}_2$ exhibit structurally similar neighbor relationships but they are not candidates for alignment. ***(ii)*** Anchors specifically link "*Maurice Sendak*" instances across the KGs without connecting "*Bill Evans*" in $\mathcal{G}_2$ to "*Maurice Sendak*" in $\mathcal{G}_1$. ***(iii)*** Our ASGEA is employed to filter out non-relevant neighbors, retaining only essential alignment information .

Intuitively, considering whether corresponding neighbors are aligned significantly enhances the model's ability to distinguish entities with similar neighborhoods. This is grounded in the fact that if two entities are aligned, their neighbors connected by the same relations are likely aligned as well. By treating alignment seeds as anchor links, such information can be integrated into the model.

Upon further analysis, we find that this essentially establishes a simple logic rule, framed by two identical relations surrounding an alignment relation. Notably, if the spouses of two entities are aligned (the same person), then the entities themselves are also aligned (the same person). However, existing methods overlook the cross-graph logic

rules (i.e., alignment rules) hidden behind EA.

To overcome the above challenges, we propose a novel **A**lign-**S**ub**g**raph **E**ntity **A**lignment (ASGEA) framework. For alignment rule mining, we design an Align-Subgraph (ASG) extraction algorithm to get specific ASGs for each entity pair which incorporates the entirety of potential alignment rules between them to filter out non-relevant neighbor information. We further design an **A**lign-**S**ubgraph **G**raph **N**eural **N**etwork (ASGNN) which leverages an interpretable attention mechanism to prioritize edges along significant paths and employs a unidirectional Path-based message passing strategy to preserve path-specific information. This ensures the generation of entity-specific representations that focus solely on pivotal alignment details. To accommodate multi-modal EA (MMEA) tasks, we further innovate a node-level multi-modal attention mechanism. This enhancement is crucial, especially considering the real-world scarcity of anchor links which complicates the acquisition of extensive potential paths for ASGs. By creating new multi-modal anchor links and integrating auxiliary anchor data, we facilitate the extraction of ASGs. Overall, our contributions can be summarized as:

- We propose a ASG extraction algorithm to identify subgraphs encapsulate all possible alignment rule paths, effectively reducing noise from non-relevant neighbors during the alignment process. Additionally, our ASGNN, an interpretable Path-based GNN for rule mining in ASGs, leverages dynamic programming to improve efficiency in message passing.
- Our model is expanded through a node-level unified multi-modal attention mechanism, enhancing its utility for MMEA tasks.
- Our ASGEA framework achieves SOTA performance in both EA and MMEA tasks.

2 Related work

2.1 Entity Alignment

EA integrates KGs by matching entities across graphs, leveraging embedding-based methods to overcome heterogeneity (Sun et al., 2023). These methods fall into two categories: translation-based approaches, employing techniques like TransE (Bordes et al., 2013) to capture structural information within triples (Chen et al., 2017; Zhang et al., 2019), and GNN-based approaches, using networks such as GCN (Kipf and Welling, 2017) to aggregate neighborhood features for richer entity representations (Cao et al., 2019; Shi and Xiao, 2019). Enhancements include shared parameter embeddings for seed alignment consistency (Zhu et al., 2017) and iterative learning (Sun et al., 2018). PEEA's positional encoding introduces global relational perspectives, but it remains reliant on embedding techniques (Tang et al., 2023).

2.2 Multi-modal Entity Alignment

In the emerging field of MMEA, researchers address the complex nature of real-world KGs by integrating multiple modalities (Zhu et al., 2024; Chen et al., 2024). Liu et al. (2021) introduce attention mechanisms to weigh modalities differently, while Chen et al. (2022) and Lin et al. (2022) focus on integrating visual cues and enhancing intra-modal learning, respectively. Chen et al. (2023a) adopts a transformer-based multi-modal approach, and Li et al. (2023a) targets the contextual gap by ensuring attribute consistency across modalities.

Recent rule-based EA methods (Leone et al., 2022; Li et al., 2023b) primarily build upon the PARIS (Suchanek et al., 2011), a probabilistic approach that leverages relationship functionality. However, the lack of public access to code or the use of differing metrics hampers direct comparisons with ASGEA. To the best of our knowledge, we are the first to employ logic rules for EA.

2.3 Rule-based Reasoning

Knowledge graph reasoning leverages entity and relationship analysis to infer new knowledge(Liang et al., 2022). Various methods have been developed for reasoning through logic rule mining. AIME (Galárraga et al., 2013) initiates the exploration with inductive logic programming for rule discovery in KGs. Neural LP (Yang et al., 2017) and DRUM (Sadeghian et al., 2019) later apply dynamic programming to create linear combinations of logic rules. AnyBURL (Meilicke et al., 2019) introduces a bottom-up approach for efficient rule learning in large graphs. RED-GNN (Zhang and Yao, 2022) adopts a Bellman-Ford-like method for learning path representations. Rlogic (Cheng et al., 2022), employs a deductive reasoning-based scoring model for quality evaluation of rules derived from sampled paths. However, these methods mainly learn logic rules within a single KG, not addressing EA across multiple KGs.

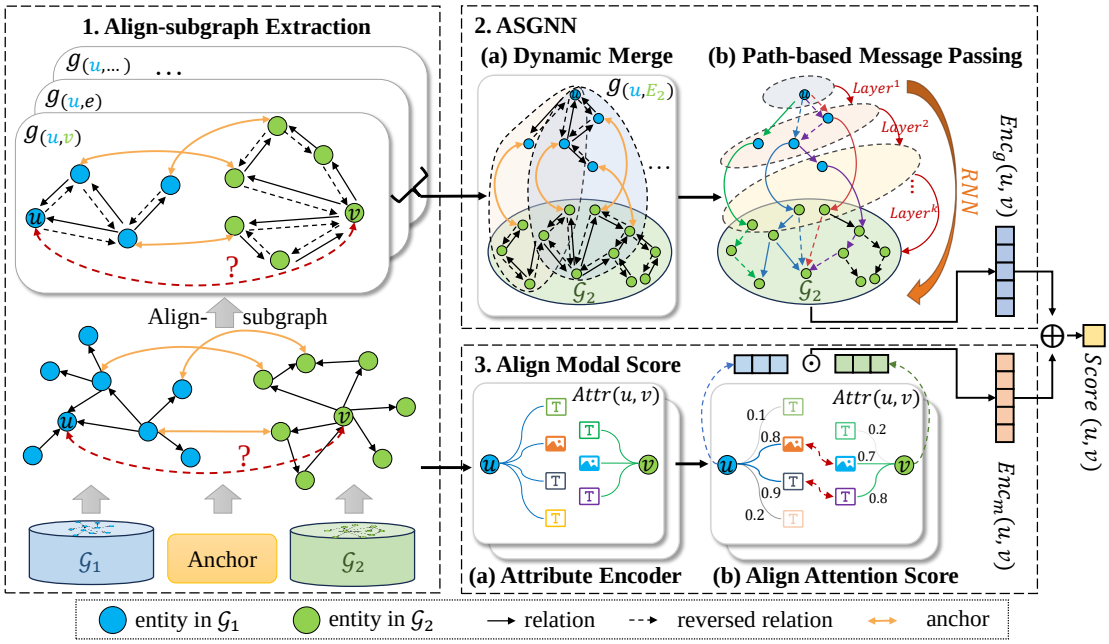

Figure 2: The overall framework of AsgEa. On the left, the ASG Extraction is utilized for constructing entity-specific Align-Subgraphs during training and testing. On the right, the AsGnn is designed for mining alignment rules and scoring, while the Align modal score is employed for matching aligned attributes to subsequently obtain node-level multi-modal scores.

3 Methodology

3.1 Preliminaries

Following Chen et al. (2023b,a), we define a MMKG as $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{A}, \mathcal{V}, \mathcal{T}\}$, where $\mathcal{E}$, $\mathcal{R}$, $\mathcal{A}$, and $\mathcal{V}$ denote the sets of entities, relations, attributes, and attribute values, respectively. The set $\mathcal{T}$ is divided into $\mathcal{T}_{\mathcal{R}}$ and $\mathcal{T}_{\mathcal{A}}$, where $\mathcal{T}_{\mathcal{R}} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ is the set of relation triples, and $\mathcal{T}_{\mathcal{A}} \subseteq \mathcal{E} \times \mathcal{A} \times \mathcal{V}$ is the set of attribute triples. For traditional KG, the sets $\mathcal{A}$, $\mathcal{V}$, and $\mathcal{T}_{\mathcal{A}}$ are empty. Given two MMKGs $\mathcal{G}_1$ and $\mathcal{G}_2$, the goal of MMEA is to identify pairs of entities (e_i^1, e_i^2), where $e_i^1 \in \mathcal{E}_1$,$e_i^2 \in \mathcal{E}_2$ that represent the same real-world entity e_i. A set of aligned entity pairs is provided, which is partitioned into a training subset (i.e., seed alignments $\mathcal{S}$) and a testing subset $\mathcal{S}_{te}$, according to a specified seed alignment ratio (R_{sa}).

