

A Dual-Way Enhanced Framework from Text Matching Point of View for Multimodal Entity Linking

Shezheng Song¹, Shan Zhao^{*1}, Chengyu Wang², Tianwei Yan²
Shasha Li^{*2}, Xiaoguang Mao², Meng Wang¹

¹Hefei University of Technology

²National University of Defense Technology

bettterszsong@gmail.com, zhaoshan@hfut.edu.cn, shashali@nudt.edu.cn

Abstract

Multimodal Entity Linking (MEL) aims at linking ambiguous mentions with multimodal information to entity in Knowledge Graph (KG) such as Wikipedia, which plays a key role in many applications. However, existing methods suffer from shortcomings, including modality impurity such as noise in raw image and ambiguous textual entity representation, which puts obstacles to MEL. We formulate multimodal entity linking as a neural text matching problem where each multimodal information (text and image) is treated as a query, and the model learns the mapping from each query to the relevant entity from candidate entities. This paper introduces a dual-way enhanced (DWE) framework for MEL: (1) our model refines queries with multimodal data and addresses semantic gaps using cross-modal enhancers between text and image information. Besides, DWE innovatively leverages fine-grained image attributes, including facial characteristic and scene feature, to enhance and refine visual features. (2) By using Wikipedia descriptions, DWE enriches entity semantics and obtains more comprehensive textual representation, which reduces between textual representation and the entities in KG. Extensive experiments on three public benchmarks demonstrate that our method achieves state-of-the-art (SOTA) performance, indicating the superiority of our model. The code is released on <https://github.com/season1blue/DWE>.

Introduction

Entity linking (EL) (Shen, Wang, and Han 2014; Ling, Singh, and Weld 2015) has attracted increasing attention in the natural language processing (NLP) domain, which aims at linking ambiguous mentions to the referent unambiguous entities in a given knowledge base (KB) (Auer et al. 2007). It is crucial in many information retrieval applications, such as information extraction (Al-Moslmi et al. 2020), question answering (Lu et al. 2016a), and Web query (Michel et al. 2020). However, the ambiguity in mentions poses great challenges to this task. Hence, Multimodal Entity Linking (MEL) has been introduced recently (Sun et al. 2022a; Gan et al. 2021), where multimodal information (e.g., visual information) is used to disambiguate the mention. Figure 1 shows an example of MEL, wherein the mention *Trump* is associated with the text “*Trump and his wife Melania at the*

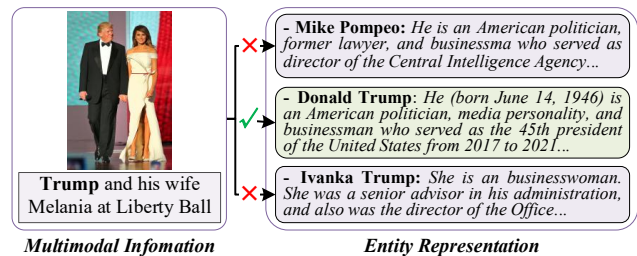


Figure 1: Example of entity linking for the mention *Trump*.

Liberty Ball”. It is challenging to distinguish the entities because the short text is ambiguous, and several entities are related to *Trump*. Thus, image is required for further disambiguation. Dozens of MEL studies have been conducted in the past few years and achieved promising performances, where delicate interaction and fusion mechanisms are designed for encoding the image and text features (Wang 2022; Zhou et al. 2021; Chen et al. 2022b). Nonetheless, existing approaches still struggle to fully exploit the potential of the feature sources from two distinct information perspectives, which may hinder further task development.

Modality impurity: In MEL task, one of the challenges lies in the fact that the image modality often contains less information and more redundant features when compared to the text modality. Indeed, not all the visual sources play positive roles. As revealed by Vempala and Preoțiu-Pietro, as high as 33.8% of visual information serves no context or even noise in task. As shown in Figure 1, only the left person is relevant while the right person is noise, causing redundant visual elements and low information density for the entity. However, for the existing MEL works, they may suffer from ineffective and insufficient modeling of visual pivot features. The majority of MEL models perform vision-language grounding by simultaneously considering the entire image and text (Zhang et al. 2019; Fei et al. 2023). Unfortunately, such coarse-grained representation learning often leads to information mismatches and compromises the subtle semantic correlations between vision and language. We argue that a fine-grained feature complementing between text features and visual information is needed.

Ambiguous Entity Representation: Wang et al. (2022)

*Corresponding Author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

and Zhou et al. (2021) concatenate the main properties of an entity into a piece of text, which is regarded as a representation of the entity in KG. Sun et al.(2022b) manually collect short text to represent entity, e.g., “Bahador Abb: Iranian footballer”. However, the set of properties or short text are insufficient to represent entities. Different people may have the same properties. For example, the entity *Donald Trump* and *Mike Pompeo* both have the properties of “*Sex: male. Religion: Presbyterianism. Occupation: businessperson, politician. Languages: English. work location: Washington, D.C.*” With so similar entity representation, it is a great challenge to link entity even with an ideal multimodal method. Thus, more distinctive and representative entity representation is required for further disambiguation. Fortunately, Wikipedia description offers a promising solution, which has been shown to enrich the semantics of raw data, especially in low-resources NLP tasks (Chen et al. 2022a).

