

Anchoring Path for Inductive Relation Prediction in Knowledge Graphs

Zhixiang Su^{1,2,3}, Di Wang^{2,3*}, Chunyan Miao^{1,2,3,4*}, Lizhen Cui^{4,5}

¹School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore

²Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY), NTU, Singapore

³WeBank-NTU Joint Research Institute on Fintech, NTU, Singapore

⁴SDU-NTU Centre for Artificial Intelligence Research (C-FAIR), Shandong University (SDU), China

⁵School of Software, SDU, China

{zhixiang002, wangdi, ascymiao}@ntu.edu.sg, clz@sdu.edu.cn

Abstract

Aiming to accurately predict missing edges representing relations between entities, which are pervasive in real-world Knowledge Graphs (KGs), relation prediction plays a critical role in enhancing the comprehensiveness and utility of KGs. Recent research focuses on path-based methods due to their inductive and explainable properties. However, these methods face a great challenge when lots of reasoning paths do not form Closed Paths (CPs) in the KG. To address this challenge, we propose Anchoring Path Sentence Transformer (APST) by introducing Anchoring Paths (APs) to alleviate the reliance of CPs. Specifically, we develop a search-based description retrieval method to enrich entity descriptions and an assessment mechanism to evaluate the rationality of APs. APST takes both APs and CPs as the inputs of a unified Sentence Transformer architecture, enabling comprehensive predictions and high-quality explanations. We evaluate APST on three public datasets and achieve state-of-the-art (SOTA) performance in 30 of 36 transductive, inductive, and few-shot experimental settings.

Introduction

As a structured representation of knowledge (e.g., real-world facts), Knowledge Graphs (KGs) serve as the basis for various downstream tasks requiring reasoning, such as question answering. KGs generally consist of a collection of triplets, denoted as $G(E_G, R_G) = \{(h_i, r_i, t_i) | i = 1, 2, 3, \dots, m\}$, where entities (head and tails, $h_i, t_i \in E_G$) and relations ($r_i \in R_G$) are represented as nodes and edges, respectively. However, real-world KGs are often incomplete, making relation prediction a necessary step to validate missing relations and subsequently enhance the comprehensiveness and reasoning ability of KGs. Given a query triplet (h, r, t) with h (or t) masked, relation prediction aims to predict the masked entity, utilizing the incomplete knowledge stored in the KG.

Relation prediction approaches can generally be classified into three different types: embedding-based, Graph Neural Network (GNN)-based, and path-based methods. Recent research has focused on path-based methods (e.g., KRST (Su et al. 2023)) due to their notable advantages as follows:

Inductiveness: Unlike traditional approaches that focus on transductive situations (where the entity set is static and all

entities are observed during training), real-world dynamic graphs are constantly evolving. This necessitates models to possess inductive capabilities to deal with unseen entities. Most embedding-based and GNN-based methods generate embeddings, which are shown as difficult to represent unseen entities. Comparatively, path-based methods are well suited to handle the inductive cases by making predictions based on a chain of relations between the head and tail entities. Such a chain of relations is usually regarded as the **supporting evidence** in this reasoning process. Once the path information is learned, similar chains of relations can be extrapolated for predictions of unseen entities.

Explainability: Explainability holds significant importance in relation prediction, encompassing a model’s ability to justify its predictive decisions by illustrating the essential reasoning process. Embedding-based and GNN-based methods have difficulties in this aspect because their encoded information is incomprehensible to humans. Comparatively, path-based methods inherently offer explainability, because the paths themselves can serve as the supporting evidence for predictions with explanations.

Recent approaches (e.g., BERTRL (Zha, Chen, and Yan 2022) and KRST (Su et al. 2023)) combined path-based methods with pre-trained language models (PLMs). PLMs take text descriptions as input, offering valuable auxiliary information for encoding paths and triplets. These text descriptions may significantly enrich the semantic meanings of entities and relations. By leveraging text descriptions, these recent approaches achieved state-of-the-art (SOTA) performance on the most widely benchmarked datasets.

Despite the great success of path-based methods with PLMs in achieving inductive explainable relation prediction, their performance is still limited. Specifically, path-based methods make predictions by leveraging Closed Paths (CPs), where paths are connected to both query head and query tail entities. Therefore, when the given KGs are highly incomplete, i.e., CPs between many head and tail entities are absent, the performance of the existing approaches may be greatly affected. On the other hand, a large portion of supporting evidence is not explicitly provided in the form of CPs. For example, as shown in Figure 1, consider the query triplet $(Morgan, Profession, Actor)$ with *Actor* masked, a piece of evidence that aligns with human commonsense

*Corresponding Authors

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

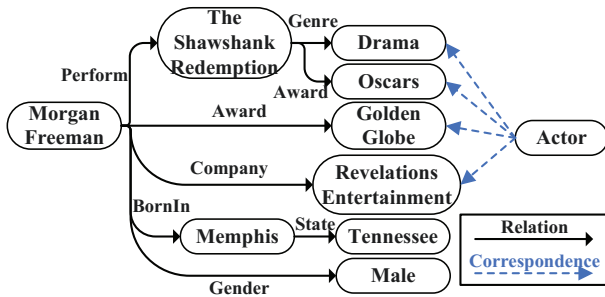


Figure 1: An incomplete KG with paths easily deducible by applying human commonsense knowledge.

knowledge is shown in Path 1 as follows:

Path 1: $Morgan \xrightarrow{\text{Perform}} \text{TheShawshankRedemption} \xrightarrow{\text{Genre}^{-1}} \text{Drama}$

The fact that *Morgan (Freeman)* performs in a *Drama* implies with a high probability that he is an *Actor*. This evidence connects to the query head (*Morgan*) but does not connect to the tail (*Actor*) as shown in Figure 1. Capable of capturing this kind of evidence will greatly enhance the reasoning ability of path-based methods. However, in the realm of KG, researchers are still exploring systematic approaches to effectively capture such commonsense knowledge.