Definition 1 ***Anchor Link.*** *For $e_u \in \mathcal{E}_1$ and $e_v \in \mathcal{E}_2$, an anchor link between them is denoted as $A(e_u, e_v)$, which means e_u is aligned with e_v.*

3.2 Align-Subgraph Extraction

3.2.1 Alignment Rules

Alignment rules are logical constructs that identify aligned entities across KGs using structural and semantic information. We introduce three rules:

One-hop Rule: As shown in Figure 3a, for entities e_u^1 and e_v^2 with neighbors e_x^1 and e_y^2 in KGs $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively, if e_x^1 and e_y^2 are aligned ($A(e_x^1, e_y^2)$) and the relations $r^1(e_u^1, e_x^1)$ and $r^2(e_v^2, e_y^2)$ are identical, then e_u^1 and e_v^2 are aligned. This is summarized as $A(e_u^1, e_v^2) \leftarrow r^1(e_u^1, e_x^1) \wedge r^2(e_v^2, e_y^2) \wedge r^1 = r^2 \wedge A(e_x^1, e_y^2)$.

Symmetric k-hop Rule: As shown in Figure 3b, e_u^1 and e_v^2 are connected to their respective k-hop neighbors e_x^1 and e_y^2 through identical sequences of relations. e_u^1 and e_v^2 are aligned if e_x^1 and e_y^2 are aligned, which is expressed as $A(e_u^1, e_v^2) \leftarrow r_1^1(e_u^1, e_{i_1}^1) \wedge \ldots \wedge r_k^1(e_{i_{k-1}}^1, e_x^1) \wedge r_1^2(e_v^2, e_{j_1}^2) \wedge \ldots \wedge r_k^2(e_{j_{k-1}}^2, e_y^2) \wedge r_1^1 = r_1^2 \wedge \ldots \wedge r_k^1 = r_k^2 \wedge A(e_x^1, e_y^2)$.

Asymmetric k-hop Rule: In Figure 3c, e_u^1 and e_v^2 connected to e_x^1 and e_y^2 through different numbers of hops. If the composite relations R^1 and R^2 between e_u^1 and e_x^1, and e_v^2 and e_y^2 are semantically equivalent and e_x^1, e_y^2 are aligned, then e_u^1 and e_v^2 are aligned, which is expressed as $A(e_u^1, e_v^2) \leftarrow R^1(e_u^1, e_x^1) \wedge R^2(e_v^2, e_y^2) \wedge R^1 = R^2 \wedge A(e_x^1, e_y^2)$.

3.2.2 Alignment Path

Expanding on the Alignment Rule analysis, we integrate anchor link relations (r_{anchor}) to connect the two KGs. We categorize anchor links into: ***(i)*** **Anchor**: Utilizes established alignment seeds as anchors. Training strategies are detailed in § 3.5.1. ***(ii)*** **Multi-modal Anchor**: Forms connections between nodes with similar multi-modal attributes.

We reverse relations within KGs (remarked as

r') for path-finding. The alignment rule in § 3.2.1 is rewritten as $R^1(e_u^1, e_x^1) \wedge r_{anchor}(e_x^1, e_y^2) \wedge R^{2'}(e_y^2, e_v^2) \rightarrow A(e_u^1, e_v^2)$. Based on this rule, we define the Alignment Path below:

Definition 2 ***Alignment Path.*** *An alignment path $P(e_u^1, e_v^2)$ with length L is a sequence of relations $e_u^1 \xrightarrow{r_1^1} \dots \xrightarrow{r_{k_1}^1} e_x^1 \xrightarrow{r_{anchor}} e_y^2 \xrightarrow{r_{k_2}^{2\,\prime}} \dots \xrightarrow{r_1^{2\prime}} e_v^2$ from e_u^1 to e_v^2 that follows alignment rules.*

3.2.3 Align-Subgraph

Considering that a solitary alignment path might not suffice for definitive entity alignment, it's crucial to aggregate all potential alignment paths between entities e_u and e_v. Thus, we define an Align-Subgraph (ASG) for a pair (e_u, e_v).

Definition 3 ***Align-Subgraph.*** *For entity pair (e_u, e_v), Align-Subgraph $g(e_u, e_v)$ is the subgraph comprised of all alignment paths between e_u and e_v, formally denoted as $g(e_u, e_v) = \bigcup P(e_u, e_v)$.*

Definition 4 ***K-hop Align-Subgraph.*** *K-hop Align-Subgraph is an Align-Subgraph formed by integrating all alignment paths with length $\leq K$, formally denoted as $g^K(e_u, e_v) = \bigcup_{i \leq K} P^i(e_u, e_v)$.*

We designed an ASG extraction algorithm to obtain the k-hop ASG for the entity pair (e_u, e_v), see Appendix for more details.

3.3 ASGNN

3.3.1 Dynamic Merge

To address the computational challenge of getting alignment scores between an entity e_u from $\mathcal{G}_1$ and all entities in $\mathcal{G}_2$, we employ a dynamic merge strategy, inspired by Zhang and Yao (2022).

Our method facilitates unidirectional message passing from e_u towards entities in $\mathcal{G}_2$, including specific instances like e_v and $e_{v'}$ as shown in Figure 4. When e_v and $e_{v'}$ share a common ancestor d, we utilize dynamic programming to share d's representation, significantly reducing computational efforts by merging convergent alignment paths.

To evaluate e_u's alignment with all entities in $\mathcal{G}_2$, we construct a merged ASG $g(e_u, \mathcal{E}_2)$ encapsulating all paths from e_u to $\mathcal{E}_2$, formalized as:

$$g(e_u, \mathcal{E}_2) = \bigcup\nolimits_{e_v \in \mathcal{E}_2} g(e_u, e_v). \quad (1)$$

3.3.2 Path-based Message Passing

We introduce a Path-based Message Passing technique within an Align-Subgraph Graph Neural Network (ASGNN) to facilitate directed information flow from entity e_u to $\mathcal{E}_2$. In the merged ASG $g^L(e_u, \mathcal{E}_2)$, we ensure each node, apart from e_u, is involved in at least one alignment path, thus incorporating incoming edge information effectively. The computation of the path representation from e_u to any entity e_i is defined as follows:

$$h_{p(e_u,e_i)}^l = \sum\nolimits_{e_j \in N(e_i)} \alpha_{e_j,e_i}^{e_u,l} (h_{e_u,e_j}^{l-1} + r_{e_j,e_i}^l), \quad (2)$$

where $e_j \in N(e_i)$ denotes the neighbors with edges pointing to entity e_i, and r_{e_j,e_i}^l represents the relation representation between e_j and e_i. $\alpha_{e_j,e_i}^{e_u,l}$ is determined via an interpretable attention mechanism, using an MLP to prioritize the most significant alignment paths:

$$\alpha_{e_j,e_i}^{e_u,l} = \text{MLP}_\alpha^l (h_{e_u,e_j}^{l-1} \,||\, r_{e_j,e_i}^l \,||\, f_g(e_j, e_i)), \quad (3)$$

where $f_g(e_j, e_i)$ incorporates graph knowledge, differentiating edges within or across KGs to accommodate their heterogeneity, calculated as:

$$f_g(e_j, e_i) = \text{MLP}_{kg}(\text{EMB}(e_j) \,||\, \text{EMB}(e_i)), \quad (4)$$

where EMB is a learnable embedding function.We employ an RNN for information updates to preserve crucial path information:

$$h_{e_u,e_i}^l = \text{RNN}(h_{e_u,e_i}^{l-1}, W^l h_{p(e_u,e_i)}^l). \quad (5)$$

After L layers of updates on $g^L(e_u, \mathcal{E}_2)$, we obtain the final representations h_{e_u,e_v}^L for all $e_v \in \mathcal{E}_2$ which encapsulate critical alignment path information for pair (e_u, e_v). We employ a MLP for scoring, defined as $\text{score}_g(e_u, e_v) = \text{MLP}_g(h_{e_u,e_v}^L)$.

3.4 Align Modal Score

3.4.1 Attribute Encoder

To integrate text and image attributes from MMKGs into a unified representation space, we utilize modality-specific fully connected layers, $W_\mathcal{M} \in \mathbb{R}^{d_\mathcal{M} \times d}$, to project features into a unified space via $h_\mathcal{A}(v) = W_\mathcal{M} h_\mathcal{A}^0(v)$. Initial embeddings $h_\mathcal{A}^0(v)$ are derived from BERT for text and a pre-trained model for images.

3.4.2 Align Attention Score

To mitigate the noise from varying attribute types and quantities in EA, we introduce an align attention mechanism, which prioritizes attributes shared between entities e_u and e_v, diminishing the influence of mismatched attributes. We calculate pair-aware multi-modal representations $h_u^{e_u,e_v}$ as:

$$h_u^{e_u,e_v} = \sum_{e_u, a_i, v_i \in \mathcal{T}_\mathcal{A}(e_u)} \alpha_{a_i}^{e_u,e_v} h_\mathcal{A}(v_i). \quad (6)$$

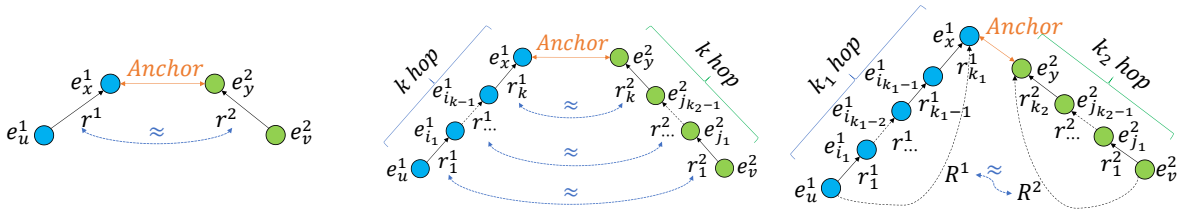

(a) One-hop rule (b) Symmetric k-hop rule (c) Asymmetric k-hop rule

Figure 3: Illustration of different alignment rules.

