

Progressively Modality Freezing for Multi-Modal Entity Alignment

Yani Huang¹, Xuefeng Zhang¹, Richong Zhang^{1,2*}, Junfan Chen³, Jaein Kim¹

¹CCSE, School of Computer Science and Engineering, Beihang University, Beijing, China

²Zhongguancun Laboratory, Beijing, China

³School of Software, Beihang University, Beijing, China

{huangyn, zhangxf, jaein}@buaa.edu.cn

{zhangrc, chenjyf}@act.buaa.edu.cn

Abstract

Multi-Modal Entity Alignment aims to discover identical entities across heterogeneous knowledge graphs. While recent studies have delved into fusion paradigms to represent entities holistically, the elimination of features irrelevant to alignment and modal inconsistencies is overlooked, which are caused by inherent differences in multi-modal features. To address these challenges, we propose a novel strategy of progressive modality freezing, called PMF, that focuses on alignment-relevant features and enhances multi-modal feature fusion. Notably, our approach introduces a pioneering cross-modal association loss to foster modal consistency. Empirical evaluations across nine datasets confirm PMF’s superiority, demonstrating state-of-the-art performance and the rationale for freezing modalities. Our code is available at <https://github.com/ninibymilk/PMF-MMEA>.

1 Introduction

Multi-modal Knowledge Graphs (MMKGs) integrate various modalities, including text, vision, and structural data, to provide comprehensive insights into knowledge. This integration underpins a range of applications, from question answering (Zhu et al., 2015) and information retrieval (Dietz et al., 2018; Yang, 2020), to recommendation systems (Sun et al., 2020a). Multi-Modal Entity Alignment (MMEA) aims to identify identical entities across heterogeneous MMKGs, which is essential for the integrity of the knowledge represented within these KGs. The essence of MMEA lies in identifying feature commonalities across entities from varied modalities to determine their alignment. The diversity in KG construction introduces potential mismatches in multi-modal features of entities meant to be aligned. For example,

*Corresponding author.




Figure 1: Illustration of irrelevant vs. relevant features in multi-modal knowledge graphs.

Figure 1 depicts how the entity “Interstellar (film)” might be represented in one knowledge graph \mathcal{G} by a poster image of a spaceship, whereas in another knowledge graph \mathcal{G}' , by a portrait of Anne Hathaway, the starring actress. Although both images are related to “Interstellar (film)”, their disparate perspectives can weaken the semantic relationship, challenging the alignment task. This scenario underscores the problem of *alignment-irrelevant features*, distinctly different features that complicate accurate entity alignment. While recent studies (Lin et al., 2022; Chen et al., 2023a) have employed attention mechanisms to calculate cross-modal weights at both the KG and entity levels, they overlook the critical step of evaluating modality relevance, thereby neglecting to exclude these irrelevant features.

Another obstacle in MMEA involves ensuring semantic consistency across different modalities. Achieving consistent representations across modalities is crucial for effective entity alignment. Existing methods (Lin et al., 2022; Chen et al., 2023a,b) utilize contrastive learning to minimize the intra-modal inconsistencies among aligned entities, yet overlook the *cross-modal inconsistencies*. Furthermore, the presence of alignment-irrelevant features exacerbates these inconsisten-

cies, complicating the task of achieving consistent representations across modalities.

To address the aforementioned challenges, we propose a novel method for multi-modal entity alignment, named Progressive Modality Freezing (PMF), designed to effectively filter out irrelevant features. PMF is structured around three key parts: multi-modal entity encoders, progressive multi-modality feature integration, and a unified training objective for MMEA. Initially, PMF separately encodes entity features from each modality, allowing for flexible adjustment of input modalities. Following this, the method employs a progressive approach to selectively freeze features deemed less relevant for alignment, simultaneously integrating valuable multi-modal features based on their alignment relevance. The culmination of our strategy is a unified training objective, aimed at minimizing the discrepancies both within individual KGs and across different modalities.

To encapsulate, our research contributes in the following three-folds:

- We propose the Progressive Modality Freezing (PMF) strategy to tackle the challenge of irrelevant features in MMEA. By assigning alignment-relevance scores to modalities, our method progressively freezes features that are alignment-irrelevant, while seamlessly integrating beneficial multi-modal information, ensuring the emphasis remains on features of utmost relevance.
- We are at the forefront of employing contrastive learning to address modality consistency across multiple modalities, accompanied by a unified training objective to enhance cross-modal consistency.
- We confirm the effectiveness of PMF across diverse datasets and experimental settings, where it demonstrates superior performance, achieving state-of-the-art in the field. Our thorough analysis further elucidates the advantage of the feature-freezing strategy.

2 Related Work

2.1 Multi-modal Knowledge Graph Representation Learning

The field of MMKG representation learning forms a critical foundation for tasks such as entity alignment and link prediction, transitioning from

traditional single-modal knowledge graph methods to more complex that leverage diverse data types. Early efforts expanded upon established KG representation learning methods for single-modal knowledge graphs, like those introduced by Xie et al. (2016) and Mousselly-Sergieh et al. (2018), adapting translational KG embedding techniques (Bordes et al., 2013) to incorporate multi-modal data effectively. Subsequent research has delved deeper into the impact of multi-modal information in representation learning. Pezeshkpour et al. (2018) implemented an adversarial method to tackle the challenge of missing data within MMKGs, while Fang et al. (2022) explored a contrastive learning approach to leverage the rich, varied data and heterogeneous structures in MMKGs. These advancements highlight the field’s move towards more sophisticated, data-inclusive methodologies.

2.2 Multi-Modal Entity Alignment

The exploration of MMEA began with MMKG (Liu et al., 2019), which provided a dataset facilitating entity embedding through linked images. Building on this, MMEA (Chen et al., 2020) advanced the field by fusing relational triples, images, and numerical attributes to generate entity embeddings.

Given the distinctiveness of images as a modality, utilizing visual data in alignment processes has emerged as a key focus in MMEA. For instance, EVA (Liu et al., 2021) exploited visual similarities for unsupervised entity alignment, while MSNEA (Chen et al., 2022) and UMAEA (Chen et al., 2023b) enhanced relational feature learning using visual cues to tackle the challenges posed by missing and ambiguous visual data.

Central to MMEA is the integration of multi-modal information. MEAformer (Chen et al., 2023a) employed a dynamic weighting transformer for integration, whereas MCLEA (Lin et al., 2022) aimed at bridging the modality gap via contrastive learning. ACK-MMEA (Li et al., 2023) introduced a strategy for normalizing multi-modal attributes, ensuring consistency across modalities.