To further improve the performance of path-based methods with PLMs by incorporating the aforementioned deducible evidence for relation prediction, we propose **Anchoring Path Sentence Transformer (APST)**, which comprises the following key innovations:

Enriched text description: We develop a search-based description enrichment approach, aiming to systematically capture relevant, more detailed descriptions from designated public websites.

Anchoring Path (AP) extraction and filtering: To handle the situations when CPs are absent and alternatively provide multiple pieces of supporting evidence, we define AP to identify reliable paths that connect to only the head or tail entity. APs are easy to obtain, while most of them are logically irrelevant w.r.t the given query triplet. Therefore, we propose AP accuracy and AP recall to assess the rationality of APs and filter out the logically irrelevant ones. The remaining APs are used for relation prediction, well preserving the inductive and explainable properties of our method.

Unified Sentence Transformer architecture: We fine-tune a unified Sentence Transformer to encode inputs comprising both APs and CPs and subsequently make comprehensive relation predictions with intuitive explanations.

Our key contributions in this work are as follows:

(i) We define Anchoring Path (AP) and two AP evaluation metrics, namely AP accuracy and AP recall, to capture and assess supporting evidence other than CPs in KGs.

(ii) We propose a unified Sentence Transformer named APST, which takes inputs of both APs and CPs with the complement of external textual knowledge captured from public websites.

(iii) We evaluate the performance of APST on three public datasets. APST achieves SOTA performance in 30 of 36

transductive, inductive, and few-shot experimental settings.

Related Work

In this section, we review the three main relation prediction approaches, namely embedding-based, GNN-based, and path-based methods.

Embedding-based methods originate from Word2Vec (Mikolov et al. 2013) and represent entities and relations as dense, low-dimensional vectors in a continuous vector space. By associating entities and relations with these embeddings, embedding-based methods measure the distance between entities and relations, facilitating relation predictions based on proximity. Such methods (e.g., TransE (Bordes et al. 2013), TransH (Wang et al. 2014), TransD (Ji et al. 2015), TransR (Lin et al. 2015), ComplEx (Trouillon et al. 2017), ConvE (Dettmers et al. 2018), and TuckER (Balažević, Allen, and Hospedales 2019)) learn embeddings by optimizing certain scoring functions, which are used to capture the relationships between entities. Embedding-based methods offer the advantage of efficient computation and the ability to capture complex relationships between entities in a compact representation. However, they struggle with the generalization to unseen entities and relations. Additionally, these methods are often not explainable because the low-dimensional vectors are not comprehensible to humans.

GNN-based methods leverage GNNs to encode the structural information and propagate information through the graph. GNNs operate by aggregating and updating entity representations based on the features of their neighboring entities. By iteratively propagating information through the graph, GNNs can capture complex dependencies and higher-order relationships in the KG. Notable GNN-based models for relation prediction include R-GCN (Schlichtkrull et al. 2018), CompGCN (Vashishth et al. 2019), and GRAIL (Teru, Denis, and Hamilton 2020). However, recent studies (Li et al. 2022; Zhang et al. 2022) explored the effective components of GNN-based methods in KGs and found that aggregation actually does not significantly affect the model performance. Consequently, GNN-based explanations relying on neighboring weights may not be always convincing. Moreover, while several GNN-based methods (e.g., GRAIL) are inductive, the majority of GNN-based methods still struggle with inductive cases.

Path-based methods leverage CPs in KGs to infer relationships between entities. These methods identify informative paths in CPs connecting two entities and subsequently predict missing relations. Path-based methods often employ path-finding algorithms, such as random walk or breadth-first search, to explore the KG and identify relevant CPs. Notable path-based methods include DeepPath (Xiong, Hoang, and Wang 2017), MINERVA (Das et al. 2017), BERTRL (Zha, Chen, and Yan 2022), and KRST (Su et al. 2023). Compared with other methods, the distinct advantage of path-based ones is the explicit incorporation of the graph structure with rich information encoded in the paths.

Methods with PLMs: Leveraging PLMs (e.g., BERT (Devlin et al. 2018), GPT (Radford et al. 2018), and RoBERTa (Liu et al. 2019)) for text-related tasks has shown

remarkable success in Natural Language Processing (NLP). Recent methods also applied PLMs to enhance the performance of KG-related tasks. Such methods (e.g., KG-BERT (Yao, Mao, and Luo 2019), BERTRL (Zha, Chen, and Yan 2022), and KRST (Su et al. 2023)) utilize the contextualized embeddings generated by PLMs. These embeddings capture the semantics of entities and relations from their text descriptions within text corpora. By encoding KG information from text descriptions, embeddings can be obtained by feeding triplets or CPs to the PLM and subsequently used as input features for relation prediction. In summary, path-based methods with PLMs offer the advantage of capturing both structural and semantic information.

