

# HAHE: Hierarchical Attention for Hyper-Relational Knowledge Graphs in Global and Local Level

Haoran Luo<sup>1</sup>, Haihong E<sup>1\*</sup>, Yuhao Yang<sup>2</sup>, Yikai Guo<sup>3</sup>, Mingzhi Sun<sup>1</sup>,  
Tianyu Yao<sup>1</sup>, Zichen Tang<sup>1</sup>, Kaiyang Wan<sup>1</sup>, Meina Song<sup>1</sup>, Wei Lin<sup>4</sup>

<sup>1</sup>School of Computer Science, Beijing University of Posts and Telecommunications, China

<sup>2</sup>School of Automation Science and Electrical Engineering, Beihang University, China

<sup>3</sup>Beijing Institute of Computer Technology and Application, China

<sup>4</sup>Inspur Group Co., Ltd., China

{luohaoran, ehaihong}@bupt.edu.cn

## Abstract

Link Prediction on Hyper-relational Knowledge Graphs (HKG) is a worthwhile endeavor. HKG consists of hyper-relational facts (H-Facts), composed of a main triple and several auxiliary attribute-value qualifiers, which can effectively represent factually comprehensive information. The internal structure of HKG can be represented as a hypergraph-based representation globally and a semantic sequence-based representation locally. However, existing research seldom simultaneously models the graphical and sequential structure of HKGs, limiting HKGs' representation. To overcome this limitation, we propose a novel Hierarchical Attention model for HKG Embedding (HAHE), including global-level and local-level attention. The global-level attention can model the graphical structure of HKG using hypergraph dual-attention layers, while the local-level attention can learn the sequential structure inside H-Facts via heterogeneous self-attention layers. Experiment results indicate that HAHE achieves state-of-the-art performance in link prediction tasks on HKG standard datasets. In addition, HAHE addresses the issue of HKG multi-position prediction for the first time, increasing the applicability of the HKG link prediction task. Our code is publicly available<sup>1</sup>.

## 1 Introduction

Knowledge graphs (KGs) are semantic networks that define entity relationships. Early KG research (Bordes et al., 2013; Sun et al., 2019; Balazevic et al., 2019) use binary relationships, often expressed as a triple-based fact (*subject, relation, object*). Yet, n-ary relational facts (containing more than two entities) are abundant in real-world KGs like Freebase (Bollacker et al., 2008) and Wikidata (Vrandečić and Krötzsch, 2014). Rosso et al. (2020) represent an n-ary relational fact as a

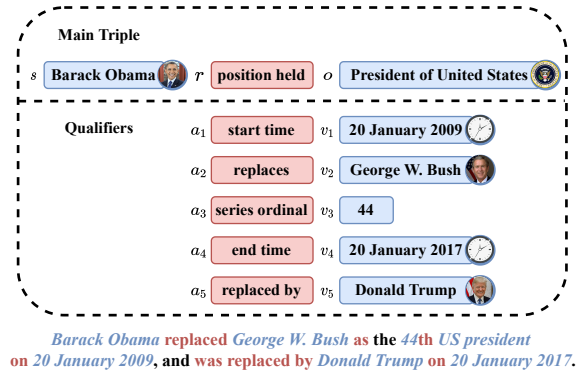


Figure 1: An example of hyper-relational fact structure.

hyper-relational fact (H-Fact) consisting of a main triple( $s,r,o$ ) and several auxiliary attribute-value qualifiers  $\{(a_i :v_i)\}$ , and KGs composed of H-Facts are called hyper-relational knowledge graphs (HKGs).

As shown in Figure 1, an H-fact can describe a real-world fact. Unlike traditional triple-based facts, H-Facts do not just raise the number of entities in facts from two to  $n$ . It structurally and effectively represents the  $n$ -ary relational facts prevalent in reality. Globally, it extends ordinary graph structure to hypergraph (Zhou et al., 2006) structure. Locally, it defines five heterogeneous roles of  $s,r,o,a,v$  within facts to capture the semantic information of the fact 'Barack Obama held position as US president', as illustrated in Figure 2.

Recent research has demonstrated various embedding strategies for hyper-relational representations. However, current approaches only consider global hypergraph structures or local semantic sequence structures. For instance, StarE (Galkin et al., 2020) employs the information transfer function of graph neural networks (GNN) to unidirectionally pass auxiliary key-value pair information into the main triples' relations, thereby capturing the graph structure but insufficiently between multiple entities and relations within the H-facts. In

\*Corresponding author.

<sup>1</sup><https://github.com/LHRLAB/HAHE>

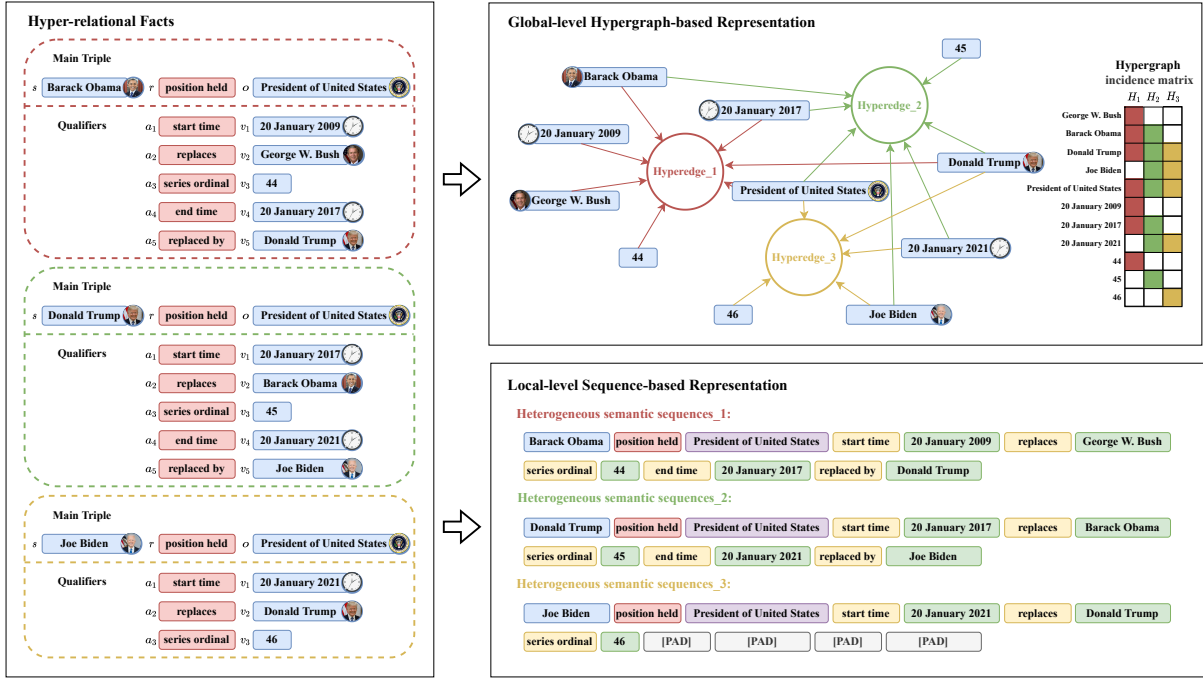


Figure 2: The Global-level Hypergraph-based representation and Local-level Sequence-based representation based on three examples of H-Facts in HKGs.

contrast, GRAN (Wang et al., 2021) initially incorporates the Transformer encoder (Vaswani et al., 2017) into the HKG embedding, capturing the fully connected semantic information locally inside H-facts, while disregarding the global structure. Consequently, representing the global and local structure of HKG simultaneously with hierarchical attention becomes a promising research direction, but an inadequate representation of HKG structure constrains HKG embeddings.