Iterative learning has emerged as a powerful strategy for refining entity alignment with methods like XGEA (Xu et al., 2023) and PSNEA (Ni et al., 2023) employing pseudo-labeling to mitigate scarcity of initial seeds.

Despite these advancements, existing methods neglect the necessity of filtering out alignment-irrelevant features and addressing cross-modal inconsistencies. Our method addresses these gaps by selectively freezing less pertinent modalities within an entity to enhance the alignment process.

3 Model

We propose a multi-modal entity alignment framework based on progressively freezing entities with irrelevant features and contrastive learning. As shown in Figure 2, the framework consists of three parts: multi-modal entity encoder, progressive multi-modality feature integration, and cross-graph contrastive learning. First, a multi-modal encoder is used to transform the raw input of each modality graph into entity embeddings. Then, we integrate modality features progressively during training, which gradually freezes alignment-irrelevant features and fuse the modality graphs. Finally, the model is optimized through contrastive loss between different knowledge graphs and different modality graphs.

3.1 Preliminaries

MMKG Definition A multi-modal knowledge graph (MMKG) is defined as $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{T}, \mathcal{A}\}$, where \mathcal{E} is a set of entities, \mathcal{R} is a set of relations between entities, $\mathcal{T} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ denotes a set of triples in the knowledge graph, and \mathcal{A} is a set of multi-modal attributes of entities (e.g., textual descriptions, images, etc.). For modality $m \in \mathcal{M}$, where \mathcal{M} is the set of modalities, we define a modality graph \mathcal{G}_m that captures a unique aspect pertinent to the knowledge graph. More specifically, $\mathcal{G}_m = \{\mathcal{E}, \mathcal{R}, \mathcal{T}, \mathcal{A}_m\}$, where \mathcal{A}_m is the set of attributes associated with modality m .

Task Description In the context of multi-modal entity alignment (MMEA), consider two MMKGs $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{T}, \mathcal{A}\}$ and $\mathcal{G}' = \{\mathcal{E}', \mathcal{R}', \mathcal{T}', \mathcal{A}'\}$. The set $\mathcal{S} = \{(e_i, e_j) | e_i \in \mathcal{E}, e_j \in \mathcal{E}'\}$ denotes pre-established pairs of entities known to correspond to the same real-world object across both graphs. In a supervised learning setting, a model is trained with part of \mathcal{S} and then tasked to predict the alignments of the remaining entity pairs.

3.2 Multi-Modal Entity Encoder

In the Multi-modal Entity Encoder module, we process the features of entities from each KG according to their respective modalities. Given an

entity $e_i \in \mathcal{G}_m$, its feature is denoted by e_i^m . The encoding process for each modality of an entity is defined by the following equation:

$$h_i^m = \text{ENC}_m(\Theta_m, e_i^m), \quad (1)$$

where ENC_m represents the encoder designed for modality m with its associated learnable parameter Θ_m . While our framework allows the integration of any modality, for a fair comparison, we incorporate four modalities: structure (*str*), relation (*rel*), attribute (*att*), and image (*img*). Details of ENC_m are attached in Appendix A.1

3.3 Progressive Multi-Modality Feature Integration

The process of integrating multi-modal features in a stepwise manner unfolds in three distinct phases. Initially, the process involves calculating an alignment-relevant score for each modality graph (e.g., \mathcal{G}_m) to discern features that are irrelevant for alignment from those that are crucial. Following this, certain features are selectively ‘frozen’ based on modality scores to mitigate the inclusion of extraneous features. Finally, features from various modality graphs are fused to align multi-modal entities. This approach is executed iteratively and progressively to identify and incorporate beneficial features into the alignment process throughout the training phase, ensuring a refined alignment outcome.

3.3.1 Feature Relevance Measuring

We assess the likelihood that features from a given modality of an entity will contribute to the alignment process, thereby enabling the identification of potential alignment-irrelevant features within each modality graph. The alignment-relevant scores for a modality graph \mathcal{G}_m are denoted as \mathcal{W}_m , calculated by the following equation:

$$\mathcal{W}_m = \text{SCORE}(\mathcal{G}_m, \mathcal{G}'_m), \quad (2)$$

where $\text{SCORE}(\cdot)$ represents a score function. Specifically, for an entity e_i^m from \mathcal{G}_m and considering a corresponding modality graph \mathcal{G}'_m from target knowledge graph \mathcal{G}' , the alignment-relevant score is computed as:

$$w_i^m = \text{Relu} \left(\frac{\alpha_i^m - \delta}{\max_{e_j \in \mathcal{E}}(\alpha_j^m) - \delta} \right), \quad (3)$$

$$\alpha_i^m = \max_{e_k \in \mathcal{E}'} \left(\frac{h_i^m \cdot h_k^m}{|h_i^m| |h_k^m|} \right), \quad (4)$$

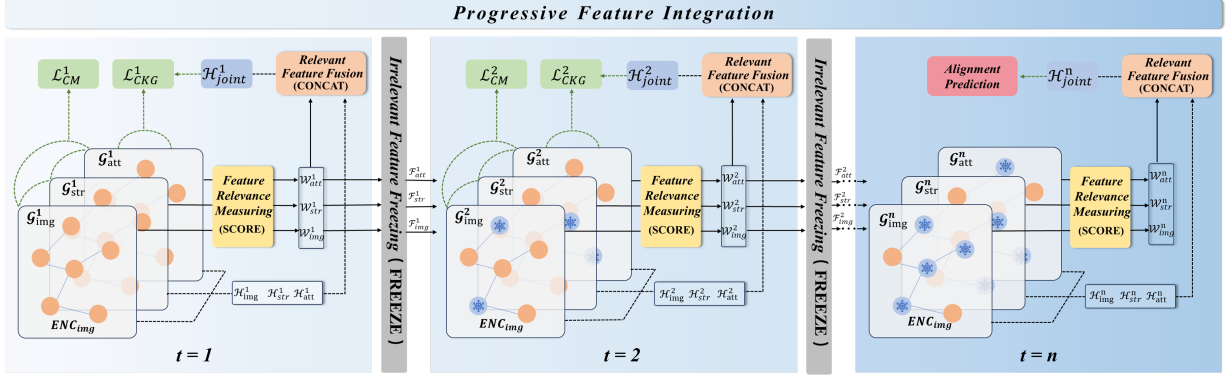


Figure 2: Overview of the PMF Model. The framework consists of three components: the Multi-Modal Entity Encoder (&3.2), Progressive Multi-Modality Feature Integration (&3.3), and Cross-Graph Contrastive Learning (&3.4). At epoch t , the multi-modal encoder ENC_m transforms raw inputs from each modality graph into entity embeddings \mathcal{H}_m^t . *Progressive Feature Integration* (&3.3.4) occurs during training, employing *Irrelevant Feature Freezing* (&3.3.2) and performing *Relevant Feature Fusion* (&3.3.3) guided by *Relevant Feature Measuring* (&3.3.1). The model is optimized using Cross-KG Alignment Loss \mathcal{L}_{CKG}^t (&3.4.2) and Cross-Modality Association Loss \mathcal{L}_{CM}^t (&3.4.1).

