

Hierarchical Aligned Multimodal Learning for NER on Tweet Posts

Peipei Liu^{1,2}, Hong Li^{1,2*}, Yimo Ren^{1,2}, Jie Liu^{1,2}, Shuaizong Si¹, Hongsong Zhu^{1,2}, Limin Sun^{1,2}

¹Institute of Information Engineering, Chinese Academy of Sciences
19 Shucun Road, Haidian District, Beijing 100085 P.R.China

²School of Cyber Security, University of Chinese Academy of Sciences
19 Yuquan Road, Shijingshan District, Beijing 100049 P.R.China
{liupeipei, lihong, zhuhongsong}@iie.ac.cn

Abstract

Mining structured knowledge from tweets using named entity recognition (NER) can be beneficial for many downstream applications such as recommendation and intention understanding. With tweet posts tending to be multimodal, multimodal named entity recognition (MNER) has attracted more attention. In this paper, we propose a novel approach, which can dynamically align the image and text sequence and achieve the multi-level cross-modal learning to augment textual word representation for MNER improvement. To be specific, our framework can be split into three main stages: the first stage focuses on intra-modality representation learning to derive the implicit global and local knowledge of each modality, the second evaluates the relevance between the text and its accompanying image and integrates different grained visual information based on the relevance, the third enforces semantic refinement via iterative cross-modal interactions and co-attention. We conduct experiments on two open datasets, and the results and detailed analysis demonstrate the advantage of our model.

Introduction

In recent years, social media platforms like Twitter and Facebook have made great development, and they provide important sources (i.e., tweets) for various applications such as the identification of cyber-attacks or natural disasters, analyzing public opinion, and mining disease outbreaks (Bruns and Liang 2012; Ritter et al. 2015). Named entity recognition, the task of detecting and classifying named entities from unstructured free-form text, is a crucial step to extract the structured information for those downstream applications (Perera et al. 2018; Dionísio et al. 2019). With tweets tending to be multimodal and traditional unimodal NER methods having challenges to understand these multimedia contents perfectly, multimodal named entity recognition (MNER) has become a new direction and it improves conventional text-based NER by considering images as additional inputs (Lu et al. 2018; Yu et al. 2020; Sun et al. 2021; Chen et al. 2022b). For example, as shown in Figure 1(a), the supplement of visual information can alleviate the semantic ambiguity and inadequacy problem caused by only text in-



Figure 1: The samples for MNER task, where the named entities and their types are highlighted. a: fully relevant (explicit support information), b: partially relevant (implicit support information), c: entity irrelevant (no entities, no support).

formation while classifying the named entity *Baby Jasmine* to **OTHER** instead of **PER**.

The core of existing MNER methods is to achieve the fusion and alignment of visual information and textual information through different cross-modal technologies. The methods can roughly be divided into four main streams: (1) (Lu et al. 2018; Zhang et al. 2018; Arshad et al. 2019; Durant 2021; Wang et al. 2022c) employ pre-trained CNN models such as ResNet (He et al. 2016) to encode the whole images into a global feature vector, and then augment each word representation with the global image vector by effective attention mechanism. (2) (Yu et al. 2020; Xu et al. 2022b; Liu et al. 2022) divide the feature map obtained from the whole image into multiple blocks averagely, and subsequently learn the most valuable vision-aware word representation by modeling the interaction between text sequence and the visual regions with Transformer or gating mechanism. (3) Some researchers apply object detection models like Mask RCNN (He et al. 2017) to obtain the visual objects from the associated image, and then combine the object-level visual information and textual word information based on GNN or cross-modal attention (Wu et al. 2020; Zhang et al. 2021; Lu et al. 2022; Chen et al. 2022a,b; Zhao et al. 2022b). (4) There are also some works to explore the derivative knowledge of image content including OCR, image caption, query guidance and other image attributes (Chen et al. 2021b; Jia et al. 2022; Wang et al. 2022a; Zhao et al. 2022a;

*The corresponding author

Wang et al. 2022b; Jia et al. 2023), which is used to guide words to get the expanded visual semantic information.

Despite the impressive results of these existing methods, there are still several evident limitations remained. Firstly, the current methods heavily relied on the argument that the accompanying image of posted text is entity-related and the visual information is helpful for textual entity extraction. However, the viewpoint is not always valid and the relevance between the text and image is in various situations: fully relevant, partially relevant and irrelevant. As a result, the noise of irrelevant visual content would lead to misleading interaction representation and further affect the MNER performance. In fact, we can observe that the image adds no additional content to the text in 33.8% of tweets from the report of (Vempala and Preotjuc-Pietro 2019). Secondly, the current methods usually leverage the image or object representation directly extracted from the original vision view but the implicit knowledge such as image scene and interactive relationship between different objects is neglected. Take Figure 1(b) as an example, the visual stadium scene can easily help us make a correct prediction for the entities *Penguins* and *Sharks* with regarding them as **ORG** entities rather than **OTHER** entities (i.e., not animals). Lastly, although they have achieved state-of-the-art results with cross-modal interaction in various ways, seldom of them have explored the multi-level semantic alignments between the vision and text modality. In practice, there are two key points that a word is the basic unit and several words make up a sentence while an image is composed of a number of objects and attributes. Constructing different level alignments not only captures fine-to-coarse correlations between images and texts, but also takes the advantage of the complementary information among these semantic levels.