To tackle the above issues, in this work, we propose a novel framework to improve MEL. Inspired by neural text matching task, we view each mention as a query and attempt to learn the mapping from each mention to the relevant entity. To achieve this, we design a dual-way enhanced (DWE) framework: enhancing query with refined multimodal information; enriching the semantics of entity by Wikipedia description as shown in Figure 2. Specifically, (1) for the enhancing query (e.g., mention), DWE employs a pretrained visual encoder and visual tools to obtain the image representation and fine-grained visual attributes, then purify it into the learnable visual characteristics. The visual characteristics and text are regarded as two types of refined information, which are separately employed to enhance the mention. Considering the semantic gap between vision and text in the language model, we design three enhanced units based on the cross-modal enhancer to bridge the semantic gap between (i) facial feature and object visual feature (ii) mention and vision-enhanced feature and (iii) mention and text-enhanced feature. By introducing the cross-modal alignment, our model makes use of the semantics of text and objective image information. (2) On the other hand, we utilize a knowledge base (i.e. Wikipedia) to enrich the semantics of entity representation. This process reduces the disparity between textual representation and the entities in the knowledge graph, resulting in a more comprehensive textual representation. To sum up, we enhance entity representations for 17805, 17291, and 57007 entities in the Richpedia, WikiMEL, and Wikidiverse datasets, respectively.

Extensive experiments show that DWE pushes the state-of-the-art (SOTA) by a large margin, indicating the efficiency of DWE. Our contributions are as follows:

- We explore Multimodal Entity Linking in a neural matching formulation. We design a dual-way enhanced framework: query is enhanced by refined multimodal information and we utilize the Wikipedia to enrich the semantics of entity representation.
- We innovatively introduce fine-grained image attributes (such as facial characteristics and scene feature) which are leveraged to enhance and refine the visual features and take cross-modal alignment to bridge the semantic

gap between textual and visual features.

- Experimental results on three benchmark datasets Richpedia, WikiMEL and Wikidiverse demonstrate the superiority of our proposed model over the SOTA method.

Related Work

MEL has been widely studied in recent years, which leverages the associated multimodal information to better link mention to entity in KG. Previous studies mainly focus on the following two aspects:

Multimodal Information Refinement Recently, dozens of works have focused on utilizing multimodal information for disambiguation. Wang et al. (2022) propose to use gated multimodal feature fusion and contrastive learning for MEL. Wang et al. (2022) use ResNet(He et al. 2016) as a visual encoder and BERT(Devlin et al. 2019) as a textual encoder to extract the joint feature for MEL. However, these studies (Zhao et al. 2023, 2021b,a) ignore the noise of the image, i.e., some objects in the image are irrelevant to the mention. Chen et al. (2022b) notice the visual and textual noise and take image segmentation and attention mechanism as refinement method, but ignore the global relationship between text and image. Wang et al. (2021) also realize the noise, but they pay too much attention to extracting text such as OCR text or caption from raw image. To conclude, these work (Ma et al. 2023, 2022) can not fully use multimodal information and filter the irrelevant part (i.e., noise) of raw image effectively.

Ambiguous Entity Representation It is still controversy over how to represent the entity in KG, so there are different methods of entity representation. Adjali et al. (2020a) propose a framework for automatically building the MEL dataset from Twitter(Wagh and Punde 2018), taking profile on Twitter as the user entity representation. However, the profile is edited by user and can not fully represent the user entity. Zhou et al. (2021) propose three MEL datasets, built from Weibo, Wikipedia, and Richpedia information, collecting the property of entity as representation, such as “Donald Trump: male, 1946, politician” and using Wikidata (Vrandečić 2012) and Richpedia (Wang et al. 2020) as the corresponding KGs. Similarly, Wikiperson(Sun et al. 2022b) dataset takes a manually organized short text as entity representation such as “Bahador Abdi: Iranian footballer”. However, the entity representation of the above datasets is ambiguous, making it hard to distinguish the MEL target (entity) and affecting the plausibility of performance evaluation. Our work differs from the above papers in that we use the textual Wikipedia description as entity representation to guide the learning of multimodal feature, which are proven to be more efficient.

Methodology

Notation and Task Definition

Let $X = \{x_i\}_{i=1}^N$ be a set of N input multimodal samples, with corresponding entities set $Y = \{e_i\}_{i=1}^M$ in a KG. Each input is composed of three parts: $x = \{x_m; x_t; x_v\}$, where x_t denotes a text consisting of a sequence of tokens with length l_t , x_v is an image associated with the text, and x_m is a set of mentions in text x_t .

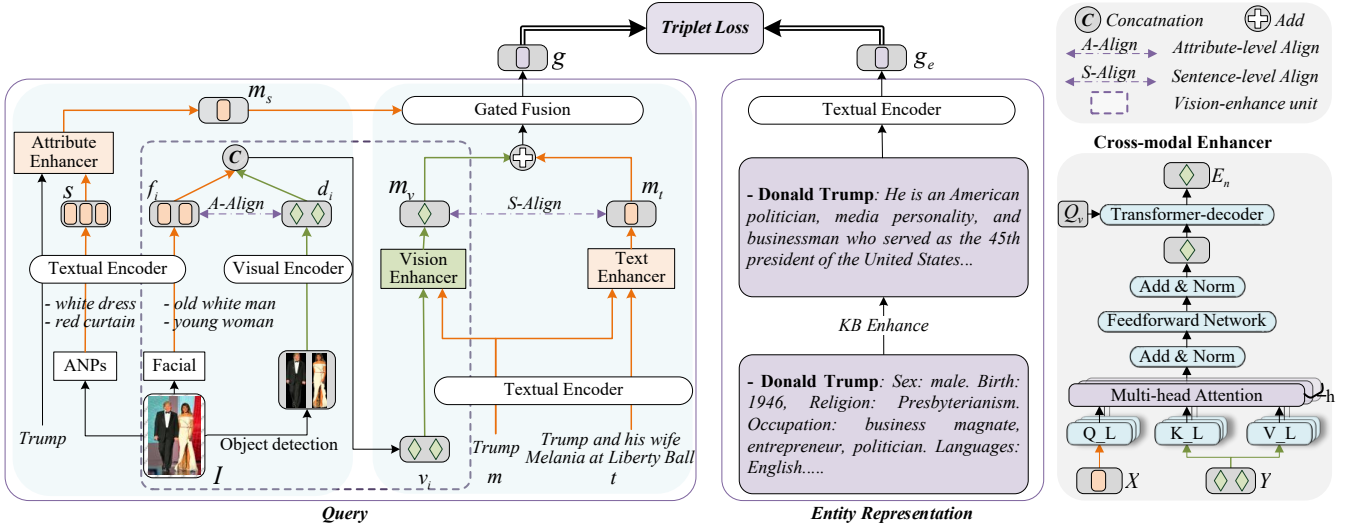


Figure 2: Overview of DWE. Input consists of image I , text t , and mention m . Object detection is applied to extract object feature d_i from image. Facial feature f and scene feature s (i.e. Adjective-Noun Pairs(ANPs)) is retrieved from image. The fusion feature g is combination of image-enhanced feature m_v , text-enhanced feature m_t and attribute-enhanced feature m_s .

