

# Joint Multimodal Entity-Relation Extraction Based on Edge-Enhanced Graph Alignment Network and Word-Pair Relation Tagging

Li Yuan<sup>1, 2</sup>, Yi Cai<sup>1, 2, 3 \*</sup>, Jin Wang<sup>4</sup>, Qing Li<sup>5</sup>

<sup>1</sup>School of Software Engineering, South China University of Technology, Guangzhou, China

<sup>2</sup>Key Laboratory of Big Data and Intelligent Robot (SCUT), MOE of China

<sup>3</sup>The Peng Cheng Laboratory, Shenzhen, China

<sup>4</sup>School of Information Science and Engineering, Yunnan University, Yunnan, P.R. China

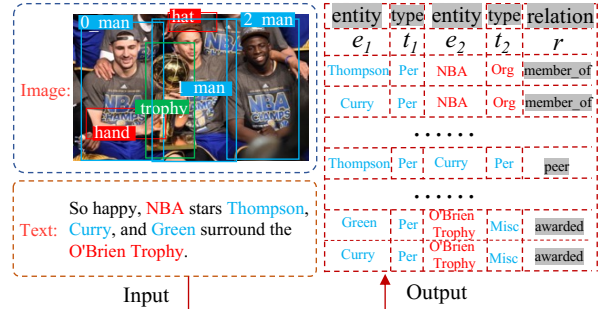
<sup>5</sup>Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China  
seyuanli@mail.scut.edu.cn, ycai@scut.edu.cn, wangjin@ynu.edu.cn, qing-prof.li@polyu.edu.hk

## Abstract

Multimodal named entity recognition (MNER) and multimodal relation extraction (MRE) are two fundamental sub-tasks in the multimodal knowledge graph construction task. However, the existing methods usually handle two tasks independently, which ignores the bidirectional interaction between them. This paper is the first to propose jointly performing MNER and MRE as a joint multimodal entity-relation extraction (JMERE) task. Besides, the current MNER and MRE models only consider aligning the visual objects with textual entities in visual and textual graphs but ignore the entity-entity relationships and object-object relationships. To address the above challenges, we propose an edge-enhanced graph alignment network and a word-pair relation tagging (EEGA) for the JMERE task. Specifically, we first design a word-pair relation tagging to exploit the bidirectional interaction between MNER and MRE and avoid error propagation. Then, we propose an edge-enhanced graph alignment network to enhance the JMERE task by aligning nodes and edges in the cross-graph. Compared with previous methods, the proposed method can leverage the edge information to auxiliary alignment between objects and entities and find the correlations between entity-entity relationships and object-object relationships. Experiments are conducted to show the effectiveness of our model<sup>1</sup>.

## Introduction

Multimodal named entity recognition (MNER) and multimodal relation extraction (MRE) are two fundamental sub-tasks for the multimodal knowledge graph construction (Liu et al. 2019; Chen, Jia, and Xiang 2020), which aims to extend the text-based models by taking images as additional inputs. Previous works usually consider MNER and MRE as two independent tasks (Lu et al. 2018; Moon, Neves, and Carvalho 2018; Wu et al. 2020b; Yu et al. 2020; Zheng et al. 2021c; Zhang et al. 2021a), which ignore the interaction between these two tasks. Recently, jointing NER and RE as joint entity-relation extraction tasks have attracted much attention in text scenarios, which can exploit the bidirectional interaction between tasks and improve their performance (Wei



Joint Multimodal Entity-Relation Extraction (JMERE)

Figure 1: An illustrative example of the joint multimodal entity-relation extraction (JMERE) task, where the Per, Org, and Misc are denoted as the entity types of person, organization, and miscellaneous.

et al. 2020; Yuan et al. 2020a,b). As shown in Figure 1, if we extract the entity type of (**Curry**, **NBA**) is *Per* and *Org*, then their relation should not be *peer*. Otherwise, if we know the relation of entity pair (**Curry**, **Thompson**) is the *peer*, then their entity types should be *Per* and *Per*. Thus, the NER can facilitate RE. Meanwhile, RE is also beneficial for NER.

However, to the best of our knowledge, joining the MNER and MRE as a joint multimodal entity-relation extraction (JMERE) task has not been studied in the multimodal scenario. Compared with separate tasks, the JMERE task requires extracting different characteristic information from vision. As shown in Figure 1, for the MNER task, if the model can capture the people objects from the image, e.g., outlines of multiple people (blue boxes), it helps to identify the person entity in the text. Meanwhile, the MRE task needs to extract the object-object relationships, e.g., if we know the *holding* is the relationship between **man<sub>0</sub>** and **trophy**, we can understand the relation *awarded* between entities **Thompson** and **O'Brien Trophy**. Thus, we consider that the JMERE task should align entities with objects and entity-entity relationships (in text) with object-object relationships (in image). Most recent MNER and MRE studies (Zhang et al. 2021a; Zheng et al. 2021a) align the entities with objects in the visual and textual graphs constructed by the latent relationships of objects and words, as shown by the

\*Corresponding author: Yi Cai (ycai@scut.edu.cn)

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>The code and appendix are available at <https://github.com/YuanLi95/EEGA-for-JMERE>

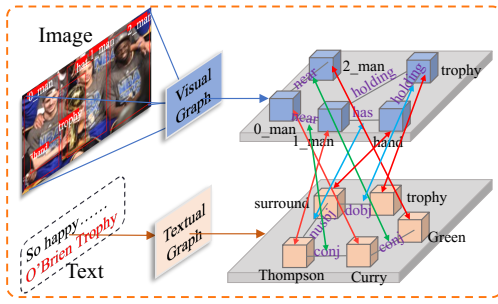


Figure 2: An example for illustrating node (red line) and edge alignment (blue and green lines) in the cross-graph.

red lines in Figure 2. However, this method only considers node alignment in the cross-graph but ignores edge alignment. As shown in the blue and green lines in Figure 2, the edge information in the cross-graph can auxiliary align the nodes and contain clues about the relations between textual entities.