There is some work carried out for the partial limitations: (Xu et al. 2022b; Sun et al. 2021, 2020; Xu et al. 2022a) design additional classification tasks to measure the text-image relationship by introducing external tools and datasets, (Liu et al. 2022; Chen et al. 2022a; Jia et al. 2023) initially employ the multi-level vision information. However, their solutions are imperfect since additional tasks usually need to be built and the multi-level textual content is not used. Motivated by such findings, we propose the novel **Hierarchical aligned multimodal Learning (HamLearning)**, which aims to end-to-end model multi-level semantics for both modalities and enforce cross-modal interactions at different semantic levels with the basis of measuring the relevance of text and image. Specifically, we perform the model within three stages: 1) Firstly, we use the textual Transformer (Vaswani et al. 2017) to learn the global representation of sentence and contextual representations of words. At the same time, two separate visual encoders are deployed to capture object-to-object relations and consider the global scene of image from semantic and spacial view, respectively. 2) Secondly, the relevance between text sentence and its accompanying image is measured through the global content alignment. Then, we integrate the object-level and image-level visual representation to acquire local-to-global sufficient visual feature based on the relevance score. 3) Finally, we implement cross-modal interaction between word representations and

the fused visual feature iteratively to refine the most effective multimodal clues for decoding.

We conduct the extensive experiments on two popular MNER datasets, Twitter2015 (Lu et al. 2018) and Twitter2017 (Zhang et al. 2018), to evaluate the performance of the **HamLearning** framework, and results show the superiority of our approach. Moreover, the full experimental analyses help us understand the advantages and details of the model comprehensively.

The main contributions of this paper can be summarized as:

- We construct multi-level alignments to capture coarse-to-fine interactions between vision and language, and take the advantage of the complementary information among these semantic levels.
- We introduce the new spatial and semantic learning of visual scene for the MNER task and directly utilize learned visual feature without using the existing generation tools for visual semantic.
- We design the end-to-end dynamic relevance measuring on image-text for specific MNER task instead of performing additional text-image relationship classification tasks based on the external tools and datasets.
- Through detailed experiments and analyses, we demonstrate the competitive performance of our method in comparison with the current excellent models.

Related Work

As social media posts become more multimodal, MNER is attracting researchers' attention. (Yu et al. 2020) extends the vanilla Transformer to cross-modality Transformer for capturing multimodal interactions between text words and image regions, and further designs an additional entity-span detection module to alleviate the bias of visual factor and improve the performance of MNER. (Chen et al. 2021a; Wang et al. 2022c) enhance the text representation by integrating the knowledge and attributes of corresponding images. Different from above works of using the whole image, (Wu et al. 2020; Zhang et al. 2021; Zhao et al. 2022b; Chen et al. 2022b) believe that the textual entities are determined by the visual objects, so the object-level visual representation is used to guide the entity recognition in the text. In addition to improvements of MNER methods, (Sun et al. 2020, 2021) control the effect of images on text at two different stages through gate- and attention- mechanisms, and pre-train multimodal BERT models for MNER based on text-image relationship inference.

Nevertheless, these researches just focus on the certain visual grained feature (i.e., fine or coarse), but ignore the effect of multi-level interaction. (Chen et al. 2022a,b; Jia et al. 2023) make a few preliminary attempts on multi-level visual information. In their works, they simply exploit the output features from different stages of pre-trained vision models for cross-modal fusion. There are some other approaches that do not directly use the visual information from the images, but they open the new paths to mine the hidden information behind the image. (Wang et al. 2022b) designs several prompt templates for each image to bridge the gap

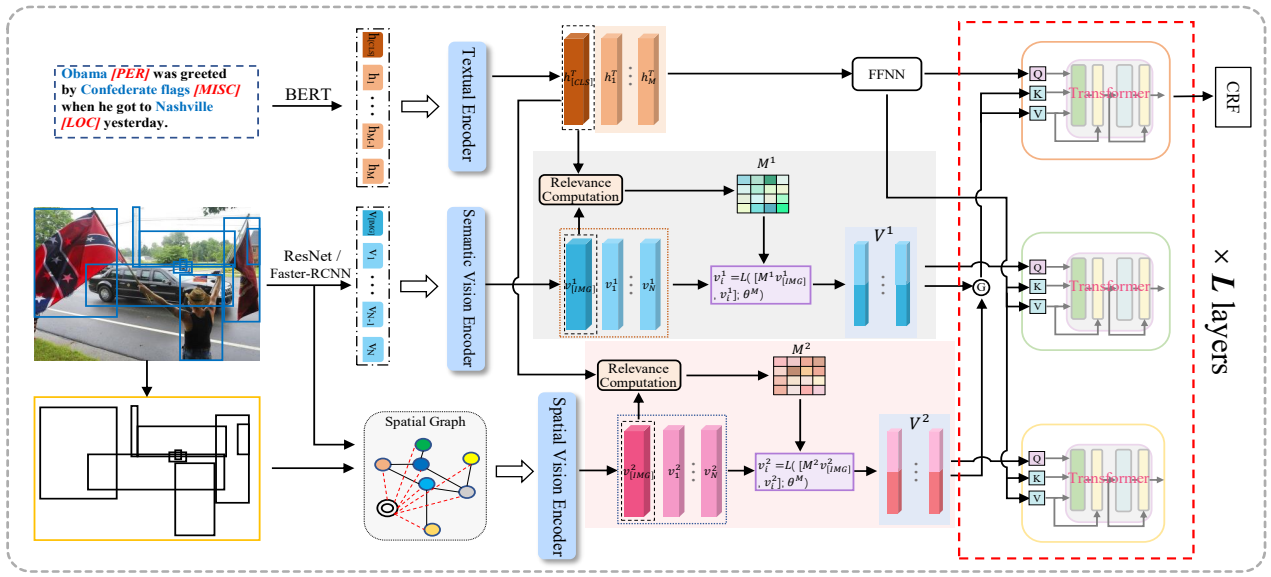


Figure 2: The overview of our proposed method.

between vision space and text space while (Jia et al. 2022, 2023) construct the queries for the image description. Not only that, (Wang et al. 2022a; Zhao et al. 2022a) introduce the image caption and OCR information to compensate for the irrationality of original visual information.