To overcome this limitation, we propose a novel **Hierarchical Attention model for HKG Embedding (HAHE)** that incorporates global-level and local-level attention. We update the global node embeddings using the HKG hypergraph structure. However, by complete connectivity, the previous hypergraph attention network (Bai et al., 2021) just converts all hypergraph nodes into a regular graph and then utilizes the GAT (Veličković et al., 2018) layer for node embedding updates, rendering it unable to distinguish which nodes comprise a hyperedge. Consequently, we design hypergraph dual-attention layers to aggregate node embedding information into hyper-edge embedding through the attention mechanism. After obtaining the hyper-edge embedding, we update the node embedding by feeding it back to the node through the attention mechanism. In this way, nodes are allowed to learn more distant

information from the whole HKG. This hypergraph dual-attention method significantly enhances learning capacity. It then transfers the updated node information to the local level’s attention. Inspired by GRAN’s heterogeneous attention (Wang et al., 2021), we define five types of nodes and fourteen types of edges in a single H-Fact and develop heterogeneous self-attention layers with both node-bias and edge-bias attention to learn the semantic content of H-Facts. The last step is to output the link prediction findings using an MLP-based decoding process for one-position or multi-position link prediction tasks on HKGs.

Experiments on link prediction were performed on three HKG standard datasets, JF17K (Wen et al., 2016), Wikipeople (Guan et al., 2019), and WD50K (Galkin et al., 2020). The state-of-the-art results indicate that HAHE is effective in the link prediction task. In addition, adequate ablation experiments were designed to highlight the importance of global and local focus, and HAHE is also used for the HKG multi-position prediction task, i.e., predicting two or more entities or relations simultaneously in a single H-fact, hence increasing the applicability of the HKG link prediction task. Ultimately, we make our code publicly available and discuss the limitations and future work of HKG embedding representation.

## 2 Related Work

Early approaches consider hyper-relational fact(H-Fact) mainly as graph structure, focusing more on the topological relations of entities. For example, m-TransH (Wen et al., 2016) projects the entities onto the relation hyperplane. RAE, NaLP, NeuInfer, and N-TuckER (Zhang et al., 2018; Guan et al., 2019, 2020; Liu et al., 2020) optimize the method of the Hyper-relational knowledge graph (HKG) embedding based on m-TransH. However, none of them adopts the hyper-relational structure.

HINGE (Rosso et al., 2020) firstly proposes the attribute-value qualifiers for embedding hyper-relational representation using CNN, and Hyper2 (Yan et al., 2022) initializes the relation and entity in the Poincaré ball vectors to improve the model accuracy. Yet, neither of these methods considers graphical structure or semantic sequences of HKGs.

StarE (Galkin et al., 2020) employs GNN as the message-passing mechanism to encode entities and relationships, while Transformer is the decoder to get the result. HyTransformer (Yu and Yang, 2021) applies layer normalization and dropout methods to replace StarE’s encoder. MSeaHKG (Di and Chen, 2022) proposed that a message-passing function significantly impacts the model performance, so it replaced the static message-passing function in StarE with a dynamic one. These models consider the graph structure within the hyper-relational facts, but disregard the semantic sequences.

GRAN (Wang et al., 2021) is an improvement for Transformer (Vaswani et al., 2017). It replaces Transformer’s self-attention with edge-biased fully-connected attention and accurately collects semantic information. Despite this, it ignores the graph structure.

Unlike earlier models, HAHE considers graph structure and semantic sequences simultaneously and employs hierarchical attention for link prediction on HKGs. Moreover, it is the first to improve previous methods by modeling the structure of HKG via global-level embedding and can perform multi-position prediction.

## 3 Preliminaries

This section presents important concepts and techniques in Hyper-relational Knowledge Graph (HKG), including definitions of HKGs, hypergraph learning, global and local structure of HKGs, and multi-position prediction on HKGs.

### 3.1 Hyper-relational Knowledge Graphs

HKGs comprise hyper-relational facts (H-Facts). Typically, an H-Fact can be represented as  $\mathcal{H} = \{((s, r, o), \{(a_i : v_i)\}_{i=1}^m) | s, o, v_1, \dots, v_m \in \mathcal{E}, r, a_1, \dots, a_m \in \mathcal{R}\}$ , where  $(s, r, o)$  represents the main triple and  $\{(a_i : v_i)\}_{i=1}^m$  represents  $m$  auxiliary attribute-value qualifiers.

The link prediction (LP) task on HKGs is to predict missing elements from H-Facts, where missing elements can be entities  $\in \{s, o, v_1, \dots, v_m\}$  or relations  $\in \{r, a_1, \dots, a_m\}$ .

### 3.2 Hypergraph Learning on HKGs

Since there are more than two entities in an H-Fact, we introduce hypergraph learning (Feng et al., 2019). A hypergraph of HKG,  $\mathcal{G}_H = \{\mathcal{E}_H, \mathcal{H}_H, \mathcal{I}_H\}$ , contains a node set  $\mathcal{E}_H = \mathcal{E}$ , a hyperedge set  $\mathcal{H}_H = \mathcal{H}$ , and a special incidence matrix  $\mathcal{I}_H$  that records the weights of each hyperedge.  $\mathcal{I}_H$  is a  $|\mathcal{E}_H| \times |\mathcal{H}_H|$  matrix defined as follows:

$$\begin{aligned} h(v, e) &= 1, \text{ if } v \in e, \\ h(v, e) &= 0, \text{ if } v \notin e, \end{aligned} \quad (1)$$

where  $v \in \mathcal{E}_H, e \in \mathcal{H}_H$ .  $h$  is a function to represent the value in  $\mathcal{I}_H$ . For a node  $v \in \mathcal{E}_H$ , its degree is defined as  $d(v) = \sum_{e \in \mathcal{H}_H} h(v, e)$ , which represents the number of times that a node (entity) appears in different hyperedges (H-Facts) in the whole HKG.