In addition, the pipeline framework method is an intuitive way to solve the JMERE task. It extracts the entities using MNER method and then classifies their relations by MRE method. However, the pipeline framework only benefits MRE through the results of MNER and suffers from error propagation (Ju et al. 2021). As shown in Figure 1, if the MNER extracts the entity type of (**Curry**, **NBA Stars**) is *Per* and *Misc*, the result should be incorrect. Inspired by the grid tagging scheme in the aspect-based sentiment triplet extraction task (Wu et al. 2020a), we first adopt a word-pair ( $w_i, w_j$ ) classification scheme for the JMERE task, namely the word-pair relation tagging. This scheme simultaneously trains MNER and MRE tasks to exploit the interaction between them and avoid the error propagation caused by the pipeline framework. As shown in Figure 3, the word-pair relation tagging of word pairs (**Curry**, **Curry**) and (**Thompson**, **Curry**) is *Per* and *Peer*, which denotes **Curry** belongs to a person and the *peer* denotes this the relation between **Curry** and **Thompson**.

To address the above challenges, we propose an edge-enhanced graph alignment network (EEGA) and word-pair relation tagging to enhance JMERE by simultaneously aligning objects with entities (e.g., **0\_man** with **Curry** and **trophy** with **trophy**) and object-object relationships with entity-entity relationships (e.g., *near* with *conj* and *holding* with *dobj*) in the cross-graph. The overall framework of EEGA is shown in Figure 4. Specifically, we use a graph encoder layer to construct the textual and visual graphs from the input text-image using pre-trained models. Then, we propose an edge-enhanced graph alignment module with Wasserstein distance to align the nodes and edges in the cross-graph. Meanwhile, the module can leverage the edge information to auxiliary alignment between objects and entities and find the correlations between entity-entity relationships and object-object relationships. Finally, we design a multi-channel layer by mining word-word relationships from different perspectives to obtain the final word pair representations.

Our main contributions can be summarized as follows:

- We are the first to propose the joint multimodal entity relation extraction (JMERE) task to handle the multimodal

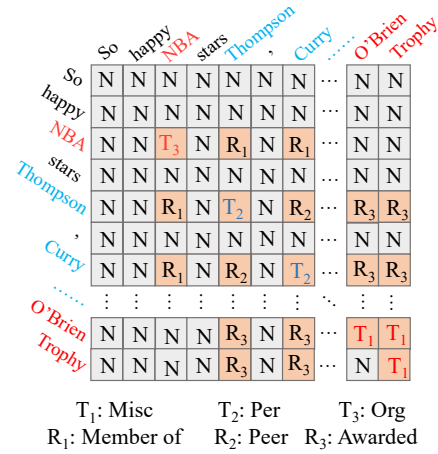


Figure 3: A word-pair relation tagging example for the JMERE task.

NER and RE tasks. Meanwhile, we design a word-pair relation tagging for JMERE. This scheme can exploit the bidirectional interaction between MNER and MRE and avoid the error propagation caused by the pipeline framework.

- We propose an edge-enhanced graph alignment network (EEGA) to enhance the JMERE task by aligning nodes and edges simultaneously in the cross-graph. Compared with previous methods, the EEGA can leverage the edge information to auxiliary alignment between objects and entities and find the correlations between entity-entity relationships and object-object relationships.
- We conduct extensive experiments on the collected JMERE dataset, and the experimental results demonstrate the effectiveness of our proposed model.

## Related Works

The crucial components of the knowledge graphs construction task (Chen, Jia, and Xiang 2020; Chen et al. 2022b,a), named entity recognition (NER) and relation extraction (RE), have attracted much attention from researchers (Vashishth, Joshi, and Suman 2018; Wen et al. 2020; Li et al. 2020; Ren et al. 2020; Nasar, Jaffry, and Malik 2021; Zhao et al. 2021). Previous researches have mainly focused on a single modality. With the increasing popularity of multimodal data on social platforms, some studies have begun to focus on multimodal NER (MNER) and multimodal RE (MRE), which aim to consider the image as a supplement to text and better recognize the entities and their relations. According to the object of image-text alignment, the current methods of MNER and MRE can be divided into image alignment methods, object alignment methods, and node alignment methods.

## Image Alignment Methods

Previous studies usually used RNN (Recurrent Neural Networks) to encode text and CNN (Convolution Neural Network) to encode image as a vector. Then, designing an implicit interaction module to model the information between

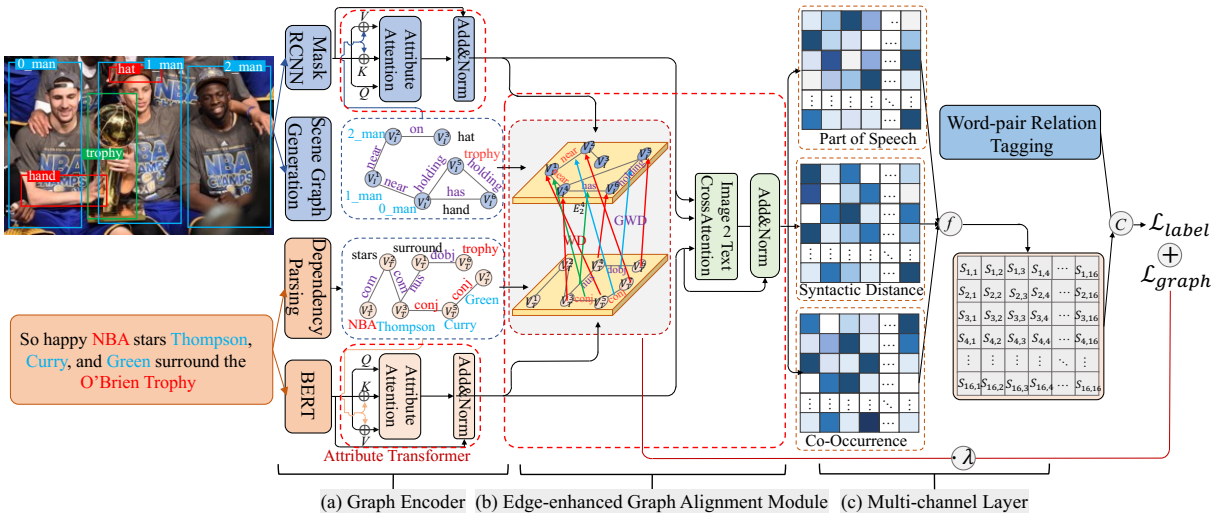


Figure 4: The overall framework of EEGA for JMERE, where com and nus denote the syntactic dependency relationships compound and nsubj.