Some works (Xu et al. 2022b; Sun et al. 2021, 2020; Xu et al. 2022a) also question the matching problem between images and text information, and they have paid attention to the solution. They mainly use the pre-training Vision-Language Model (VLM) such as CLIP (Radford et al. 2021) to design additional text-image relationship classification tasks, and add vision assistance to text information based on prediction probability. However, this line of thinking depends too much on the pre-trained VLM models, and the performance has a great correlation with VLM. In addition, the computed vector seems to refine the visual information related to the entities, but in fact, the content of whole image is the main one. If the image does not match with the text, there is still large visual misleading noise.

In this study, we consider such problems comprehensively and create the **HamLearning** for the solution.

Method

In this section, we firstly define the MNER task and then we take the image and text sequence shown in Figure 2 as a running example to introduce details of our proposed approach.

Task Definition: Given the input pair containing a text sentence \mathbf{X} and an image \mathbf{I} , the goal of MNER is to detect a set of entities from \mathbf{X} , and classify them into the pre-defined types (Yu et al. 2020). As with other works in the literature (Moon, Neves, and Carvalho 2018; Lu et al. 2018; Zhang et al. 2018; Arshad et al. 2019; Yu et al. 2020; Jia et al. 2022; Chen et al. 2022b), we regard the MNER as a sequence labeling task. Let $\mathbf{X} = \{x_1, \dots, x_M\}$ denote the input sequence with M words and $\mathbf{Y} = \{y_1, \dots, y_M\}$ indicate the

corresponding label sequence, where $y_i \in \zeta$ and ζ is a pre-defined label set in standard BIO2 formats (Li et al. 2022).

Intra-modality Learning

In this section, we use the Transformer and modified R-GCN to learn local and global representations of text and vision respectively.

Text Encoding BERT (Devlin et al. 2019) benefits from a large external corpus and has a strong dynamic feature extraction capability for the same word in different contexts. In this work, each text sequence $\mathbf{X} = \{x_1, \dots, x_M\}$ is fed into the pre-trained 12-layers BERT to get the sequence representations. As we all know, the additional special token “[CLS]” should be added to the first position to represent the global semantic of entire sentence. Therefore, we can obtain the factual output $H = \{h_{[CLS]}, h_1, \dots, h_M\}$, where $h_{[CLS]}$ is the global sentence feature and h_i ($i \in \{1, \dots, M\}$) is the extracted word representation for x_i .

$$h_i = BERT(x_i; \theta^{bert}) \quad (1)$$

where θ^{bert} is the BERT parameter. Particularly, if x_i is split into several sub-tokens through the tokenizer, we get h_i by summing the sub-tokens.

Next, we feed the H into a textual Transformer for the further encoding. As a result, we can receive the output $\{h_{[CLS]}^T, h_1^T, \dots, h_M^T\}$

Vision Encoding To our intuition, image-level feature represents the global visual information (containing the scene and category, etc) of each image. But the original features directly extracted from general vision models can usually not satisfy the MNER requirement due to their specific task objectives. Following (Yao et al. 2018; Yang, Li, and Yu 2021; Han et al. 2022), we enhance the global representation of the image-level by analyzing the objects and the




Figure 3: We have the relation between object i (red region) and object j (black region).

relationship between them in the image. Also, the local representation of each object can benefit from their companions and superiors through the interaction learning.

Here, for both forms of encoding, we firstly get the initial image-level and object-level representations. Then, the encoders are applied to update the representations by capturing the relationship between different elements (i.e., including image and objects) on spatial/semantic views.

For the global representation of whole image, we choose ResNet (He et al. 2016) as the feature extractor considering that it is one of the most excellent CNN models for many vision tasks. We take the feature map from the last convolutional layer in a pre-trained 152-layers ResNet to represent \mathbf{I} . Then, we transform the feature map from ResNet into a one-dimensional vector by linear function. We can denote the feature vector as follows:

$$v_{[IMG]} = L(\text{ResNet}(\mathbf{I}; \theta^{res}); \theta^I) \quad (2)$$

where θ^{res} is the ResNet parameter and θ^I is the learnable parameter of linear transformation.

For local features of objects, we employ a Faster-RCNN (Ren et al. 2017) model pre-trained on the Visual Genome (Krishna et al. 2017) to detect the objects $\{o_1, o_2, \dots, o_N\}$ in the image and output the representation of each object region proposal. What’s more, we adopt the concept classified by the Faster-RCNN for each object as another feature clue. Therefore, we can describe the feature representation of object o_i as:

$$v_i = L(f_i; \theta^v) + W_e e_i^v \quad (3)$$

where N indicates top- N detected objects with higher scores, f_i is the object region feature extracted by Faster-RCNN, θ^v is the learnable parameter of linear transformation, W_e is the projection matrix, e_i^v is the label embedding by looking up the embedding table E^v .

Semantic Encoding

As the same with **Text Encoding**, we arrange the whole image and objects into a sequence $V = \{v_{[IMG]}, v_1, v_2, \dots, v_N\}$ and encourage the interaction among them to learn contextual presentations by Vision Transformer (ViT) from a semantic view. After the semantic encoding, we can get the result:

$$\{v_{[IMG]}^1, v_1^1, v_2^1, \dots, v_N^1\} = ViT(V; \theta^{ViT}) \quad (4)$$

where θ^{ViT} is the learnable parameter of ViT.