### 3.3 Multi-position Prediction on HKGs

Multi-position prediction is a new meaningful task with more practicality than the one-position link prediction task on HKGs. For HKG link prediction, in one main triple, we can predict another element for every two of them we know, i.e., we can predict  $(s, r, ?), (s, ?, o), (? , r, o)$ . In one auxiliary attribute-value qualifier, if we know any of the attributes and values of one of the auxiliary attribute-value qualifiers, we can predict the other, i.e.,  $(a, ?), (? , v)$ . Thus in practical link prediction, there are some problems have two or more prediction points for example  $(s, r, ?, a_1, ?, a_2, v_2)$  (prediction position one in the main triple and the other in the first attribute-value qualifier) or  $(s, r, o, a_1, ?, a_2, ?)$  (both predicted positions are in the auxiliary attribute-value qualifiers). We refer to tasks with two or more prediction positions in link prediction tasks as multi-position prediction tasks on HKGs.

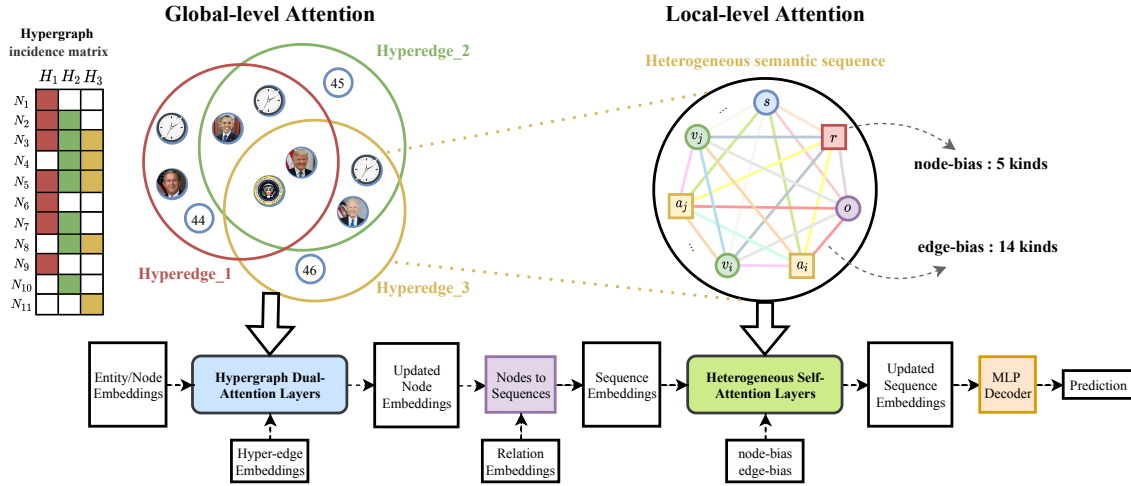


Figure 3: The overview of HAHE model for Global-level and Local-level Representation of HKGs.

## 4 Methodology

This section introduces our hyper-relational knowledge graph (HKG) embedding model HAHE, including global and local representation, two hierarchical attention layers, and MLP decoder.

### 4.1 Global and Local Representation

HKGs  $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{H}\}$  consist of multiple hyper-relational facts (H-Facts)  $\mathcal{H}$  with entities  $\mathcal{E}$  and relations  $\mathcal{R}$ . Since each H-Fact represents an n-ary relation ( $n > 2$ ) and has rich, heterogeneous semantic information, we model HKGs in terms of global-level hypergraph-based representation and local-level sequence-based representation with hypergraph dual-attention layers and heterogeneous self-attention layers respectively. The overview of HAHE is illustrated in Figure 3. For global-level representation, we define  $\mathcal{G}_H = \{\mathcal{E}_H, \mathcal{H}_H, \mathcal{I}_H\}$  to represent the graph structure of entities. Unlike the regular graph, the hyperedges can connect more than two entity nodes. Moreover, we use incidence matrix  $\mathcal{I}_H$  to represent the association information of nodes and hyperedges. For local-level representation, every H-Fact has the structure of one main triple and several auxiliary attribute-value qualifiers  $\mathcal{H} = \{((s, r, o), \{(a_i : v_i)\}_{i=1}^m) | s, o, v_1, \dots, v_m \in \mathcal{E}, r, a_1, \dots, a_m \in \mathcal{R}\}$ , which represents the semantic information of facts. We can fully connect entities and relations in H-Facts and represent them as heterogeneous semantic sequence structure, containing five kinds of nodes  $s, r, o, a, v$  and 14 kinds of edges  $s-r, s-o, r-o, s-a, s-v, r-a, r-v, o-a, o-v, a_i-a_j, o_i-o_j, a_i-o_i, a_i-o_j$ , where  $i, j$  are the serial numbers of different qualifiers.

### 4.2 Hypergraph Dual-Attention Layers

As shown in Figure 4(a), entity embedding first utilizes Hypergraph Dual-Attention Layers to learn hypergraph structural information in global level. Previous Hypergraph Attention Network methods created a transformed ordinary graph by full joining nodes within the same hyperedge and applying GAT. The hypergraph representation loses because the ordinary graph after this transformation cannot distinguish whether two nodes are within the same hyperedge or different hyperedges. We first initialize the entities as nodes with embedding as  $\mathbf{h}_{v_i} \in \mathbb{R}^d$ , where  $v_i \in \mathcal{E}_H$  and  $d$  is dimension of embedding, and initialize the H-Facts as hyperedges with embedding as  $\mathbf{h}_{e_i} \in \mathbb{R}^d$ , where  $e_i \in \mathcal{H}_H$ . Then the node and hyperedge embeddings are projected into the same space to obtain  $\mathbf{W}\mathbf{h}_{v_i}$  and  $\mathbf{W}\mathbf{h}_{e_i}$ , where  $\mathbf{W} \in \mathbb{R}^{d \times d}$ . The attention from nodes to hyper-edges (N-to-H Attention) is performed as follows:

$$\alpha_{ij} = \text{att}(\mathbf{W}\mathbf{h}_{e_i}, \mathbf{W}\mathbf{h}_{v_j}) | v_j \in e_i, \quad (2)$$

where  $\alpha_{ij}$  indicates the importance of node  $v_j$ 's features to hyperedge  $e_i$ ,  $\text{att}$  is N-to-H Attention function where we choose a single-layer neural network with concatenation operation.  $v_j \in e_i$  means we only calculate the attention where node  $v_j$  is in hypergraph  $e_i$ , which is indexed by hypergraph incidence matrix  $\mathcal{I}_H$ . Then, the information of nodes is aggregated to hyperedges:

$$\tilde{\mathbf{h}}_{e_i} = \sigma \left( \sum_{v_j \in e_i} \frac{\exp(\text{LR}(\alpha_{ik}))}{\sum_{v_k \in e_i} \exp(\text{LR}(\alpha_{ik}))} \mathbf{W}\mathbf{h}_{v_j} \right), \quad (3)$$