where δ is a dynamic threshold. The collection of w_i^m for all entities constitutes \mathcal{W}_m . In a parallel manner, the alignment-relevant scores \mathcal{W}_m' for \mathcal{G}_m' are computed similarly.

The underlying rationale for the alignment-relevant score is to ascertain whether an entity from the current modality graph \mathcal{G}_m exhibits similarity with entities from the modality graph \mathcal{G}_m' . Should the features h_i^m of an entity e_i in \mathcal{G}_m lack similar counterparts in \mathcal{G}_m' , the correlation between h_i^m and the entity aligned with e_i in \mathcal{G}_m' would consequently be weak. As such, h_i^m is less likely to contribute positively to the alignment process. To facilitate a reasonable measurement of feature similarity within a modality, we normalize this metric using the highest similarity across all entities from the same modality graph. This alignment-relevant score is subsequently utilized to guide the processes of modality feature freezing and fusion.

3.3.2 Irrelevant Feature Freezing

To minimize the detrimental impact of alignment-irrelevant features on the training process, we introduce a feature-freezing strategy. This approach selectively prevents back-propagation for specific entities within particular modality graphs, thereby preserving the integrity of the training regime. Formally, we define \mathcal{G}_m^t by extending \mathcal{G}_m with a function $\mathcal{F}_m^t : \mathcal{E} \rightarrow \{0, 1\}$, to characterize whether entity e is frozen at step t . More specifically, $\mathcal{G}_m^t = \{\mathcal{E}, \mathcal{R}, \mathcal{T}, \mathcal{A}_m, \mathcal{F}_m^t\}$, where for each entity e in \mathcal{E} , e is frozen only if $\mathcal{F}_m^t(e) = 0$. Correspond-

ingly, we refer \mathcal{H}_m^t and \mathcal{W}_m^t to the representation of all entities and their alignment-relevant score at step t , respectively. All elements in vector \mathcal{F}_m^t are initialized as 1.

The formalization of the entity freezing process is given by:

$$\mathcal{F}_m^t = \text{FREEZE}(\mathcal{W}_m^t, \mathcal{G}_m^t). \quad (5)$$

At step t , when the alignment-relevant score w_i^m for entity e_i from \mathcal{G}_m^t is zero, it is inferred that this entity does not contribute positively towards the training of the alignment model. Consequently, the gradient update for features h_i^m associated with such an entity is halted during back-propagation. The freezing operation for an entity e_i^m , contingent upon its score w_i^m , is defined as:

$$\text{FREEZE}(e_i^m) = \begin{cases} 0 & \text{if } w_i^m = 0, \\ 1 & \text{otherwise.} \end{cases}, \quad (6)$$

where 0 indicates "stop gradient", ensuring that the gradient flow is interrupted for entities deemed irrelevant, with $w_i^m \in \mathcal{W}_m^t$. $\text{FREEZE}(\cdot)$ systematically applies across all entities within \mathcal{G}_m^t . This feature freezing is enacted across all modality graphs of each KG (e.g., \mathcal{G} and \mathcal{G}') to optimize the alignment model's focus on beneficial features.

3.3.3 Relevant Feature Fusion

Recognizing that the significance of the features of each modality may vary across entities, we leverage the alignment-relevant scores introduced in

Eq. 3 as a guide during the fusion process.

$$\mathcal{H}_{joint}^t = \text{CONCAT}(\mathcal{W}_m^t, \mathcal{G}_m^t), \quad (7)$$

where $\text{CONCAT}(\cdot)$ symbolizes the fusion operator acting on all modality graphs \mathcal{G}_m^t , integrating them into a joint representation \mathcal{H}_{joint}^t .

In practice, for each entity e_i , the representations h_i^m are weighted by their respective confidence w_i^m , ensuring that modalities with higher relevance scores have a greater influence in the fused representation. Accordingly, at step t , this fusion process is defined as follows:

$$h_i^{joint} = \bigoplus_{m \in \mathcal{M}} ([w_i^m \cdot h_i^m]), \quad (8)$$

where \bigoplus denotes the vector concatenation operation, and h_i^{joint} represents the composite feature embedding for entity e_i , synthesized from all modalities. The collective h_i^{joint} across all entities at step t forms \mathcal{H}_{joint}^t .

3.3.4 Progressive Feature Integration

By integrating information from each modality graph, we can progressively optimize entity alignment. This is done by iteratively refining feature relevance, and gradually isolating alignment-irrelevant features from overt to subtle ones. Initially, basic representations are used to gauge modality confidence and selectively freeze modality features, employing a lower threshold to exclude overtly irrelevant features due to nascent modality representations. As training advances, the model accumulates a deeper understanding of modal information, systematically increasing the threshold to prune less obvious features independent of alignment goals.

The process unfolds as follows: Assuming the model has been trained for t epochs, in epoch $t+1$, we compute the alignment-relevant scores \mathcal{W}_m^{t+1} using the current stage of \mathcal{G}_m^t and \mathcal{G}_m^t . Concurrently, the threshold δ for discerning feature relevance is increased at a predefined rate. Based on the newly calculated scores \mathcal{W}_m^{t+1} , entities are frozen to reflect the updated insights, updating each modality graph to \mathcal{G}_m^{t+1} . Subsequently, these refined modality graphs are merged, leveraging the latest alignment-relevant scores to produce the joint graph representation $\mathcal{H}_{joint}^{t+1}$, as formalized in Equation (2), (5) and (7).

Comparatively, alternatives like ‘‘early integration’’ or ‘‘late integration’’ present limitations.