modalities for the MNER task (Zhang et al. 2018; Lu et al. 2018; Moon, Neves, and Carvalho 2018). For example, Zhang et al. (2018) constructed an MNER dataset and proposed a baseline model based on a bi-directional long-term memory network using an attention mechanism to align the image representation with text. However, encoding image as a vector cannot benefit extract different type entities, e.g., **Curry** (*Per*) and **O’Brien Trophy** (*Misc*).

### Object Alignment Methods

To address limitation of image alignment methods, the previous models extracted different visual objects using Mask-RNN or Fast-RNN (He et al. 2017) and aligned visual objects with the text representation (Wu et al. 2020b; Yu et al. 2020; Zheng et al. 2021c; Zhang et al. 2021a; Xu et al. 2022). Wu et al. (2020b) proposed an interactive attention structure to align text with visual objects. In addition, Zheng et al. (2021c) designed a gated bilinear attention network (Kim, Jun, and Zhang 2018) with an adversarial strategy to better extract the fine-grained objects from the image. However, the object alignment method does not consider the relations of entity-entity and object-object, and the model will ineffectively match overlapping visual objects with textual entities. For example, the **trophy** of text may align with the **1\_man**, since the **1\_man** contains the region of **trophy**.

### Node Alignment Methods

To address the above limitation, the most current researches align the entities with objects in the visual and textual graphs constructed by the latent relationships of objects and words (Zhang et al. 2021a,b; Zheng et al. 2021a). Zhang et al. (2021a) proposed a graph-based multimodal fusion model based on the syntactic dependency text graph and full connection visual graph to exploit the fine-grained semantic alignment different modalities. In the MRE task (Zheng et al. 2021b), Zheng et al. (2021a) designed a graph alignment module to align nodes in textual graph and visual graphs.

However, the node alignment methods only consider the nodes in the cross-graph and ignore the edge information. The edge information in the cross-graph can effectively improve the matching precision of nodes and contain some clues about the classifying relation between entities.

## Task Definition and Word-Pair Relation Tagging

### Task Definition

The joint multimodal entity-relation extraction task is defined: given an input text  $w = \{w_1, w_2, \dots, w_n\}$  with a corresponding image  $I$ , the model is required to extract a set of quintuples  $y = \{(e_1, t_1, e_2, t_2, r)\}_{c=1}^C$ , where  $(e_1, t_1, e_2, t_2, r)_c$  represents the  $c$ -th quintuple, consisting of two entities  $e_1$  and  $e_2$  with the corresponding entity types  $t_1$  and  $t_2$ , where  $e_1 \neq e_2$ . Furthermore,  $r$  represents the relation between the entities  $e_1$  and  $e_2$ . Figure 1 gives an example to better understand the JMERE task, which aims to extract all quintuples, e.g., (**Thompson**, *Per*, **NBA**, *Org*, *Member\_of*), where *Per* and *Org* indicate the entity types of **Thompson** and **NBA**, and *Member\_of* denotes their relation type.

### Word-Pair Relation Tagging

Inspired by the grid tagging scheme in aspect-based sentiment triplet extraction (Wu et al. 2020a), we design a word-pair relation tagging to extract all elements of JMERE in one step. By word-pair relation tagging, the JMERE task is converted into extracting the relations  $Y$  between each word-pair  $(w_i, w_j)$  and avoid the error propagation caused by the pipeline framework. These relations can be explained below, and we also give an example in Figure 3 to better understand the word-pair relation tagging.

- $N$  indicates that the word-pair does not have any relation.
- $T$  indicates that the word-pair belongs to an entity type, which is contained 4 defined types in the previous work (Zheng et al. 2021c)

- $R$  indicates that the word-pair belongs to defined relation (Zheng et al. 2021a) and each word is an entity.

## Edge-Enhanced Graph Alignment Network

Figure 4 shows the overall architecture of the proposed model, consisting of three parts: a graph encoder, an edge-enhanced graph alignment module, and a multi-channel layer. The graph encoder layer uses the pre-trained models to construct the textual and visual graphs from the input. To enhance the ability to capture edge information, we do not directly send textual and visual representations to the next module but an attribute transformer. Then, to match the objects with entities more precisely and capture the entity-entity relation clues from the visual graph, we use a cross-graph optimal transport method (Chen et al. 2020) with Wasserstein distance and Gromov-Wasserstein distance to simultaneously align the nodes and edges in the cross-graph. Finally, we propose a multi-channel layer that uses a weighted graph convolution network (W-GCN) to mine the latent relationships for word pairs from multi-perspectives. A detailed description of each component is provided below. The code of the manuscript will publish in the final version.

### Graph Encoder

**Textual Graph.** Formally, we first use the dependency parse toolkit<sup>2</sup> to construct the textual graph. As shown in Figure 4 (a), after parsing, the given sentence is converted into a textual graph  $G_T = \{V_T, E_T\}$ , where  $V_T \in \mathbb{R}^n$  and  $E_T \in \mathbb{R}^{n \times n}$  denote syntactic dependency nodes and edges, respectively, and  $G_T$  is an un-directed self-loop graph. Meanwhile, we use  $A_T \in \mathbb{R}^{n \times n}$  to denote the adjacent mask matrix for textual graph, where  $A_T^{i,j} \in \{0, 1\}$  indicate whether there is an edge between  $w_i$  and  $w_j$  or not. Furthermore, the nodes  $V_T$  are fed into BERT to obtain  $X_T \in \mathbb{R}^{n \times d_T}$ . Meanwhile, an edge transition matrix is used to map the edge type  $E_T^{i,j}$  into a trainable vector and obtained edge trainable matrix  $Z_T \in \mathbb{R}^{n \times n \times d_{zT}}$ ,

$$X_T = BERT(w) \quad (1)$$

where  $BERT$  denotes the BERT as the text encoder and  $d_T$  is the dimension of the BERT output.