Preliminary

Before introducing APST, we define the inductive relation prediction. Then we introduce Closed Path (CP), which is commonly adopted by path-based methods.

Definition 1 (Inductive Relation Prediction). *Given a training graph $G_{train}(E_{train}, R_{train})$, a testing graph $G_{test}(E_{test}, R_{test})$, and a set of query triplets $Q(E_{query}, R_{query}) = \{(h_q, r_q, t_q) | h_q, t_q \in E_{query}, r_q \in R_{query}\}$, a relation prediction is inductive if:*

- 1) $E_{train} \cap E_{test} = \emptyset, E_{query} \subseteq E_{test}$,
- 2) $R_{test} \subseteq R_{train}, R_{query} \subseteq R_{train}$.

Inductive relation prediction requires the model to handle unseen entities. To evaluate the model’s inductive ability, we adopt the same settings used in (Teru, Denis, and Hamilton 2020; Zha, Chen, and Yan 2022; Su et al. 2023). Specifically, during training, only triplets from the training graph G_{train} are observed, while the testing graph G_{test} is only used to extract supporting evidence during testing. Furthermore, to prevent the model from relying on transductive evidence, we set the strictest condition in Definition 1. Although real-world training and testing graphs may contain overlapping entities, we require E_{train} and E_{test} to have zero overlaps to avoid the model relying on transductive evidence.

Because both testing and query entities are unseen during training, most path-based inductive methods use CPs in KGs for predictions and explanations. For a given query triplet (h_q, r_q, t_q) , triplets with the same relation r_q that appear in the training graph can be represented as follows:

$$G_{train}(r_q) = \{(h, r_q, t) | (h, r_q, t) \in G_{train}\}.$$

Among triplets in $G_{train}(r_q)$, path-based methods identify paths between heads and tails in the following manner:

$$h \xrightarrow{r_1} e_1 \xrightarrow{r_2} e_2 \xrightarrow{r_3} \dots \xrightarrow{r_{n-1}} e_{n-1} \xrightarrow{r_n} t.$$

By disregarding the entities, certain chains of relations commonly appear in the paths. These paths are known as Closed Paths (CPs), which are defined as follows:

$$CP(h, r, t) : h \xrightarrow{r'_1} r'_2 \xrightarrow{r'_3} \dots \xrightarrow{r'_{n-1}} r'_n \xrightarrow{r'_n} t.$$

Such chains of relations, denoted as $R_{CP}(r_q)$ are considered to be logical w.r.t r_q , which is defined as follows:

$$R_{CP}(r_q) = \{(r'_1, r'_2, \dots, r'_n)\}.$$

Starting from h_q , path-based methods traverse the chain of relations in $R_{CP}(r_q)$ on the testing graph and consider the ending entities as potential prediction candidates.

CPs offer crucial connectivity information between heads and tails. However, most existing path-based approaches only rely on CPs for reasoning, which greatly hinders their reasoning ability. In this paper, we propose a more generalized type of path to provide logical supporting evidence w.r.t the query triplets, alleviating the reliance of CPs.

Methodology

In this section, we first introduce our key innovation, i.e., Anchoring Path (AP). We then define AP recall and AP accuracy for filtering out logically irrelevant APs. Furthermore, to fully leverage the knowledge provided by APs, we adopt the SOTA Sentence Transformer model and propose a unified architecture to retrieve detailed descriptions, encode paths into sentences, and make predictions.

Anchoring Path

Definition 2 (Anchoring Path (AP)). *Given a KG $G = \{(h_i, r_i, t_i) | i = 1, 2, 3, \dots, m\}$ and a query triplet (h, r, t) , the head-Anchoring Path (head-AP) and tail-Anchoring Path (tail-AP) are defined as follows:*

$$AP_{head}(h, r, t) : h \xrightarrow{r_1} r_2 \xrightarrow{r_3} \dots \xrightarrow{r_{n-1}} r_n \xrightarrow{r_n} e_n,$$

$$AP_{tail}(h, r, t) : e'_0 \xrightarrow{r'_1} r'_2 \xrightarrow{r'_3} \dots \xrightarrow{r'_{n-1}} r'_n \xrightarrow{r'_n} t.$$

AP_{head} (AP_{tail}) focuses on the head (tail) entity and traverse the KG to find paths connecting to it. The afore-introduced Path 1 is a typical head-AP w.r.t triplet (*Morgan, Profession, Actor*). In Theorems 1 and 2, we discuss the relationship among CPs, APs, and any paths in KG.

Theorem 1. *Let $S(p)$ denote the set of all p paths. For any triplet (h, r, t) , $S(CP(h, r, t)) = S(AP_{head}(h, r, t)) \cap S(AP_{tail}(h, r, t))$.*

Theorem 1 illustrates the relationship between APs and CPs. Specifically, the complete set of CPs is precisely the intersection of both head-APs and tail-APs. APs capture a wider range of diverse paths including all CPs. Such paths contain rich information due to their connections with head or tail entities, particularly useful when CPs are absent.