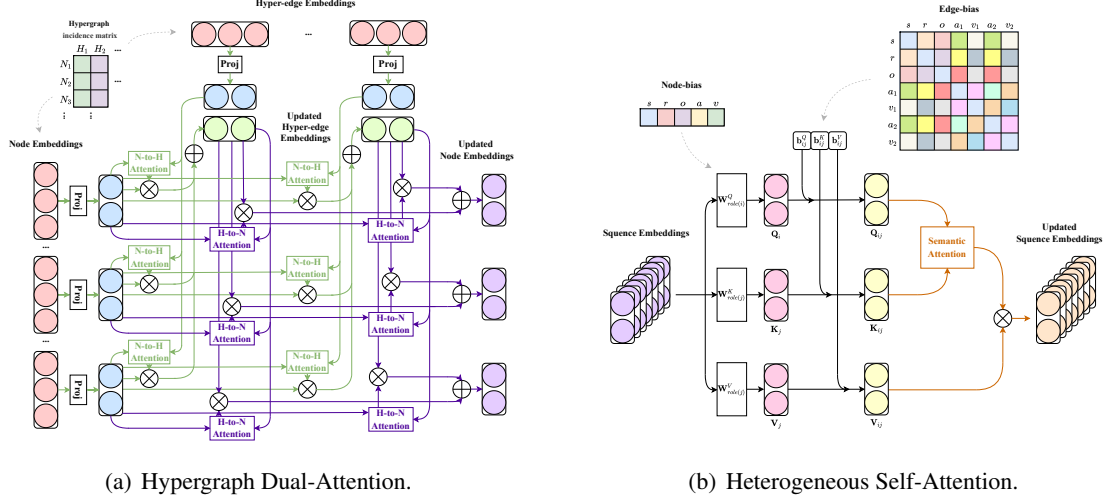


Figure 4: The structure of Hypergraph Dual-Attention Layers and Heterogeneous Self-Attention Layers in HAHE.

where  $\tilde{\mathbf{h}}_{e_i} \in \mathbb{R}^d$  is updated hyperedge embeddings and  $LR, \sigma$  are activation functions. After that, we use a similar way to aggregate the information of updated hyperedges back to nodes with Attention (H-to-N Attention) as follows:

$$\beta_{ij} = att(\mathbf{W}\mathbf{h}_{v_i}, \mathbf{W}\mathbf{h}_{e_j}) | e_j \in v_i, \quad (4)$$

$$\tilde{\mathbf{h}}_{v_i} = \sigma \left( \sum_{e_j \in v_i} \frac{\exp(LR(\beta_{ik}))}{\sum_{e_k \in v_i} \exp(LR(\beta_{ik}))} \tilde{\mathbf{h}}_{e_j} \right), \quad (5)$$

where  $\beta_{ij}$  denotes the importance between updated hyperedge  $e_j$  and node  $v_i$ , and  $\tilde{\mathbf{h}}_{v_i}$  is the updated node embedding. This way, we update the global hypergraph entity embedding and achieve node-hyperedge-node dual-attention message passing. Though we introduce more hyperedge embeddings than ordinary GNNs, they are only used as an intermediate weight variable for hypergraph attention computation. PyG makes dual-attention easy to implement, making it scalable to large graphs as GNNs.

After hypergraph dual-attention layers, H-Facts distribute updated node embeddings to sequence embeddings with relation embeddings, and fed them into Heterogeneous Self-Attention Layers.

### 4.3 Heterogeneous Self-Attention Layers

Each element in the sequence  $\mathbf{x}_i \in \mathbb{R}^d$  has five roles, including  $s, r, o, a, v$ , and a total of 14 kinds of edge with other elements in the sequence  $\mathbf{x}_j \in \mathbb{R}^d$ . So, as shown in Figure 4(b), we design a heterogeneous self-attentive layer with both **node-bias**

and **edge-bias** to learn the local semantic information of the H-Facts. The elements in the sequence pass through this attention layer and update the embedding as follows:

$$\gamma_{ij} = \frac{(\mathbf{W}_{role(i)}^Q \mathbf{x}_i + \mathbf{b}_{ij}^Q)^\top (\mathbf{W}_{role(j)}^K \mathbf{x}_j + \mathbf{b}_{ij}^K)}{\sqrt{d}}, \quad (6)$$

$$\tilde{\mathbf{x}}_i = \sum_{j=1}^n \frac{\exp(\gamma_{ik})}{\sum_{k=1}^n \exp(\gamma_{ik})} (\mathbf{W}_{role(j)}^V \mathbf{x}_j + \mathbf{b}_{ij}^V), \quad (7)$$

where  $\gamma_{ij}$  is the importance between one element in sequence  $\mathbf{x}_i$  and another  $\mathbf{x}_j$ ,  $\mathbf{W}_{role(i)}^Q, \mathbf{W}_{role(i)}^K, \mathbf{W}_{role(i)}^V \in \mathbb{R}^{d \times d}$  are the linear weight metrics of query, key, value, and five different kinds of  $\mathbf{x}_i$  pass through different weight metrics indexed by  $role$  function as the node-bias, and  $\mathbf{b}_{ij}^Q, \mathbf{b}_{ij}^K, \mathbf{b}_{ij}^V \in \mathbb{R}^d$  are designed as the edge-bias. Then the sequence embeddings are updated by learning the semantic information inside the H-Facts as  $\tilde{\mathbf{x}}_i \in \mathbb{R}^d$ .

### 4.4 MLP Decoder

Finally, the updated sequence embeddings are selected for the embedding at the position to be predicted  $\tilde{\mathbf{x}}_p$  with MLP decoder to get the prediction distribution and obtain the link prediction results.

$$\mathbf{P} = softmax(MLP(\tilde{\mathbf{x}}_p)\mathbf{E}^\top), \quad (8)$$

where  $MLP: \mathbb{R}^d \rightarrow \mathbb{R}^d$ ,  $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times d}$  shares the initial element embedding matrix,  $\mathbf{P} \in \mathbb{R}^{|\mathcal{E}|}$  is obtained after softmax operation, which denotes the similarity probability of  $\tilde{\mathbf{x}}_p$  with each element in HKG for obtaining the link prediction answers.

Dataset	H-Facts	H-Facts with Q(%)	Entities	Relations	Train	Valid	Test	Arity
JF17K	100,947	46,320(45.9%)	28,645	501	76,379	-	24,568	2-6
WikiPeople	369,866	9,482(2.6%)	34,839	178	294,439	37,715	37,712	2-7
WD50K	236,507	32,167(13.6%)	47,156	532	166,435	23,913	46,159	2-67

Table 1: Dataset statistics, where the columns respectively indicate the number of all H-Facts, H-facts with qualifiers, entities, relations, H-facts in train/valid/test sets, and the range of arity of H-facts.