Spatial Encoding

Different from the neat semantic sequence, the objects in the visual space are often scattered and irregular. We thus build the structure graph for realizing spacial modeling, resembling to the approaches of (Yao et al. 2018; Wang, Tang, and Luo 2020; Yang, Li, and Yu 2021).

Firstly, we have the spatial relations between any two objects based on their region sizes and locations. For the object o_i , we can denote its location as $(x_i^c, y_i^c, h_i^c, w_i^c)$, where (x_i^c, y_i^c) is the normalized coordinate center of o_i region box, h_i^c is the normalized box height and w_i^c is the normalized box width. The initial spacial feature of o_i is defined as $\hat{v}_i^2 = [x_i^c, y_i^c, h_i^c, w_i^c, v_i]$. Specially, for the whole image \mathbf{I} , we set its initial spacial feature as $\hat{v}_{[IMG]}^2 = [x_{[IMG]}^c, y_{[IMG]}^c, h_{[IMG]}^c, w_{[IMG]}^c, v_{[IMG]}]$, where $(x_{[IMG]}^c, y_{[IMG]}^c, h_{[IMG]}^c, w_{[IMG]}^c)$ are the centroid, height, width of the image, respectively.

Given a pair of objects o_i and o_j , their relationship r_{ij} can be depicted in the following. We compute the relative distance d_{ij} , Intersection over Union u_{ij} , and relative angle δ_{ij} (i.e., angle between vector $(x_j^c - x_i^c, y_j^c - y_i^c)$ and the positive direction of x-axis). If o_i includes o_j completely, we establish an edge $o_i \rightarrow o_j$ with a relation label “inside”; if conversely o_i is contained by o_j fully, we set the edge label as “cover”. Except for the above two cases, we define the edge $o_i \rightarrow o_j$ with a label “overlap” if $u_{ij} > 0.5$. If none of these is true, we build the relationship rely on the ratio ρ_{ij} (i.e., the ratio between d_{ij} and the diagonal length of image) and relative angle δ_{ij} . We can not have an edge for (o_i, o_j) when $\rho_{ij} > 0.5$ but assign r_{ij} into one of the relationship set $\{\text{class1, class2, class3, class4, class5, class6, class7 and class8}\}$ according to δ_{ij} when $\rho_{ij} < 0.5$. (Please see Figure 3 for more details.) By default, we think the \mathbf{I} includes all the objects.

The spatial structure graph $G=(\mathbb{V}, \mathbb{E})$ is then derived through the object-to-object relations, where $\mathbb{V} = \{\mathbf{I}, o_1, \dots, o_N\}$ is the node set, \mathbb{E} is the established edge set with relations. Imitating the operation in text (Chen et al. 2020), we regard the \mathbf{I} in the graph as a super node to collect the global visual information. Next, we update the node representations by using a modified R-GCN with the initial spacial features.

For the vertex $i \in \mathbb{V}$ in k layer of GCN, the accumulating information $v_i^{2,k'}$ from its neighbors can be formalized as follow with considering the relationship label and direction of all connection edges:

$$v_i^{2,k'} = \phi\left(\sum_{j \in \Omega(i)} W_{r_{i \leftrightarrow j}} v_j^{2,k} + b\right) \quad (5)$$

where ϕ is a nonlinear function, $\Omega(i)$ is the neighbor set of i , $r_{i \leftrightarrow j}$ denotes the directional relationship between i and j , $W_{r_{i \leftrightarrow j}}$ indicates the transformation weight with regard to the edge direction and label, $v_j^{2,k}$ is the representation of j in k layer, b is the bias term.

We update the representation of i in $k + 1$ layer through a

gate mechanism:

$$\lambda_{i,g}^{2,k} = \text{sigmoid}(W_{gcn}[v_i^{2,k'}, v_i^{2,k}]) \quad (6)$$

$$v_i^{2,k+1} = v_i^{2,k} + \lambda_{i,g}^{2,k} v_i^{2,k'} \quad (7)$$

where $\text{sigmoid}(\cdot)$ is the activation function, W_{gcn} is the trainable matrix.

After the spacial encoding and learning among graph nodes, we can finally have the local and global visual feature vectors $\{v_{[IMG]}^2, v_1^2, v_2^2, \dots, v_N^2\}$.

Relevance Measuring

Unlike the current works that compute the match degree between image and text relying on the extra image-text classification task and VLM, we dynamically measure the relevance only depend on the MNER objective. We can define the relevance score M^r ($r \in \{1, 2\}$) between global text feature $h_{[CLS]}^T$ and vision feature $v_{[IMG]}^r$ as:

$$C^r = \tanh(h_{[CLS]}^T W_{TI}^r v_{[IMG]}^r) \quad (8)$$

$$M^r = \tanh(W_T^r h_{[CLS]}^T + W_I^r v_{[IMG]}^r C^r) \quad (9)$$

where W_{TI}^r , W_T^r and W_I^r are the learnable weight matrices, $\tanh(\cdot)$ is the activation function. Based on the measure result, we get the local-global vision feature \mathbf{V}^r as follows:

$$v_i^r = L([M^r v_{[IMG]}^r, v_i^r]; \theta^M), i \in \{1, 2, \dots, N\} \quad (10)$$

$$\mathbf{V}^r = [v_1^r, \dots, v_N^r], r \in \{1, 2\} \quad (11)$$

where θ^M is the parameter of linear function.

Inter-modality Learning

During this stage, different modalities iteratively encourage each other to acquire the most powerful multimodal feature. As the Figure 2 shows, this interaction module contains three parts which all take the cross-modal Transformer as the core encoder. To reduce the heterogeneity between text and vision, we transform the text representation before the cross-modal learning:

$$\mathbf{H} = FFNN(\{h_1^T, \dots, h_M^T\}; \theta^{FFNN}) \quad (12)$$

where θ^{FFNN} is a parameter for FFNN training.

