

A Hierarchical N-Gram Framework for Zero-Shot Link Prediction

Mingchen Li^{1*}, Junfan Chen², Samuel Mensah³

Nikolaos Aletras³, Xiulong Yang¹, Yang Ye¹

¹Georgia State University ²Beihang University ³University of Sheffield

¹{mli33, xyang22, yye10}@student.gsu.edu, ²chenjf@act.buaa.edu.cn

³{s.mensah,n.aletras}@sheffield.ac.uk

Abstract

Knowledge graphs typically contain a large number of entities but often cover only a fraction of all relations between them (i.e., incompleteness). Zero-shot link prediction (ZSLP) is a popular way to tackle the problem by automatically identifying unobserved relations between entities. Most recent approaches use textual features of relations (e.g., surface names or textual descriptions) as auxiliary information to improve the encoded representation. These methods lack robustness as they are bound to support only tokens from a fixed vocabulary and are unable to model out-of-vocabulary (OOV) words. Subword units such as character n-grams have the capability of generating more expressive representations for OOV words. Hence, in this paper, we propose a **Hierarchical N-gram framework for Zero-Shot Link Prediction (HNZSLP)** that leverages character n-gram information for ZSLP. Our approach works by first constructing a hierarchical n-gram graph from the surface name of relations. Subsequently, a new Transformer-based network models the hierarchical n-gram graph to learn a relation embedding for ZSLP. Experimental results show that our proposed HNZSLP method achieves state-of-the-art performance on two standard ZSLP datasets.¹

1 Introduction

Link prediction models aim to predict relations between entities in knowledge graphs (KGs). Majority of these methods learn low-dimensional representations of entities and relations (i.e., knowledge graph embeddings (KGE)), which are then used to infer links between entities. Traditional approaches (Bordes et al., 2013) assume that all relation types are known in the KG. This assumption is however unrealistic since KGs are inherently incomplete. To tackle this issue, the zero-shot link

prediction (ZSLP) task has been introduced for identifying unseen relations by leveraging auxiliary information that bridges the gap between seen and unseen relations (Qin et al., 2020).

Little previous work exists on ZSLP as the task is relatively new (Qin et al., 2020; Geng et al., 2021; Wang et al., 2021). Most efforts focus on using textual features (Qin et al., 2020; Wang et al., 2021) or ontologies (Geng et al., 2021) as auxiliary information for the task. Particularly, Wang et al. (2021) use surface names of entities and relations while Qin et al. (2020) use the textual descriptions of relations. However, these approaches have two main limitations. First, common knowledge graphs such as WordNet (Miller, 1995) and FreeBase (Bollacker et al., 2008) often do not include textual descriptions of the relations. As such, these need to be obtained from other external sources (e.g., Wikipedia²) and are likely to be noisy, leading to poor performance. Second, manually obtaining such relation descriptions is a labor-intensive and time-consuming process due to the large size of KGs.

Alternatively, Wang et al. (2021) proposed learning word representations from the surface name of relations using a pre-trained language model such as RoBERTa (Liu et al., 2019). As surface names are readily available in the KG, this approach is more robust. However, it faces two fundamental weaknesses. First, context is an essential requirement for any text representation method. Surface names on the other hand are represented by short texts, e.g., a relation “teammate” will have a single word representation observed in training and will therefore have little to no association with an unobserved relation for zero-shot. Second, neural text encoders lack the ability to capture representations for out-of-vocabulary words. This same problem also applies to “word”-delimited models (Qin et al., 2020) that aim to learn from textual descriptions of

*Corresponding author

¹The code is available here: <https://github.com/ToneLi/HNZSLP>

²<https://www.wikipedia.org/>

relations. In such cases, the relation representation ability of current methods may be limited significantly, which inadvertently hurts the zero-shot link prediction performance.

In this paper, we follow a different direction. Instead of simply learning representations from entire words of a relation’s surface name, we hypothesize that leveraging character n-grams³ (or n-grams for brevity) information from the relation name will help in generating better representations of unseen relations in zero-shot settings. Models built on subword units (e.g., character n-grams) have the intrinsic ability of generating representations for rare or out-of-vocabulary words (Santos et al., 2021). Inspired by this, we propose a novel Hierarchical N-gram framework for Zero-Shot Link Prediction (HNZSLP) that learns auxiliary information from character n-grams of the surface name of a relation. HNZSLP consists of three main components: (1) a new hierarchical n-gram graph (or n-gram graph for brevity) for representing the relationships between all the character n-grams of a relation; (2) GramTransformer, a new transformer-based (Vaswani et al., 2017) model for encoding the relation n-gram graph; and (3) a KG Embedding model which adapts prevalent KGE models (e.g., TransE (Bordes et al., 2013), DistMult (Yang et al., 2014), TuckER (Balažević et al., 2019)) to compute a link prediction score between entities in the zero-shot setting. We perform extensive experiments on two standard datasets for zero-shot link prediction demonstrating the superiority of our method over prior state-of-the-art methods.

Our contributions are the following:

- We propose HNZSLP, a new framework that uses the character n-gram information from the relation surface name for ZSLP;
- We show that our approach outperforms previous state-of-the-art when evaluated on character and byte-level encoders;
- We conduct a thorough analysis of our method, including an ablation study, demonstrating the robustness of HNZSLP.