## 4.5 Learning Strategy

The model trains through the final loss, which is calculated by the similarity between the target of prediction and all entities:

$$\mathcal{L} = \sum_{t=1}^{|\mathcal{E}|} \mathbf{y}_t \log \mathbf{P}, \quad (9)$$

where  $\mathbf{y}_t$  is the  $t$ -th entry of the label  $\mathbf{y}$ .

For Multi-position Prediction, due to the fully connected attention mechanism, HAHE can accomplish the multi-position prediction task of HKG by masking two or more entities or relations at two or more positions in the same hyper-relational fact. After passing the MLP encoder, HAHE can get the prediction value  $\mathbf{P}_i (i \geq 2)$  of the corresponding positions, and find the entity or relation with the highest similarity among all HKG elements as the prediction results respectively.

## 5 Experiments

This section introduces the experimental settings, results and analysis. We answer the following research questions (RQs). **RQ1:** Can HAHE outperform other Hyper-relational Knowledge Graphs (HKG) embedding models on HKG datasets? **RQ2:** How does the hierarchical attention mechanism contribute to HAHE? **RQ3:** How do hypergraph dual-attention mechanisms contribute to HAHE in global level? **RQ4:** How do heterogeneity of nodes and edges in hyper-relation fact contribute to HAHE in local level? **RQ5:** How HAHE performs in multi-position prediction tasks on HKGs?

### 5.1 Experimental Setup

#### 5.1.1 Datasets

We conduct experiments on three hyper-relational datasets **JF17K** (Wen et al., 2016), **WikiPeople** (Guan et al., 2019), and **WD50K** (Galkin et al., 2020), respectively, as shown in Table 1. Among them, JF17K is extracted from Freebase (Bollacker et al., 2008). WikiPeople is obtained by filtering out the statements containing literals in the original

WikiPeople dataset, derived from Wikidata (Vrandečić and Krötzsch, 2014) concerning entities of type human, and WD50K is a high-quality hyper-relational dataset with richer hyper-relational facts with auxiliary attribute-value qualifiers.

#### 5.1.2 Baselines

We compare HAHE against a sizable collection of previous hyper-relational approaches namely: (i) m-TransH (Wen et al., 2016) (ii) RAE (Zhang et al., 2018) (iii) NaLP (Guan et al., 2019) (iv) NeuInfer (Guan et al., 2020) (v) HINGE (Rosso et al., 2020) (vi) StarE (Galkin et al., 2020) (vii) Hyper2 (Yan et al., 2022) (viii) HyTransformer (Yu and Yang, 2021) (ix) GRAN (Wang et al., 2021) (x) MSeaHKG (Di and Chen, 2022).

#### 5.1.3 Ablations

To evaluate the significance of HAHE’s three main modules, hypergraph dual-attention mechanism, node heterogeneity, and edge heterogeneity, we obtain 7 simplified model variants by removing any one or two modules from the full model (**HAHE-node**, **HAHE-edge**, **HAHE-node&edge**, **HAHE-global**, **HAHE-global&node**, **HAHE-global&edge**), and the **basic** variant by removing all three modules.

#### 5.1.4 Evaluation Metrics

Each model predicts entities and relations separately. We split each task of predictions into subject/object prediction in main triples and all entities prediction in whole H-facts to test the model’s main triple prediction ability. MRR (the average of reciprocal rankings) and Hits@K (the proportion of top K rankings) for K=1,10 are used to evaluate each link prediction task.

#### 5.1.5 Hyperparameters and Environment

The model was trained for 300 epochs using the Adam optimizer with a batch size of 1024 examples across 1 GeForce GTX 1080Ti on each dataset. Appendix A shows HAHE’s optimal hyperparameter settings. Appendix B shows training details.

Model	JF17K						Wikipedia						WD50K						
	subject / object			all entities			subject / object			all entities			subject / object			all entities			
	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	
m-TransH	0.206	0.206	0.462	0.102	0.069	0.168	0.063	0.063	0.300	-	-	-	-	-	-	-	-	-	-
RAE	0.215	0.215	0.466	0.310	0.219	0.504	0.058	0.058	0.306	0.172	0.102	0.320	-	-	-	-	-	-	-
NaLP	0.221	0.165	0.331	0.366	0.290	0.516	0.408	0.331	0.546	0.338	0.272	0.466	-	-	-	0.224	0.158	0.330	-
NeuInfer	0.449	0.361	0.624	0.473	0.397	0.618	0.476	0.415	0.585	0.333	0.259	0.477	0.243	0.176	0.377	0.228	0.162	0.341	-
HINGE	0.431	0.342	0.611	0.517	0.436	0.675	0.342	0.272	0.463	0.350	0.282	0.467	-	-	-	0.232	0.164	0.343	-
StarE	0.574	0.496	0.725	0.542	0.454	0.685	0.491	0.398	0.592	0.378	0.265	0.542	0.349	0.271	0.496	-	-	-	-
Hyper2	0.583	0.500	0.746	-	-	-	0.461	0.391	0.597	-	-	-	-	-	-	-	-	-	-
HyTransformer	0.582	0.501	0.742	-	-	-	0.501	0.426	0.634	-	-	-	0.356	0.281	0.498	-	-	-	-
GRAN	0.617	0.539	0.770	0.656	0.582	0.799	0.503	0.438	0.620	0.479	0.410	0.604	-	-	-	0.309	0.24	0.441	-
MSeaHKG	-	-	-	0.577	0.481	0.711	-	-	-	0.395	0.291	0.554	-	-	-	-	-	-	-
HAHE	<b>0.623</b>	<b>0.554</b>	<b>0.806</b>	<b>0.668</b>	<b>0.597</b>	<b>0.816</b>	<b>0.509</b>	<b>0.447</b>	<b>0.639</b>	<b>0.495</b>	<b>0.420</b>	<b>0.631</b>	<b>0.368</b>	<b>0.291</b>	<b>0.516</b>	<b>0.402</b>	<b>0.327</b>	<b>0.546</b>	-
HAHE w/o global	0.621	0.548	0.787	0.659	0.588	0.797	0.501	0.434	0.629	0.483	0.407	0.612	0.356	0.280	0.501	0.390	0.315	0.531	-
HAHE w/o node-bias	0.620	0.546	0.787	0.659	0.587	0.797	0.474	0.429	0.622	0.487	0.405	0.611	0.354	0.283	0.506	0.396	0.313	0.529	-
HAHE w/o edge-bias	0.620	0.545	0.786	0.657	0.586	0.796	0.503	0.435	0.629	0.483	0.407	0.612	0.357	0.281	0.513	0.391	0.316	0.533	-

Table 2: Comparison of HAHE with other models, composed of entity prediction accuracy on JF17K, WikiPeople and WD50K. Results of the models are mainly taken from the original paper. Best results in each tasks are in **bold**.