To effectively leverage all relation information in KGs, inverse relations $R^{-1} = \{r^{-1} | r \in R\}$ and inverse triplets $\{(t_i, r_i^{-1}, h_i) | i = 1, 2, 3, \dots, m\}$ are usually added into the graph. Consequently, paths in the following forms:

$$e_0 \xrightarrow{r_1} r_2 \xrightarrow{r_3} \dots \xrightarrow{r_{n-1}} r_n \xrightarrow{r_n} h,$$

$$t \xrightarrow{r'_1} r'_2 \xrightarrow{r'_3} \dots \xrightarrow{r'_{n-1}} r'_n \xrightarrow{r'_n} e'_n,$$

can be converted to AP_{head} and AP_{tail} , respectively:

$$h \xrightarrow{r_{n-1}^{-1}} r_{n-2}^{-1} \xrightarrow{r_{n-3}^{-1}} \dots \xrightarrow{r_1^{-1}} e_0,$$

$$e'_n \xrightarrow{r'_{n-1}^{-1}} r'_{n-2}^{-1} \xrightarrow{r'_{n-3}^{-1}} \dots \xrightarrow{r'_1^{-1}} t.$$

By incorporating these inverse relations and triplets, paths in $G'(E, R \cup R^{-1})$ could be simply categorized and represented (see Theorem 2).

Theorem 2. Given a KG with inverse relations $G'(E, R \cup R^{-1})$ and a query triplet (h, r, t) , any path in G' can be represented in one of the following exclusive categories:

- (i) A path not traversing through either h or t ;
- (ii) An AP: $AP_{head}(h, r, t)$ or $AP_{tail}(h, r, t)$;
- (iii) A concatenation of APs: $AP_1 \circ AP_2 \circ \dots \circ AP_n$.

Theorem 2 demonstrates that any path having connections with the query triplet can be represented using an AP (Type ii) or a concatenated AP (Type iii). Empirically, paths that do not traverse through either h or t (Type i) are unlikely to provide useful information supporting the reasoning. To capture supporting evidence from both APs (Type ii) and concatenated APs (Type iii), we carefully design the architecture of APST, enabling it to take multiple paths (Types ii and iii) as inputs. For a concatenated AP, we may either regard it as a single path or break it into multiple APs to input together. Based on the preliminary results, we find that breaking concatenated APs into multiple APs usually leads to better performance in practice. Therefore, for all experiments on APST in this paper, we break all concatenated APs and use the resulting APs as inputs.

Anchoring Path Filtering

APs can be extracted with ease by fixating the head or tail entity and then performing a random walk in the given KG. However, APs without restrictions, typically do not contribute logical supporting evidence. For instance, referring to Figure 1, the following two APs do not provide logical evidence w.r.t triplet $(Morgan, Profession, Actor)$:

$$\begin{aligned} Morgan &\xrightarrow{BornIn} Memphis \xrightarrow{State} Tennessee, \\ Morgan &\xrightarrow{Gender} Male. \end{aligned}$$

Such APs constitute the majority of randomly generated APs in real-world KGs. To eliminate these logically irrelevant APs, we propose AP accuracy and AP recall to measure their rationality w.r.t the given query triplet. For brevity, we only use head-APs to introduce the following techniques, because it is trivial to apply these techniques to tail-APs.

AP accuracy and AP recall are designed to identify logical APs. Given a query triplet (h_q, r_q, t_q) and a corresponding $AP_{head}(h_q, r_q, t_q)$, we focus on r_q and the chain of relations $r_{AP_{head}} = (r_1, r_2, r_3, \dots, r_n)$ on the AP. By traversing the training graph G_{train} , we can find multiple paths that satisfy the chain of relations $r_{AP_{head}}$:

$$\begin{aligned} h^1 &\xrightarrow{r_1^1} \xrightarrow{r_2^1} \dots \xrightarrow{r_{n_1-1}^1} \xrightarrow{r_{n_1}^1} e_{n_1}^1, \\ h^2 &\xrightarrow{r_1^2} \xrightarrow{r_2^2} \dots \xrightarrow{r_{n_2-1}^2} \xrightarrow{r_{n_2}^2} e_{n_2}^2, \\ &\dots \\ h^m &\xrightarrow{r_1^m} \xrightarrow{r_2^m} \dots \xrightarrow{r_{n_m-1}^m} \xrightarrow{r_{n_m}^m} e_{n_m}^m. \end{aligned}$$

Let $E_{AP}^{head} = \{h^1, h^2, \dots, h^m\}$ denote the set of all starting entities of such paths and $E_{r_q}^{head} = \{h | (h, r_q, t) \in G_{train}\}$ denote the head entities of all triplets having relation r_q . Then any entity e on G_{train} could be classified into one of the following exclusive categories:

- (i) Path-Triplet (**PT**): $e \in E_{AP}^{head}, e \in E_{r_q}^{head}$,
- (ii) Path-Only (**PO**): $e \in E_{AP}^{head}, e \notin E_{r_q}^{head}$,
- (iii) Triplet-Only (**TO**): $e \notin E_{AP}^{head}, e \in E_{r_q}^{head}$,
- (iv) No-Connection (**NC**): $e \notin E_{AP}^{head}, e \notin E_{r_q}^{head}$.

Logical APs typically possess the following favorable characteristics on G_{train} :

High accuracy: Logical APs are expected to provide supporting evidence w.r.t the query triplet. Thus, if such a logical AP exists in G_{train} , there should be a high likelihood of the query triplet's existence. To measure the existence ratio of the query triplet, we define AP accuracy as follows:

$$acc(r_q, r_{AP_{head}}) = \frac{\#PT}{\#PO + \#PT}, \quad (1)$$

where # denotes the number of all entities in the respective categories.