Early integration might only freeze superficially irrelevant features due to the premature stage of training, which can be challenging to correct in later stages. Our progressive strategy synergistically blends the benefits of both approaches, ensuring a balanced and effective integration of modality information throughout the training lifecycle.

3.4 Multi-Modal Entity Alignment Objective

Our framework adopts contrastive learning as a pivotal strategy to extract and leverage alignment information across knowledge graphs and to discern semantic associations between various modality graphs. To accomplish this, we have developed a composite loss function that integrates both cross-KG alignment and cross-modality association within a unified learning objective. The overall loss function at step t is defined as follows:

$$\mathcal{L}^t = \mathcal{L}_{\text{CM}}^t + \mathcal{L}_{\text{CKG}}^t, \quad (9)$$

where $\mathcal{L}_{\text{CKG}}^t$ is the alignment loss across different KGs, and $\mathcal{L}_{\text{CM}}^t$ denotes the association loss across modalities within each KG.

3.4.1 Cross-Modality Association Loss

One significant source of inconsistencies across modalities is alignment-irrelevant features. Freezing features within each modality graph allows our framework to more effectively draw together the remaining modal features that are beneficial for alignment.

Specifically for each entity e_i , we compute the contrastive loss between pairs of modality graphs \mathcal{G}_p and \mathcal{G}_q . The loss function aims to minimize the distance between positive pairs (e_i^p, e_i^q) and maximize the separation between negative pairs (e_i^p, e_j^q) and (e_j^p, e_i^q) . Here e_i and e_j are entities in the same knowledge graph \mathcal{G} . For each positive pair (e_i^p, e_i^q) , the cross-modality loss at step t is defined as:

$$l_{\text{CM}}^t(e_i^p, e_i^q) = \frac{\exp(h_i^p \cdot h_i^q / \tau)}{\exp(h_i^p \cdot h_j^q / \tau) + \sum_{(e_j^p, e_k^q) \in N_{i,p,q}^-} \exp(h_i^p \cdot h_j^q / \tau)}, \quad (10)$$

where $h_i^p \in \mathcal{H}_p^t$, $h_j^q \in \mathcal{H}_q^t$, and $\exp(h_i^p \cdot h_j^q / \tau)$ are scaled by a temperature factor τ , and set $N_{i,p,q}^-$ represents negative samples, encompassing all other entity pairs from the modality graphs \mathcal{G}_p and \mathcal{G}_q .

The framework computes the contrastive loss for entities across each modality graph of the two KGs. The final cross-modality loss is then deter-

mined by:

$$\mathcal{L}_{CM}^t = \sum_{p,q \in \mathcal{M}} \sum_{e_i \in \mathcal{E} \cup \mathcal{E}'} \beta_p \beta_q (-\log l_{CM}^t(e_i^p, e_i^q)), \quad (11)$$

where β_p and β_q are hyper-parameters associated with modality p and q , adjusting the learning rate for features of the corresponding modality.

3.4.2 Cross-KG Alignment Loss

Pre-aligned entity pairs serve as crucial supervision signals in multi-modal entity alignment. Our approach employs supervised contrastive learning to elucidate the alignment relations between entities from corresponding modality graphs of different KGs. We formalized this supervised contrastive learning process at step t as follows:

$$l_{ckg}^t(e_i^m, e_j^m) = \frac{\exp(h_i^m \cdot h_j^m / \tau)}{\exp(h_i^m \cdot h_j^m / \tau) + \sum_{e_k \in N_i^-} \exp(h_i^m \cdot h_k^m / \tau)}, \quad (12)$$

where N_i^- encompasses a set of in-batch negative samples for e_i^m .

Given the undirected nature of entity alignment, we compute the contrastive loss considering both directions across KGs, which is defined as:

$$\mathcal{L}_{CKG}^t = \sum_{m \in \mathcal{M}^+} \sum_{(e_i, e_j) \in \mathcal{S}} -\frac{1}{2} \log (l_{ckg}^t(e_i^m, e_j^m) + l_{ckg}^t(e_j^m, e_i^m)), \quad (13)$$

where \mathcal{M}^+ is the set of modalities, including \mathcal{M} and “joint” as a special modality that encapsulates the fusion features of multi-modal information. Note that the representations of modality “joint” is taken from \mathcal{H}_{joint}^t .

3.5 Alignment Prediction

The alignment of entities is executed sequentially, focusing on the similarity between their fused multi-modal embeddings. Specifically, for each entity e_i within the joint modality graph \mathcal{G}_{joint} that remains unaligned, we seek its counterpart, e_j , in the corresponding joint modality graph \mathcal{G}'_{joint} . The sequence in which entities are aligned follows a greedy strategy. The chosen e_j is the unaligned entity that exhibits the highest cosine similarity score with e_i .

4 Experiment

4.1 Experiment Setup

Datasets To evaluate the effectiveness of our proposed method, we utilized three publicly avail-

able MMEA datasets. **MMKG** (Liu et al., 2019) features two subsets extracted from Freebase, YAGO, and DBpedia. **Multi-OpenEA** (Sun et al., 2020b), augmented with entity images from Google search queries, includes two multilingual subsets and two monolingual subsets. **DBP15K** (Sun et al., 2017; Liu et al., 2021) comprises three subsets from DBpedia’s multilingual version. Seed alignments are designated for 20% of MMKG and Multi-OpenEA entity pairs and 30% for DBP15K entity pairs, aligning with proportions used in prior studies (Chen et al., 2020; Liu et al., 2021; Lin et al., 2022; Chen et al., 2023a,b).

Baselines We compare our approach against six leading multi-modal alignment methods: MMEA (Chen et al., 2020), EVA (Liu et al., 2021), MSNEA (Chen et al., 2022), MCLEA (Lin et al., 2022), MEAformer (Chen et al., 2023a), and UMAEA (Chen et al., 2023b). We replicated MEAformer and UMAEA using their publicly available code to establish robust baselines. Although other notable methods like ACK-MMEA (Li et al., 2023) exist, their exclusion was due to the lack of open-source code, ensuring fairness in our comparisons.

Evaluation Metrics We employ two metrics for evaluation. Hits@1 (abbreviated as H@1) measures the accuracy of top-one predictions. Mean Reciprocal Rank (MRR) assesses the average inverse rank of the correct entity. Higher values of H@1 and MRR indicate better performance.