**Visual Graph.** In previous multimodal tasks, the objects were considered the semantic information of images. As shown in Figure 4 (a), we convert the image into a visual graph  $G_I = \{V_I, E_I\}$  by using the scene graph generation model (Tang et al. 2020) (the Mask-RCNN used as the backbone). Furthermore, we only consider the top  $k$  salient objects with the highest object classification scores as the valid visual objects and exploit the adequate visual information while ignoring the irrelevant ones. Thus, the final node represents  $V_I \in \mathbb{R}^k$  salient objects detected by Mask-RCNN and  $E_I \in \mathbb{R}^{k \times k}$  denotes the visual relationship set, such as position relationships (e.g., *near* and *in front of*) and affiliation relationships (e.g., *holding* and *wearing*). We use  $A_I \in \mathbb{R}^{k \times k}$  to denote the adjacent matrix of the visual graph. Thus, the

final node vectors in visual graph  $X_I \in \mathbb{R}^{k \times d_I}$  are defined as,

$$X_I = Mask-RCNN(I) \quad (2)$$

where  $d_I$  is the hidden dimension of Mask-RCNN and the final edge vectors  $Z_I \in \mathbb{R}^{k \times k \times d_{zI}}$  are obtained in the same way as the textual graph.

**Attribute Transformer.** we propose the attribute attention (*At-Att*) by incorporating the edge types into keys and values in the self-attention of the transformer as the attribute transformer, which can update the node state in the inter-model while effectively incorporating relationship edges in the cross-graph (e.g., *nsubj* and *acomp* in the textual graph, and *holding* and *wearing* in the visual graph).

Since the visual and textual graphs are two semantic units containing information in different modalities, we model them using similar operations but with different parameters. Thus, the hidden representation of the  $i$ -th token  $\tilde{H}_T^i \in \mathbb{R}^{n \times d_T}$  in text modalities is defined as,

$$\begin{aligned} \tilde{H}_T^i &= At-Att(X_T^i, Z_T^i, A_T^i) \\ &= \text{Softmax}\left(A_T^i \cdot \frac{Q_T^i (K_T^i)^\top}{\sqrt{d_T}}\right) V_T^i \end{aligned} \quad (3)$$

where  $A_T^i \in \mathbb{R}^{1 \times n}$  is the adjacency mask set of the  $i$ -th node,  $Q_T^i \in \mathbb{R}^{1 \times d_T}$ ,  $K_T^i \in \mathbb{R}^{n \times d_T}$ , and  $V_T^i \in \mathbb{R}^{n \times d_T}$  are matrices that package the queries, keys, and values for the  $i$ -th word in text correspondingly, which are defined as,

$$\begin{aligned} Q_T^i &= W_Q X_T^i \\ K_T^i &= W_K X_T + W_z Z_T^i \\ V_T^i &= W_V X_T + W_r Z_T^i \end{aligned} \quad (4)$$

The other operations of the attribute transformer are consistent with the vanilla transformer:  $\tilde{H}_T$  is added to  $X_T$  using a feed-forward network (FFN) and layer normalization (Layer-Norm) to obtain the text representation  $H_T \in \mathbb{R}^{n \times d_T}$ .

We use similar operations to obtain the visual representation. In particular, we use a variable-dimensional FFN to match the dimension of object  $H_I$  and token  $H_T$ . Thus, the image representation of graph encoder can be denoted as  $H_I \in \mathbb{R}^{k \times d_T}$ .

### Edge-Enhanced Graph Alignment Module

Given the textual graph  $G_T = \{H_T, Z_T\}$ , and the visual graph  $G_I = \{H_I, Z_I\}$ . We aim to align the nodes and edges of the cross-graphs simultaneously and to transfer the matched semantic information from objects into entities. Formally, an optimal transport method (Chen et al. 2020) is first used to match the nodes and edges in the cross-graph. Furthermore, we use image2text attention to transfer matched semantic information from the visual object to text modality and obtain the refined textual representation.

**Edge-Enhanced Graph Optimal Transport.** To explicitly encourage simultaneously aligning the nodes and edges in the cross-graph, we apply the optimal transport method initially proposed in transfer learning. As illustrated in Figure 4 (b), unlike the original optimal transport method considering

<sup>2</sup><https://spacy.io/models>.




Figure 5: The multi-channel matrices of the given sentence.

text and image as full-connected graphs, we only consider the nodes and edges of having adjacency relationships in the cross-graph. Particularly, two types of distance are adopted for cross-graph matching: (1) Wasserstein Distance (WD) (Peyré, Cuturi et al. 2019) for node matching (the red lines); (2) Gromov-Wasserstein Distance (GWD) (Peyré, Cuturi, and Solomon 2016) for edge matching (the blue and green lines). Formally, the  $D_{wd}(H_I, H_T)$  is measured the optimal transport distance to match the nodes  $H_I$  to  $H_T$ , which is defined as:

$$D_{wd}(H_I, H_T) = \min \sum_{i=1}^k \sum_{j=1}^n T_{i,j} \cdot c(H_I^i, H_T^j) \quad (5)$$

where  $c(H_I^i, H_T^j)$  denotes the cosine distance between  $X_I^i$  to  $X_T^j$ , which is defined as  $c(H_I^i, H_T^j) = 1 - \frac{(H_I^i)^T H_T^j}{\|H_I^i\| \|H_T^j\|}$ . The matrix  $T$  is the transport information flow, where  $T_{i,j}$  represents the amount of cost shifted from node  $X_I^i$  to  $X_T^j$ .