Subsequently, the detailed illustrations of these three parts are given (top-down). At the first part, we refine representations of textual words by introducing the fused spatial and semantic visual information. The comprehensive vision feature can be resulted by a gate control:

$$\alpha = W^V \sigma(W^{V1} \mathbf{V}^1 + W^{V2} \mathbf{V}^2) \quad (13)$$

$$\mathbf{V} = \alpha \odot \mathbf{V}^1 + (\mathbf{1} - \alpha) \odot \mathbf{V}^2 \quad (14)$$

where W^V , W^{V1} and W^{V2} are the trainable matrices, $\sigma(\cdot)$ is the activation function, \odot represents element-wise multiplication, $\mathbf{1}$ stands for an all-1 vector. For inputs (Q, K, V) of the Transformer, we assign \mathbf{H} to Q , \mathbf{V} to (K, V) for the update of \mathbf{H} . That is:

$$\mathbf{H} = \text{Transformer}(\mathbf{H}, \mathbf{V}; \theta^{FT}) \quad (15)$$

where θ^{FT} is the learnable parameter for final multimodal text encoding.

Similar to the first part, we can update \mathbf{V}^1 through feeding \mathbf{V}^1 to Q , \mathbf{H} to (K, V) in the second part. Also in the third part, we update \mathbf{V}^2 by taking $\mathbf{V}^2 = Q$, $\mathbf{H} = K = V$.

After L synchronous iterations of all parts, we use the final \mathbf{H} for MNER decoding.

MNER Decoding

Following (Moon, Neves, and Carvalho 2018), we produce the probability of a predicted label sequence y by feeding \mathbf{H} to a CRF layer:

$$p(y|\mathbf{H}; \theta^{CRF}) = \frac{\prod_{i=1}^{M-1} \varphi_i(y_i, y_{i+1}; \mathbf{H})}{\sum_{y' \in \mathbb{Y}} \prod_{i=1}^{M-1} \varphi_i(y'_i, y'_{i+1}; \mathbf{H})} \quad (16)$$

where $\varphi_i(y_i, y_{i+1}; \mathbf{H})$ is a potential function, \mathbb{Y} is a set of all possible label sequences, θ^{CRF} is a set of parameters which define the potential function and the transition score from the label y_i to the label y_{i+1} .

We train the model by maximizing conditional likelihood estimation for the training set $\{(\mathbf{X}, \mathbf{Y})_t\}$:

$$\text{Loss} = \sum_t \log p(\mathbf{Y}|\mathbf{X}; \theta^{CRF}) \quad (17)$$

In the decoding phase, we output the label sequence prediction y^* for given \mathbf{X} based on maximizing the following score:

$$y^* = \arg \max_{y \in \mathbb{Y}} p(y|\mathbf{X}; \theta^{CRF}) \quad (18)$$

Experiments

We evaluate our model on two publicly MNER datasets referring to (Yu et al. 2020; Wang et al. 2022a; Chen et al. 2022b; Wang et al. 2022c) and compare it with a number of approaches. (Due to limitations of the format and pages, we have to move some content into the Appendices¹.)

Datasets

The experiments are carried out on the datasets TWITTER-2015 and TWITTER-2017, which are constructed based on Twitter by (Lu et al. 2018) and (Zhang et al. 2018) separately. Please see the Appendices for more details about these two datasets.

Implementation Details

For both datasets, we have the same hyperparameters and the specific parameter content can be found in the Appendices.

Baselines

We compare our model with some typical excellent approaches for NER, including unimodal approaches (only text as inputs) and multimodal approaches (text-image pairs as inputs). For unimodal approaches, we consider: **BiLSTM-CRF** (Huang, Xu, and Yu 2015). **CNN-BiLSTM-CRF** (Ma and Hovy 2016), extends the work of BiLSTM-CRF and incorporates the character-level word representation learned by CNN into input. **HBiLSTM-CRF** (Lample et al. 2016), similar to CNN-BiLSTM-CRF, but get

¹<https://arxiv.org/abs/2305.08372>

Modality	Methods	TWITTER-2015			TWITTER-2017		
		Overall			Overall		
		P	R	F1	P	R	F1
Text	BiLSTM-CRF	68.14	61.09	64.42	79.42	73.43	76.31
	CNN-BiLSTM-CRF	66.24	68.09	67.15	80.00	78.76	79.37
	HBiLSTM-CRF	70.32	68.05	69.17	82.69	78.16	80.37
	BERT	68.30	74.61	71.32	82.19	83.72	82.95
	BERT-CRF	69.22	74.59	71.81	83.32	83.57	83.44
Text+Image	GVATT-BERT-CRF	69.15	74.46	71.70	83.64	84.38	84.01
	AdaCAN-BERT-CRF	69.87	74.59	72.15	85.13	83.20	84.10
	UMT	71.67	75.23	73.41	85.28	85.34	85.31
	UMGF	74.49	75.21	74.85	86.54	84.50	85.51
	GEI	73.39	75.51	74.43	87.50	86.01	86.75
	MAF	71.86	75.10	73.42	86.13	86.38	86.25
	RDS	71.96	75.00	73.44	86.25	86.38	86.32
	ITA	-	-	75.60	-	-	85.72
	MRC-MNER	78.10	71.45	74.63	88.78	85.00	86.85
	MNER-QG	77.43	72.15	74.70	88.26	85.65	86.94
	MGCMT	73.57	75.59	74.57	86.03	86.16	86.09
	HVPNeT	73.87	76.82	75.32	85.84	87.93	86.87
	HamLearning	77.25	75.75	76.49	86.99	87.28	87.13