Implement Details The total training epochs are set to 250, with an option for an additional 500 epochs using an iterative training strategy. Our training regimen incorporates a cosine warm-up schedule (15% step for LR warm-up), early stopping, and gradient accumulation, utilizing the AdamW optimizer with a consistent batch size of 3500. Details of the experimental setup are available in the Appendix A.

4.2 Overall Results

Table 1 shows a comprehensive comparison between iterative and non-iterative methods across the three datasets, showcasing our model outperforms all nine sub-datasets.

Our approach significantly surpasses other non-iterative methods across various datasets. Specifically, on DBP15K, we see an improvement in

Table 1: Comparison of overall performance presenting both non-iterative and iterative results. The highest-performing baseline results are underlined, and any instances where our method sets a new state-of-the-art are highlighted in bold. Results with an asterisk (*) indicate our reproduction under identical settings.

Model	MMKG				Multi-OpenEA								DBP15K					
	FBDB15K		FBYG15K		EN-FR-V1		EN-DE-V1		D-W-V1		D-W-V2		ZH-EN		JA-EN		FR-EN	
	H@1	MRR	H@1	MRR	H@1	MRR	H@1	MRR	H@1	MRR	H@1	MRR	H@1	MRR	H@1	MRR	H@1	MRR
MMEA	.265	.357	.234	.317	-	-	-	-	-	-	-	-	-	-	-	-	-	-
MSNEA	.114	.175	.103	.153	.692	.734	.753	.804	.800	.826	.838	.873	.609	.685	.541	.620	.557	.643
EVA	.199	.283	.153	.224	.785	.836	.922	.945	.858	.891	.890	.922	.683	.762	.669	.752	.686	.771
MCLEA	.295	.393	.254	.332	.819	.864	.939	.957	.881	.908	.928	.949	.726	.796	.719	.789	.719	.792
MEAformer*	.418	.519	.327	.418	.836	.882	.954	.971	.909	.933	.944	.962	.771	.835	.764	.834	.772	.841
UMAEA*	.454	.552	.355	.451	.847	.891	.955	.970	.905	.930	.948	.967	.800	.860	.801	.862	.818	.877
PMF	.539	.620	.459	.539	.912	.942	.973	.983	.955	.970	.981	.989	.835	.884	.835	.885	.850	.898
improve	8.5%	6.8%	10.4%	8.8%	6.5%	5.1%	1.8%	1.2%	4.6%	3.7%	3.3%	2.2%	3.5%	2.4%	3.4%	2.3%	3.2%	2.1%
MSNEA	.149	.232	.138	.210	.699	.742	.788	.835	.809	.836	.862	.894	.648	.728	.557	.643	.583	.672
EVA	.231	.318	.188	.260	.849	.896	.956	.968	.915	.942	.925	.951	.750	.810	.741	.807	.765	.831
MCLEA	.395	.487	.322	.400	.888	.924	.969	.979	.944	.963	.969	.982	.811	.865	.805	.863	.808	.867
MEAformer*	.578	.661	.444	.529	.903	.935	.963	.977	.954	.970	.969	.981	.847	.892	.842	.892	.845	.894
UMAEA*	.561	.660	.463	.560	.895	.931	.974	.984	.945	.965	.973	.984	.856	.900	.857	.904	.873	.917
PMF	.624	.702	.543	.620	.923	.950	.980	.988	.960	.973	.986	.992	.867	.908	.866	.909	.879	.921
improve	4.6%	4.1%	8.1%	6.0%	2.0%	1.5%	0.6%	0.4%	0.6%	0.3%	1.3%	0.8%	1.1%	0.8%	0.9%	0.5%	0.6%	0.4%

Table 2: Non-iterative model performance comparison utilizing surface information on DBP15K.

Model	DBP15K					
	ZH-EN		JA-EN		FR-EN	
	H@1	MRR	H@1	MRR	H@1	MRR
MSNEA	.929	.951	.964	.976	.990	.994
EVA	.887	.913	.938	.955	.969	.980
MCLEA	.926	.946	.961	.973	.987	.992
MEAformer*	.944	.962	.976	.985	.990	.994
UMAEA*	.945	.963	.978	.986	.990	.994
PMF	.958	.973	.985	.990	.992	.995

H@1 of up to 3.5% over the nearest competitor. The Multi-OpenEA datasets show enhancements in the range of 1.8% to 6.5%, while the FBDB15K and FBYG15K datasets experience substantial gains, with increases up to 10.4%. These improvements highlight our method’s capability to effectively filter out conflicting and lower-quality modality information, enhancing overall alignment accuracy.

To assess our model’s performance with added textual data, we initialized text modality embeddings using surface representations from (Lin et al., 2022). The incorporation of textual information markedly enhances entity alignment, as evidenced by the results in Table 2 for the DBP15K dataset. Despite all models benefiting from this supervisory signal, our method outperforms the baselines and demonstrates its robustness in leveraging textual data for improved alignment.

4.3 Ablation Study

To assess the contribution of various components within our framework, we introduce five variants: Feature Relevance Measuring (FRM), Irrelevant Feature Freezing (IFF), Relevant Feature




Figure 3: Comparison of variants of PMF on DBP15K in terms of H@1.

Fusion (RFF), Progressive Integration (PI), and Cross-Modality Association Loss (\mathcal{L}_{CM}). By individually removing or modifying these components, their contributions to the performance were assessed. As depicted in Figure 3, the omission of the FRM component leads to a significant drop in H@1 performance across all datasets. For example, H@1 fell from 0.835 to 0.749 on DBP15K_{ZH-EN}. The influence of FRM extends to both IFF and RFF components, with RFF showing a more significant impact on performance due to its integral role in modality fusion and its effect on the overall training loss.

Further, replacing our "progressive" integration strategy with a "static" one, where FRM occurs only once, resulted in a marked performance downturn. This underscores the significance of progressively identifying alignment-irrelevant features, a task challenging to accomplish in a single iteration. Eliminating \mathcal{L}_{CM} also led to a decline in model performance. This underlines the effectiveness of cross-modality loss in fostering consistent entity representation learning across modalities.

4.4 Modality Analysis

4.4.1 Impact of Different Modality




Figure 4: Impact of various modalities on DBP15K in terms of H@1.

We examined the influence of modality information on entity alignment on DBP15K, as depicted in Figure 4. The absence of any modality notably hinders performance, with structure and image modalities showing the most significant effect. The foundational role of structures in KGs and the distinctive features provided by images underscore their importance in alignment.