Then, we use the Gromov-Wasserstein distance (Peyré, Cuturi, and Solomon 2016) to measure the similarity scores  $D_{gwd}$  of the edge in the cross-graph by calculating distances between node-pairs, which is defined as,

$$D_{gwd}(H_I, H_T, H_I', H_T') = \min \sum_{i=1}^k \sum_{j=1}^n \hat{T}_{i,j} \hat{T}_{i',j'} \cdot L(H_I^i, H_I'^{i'}, H_T^j, H_T'^{j'}) \quad (6)$$

where  $H_T^j$  and  $H_I'^{i'}$  are adjacent nodes sets in the textual and visual graphs of  $H_T^j$  and  $H_I'^{i'}$  respectively, and  $L(\cdot)$  is considered as the distance cost of the cross-graph edges  $(H_I^i, H_I'^{i'})$  to  $(H_T^j, H_T'^{j'})$ , i.e.  $L(H_I^i, H_I'^{i'}, H_T^j, H_T'^{j'}) = \|c(H_I^i, H_I'^{i'}) - c(H_T^j, H_T'^{j'})\|$ . The learned matrix  $\hat{T}$  now denotes a transport plan that aids in aligning edges in the cross-graph.

We use a unified solver and use the Sinkhorn algorithm (Cuturi 2013) with an entropic-regularizer (Benamou et al. 2015) to iteratively optimize costs  $D_{wd}$  and  $D_{gwd}$ . Thus, the object loss function of optimizing the cross-graph is,

$$\mathcal{L}_{graph} = \alpha \cdot D_{wd}(H_I, H_T) + (1 - \alpha) \cdot D_{gwd}(H_I, H_I', H_T, H_T') \quad (7)$$

where  $\alpha$  is the hyper-parameter for balancing the importance of costs. Then, we use image2text attention to effectively

transform the visual semantic information into the textual representation, which is denoted as,

$$\tilde{O} = ATT_{cross}(H_T, H_I, H_I) \quad (8)$$

where  $ATT_{cross}$  denotes the cross-modal multihead attention (Ju et al. 2020). Then, the  $\tilde{O}$  is added with  $H_T$  and sends a layer normalization to obtain the final contextual representation  $O$ .

## Multi-Channel Layer

In this subsection, we aim to mine the different dependency features between  $w_i$  and  $w_j$  to help detect relations between them. As shown in Figures 5: (a) we consider that part of speech (Pos) can provide lexical information for word pairs. For example, the Pos of most entities belong to the NOUN and PEROPN, e.g., **NBA**, **Curry**, and **Thompson**; (b) encoding the syntactic distance (Sd) between word-pair can improve the ability of the model to capture long-range syntactic information; (c) the word co-occurrences matrix (Co) can provide corpus-level information between word pairs. For example, **Curry** and **NBA** appear some times in the corpus. The details about constructing each feature matrix are added in Appendix-A.

After data preprocessing, three feature matrices are obtained,  $M^l \in \mathbb{R}^{n \times n}, l \in \{Pos, Sd, Co\}$ . We propose a W-GCN module to model each matrix, obtaining each channel representation. Each matrix  $M^l$  firstly sends an embedding layer yielding a trainable representation  $R^l \in \mathbb{R}^{n \times n \times d_l}$  and  $d_l$  is the dimension of representation. The calculation W-GCN process of  $i$ -th word in  $l$ -th matrix is shown as,

$$S_i^l = \text{W-GCN}_l(R_i^l, O) = \text{Softmax}(\text{ReLU}(W_l^{r1} R_i^{lT} + b_l)) \cdot (W_l^{r2}(O)) \quad (9)$$

where  $R_i^l \in \mathbb{R}^{n \times d_l}$  is the  $i$ -th word in  $l$ -th linguistic matrix and  $W_l^{r1} \in \mathbb{R}^{1 \times d_l}$  and  $W_l^{r2} \in \mathbb{R}^{d_T \times d_T}$  are shared weights used to perform a linear layer to learn linguistic features and representational abilities. We combine the representations and send them to the MLP (MultiLayer Perception) layer for obtaining the final word representation,

$$S_i = \text{MLP}[S_i^{Pos}; S_i^{Sd}; S_i^{Co}] \quad (10)$$

where  $S_i \in \mathbb{R}^{d_T}$  is the  $i$ -th word representation. Thus, the output representation of the multi-channel layer is denoted as  $S = [S_1, S_2 \dots, S_n]$ . Finally, we concatenate the enhanced representations of  $S_i$  and  $S_j$  to represent the word-pair  $(w_i, w_j)$ , i.e.,  $r_{i,j} = [S_i; S_j]$ . Then, send the  $r_{i,j}$  to a linear prediction layer and obtain the probability distribution,

$$p_{i,j} = \text{Softmax}(W_p r_{i,j} + b_p) \quad (11)$$

where  $W_p \in \mathbb{R}^{d_y \times 2d_T}$  and  $b_p \in \mathbb{R}^{d_y}$  are trainable parameters and  $d_y$  is the number of tags. Then, we used the cross-entropy error to measure the ground truth distribution  $\mathcal{Y}$  and predicted tagging distribution,

$$\mathcal{L}_{main}(\theta) = - \sum_{s=1}^S \sum_{i=1}^n \sum_{j=1}^n \mathcal{Y}_{i,j}^s \log(p_{i,j}^s \theta) \quad (12)$$

where  $S$  and  $\theta$  denote the number of training samples and all trainable parameters, respectively.