2 Related Work

2.1 Link Prediction

So far, a variety of works have been proposed for link prediction, and the difference in their architec-

³A character n-gram is defined as a contiguous sequence of n characters.

ture ranges from the scoring function to how the optimization problem is modeled to learn entities and relation embeddings. As current work is vast and fast growing, we restrict ourselves to reviewing only those closely related to our work. Some of the well-known methods include the translation-based model TransE (Bordes et al., 2013), which requires that the tail entity embedding is close to the sum of the head and relation embeddings; the bilinear model DistMult (Yang et al., 2014) that interprets the relation as a bilinear map and uses multiplicative interactions to learn entity and relation embeddings; the non-bilinear model TuckER (Balažević et al., 2019) utilizes the tucker decomposition (Malik and Becker, 2018) to build the connection between different knowledge graph triples. Although performance has been achieved incrementally, these approaches in their original form are unable to learn embeddings for unseen relations. This is due to the fact that they learn entities and relation embeddings using the topological structure of the KG. We refer the reader to the work by (Rossi et al., 2021) for further background on such methods.

2.2 Zero-shot Link Prediction

The zero-shot link prediction (Qin et al., 2020) is a new task that aims to predict unseen relations between entities by using auxiliary information to bridge the gap between seen and unseen relations. Qin et al. (2020) uses textual information of the relation as auxiliary information and applies a Zero-Shot Generative Adversarial Network (ZS-GAN) to learn the unseen relation embedding for the task. An Ontology-enhanced Zero-Shot Learning (OntoZSL) (Geng et al., 2021) obtains structural information of relations from the ontology and combines it with the textual descriptions of the relations for zero-shot learning. Despite the success, these textual descriptions are typically not present in knowledge graphs and therefore these methods rely on external sources to collect such data. This makes it labor-intensive and time-consuming to obtain the most representative descriptions of entities and relations.

2.3 Character-level Information for Zero-shot Learning

An emerging trend is to use the character-level information of the raw text in zero-shot learning. Byt5 (Xue et al., 2021) is one of such models that uses a language model T5 (Raffel et al.,

2019) to process byte or character sequences. Charformer (Tay et al., 2021) improves upon Byt5 by introducing a gradient-based subword tokenization module to learn the character information. Our proposed model is somewhat aligned with these models in the sense that we consider the character information in the text. However, we consider the process of how words are formed. That is, being considered as a sequential combination of characters/n-grams or a hierarchical structure, whereby different n-grams aggregate up to a complete word. We model this structure, referred to as a n-gram graph structure, using a self-attention based Transformer (Vaswani et al., 2017) due to its success in graph learning (Ahmad et al., 2021; Cai and Lam, 2020; Lyu et al., 2021; Yao et al., 2020).

3 Problem Statement

A Knowledge Graph (KG) is defined as a graph $\mathcal{G} = (\mathcal{R}, \mathcal{E})$, where \mathcal{E} denotes a set of entities and \mathcal{R} denotes the set of relations among these entities. In a KG, the entities and relations are usually organized as *facts* and each fact is defined as a triplet (h, r, t) where $h, t \in \mathcal{E}$ and $r \in \mathcal{R}$ denote the head entity, tail entity and the relation between the two entities, respectively.

In the zero-shot link prediction problem, we assume that there are two disjoint relation sets in the KG, a seen relation set \mathcal{R}_s and an unseen relation set \mathcal{R}_u , where $\mathcal{R}_s \cap \mathcal{R}_u = \emptyset$. We are given a training set $\mathcal{D}_s = \{(h, r_s, t) | h, t \in \mathcal{E}, r_s \in \mathcal{R}_s\}$ in which the facts are involved with observed relations $r_s \in \mathcal{R}_s$. Meanwhile, we define the test set as $\mathcal{D}_u = \{(h, r_u, t, C_{(h, r_u)}) | h, t \in \mathcal{E}, C_{(h, r_u)} \subseteq \mathcal{E}, r_u \in \mathcal{R}_u\}$, where t is the ground-truth tail entity and $C_{(h, r_u)}$ denotes a candidate set corresponding to a query (h, r_u) . Given a query (h, r_u) , the objective of zero-shot link prediction is to find the ground-truth tail entity t from the candidate set $C_{(h, r_u)}$.

4 HNZSLP

Figure 1 gives an overview of our HNZSLP framework, which consists of three major parts: (1) Hierarchical N-gram Graph Constructor constructs a hierarchical n-gram graph from each relation surface name, where the n-gram graph can be further decomposed into an *adjoin* graph and *compositional* graph to simplify the learning of the relation representation; (2) GramTransformer constructs the relation representation by modeling over the ad-

join and compositional graphs; (3) KG Embedding Learning Module combines the embeddings of the head entity and relation to predict the tail entity and update the embeddings.

4.1 Hierarchical N-gram Graph Constructor

Node Construction For each word token in the relation surface name, we first collect all possible n-grams, where n is valued from 1 up to the maximum gram size M of a word. For example, $M = 3$ for the relation *has* in Figure 1. All n-grams are treated as nodes in the hierarchical n-gram graph. Suppose the relation surface name contains multiple words, the n-grams of each word are composed into a unified n-gram graph. For each hierarchical n-gram graph, we denote all its nodes as a sequence $X = \{x_1, x_2, \dots, x_b\}$, where $b = \frac{M(M+1)}{2}$. Let $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_b\}$ be the corresponding node embeddings for the graph.

Edge Construction We define two types of edges among n-grams in the hierarchical n-gram graph: *adjoin* edge and *compositional* edge. The adjoin edge implies that two n-grams at the same hierarchical level are neighbors, e.g., the edge between nodes “h” and “a” in Figure 1. The compositional edge implies that the n-gram node at a higher-level (i.e., a superior node) is a composition of the adjacent n-gram nodes at the immediate lower-level, e.g., the edge between node “h” and “ha”, and “a” and “ha” in Figure 1. According to the two edge types, we can decompose the n-gram graph into the adjoin graph and compositional graph, as shown in Figure 1.