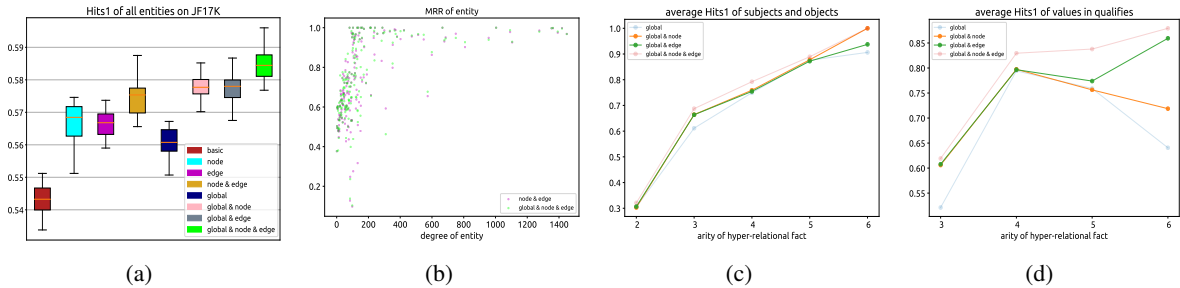


Figure 5: Ablation results. (a) Hits@1 results for ablation study with all HAHE variants. (b) MRR results with different degrees of entity for global-level analysis. (c) Hits@1 results of subject/object prediction with different arity for local-level analysis. (d) Hits@1 results of prediction in qualifiers with different arity for local-level analysis.

## 5.2 Main Results (RQ1)

In this experiment, we evaluate our model on the link prediction task. For entity prediction, the results of our model and each variant of our model can be found in Table 2. For relation prediction, the result is shown in Appendix C. We can observe that the HAHE outperforms the other current methods on all three datasets. On JF17K, for the prediction of subject and object, HAHE reports an improvement of 0.6 (0.9%) MRR points, 1.5 (2.7%) H@1, and 3.6 (4.6%) H@10 compared with the best approach. For the prediction of all entities, HAHE reports a gain of 1.2 (1.8%) MRR, 1.5 (2.5%) H@1, and 1.7 (2.1%) H@10 compared with the next-best approach. For the other two datasets, we also have different degrees of improvement. For WD50K, the latest high-quality HKG dataset, our model has the largest improvement over the existing SOTA model GRAN, with an MRR improvement of about 10 points, which proves that this model is more suitable for hypergraph-structured knowledge graphs with hyper-relational facts beyond binary relation.

## 5.3 Ablation Study (RQ2)

The hypergraph dual-attention mechanism, node heterogeneity, and edge heterogeneity are the three components of HAHE that are required for its operation. We evaluate seven different variants of HAHE, irrespective of whether or not each component is helpful. When evaluating each model variant with a variety of hyperparameters, the results of the optimal prediction were recorded. For different HAHE variants in Figure 5(a), it can be observed that hypergraph dual-attention, node heterogeneity, and edge heterogeneity all contribute to the accurate result of our complete model. In addition, we have outlined the specific results of three primary HAHE variants in Table 2. Each variant lacks a necessary component that is required. Through comparison, our experiment results intuitively demonstrated the effectiveness of HAHE.

Then, we did more refined ablation analysis to explore the significance of the hypergraph dual-attention mechanism in global level and heterogeneity of nodes and edges in local level, respectively.

## 5.4 Analysis of Hypergraph Dual-attention Mechanism in Global Level (RQ3)

We statistically displayed the entity evaluation results of JF17K by the degree of entities in the HKG hypergraph structure to investigate the hypergraph dual attention mechanism. As shown in Figure 5(b), the hypergraph dual-attention mechanism improves the prediction accuracy of entities with different degree. Due to the presence of the attention mechanism, the entity feature information will help in message passing of global information by the hyper-relational facts (hyperedges). For entities with higher degree, the hypergraph dual-attention mechanism can better capture the global features of the hyperrelational facts. Entities with fewer degrees have little impact on capturing global information, but entities with more degrees compensate for this and improve their prediction accuracy.

## 5.5 Analysis of Heterogeneity of Nodes and Edges in Local Level (RQ4)

The experimental results show that both node-bias and edge-bias with heterogeneity can better distinguish different types of entities in an H-fact and different types of relationships between them in local level. Moreover, we find that node-bias plays a more prominent role than edge-bias in the subject/object prediction tasks as shown in Figure 5(c), because the entity roles in the main triplet are more diverse than those in the qualifier. And on the element prediction task in qualifiers, edge-bias is better than node-bias as shown in Figure 5(d), because edge-bias can distinguish the relationship between corresponding attribute-value pairs  $(a_i, v_i)$  and non-corresponding attribute-value pairs  $(a_j, v_i)$  rather than node-bias.

## 5.6 Results of Multi-position Prediction Tasks on HKGs (RQ5)

JF17K, WikiPeople, and WD50K were applied to test our multi-position prediction model. As an example, in Figure 6, our model outputs the embedding for each position and calculates the joint probability distribution for each candidate answer tuple. Because the increase of predicted positions leads to more answer sets than predicted by the unit placement link, we set the evaluation threshold to keep only the higher scoring answer tuples to obtain the evaluation results.

As shown in Table 3, we evaluated 2-position prediction and 3-position prediction on three HKG

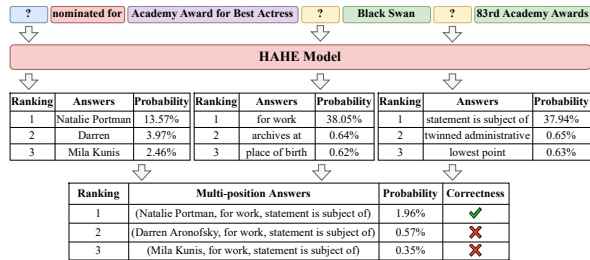


Figure 6: Case study of Multi-position Prediction on WD50K.

	JF17K			Wikipole			WD50K		
<b>2-position prediction</b>									
	av	sro/av	all	av	sro/av	all	av	sro/av	all
Ent-Ent	0.039	0.220	0.198	0.265	0.129	0.140	0.326	0.157	0.184
Ent-Rel	0.261	0.592	0.540	0.405	0.409	0.409	0.527	0.481	0.490
Rel-Rel	0.350	0.106	0.159	0.656	0.544	0.561	0.940	0.210	0.407
<b>3-position prediction</b>									
	av	sro/av	all	av	sro/av	all	av	sro/av	all
Ent-Ent	0.011	0.012	0.011	0.329	0.036	0.046	0.169	0.350	0.052
Ent-Rel	0.022	0.126	0.115	0.432	0.132	0.144	0.332	0.166	0.193
Rel-Rel	0.029	0.005	0.009	0.647	0.162	0.192	0.884	0.039	0.230

Table 3: The MRR results of Multi-position Prediction.

datasets, respectively, and used MRR to rank the answer tuples. The test set has three categories: Ent-Ent, Ent-Rel, and Rel-Rel. Ent-Ent and Rel-Rel indicate that all predicted locations are entities or relations, and Ent-Rel indicates that entities and relations are missing jointly. In addition, whether the main triple is complete divides the sample into two categories. "av" indicates that all predicted positions are in the auxiliary qualifiers, while "sro/av" indicates that the main triple has lost a position and others in qualifiers. "all" includes both categories. According to the results of multi-position prediction tasks, HAHE performs better on the high-quality WD50K dataset, and more positions predicted or one position in the main triple makes the prediction more difficult.