A KG is composed of a set of entities $Y = \{e_i\}_{i=1}^M$. Each entity e_i is described by a sequence of entity representations $e_i = u_1, u_2, \dots, u_w$ (i.e., textual Wikipedia description about the entity). w is the token length of entity representation. The MEL task aims to link the ambiguous mention (i.e. query) to the entity in KG by calculating the similarity between the joint mention feature and textual entity representation. The entity with the highest similarity is regarded as the corresponding linking entity. In this process, textual and visual information is used to disambiguate and further clarify semantics. Finally, the KG entity y linked to mention is selected from λ candidate entities $C_e = \{e_i\}_{i=1}^\lambda$. It can be formulated as follows:

$$y = \underset{e \in C_e}{\operatorname{argmax}} \Gamma(\Phi(x), \Psi(e)) \quad (1)$$

where Φ represents the multimodal feature function, Ψ represents the KG entity representation feature function, and $\Gamma()$ produces a similarity score between the multimodal representation and entity representation in KG. y is the predicted linking entity.

In evaluation, we take cosine similarity to measure the correlation between the joint multimodal feature g and the candidate entities C_e . The entity with the highest similarity to g is regarded as the corresponding linking entity. For the example shown in Figure 1, The similarities between the mention and entities are 0.47, 0.81 and 0.13. Therefore, candidate entity *Donald Trump* is chosen as the linked entity.

Preliminary

Knowledge base Enhanced Entity Representation. Previous works (Wang 2022; Zhou et al. 2021) tend to concatenate the main properties of an entity into a piece of text and regard the properties as textual entity representation. However, a set of properties is not clear enough to distinguish the

linking entity from candidate entities. Different entity may be the same in some fields. As shown in the entity representation in Figure 2, the properties of *Mike Pompeo* is similar with *Donald Trump* in **previous** representation. Thus it is challenging to distinguish the two people with nearly the same properties.

Therefore, we leverage Wikipedia to get **enhanced** representation, with the textual Wikipedia description from the text on Wikipedia pages. In summary, 17805, 17391, and 57007 entity representation are enhanced on Richpedia, WikiMEL, and Wikidiverse datasets respectively.

Textual feature extraction. Given a sentence x_t , mention x_m and Wikipedia description of entity x_e , we follow CLIP (Radford et al. 2021) to tokenize it into a sequence of word embeddings. Then the special tokens *startoftext* and *endoftext* are added at the beginning and end positions of word embeddings. As a result, with N sentences and N_e candidate entities, we feed sentence representation $t \in \mathbb{R}^{N \times d}$, mention representation $m \in \mathbb{R}^{N \times d}$ and enhanced entity representation $g_e \in \mathbb{R}^{N_e \times d}$ into model.

Multimodal Information Enhanced Query

We first present a general architecture: cross-modal enhancer, as shown on the right of Figure 2. Then we build three enhanced units based on this architecture to capture dense interaction over multimodal information, namely a text-enhanced unit, an attribute-enhanced and a vision-enhanced unit.

Cross-modal enhancer. Given two types of features, X and Y , the enhanced features is calculated as:

$$E_n = E_n(X, Y) \quad (2)$$

Specifically, the calculation of multi-head attention for hidden state h_t is first performed on the hybrid key and value:

$$h_t = \text{Softmax}\left(\frac{(W^Q X)(W^K Y)^T}{\sqrt{d}}\right)(W^V Y) \quad (3)$$

Then we utilize the transformer decoder for decode hidden state h_t with the learnable textual query $Q_v \in \mathbb{R}^{N_v \times d}$:

$$E_n = \text{Decoder}(h_t, Q_v) \quad (4)$$

In detail, we formulate the attention of transformer decoders in Decoder as:

$$E_n = \text{Softmax}\left(\frac{(W^Q Q_v)(W^K h_t)^T}{\sqrt{d}}\right)(W^V h_t) \quad (5)$$

where $W^Q \in \mathbb{R}^{d \times d_q}$, $W^K \in \mathbb{R}^{d \times d_k}$, $W^V \in \mathbb{R}^{d \times d_v}$ are randomly initialized projection matrices. We set $d_q = d_k = d_v = d/h$. h is the number of heads of attention layer.

Text-enhanced unit. The textual-enhanced is designed to model sentence-to-mention correlations. It takes two groups of feature sentence representation $t \in \mathbb{R}^{N \times d}$, mention representation $m \in \mathbb{R}^{N \times d}$ as inputs. Next, the text-enhanced unit models pairwise relationship between each paired sample (m, t) and finally outputs enhanced mention features:

$$m_t = E_n(m, t) \quad (6)$$

Attribute-enhanced unit. Sentibank(Borth et al. 2013) is applied to extract adjective-noun pairs (ANPs) x_s such as “nice clouds, white dress, happy man.” from image, which is relevant to the attributes of scene feature. Similarly, we integrate and encode x_s to get attribute feature s . We apply cross-modal enhancer to get the attribute-enhanced mention feature m_s .

$$m_s = E_n(m, s) \quad (7)$$

where $s \in \mathbb{R}^{n_s \times d}$ and $m_s \in \mathbb{R}^{1 \times d}$, d indicates the hidden dimension and n_s is the number of adjective-noun pairs.