Table 1: Performance comparison of different competitive uni-modal and multi-modal approaches.

the character-level word representation from LSTM. **BERT** (Devlin et al. 2019) and its variant **BERT+CRF**. For multimodal approaches, we consider: **AdaCAN-BERT-CRF** (Zhang et al. 2018) and **GVATT-BERT-CRF** (Lu et al. 2018), which combine the whole image feature through visual attention, and we replace their original sentence encoders BiLSTM with BERT. **UMT** (Yu et al. 2020) creates the cross-modal Transformer to encode the image region and text for MNER. **MAF** (Xu et al. 2022b) and **RDS** (Xu et al. 2022a), design the extra contrastive tasks to make text and image more consistent and assign visual feature to assist words. **UMGF** (Zhang et al. 2021) and **GEI** (Zhao et al. 2022b), which enhance cross-modal interaction between visual objects and textual words with GNN. **MRC-MNER** (Jia et al. 2022), **MNER-QG** (Jia et al. 2023) and **ITA** (Wang et al. 2022a) which leverage the prior knowledge, caption and OCR of the image to guide vision-aware word information. **MGCMT** (Liu et al. 2022) and **HVPNeT** (Chen et al. 2022a), the initial attempts for multi-level semantic alignment at different vision levels.

Main Results

Following the literature (Xu et al. 2022a; Zhao et al. 2022b; Chen et al. 2022a), we compute experimental results of overall precision (**P**), recall (**R**), and F1 score (**F1**). For a fair comparison, we refer to the results of all baselines introduced in their papers. Table 1 shows all comparison results. The results suggest that both overall F1s of our method on the two benchmark datasets outperform the published state-of-the-art (SOTA) performance. We also have several findings below:

(1) It is clear that BERT-based methods perform better compared with BiLSTM-based encoders ((BiLSTM-CRF,

CNN-BiLSTM-CRF) vs (BERT, BERT-CRF)), which indicates that the pre-trained model is quite effective due to its large external knowledge support.

(2) Through all multimodal and unimodal approaches, we can see that unimodal approaches generally produce poor performance compared to the multimodal ones, which suggests that visual information of either the global image or the local objects is valuable for MNER. Comparing UMT and UMGF/GEI, we find that UMGF/GEI can get the better result than UMT, which may be that more entities can be guided by fine-grained visual objects. With the performances of UMT and MAF/RDS, we can notice the importance of evaluating the text-image matching degree and eliminating the vision noise.

(3) As the results of ITA, MRC-MNER and MNER-QG show, transforming and guiding the image content into text knowledge to assist NER can also be effective. However, in fact, this method depends on the effect of external energies such as OCR model and hand-crafted templates.

(4) From HamLearning, HVPNeT, MGCMT and others, we can observe that the multi-level semantic interaction is usually more advantageous than single grained information. Considering HamLearning and HVPNeT/MGCMT, we probably think that the more comprehensive and in-depth interaction could lead to better results since HVPNeT and MGCMT focus on the multi-level vision while HamLearning performs on both modalities.

(5) The performance of MAF/RDS and HamLearning presents us the difference of two text-image matching methods, we believe that the dynamic measuring is more reasonable and reliable because of the direct connection of end-to-end to the specific MNER task instead of the parallel classification task.

Settings	TWITTER-2015		
	P	R	F
Default	77.25	75.75	76.49
w/o RGCN	74.73	75.42	75.07
w/o ViT	74.45	75.92	75.18
w/o RM	75.35	73.89	74.61

Table 2: The ablation study on TWITTER-2015.

Settings	TWITTER-2017		
	P	R	F
Default	86.99	87.28	87.13
w/o RGCN	86.17	85.94	86.05
w/o ViT	85.75	86.53	86.14
w/o RM	86.02	85.19	85.60

Table 3: The ablation study on TWITTER-2017.

Ablation Study

To investigate the contribution of main modules in our model, we have an ablation study. Table 2 and Table 3 report comparison results between the full model and its ablation methods. (For convenience, we here use R-GCN to represent the spatial visual encoding, ViT to represent the semantic visual encoding, and RM to represent the relevance measuring.) We find that: 1) All the modules have contributions to our final optimal model, and any removal of the three modules would result in the inferior performance. 2) When we remove the R-GCN, the hidden information behind the image including the scene knowledge and the relationship between objects will be difficult to be captured for multimodal reasoning, and we only use the semantic feature extracted by ViT as the visual supplement. As the results show, the R-GCN plays an more important role compared to the ViT. 3) From the table, we observe that the elimination of the relevance measuring leads to the significant performance degradation. Especially, the recall on both datasets drop a lot. The reason may be that, with the removal of relevance measuring, a large amount of redundant visual feature has produced misleading noise, which damages the ability of the model to detect the correct entities. Furthermore, as described in our method, the absence of RM can also affect the role of R-GCN and ViT.

Further Analysis

Parameter Sensitivity In this section, we evaluate our model on different number L to find the optimal parameter. The Table 4 shows the experiment results. As the results

L	TWITTER-2015			TWITTER-2017		
	P	R	F	P	R	F
1	74.97	75.49	75.23	85.30	86.16	85.73
2	76.61	75.36	75.98	86.97	86.38	86.67
3	77.25	75.75	76.49	86.99	87.28	87.13
4	76.33	75.21	75.77	87.01	85.79	86.40
5	75.25	76.03	75.64	85.43	86.57	86.00

Table 4: The performance of HamLearning by different L numbers.