## 6 Conclusion

In this paper, we present HAHE, a model with hierarchical attention in global and local level. HAHE outperforms other baselines link prediction tasks on hyper-relational knowledge graphs (HKGs). The experimental results demonstrate that our hypergraph dual-attention layers and heterogeneous self-attention layers are effective in learning the global and local structure of HKGs. We also use HAHE to solve the HKG multi-location prediction task and analyze the results for the first time.



## Acknowledgements

This work is supported by the National Science Foundation of China (Grant No. 62176026) and Beijing Natural Science Foundation (M22009). This work is also supported by the BUPT Postgraduate Innovation and Entrepreneurship Project led by Haoran Luo.

## Limitations

For HKG one-position link prediction tasks, HAHE shows the best performance in all three datasets. However, because HAHE is based on hypergraph learning, it improves more on the WD50K high quality hyper-relational knowledge graph link prediction dataset, and less on the Wikipeople dataset where triples are the majority, so HAHE prefers the fact with more arity numbers. In the future, we will consider extending our approach to triples as a unified architecture.

For HKG multi-position link prediction tasks, it can be seen that our model is effective when predicting multiple missing auxiliary information, which is a frequent situation in practical applications. However, the prediction accuracy of our model needs to be further improved in the case of missing primary relations.

## Ethics Statement

This paper investigates the problem of knowledge graph link prediction, aiming at complementing incomplete hyper-relational knowledge graphs using deep learning methods to better promote knowledge graphs for assisted decision-making and intelligent question-and-answer applications. Therefore, we believe it does not violate any ethics.

## References

- Song Bai, Feihu Zhang, and Philip H.S. Torr. 2021. [Hypergraph convolution and hypergraph attention](#). *Pattern Recognition*, 110:107637.
- Ivana Balazevic, Carl Allen, and Timothy Hospedales. 2019. [Tucker: Tensor factorization for knowledge graph completion](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5185–5194, Hong Kong, China. Association for Computational Linguistics.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: A collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, page 1247–1250, New York, NY, USA. Association for Computing Machinery.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Shimin Di and Lei Chen. 2022. [Message function search for hyper-relational knowledge graph](#).
- Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. 2019. [Hypergraph neural networks](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):3558–3565.
- Mikhail Galkin, Priyansh Trivedi, Gaurav Maheshwari, Ricardo Usbeck, and Jens Lehmann. 2020. [Message passing for hyper-relational knowledge graphs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7346–7359, Online. Association for Computational Linguistics.
- Saiping Guan, Xiaolong Jin, Jiafeng Guo, Yuanzhuo Wang, and Xueqi Cheng. 2020. [NeuInfer: Knowledge inference on N-ary facts](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6141–6151, Online. Association for Computational Linguistics.
- Saiping Guan, Xiaolong Jin, Yuanzhuo Wang, and Xueqi Cheng. 2019. [Link prediction on n-ary relational data](#). In *The World Wide Web Conference, WWW '19*, page 583–593, New York, NY, USA. Association for Computing Machinery.
- Yu Liu, Quanming Yao, and Yong Li. 2020. [Generalizing tensor decomposition for n-ary relational knowledge bases](#). In *Proceedings of The Web Conference 2020, WWW '20*, page 1104–1114, New York, NY, USA. Association for Computing Machinery.
- Paolo Rosso, Dingqi Yang, and Philippe Cudré-Mauroux. 2020. [Beyond triplets: Hyper-relational knowledge graph embedding for link prediction](#). In *Proceedings of The Web Conference 2020, WWW '20*, page 1885–1896, New York, NY, USA. Association for Computing Machinery.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. [Rotate: Knowledge graph embedding by relational rotation in complex space](#). In *International Conference on Learning Representations*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). In *International Conference on Learning Representations*.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: A free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.
- Quan Wang, Haifeng Wang, Yajuan Lyu, and Yong Zhu. 2021. [Link prediction on n-ary relational facts: A graph-based approach](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 396–407, Online. Association for Computational Linguistics.
- Jianfeng Wen, Jianxin Li, Yongyi Mao, Shini Chen, and Richong Zhang. 2016. [On the representation and embedding of knowledge bases beyond binary relations](#). In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 1300–1307. IJCAI/AAAI Press.
- Shiyao Yan, Zequn Zhang, Xian Sun, Guangluan Xu, Li Jin, and Shuchao Li. 2022. [Hyper2: Hyperbolic embedding for hyper-relational link prediction](#). *Neurocomputing*, 492:440–451.
- Donghan Yu and Yiming Yang. 2021. [Improving hyper-relational knowledge graph completion](#). *CoRR*, abs/2104.08167.
- Richong Zhang, Junpeng Li, Jiajie Mei, and Yongyi Mao. 2018. [Scalable instance reconstruction in knowledge bases via relatedness affiliated embedding](#). In *Proceedings of the 2018 World Wide Web Conference, WWW '18*, page 1185–1194, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Dengyong Zhou, Jiayuan Huang, and Bernhard Schölkopf. 2006. [Learning with hypergraphs: Clustering, classification, and embedding](#). In *Advances in Neural Information Processing Systems*, volume 19. MIT Press.

## Appendix

### A Hyperparameter Settings

We use the grid search method to select the optimal hyperparameter settings for the network. The average Hits1 predicted by all entities is chosen as the evaluation metric. The hyperparameters that we can adjust and the possible values of the hyperparameters are first determined according to the structure of our model in Table 4.

Afterwards, the different hyperparameter choices are combined and the predictive metrics after 50 epochs of training are used to judge the merit of the hyperparameter combinations. The optimal hyperparameter combinations of the model are obtained by circular traversal of all hyperparameter combinations. The optimal hyperparameter combinations are shown in bold.

### B Model Training Details

We train 300 epochs in each dataset with the optimal combination of hyperparameters. The prediction results of the network are evaluated using the test sets with the best model. HAHE and all its variants have been trained on a single 11G 1080Ti GPU. Using our optimal hyperparameter settings, the time required to complete the training on the three datasets JF17K, Wikipeople, and WD50K is 5h, 12h and 8h, respectively.