Vision-enhanced unit. Multimodal information in the MEL task generally consists of textual and visual information. Text is artificial, and thus its semantics is intensive. In contrast, visual information tends to be a photo of objects or scenes. There is inevitably some redundant information in the raw image. For an image with multiple mentions, only some interested objects in the image are related to one mention, while the rest part is related to the other mention. Some objects or scenes are still irrelevant and disturbing, even for an image with a mention. Therefore, refining the raw image by detecting and extracting the related visual object is necessary. As shown in Figure 2, the text is *Trump and his wife Melania at the Liberty Ball*, and the mention is *Trump*. Only the local region of the left person contributes to linking mention *Trump* and entity *Donald Trump* in Wikidata, while the rest of the image is of less significance.

An image I can be decomposed into l objects in local regions. Each mention only describes a part of the image. Therefore, we take an object detector D (Olafenwa 2021) to identify all possible objects in the image. The object visual feature d is extracted through the image encoder of CLIP.

$$x_d = \{x_d^i\}_{i=1}^l = D(I) \quad (8)$$

$$d = \text{clip}(x_d) \quad (9)$$

where $d = \{d_i\}_{i=1}^l$ and l is the number of the detected objects in the image.

Deepface¹ is applied to extract facial appearance x_f such as gender, race, age and etc, which are relevant to the attributes of facial feature in the image. x_f is integrated into prompt sentence such as “Trump, gender: male, race: white, age: 50...” and encoded through image encoder of CLIP to get facial feature $f = \{f_i\}_{i=1}^l$. Due to f_i and d_i is facial feature and local visual feature of i -th object of image, the visual feature of i -th object v_i is enhanced by $v_i = [d_i, f_i]$, where refined visual feature $v = \{v_i\}_{i=1}^l \in \mathbb{R}^{l \times d}$ and $[\cdot]$ denotes concatenation.

Only some objects in the image are related to the mention. Thus, we utilize the visual enhancer to determine related visual feature and obtain vision-enhanced mention feature m_v .

$$m_v = E_n(m, v) \quad (10)$$

Gated Fusion and Training Loss

we get the feature of mention m , text-enhanced mention m_t , vision-enhanced mention m_v and attribute-enhanced mention m_s and apply gated fusion to restrain noise in m_s .

$$g_m = m + m_t + m_v \quad (11)$$

$$g = (g_m + \alpha(m_s - g_m))W_g + b_g \quad (12)$$

where $W_g \in \mathbb{R}^{2d \times d_m}$, d is hidden dimension, d_m is the dimension of joint feature g . b_g and α is learnable parameters.

During training, we attempt to maximize the similarity between the mention and its corresponding entity (positive sample θ_+) in KG and minimize the similarity between the mention and other entities (negative sample θ_-). Thus we take triplet loss(Koch, Zemel, and Salakhutdinov 2015) to maximize the distance between multimodal feature g and negative samples while minimizing the distance between g and positive samples. The triplet loss \mathcal{L}_t is defined as:

$$\mathcal{L}_t = \max(\Gamma(g, \theta_+) - \Gamma(g, \theta_-) + \varepsilon, 0) \quad (13)$$

Γ is cosine similarity between features. ε denotes margin, a hyperparameter for learning.

Cross-modal alignment. To bridge the semantic gap between features, we design (1)attribute-level alignment for facial feature f and object visual feature d , and (2)sentence-level alignment for text-enhanced feature m_t and vision-enhanced feature m_v

$$\mathcal{L}_c = L_m(f, d) + L_m(m_t, m_v) \quad (14)$$

where L_m is Mean-Shifted Contrastive Loss (MSC) (Reiss and Hoshen 2021).

The overall loss is the combination of triplet loss and crossmodal alignment loss. β is the coefficient factor for balancing the losses of two tasks: $\mathcal{L} = \mathcal{L}_t + \beta \cdot \mathcal{L}_c$.

Experiments

Datasets and Settings

We take three public datasets: Richpedia(Wang et al. 2020), WikiMEL(Zhou et al. 2021), and Wikidiverse(Wang et al. 2022) as benchmark. Table 2 shows the statistics of datasets. Besides, we apply enhanced entity representation on three datasets and reproduce several baselines for comparison.

¹<https://github.com/serengil/deepface>

Models	Richpedia				WikiMEL				Wikidiverse			
	Top-1	Top-5	Top-10	Top-20	Top-1	Top-5	Top-10	Top-20	Top-1	Top-5	Top-10	Top-20
BLINK	30.8	38.8	44.5	53.6	30.8	44.6	56.7	66.4	-	71.2	-	-
DZMNED	29.5	41.6	45.8	55.2	30.9	50.7	56.9	65.1	-	39.1	-	-
JMEL	29.6	42.3	46.6	54.1	31.3	49.4	57.9	64.8	21.9	54.5	69.9	76.3
BERT	31.6	42.0	47.6	57.3	31.7	48.8	57.8	70.3	22.2	53.8	69.8	82.8
ARNN	31.2	39.3	45.9	54.5	32.0	45.8	56.6	65.0	22.4	50.5	68.4	76.6
MEL-HI	34.9	43.1	50.6	58.4	38.7	55.1	65.2	75.7	27.1	60.7	78.7	89.2
HicCoAtt	37.2	46.8	54.2	62.4	40.5	57.6	69.6	78.6	28.4	63.5	84.0	92.6
GHMFC	38.7	50.9	58.5	66.7	43.6	64.0	74.4	85.8	-	-	-	-
MMEL	-	-	-	-	71.5	91.7	96.3	98.0	-	-	-	-
CLIP	60.4	96.1	98.3	99.2	36.1	81.3	92.8	98.3	42.4	80.5	91.7	96.6
CLIP-text [†]	59.2	91.1	94.6	97.4	50.2	87.4	93.5	97.2	35.8	79.6	89.1	95.3
CLIP [†]	63.9	96.2	98.3	99.2	58.0	95.9	98.3	99.5	43.1	80.2	90.7	96.2
DWE	67.6	97.1	98.6	99.5	44.7	65.9	80.8	93.2	47.5	81.3	92.0	96.9
STDEV	±0.74	±0.11	±0.07	±0.05	±0.08	±0.15	±0.18	±0.20	±0.11	±0.15	±0.29	±0.23
DWE [†]	72.5	97.3	98.8	99.6	72.8	97.5	98.9	99.7	51.2	91.0	96.3	98.9
STDEV	±0.28	0.06	±0.05	±0.07	±0.63	±0.06	±0.05	±0.05	±1.04	±0.23	±0.11	±0.07