Methods		JMERE			#MNER		
		#P	#R	F1	#P	#R	F1
Pipeline Methods	AdapCoAtt+MEGA	48.44	47.06	47.74	74.32	72.11	73.20
	OCSGA+MEGA	48.21	47.99	48.10	75.27	72.32	73.77
	AGBAN+MEGA	47.87	48.28	48.57	74.78	73.69	74.23
	UMGF+MEGA	49.28	50.76	50.01	75.02	76.77	75.88
Word-pair Relation Tagging Methods	AdapCoAtt*	50.22	47.67	48.91	77.32	73.28	75.25
	OCSGA*	52.11	47.41	49.64	77.13	75.03	76.07
	AGBAN*	51.07	48.89	49.95	76.57	75.82	76.19
	UMGF*	52.76	50.22	51.45	77.51	76.01	76.75
	MEGA*	55.08	51.40	53.18	77.78	76.67	<u>77.22</u>
	MAF*	52.56	<b>54.73</b>	<u>53.62</u>	76.07	77.57	76.81
	EEGA( <i>ours</i> )	<b>58.26</b> †	<u>52.61</u>	<b>55.29</b> †	<b>78.27</b>	<b>78.91</b> †	<b>78.59</b> †

Table 1: The experiment results on the JMERE task (%) and the #MNER denote the MNER results computed by the JMERE results. AGBAN\* means using the word-pair relation tagging in the AGBAN model. The marker † refers to significant test  $p - value < 0.05$ . The best result is in bold and #P, #R, and F1 denote the precision, recall, and F1-score.

Methods	#P	#R	F1
EEGA( <i>All</i> )	58.26	52.61	55.29
<i>w/o edge-enhanced</i>	51.85	50.31	51.07
<i>w/o attribute transformer</i>	55.69	51.07	53.28
<i>w/o multi-channel</i>	55.48	52.13	53.75

Table 2: Results of ablation study for the JMERE task.

## Join Training

The final objective is a combination of the main task and optimizing the cross-graph as follows,

$$\mathcal{L} = \mathcal{L}_{main} + \lambda \cdot \mathcal{L}_{graph} \quad (13)$$

where  $\lambda$  are trade off hyper-parameters to control the contribution of optimizing the cross-graph.

## Experiments

Comparative experiments were conducted to evaluate and compare the performance of the EEGA method against several prior works. Furthermore, more detailed experiments (e.g., datasets, setting, and parameter sensitivity) are presented in Appendix B-D.

### Comparative Results

**Compared Methods.** We summarize the MNER and MRE studies and combine the state-of-the-art methods as our strong JMERE baselines, as shown in Table 1. They include AdapCoAtt (Zhang et al. 2018), OCSGA (Wu et al. 2020b), AGBAN (Zhang et al. 2021a), and MAF (Xu et al. 2022) for extraction of the entity and the corresponding type, and UMGF (Zheng et al. 2021a) for relation extraction of entities. In addition, we apply the word-pair relation tagging to the above baseline models to investigate the effectiveness of the word-pair relation tagging and the proposed method, e.g., AdapCoAtt\*, OCSGA\*, and MEGA\*.

**Overall Results.** Observing pipeline methods, we find that the UMGF+MEGA performs better than other pipeline methods, which shows that aligning the nodes in the cross-graph

can benefit matching the entities with objects. The word-pair relation tagging methods outperform the pipeline methods in JMERE and #MNER, such as OCSGA, AGBAN, and UMGF, showing that the word-pair relation tagging can improve performance by leveraging task relationships and reducing error propagation issues caused by the pipeline framework.

Furthermore, the EEGA surpasses all baselines. Compared with the best results of existing baselines, EEGA still achieves absolute F1-score increases of 1.67% and 1.37% on JMERE and #MNER. The experimental results strongly prove that simultaneously aligning the nodes and edges in the cross-graph can effectively improve the precision of matching the objects with entities and capture more relationships between entities. In addition, the proposed attribute transformer enhances the ability to mine relations between nodes by incorporating edge information into the key and value in the transformer. Meanwhile, the multi-channel layer can take linguistic relations between word pairs to refine the final representation and improve prediction performance.

**Ablation Study.** To investigate the effectiveness of different components in EEGA, edge-enhanced graph optimal transport (edge-enhanced), attribute transformer, and multi-channel layer, we conduct an ablation study for the JMERE task in Table 2. *W/o attribute transformer* means that a vanilla transformer replaces the attribute transformer. The F1-score dropped 2.01%, indicating that integrating the edge information into the key and value in the transformer can enhance the ability to capture the relations between nodes and benefit the edge alignment in the cross-graph. The performance is decreased after removing the multi-channel layer (*w/o multi-channel*), indicating the multi-channel layer can mine relationships of word pairs from different perspectives and refine the final representation.

*W/o edge-enhanced* means removing the edge-enhanced graph optimal transport from EEGA. The performance of the model is highly degraded after removing the edge-enhanced, showing that simultaneously aligning nodes and edges in the cross-graph can be beneficial for matching the visual objects with textual entities more precisely and finding the entity classification clues from the relationship between objects.

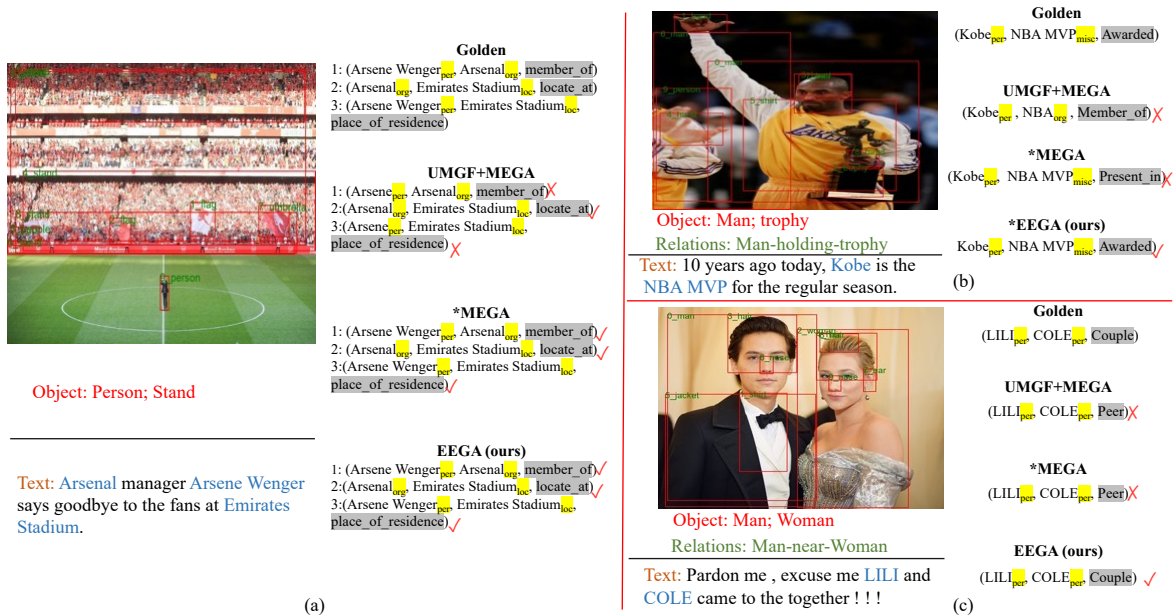


Figure 6: Three cases of the predictions by UMGF+MEGA, \*MEGA, and EEGA.