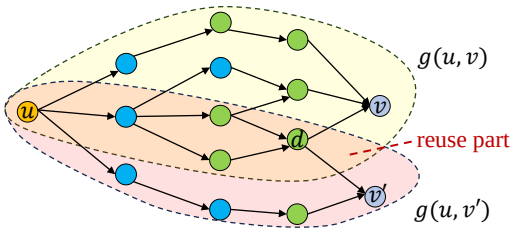

Figure 4: An example for dynamic merge. The yellow portion denotes ASG $g(e_u, e_v)$ while the red portion denotes ASG $g(e_u, e'_v)$. The overlapping area signifies the parts that can be reused through dynamic programming.

Attention scores $\alpha_{a_i}^{e_u,e_v}$ amplify the weights of more relevant attributes, calculated using a softmax function over the similarity between attribute types:

$$\alpha_{a_i}^{e_u,e_v} = \sigma\left(\sum\nolimits_{e_v,a_j,v_j \in \mathcal{T}_{\mathcal{A}}(e_v)} \text{sim}(a_i, a_j)\right). \quad (7)$$

Finally, the multi-modal attribute alignment score, which quantifies the similarity between e_u and e_v, is determined by processing their combined multi-modal representations through an MLP: $\text{score}_m(e_u, e_v) = \text{MLP}_m(h_u^{e_u,e_v} \odot h_v^{e_v,e_u})$.

3.5 Training Objective

3.5.1 Training Strategy

During the test phase, we utilize all alignment seeds as anchor links to extract Align-Subgraphs for evaluation. To emulate this process in training, we randomly split the alignment seeds S into two subsets: S_{anchor} (75%) and S_{train} (25%). S_{anchor} acts as known anchor links to extract Align-Subgraphs for predicting the remaining S_{train} subset.

3.5.2 Loss Function

We present two variants: AsGEa-Stru, focusing on graph structure, and AsGEa-MM, incorporating multimodal data. AsGEa-Stru only uses $\text{score}_g(e_u, e_v)$ from AsGNN for alignment scores, while AsGEa-MM adds a multimodal score, yielding $s(e_u, e_v) = \text{score}_g(e_u, e_v) + \text{score}_m(e_u, e_v)$.

We treat e_v, the entity aligned with e_u, as a positive example against all other entities as negatives. For each pair of S_{train}, we utilize the following loss for optimization (Lacroix et al., 2018):

$$l(e_u) = \log\left(\sum\nolimits_{e \in \hat{\mathcal{G}}} e^{s(e_u,e)}\right) - s(e_u, e_v), \quad (8)$$

$$\mathcal{L} = \sum\nolimits_{(e_u,e_v) \in S_{train}} (l(e_u) + l(e_v)), \quad (9)$$

where $\hat{\mathcal{G}}$ represents the KG opposite to the one containing e_u. We use stochastic gradient descent (Kingma and Ba, 2015) to minimize (9).

4 Experiment

4.1 Dataset

(i) MMKG (Liu et al., 2019) include two subsets: FB15K-DB15K (FBDB15K) and FB15K-YAGO15K (FBYG15K). These datasets are split into three subsets based on the ratio of seed alignments: $R_{sa} \in \{0.2, 0.5, 0.8\}$.

(ii) DBP15K (Sun et al., 2017) are derived from the multilingual versions of DBpedia and include three subsets: $\text{DBP15K}_{\text{ZH-EN}}$, $\text{DBP15K}_{\text{JA-EN}}$, and $\text{DBP15K}_{\text{FR-EN}}$. We incorporate the multi-modal variant of these datasets (Liu et al., 2021), which enhances the entities with corresponding images.

(iii) Multi-OpenEA (Li et al., 2023c) are the multi-modal variants of the OpenEA benchmarks (Sun et al., 2020b) which enhance entity images through Google search. We use the datasets {EN-FR-15K, EN-DE-15K, D-W-15K-V1, D-W-15K-V2}, following Chen et al. (2023b) to select only the rank one image for each entity.

4.2 Implementation Details

We use pre-trained BERT for text attribute initialization and the dimension of text d_t is 768. Following (Chen et al., 2020; Lin et al., 2022; Chen et al., 2023b), the vision encoders are set to ResNet-152 (He et al., 2016) on DBP15K with the vision dimension $d_v = 2048$, and set to VGG-16 (Simonyan and Zisserman, 2015) on MMKG with

$d_v = 4096$, and set to CLIP (Radford et al., 2021) on Multi-OpenEA with $d_v = 512$.

4.3 Main Results

Across all datasets, surface information is excluded for fair comparison, following Chen et al. (2023b).

4.3.1 MMKG Dataset

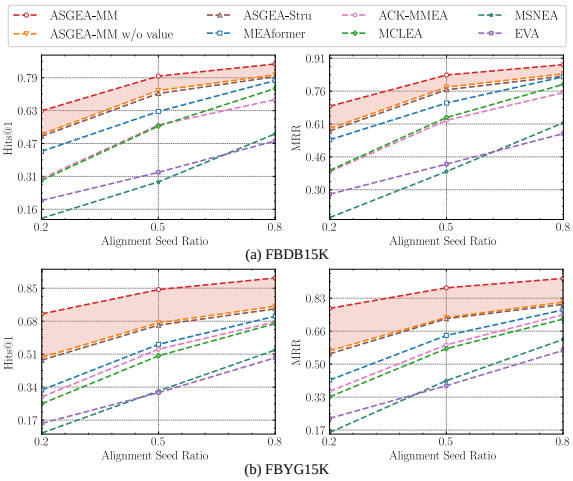

Figure 5: Hits@1 and MRR performance on FBDB15K and FBYG15K with different R_{sa}. **"-MM"** indicates using multi-modal information(abbreviated as MM), while **"-Stru"** indicates reliance on graph structure alone (i.e., triples). "w/o value" indicates without attribute values. For complete results and baseline details, see Appendix.

Results on FBDB15K and FBYG15K datasets are detailed in Figure 5, providing insights as follows: ***(i)*** ASGEA-MM exceeds baselines in all metrics and data proportions, achieving MRRs up to 0.881 on FBDB15K and 0.926 on FBYG15K, demonstrating significant and consistent advantages. ***(ii)*** Delving deeper into the results, we observe that in the H@1 metric, ASGEA-MM exhibits substantial improvements of ↑19.4%, ↑16.9%, and ↑7.9% for the 20%, 50%, and 80% training data splits on FBDB15K (averaging **↑14.7%**) and ↑39.2%, ↑28.2%, and ↑19.7% on the FBYG15K (averaging **↑29.0%**). In terms of the MRR metric, the model shows a steady enhancement with increases of ↑15.5%, ↑12.9%, and ↑5.6% for the 20%, 50%, and 80% training data proportions on FBDB15K (averaging **↑11.3%**), and ↑36.0%, ↑23.9%, and ↑15.8% on FBYG15K (averaging **↑25.2%**). ***(iii)*** ASGEA-Stru, utilizing solely graph structure data and **excluding any auxiliary information like images or texts**, outperforms all existing traditional EA and MMEA baselines. This performance indicates our model's strong capability in mining alignment rules. In terms of architecture, ASGNN functions similarly to a subgraph GNN, which is shown to possess greater expressive power than traditional GNN models (Frasca et al., 2022). ***(iv)*** Previous studies have concentrated on one-hot encoding of attribute types (Chen et al., 2023a), overlooking the significance of attribute values. To ensure a fair comparison with them, we present a model variant "w/o value" that solely relies on attribute types. Our findings demonstrate the robust performance of ASGEA and underscore the importance of incorporating attribute values. ***(v)*** Considering the absence of models that apply rule-based reasoning techniques to EA tasks, we adapt several reasoning methods Neural LP (Yang et al., 2017), AnyBURL (Meilicke et al., 2019) and RED-GNN (Zhang and Yao, 2022) in our framework to tackle EA challenges. The adaptations for the alignment tasks are denoted by the suffix "∗". Notably, they achieve commendable results, underscoring the efficacy of ASGEA.

4.3.2 DBP15K Dataset

Results on three datasets of DBP15K with $R_{sa} = 0.3$ following (Liu et al., 2021) are detailed in Table 1, ASGEA-MM outperforms all other baselines on three datasets in both EA and MMEA settings. Notably, ASGEA-MM achieves up to ↑9.3% improvement in H@1 and ↑6.1% in MRR. Different from MMKG Dataset, these models generally perform better than those without multi-modal information, indicating the importance of multi-modal information in MMEA and the potential low quality of images in the MMKG dataset.

	Models	$DBP15K_{ZH-EN}$		$DBP15K_{JA-EN}$		$DBP15K_{FR-EN}$	
		H@1	MRR	H@1	MRR	H@1	MRR
EA	AlignEA	.472	.581	.448	.563	.481	.599
	KECG	.478	.598	.490	.610	.486	.610
	MUGNN	.494	.611	.501	.621	.495	.621
	AliNet	.539	.628	.549	.645	.552	.657
	ASGEA-Stru	**.560**	**.660**	**.595**	**.690**	**.653**	**.745**
MMEA	MSNEA	.609	.685	.541	.620	.557	.643
	EVA	.683	.762	.669	.752	.686	.771
	MCLEA	.726	.796	.719	.789	.719	.792
	MEAformer	.771	.835	.764	.834	.770	.841
	UMAEA	.800	.860	.801	.862	.818	.877
	ASGEA-MM	**.815**	**.862**	**.849**	**.889**	**.911**	**.938**

Table 1: Results on DPB15K datasets. For complete results and baseline details, see the Appendix.