Table 1: MEL results of the compared models in terms of Top-1, 5, 10, 20 accuracy (%) from λ candidate entities, where $\lambda = 100$. For the model, the results with [†] are using enhanced entity representations (textual Wikipedia demonstration), whereas the others represent the results using previous entity representation. We not only report the best performance but also the standard deviation (STDEV) over 10 runs with random initialization. The best results are highlighted in bold.

Dataset	Sample	Entity	Mention	Text length
Richpedia	17805	17804	18752	13.6
WikiMEL	18880	17391	25846	8.2
Wikidiverse	13765	57007	16097	10.1

Table 2: The statistics of three multimodal datasets

Metrics We use Top-k accuracy as the metric:

$$Acc_{top-k} = \frac{1}{N} \sum_{i=1}^N \eta\{I(\cos(g, gt), \cos(g, C_e)) \leq k\} \quad (15)$$

where N represents the total number of samples, and η is the indicator function. When the receiving condition is satisfied, η is set to 1, and 0 otherwise. gt is ground truth entity while C_e is a set of candidate entities. \cos means cosine similarity function. I is function to calculate the rank of similarity between joint feature g and ground truth gt among a set of candidate entities C_e .

Hyperparameters and Training Details During the experiments, the dimension of text representation t and visual representation v are set to 512, the dimension of related visual object feature R is set to 768, and the heads of Multi-head attention is set to 8. The dropout, training epoch, evaluating steps, weight decay, and triplet loss interval are set to 0.4, 300, 2000, 0.001, and 0.5, respectively. We optimize the parameters using AdamW (Loshchilov and Hutter 2017) optimizer with a batch size of 64, the learning rate of 5×10^{-5} and coefficient factor β of 0.5. All the experiments is processed on RTX3090 and Pytorch 2.0.

Candidate Retrieval Following previous work (Zhou et al. 2021; Wang et al. 2022), we set the number λ of candidate entities to 100. The strategy of candidate entity selection varies on different datasets. In Richpedia and WikiMEL, we apply $fuzz^2$ to search 100 candidate entities whose name is similar with mention. In Wikidiverse, the dataset provides 10 similar entities. We first divide entities into different collections according to entity types, such as person, location, and so on. According to mention type, we use $fuzz$ to search 90 entities with names similar to mention.

Negative Samples We employ two negative sampling strategies for triplet loss: hard negatives and in-batch negatives. The hard negatives are the candidate entities retrieved in the Candidate Retrieval step except for the gold entity. The in-batch negatives are gold entities of other $B - 1$ mentions in the batch, where B is batch size.

Baseline

we conduct comparative experiments to evaluate the effectiveness of our model on the previous entity representation (combination of properties) for comparison. Furthermore, previous entity is ambiguous and is not representative for entity. Thus, we propose the enhanced entity representation (textual Wikipedia description) and conduct extensive experiments on it.

(1) *BLINK* (Wu et al. 2019) is an BERT-based EL model with a two-stage zero-shot linking. (2) *BERT* (Devlin et al. 2019) stacks several layers of transformer to encode each token in text. (3) *ARNN* (Eshel et al. 2017) utilizes Attention-RNN to predict associations with candidate entity textual features. (4) *DZMNED* (Moon, Neves, and Carvalho 2018)

²<https://github.com/seatgeek/fuzzywuzzy>

	m_t	m_v	\mathcal{L}_c	m_s	f	Top-1	Top-5	Top-10	Top-20
0	✓	✓	✓	✓	✓	72.5	97.3	98.8	99.6
1	✓	✓	✓	✓	✗	70.9	97.5	98.8	99.5
2	✓	✓	✓	✗	✗	69.8	97.3	98.4	99.4
3	✓	✓	✗	✗	✗	67.3	93.9	98.4	99.4
4	✓	✗	✗	✗	✗	63.5	92.7	96.0	98.3

Table 3: Ablation study on Richpedia (m_t : textual enhanced feature; m_v : visual enhanced feature; \mathcal{L}_c : crossmodal alignment; m_s : attribute-enhanced mention; f : facial characteristics). Method 0 represents our model.

Features	Top-1	Top-5	Top-10	Top-20
primitive on rich	66.7	97.2	98.8	99.6
refined on rich	72.5	97.3	98.8	99.6
primitive on wiki	64.0	96.2	98.4	99.5
refined on wiki	72.8	97.5	98.9	99.7

Table 4: Ablation study on using refined visual feature v or primitive visual feature which is extracted directly from image over public datasets: Richpedia (rich), WikiMEL (wiki)

takes concatenated multimodal attention mechanism to fuse visual, textual, and character features of mentions. (5) *JMEL* (Adjali et al. 2020b) utilizes fully connected layers to project the visual and textual features into an implicit joint space. (6) *MEL-HI* (Zhang, Li, and Yang 2021) uses multiple attention to get richer information and refine the negative impact of noisy image. (7) *HieCoAtt* (Lu et al. 2016b) is a multimodal fusion mechanism, using alternating co-attention and three textual levels (tokens, phrases, and sentences) to calculate co-attention maps. (8) *GHMFC* (Wang 2022) is a advanced baseline proposed by Wang et.al (Wang 2022). It takes the gated multimodal fusion and novel attention mechanism to link multimodal entities. (9) *MMEL* (Yang et al. 2023) is a joint feature extraction module to learn the representations of context and entity candidates, from both the visual and textual perspectives. (10) *CLIP-text* (Radford et al. 2021) only use textual information such as text and mention in the dataset. It mainly focuses on the ability to understand textual relationship among text, mention, and entity representation. (11) *CLIP* (Radford et al. 2021) take both textual and visual features into consideration. The model concatenates multimodal features and calculates the similarity between features and ground-truth.