Methods	TWITTER-17→TWITTER-15		
	P	R	F
UMT \ddagger	64.67	63.59	64.13
UMGF \ddagger	67.00	62.81	66.21
HamLearning	69.17	66.84	67.98

Table 5: The performance comparison of generalization ability between HamLearning and other methods. Results with \ddagger are from (Zhang et al. 2021).

Methods	TWITTER-15→TWITTER-17		
	P	R	F
UMT \ddagger	67.80	55.23	60.87
UMGF \ddagger	69.88	56.92	62.74
HamLearning	71.03	59.40	64.70

Table 6: The same meaning with Table 5 but the opposite generalization.

show, with the increase of L , the performance of the model becomes better. When the number is 3, we can obtain the best model. However, when the number is greater than 3, the performance begins to decline. This may be because, with the deepening of cross-modal interaction learning, the differences between modal information are decreasing, leading to the lack of valuable features.

Generalization Analysis Considering the different data characteristics of the two datasets, we conducted cross validation on them to test the generalization ability of our model and the comparison methods. As shown in Table 5, TWITTER-17→TWITTER-15 indicates that the model trained on TWITTER-2017 dataset is used to test the TWITTER-2015 dataset, and vice versa in Table 6. From the results, we can discover that, our model significantly outperforms its comparisons by a large margin. This phenomenon may potentially confirm the transfer and adaptability ability of multimodal hierarchical semantics in model generalization.

Different Object Detectors and Case Study Please see the Appendices.

Conclusion

In this paper, we introduce a novel hierarchical neural network HamLearning to achieve the multimodal learning for MNER. Our model contains three main modules: intra-modality learning, image-text relevance measuring and iterative cross-modal learning. The intra-modality learning aims to learn unimodal representations of tokens through their inherent attributes and contextual neighbors. In the image-text relevance measuring module, we use the global representations of both text sentence and vision image from previous stage to compute text-image matching score, and then have a local-global visual feature for the text based on the score. In the last, we iteratively perform cross-modal learning between vision and text to refine the most valuable feature for MNER. We conduct the extensive experiments and analyses to demonstrate the advantage of HamLearning.

Acknowledgements

This work was supported by the National Key Research and Development Program of China (2022YFB3103904), the National Natural Science Foundation of China (No.61931019), and the Strategic Priority Research Program of the Chinese Academy of Sciences (No. XDC02011200).

References

- Arshad, O.; Gallo, I.; Nawaz, S.; and et al. 2019. Aiding Intra-Text Representations with Visual Context for Multimodal Named Entity Recognition. In *2019 International Conference on Document Analysis and Recognition*, 337–342. IEEE.
- Bruns, A.; and Liang, Y. E. 2012. Tools and methods for capturing Twitter data during natural disasters. *First Monday*, 17.
- Chen, D.; Li, Z.; Gu, B.; and Chen, Z. 2021a. Multimodal Named Entity Recognition with Image Attributes and Image Knowledge. In *Database Systems for Advanced Applications*, 186–201. Springer International Publishing.
- Chen, S.; Aguilar, G.; Neves, L.; and et al. 2021b. Can images help recognize entities? A study of the role of images for Multimodal NER. In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, 87–96. Association for Computational Linguistics.
- Chen, S.; Zhao, Y.; Jin, Q.; and et al. 2020. Fine-Grained Video-Text Retrieval With Hierarchical Graph Reasoning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10635–10644.
- Chen, X.; Zhang, N.; Li, L.; and et al. 2022a. Good Visual Guidance Make A Better Extractor: Hierarchical Visual Prefix for Multimodal Entity and Relation Extraction. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 1607–1618. Association for Computational Linguistics.
- Chen, X.; Zhang, N.; Li, L.; and et al. 2022b. Hybrid Transformer with Multi-level Fusion for Multimodal Knowledge Graph Completion. *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and et al. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4171–4186. Association for Computational Linguistics.
- Dionísio, N.; Alves, F.; Ferreira, P. M.; and et al. 2019. Cyberthreat Detection from Twitter using Deep Neural Networks. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–8.
- Durant, K. 2021. Multi-Granularity Contrastive Knowledge Distillation for Multimodal Named Entity Recognition.
- Han, K.; Wang, Y.; Chen, H.; and et al. 2022. A Survey on Vision Transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 87–110.
- He, K.; Gkioxari, G.; Dollár, P.; and et al. 2017. Mask R-CNN. In *2017 IEEE International Conference on Computer Vision*, 2980–2988. IEEE.
- He, K.; Zhang, X.; Ren, S.; and et al. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, 770–778. IEEE.
- Huang, Z.; Xu, W.; and Yu, K. 2015. Bidirectional LSTM-CRF Models for Sequence Tagging. *ArXiv*, abs/1508.01991.
- Jia, M.; Shen, L.; Shen, X.; and et al. 2023. MNER-QG: An End-to-End MRC Framework for Multimodal Named Entity Recognition with Query Grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8032–8040. Association for the Advancement of Artificial Intelligence.
- Jia, M.; Shen, X.; Shen, L.; and et al. 2022. Query Prior Matters: A MRC Framework for Multimodal Named Entity Recognition. In *Proceedings of the 30th ACM International Conference on Multimedia*, 3549–3558. Association for Computing Machinery.
- Krishna, R.; Zhu, Y.; Groth, O.; and et al. 2017. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *Int. J. Comput. Vision*, 32–73.
- Lample, G.; Ballesteros, M.; Subramanian, S.; and et al. 2016. Neural Architectures for Named Entity Recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 260–270. Association for Computational Linguistics.
- Li, J.; Sun, A.; Han, J.; and et al. 2022. A Survey on Deep Learning for Named Entity Recognition. *IEEE Transactions on Knowledge and Data Engineering*, 50–70.
- Liu, P.; Wang, G.; Li, H.; and et al. 2022. Multi-Granularity Cross-Modality Representation Learning for Named Entity Recognition on Social Media. *ArXiv*, abs/2210.14163.
- Lu, D.; Neves, L.; Carvalho, V.; and et al. 2018. Visual Attention Model for Name Tagging in Multimodal Social Media. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 1990–1999. Association for Computational Linguistics.
- Lu, J.; Zhang, D.; Zhang, J.; and et al. 2022. Flat Multimodal Interaction Transformer for Named Entity Recognition. In *Proceedings of the 29th International Conference on Computational Linguistics*, 2055–2064. International Committee on Computational Linguistics.
- Ma, X.; and Hovy, E. 2016. End-to-end Sequence Labeling via Bi-directional LSTM-CNNs-CRF. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 1064–1074. Association for Computational Linguistics.
- Moon, S.; Neves, L.; and Carvalho, V. 2018. Multimodal Named Entity Recognition for Short Social Media Posts. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 852–860. Association for Computational Linguistics.