A slight drop in performance is observed when either relation or attribute is excluded, attributed to the typically lower count of attributes and relations compared to entities, which may lead to attribute or relation overlaps among different entities. Conversely, a notable performance drop occurs when both are removed. This suggests a complementary relationship between attributes and relations that enhances alignment.

4.4.2 Distribution of Modality Scores

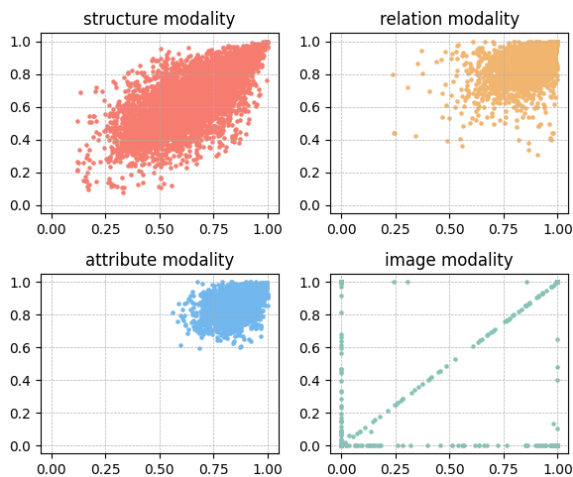


Figure 5: Distribution of relevance scores across modalities for pre-aligned entity pairs in DBP15K_{FR-EN}.

We examined the distribution of modality scores to understand the contribution of different modalities to entity alignment, as depicted

in Figure 5. Each point in the figure represents an aligned entity pair, with axes reflecting their modality scores across source and target KGs.

Attributes, relations, and structures show a widespread and increased distribution of scores, indicating their positive impact on alignment. Conversely, image modality shows a distinct, concentrated pattern, highlighting the variability in image features between aligned entities due to the limited scope of images and inconsistencies in KG construction. This underscores the benefit of excluding less reliable image features. Moreover, the presence of data points near the axes suggests challenges with one-sided visual information. Interestingly, a strong positive correlation emerges between modality scores of aligned entities, validating our modality scoring approach. This correlation implies that when a modality is less helpful for one entity in an aligned pair, it tends to be similarly low for the counterpart.

4.5 Parameters Analysis




Figure 6: Impact of varying threshold δ on frozen ratio and H@1 performance: the x-axis represents δ values ranging from 0.1 to 0.95.

To further investigate the impact of parameters on integration, we analyze how varying the threshold parameter δ impacts the frozen ratio in image modality and alignment performance on DBP15K. Figure 6 shows that when δ is below 0.4, minimal changes are observed in both the ratio and H@1 performance, likely due to the absence of images for certain entities. Beyond this point, as δ ascends, we note a gradual increase in the frozen ratio, paralleled by an improvement in H@1, despite the initial decline. This underscores the benefits of selectively and gradually freezing entities. A notable drop in H@1 performance for some datasets occurs at $\delta = 0.95$, indicating that images are beneficial for alignment and should not be frozen excessively.

Further experimental analysis results of the

Table 3: Case study of representative frozen images on DBP15K_{FR-EN} during training.

	Source-KG(French)		Target-KG(English)	
	Entity	Image	Entity	Image
Early frozen	Louis XV		Louis XV of France	
	Président du Mexique		President of Mexico	
	Université Stanford		Stanford University	
	Disque compact		Compact disc	
Late frozen	Juliana of the Netherlands		Juliana of the Netherlands	
	Brampton (Ontario)		Brampton	
	STS-129		STS-129	
	Queens of the Stone Age		Queens of the Stone Age	

feature-freezing process of PMF, including additional experiments, are available in Appendix B.

4.6 Case Study

To illustrate differences in features frozen at different stages, in Table 3, we select frozen image information during training in the DBP15K_{FR-EN} dataset arranged in a time series. Early frozen image pairs show significant differences. For example, the "President_of_Mexico" entity corresponds to a presidential portrait in the DBP_{FR} but to a national flag image in the DBP_{EN}. Images will not benefit the process of alignment in this case. After several epochs of training, the differences between frozen image pairs frozen later are smaller. For example, the "Juliana_of_the_Netherlands" entity, where the image in DBP_{FR} is a photo of the queen in her youth, while in DBP_{EN}, it is a photo of the queen in her old age. These cases reflect the process of effectively detecting alignment-irrelevant features based on the principle of progressing from easy to difficult cases.

5 Conclusion

In this study, we presented the Progressive Modality Freezing (PMF) model to advance Multi-Modal Entity Alignment. By measuring and evaluating the relevance of various modalities, PMF progressively freezes features deemed less critical, thereby facilitating the integration and consistency of multi-modal features. Furthermore, we introduced a unified training objective tailored to foster a harmonious contrast between KGs and modalities. Empirical evaluations on nine sub-datasets affirm the superiority of PMF and validate the rationale behind the selective freezing of modalities.

Limitations

PMF marks a significant advancement in MMEA, yet it encounters limitations in two main aspects.

First, the method for measuring the alignment relevance of modal features by evaluating their usefulness in identifying similar traits across KGs sometimes shows inconsistent reliability. As shown in Figure 6, a minor decline in H@1 when δ is below 0.5 suggesting that essential modal information may be prematurely frozen. Future work should aim to refine the approach for scoring alignment relevance, ensuring a more consistent and accurate identification of pertinent features.

Second, our model predominantly exhibits efficiency in non-iterative learning settings, with its performance often unstable under iterative learning frameworks. This limitation stems from the model's initial design, which is inherently optimized for non-iterative settings, limiting the effectiveness of incorporating iterative strategies like pseudo seeds for enhancing progressive freezing. Further attempts will explore improved ways to integrate iterative learning with progressive freezing for enhanced results.

Acknowledgments

This work was supported by the National Science and Technology Major Project under Grant 2022ZD0120202, in part by the National Natural Science Foundation of China (No. U23B2056), in part by the Fundamental Research Funds for the Central Universities, in part by the State Key Laboratory of Complex & Critical Software Environment, and in part by the Shanxi Province Special Support for Science and Technology Cooperation and Exchange (202204041101020).