Main Results

On Richpedia(Wang et al. 2020), WikiMEL(Zhou et al. 2021) and Wikidiverse(Wang et al. 2022) datasets, we compare the proposed model with several competing approaches and show the results in Table ?? under different entity representation. Under previous entity representation, our method could reach 67.6%, 44.7%, and 47.5% on three datasets. Besides, under enhanced entity representation, our method also

	WikiMEL	Wikidiverse
GEMEL	75.5	82.7
DWE(ours)	81.0	87.1

Table 5: Top-1 accuracy for the model from 16 candidates.

achieves state-of-the-art (SOTA) performance by obtaining Top-1 accuracy of 72.5%, 72.8%, and 51.2%, respectively, which demonstrates the superiority over the previous state-of-the-art method. Apparently, with pre-trained on massive multimodal data, the text and image encoder of clip facilitates the understanding of complicated multimodal information. For the gated fusion, it benefits models by restraining noise while using information. Besides, the enhancer for text and vision contributes to fusion between features.

Enhanced Entity Representation. Previous works (Zhou et al. 2021; Wang 2022) consider the properties as entity representation, which is ambiguous and limited. we take textual Wikipedia description as enhanced representation. Results show that our method is not only useful for previous entity representation but also the enhanced representation. Besides, on CLIP and our method, the improvement over different representations shows that the enhanced representation is helpful for disambiguation. In detail, on Richpedia, by using enhanced representation, CLIP obtain the 3.5% improvement and our method obtains 4.9% improvement. The boost proves that the enhanced entity representation makes the entity more distinctive.

Comparison with the SOTA Generative Method. To investigate the comprehensiveness of the advantages of our model, we compare our model with the SOTA generative model, GEMEL (Shi et al. 2023) on accuracy. Smaller values of the quantity of candidate entity λ make the task of selecting link entity from λ candidates easier. Following the configuration of GEMEL, we set λ to 16 and the performance is presented in Table 5. As we can see, our model can still show superiority. This is due to our method leveraging cross-modal enhancer and alignment to capture dense interaction over multimodal information while GEMEL mainly relies on the ability of LLM. In fact, the LLM of GEMEL cannot truly learn visual information. Instead, it only depends on CLIP to get visual prefix, inevitably resulting in noise and mismatch between text and image.

Ablation Study and Discussion

Attribute enhancement The visual attributes extracted from the image depict the entity’s outline and some prominent features. The extracted visual attributes from the image include facial features and adjective-noun pair (ANPs), which play different roles in the model. (1) Facial features: When the entity in the image is a human subject, the extracted facial features serve as a supplement and correction to the feature extractor, such as in cases where the feature extractor fails to recognize a *white male* in the image. As shown in Table 3, by removing the facial features (method 0→1), the model’s performance decreased by 1.6%, demonstrating the effectiveness of facial features. (2) ANPs: When the image depicts non-human entities like places, ANPs are





Case	1	2	3	4
Image				
Text	Khanna and Manek at the music launch of Ra.One	Khanna and Manek at the music launch of Ra.One	Miller and Saxby Chambliss in 2004	Messi has scored over 700 senior career goals
Mention Entity	Manek <i>Giaa Manek</i> is an Indian television actress, who played a role in film Na Ghar Ke.. Giaa Manek 0.91	Khan <i>Shah Rukh Khan</i> also known by the initialism SRK, is an Indian actor, film producer.. Rajesh Khanna 0.76 Shah Rukh Khan 0.72	Miller <i>Zell Bryan Miller</i> (1932-2013) was an American author and politician.. Sienna Miller 0.79	Messi <i>Lionel Messi</i> , also known as Leo Messi, is an Argentine professional footballer.. Lionel Messi 0.85
Top-5 candidate	Gia Kancheli 0.34 Mihir Bipin Manek 0.29 Brady Manek 0.28 Rama 0.04	Zareen Khan 0.43 Gauhar Khan 0.33 Aamir Khan 0.19	Marisa Miller 0.61 Kristine Miller 0.56 Ezra Miller 0.31 Dean Miller 0.30	Antonello da Messina 0.76 Charles Messier 0.56 Massiv 0.29 Ramesses II 0.06
Predict	✓	✗	✓	✓

Table 6: Multimodal entity linking cases. Entity is depicted by textual Wikipedia description. Top-5 entities are selected from λ candidate entities. We take the bold text to show the correct entity in Top-5 candidate. Case 1 and 2 are multiple object samples for different mentions with the same image and text.

coarse-grained concepts that describe scenes, such as “nice clouds,” and “white dress.” Due to the presence of noise in ANPs, we use a gated fusion to integrate them. By removing (method 1→2) attribute-enhanced mention m_s (ANPs), the accuracy decreased by 1.1%.

Refined visual feature The noise in the image brings challenge to the multimodal model. In this perspective of view, DWE takes visual information refinement to suppress the noise in the image. Table 4 shows the differences between using refined visual feature v and using primitive visual feature that are extracted from the whole image. It could be seen that the refinement brings 2.1% improvement over Top-1 accuracy. Besides, Tables 4 that the effectiveness of the refined visual feature varies on datasets. The 8.8% improvement on WikiMEL (64.0%→72.8%) is more significant than 5.8% on Richpedia (66.7%→72.5%). The reason is that the image in WikiMEL tends to be noisy with multiple objects, while the object number in Richpedia is less, suggesting that the mechanism is more applicable for noisy datasets (such as WikiMEL). Apart from this, Table 3 shows our method without visual feature v : the 3.8% drop (method 3→4) of Top-1 accuracy shows the importance of vision.