4.3.3 Multi-OpenEA Dataset

Results on Multi-OpenEA with $R_{sa} = 0.2$ are detailed in Table 2. ASGEA-MM surpasses all baseline models in both EN-FR-15K and EN-DE-15K

dataset. Specifically, ASGEA-MM demonstrates an improvement of ↑6.3% in H@1 and ↑4.1% in MRR for the EN-FR-15K dataset.

Models	OpenEA$_{\text{EN-FR}}$			OpenEA$_{\text{EN-DE}}$		
	H@1	H@10	MRR	H@1	H@10	MRR
MSNEA (Chen et al., 2022)	.692	.813	.734	.753	.895	.804
EVA (Liu et al., 2021)	.785	.932	.836	.922	.983	.945
MCLEA (Lin et al., 2022)	.819	.943	.864	.939	.988	.957
UMAEA (Chen et al., 2023b)	.848	.966	.891	.956	.994	.971
ASGEA-MM	**.911**	**.968**	**.932**	**.980**	**.996**	**.986**

Table 2: Results on Multi-OpenEA datasets. For complete results and baseline details, see the Appendix.

	Models	FBDB15K			FBYG15K		
		H@1	H@10	MRR	H@1	H@10	MRR
20%	**ASGEA-MM**	**.628**	**.799**	**.689**	**.717**	**.848**	**.776**
	w/o AMS	.531	.734	.600	.500	.705	.572
	ASGEA-Stru	.509	.712	.579	.486	.688	.555
	symmetric	.499	.702	.569	.473	.674	.542
50%	**ASGEA-MM**	**.794**	**.899**	**.833**	**.842**	**.920**	**.879**
	w/o AMS	.738	.876	.788	.690	.852	.749
	ASGEA-Stru	.712	.855	.764	.671	.836	.730
	symmetric	.701	.848	.753	.659	.823	.717
80%	**ASGEA-MM**	**.852**	**.927**	**.881**	**.902**	**.965**	**.926**
	w/o AMS	.818	.916	.854	.796	.916	.839
	ASGEA-Stru	.803	.904	.841	.758	.900	.808
	symmetric	.788	.891	.825	.747	.885	.798

Table 3: Variant experiment results on MMKG datasets. **"ASGEA-MM"** indicates our complete model that uses multi-modal information(abbreviated as MM), while **"ASGEA-Stru"** indicates reliance on graph structure alone (i.e., triples), **"w/o AMS"** denotes the exclusion of the AMS module from the complete model, and **"symmetric"** indicates employing only symmetric alignment rule based on the variant "ASGEA-Stru".

4.4 Discussions for Model Variants

To investigate the effectiveness of each part of ASGEA, we conduct variant experiments, as detailed in Table 3. ***(i)*** "ASGEA-Stru" Variant: Focusing solely on graph structure, this variant demonstrates a drop in performance compared to the full ASGEA-MM model yet still surpasses other MMEA methods, indicating the practicality of the ASG extraction algorithm and the capability of ASGNN to effectively identify alignment rules. ***(ii)*** "symmetric" Variant: Incorporating symmetric alignment rules, this variant sees further performance reductions across all metrics compared to the "ASGEA-Stru" variant, suggesting asymmetric rules also play a role in enhancing model performance. ***(iii)*** "W/o AMS" Variant: By averaging multi-modal information without generating node-level attention weights, this variant shows a slight decrease in performance compared to ASGEA-MM. However, it still outperforms the "w/o value" variant as shown in Table 6, indicating the critical importance of attribute values in EA.

This ablation study highlights the vital role of graph structure, while also showing the contributions of asymmetric rules and the Align Modal Score (AMS) module to the overall effectiveness of the ASGEA-MM model.

4.5 Discussions for Model Depth

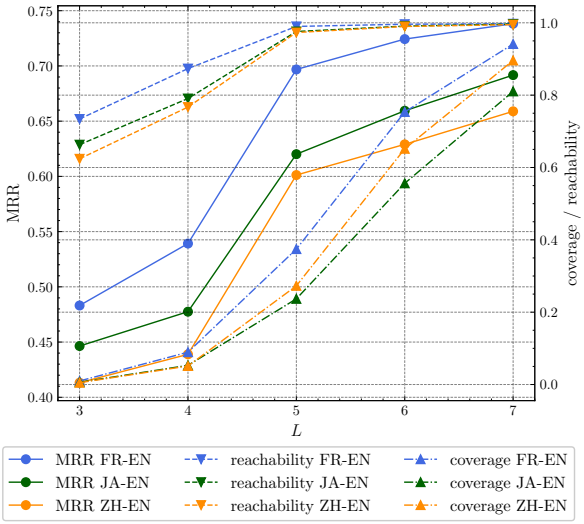

Figure 6: The MRR performance, coverage and reachability on DBP15K with different model depths (L).

We assessed the influence of varying model depths L on performance metrics MRR, coverage (i.e., the percentage of the target KG's nodes we can reach), and reachability (i.e., the proportion of target nodes that are among those reached) across different layers on the DBP15K dataset. The results, depicted in Figure 6, highlight several key insights: ***(i)*** For $L < 5$, performance is generally suboptimal. This is attributed to the model primarily capturing basic 1-hop alignment rules at $L = 3$ and asymmetric 2:1 hop rules at $L = 4$ (i.e., 2 hops in one KG and 1 hop in another), which are insufficient for comprehensive alignment. ***(ii)*** A notable performance boost is observed at $L = 5$, indicating the effectiveness of 2-hop symmetric alignment rules. ***(iii)*** For $L > 5$, the incorporation of more complex alignment rules leads to progressive improvements in model performance, underscoring the efficacy of the alignment rules proposed in § 3.2.1.

The analysis also indicates a limitation in information propagation for aligned nodes beyond L hops, potentially affecting predictive accuracy. However, increasing L enriches the informational

hop=3	A(X,Y) ← directed_by(A,X) ∧ A(A,B) ∧ director(B,Y)
	A(X,Y) ← capital(A,X) ∧ A(A,B) ∧ capital(B,Y)
	A(X,Y) ← children(X,A) ∧ A(A,B) ∧ parent(B,Y)
hop=4	A(X,Y) ← children(A,X) ∧ A(A,B) ∧ spouse(C,B) ∧ child(C,Y)
	A(X,Y) ← geographic_scope(A,X) ∧ A(A,B) ∧ headquarter(B,C) ∧ country(C,Y)
	A(X,Y) ← country(A,X) ∧ citytown(B,A) ∧ A(B,C) ∧ country(C,Y)
hop=5	A(X,Y) ← capital(A,X) ∧ administrative_parent(B,A) ∧ A(B,C) ∧ state(C,D) ∧ capital(D,Y)
	A(X,Y) ← currency_used(A,X) ∧ country(B,A) ∧ A(B,C) ∧ country(C,D) ∧ currency(D,Y)
	A(X,Y) ← children(X,A) ∧ children(B,A) ∧ A(B,C) ∧ parent(D,C) ∧ parent(D,Y)

Table 4: Alignment rules learned from FBDB15K.

content available for analysis, enhancing both the coverage and reachability of the alignment subgraph. Specifically, this is demonstrated when from $L = 4$ to $L = 5$, where the pronounced improvements in coverage and reachability vividly illustrate the model's enhanced performance due to the enriched information base. What's more, we find that when $L = 3$, despite a mere 0.05% coverage, it achieves a reachability of 60%. This indicates the algorithm's ability to identify and focus on potentially aligned entities while filtering out non-relevant information.

4.6 Visualization of interpretable

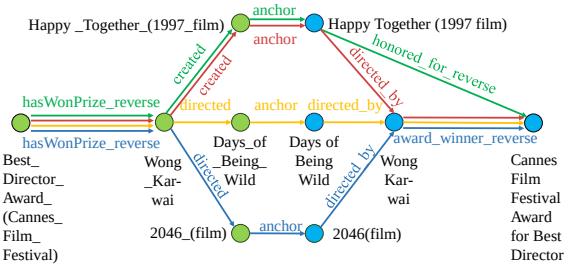

Figure 7: Visualization of learned Align-Subgraph

We can extract the alignment paths from the attention scores in (3) in ASGEA for interpretation. We prune edges below a threshold to highlight identified rules. Figure 7 shows the alignment paths learned by ASGEA for a test sample. Given a pair of entities to align ("*Best_Director_Award_(Cannes_Film_Festival)*", "*Cannes Film Festival Award for Best Director*"), ASGEA identified four alignment paths, marked in different colors, which encompass three alignment rules. Notably, two paths (in yellow and blue) adhere to the same rule, represented as $A(e_u^1, e_v^2) \leftarrow \textit{hasWonPrize}(e_i^1, e_u^1) \wedge \textit{directed}(e_i^1, e_x^1) \wedge \textit{award_winner}(e_v^2, e_j^2) \wedge \textit{directed_by}(e_y^2, e_j^2) \wedge \textit{directed} = \textit{directed_by} \wedge \textit{hasWonPrize} = \textit{award_winner} \wedge A(e_x^1, e_y^2)$. This indicates that through the "*directed*" and "*directed_by*" relations, as well as the alignment of "*Days of Being Wild*" with "*Days_of_Being_Wild*", it can be inferred that "*Wong_Kar-wai*" and *Wong Kar-wai*" are the same entity. Further analysis using "*hasWonPrize*" and "*award_winner*" relations concludes that the award is the same. See the Appendix for more visualization results. Additionally, Table 4 enumerates some of the learned alignment rules across various hops: 3-hop rules is the One-hop rule detailed in § 3.2.1. The 4-hop corresponds to an asymmetric rule where $k_1, k_2 = 1, 2$ (e.g., $A(e_x, e_y) \leftarrow \textit{children}(e_a, e_x) \wedge A(e_a, e_b) \wedge \textit{spouse}(e_c, e_b) \wedge \textit{child}(e_c, e_y)$), whereas the 5-hop encompasses a symmetric 2-hop rule, along with an asymmetric rule where $k_1, k_2 = 1, 3$ (e.g., $A(e_x, e_y) \leftarrow \textit{state_province_region}(e_a, e_x) \wedge A(e_a, e_b) \wedge \textit{city}(e_b, e_c) \wedge \textit{largestCity}(e_d, e_c) \wedge \textit{state}(e_d, e_y)$). See Appendix for more learned alignment rules.