### C Results of Relation Prediction

As shown in Table 5, the result of relation prediction of previous models for HKG embedding has achieved very good results in relation prediction, some indicators have even reached 99%.

Hyperparameter	JF17K	Wikepeople	WD50K
Embedding dimension	{128, <b>256</b> , 512, 1024}	{128, <b>256</b> , 512, 1024}	{128, <b>256</b> , 512, 1024}
Global layers	{0, 1, <b>2</b> , 3, 4}	{0, 1, <b>2</b> , 3, 4}	{0, 1, <b>2</b> , 3, 4}
Global dropout	{0.0, 0.1, <b>0.2</b> , 0.3, 0.4, 0.5}	{0.0, <b>0.1</b> , 0.2, 0.3, 0.4, 0.5}	{0.0, <b>0.1</b> , 0.2, 0.3, 0.4, 0.5}
Global activation	{ <i>relu</i> , <b>elu</b> , <i>gelu</i> , <i>tanh</i> }	{ <i>relu</i> , <b>elu</b> , <i>gelu</i> , <i>tanh</i> }	{ <i>relu</i> , <b>elu</b> , <i>gelu</i> , <i>tanh</i> }
Global attention heads	{ <b>4</b> , 8, 12, 16}	{ <b>4</b> , 8, 12, 16}	{ <b>4</b> , 8, 12, 16}
Local layers	{4, 8, <b>12</b> , 16, 24}	{4, 8, <b>12</b> , 16, 24}	{4, 8, <b>12</b> , 16, 24}
Local dropout	{0.0, 0.1, <b>0.2</b> , 0.3, 0.4, 0.5}	{0.0, <b>0.1</b> , 0.2, 0.3, 0.4, 0.5}	{0.0, <b>0.1</b> , 0.2, 0.3, 0.4, 0.5}
Local attention heads	{ <b>4</b> , 8, 12, 16}	{ <b>4</b> , 8, 12, 16}	{ <b>4</b> , 8, 12, 16}
Decoder activation	{ <i>relu</i> , <i>elu</i> , <b>gelu</b> , <i>tanh</i> }	{ <i>relu</i> , <i>elu</i> , <b>gelu</b> , <i>tanh</i> }	{ <i>relu</i> , <i>elu</i> , <b>gelu</b> , <i>tanh</i> }
Hidden size	{128, <b>256</b> , 512, 1024, 2048, 4096}	{128, <b>256</b> , 512, 1024, 2048, 4096}	{128, <b>256</b> , 512, 1024, 2048, 4096}
Batch size	{128, 256, 512, <b>1024</b> , 2046}	{128, 256, 512, <b>1024</b> , 2046}	{ <b>64</b> , 128, 256, 512, 1024}
Learning rate	{0.0001, <b>0.0005</b> , 0.001}	{0.0001, <b>0.0005</b> , 0.001}	{0.0001, <b>0.0005</b> , 0.001}
Weight decay	{ <b>0.01</b> , 0.02}	{ <b>0.01</b> , 0.02}	{ <b>0.01</b> , 0.02}
Soft label for entity	{0.5, 0.6, 0.7, 0.8, <b>0.9</b> }	{0.0, 0.1, <b>0.2</b> , 0.3, 0.4}	{0.0, 0.1, <b>0.2</b> , 0.3, 0.4}
Soft label for relation	{ <b>0.0</b> , 0.3, 0.6, 0.9}	{0.0, <b>0.1</b> , 0.2, 0.3}	{0.0, <b>0.1</b> , 0.2, 0.3}

Table 4: Hyperparameter Search.

Model	JF17K						Wikepeople						WD50K					
	main relation			all relations			main relation			all relations			main relation			all relations		
	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10
m-TransH	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
RAE	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
NaLp	0.639	0.547	0.822	0.825	0.762	0.927	0.482	0.320	0.482	0.735	0.595	0.938	-	-	-	-	-	-
NeuInfer	0.936	0.901	0.989	-	-	-	0.950	0.915	<b>0.997</b>	-	-	-	-	-	-	-	-	-
HINGE	-	-	-	0.861	0.832	0.910	-	-	-	0.765	0.686	0.900	-	-	-	-	-	-
StarE	-	-	-	0.901	0.884	0.963	-	-	-	0.378	0.265	0.542	-	-	-	-	-	-
Hyper2	0.950	0.933	0.976	-	-	-	0.947	0.914	0.987	-	-	-	-	-	-	-	-	-
HyTransformer	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRAN	0.992	0.988	<b>0.988</b>	<b>0.996</b>	0.993	<b>0.999</b>	<b>0.957</b>	<b>0.942</b>	0.976	<b>0.960</b>	<b>0.946</b>	0.977	-	-	-	-	-	-
MSeaHKG	-	-	-	0.933	0.894	0.972	-	-	-	0.831	0.787	0.972	-	-	-	-	-	-
HAHE	<b>0.993</b>	<b>0.989</b>	<b>0.998</b>	<b>0.996</b>	<b>0.994</b>	<b>0.999</b>	<b>0.957</b>	0.941	0.978	0.958	0.942	<b>0.978</b>	<b>0.916</b>	<b>0.885</b>	<b>0.964</b>	<b>0.927</b>	<b>0.900</b>	<b>0.969</b>
HAHE w/o global	0.992	0.988	0.997	<b>0.996</b>	<b>0.994</b>	0.998	0.946	0.930	0.967	0.947	0.931	0.968	0.915	0.884	0.961	<b>0.927</b>	<b>0.900</b>	0.966
HAHE w/o node-bias	0.989	0.984	0.995	0.994	0.991	0.997	0.926	0.906	0.952	0.926	0.906	0.953	0.915	0.883	0.963	0.922	0.899	0.969
HAHE w/o edge-bias	0.992	0.987	<b>0.998</b>	0.995	0.993	<b>0.999</b>	0.948	0.932	0.970	0.949	0.933	0.970	0.915	0.882	0.963	0.926	0.898	0.968

Table 5: Comparison of HAHE with other models, composed of relation prediction accuracy on JF17K, WikiPeople and WD50K. Results of the models are mainly taken from the original paper. Best results in each tasks are in **bold**.

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*Limitations after section 6.*
- A2. Did you discuss any potential risks of your work?  
*Ethics Statement.*
- A3. Do the abstract and introduction summarize the paper’s main claims?  
*Abstract and Section 1: Introduction.*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*Left blank.*

- B1. Did you cite the creators of artifacts you used?  
*No response.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*No response.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*No response.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*No response.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*No response.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*No response.*

### C Did you run computational experiments?

*Section 5: Experiments.*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*Appendix B: Model training details*

---

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*Section 5.1: Experimental Setup*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*Section 5.2: Main Results, Appendix B: Model training details*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*Appendix A: Hyperparameter Settings*

**D  Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*No response.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*No response.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*No response.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*No response.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*No response.*