For the adjoin and compositional edges, we first define their textual definitions based on Wikidata,⁴ and calculate their embeddings using SentenceBERT (Reimers and Gurevych, 2019). For later use, we define the embeddings of the adjoin and compositional edges as $\mathbf{r}_a \in \mathbb{R}^d$ and $\mathbf{r}_c \in \mathbb{R}^d$, respectively.

Note that some surface names of the relations may be a long sequence of words, which may result in a large set of nodes in the n-gram graph, making it hard to process. To boost the graph construction process, we reduce the number of nodes in the hierarchical n-gram graph using two strategies (see details in Appendix Section 9.1).

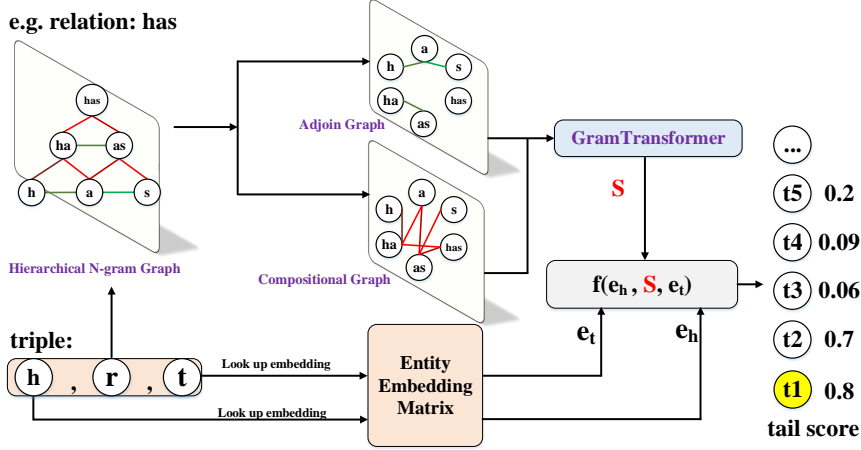


Figure 1: Overview of HNZSLP, we take a relation *has* as an example. In the adjoin Graph, the green line denotes the adjoin relationship. In the compositional Graph, the red line indicates the compositional relationship. In these two graphs, different color depths represent different attention weights. The node in the adjoin graph and the compositional graph is called neighbor node and superior node separately.

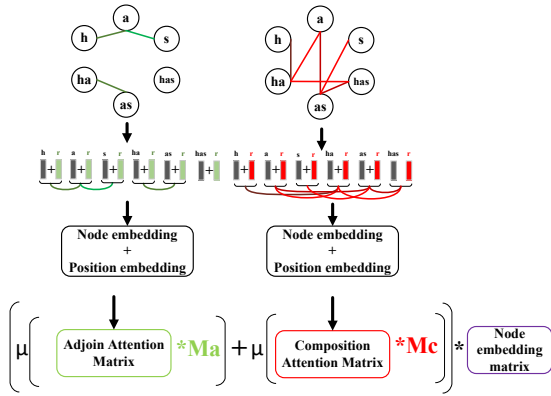


Figure 2: Overview of the relation enhanced mask attention used in the GramTransformer. M_a and M_c are the mask matrix that models the relationship among nodes in the adjoin and compositional graphs, respectively. μ is the softmax function. The green rectangle refers to the embedding $r_a \in \mathbb{R}^d$. The red rectangle refers to the embedding $r_c \in \mathbb{R}^d$.

4.2 GramTransformer

We propose a GramTransformer to efficiently extract hierarchical n-gram features. The difference between the standard Transformer and our GramTransformer is the attention calculation, as shown in Figure 2. The most important characteristic of the GramTransformer is that it can encode the neighbor node and superior node information while taking the edge information in the adjoin and compositional graphs into account. In this way, a node can directly learn the information from different

neighborhoods. These operations are achieved by our proposed relation enhanced mask attention mechanism. Specifically, we decompose the original n-gram graph into the adjoin graph and compositional graph based on adjoin edge and compositional edge. We then initialize the node embeddings in each subgraph as the sum of the node embedding, edge embedding r_a (or r_c), and position⁵ embedding (Vaswani et al., 2017). Next, the initialized subgraph elements with mask matrix is used to learn the attention matrix, highlighting the relevant features of the n-gram graph. Multiple layers of relation enhanced mask attention networks are stacked to calculate the final node representation. At each layer, a node vector is updated based on neighbor nodes and associated edge types. The nodes' vectors at the last layer is considered as the final n-gram graph representation.

4.2.1 N-gram Graph Encoder

Our graph encoder aims to transform an input n-gram graph into a set of node embeddings. To calculate the node information in our graph, the central problem is how to calculate the node vectors based on the different subgraphs (adjoin graph or compositional graph). To this end, we propose a relation enhanced mask attention mechanism, which is an extension of the self-attention mechanism to relate the different nodes across the subgraphs.