High recall: Certain APs with high accuracy may not contribute much towards relation predictions because these APs rarely appear in G_{test} . Empirically speaking, such APs usually also appear infrequently in G_{train} . To better preserve the generalization from G_{train} to G_{test} , logical APs corresponding to the query triplet should have high occurrence in G_{train} . To measure the occurrence ratio of an AP, we define AP recall as follows:

$$rec(r_q, r_{AP_{head}}) = \frac{\#PT}{\#TO + \#PT}. \quad (2)$$

With AP accuracy and AP recall defined, the rationality of $r_{AP_{head}}$ corresponding to r_q can be quantified.

During training, we extract logical APs as follows: (i) Choose a relation r_q and extract a set of potential APs corresponding to r_q from the given KG; (ii) Evaluate the extracted APs using (1) and/or (2); (iii) Exclude APs falling below a pre-determined threshold for either of the two metrics from the potential set; (iv) Store remaining logical APs for r_q ; and (v) Choose another relation r'_q and repeat Steps (i) to (iv) until APs for all relations are extracted.

During testing, we can find many APs corresponding to the testing query triplets. However, only APs that match the chain of relations in the stored logical AP sets are extracted and utilized as supporting evidence.

Anchoring Path Sentence Transformer

To fully utilize the supporting evidence provided by the logical APs and incorporate text descriptions for more accurate relation prediction, we propose Anchoring Path Sentence Transformer (APST) by extending the SOTA relation prediction approach (Su et al. 2023). Figure 2 shows the architecture of APST. Specifically, APST sequentially processes the following steps: (i) Detailed description retrieval; (ii) AP extraction and filtering; (iii) Sentence formation and AP encoding; and (iv) APST prediction.

Detailed description retrieval: To effectively utilize the semantic knowledge in PLMs, rich text descriptions of entities and relations are necessary. However, descriptions are not mandatory when constructing KGs, hence are not provided in many KGs. Even the commonly adopted KG datasets, such as FB15k-237 and NELL-995, only provide the entity

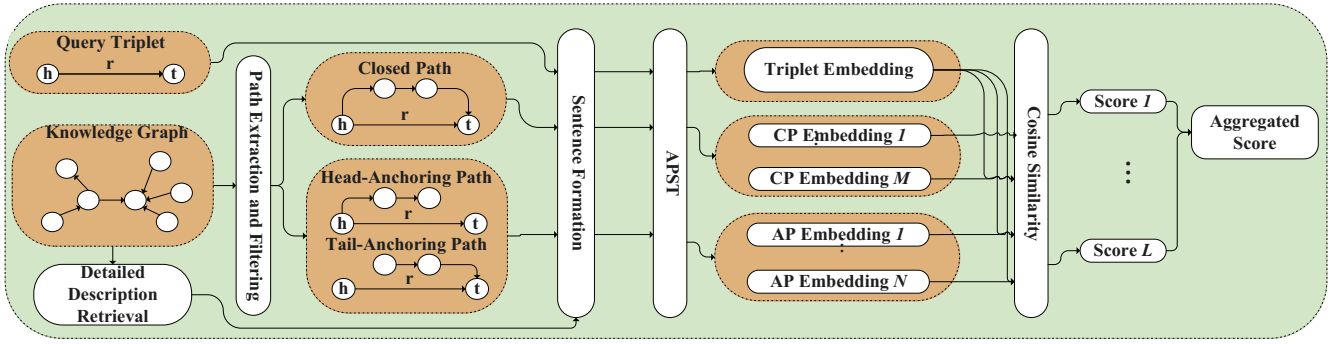


Figure 2: Architecture of Anchoring Path Sentence Transformer (APST).

names with a couple of words as descriptions, making it difficult for the model to accurately capture the semantic meaning of entities. In APST, we deploy Google Custom Search API to search on specialized websites and automatically retrieve detailed descriptions. Google Custom Search allows us to search on specific websites and return pages relevant to the given keywords. Such descriptions on the returned pages are detailed and can be formulated into input sentences.

AP extraction and filtering: As discussed in the context of Theorem 1, CPs are subsets of APs. Because CPs connect to both the head and tail entities, they are usually more informative than other APs. Therefore, APST first generates APs and filters out logically irrelevant ones using (1) and/or (2) as previously introduced. Subsequently, a number of L paths are chosen as inputs for APST, including all CPs and several randomly selected logical APs.

Sentence formation and AP encoding: To take advantage of the text descriptions for entities and relations, we convert the query triplets and the filtered APs into sentences. With detailed descriptions for entities, our input sentences for CPs, APs, and query triplets are formulated as follows:

$$\begin{aligned} \text{Sentence}_{CP} &: [e_1][D_{e_1}][r_1][e_2][r_2] \dots [e_{n-1}][r_n][e_n][D_{e_n}], \\ \text{Sentence}_{AP} &: [e_1][D_{e_1}][r_1][r_2] \dots [r_n][e_n][D_{e_n}], \\ \text{Sentence}_q &: [h_q][D_{h_q}][r_q][t_q][D_{t_q}], \end{aligned}$$

where D_e denotes the detailed descriptions for the corresponding entity e . We only use detailed descriptions for the first and last entities because the descriptions for these two entities are usually the most important ones. Moreover, including detailed descriptions for all entities often results in surpassing the maximum input length permitted by the adopted Sentence Transformer, leading to the unwanted truncation of the overall description. Then we apply Sentence Transformer to encode paths and triplets. Specifically, the formed sentences are encoded into embeddings. These embeddings contain textual information of triplets and APs, which can be used for the next APST prediction step.