5 Conclusion

In conclusion, our study presents the Align-Subgraph Entity Alignment (ASGEA) model, a significant shift from traditional embedding-based entity alignment methods in KGs. ASGEA, through its innovative Path-based Graph Neural Network (ASGNN) and a novel multi-modal attention mechanism, addresses interpretability challenges and enhances alignment accuracy. Our experimental results validate ASGEA's effectiveness in both EA and MMEA tasks, marking a notable advancement in KG entity alignment research. This work not only broadens the scope of entity alignment methods but also sets a foundation for future explorations in integrating multi-modal data and improving interpretability in KG applications.

Limitations

We introduce a new paradigm for entity alignment based on logic rules. Despite our best efforts, ASGEA still has certain limitations.

Limitations of Inference Complexity. Our method ASGEA shifts away from conventional embedding-based methods, employing Align-Subgraph extraction algorithms and ASGNN for scoring to derive node-level alignment features. Despite its effectiveness in outperforming existing MMEA approaches without the need for multi-modal data, this strategy entails a higher time complexity due to the necessity of Align-Subgraph extraction for each entity pair. Future work will focus on developing similarly effective models with lower time complexity.

Limitations in Low-Source Scenarios. Our approach utilizes partial anchors to connect two KGs for Align-Subgraph extraction. However, in scenarios with limited anchors, it is challenging to propagate information to the aligned nodes as they are not reachable easily. Although we propose the inclusion of modal anchors as a remedy, this issue persists in cases of poor multi-modal data quality. Future efforts will explore iterative methods to improve performance in such low-source settings.

The Scope of Modalities. Our research covers three MMKG datasets, limited to text and image modalities. However, MMKGs also include other modalities like temporal information. Although our Align modal score is designed to handle such diversity theoretically, we lack empirical validation. Moving forward, we plan to investigate more modalities to assess our model's limits and increase its adaptability across various data types.

Ethical Considerations

Our experiments and model training strictly use publicly available datasets (as detailed in Appendix), mitigating ethical issues regarding privacy, confidentiality, or the misuse of personal biological information.

References

Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *NIPS*, pages 2787–2795.

Yixin Cao, Zhiyuan Liu, Chengjiang Li, Zhiyuan Liu, Juanzi Li, and Tat-Seng Chua. 2019. Multi-channel graph neural network for entity alignment. In *ACL (1)*, pages 1452–1461. Association for Computational Linguistics.

Liyi Chen, Zhi Li, Yijun Wang, Tong Xu, Zhefeng Wang, and Enhong Chen. 2020. MMEA: entity alignment for multi-modal knowledge graph. In *KSEM (1)*, volume 12274 of *Lecture Notes in Computer Science*, pages 134–147. Springer.

Liyi Chen, Zhi Li, Tong Xu, Han Wu, Zhefeng Wang, Nicholas Jing Yuan, and Enhong Chen. 2022. Multi-modal siamese network for entity alignment. In *KDD*, pages 118–126. ACM.

Muhao Chen, Yingtao Tian, Mohan Yang, and Carlo Zaniolo. 2017. Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In *IJCAI*, pages 1511–1517. ijcai.org.

Zhuo Chen, Jiaoyan Chen, Wen Zhang, Lingbing Guo, Yin Fang, Yufeng Huang, Yichi Zhang, Yuxia Geng, Jeff Z. Pan, Wenting Song, and Huajun Chen. 2023a. Meaformer: Multi-modal entity alignment transformer for meta modality hybrid. In *ACM Multimedia*, pages 3317–3327. ACM.

Zhuo Chen, Lingbing Guo, Yin Fang, Yichi Zhang, Jiaoyan Chen, Jeff Z. Pan, Yangning Li, Huajun Chen, and Wen Zhang. 2023b. Rethinking uncertainly missing and ambiguous visual modality in multi-modal entity alignment. In *ISWC*, volume 14265 of *Lecture Notes in Computer Science*, pages 121–139. Springer.

Zhuo Chen, Yichi Zhang, Yin Fang, Yuxia Geng, Lingbing Guo, Xiang Chen, Qian Li, Wen Zhang, Jiaoyan Chen, Yushan Zhu, Jiaqi Li, Xiaoze Liu, Jeff Z. Pan, Ningyu Zhang, and Huajun Chen. 2024. Knowledge graphs meet multi-modal learning: A comprehensive survey. *CoRR*, abs/2402.05391.

Kewei Cheng, Jiahao Liu, Wei Wang, and Yizhou Sun. 2022. Rlogic: Recursive logical rule learning from knowledge graphs. In *KDD*, pages 179–189. ACM.

Junyoung Chung, Çaglar Gülçehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*, abs/1412.3555.

Fabrizio Frasca, Beatrice Bevilacqua, Michael M. Bronstein, and Haggai Maron. 2022. Understanding and extending subgraph gnns by rethinking their symmetries. In *NeurIPS*.

Luis Antonio Galárraga, Christina Teflioudi, Katja Hose, and Fabian M. Suchanek. 2013. AMIE: association rule mining under incomplete evidence in ontological knowledge bases. In *WWW*, pages 413–422. International World Wide Web Conferences Steering Committee / ACM.

Yunjun Gao, Xiaoze Liu, Junyang Wu, Tianyi Li, Pengfei Wang, and Lu Chen. 2022. Clusterea: Scalable entity alignment with stochastic training and normalized mini-batch similarities. In *KDD*, pages 421–431. ACM.

Lingbing Guo, Zhuo Chen, Jiaoyan Chen, and Huajun Chen. 2023. Revisit and outstrip entity alignment: A perspective of generative models. *CoRR*, abs/2305.14651.

Fuzhen He, Zhixu Li, Qiang Yang, An Liu, Guanfeng Liu, Pengpeng Zhao, Lei Zhao, Min Zhang, and Zhigang Chen. 2019. Unsupervised entity alignment using attribute triples and relation triples. In *DASFAA (1)*, volume 11446 of *Lecture Notes in Computer Science*, pages 367–382. Springer.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778. IEEE Computer Society.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.

Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR (Poster)*. OpenReview.net.

Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. 2018. Canonical tensor decomposition for knowledge base completion. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 2869–2878. PMLR.

Manuel Leone, Stefano Huber, Akhil Arora, Alberto García-Durán, and Robert West. 2022. A critical re-evaluation of neural methods for entity alignment. *Proc. VLDB Endow.*, 15(8):1712–1725.

Chengjiang Li, Yixin Cao, Lei Hou, Jiaxin Shi, Juanzi Li, and Tat-Seng Chua. 2019. Semi-supervised entity alignment via joint knowledge embedding model and cross-graph model. In *EMNLP/IJCNLP (1)*, pages 2723–2732. Association for Computational Linguistics.

Qian Li, Shu Guo, Yangyifei Luo, Cheng Ji, Lihong Wang, Jiawei Sheng, and Jianxin Li. 2023a. Attribute-consistent knowledge graph representation learning for multi-modal entity alignment. In *WWW*, pages 2499–2508. ACM.

Yangning Li, Jiaoyan Chen, Yinghui Li, Yuejia Xiang, Xi Chen, and Hai-Tao Zheng. 2023b. Vision, deduction and alignment: An empirical study on multi-modal knowledge graph alignment. In *ICASSP*, pages 1–5. IEEE.

Yangning Li, Jiaoyan Chen, Yinghui Li, Yuejia Xiang, Xi Chen, and Haitao Zheng. 2023c. Vision, deduction and alignment: An empirical study on multi-modal knowledge graph alignment. *CoRR*, abs/2302.08774.

Ke Liang, Lingyuan Meng, Meng Liu, Yue Liu, Wenxuan Tu, Siwei Wang, Sihang Zhou, Xinwang Liu, and Fuchun Sun. 2022. Reasoning over different types of knowledge graphs: Static, temporal and multi-modal. *CoRR*, abs/2212.05767.

Zhenxi Lin, Ziheng Zhang, Meng Wang, Yinghui Shi, Xian Wu, and Yefeng Zheng. 2022. Multi-modal contrastive representation learning for entity alignment. In *COLING*, pages 2572–2584. International Committee on Computational Linguistics.

Fangyu Liu, Muhao Chen, Dan Roth, and Nigel Collier. 2021. Visual pivoting for (unsupervised) entity alignment. In *AAAI*, pages 4257–4266. AAAI Press.