References

- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.
- Liyi Chen, Zhi Li, Yijun Wang, Tong Xu, Zhefeng Wang, and Enhong Chen. 2020. Mmea: entity alignment for multi-modal knowledge graph. In *Knowledge Science, Engineering and Management: 13th International Conference, KSEM 2020, Hangzhou, China, August 28–30, 2020, Proceedings, Part I 13*, 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 *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 118–126.
- Zhuo Chen, Jiaoyan Chen, Wen Zhang, Lingbing Guo, Yin Fang, Yufeng Huang, Yichi Zhang, Yuxia Geng, Jeff Z Pan, Wenting Song, et al. 2023a. Meaformer: Multi-modal entity alignment transformer for meta modality hybrid. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3317–3327.
- 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 *International Semantic Web Conference*, pages 121–139. Springer.
- Laura Dietz, Alexander Kotov, and Edgar Meij. 2018. Utilizing knowledge graphs for text-centric information retrieval. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 1387–1390.
- Quan Fang, Xiaowei Zhang, Jun Hu, Xian Wu, and Changsheng Xu. 2022. Contrastive multi-modal knowledge graph representation learning. *IEEE Transactions on Knowledge and Data Engineering*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Qian Li, Shu Guo, Yangyifei Luo, Cheng Ji, Lihong Wang, Jiawei Sheng, and Jianxin Li. 2023. Attribute-consistent knowledge graph representation learning for multi-modal entity alignment. *arXiv preprint arXiv:2304.01563*.
- Zhenxi Lin, Ziheng Zhang, Meng Wang, Yinghui Shi, Xian Wu, and Yefeng Zheng. 2022. Multi-modal contrastive representation learning for entity alignment. *arXiv preprint arXiv:2209.00891*.
- Fangyu Liu, Muhao Chen, Dan Roth, and Nigel Collier. 2021. Visual pivoting for (unsupervised) entity alignment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 4257–4266.
- Ye Liu, Hui Li, Alberto Garcia-Duran, Mathias Niepert, Daniel Onoro-Rubio, and David S Rosenblum. 2019. Mmkg: multi-modal knowledge graphs. In *The Semantic Web: 16th International Conference, ESWC 2019, Portorož, Slovenia, June 2–6, 2019, Proceedings 16*, pages 459–474. Springer.
- Xin Mao, Wenting Wang, Yuanbin Wu, and Man Lan. 2021. From alignment to assignment: Frustratingly simple unsupervised entity alignment. *arXiv preprint arXiv:2109.02363*.
- Hatem Mousselly-Sergieh, Teresa Botschen, Iryna Gurevych, and Stefan Roth. 2018. A multimodal translation-based approach for knowledge graph representation learning. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 225–234.
- Wenxin Ni, Qianqian Xu, Yangbangyan Jiang, Zongsheng Cao, Xiaochun Cao, and Qingming Huang. 2023. Psnea: Pseudo-siamese network for entity alignment between multi-modal knowledge graphs. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3489–3497.
- Pouya Pezeshkpour, Liyan Chen, and Sameer Singh. 2018. Embedding multimodal relational data for knowledge base completion. *arXiv preprint arXiv:1809.01341*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Rui Sun, Xuezhi Cao, Yan Zhao, Junchen Wan, Kun Zhou, Fuzheng Zhang, Zhongyuan Wang, and Kai Zheng. 2020a. Multi-modal knowledge graphs for recommender systems. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pages 1405–1414.
- Zequan Sun, Wei Hu, and Chengkai Li. 2017. Cross-lingual entity alignment via joint attribute-preserving embedding. In *The Semantic Web–ISWC 2017: 16th International Semantic Web Conference, Vienna, Austria, October 21–25, 2017, Proceedings, Part I 16*, pages 628–644. Springer.
- Zequan 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. *arXiv preprint arXiv:2003.07743*.

Ruobing Xie, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Image-embodied knowledge representation learning. *arXiv preprint arXiv:1609.07028*.

Baogui Xu, Chengjin Xu, and Bing Su. 2023. Cross-modal graph attention network for entity alignment. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3715–3723.

Zuoxi Yang. 2020. Biomedical information retrieval incorporating knowledge graph for explainable precision medicine. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2486–2486.

Yuke Zhu, Ce Zhang, Christopher Ré, and Li Fei-Fei. 2015. Building a large-scale multimodal knowledge base system for answering visual queries. *arXiv preprint arXiv:1507.05670*.

Appendix

A Detailed setup

A.1 Encoders of each modality

We incorporate four modalities, including structure, image, relationship, and attribute, to implement and evaluate the proposed method. The settings of encoders for each modality are as follows.

The modality of structures describes the topological structure of entities in the knowledge graph. Specifically, the input of structural modal is the connections between various entities in the knowledge graph. The connections can be described by the adjacency matrix of the knowledge graph. Let the adjacency matrix of \mathcal{G} be A , we use graph attention networks (GAT) to encode structural information. Therefore, we define:

$$h_i^{str} = \mathbf{GAT}(A, \mathbf{e}_i^{str}), \quad (14)$$

Where \mathbf{e}_i^{str} is the randomly initialized representation of the entity e_i .

The modality of relations consists of the relationships related to each entity in text form. The relational feature of entity e_i is the set of relation names from all the triples in which e_i occurs. Since the number of relations is limited, we use the bag-of-words method to process input features. \mathbf{e}_i^{rel} is defined as a one-hot vector, and the dimension is determined by the number of relations. ENC_{rel} consists of a fully connected layer.

Similarly, textual entity attributes constitute an independent modality. Since the number of attributes is also limited, we adopt a processing method similar to the modality of relations. Specifically, ENC_{att} is a fully connected layer, and the input \mathbf{e}_i^{att} is defined as a one-hot vector, and the dimension depends on the number of attributes.

The image modality is composed of the visual entity attributes. Encoding images is usually difficult and requires a large amount of images for pre-training. Therefore, we first use a fixed pre-trained vision model (Simonyan and Zisserman, 2014; He et al., 2016; Radford et al., 2021), to convert the original pixel information into a vector representation \mathbf{e}_i^{img} . Then, a trainable fully connected layer is used as ENC_{img} .

The encoder of relations, attributes, and images can be expressed by:

$$h_i^m = \mathbf{FC}_m(\mathbf{e}_i^m), \quad (15)$$

where \mathbf{FC}_m is a trainable fully connected layer, $m \in \{rel, attr, img\}$.

A.2 Dataset Statistics

Statistics for DBP15K, Multi-OpenEA, and MMKG are presented in Table 4. The column named "EA pairs" indicates pre-aligned entity pairs. Note that not every entity is paired with images or counterparts in the alternate KG.

A.3 Metric Details

We evaluate the performance of entity alignment tasks using two common metrics: Hits@n and MRR.