APST prediction: During training, triplets with labels 1 and -1 are regarded as positive and negative triplets, respectively. In APST, APs are considered as having similar semantic meanings with the corresponding positive triplet. A query triplet is predicted as positive if at least one of the corresponding APs has a high similarity score with it. To measure

the similarity score between an AP and the corresponding query triplet, we apply cosine similarity as follows:

$$\text{Sim}(s_{hrt}, s_{AP}) = \frac{s_{hrt}s_{AP}}{|s_{hrt}||s_{AP}|}, \quad (3)$$

where s_{hrt} and s_{AP} denote the sentence embeddings of the query triplet and the corresponding AP, respectively. APST allows up to a number of L APs as inputs to compute the similarity scores. The AP with the highest similarity score is considered to be the most reasonable path and the highest similarity score is considered to be the final score for the corresponding triplet:

$$\text{Score}(h, r, t) = \max\{\text{Sim}(s_{hrt}, s_{AP_i}) | i = 1, 2, \dots, L\}. \quad (4)$$

During training, we use cosine embedding loss as the training signal, which is formulated as follows:

$$\mathcal{L}(h, r, t, y) = \begin{cases} 1 - \text{Score}(h, r, t), & y = 1, \\ \max(0, \text{Score}(h, r, t) - M), & y = -1, \end{cases} \quad (5)$$

where $M \in (-1, 1)$ is a pre-defined margin and $y \in \{1, -1\}$ denotes the label.

Experimental Results

We conduct extensive experiments using three datasets to comprehensively evaluate the performance of APST. Additionally, we perform an ablation study to assess the impact and effectiveness of APs and detailed descriptions.

Experimental Setup and Implementation Details

Following the setup used in (Zha, Chen, and Yan 2022) and (Su et al. 2023), we conduct experiments using the commonly benchmarked transductive and inductive datasets introduced by (Teru, Denis, and Hamilton 2020), which are the subsets of WN18RR, FB15k-237, and NELL-995. Moreover, for each query triplet, 1 positive and 49 negative candidate entities are provided by (Zha, Chen, and Yan 2022) for evaluation. We measure the mean reciprocal rank (MRR) and hit rate (Hit@1) among the 50 candidate entities. For inductive experiments, there is zero overlap between training and testing entities. Few-shot experiments are conducted on both inductive and transductive datasets, with 1000 or 2000 training triplets randomly selected.

We implement APST¹ based on the SOTA Sentence

¹github.com/ZhixiangSu/APST

		Transductive			Inductive		
		WN18RR	FB15k-237	NELL-995	WN18RR	FB15k-237	NELL-995
MRR	RuleN	0.669	0.674	0.736	0.780	0.462	0.710
	GRAIL	0.676	0.597	0.727	0.799	0.469	0.675
	MINERVA	0.656	0.572	0.592	-	-	-
	TuckER	0.646	0.682	<u>0.800</u>	-	-	-
	KG-BERT	-	-	-	0.547	0.500	0.419
	BERTRL	0.683	0.695	0.781	0.792	0.605	0.808
	KRST	<u>0.899</u>	<u>0.720</u>	<u>0.800</u>	<u>0.890</u>	<u>0.716</u>	<u>0.769</u>
	APST	0.902	0.774	0.801	0.908	0.764	<u>0.769</u>
Hit@1	RuleN	0.646	0.603	0.636	0.745	0.415	0.638
	GRAIL	0.644	0.494	0.615	0.769	0.390	0.554
	MINERVA	0.632	0.534	0.553	-	-	-
	TuckER	0.600	0.615	0.729	-	-	-
	KG-BERT	-	-	-	0.436	0.341	0.244
	BERTRL	0.655	0.620	0.686	0.755	0.541	0.715
	KRST	<u>0.835</u>	<u>0.639</u>	0.694	<u>0.809</u>	<u>0.600</u>	0.649
	APST	0.839	0.694	<u>0.698</u>	0.837	0.643	<u>0.663</u>

Table 1: Transductive and inductive results on WN18RR, FB15k-237 and NELL-995

Transformer (all-mpnet-base-v2²) using PyTorch and train it on two NVIDIA Tesla V100 GPUs with 32GB RAM.

In APST, we apply both (1) and (2) with thresholds for filtering. To automatically retrieve detailed descriptions, we use the entity names along with short descriptions in all datasets as keywords for Google Custom Search on entity deciphering websites. Because FB15k-237 and NELL-995 only contain commonsense knowledge, we search their entities on Wikipedia. For multilingual words in WN18RR, we search them on Wiktionary.

Transductive and Inductive Relation Predictions

Among these baselines compared to APST, TuckER, GRAIL, and KRST are the SOTA embedding-based, GNN-based, and path-based methods, respectively.

For both transductive and inductive settings, training graph G_{train} is used to extract logical APs for training APST. However, during testing in inductive settings, APs for query triplets do not exist in G_{train} due to zero overlap between E_{train} and E_{query} (see Definition 1). Therefore, we follow the configuration applied in (Su et al. 2023) and extract APs from G_{test} in inductive settings.