Ye Liu, Hui Li, Alberto García-Durán, Mathias Niepert, Daniel Oñoro-Rubio, and David S. Rosenblum. 2019. MMKG: multi-modal knowledge graphs. In *ESWC*, volume 11503 of *Lecture Notes in Computer Science*, pages 459–474. Springer.

Zhiyuan Liu, Yixin Cao, Liangming Pan, Juanzi Li, and Tat-Seng Chua. 2020. Exploring and evaluating attributes, values, and structures for entity alignment. In *EMNLP (1)*, pages 6355–6364. Association for Computational Linguistics.

Christian Meilicke, Melisachew Wudage Chekol, Daniel Ruffinelli, and Heiner Stuckenschmidt. 2019. Anytime bottom-up rule learning for knowledge graph completion. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 3137–3143. ijcai.org.

Shichao Pei, Lu Yu, Robert Hoehndorf, and Xiangliang Zhang. 2019. Semi-supervised entity alignment via knowledge graph embedding with awareness of degree difference. In *WWW*, pages 3130–3136. ACM.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.

Ali Sadeghian, Mohammadreza Armandpour, Patrick Ding, and Daisy Zhe Wang. 2019. DRUM: end-to-end differentiable rule mining on knowledge graphs. In *NeurIPS*, pages 15321–15331.

Xiaofei Shi and Yanghua Xiao. 2019. Modeling multi-mapping relations for precise cross-lingual entity alignment. In *EMNLP/IJCNLP (1)*, pages 813–822. Association for Computational Linguistics.

Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *ICLR*.

Fabian M. Suchanek, Serge Abiteboul, and Pierre Senellart. 2011. PARIS: probabilistic alignment of relations, instances, and schema. *Proc. VLDB Endow.*, 5(3):157–168.

Zequn Sun, Wei Hu, and Chengkai Li. 2017. Cross-lingual entity alignment via joint attribute-preserving embedding. In *ISWC (1)*, volume 10587 of *Lecture Notes in Computer Science*, pages 628–644. Springer.

Zequn Sun, Wei Hu, Chengming Wang, Yuxin Wang, and Yuzhong Qu. 2023. Revisiting embedding-based entity alignment: A robust and adaptive method. *IEEE Trans. Knowl. Data Eng.*, 35(8):8461–8475.

Zequn Sun, Wei Hu, Qingheng Zhang, and Yuzhong Qu. 2018. Bootstrapping entity alignment with knowledge graph embedding. In *IJCAI*, pages 4396–4402. ijcai.org.

Zequn Sun, Chengming Wang, Wei Hu, Muhao Chen, Jian Dai, Wei Zhang, and Yuzhong Qu. 2020a. Knowledge graph alignment network with gated multi-hop neighborhood aggregation. In *AAAI*, pages 222–229. AAAI Press.

Zequn Sun, Qingheng Zhang, Wei Hu, Chengming Wang, Muhao Chen, Farahnaz Akrami, and Chengkai Li. 2020b. A benchmarking study of embedding-based entity alignment for knowledge graphs. *Proc. VLDB Endow.*, 13(11):2326–2340.

Wei Tang, Fenglong Su, Haifeng Sun, Qi Qi, Jingyu Wang, Shimin Tao, and Hao Yang. 2023. Weakly supervised entity alignment with positional inspiration. In *WSDM*, pages 814–822. ACM.

Zhichun Wang, Qingsong Lv, Xiaohan Lan, and Yu Zhang. 2018. Cross-lingual knowledge graph alignment via graph convolutional networks. In *EMNLP*, pages 349–357. Association for Computational Linguistics.

Fan Yang, Zhilin Yang, and William W. Cohen. 2017. Differentiable learning of logical rules for knowledge base reasoning. In *NIPS*, pages 2319–2328.

Qingheng Zhang, Zequn Sun, Wei Hu, Muhao Chen, Lingbing Guo, and Yuzhong Qu. 2019. Multi-view knowledge graph embedding for entity alignment. In *IJCAI*, pages 5429–5435. ijcai.org.

Yongqi Zhang and Quanming Yao. 2022. Knowledge graph reasoning with relational digraph. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, pages 912–924. ACM.

Hao Zhu, Ruobing Xie, Zhiyuan Liu, and Maosong Sun. 2017. Iterative entity alignment via joint knowledge embeddings. In *IJCAI*, pages 4258–4264. ijcai.org.

Xiangru Zhu, Zhixu Li, Xiaodan Wang, Xueyao Jiang, Penglei Sun, Xuwu Wang, Yanghua Xiao, and Nicholas Jing Yuan. 2024. Multi-modal knowledge graph construction and application: A survey. *IEEE Trans. Knowl. Data Eng.*, 36(2):715–735.

A Appendix

A.1 Dataset Statistics

The detailed statistics of our dataset are delineated in Table 5, aligning with Chen et al. (2023a,b); Guo et al. (2023). It is important to note that our dataset includes a collection of pre-aligned entity pairs. These pairs are systematically divided into a training set (denoted as seed alignments $\mathcal{S}$) and a testing set (denoted as $\mathcal{S}_{te}$), under the specified seed alignment ratio (R_{sa}).

A.2 Baseline Details

A.2.1 Comparision Methods

For traditional EA methods, we select IPTransE (Zhu et al., 2017), GCN-Align (Wang et al., 2018), KECG (Li et al., 2019), AlignEA (Sun et al., 2018), SEA (Pei et al., 2019) and IMUSE (He et al., 2019),MUGNN (Cao et al., 2019),AliNet (Sun et al., 2020a) as baselines. For Multi-modal methods, we select PoE (Liu et al., 2019), MMEA (Chen et al., 2020), EVA (Liu et al., 2021), MCLEA (Lin et al., 2022), MSNEA (Chen et al., 2022), MEAformer (Chen et al., 2023a), ACK-MMEA (Li et al., 2023a), UMAEA (Chen et al., 2023b). For a clear comparison, all the baselines are not iterable and do not use the surface information.

A.2.2 Traditional methods

Traditional methods primarily utilize relational modalities, with some incorporating attribute modalities. Key methods in this category include:

IPTransE (Zhu et al., 2017): It iteratively expands a set of softly aligned entity pairs, using a shared parameter strategy for different KGs.

GCN-Align (Wang et al., 2018): This approach integrates structural and attribute information through graph convolutional networks.

BootEA (Sun et al., 2018): It labels probable alignments as training data, reducing error accumulation through an iterative alignment method.

SEA (Pei et al., 2019): This method leverages both labeled and abundant unlabeled entity information for alignment, incorporating adversarial training to account for differences in point degrees.

IMUSE (He et al., 2019): It employs a bivariate regression model for combining alignment results, focusing on relation and attribute similarities.

A.2.3 Multi-modal methods

Multi-modal methods involve relational, attribute, and visual modalities. Prominent methods include:

EVA (Liu et al., 2021): Utilizes visual semantic representations for unsupervised entity alignment in heterogeneous KGs, excelling in aligning long-tail entities by creating comprehensive entity representations.

MMEA (Chen et al., 2020): Generates entity representations from relational, visual, and numerical knowledge, and then integrates these modalities.

MCLEA (Lin et al., 2022): Proposes a multi-modal contrastive learning framework for entity representation learning across different modalities.

MSNEA (Chen et al., 2022): Introduces a multi-modal siamese network approach for entity representation learning.

MEAformer (Chen et al., 2023a): Employs a multi-modal transformer-based approach for learning entity representations.

ACK-MMEA (Li et al., 2023a): Focuses on contextual gap problems, and introduces an attribute-consistent KG entity alignment framework considering multiple modalities.

UMAEA (Chen et al., 2023b): A robust multi-modal entity alignment method designed to handle uncertainties and ambiguities in visual modalities.

A.3 Implementation Details

We tune the learning rate in [10^{-4},10^{-2}], weight decay in [10^{-5},10^{-2}], dropout rate in [0,0.3], batch size in {4, 8, 16}, dimension d in {32,48,64}, Adam (Kingma and Ba, 2015) is used as the optimizer, the training epoch is 50. For the GNN layer L, we use $L = 5$ for MMKG and $L = 7$ for DBP15K and Multi-OpenEA depending on the density of KG. All experiments were conducted on Tesla V100. We employ the Gated Recurrent Unit (GRU) (Chung et al., 2014) as our RNN model. Our work is based on the open-source codebase of RED-GNN[1] and MEAformer[2]. For other baselines AnyBURL[3] and Neural LP[4], we use their open-source codebase.

A.4 Complete experiment results

We provide the complete results of all experiments on MMKG, DBP15K, and Multi-OpenEA datasets as shown in Tables 6, 7, and 8.

[1] https://github.com/LARS-research/RED-GNN.

[2] https://github.com/zjukg/MEAformer. MIT license.

[3] https://web.informatik.uni-mannheim.de/AnyBURL. 3-clause BSD license.

[4] https://github.com/kexinyi71/Neural-LP. MIT license.