To maintain the graph structure information, our

⁴https://www.wikidata.org/wiki/Wikidata:Main_Page

⁵for the position about nodes, please refer Appendix Section 9.1

idea is to introduce the explicit edge information and incorporate it into the subgraph attention score computation. We introduce the mask matrix to model the relationship among nodes in each subgraph. Recall, for the standard self-attention, an L-layer Transformer takes \mathbf{X} as input and produces the latent representation $\mathbf{H}^l = (\mathbf{h}_1^l, \mathbf{h}_2^l, \dots, \mathbf{h}_b^l)$ of relations. To enhance the semantic representation of the input, multi-head self-attention is used in each Transformer Layer. Specifically, the output of $(l-1)^{\text{th}}$ Transformer layer is projected to a query matrix \mathbf{Q}_l and a set of key-value $(\mathbf{K}_l, \mathbf{V}_l)$ pairs,

$$\mathbf{Q}_l = \mathbf{H}^{l-1} \mathbf{W}_l^q, \mathbf{K}_l = \mathbf{H}^{l-1} \mathbf{W}_l^k, \mathbf{V}_l = \mathbf{X}^{l-1} \mathbf{W}_l^v$$

where $\mathbf{W}_l^q, \mathbf{W}_l^k, \mathbf{W}_l^v \in \mathbb{R}^{d_{\text{model}} \times d_k}$ denote the learnable weight matrices. d_{model} is the model dimension, d_k is the head dimension. The output of a self-attention head $\bar{\mathbf{H}}^l$ is calculated by:

$$\bar{\mathbf{H}}^l = \mu\left(\frac{\mathbf{Q}_l \mathbf{K}_l^T}{\sqrt{d_k}}\right) \mathbf{V}_l \quad (1)$$

where μ denotes the *softmax* function. The self-attention learns the implicit relationships between nodes in the hierarchical n-gram graph.

4.2.2 Relation Enhanced Mask Attention

As our n-gram graph consists of adjoin graph and compositional graph, we respectively compute the attention head as follows,

$$\mathbf{Q}_l^a = \mathbf{K}_l^a = \mathbf{V}_l^a = \mathbf{Q}_l + \mathbf{r}_a \quad (2)$$

$$\mathbf{Q}_l^c = \mathbf{K}_l^c = \mathbf{V}_l^c = \mathbf{Q}_l + \mathbf{r}_c \quad (3)$$

$$\tilde{\mathbf{H}}^l = \left[\mu\left(\frac{\mathbf{Q}_l^a (\mathbf{K}_l^a)^T}{\sqrt{d_k}}\right) + \mu\left(\frac{\mathbf{Q}_l^c (\mathbf{K}_l^c)^T}{\sqrt{d_k}}\right) \right] \mathbf{V}_l \quad (4)$$

where we split the node embedding Q_l into neighbor node embedding Q_l^a and superior node embedding Q_l^c . Then we compute the attention score based on the embeddings of the nodes and edges in the subgraph. In Eq.4, the first term represents the node weight calculated from its adjoin neighbors. The second term represents the node weight calculated using the compositional edge information. In comparison, our model can calculate the node embedding respectively based on the different edge types, it can compute the node embedding more precisely than the standard mask self-attention which does not consider the edge type. The comparative experiment is in our Section 6.2.

To denote the node connection in subgraphs, the central idea is to incorporate the mask matrix into the attention matrix of self-attention, which can impose the structure of the n-gram graph and reassign the attention weight for each relation.

We denote the mask matrix $\mathbf{M} \in \mathbb{R}^{m \times m}$, where $\mathbf{M}_{ij} \in [0, 1]$ denotes the connection between node at position i and j in the input n-gram node list. 1 denotes there is a connection between two nodes. So our proposed attention strategy can be redefined as,

$$\mathbf{H}^l = \left[\mu\left(\frac{\mathbf{Q}_l^a (\mathbf{K}_l^a)^T}{\sqrt{d_k}}\right) \mathbf{M}_a + \mu\left(\frac{\mathbf{Q}_l^c (\mathbf{K}_l^c)^T}{\sqrt{d_k}}\right) \mathbf{M}_c \right] \mathbf{V}_l \quad (5)$$

where \mathbf{M}_a indicates the relationship among nodes in the adjoin graph, \mathbf{M}_c indicates the relationship among nodes in the compositional graph. Hence, for an input n-gram graph, our graph encoder module produces the attention head \mathbf{H}^l which is fed to subsequent layers of the GramTransformer to output the node representations $\mathbf{h}_1^l, \mathbf{h}_2^l, \dots, \mathbf{h}_b^l$. We then apply a mean pooling over these representations to obtain the relation embedding \mathbf{S} of the surface name.

4.3 Embedding Learning Module

We randomly initialize entity embedding matrix $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times d^e}$, where each row vector is the embedding of an entity and d^e is the dimension of the entity embedding. In a triplet (h, r, t) , we define \mathbf{e}_h and \mathbf{e}_t as the embedding of the head entity and tail entity retrieved from the embedding matrix \mathbf{E} . Given the entity embeddings \mathbf{e}_h and \mathbf{e}_t , and the relation embedding \mathbf{S} computed by our proposed GramTransformer, we then define a scoring function $f(\cdot)$ that assigns a score η to each triple (h, r, t) ,

$$\eta = f(\mathbf{e}_h, \mathbf{S}, \mathbf{e}_t) \quad (6)$$

where the scoring function f can be replaced by any knowledge graph embedding model, e.g., TransE, DistMult. The model is trained with cross-entropy loss.

During the inference process, the trained HNZSLP scores each candidate tail entity $t' \in C_{(h, r_u)}$ given the query (h, r_u) . Let \mathbf{S}_u be the embedding of r_u computed by GramTransformer, the entity with the highest score in the candidate entity set is selected as the predicted tail entity:

$$t^* = \arg \max_{t' \in C_{(h, r_u)}} f(\mathbf{e}_h, \mathbf{S}_u, \mathbf{e}_{t'}) \quad (7)$$

where t^* refers to the predicted tail entity.