Table 1 reports the results of transductive and inductive relation predictions. APST achieves the best (bolded) MRR performance in 5 out of 6 settings and the best Hit@1 rate in 4 out of 6. For the remaining settings, APST achieves the second-best (underlined) performance. Notably, APST’s performance is significantly better in the transductive setting of FB15k-237 (+0.054 for MRR and +5.5% for Hit@1), inductive setting of WN18RR (+0.018 for MRR and +2.8% for Hit@1), and inductive setting of FB15k-237 (+0.048 for MRR and +4.3% for Hit@1). In both the transductive and inductive settings of NELL-995, APST outperforms KRST. However, APST only achieves the second-best performance on NELL-995, we hold the view that APST’s performance

is limited by the retrieved detailed descriptions (see ablation study results for more details).

Few-shot Relation Predictions

We use the subsets, containing 1000 (or 2000) training triplets for all three datasets provided by (Zha, Chen, and Yan 2022), to further benchmark the performance of APST in the few-shot experiments. As shown in Table 2, APST demonstrates superior performance in majority few-shot scenarios (10 out of 12 for MRR and 11 out of 12 for Hit@1). For the remaining scenarios, APST achieves the second-best performance. These results underscore the robust transferability of APST, showcasing its remarkable performance even with a limited number of training pairs.

Ablation Study Results

To further investigate the effectiveness of the key components in APST, we conduct ablation studies using all three datasets in both the transductive and inductive settings. The various testing scenarios are listed as follows:

- (i) **SC**: Short descriptions with CPs only,
- (ii) **SA**: Short descriptions with APs,
- (iii) **DC**: Detailed descriptions with CPs only,
- (iv) **DA**: Detailed descriptions with APs.

Specifically, for SC and SA, we remove the retrieved detailed descriptions and only apply the short descriptions provided in the respective datasets. For SC and DC, we only use CPs and remove all other APs for encoding and prediction.

The results of the ablation studies are visualized in Figure 3. As shown, DA achieves the best performance in 5 out of 6 cases in terms of MRR and Hit@1, respectively. This finding demonstrates that both long descriptions and APs contribute to elevated performance. DA leads to a better performance than DC in all settings, suggesting that with long descriptions, APs always lead to performance gain. This phenomenon does not apply between SA and SC. Thus,

²huggingface.co/sentence-transformers/all-mpnet-base-v2

	Transductive						Inductive					
	WN18RR		FB15k-237		NELL-995		WN18RR		FB15k-237		NELL-995	
	1000	2000	1000	2000	1000	2000	1000	2000	1000	2000	1000	2000
MRR												
RuleN	0.567	0.625	0.434	0.577	0.453	0.609	0.681	0.773	0.236	0.383	0.334	0.495
GRAIL	0.588	0.673	0.375	0.453	0.292	0.436	0.652	0.799	0.380	0.432	0.458	0.462
MINERVA	0.125	0.268	0.198	0.364	0.182	0.322	-	-	-	-	-	-
TuckER	0.258	0.448	0.457	0.601	0.436	0.577	-	-	-	-	-	-
KG-BERT	-	-	-	-	-	-	0.471	0.525	0.431	0.460	0.406	0.406
BERTRL	0.662	0.673	0.618	0.667	0.648	0.693	0.765	0.777	0.526	0.565	0.736	.744
KRST	.871	0.882	.696	.701	.743	0.781	.886	.878	.679	.680	.745	0.738
APST	0.874	.880	0.724	0.753	0.745	.767	0.894	0.879	0.697	0.747	0.765	0.747
Hit@1												
RuleN	0.548	0.605	0.374	0.508	0.365	0.501	0.649	0.737	0.207	0.344	0.282	0.418
GRAIL	0.489	0.633	0.267	0.352	0.198	0.342	0.516	0.769	0.273	0.351	0.295	0.298
MINERVA	0.106	0.248	0.170	0.324	0.152	0.284	-	-	-	-	-	-
TuckER	0.230	0.415	0.407	0.529	0.392	0.520	-	-	-	-	-	-
KG-BERT	-	-	-	-	-	-	0.364	0.404	0.288	0.317	0.236	0.236
BERTRL	0.621	0.637	0.517	0.583	0.526	0.582	0.713	0.731	0.441	0.493	0.622	0.628
KRST	.790	.810	.611	.602	.628	0.678	.811	.793	.537	.524	.637	.629
APST	0.798	0.813	0.632	0.665	0.640	.663	0.822	0.798	0.561	0.627	0.654	0.637

Table 2: Few-shot results on WN18RR, FB15k-237 and NELL-995

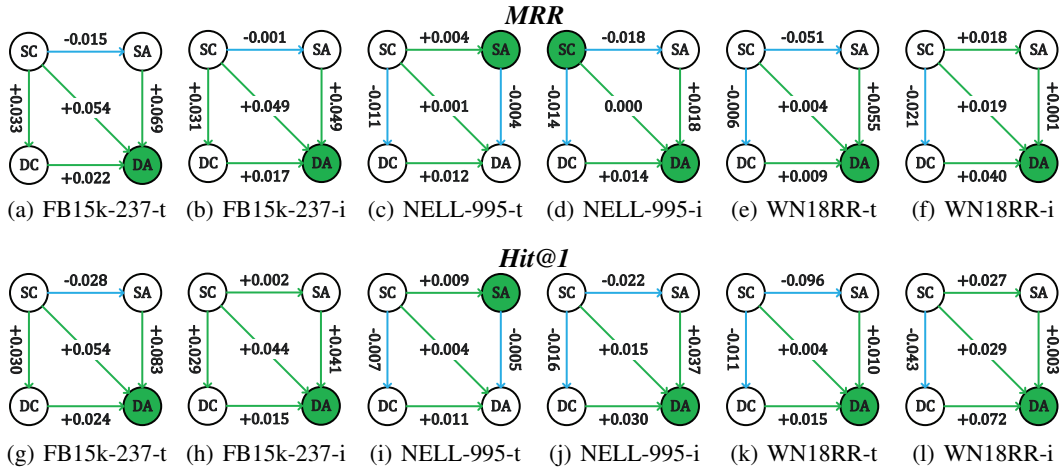


Figure 3: Ablation study results on FB15k-237, NELL-995 and WN18RR (-t: transductive; -i: inductive).