Dataset	KG	# Ent.	# Rel.	# Attr.	# Rel. Triples	# Attr. Triples	# Image	# EA pairs
FBDB15K	FB15K	14,951	1,345	116	592,213	29,395	13,444	12,846
	DB15K	12,842	279	225	89,197	48,080	12,837	
FBYG15K	FB15K	14,951	1,345	116	592,213	29,395	13,444	11,199
	YAGO15K	15,404	32	7	122,886	23,532	11,194	
$\text{DBP15K}_{\text{ZH-EN}}$	ZH (Chinese)	19,388	1,701	8,111	70,414	248,035	15,912	15,000
	EN (English)	19,572	1,323	7,173	95,142	343,218	14,125	
$\text{DBP15K}_{\text{JA-EN}}$	JA (Japanese)	19,814	1,299	5,882	77,214	248,991	12,739	15,000
	EN (English)	19,780	1,153	6,066	93,484	320,616	13,741	
$\text{DBP15K}_{\text{FR-EN}}$	FR (French)	19,661	903	4,547	105,998	273,825	14,174	15,000
	EN (English)	19,993	1,208	6,422	115,722	351,094	13,858	
$\text{OpenEA}_{\text{EN-FR}}$	EN (English)	15,000	267	308	47,334	73,121	15,000	15,000
	FR (French)	15,000	210	404	40,864	67,167	15,000	
$\text{OpenEA}_{\text{EN-DE}}$	EN (English)	15,000	215	286	47,676	83,755	15,000	15,000
	DE (German)	15,000	131	194	50,419	156,150	15,000	
$\text{OpenEA}_{\text{D-W-V1}}$	DBpedia	15,000	248	342	38,265	68,258	15,000	15,000
	Wikidata	15,000	169	649	42,746	138,246	15,000	
$\text{OpenEA}_{\text{D-W-V2}}$	DBpedia	15,000	167	175	73,983	66,813	15,000	15,000
	Wikidata	15,000	121	457	83,365	175,686	15,000	

Table 5: Statistics for three datasets, where "EA pairs" refers to the pre-aligned entity pairs.

	Models	FBDB15K(20%)			FBDB15K(50%)			FBDB15K(80%)			FBYG15K(20%)			FBYG15K(50%)			FBYG15K(80%)		
		H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR
EA	IPTransE	.065	.215	.094	.210	.421	.283	.403	.627	.469	.047	.169	.084	.201	.369	.248	.401	.602	.458
	GCN-align	.053	.174	.087	.226	.435	.293	.414	.635	.472	.081	.235	.153	.233	.424	.294	.406	.643	.477
	SEA	.170	.425	.255	.373	.657	.470	.512	.784	.505	.141	.371	.218	.294	.577	.388	.514	.773	.605
	IMUSE	.176	.435	.264	.309	.576	.400	.457	.726	.551	.081	.257	.142	.398	.601	.469	.512	.707	.581
	Neural LP$*$	.098	.297	.161	.233	.552	.339	.314	.604	.412	.106	.285	.163	.281	.531	.362	.343	.631	.441
	AnyBURL$*$	.213	.515	.312	.386	.698	.491	.437	.727	.533	.210	.519	.311	.356	.662	.459	.397	.678	.489
	RED-GNN$*$	.482	.687	.552	.692	.843	.746	.765	.882	.817	.464	.659	.532	.647	.815	.707	.728	.871	.779
	ASGEA-Stru	**.506**	**.706**	**.575**	**.711**	**.860**	**.765**	**.791**	**.897**	**.828**	**.479**	**.681**	**.548**	**.659**	**.827**	**.724**	**.743**	**.888**	**.796**
MMEA	MMEA	.265	.541	.357	.417	.703	.512	.590	.869	.685	.234	.480	.317	.403	.645	.486	.598	.839	.682
	EVA	.199	.448	.283	.334	.589	.422	.484	.696	.563	.153	.361	.224	.311	.534	.388	.491	.692	.565
	MSNEA	.114	.296	.175	.288	.590	.388	.518	.779	.613	.103	.249	.153	.320	.589	.413	.531	.778	.620
	MCLEA	.295	.582	.393	.555	.784	.637	.735	.890	.790	.254	.484	.332	.501	.705	.574	.667	.824	.722
	ACK-MMEA	.304	.549	.387	.560	.736	.624	.682	.874	.752	.289	.496	.360	.535	.699	.593	.676	.864	.744
	MEAformer	.434	.728	.534	.625	.847	.704	.773	.918	.825	.325	.598	.416	.560	.780	.640	.705	.874	.768
	ASGEA-MM w/o value	**.516**	**.718**	**.586**	**.725**	**.867**	**.777**	**.800**	**.908**	**.839**	**.493**	**.696**	**.563**	**.671**	**.835**	**.730**	**.756**	**.893**	**.807**
	ASGEA-MM	**.628**	**.799**	**.689**	**.794**	**.899**	**.833**	**.852**	**.927**	**.881**	**.717**	**.848**	**.776**	**.842**	**.920**	**.879**	**.902**	**.965**	**.926**

Table 6: Performance comparison on FBDB15K and FBYG15K datasets at various reference R_{sa} for training, between EA and MMEA models. "ASGEA-MM" indicates our complete model that uses multi-modal information(abbreviated as MM), while "ASGEA-Stru" indicates reliance on graph structure alone (i.e., triples), "w/o value" denotes the exclusion of attribute value. "$*$" indicates the reasoning method adapted in our framework to EA task.

	Models	$\text{DBP15K}_{\text{ZH-EN}}$			$\text{DBP15K}_{\text{JA-EN}}$			$\text{DBP15K}_{\text{FR-EN}}$		
		H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR
EA	AlignEA (Sun et al., 2018)	.472	.792	.581	.448	.789	.563	.481	.824	.599
	KECG (Li et al., 2019)	.478	.835	.598	.490	.844	.610	.486	.851	.610
	MUGNN (Cao et al., 2019)	.494	.844	.611	.501	.857	.621	.495	.870	.621
	AliNet (Sun et al., 2020a)	.539	.826	.628	.549	.831	.645	.552	.852	.657
	ASGEA-Stru	**.560**	**.844**	**.660**	**.595**	**.866**	**.690**	**.653**	**.906**	**.745**
MMEA	MSNEA (Chen et al., 2022)	.609	.831	.685	.541	.776	.620	.557	.820	.643
	EVA (Liu et al., 2021)	.683	.906	.762	.669	.904	.752	.686	.928	.771
	MCLEA (Lin et al., 2022)	.726	.922	.796	.719	.915	.789	.719	.918	.792
	MEAformer (Chen et al., 2023a)	.771	.951	.835	.764	.959	.834	.770	.961	.841
	UMAEA (Chen et al., 2023b)	.800	.962	.860	.801	.967	.862	.818	.973	.877
	ASGEA-MM	**.815**	**.941**	**.862**	**.849**	**.956**	**.889**	**.911**	**.982**	**.938**

Table 7: Complete results on DPB15K (Sun et al., 2017) datasets.

Models	$\text{OpenEA}_{\text{EN-FR}}$			$\text{OpenEA}_{\text{EN-DE}}$			$\text{OpenEA}_{\text{D-W-V1}}$			$\text{OpenEA}_{\text{D-W-V2}}$		
	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR
MSNEA* (Chen et al., 2022)	.692	.813	.734	.753	.895	.804	.800	.874	.826	.838	.940	.873
EVA* (Liu et al., 2021)	.785	.932	.836	.922	.983	.945	.858	.946	.891	.890	.981	.922
MCLEA* (Lin et al., 2022)	.819	.943	.864	.939	.988	.957	.881	.955	.908	.928	.983	.949
UMAEA* (Chen et al., 2023b)	.848	.966	.891	.956	.994	.971	**.904**	**.971**	**.930**	.948	**.996**	.967
ASGEA-MM	**.911**	**.968**	**.932**	**.980**	**.996**	**.986**	.898	.968	.924	**.958**	.994	**.972**

Table 8: Complete results on four Multi-OpenEA datasets.

A.5 Align-Subgraph Extracting Algorithm

We have developed various versions of the extracting algorithm, including a non-merging version, a merging version, and a version focused solely on symmetric rules.

The non-merging version, as delineated in Algorithm 1, extracts the ASG between e_u and e_v.

The merging version, outlined in Algorithm 2, extracts a combined alignment subgraph from entity e_u to all entities in $\mathcal{G}_2$.

The version concentrating only on symmetric rules filters the rules accordingly, ensuring that the Align-Subgraph comprises only paths based on symmetric rules, as shown in Algorithm 3.

A.6 More Visualization of Learned ASG and alignment rules

Figure 8 visualizes more learned ASG with alignment paths on different samples. We also provide more learned rules in Table 9.