5 Experiments

We validate HNZSLP by comparing HNZSLP’s performance with several recent works, including ZSGAN (Qin et al., 2020) and OntoZSL (Geng et al., 2021). ZSGAN exploits the generated description embeddings of unseen relations to predict the tail entity while OntoZSL introduces the ontology strategy in the task. Following Geng et al. (2021), we further compare with the baselines ZSL-TransE and ZSL-DistMult that use Word2vec (Vinyals and Le, 2015), and respectively employ TransE (Bordes et al., 2013) and DisMult (Yang et al., 2014) as KGE models for zero-shot link prediction. For a fair comparison, we exclude the results of (Wang et al., 2021) as this approach does not show the results in our used datasets. We use four commonly used metrics, mean reciprocal ranking (MRR), hits@10, hits@5, hits@1, and evaluate on two benchmark datasets, including NELL-ZS and Wikidata-ZS (Wiki-ZS) proposed by Qin et al. (2020). A summary of dataset statistics is given in Table 1.

Dataset	# Entities	# Triples	# Train/ # Dev/ # Test
NELL-ZS	65,567	188,392	139/ 10/ 32
Wiki-ZS	605,812	724,967	469/ 20/ 48

Table 1: Datasets Statistics, column 4 refers to the number of relations in the different set.

5.1 Implementation details

On NELL-ZS (or Wiki-ZS), each word in the surface name of relation is set to 13-gram (or 15-gram) and the number of nodes in the n-gram graph is set to 90 (or 70). Each node in the n-gram graph is randomly initialized with a 200-dim embedding. The GramTransformer contains one multi-head attention block with three attention heads and a 200-dim feed-forward layer. The dropout rate in the multi-head attention and feed-forward layer is set to 0.5. Entities are also randomly initialized with 200-dim embeddings. During training, we use Adam (Kingma and Ba, 2014) as the optimizer and a Cross-entropy is used as the loss function with a learning rate of 0.0005. We use label smoothing to prevent the model from becoming over-confident. All embeddings are fine-tuned during training.

5.2 Main Results

Evaluation results are shown in Table 2. Results are the average over 5 runs. We find that our approach outperforms all previous methods on the different KGE models. With TransE and DistMult, our model achieves hits@1 scores of 0.222 and 0.216 respectively on NELL-ZS, outperforming the previous best-performing network OntoZSL by a margin of 0.05 and 0.028. With DistMult, we also find that our model achieves the best performance on hits@1, but slightly underperforms the best-performing network OntoZSL on NELL-ZS. It is worth noting that OntoZSL utilizes external ontology resources, thus, they present an additional advantage over ours that do not consider external knowledge. In real applications, ontology is not always available, which limits the scalability of their method. On Wiki-ZS, our model also sets a new hits@1 score. Particularly, for TransE, we improve over the state-of-the-art OntoZSL by about 0.081 points. Similar performance is achieved on other metrics for the different KGE models. These results indicate that sufficient information is contained in the relation surface name to achieve zero-shot link prediction, and our proposed method is effective, it can utilize the n-gram graph to transfer the knowledge between seen relation and unseen relation is effective.

6 More Analysis

In order to further explore the effectiveness of our framework, we perform a series of analyses based on different characteristics of our model. First, we explore the effectiveness of our proposed GramTransformer with two latest works that learn text information from the character level or byte level. The contribution of our model components can also be learned from ablated models. So we propose two model variants to help us validate the advantages of the n-gram graph information and GramTransformer. Next, we explore the performance of our model with a different number of nodes in the n-gram graph. In the last, we did a comparison with the method which applies the language model in this task.

For more experiments about the HNZSLP performance in the OOV problem, the impact of different node selection strategies, please refer to Appendix Sections 9.2, 9.3.

KGE model	Method	NELL-ZS				Wiki-ZS			
		MRR	hits@10	hits@5	hits@1	MRR	hits@10	hits@5	hits@1
TransE	ZSL-TransE(Qin et al., 2020)	0.097	0.203	0.147	0.043	0.053	0.119	0.081	0.018
	OntoZSL(Geng et al., 2021)	0.250	0.399	0.327	0.172	0.184	0.265	0.215	0.138
	ZSGAN(Qin et al., 2020)	0.234	0.373	0.304	0.160	0.177	0.258	0.207	0.131
	HNZSLP	0.289	0.413	0.359	0.222	0.252	0.307	0.281	0.219
DistMult	ZSL-DistMult(Qin et al., 2020)	0.235	0.326	0.284	0.185	0.189	0.236	0.210	0.161
	OntoZSL(Geng et al., 2021)	0.256	0.385	0.318	0.188	0.211	0.289	0.238	0.167
	ZSGAN(Qin et al., 2020)	0.249	0.376	0.306	0.183	0.207	0.284	0.235	0.164
	HNZSLP	0.276	0.383	0.333	0.216	0.232	0.279	0.254	0.204

Table 2: Zero-shot link prediction results in NELL-ZS and Wiki-ZS. The baseline results were obtained from (Geng et al., 2021). The KGE models in the first column correspond to $f(\cdot)$ in equation 6.