we hold the view that detailed descriptions are significantly helpful when APs are used for predictions. Comparing SC with DC and comparing SA with DA, results show that detailed descriptions do not always lead to better performance (5 of 12 for both MRR and Hit@1). Probably the most plausible reason is that our detailed descriptions are obtained from Wikipedia and Wiktionary. FB15k-237 originates from Wikipedia, and its entities exactly match the objects found in Wikipedia. However, for entities in NELL-995 and WN18RR, their meanings may not exactly match the descriptions, which leads to a performance decrease. In summary, although using detailed descriptions may negatively impact the effectiveness of CPs, utilizing APs with detailed descriptions will result in a notable enhancement.

Conclusion

In this paper, we propose APST, which is a novel path-based method utilizing Sentence Transformer. APST retrieves detailed descriptions, captures supporting evidence in the form of Anchoring Paths (APs), and incorporates both AP and CP inputs for comprehensive predictions and convincing explanations. Through experiments on commonly applied datasets, APST demonstrates the potential for leveraging textual knowledge from external sources to enrich KGs and take advantage of APs to improve prediction accuracy.

Nonetheless, the current APST model primarily focuses on relatively small datasets. Going forward, we aim to develop a more memory-efficient variant to handle increased data complexity and apply APST to larger datasets.

Acknowledgements

This research is supported, in part, by the National Research Foundation, Prime Minister’s Office, Singapore under its NRF Investigatorship Programme (NRFI Award No. NRF-NRFI05-2019-0002). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore. This research is also supported by the National Key R&D Program of China No. 2021YFF0900800 and Shandong Provincial Key Research and Development Program (Major Scientific and Technological Innovation Project) (No. 2021CXGC010108).

References

- Balažević, I.; Allen, C.; and Hospedales, T. M. 2019. Tucker: Tensor factorization for knowledge graph completion. *arXiv preprint arXiv:1901.09590*.
- Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In *Proceedings of Advances in Neural Information Processing Systems*, 2787–2795.
- Das, R.; Dhuliawala, S.; Zaheer, M.; Vilnis, L.; Durugkar, I.; Krishnamurthy, A.; Smola, A.; and McCallum, A. 2017. Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning. *arXiv preprint arXiv:1711.05851*.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1811–1819.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ji, G.; He, S.; Xu, L.; Liu, K.; and Zhao, J. 2015. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, 687–696.
- Li, J.; Shomer, H.; Ding, J.; Wang, Y.; Ma, Y.; Shah, N.; Tang, J.; and Yin, D. 2022. Are graph neural networks really helpful for knowledge graph completion? *arXiv preprint arXiv:2205.10652*.
- Lin, Y.; Liu, Z.; Sun, M.; Liu, Y.; and Zhu, X. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2181–2187.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training. <https://openai.com/blog/language-unsupervised/>.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Berg, R. v. d.; Titov, I.; and Welling, M. 2018. Modeling relational data with graph convolutional networks. In *Proceedings of European Semantic Web Conference*, 593–607.
- Su, Z.; Wang, D.; Miao, C.; and Cui, L. 2023. Multi-Aspect Explainable Inductive Relation Prediction by Sentence Transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 6533–6540.
- Teru, K.; Denis, E.; and Hamilton, W. 2020. Inductive relation prediction by subgraph reasoning. In *Proceedings of International Conference on Machine Learning*, 9448–9457.
- Trouillon, T.; Dance, C. R.; Welbl, J.; Riedel, S.; Gaussier, E.; and Bouchard, G. 2017. Knowledge graph completion via complex tensor factorization. *arXiv preprint arXiv:1702.06879*.
- Vashishth, S.; Sanyal, S.; Nitin, V.; and Talukdar, P. 2019. Composition-based multi-relational graph convolutional networks. *arXiv preprint arXiv:1911.03082*.
- Wang, Z.; Zhang, J.; Feng, J.; and Chen, Z. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1112–1119.
- Xiong, W.; Hoang, T.; and Wang, W. Y. 2017. DeepPath: A reinforcement learning method for knowledge graph reasoning. *arXiv preprint arXiv:1707.06690*.
- Yao, L.; Mao, C.; and Luo, Y. 2019. KG-BERT: BERT for knowledge graph completion. *arXiv preprint arXiv:1909.03193*.
- Zha, H.; Chen, Z.; and Yan, X. 2022. Inductive relation prediction by BERT. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 5923–5931.
- Zhang, Z.; Wang, J.; Ye, J.; and Wu, F. 2022. Rethinking graph convolutional networks in knowledge graph completion. In *Proceedings of the ACM Web Conference 2022*, 798–807.