Crossmodal alignment The 2.5% decrease of improvement (method 2→3) shows the crossmodal alignment helps to reduce the degree of difference between two features (e.g. partial facial feature f and partial visual feature d , text feature t and refined visual feature v).

Case Study

We further analyze the effectiveness of our model DWE by cases in Table 6, including correct and wrong predictions.

We first analyze the **correct** predictions. The correctness of case 1 and 4 shows the case with multiple visual objects on different datasets. As for the correct case 3, the mention is *Miller* with two old men in image. The result shows the effectiveness of the attributes. The model could mine the relationship between *two old men*, and entity representation

of *Sienna Miller*, while the others are female. On the other hand, **wrong** predictions can still show the effectiveness of our model. In case 2, the visual and textual information includes many objects such as *Khan* and *Khanna*. It is difficult to distinguish these two mentions by the same multimodal information. However, it could be seen that the similarity of *Rajesh Khanna* and *Shah Rukh Khan* is much larger than the rest of candidate entity, which still shows learning ability.

The above cases illustrate the ability of DWE to mine the semantic relationship between textual information (e.g., text and mention) and entity representation. Besides, the supplementary visual feature and attribute extracted from image are leveraged for further disambiguation. In summary, the proposed mechanism helps DWE to understand multimodal information and learn joint feature for entity linking.

Conclusion

In this paper, we explore Multimodal Entity Linking in a neural matching formulation and design a dual-way enhanced framework: query is enhanced by refined multimodal information and the semantics of entity representation is enriched by Wikipedia. DWE utilizes a pretrained visual encoder and vision tools to acquire the visual representation and visual attributes, which are subsequently refined into learnable visual characteristics. It regards learnable visual characteristics and textual content as refined information, each contributing to the enhancement of the mention via three enhanced units. During the process, we apply the crossmodal alignment to bridge the semantic gap among visual and textual features to gain the multi-view alignment representation. Extensive experiments show that DWE outperforms the state-of-the-art methods on three public datasets. Furthermore, we introduce an enhanced entity representation, i.e., textual Wikipedia description. The enhanced entity representation is more representative and no longer restricts the learning of multimodal information.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No.72188101, No.62302144). This work was supported by the Hunan Provincial Natural Science Foundation Project (No.2022JJ30668, 2022JJ30046); and the science and technology innovation program of Hunan province under grant No. 2021GK2001.

References

- Adjali, O.; Besançon, R.; Ferret, O.; Le Borgne, H.; and Grau, B. 2020a. Building a Multimodal Entity Linking Dataset From Tweets. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 4285–4292. Marseille, France: European Language Resources Association. ISBN 979-10-95546-34-4.
- Adjali, O.; Besançon, R.; Ferret, O.; Le Borgne, H.; and Grau, B. 2020b. Multimodal entity linking for tweets. In *European Conference on Information Retrieval*, 463–478. Springer.
- Al-Moslemi, T.; Ocaña, M. G.; Opdahl, A. L.; and Veres, C. 2020. Named entity extraction for knowledge graphs: A literature overview. *IEEE Access*, 8: 32862–32881.
- Auer, S.; Bizer, C.; Kobilarov, G.; Lehmann, J.; Cyganiak, R.; and Ives, Z. 2007. Dbpedia: A nucleus for a web of open data. In *The Semantic Web: 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007+ ASWC 2007, Busan, Korea, November 11-15, 2007. Proceedings*, 722–735. Springer.
- Borth, D.; Ji, R.; Chen, T.; Breuel, T.; and Chang, S.-F. 2013. ANPs extractor: Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *Proceedings of the 21st ACM international conference on Multimedia*.
- Chen, J.; Liu, Q.; Lin, H.; Han, X.; and Sun, L. 2022a. Few-shot named entity recognition with self-describing networks. *arXiv preprint arXiv:2203.12252*.
- Chen, X.; Zhang, N.; Li, L.; Deng, S.; Tan, C.; Xu, C.; Huang, F.; Si, L.; and Chen, H. 2022b. Hybrid Transformer with Multi-level Fusion for Multimodal Knowledge Graph Completion. *arXiv preprint arXiv:2205.02357*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North*.
- Eshel, Y.; Cohen, N.; Radinsky, K.; Markovitch, S.; Yamada, I.; and Levy, O. 2017. Named entity disambiguation for noisy text. *arXiv preprint arXiv:1706.09147*.
- Fei, H.; Liu, Q.; Zhang, M.; Zhang, M.; and Chua, T.-S. 2023. Scene Graph as Pivoting: Inference-time Image-free Unsupervised Multimodal Machine Translation with Visual Scene Hallucination. *arXiv preprint arXiv:2305.12256*.
- Gan, J.; Luo, J.; Wang, H.; Wang, S.; He, W.; and Huang, Q. 2021. Multimodal entity linking: a new dataset and a baseline. In *Proceedings of the 29th ACM International Conference on Multimedia*, 993–1001.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Koch, G.; Zemel, R.; and Salakhutdinov, R. 2015. Siamese Neural Networks for One-shot Image Recognition.
- Ling, X.; Singh, S.; and Weld, D. S. 2015. Design challenges for entity linking. *Transactions of the Association for Computational Linguistics*, 3: 315–328.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled Weight Decay Regularization. *arXiv e-prints*, arXiv:1711.05101.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2016a. Hierarchical question-image co-attention for visual question answering. *Advances in neural information processing systems*, 29.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2016b. Hierarchical Question-Image Co-Attention for Visual Question Answering.
- Ma, W.; Chen, Q.; Zhou, T.; Zhao, S.; and Cai, Z. 2023. Using Multimodal Contrastive Knowledge Distillation for Video-Text Retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Ma, W.; Zhou, T.; Qin, J.; Zhou, Q.; and Cai, Z. 2022. Joint-attention feature fusion network and dual-adaptive NMS for object detection. *Knowledge-Based Systems*, 241: 108213.
- Michel, F.; Gandon, F.; Ah-Kane, V.; Bobasheva, A.; Cabrio, E.; Corby, O.; Gazzotti, R.; Giboin, A.; Marro, S.; Mayer, T.; et al. 2020. Covid-on-the-Web: Knowledge graph and services to advance COVID-19 research. In *The Semantic Web- ISWC 2020: 19th International Semantic Web Conference, Athens, Greece, November 2-6, 2020, Proceedings, Part II 19*, 294–310. Springer.
- Moon, S.; Neves, L.; and Carvalho, V. 2018. Multimodal named entity disambiguation for noisy social media posts. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2000–2008.
- Olafenwa, A. 2021. Simplifying object segmentation with pixellib library. *Online.(2021)*. <https://vixra.org/abs/2101.0122>.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. ArXiv:2103.00020 [cs].
- Reiss, T.; and Hoshen, Y. 2021. Mean-shifted contrastive loss for anomaly detection. *arXiv preprint arXiv:2106.03844*.
- Shen, W.; Wang, J.; and Han, J. 2014. Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering*, 27(2): 443–460.
- Shi, S.; Xu, Z.; Hu, B.; and Zhang, M. 2023. Generative Multimodal Entity Linking. *arXiv preprint arXiv:2306.12725*.
- Sun, W.; Fan, Y.; Guo, J.; Zhang, R.; and Cheng, X. 2022a. Visual Named Entity Linking: A New Dataset and A Baseline. *arXiv preprint arXiv:2211.04872*.