- Perera, I.; Hwang, J. D.; Bayas, K.; and et al. 2018. Cyber-attack Prediction Through Public Text Analysis and Mini-Theories. *2018 IEEE International Conference on Big Data (Big Data)*, 3001–3010.
- Radford, A.; Kim, J. W.; Hallacy, C.; and et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, 8748–8763. PMLR.
- Ren, S.; He, K.; Girshick, R.; and et al. 2017. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1137–1149.
- Ritter, A.; Wright, E.; Casey, W.; and et al. 2015. Weakly Supervised Extraction of Computer Security Events from Twitter. In *Proceedings of the 24th International Conference on World Wide Web*, 896–905. International World Wide Web Conferences Steering Committee.
- Sun, L.; Wang, J.; Su, Y.; and et al. 2020. RIVA: A Pre-trained Tweet Multimodal Model Based on Text-image Relation for Multimodal NER. In *Proceedings of the 28th International Conference on Computational Linguistics*, 1852–1862. International Committee on Computational Linguistics.
- Sun, L.; Wang, J.; Zhang, K.; and et al. 2021. RpBERT: A Text-image Relation Propagation-based BERT Model for Multimodal NER. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, 13860–13868. AAAI Press.
- Vaswani, A.; Shazeer, N.; Parmar, N.; and et al. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6000–6010. Curran Associates, Inc.
- 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. Association for Computational Linguistics.
- Wang, J.; Tang, J.; and Luo, J. 2020. Multimodal Attention with Image Text Spatial Relationship for OCR-Based Image Captioning. In *Proceedings of the 28th ACM International Conference on Multimedia*, 4337–4345. Association for Computing Machinery.
- Wang, X.; Gui, M.; Jiang, Y.; and et al. 2022a. ITA: Image-Text Alignments for Multi-Modal Named Entity Recognition. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3176–3189. Association for Computational Linguistics.
- Wang, X.; Tian, J.; Gui, M.; and et al. 2022b. PromptMNER: Prompt-Based Entity-Related Visual Clue Extraction and Integration for Multimodal Named Entity Recognition. In *Database Systems for Advanced Applications: 27th International Conference, DASFAA 2022*, 297–305. Springer-Verlag.
- Wang, X.; Ye, J.; Li, Z.; and et al. 2022c. CAT-MNER: Multimodal Named Entity Recognition with Knowledge-Refined Cross-Modal Attention. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6.
- Wu, Z.; Zheng, C.; Cai, Y.; and et al. 2020. Multimodal Representation with Embedded Visual Guiding Objects for Named Entity Recognition in Social Media Posts. In *Proceedings of the 28th ACM International Conference on Multimedia*, 1038–1046. Association for Computing Machinery.
- Xu, B.; Huang, S.; Du, M.; and et al. 2022a. Different Data, Different Modalities! Reinforced Data Splitting for Effective Multimodal Information Extraction from Social Media Posts. In *Proceedings of the 29th International Conference on Computational Linguistics*, 1855–1864. International Committee on Computational Linguistics.
- Xu, B.; Huang, S.; Sha, C.; and et al. 2022b. MAF: A General Matching and Alignment Framework for Multimodal Named Entity Recognition. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 1215–1223. Association for Computing Machinery.
- Yang, S.; Li, G.; and Yu, Y. 2021. Relationship-Embedded Representation Learning for Grounding Referring Expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2765–2779.
- Yao, T.; Pan, Y.; Li, Y.; and et al. 2018. Exploring Visual Relationship for Image Captioning. In *Computer Vision – ECCV 2018*, 711–727. Springer International Publishing.
- Yu, J.; Jiang, J.; Yang, L.; and et al. 2020. Improving Multimodal Named Entity Recognition via Entity Span Detection with Unified Multimodal Transformer. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 3342–3352. Association for Computational Linguistics.
- Zhang, D.; Wei, S.; Li, S.; and et al. 2021. Multi-modal Graph Fusion for Named Entity Recognition with Targeted Visual Guidance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 14347–14355.
- Zhang, Q.; Fu, J.; Liu, X.; and et al. 2018. Adaptive co-attention network for named entity recognition in tweets. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 5674–5681. AAAI Press.
- Zhao, F.; Li, C.; Wu, Z.; and et al. 2022a. Learning from Different Text-Image Pairs: A Relation-Enhanced Graph Convolutional Network for Multimodal NER. 3983–3992. Association for Computing Machinery.
- Zhao, G.; Dong, G.; Shi, Y.; and et al. 2022b. Entity-level Interaction via Heterogeneous Graph for Multimodal Named Entity Recognition. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 6345–6350. Association for Computational Linguistics.