Figure 7: A comparison of Attention visualization on entity-object pair from  $\alpha = 1$  (only node alignment) and  $\alpha = 0.4$  (best performance setting) in Eq. (7)

### Case Study

To understand the effectiveness of our proposed model, Figure 6 presents three examples with the predicted results. Meanwhile, the important objects and relations are detected from images. In example (a), only the pipeline-based model UMGF+MEGA extracts incorrectly since the pipeline model easily suffers from error propagation, i.e., the extraction entity **Arsene** by UMGF is incomplete, and the final model UMGF+MEGA extract the incorrect quintuple. In example (b), UMGF+MEGA imprecisely extracts **NBA** as an entity, and \*MEGA incorrectly predicts the relationship between **Kobe** and **NBA MVP** as the *Present\_in*. In example (c), the situation is similar to example (b). Since lacking effective ways to map the semantic relationship of objects *Man-near-Woman* to entities (*LILI-COLE*), the UMGF+MEGA and \*MEGA incorrectly predicts relations *peer* between entities.

For these three examples, the proposed EEGA makes accurate judgments. Benefiting from the edge-enhanced graph optimal transport module, the EEGA can align the nodes and edges in the cross-graph to match the entities with objects more precisely. Meanwhile, the EEGA also effectively captures the relation clue from the visual graph to the textual graph shown in examples (b) and (c). In addition, the attribute transformer and multi-channel layer can further enhance the

ability to model the relationships of objects and word pairs.

### Visualization Analysis

In this section, we visualize the example (b) of Figure 6 at  $\alpha = 1$  and  $\alpha = 0.4$  to test whether our edge alignment strategy helps to learn fine-grained entity-object matching. As shown in Figure 7 (a), when only node alignment means  $\alpha = 1$ , since the proposed model lacks the edge constraints, the attention weight is relatively scattered and affects the precision of matching entities with objects. Particularly, the model easily classifies the type of entity **NBA** as *Per*. Meanwhile, as shown in the **NBA** and **Kobe** in Figure 7 (b), benefiting from the edge alignment can find the mapping between the object-object relationships and entity-entity relationship; the EEGA effectively reduces the ambiguity and match objects with entities more precisely.

### Conclusion

In this paper, we are the first to propose a joint multimodal entity relation extraction (JMERE) task to handle the multimodal NER and RE tasks. To tackle this task, we propose an edge-enhanced graph alignment network and a word-pair relation tagging. Specifically, we design a word-pair relation tagging to avoid the error propagation caused by the pipeline framework. Then, we propose an edge-enhanced graph alignment network (EEGA) to enhance the JMERE task by aligning nodes and edges simultaneously in the cross-graph. The EEGA can leverage the edge information to auxiliary alignment between objects and entities and find the correlations between entity-entity relationships and object-object relationships. The detailed evaluation demonstrates that our proposed model significantly outperforms several state-of-the-art baselines. We will extend our approach to multi-label multimodal tasks in our future work and investigate other methods (e.g., the self-supervised model) to better model JMERE.

## Acknowledgments

This work was supported by National Natural Science Foundation of China (62076100, 61966038), and Fundamental Research Funds for the Central Universities, SCUT (x2rjD2220050), the Science and Technology Planning Project of Guangdong Province (2020B0101100002), the CAAI-Huawei MindSpore Open Fund, the Hong Kong Research Grants Council (project no. PolyU 11204919 and project no. C1031-18G) and an internal research grant from the Hong Kong Polytechnic University (project 1.9B0V).