KGE model	Method	NELL-ZS				Wiki-ZS			
		MRR	hits@10	hits@5	hits@1	MRR	hits@10	hits@5	hits@1
TransE	ZSL-ByT5 ∇	0.103	0.173	0.129	0.064	0.110	0.193	0.134	0.064
	ZSL-CharFormer ∇	0.269	0.383	0.337	0.202	0.170	0.236	0.201	0.132
	OntoZSL(Geng et al., 2021)	0.250	0.399	0.327	0.172	0.184	0.265	0.215	0.138
	HNZSLP-WNG \circ	0.278	0.409	0.350	0.203	0.231	0.273	0.262	0.194
	HNZSLP-WG \circ	0.244	0.382	0.322	0.173	0.232	0.284	0.258	0.199
	HNZSLP	0.289	0.413	0.359	0.222	0.252	0.307	0.281	0.219
DistMult	ZSL-ByT5 ∇	0.233	0.382	0.317	0.150	0.196	0.255	0.232	0.160
	ZSL-CharFormer ∇	0.251	0.380	0.320	0.183	0.205	0.270	0.251	0.153
	OntoZSL(Geng et al., 2021)	0.256	0.385	0.318	0.188	0.211	0.289	0.238	0.167
	HNZSLP-WNG \circ	0.274	0.381	0.330	0.195	0.178	0.208	0.192	0.157
	HNZSLP-WG \circ	0.259	0.372	0.321	0.195	0.221	0.264	0.242	0.197
	HNZSLP	0.276	0.383	0.333	0.216	0.232	0.279	0.254	0.204

Table 3: The performance of HNZSLP with different character/byte learning models (denote as ∇) and its variants (denote as \circ). WNG refers to without n-gram graph, WG refers to without GramTransformer.

6.1 Comparison with Character/Byte Models

To evaluate the effectiveness of our proposed GramTransformer, we explore two state-of-the-art methods for character/byte level learning, including CharFormer (Tay et al., 2021) and ByT5 (Xue et al., 2021) to calculate the relation surface name embedding S (as shown in (6)) at the character or byte level, respectively. Accordingly, we propose **ZSL-CharFormer** and **ZSL-ByT5** for ZSLP using the same experimental setup as our method for a fair comparison with our method. Table 3 shows the performance comparison. On the dataset NELL-ZS, our model can achieve the hits@1 score of 0.222, outperforming the best model ZSL-CharFormer by a large margin of 0.02 hits@1 score using the TransE KGE model. On the dataset Wiki-ZS, our model also outperforms the best model ZSL-CharFormer by an impressive margin of 0.087 on the same TransE KGE model.

The advantages of our model are also verified by MRR, hits@10, and hits@5. Otherwise, the model performance under different KGE models also is compared.

In CharFormer, Tay et al. (2021) lists a fixed number of subword blocks and uses an attention-based method to choose the best subword at each character position. By utilizing the stride window to get the subwords, this work ignores the semantic information of other subwords outside the window. Meanwhile, in our work, we propose to use n-gram to help calculate the relation information. Our approach is much more conducive to preserve semantic information as it considers the effective modeling of rare words or OOV words. Again, our results demonstrate that the hierarchical n-gram graph information is important to express the semantic information of the relation.

6.2 Ablation Experiments

The contribution of our model components can also be learned from ablated models. We introduce two ablated models of HNZSLP, (1) **HNZSLP-WNG** uses a traditional Transformer to learn the information of relation surface name from the 1-gram level; (2) **HNZSLP-WG** uses the standard self-attention to learn the information of n-gram graph, instead of our proposed GramTransformer that incorporates different sub-graph information. We find that the performance of HNZSLP degrades as we remove important model components. Specifically, both HNZSLP-WNG and HNZSLP-WG perform poorly when compared to HNZSLP, indicating the importance of modeling the information of the n-gram graph.

6.3 Impact of Node Number

Nell-ZS No.	n	l	T (min)	L
1	13	30	18	0.1
2	13	70	37	0.5
3	13	90	45	0.63
4	13	110	66	0.74
Wiki-ZS No.	n	l	T (min)	L
1	15	30	32	0.15
2	15	70	46	0.63
3	15	90	48	0.78
4	15	110	52	0.89

Table 4: Different node number setting, L rate denotes that there are $\%L$ relations that can cover the whole nodes in the n-gram graph. T refers to the training time.

In our proposed model, there are two kinds of parameters controlling the size of the n-gram graph. one is the gram number n , and the other is the node number l in the n-gram graph. In HNZSLP, we treat these two parameters with the same importance. For the n parameter, we use the maximum word length about the seen relation set as the gram number. In NELL-ZS, $n = 13$, in Wiki-ZS, $n = 15$. If the word length is smaller than the maximum word length, the gram number is its word length.

The performance of our models differs in terms of accuracy and training time under the nodes with different numbers. To investigate the influence of different node sizes, we conduct experiments using HNZSLP with different parameter settings which are shown in Table 4. We set the batch size to 32 both in NELL-ZS and Wiki-ZS. The training epoch in NELL-ZS is 80, and the training epoch in Wiki-ZS is 70. Other configurations are the same for HNZSLP with different node number settings. We

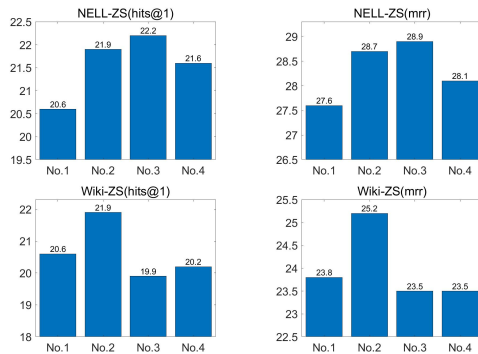


Figure 3: Mrr and hits@1 results on the dataset NELL-ZS and Wiki-ZS with different node numbers

list the training time of each running on GPU. The GPU computations were run on a single Nvidia TITAN RTX. In this section, we experiment with the models under the node number up to 110 for a relation, the KGE model is TransE.

