

A Joint Framework for Inductive Representation Learning and Explainable Reasoning in Knowledge Graphs

Rajarshi Bhowmik

Rutgers University - New Brunswick
rajarshi.bhowmik@rutgers.edu

Gerard de Melo

Rutgers University - New Brunswick
gdm@demelo.org

Abstract

Despite their large-scale coverage, existing cross-domain knowledge graphs invariably suffer from inherent incompleteness and sparsity, necessitating link prediction that requires inferring a target entity, given a source entity and a query relation. Recent approaches can broadly be classified into two categories: embedding-based approaches and path-based approaches. In contrast to embedding-based approaches, which operate in an uninterpretable latent semantic vector space of entities and relations, path-based approaches operate in the symbolic space, making the inference process explainable. However, traditionally, these approaches are studied with static snapshots of the knowledge graphs, severely restricting their applicability for dynamic knowledge graphs with newly emerging entities. To overcome this issue, we propose an inductive representation learning framework that is able to learn representations of previously unseen entities. Our method finds reasoning paths between source and target entities, thereby making the link prediction for unseen entities interpretable and providing support evidence for the inferred link.

1 Introduction

Recent years have seen a surge in the usage of large-scale cross-domain knowledge graphs for various natural language processing tasks, including factoid question answering, fact-based dialogue engines, and information retrieval. Knowledge graphs serve as a source of background factual knowledge for a wide range of applications. For example, Google’s knowledge graph is tightly integrated into its search engine, while Apple adopted Wikidata as a source of background knowledge for its virtual assistant Siri. Many such applications deal with natural language queries that can be transformed to a structured relational query of the form

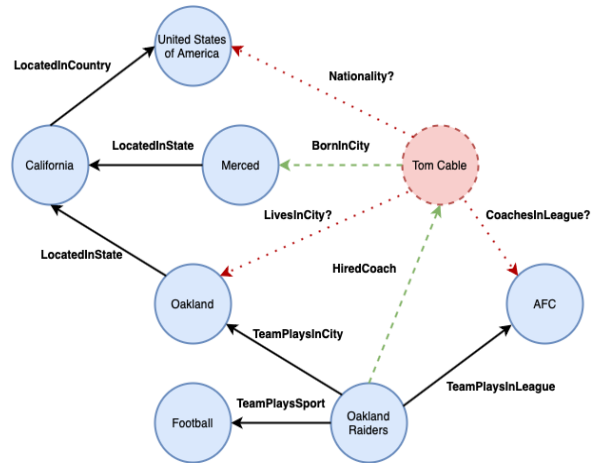


Figure 1: A subgraph of NELL with *Tom Cable* as an emerging entity. The solid-lined circles and arrows represent the existing entities and relations in the knowledge graph. The dashed-lined circles and arrows denote the emerging entity and some of its known relationships to other existing entities. The unknown relationships that need to be inferred through inductive representation learning and explainable reasoning are shown as dotted arrows.

$(e_s, r_q, ?)$, where e_s is the source entity and r_q is the query relation. For example, the query “Who is the director of World Health Organization?” can be mapped to the structured query $(World\ Health\ Organization, director, ?)$, while executing it on a knowledge graph. Unfortunately, due to the inherent sparsity and incompleteness of knowledge graphs, answers to many such queries cannot be fetched directly from the existing facts, but instead need to be inferred indirectly.

Furthermore, with the ever-increasing volume of the knowledge graphs, the number of emerging entities also increases. Many of these emerging entities have a small number of known facts at the time they are integrated into the knowledge graphs. Therefore, their connectivity to pre-existing entities in the knowledge graph is often too sparse.

Typically, embedding-based models (Nguyen, 2017) are used to infer missing relationships in a knowledge graph. In such embedding-based models, distributed vector representations of entities and relations in the knowledge graph are used to learn a scoring function $f(e_s, r_q, e_t)$ in the embedding space to determine the plausibility of inferring a new fact. However, these models are lacking in terms of the interpretability and explainability of the decisions they make. One does not obtain any clear explanation of why a specific inference is warranted. For example, from the embeddings of facts ($A, \text{born.in, California}$) and ($\text{California located.in, US}$), the fact ($A, \text{born.in, US}$) could be deduced. But logical composition steps like this one are learned implicitly by knowledge graph embeddings. This means that this approach cannot offer such logical inference paths as support evidence for an answer.

In contrast, path-based reasoning approaches (Lao et al., 2011; Gardner et al., 2013, 2014; Neelakantan et al., 2015; Guu et al., 2015) operate in the symbolic space of entities and relations, leveraging the symbolic compositionality of the knowledge graph relations, thus making the inference process explainable. This means that the user can inspect the inference path, consisting of existing edges in the knowledge graph, as support evidence. Recent path-based reasoning approaches (Das et al., 2017; Lin et al., 2018) formulate the graph-walking problem as a Partially Observable Markov Decision Process (POMDP) in which the agent learns a policy to find the inference path from the source entity to the target entity using REINFORCE (Williams, 1992), a policy gradient based reinforcement learning algorithm.