## References

- Benamou, J.-D.; Carlier, G.; Cuturi, M.; Nenna, L.; and Peyré, G. 2015. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2): 1111–1138.
- Chen, L.; Gan, Z.; Cheng, Y.; Li, L.; Carin, L.; and Liu, J. 2020. Graph optimal transport for cross-domain alignment. In *Proceedings of the ICML 2020*, 1520–1531.
- Chen, X.; Jia, S.; and Xiang, Y. 2020. A review: Knowledge reasoning over knowledge graph. *Expert Systems with Applications*, 141: 112948–112966.
- Chen, X.; Zhang, N.; Li, L.; Deng, S.; Tan, C.; Xu, C.; Huang, F.; Si, L.; and Chen, H. 2022a. Hybrid Transformer with Multi-Level Fusion for Multimodal Knowledge Graph Completion. In *Proceedings of the SIGIR 2022*, 904–915.
- Chen, X.; Zhang, N.; Li, L.; Yao, Y.; Deng, S.; Tan, C.; Huang, F.; Si, L.; and Chen, H. 2022b. Good Visual Guidance Make A Better Extractor: Hierarchical Visual Prefix for Multimodal Entity and Relation Extraction. In *Findings of the NAACL 2022*, 1607–1618.
- Cuturi, M. 2013. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Proceedings of the NIPS 2013*, 1–9.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the ICCV 2017*, 2961–2969.
- Ju, X.; Zhang, D.; Li, J.; and Zhou, G. 2020. Transformer-based label set generation for multi-modal multi-label emotion detection. In *Proceedings of the ACM MM 2020*, 512–520.
- Ju, X.; Zhang, D.; Xiao, R.; Li, J.; Li, S.; Zhang, M.; and Zhou, G. 2021. Joint multi-modal aspect-sentiment analysis with auxiliary cross-modal relation detection. In *Proceedings of the EMNLP 2021*, 4395–4405.
- Kim, J. H.; Jun, J.; and Zhang, B. T. 2018. Bilinear attention networks. In *Proceedings of the NIPS 2018*, 1564–1574.
- Li, X.; Feng, J.; Meng, Y.; Han, Q.; Wu, F.; and Li, J. 2020. A unified MRC framework for named entity recognition. In *Proceedings of the ACL 2020*, 5849–5859.
- Liu, Y.; Li, H.; Garcia-Duran, A.; Niepert, M.; Onoro-Rubio, D.; and Rosenblum, D. S. 2019. MMKG: multi-modal knowledge graphs. In *Proceedings of the ESWC 2019*, 459–474.
- Lu, D.; Neves, L.; Carvalho, V.; Zhang, N.; and Ji, H. 2018. Visual attention model for name tagging in multimodal social media. In *Proceedings of the ACL 2018*, 1990–1999.
- Moon, S.; Neves, L.; and Carvalho, V. 2018. Multimodal named entity recognition for short social media posts. In *Proceedings of the NAACL 2018*, 852–860.
- Nasar, Z.; Jaffry, S. W.; and Malik, M. K. 2021. Named entity recognition and relation extraction. *ACM Computing Surveys*, 54: 1–39.
- Peyré, G.; Cuturi, M.; and Solomon, J. 2016. Gromov-wasserstein averaging of kernel and distance matrices. In *Proceedings of the ICML 2016*, 2664–2672.
- Peyré, G.; Cuturi, M.; et al. 2019. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5-6): 355–607.
- Ren, H.; Cai, Y.; Chen, X.; Wang, G.; and Li, Q. 2020. A Two-phase Prototypical Network Model for Incremental Few-shot Relation Classification. In *Proceedings of the COLING 2020*, 1618–1629.
- Tang, K.; Niu, Y.; Huang, J.; Shi, J.; and Zhang, H. 2020. Unbiased scene graph generation from biased training. In *Proceedings of the CVPR 2020*, 3713–3722.
- Vashishth, S.; Joshi, R.; and Suman, S. 2018. RESIDE: Improving distantly-supervised neural relation extraction using side information. In *Proceedings of EMNLP 2018*, 1257–1266.
- Wei, Z.; Su, J.; Wang, Y.; Tian, Y.; and Chang, Y. 2020. A novel cascade binary tagging framework for relational triple extraction. In *Proceedings of the ACL 2020*, 1476–1488.
- Wen, Y.; Fan, C.; Chen, G.; Chen, X.; and Chen, M. 2020. A Survey on Named Entity Recognition. *Lecture notes in electrical engineering*, 571: 1803–1810.
- Wu, Z.; Ying, C.; Zhao, F.; Fan, Z.; Dai, X.; and Xia, R. 2020a. Grid tagging scheme for aspect-oriented fine-grained opinion extraction. In *Findings of the EMNLP 2020*, 2576–2585.
- Wu, Z.; Zheng, C.; Cai, Y.; Chen, J.; Leung, H. F.; and Li, Q. 2020b. Multimodal representation with embedded visual guiding objects for named entity recognition in social media posts. In *Proceedings of the ACM MM 2020*, 1038–1046.
- Xu, B.; Huang, S.; Sha, C.; and Wang, H. 2022. MAF: A general matching and alignment framework for multimodal named entity recognition. In *Proceedings of the WSDM 2022*, 1215–1223.
- Yu, J.; Jiang, J.; Yang, L.; and Xia, R. 2020. Improving multimodal named entity recognition via entity span detection with unified multimodal transformer. In *Proceedings of the ACL 2020*, 3342–3352.
- Yuan, L.; Wang, J.; Yu, L.-C.; and Zhang, X. 2020a. Graph Attention Network with Memory Fusion for Aspect-level Sentiment Analysis. In *Proceedings of the AACL 2020*, 27–36.
- Yuan, Y.; Zhou, X.; Pan, S.; Zhu, Q.; Song, Z.; and Guo, L. 2020b. A relation-specific attention network for joint entity and relation extraction. In *Proceedings of the IJCAI 2020*, 4054–4060.
- Zhang, D.; Wei, S.; Li, S.; Wu, H.; Zhu, Q.; and Zhou, G. 2021a. Multi-modal graph fusion for named entity recognition with targeted visual guidance. In *Proceedings of the AAAI 2021*, 14347–14355.

- Zhang, Q.; Fu, J.; Liu, X.; and Huang, X. 2018. Adaptive co-attention network for named entity recognition in tweets. In *Proceedings of the AAAI 2018*, 5674–5681.
- Zhang, Y.; Wang, J.; Yu, L.-C.; and Zhang, X. 2021b. MA-BERT: Learning representation by incorporating multi-attribute knowledge in transformers. In *Findings of the ACL 2021*, 2338–2343.
- Zhao, T.; Yan, Z.; Cao, Y.; and Li, Z. 2021. Asking effective and diverse questions: a machine reading comprehension based framework for joint entity-relation extraction. In *Proceedings of the IJCAI 2021*, 3948–3954.
- Zheng, C.; Feng, J.; Fu, Z.; Cai, Y.; Li, Q.; and Wang, T. 2021a. Multimodal relation extraction with efficient graph alignment. In *Proceedings of the ACM MM 21*, 5298–5306.
- Zheng, C.; Wu, Z.; Feng, J.; Fu, Z.; and Yi, C. 2021b. MNRE: A challenge multimodal dataset for neural relation extraction with visual evidence in social media posts. In *Proceedings of the ICME 2021*, 3–8.
- Zheng, C.; Wu, Z.; Wang, T.; Cai, Y.; and Li, Q. 2021c. Object-aware multimodal named entity recognition in social media posts with adversarial learning. *IEEE Transactions on Multimedia*, 23: 2520–2532.