Algorithm 1 ASG Extraction

Require: $\mathcal{G}_1 = \{\mathcal{E}_1, \mathcal{R}_1, \mathcal{T}_1\}, \mathcal{G}_2 = \{\mathcal{E}_2, \mathcal{R}_2, \mathcal{T}_2\}$, Anchor Links A, Node pair (e_u, e_v), hop K
Ensure: K-hop ASG $g^K(e_u, e_v)$ for node pair (e_u, e_v)

Initialize $\mathcal{E}_s^0 \leftarrow \{e_u\}$, $\mathcal{E}_t^0 \leftarrow \{e_v\}$, $\mathcal{T} \leftarrow \mathcal{T}_1 \cup \mathcal{T}_2 \cup A \cup \mathcal{T}_1^{rev} \cup \mathcal{T}_2^{rev} \cup A^{rev}$
for $i = 1$ to K **do**
  $\mathcal{E}_s^i \leftarrow \{e \mid (e_u, r, e) \in \mathcal{T}, e_u \in \mathcal{E}_s^{i-1}\}$
  $\mathcal{T}^i \leftarrow \{(e_u, r, e) \mid (e_u, r, e) \in \mathcal{T}, e_u \in \mathcal{E}_s^{i-1}\}$
end for
for $i = 1$ to K **do**
  $\mathcal{E}_t^i \leftarrow \{e \mid (e, r, e_v) \in \mathcal{T}^{K-i}, e_v \in \mathcal{E}_t^{i-1}\} \cup \mathcal{E}_t^{i-1}$
  $\mathcal{T}_t^i \leftarrow \{(e, r, e_v) \mid (e, r, e_v) \in \mathcal{T}^{K-i}, e_v \in \mathcal{E}_t^{i-1}\}$
end for
return $\mathcal{T}_t^0 \cup \cdots \cup \mathcal{T}_t^K$

Algorithm 2 Merged ASG Extraction

Require: $\mathcal{G}_1 = \{\mathcal{E}_1, \mathcal{R}_1, \mathcal{T}_1\}, \mathcal{G}_2 = \{\mathcal{E}_2, \mathcal{R}_2, \mathcal{T}_2\}$, Anchor Links A, Node e_u, hop K
Ensure: K-hop ASG g_u^K for Node e_u

Initialize $E_s^0 \leftarrow \{e_u\}$, $\mathbf{E}_t^0 \leftarrow \begin{cases} \mathcal{E}_1, & e_u \in \mathcal{E}_2 \\ \mathcal{E}_2, & e_u \in \mathcal{E}_1 \end{cases}$,$\mathcal{T} \leftarrow \mathcal{T}_1 \cup \mathcal{T}_2 \cup A \cup \mathcal{T}_1^{rev} \cup \mathcal{T}_2^{rev} \cup A^{rev}$
for $i = 1$ to K **do**
  $\mathcal{E}_s^i \leftarrow \{e \mid (e_u, r, e) \in \mathcal{T}, e_u \in \mathcal{E}_s^{i-1}\}$
  $\mathcal{T}^i \leftarrow \{(e_u, r, e) \mid (e_u, r, e) \in \mathcal{T}, e_u \in \mathcal{E}_s^{i-1}\}$
end for
for $i = 1$ to K **do**
  $\mathcal{E}_t^i \leftarrow \{e \mid (e, r, e_u) \in \mathcal{T}^{K-i}, e_u \in \mathcal{E}_t^{i-1}\} \cup \mathcal{E}_t^{i-1}$
  $T_t^i \leftarrow \{(e, r, e_u) \mid (e, r, e_u) \in \mathcal{T}^{K-i}, e_u \in \mathcal{E}_t^{i-1}\}$
end for
return $\mathcal{T}_t^0 \cup \cdots \cup \mathcal{T}_t^K$

Algorithm 3 Symmetric ASG Extraction

Require: $\mathcal{G}_1 = \{\mathcal{E}_1, \mathcal{R}_1, \mathcal{T}_1\}, \mathcal{G}_2 = \{\mathcal{E}_2, \mathcal{R}_2, \mathcal{T}_2\}$, Anchor Links A, Node e_u, hop K
Ensure: K-hop Align-Subgraph g_u^K for Node e_u

Initialize $\mathcal{E}_s^0 \leftarrow \{e_u\}$, $\mathcal{E}_t^0 \leftarrow \begin{cases} \mathcal{E}_1, & e_u \in \mathcal{E}_2 \\ \mathcal{E}_2, & e_u \in \mathcal{E}_1 \end{cases}$, $\mathcal{T} \leftarrow \mathcal{T}_1 \cup \mathcal{T}_2 \cup A \cup \mathcal{T}_1^{rev} \cup \mathcal{T}_2^{rev} \cup A^{rev}$
for $i = 1$ to K **do**
  $\mathcal{E}_s^i \leftarrow \{e \mid (e_u, r, e) \in \mathcal{T}, e_u \in \mathcal{E}_s^{i-1}\}$
  $\mathcal{T}^i \leftarrow \{(e_u, r, e) \mid (e_u, r, e) \in \mathcal{T}, e_u \in \mathcal{E}_s^{i-1}\}$
end for
for $i = 1$ to K **do**
  $\mathcal{E}_t^i \leftarrow \{e \mid (e, r, e_u) \in \mathcal{T}^{K-i}, e_u \in \mathcal{E}_t^{i-1}\} \cup \mathcal{E}_t^{i-1}$
  $\mathcal{T}_t^i \leftarrow \{(e, r, e_u) \mid (e, r, e_u) \in \mathcal{T}^{K-i}, e_u \in \mathcal{E}_t^{i-1}\}$
  if $i < K//2$ **then**
    $\mathcal{E}_t^i \leftarrow \mathcal{E}_t^i \cap \mathcal{E}_2$
  else
    $\mathcal{E}_t^i \leftarrow \mathcal{E}_t^i \cap \mathcal{E}_1$
  end if
end for
return $\mathcal{T}_t^0 \cup \cdots \cup \mathcal{T}_t^K$

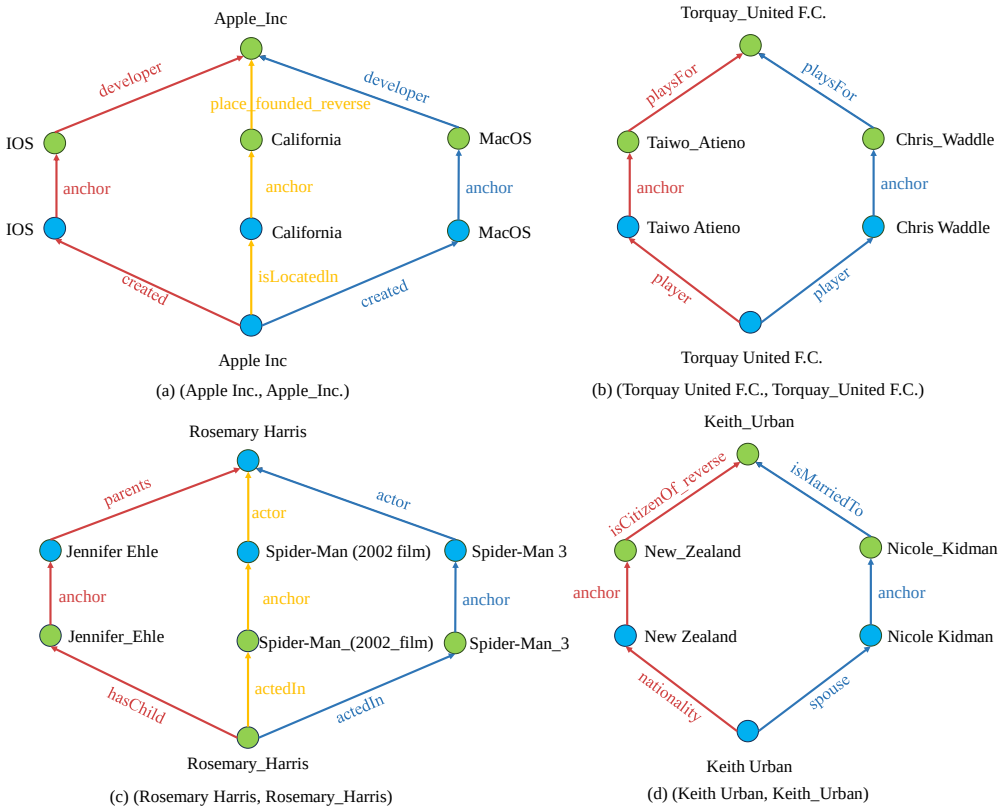

(a) (Apple Inc., Apple_Inc.)

(b) (Torquay United F.C., Torquay_United F.C.)

(c) (Rosemary Harris, Rosemary_Harris)

(d) (Keith Urban, Keith_Urban)

Figure 8: More visualization of learned ASG

hop=3	A(X,Y) ← directed_by(A,X) ∧ A(A,B) ∧ director(B,Y)
	A(X,Y) ← capital(A,X) ∧ A(A,B) ∧ capital(B,Y)
	A(X,Y) ← children(X,A) ∧ A(A,B) ∧ parent(B,Y)
	A(X,Y) ← written_by(A,X) ∧ A(A,B) ∧ writer(B,Y)
	A(X,Y) ← player(A,X) ∧ A(A,B) ∧ draftTeam(Y,B)
hop=4	A(X,Y) ← children(A,X) ∧ A(A,B) ∧ spouse(C,B) ∧ child(C,Y)
	A(X,Y) ← capital(X,A) ∧ county_seat(B,A) ∧ A(B,C) ∧ state(C,Y)
	A(X,Y) ← location(A,X) ∧ school(A,B) ∧ A(B,C) ∧ city(C,Y)
	A(X,Y) ← geographic_scope(A,X) ∧ A(A,B) ∧ headquarter(B,C) ∧ country(C,Y)
	A(X,Y) ← country(A,X) ∧ citytown(B,A) ∧ A(B,C) ∧ country(C,Y)
hop=5	A(X,Y) ← capital(A,X) ∧ administrative_parent(B,A) ∧ A(B,C) ∧ state(C,D) ∧ capital(D,Y)
	A(X,Y) ← parents(A,X) ∧ award_nominee(B,A) ∧ A(B,C) ∧ starring(C,D) ∧ parent(D,Y)
	A(X,Y) ← state_province_region(A,X) ∧ A(A,B) ∧ city(B,C) ∧ largestCity(D,C) ∧ state(D,Y)
	A(X,Y) ← currency_used(A,X) ∧ country(B,A) ∧ A(B,C) ∧ country(C,D) ∧ currency(D,Y)
	A(X,Y) ← children(X,A) ∧ children(B,A) ∧ A(B,C) ∧ parent(D,C) ∧ parent(D,Y)

Table 9: More alignment rules learned from FBDB15K.