Hits@n is a widely used performance metric in entity alignment tasks. It is calculated as follows: Given a total number of entity pairs N , for each entity in graph \mathcal{G} , compute its similarity with potential matching entities in graph \mathcal{G}' and rank them according to their scores. If the correct entity's rank is within the top n positions, the Hits@n count is incremented. The value of Hits@n is the percentage of the Hits@n count out of the total number of entity pairs N . Higher Hits@n values indicate better alignment performance.

$$\text{Hits@n} = \frac{1}{N} \sum_{i=1}^N I(\text{rank}_i \leq n) \quad (16)$$

MRR (Mean Reciprocal Rank) is a general evaluation metric for search algorithms and, when ap-

Table 4: Statistical details for MMEA datasets.

Dataset	KG	# Ent.	# Rel.	# Attr.	# Rel. Triples	# Attr. Triples	# Image	# EA pairs
DBP15K _{ZH-EN}	ZH (Chinese)	19388	1701	8111	70414	248035	15912	15000
	EN (English)	19572	1323	7173	95142	343218	14125	
DBP15K _{JA-EN}	JA (Japanese)	19814	1299	5882	77214	248991	12739	15000
	EN (English)	19780	1153	6066	93484	320616	13741	
DBP15K _{FR-EN}	FR (French)	19661	903	4547	105998	273825	14174	15000
	EN (English)	19993	1208	6422	115722	351094	13858	
Multi-OpenEA _{EN-FR}	EN (English)	15000	267	308	47334	73121	15000	15000
	FR (French)	15000	210	404	40864	67167	15000	
Multi-OpenEA _{EN-DE}	EN (English)	15000	215	286	47676	83755	15000	15000
	DE (German)	15000	131	194	50419	156150	15000	
Multi-OpenEA _{DW-V1}	DBpedia	15000	248	342	38265	68258	15000	15000
	Wikidata	15000	169	649	42746	138246	15000	
Multi-OpenEA _{DW-V2}	DBpedia	15000	167	175	73983	66813	15000	15000
	Wikidata	15000	121	457	83365	175686	15000	
FBDB15K	FB15K	14951	1345	116	592213	29395	13444	12846
	DB15K	12842	279	225	89197	48080	12837	
FBYG15K	FB15K	14951	1345	116	592213	29395	13444	11199
	YAGO15K	15404	32	7	122886	23532	11194	

plied to entity alignment tasks, calculates the similarity between each entity in graph \mathcal{G} and potential matching entities in graph \mathcal{G} , ranking them according to their scores. If the correct entity is ranked n , the score is $\frac{1}{n}$. The sum of scores for all entities yields the MRR score. MRR reflects the overall performance of the entity alignment algorithm, with higher MRR scores indicating better alignment results.

$$MRR = \frac{1}{N} \sum_{i=1}^N \frac{1}{rank_i} \quad (17)$$

A.4 Implementation Details

For a fair comparison with recent works (Chen et al., 2023a,b), we have standardized our experimental setup as follows:

- Iterative Learning:** We adopt the approach of (Lin et al., 2022) and employ a probation method for iterative training. This method collects pairs of mutual nearest entities every 5 epochs. Pairs that consistently remain nearest neighbors for 10 rounds are added to the training set.
- Multi-modal Encoders:** For relations and attributes, we use the Bag-of-Words (BoW) model to encode relations (x_r) and attributes (x_a) as fixed-length vectors with a dimension of $d_r = 1000$. For entity names’ surface information, we use 300-dimensional GloVe vectors and character bigrams. We enhance this method by applying machine translations to entity names, following (Mao et al., 2021). For images, we use ResNet-152 (He et al., 2016) as pre-trained visual encoders with a vision feature dimension (d_v) of 2048 in DBP15K. For Multi-OpenEA, CLIP (Radford et al., 2021) is used with $d_v = 512$. In FBDB15K/FBYG15K, VGG-16 (Simonyan and Zisserman, 2014) is employed with $d_v = 4096$. It should be noted that we assign a zero vector as the embedding for missing images to minimize their impact on cross-modal learning.
- Hyper Parameters:** We unify the hidden layer dimensions across all networks to 300. The parameter τ is set to 0.05 to emphasize challenging negative samples in the contrast loss. The dynamic threshold δ in Eq. 3 for progressively freezing entities begins at 0.1 and increases by 1.2 until it reaches a maximum of 0.9. For the structure, relation, attribute, and image modalities, β_m values in Eq. 11 are respectively set to 0.1, 0.1, 0.1, and 10, scaling the learning rate of modality features.
- Computational Overhead:** Our experiments, conducted on a Tesla V100 SXM2 32GB GPU, demonstrate the superior efficiency of our model, which completes training in only 20 minutes. This is notably faster than the baseline model, MEAformer (Chen et al., 2023a), at 33 minutes, and UMAEA (Chen et al., 2023b) at 29 minutes. The number of learnable parameters

is 13.1M for DBP15K and Multi-OpenEA, and 9.9M for MMKG.

B Detailed analysis

B.1 Low Resource




Figure 7: H@1 performance with different ratios of seed alignments ranging from 0.01 to 0.3 on DPB15K_{ZH-EN} (left) and FBDB15K (right). The x-axis denotes ratios and the y-axis denotes H@1.

To assess the stability of the proposed method under conditions of limited seed alignments, we conducted evaluations on two distinct datasets—DBP15K_{ZH-EN} and FBDB15K, using seed alignment ratios varying from 0.01 to 0.30. Figure 7 illustrates a clear gap between performances as the ratio escalates. Notably, with merely 1% seed alignments in the DBP15K_{ZH-EN} dataset, our model was able to achieve a H@1 score of .648, markedly outperforming the MEAformer, which scored .456. This result highlights the model’s robustness and its considerable promise for applications in few-shot entity alignment.

B.2 Analysis of Freezing Process




Figure 8: The ratio of entities with images that are frozen during the training epochs in DBP15K.

To delve deeper into the impacts of the progressive manner, we take the freezing process of im-

age modality on the DBP15K dataset as an example, analyzing the trend of the proportion of entities being frozen as training epochs change. Figure 8 indicates that as training epochs increase, the proportion of entities frozen shows an exponential rising trend, ultimately stabilizing at about 55%. The high rate of frozen images is partly due to a high image missing rate of up to 30% in the DBP15K dataset, and it also reflects high ambiguity in image information within multi-modal knowledge graphs.