From Figure 3, we can find that, under the fixed setting of gram number n , HNZSLP with more nodes achieves higher accuracy, for example, in the dataset NELL-ZS, the hits@1 results with node number 90 are higher than the hits@1 results with node number 30 and 70. This situation indicates that extending the node number is effective. However, the performance of our model about No.4 (NELL-ZS, $n=13$, $l=110$) cannot reach the one of No.3 (NELL-ZS, $n=13$, $l=90$) and No.2 (NELL-ZS, $n=13$, $l=70$), despite they have the same gram length 13. Our experimental experience suggests that it is not necessary to utilize the whole n-gram graph in HNZSLP, let alone that using the whole n-gram graph may significantly increase the computational cost. In practice, we should choose a suitable number of nodes to build the n-gram graph. However, it is difficult to set appropriate node numbers because there are no systematic methods. We thus set the node numbers by experimental experience. In future work, we will design a more efficient approach to solve this problem and balance the trade-off between the hits@1, hits@5, hits@10, mrr results, and efficiency.

6.4 Comparison with Different Language Models

To quantitatively evaluate the effect of HNZSLP, we compare the performance of HNZSLP against two models which are based on the latest work about the language model in the zero-shot link prediction task, and all based on the $BERT_{base}$ model

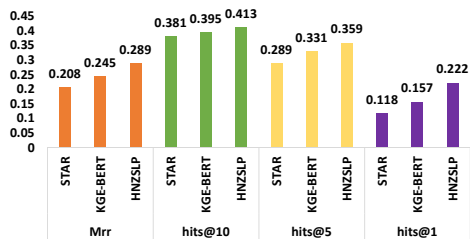


Figure 4: The model performance with different zero shot link prediction methods based on the language model.

⁶. These experiments are conducted with the test set of NELL-ZS. They are described below:

- **KGE-BERT.** We use BERT (Devlin et al., 2018) to calculate the relation embedding S in eq.6. The KGE score function $f(\cdot)$ (TransE) is utilized to predict the tail entity. We refer to this model as KGE-BERT.
- **STAR.** STAR (Wang et al., 2021) is the first work to explore the ability of language models in zero-shot link prediction by using the textual information of entity and relation. In this work, the authors use structured knowledge information and textual information of entity and relation to infer the tail entity. Moreover, they develop a self-adaptive ensemble scheme to improve the model performance by incorporating the triple scores.

We perform a detailed comparative study on different zero-shot language models to examine their impact on knowledge transfer from seen relation to unseen relation under the KGE framework. Figure 4 presents the results with a comparison to KGE-BERT and STAR.

Figure 4 shows that our model can achieve a hits@1 score of 0.214, outperforming the model KGE-BERT which uses the language model to learn the relation information by a large margin of 0.025 hits@1 scores. For the latest work which uses the language model to infer the tail entity, our model can outperform STAR by a margin of 0.096 hits@1 scores. The advantages of HNZZSLP are also verified by metrics MRR, hits@10, and hits@5.

In KGE-BERT, we use the language model to learn the relational textual information, but this way ignores the importance of n-gram graph information. In the model STAR, the authors use the

⁶<https://huggingface.co/bert-base-uncased>

language model to enhance the inference ability of link prediction, they combine the textual information of the head entity and relation by a special token $[SEP]$ and then build the triple score with learned tail entity information by the language model. Unfortunately, this model cannot be used to solve the out-of-vocabulary problem for the current word, though the language model can use its previous knowledge. More importantly, the n-gram graph information is also an important source to improve the performance of tail entity inference.

7 Conclusion

In this paper, we propose a novel ZSL framework HNZZSLP for link prediction. Specifically, we proposed a GramTransformer to learn the n-gram graph information of the relation surface name and utilize the KGE model to infer the tail entity. Experimental results show that our framework achieves consistent improvements over various baselines in two ZSLP datasets. As the GramTransformer can be considered as a text representation method, in the future, we intend to explore its effectiveness on other NLP tasks including text classification.

8 Limitations

When the surface name of relations is too long, it means the scale of our build n-gram graph is large, this way will influence the efficiency of graph computation. In our work, we just proposed two strategies to select the fixed number of nodes, this way ignores the semantic information about the nodes which are not selected. So in the future, we will design a novel strategy to dynamically select the node, and consider the computational problem at the same time.

Acknowledgements

We would like to thank the anonymous reviewers for their comments and suggestions, which helped improve the quality of this paper. We also thank Prof. Richong Zhang from Beihang University for the inspiration of this topic. SM and NA are supported by a Leverhulme Trust Research Project Grant (No. RPG-2020-148).

References

Wasi Uddin Ahmad, Nanyun Peng, and Kai-Wei Chang. 2021. Gate: Graph attention transformer encoder for cross-lingual relation and event extraction. In