- Sun, W.; Fan, Y.; Guo, J.; Zhang, R.; and Cheng, X. 2022b. Visual Named Entity Linking: A New Dataset and A Baseline. *arXiv preprint arXiv:2211.04872*.
- Vempala, A.; and Preotjuc-Pietro, D. 2019. Categorizing and inferring the relationship between the text and image of twitter posts. In *Proceedings of the 57th annual meeting of the Association for Computational Linguistics*, 2830–2840.
- Vrandečić, D. 2012. Wikidata: A new platform for collaborative data collection. In *Proceedings of the 21st international conference on world wide web*, 1063–1064.
- Wagh, R.; and Punde, P. 2018. Survey on sentiment analysis using twitter dataset. In *2018 Second International Conference on Electronics, Communication and Aerospace Technology (ICECA)*, 208–211. IEEE.
- Wang, M.; Wang, H.; Qi, G.; and Zheng, Q. 2020. Richpedia: a large-scale, comprehensive multi-modal knowledge graph. *Big Data Research*, 22: 100159.
- Wang, P. 2022. Multimodal Entity Linking with Gated Hierarchical Fusion and Contrastive Training. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 938–948. Madrid Spain: ACM. ISBN 978-1-4503-8732-3.
- Wang, X.; Gui, M.; Jiang, Y.; Jia, Z.; Bach, N.; Wang, T.; Huang, Z.; Huang, F.; and Tu, K. 2021. ITA: Image-Text Alignments for Multi-Modal Named Entity Recognition. *arXiv preprint arXiv:2112.06482*.
- Wang, X.; Tian, J.; Gui, M.; Li, Z.; Wang, R.; Yan, M.; Chen, L.; and Xiao, Y. 2022. WikiDiverse: A Multimodal Entity Linking Dataset with Diversified Contextual Topics and Entity Types. ArXiv:2204.06347 [cs].
- Wu, L.; Petroni, F.; Josifoski, M.; Riedel, S.; and Zettlemoyer, L. 2019. Scalable Zero-shot Entity Linking with Dense Entity Retrieval.
- Yang, C.; He, B.; Wu, Y.; Xing, C.; He, L.; and Ma, C. 2023. MMEL: A Joint Learning Framework for Multi-Mention Entity Linking. In *Uncertainty in Artificial Intelligence*, 2411–2421. PMLR.
- Zhang, L.; Li, Z.; and Yang, Q. 2021. Attention-based multimodal entity linking with high-quality images. In *International Conference on Database Systems for Advanced Applications*, 533–548. Springer.
- Zhang, Z.; Chen, K.; Wang, R.; Utiyama, M.; Sumita, E.; Li, Z.; and Zhao, H. 2019. Neural machine translation with universal visual representation. In *International Conference on Learning Representations*.
- Zhao, S.; Hu, M.; Cai, Z.; and Liu, F. 2021a. Dynamic modeling cross-modal interactions in two-phase prediction for entity-relation extraction. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zhao, S.; Hu, M.; Cai, Z.; Zhang, Z.; Zhou, T.; and Liu, F. 2021b. Enhancing Chinese character representation with lattice-aligned attention. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zhao, S.; Wang, C.; Hu, M.; Yan, T.; and Wang, M. 2023. MCL: multi-granularity contrastive learning framework for Chinese NER. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14011–14019.
- Zhou, X.; Wang, P.; Li, G.; Xie, J.; and Wu, J. 2021. Weibo-MEL, Wikidata-MEL and Richpedia-MEL: Multimodal Entity Linking Benchmark Datasets. In Qin, B.; Jin, Z.; Wang, H.; Pan, J.; Liu, Y.; and An, B., eds., *Knowledge Graph and Semantic Computing: Knowledge Graph Empowers New Infrastructure Construction*, volume 1466, 315–320. Singapore: Springer Singapore. ISBN 9789811664700 9789811664717. Series Title: Communications in Computer and Information Science.