- The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21).*
- Ivana Balažević, Carl Allen, and Timothy M Hospedales. 2019. Tucker: Tensor factorization for knowledge graph completion. *arXiv preprint arXiv:1901.09590*.
- 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*, pages 1247–1250.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.
- Deng Cai and Wai Lam. 2020. Graph transformer for graph-to-sequence learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7464–7471.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Yuxia Geng, Jiaoyan Chen, Zhuo Chen, Jeff Z Pan, Zhiquan Ye, Zonggang Yuan, Yantao Jia, and Huajun Chen. 2021. Ontozsl: Ontology-enhanced zero-shot learning. In *Proceedings of the Web Conference 2021*, pages 3325–3336.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Boer Lyu, Lu Chen, Su Zhu, and Kai Yu. 2021. Let: Linguistic knowledge enhanced graph transformer for chinese short text matching. *arXiv preprint arXiv:2102.12671*.
- Osman Asif Malik and Stephen Becker. 2018. Low-rank tucker decomposition of large tensors using tensors-ketch. *Advances in neural information processing systems*, 31:10096–10106.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Pengda Qin, Xin Wang, Wenhui Chen, Chunyun Zhang, Weiran Xu, and William Yang Wang. 2020. Generative adversarial zero-shot relational learning for knowledge graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8673–8680.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Andrea Rossi, Denilson Barbosa, Donatella Firmani, Antonio Matinata, and Paolo Merialdo. 2021. Knowledge graph embedding for link prediction: A comparative analysis. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(2):1–49.
- Flávio Arthur O Santos, Thiago Dias Bispo, Hendrik Teixeira Macedo, and Cleber Zanchettin. 2021. Morphological skip-gram: Replacing fasttext characters n-gram with morphological knowledge. *Inteligencia Artificial*, 24(67):1–17.
- Yi Tay, Vinh Q Tran, Sebastian Ruder, Jai Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, and Donald Metzler. 2021. Charformer: Fast character transformers via gradient-based subword tokenization. *arXiv preprint arXiv:2106.12672*.
- 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*, pages 5998–6008.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Bo Wang, Tao Shen, Guodong Long, Tianyi Zhou, Ying Wang, and Yi Chang. 2021. Structure-augmented text representation learning for efficient knowledge graph completion. In *Proceedings of the Web Conference 2021*, pages 1737–1748.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2021. Byt5: Towards a token-free future with pre-trained byte-to-byte models. *arXiv preprint arXiv:2105.13626*.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*.
- Shaowei Yao, Tianming Wang, and Xiaojun Wan. 2020. Heterogeneous graph transformer for graph-to-sequence learning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7145–7154.

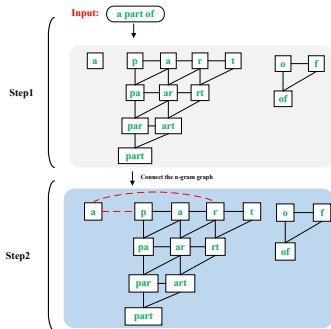


Figure 5: The building progress of n-gram graph about relation "a part of"

9 Appendix

9.1 Hierarchical N-gram Graph Building and Node Selection

Strategy	node order
Strategy1	(alp,a,r,t,pa,ar,rt,par,art,partlo,f,ofl)
Strategy2	(a,p,a,r,t,o,f pa,ar,rt,of par,art part)

Table 5: Node order of relation "a part of"

In our work, the n-gram graph from the word level is called the word n-gram graph, and the n-gram graph from the relation level is called the relational n-gram graph (the surface name of the relation contains more than one word). When the surface name of the relation contains more than one word, there are two challenges that need to be solved. The first challenge is how to connect the word n-gram graph to the relation n-grams graph. For the second challenge, when the length of a word is big or the number of words in a relation is big (such as relation "concept: agricultural product growing in state or province" in NELL), the relational n-gram graph will become very big, it is hard for the machine to progress this graph.

To solve the first issue, we connect each word n-gram graph by our proposed two relations. There are two steps in the building of a relational n-gram graph. Firstly, we split the relation by the space, and build the n-gram for each word. In the second step, we connect each word n-gram graph by relation *adjoint* and *compositional*. As shown in Figure.5, for the n-gram graph of first word "a", it appears in the n-gram graph of "part", so it also connects to "p" and "r".

For the second issue, in the relational n-gram graph, we should consider the node order in the first, this way can make sure the word does not

lose its internal order information, the node order of the n-gram graph of the word "part" should be "p,a,r,t,pa,ar,...", the order "p,pa,a,ar,r,t,..." is wrong. After that, for considering the run effectiveness and GPU memory, we select the fixed number of nodes from left to right based on the order. Based on the above discussion, we propose two strategies to help order the node. As shown in Table 5.

In Strategy1, for each word n-gram graph, we rank the node position based on the n in n -gram, the position of all 1-gram nodes all come before the 2-gram nodes. So, in the first, we list the 1-gram nodes of relation "a part of", in the second, we list the 2-gram nodes,.. In Strategy2, we list all n -grams nodes for each word and then concatenate these nodes together.

9.2 Out-of-Vocabulary

Strategy	MRR	hits@10	hits@5	hits@1
KGE-word	0.273	0.396	0.347	0.202
HNZSLP	0.289	0.413	0.359	0.222

Table 6: The model performance about the evaluation of out-of-vocabulary problem in NELL-ZS

By analyzing the dataset of zero-shot link prediction, we found some words in unseen relation are not in the seen relation set, this issue will reduce the performance of tail entity inference. To evaluate the effectiveness of our model, we propose a compared model KGE-word which directly uses the traditional transformer to learn the word information in the surface name of the relation and then uses the KGE model TransE to infer the tail entity. The results are shown in Table 6. By comparing with HNZSLP, we can see that using the n-gram graph is better than the way which uses the word information to calculate the relation information. Our n-gram graph can capture the information at different granularities, which is helpful for the knowledge transfer between the seen relation and unseen relation.

9.3 Node Selection Strategies

Strategy	MRR	hits@10	hits@5	hits@1
strategy1	0.282	0.403	0.350	0.216
strategy2	0.289	0.413	0.359	0.222

Table 7: The performance of HNZSLP in NELL-ZS with different Node order strategies

In this section, we evaluate the performance of

different node order strategies. For a fair comparison, we set 13-gram, the maximum node number is 90, the epoch of training is 80, and the KGE model is transE. In Table 7, the results show that the performance of strategy2 is better than strategy1, which shows that the lower grams are more important than the higher grams.