However, traditionally, these approaches are studied with static snapshots of the knowledge graphs, thus severely restricting their applicability for a dynamic knowledge graph with many emerging entities.

To overcome this issue, we propose a joint framework for representation learning and reasoning in knowledge graphs that aims at achieving inductive node representation learning capabilities applicable to a dynamic knowledge graph with many emerging entities while preserving the unique advantage of the path-based approaches in terms of explainability. For inductive node representation learning, we propose a variant of *Graph Transformer* encoder (Koncel-Kedziorski et al., 2019; Yun et al., 2019)

that aggregates neighborhood information based on its relevance to the query relation. Furthermore, we use policy gradient-based reinforcement learning (REINFORCE) to decode a reasoning path to the answer entity. We hypothesize that the inductively learned embeddings provide prior semantic knowledge about the underlying knowledge environment to the reinforcement learning agent.

We summarize the contributions of this paper as follows: (1) We introduce a joint framework for inductive representation learning and explainable reasoning that is capable of learning representations for unseen emerging entities by leveraging only a small number of known connections to the other existing entities in the knowledge graph. Our approach can not only infer new connections between an emerging entity and any other existing entity in the knowledge graph, but also provides an explainable reasoning path as support evidence for the inference. (2) We introduce new train/development/test set splits of existing knowledge graph completion benchmark datasets that are appropriate for inductive representation learning and reasoning.

2 Related Work

2.1 Embedding-based Methods

Embedding-based methods are the most popular approach for knowledge graph completion. Such methods learn d -dimensional distributed vector representations of entities and relations in a knowledge graph. To this end, the translation embedding model TransE (Bordes et al., 2013) learns the embeddings of the relation as a translation vector from the source entity to the target entity such that $\mathbf{e}_s + \mathbf{e}_r \approx \mathbf{e}_o$. Its variants, e.g., TransH (Wang et al., 2014), TransR (Lin et al., 2015), TransD (Ji et al., 2015) consider similar objectives. Tri-linear models such as DistMult (Yang et al., 2014), and its variant in the complex embedding space ComplEx (Trouillon et al., 2016) use a multiplicative scoring function $f(s, r, o) = \mathbf{e}_s^T \mathbf{W}_r \mathbf{e}_o$, where \mathbf{W}_r is a diagonal matrix representing the embedding of relation r . Recent convolutional neural network based models such as ConvE (Dettmers et al., 2018) and ConvKB (Nguyen et al., 2018) apply convolutional kernels over entity and relation embeddings to capture the interactions among them across different dimensions. These models obtain state-of-the-art results on the benchmark KBC datasets. However, none of the above-mentioned approaches deliver

the full reasoning paths that license specific multi-hop inferences, and hence they either do not support multi-hop inference or do not support it in an interpretable manner. Moreover, these approaches assume a static snapshot of the knowledge graph to train the models and are not naturally extensible to inductive representation learning with previously unseen entities.

2.2 Graph Convolution-based Methods

Kipf and Welling (2017) introduced Graph Convolution Networks (GCN) for node classification in a homogeneous graph. They are an instance of Message Passing Neural Networks (MPNN), in which the node representations are learned by aggregating information from the node’s local neighborhood. GraphSAGE (Hamilton et al., 2017) attempts to reduce the memory footprint of GCN by random sampling of the neighborhood. Graph Attention Networks (GAT) (Veličković et al., 2018) are a variant of GCN that learn node representations as weighted averages of the neighborhood information. However, GCN and its variants such as GAT and GraphSAGE are not directly applicable for link prediction in knowledge graphs, as they ignore the edge (relation) information for obtaining the node embeddings. To alleviate this issue, Schlichtkrull et al. (2018) introduced R-GCN for relational multi-graph. However, similar to GCN, R-GCN also needs all nodes of the graphs to be present in the memory and therefore is not scalable to large-scale knowledge graphs. Hamaguchi et al. (2017) proposed a model for computing representations for out-of-KG entities using graph neural networks. Recently, Shang et al. (2019) proposed SACN, in which the model learns the representation of the nodes (entities) using the structure of the graph. However, unlike our model, these methods are not explainable.

2.3 Path-based Methods.

An alternative stream of research has explored means of identifying specific paths of inference, which is the task we consider in this paper. To this end, the Path Ranking Algorithm (PRA) by (Lao et al., 2011) uses random walks with restarts for multi-hop reasoning. Following PRA, other approaches (Gardner et al., 2013, 2014; Neelakantan et al., 2015; Guu et al., 2015) also leverage random walk based inference. However, the reasoning paths that these methods follow are gathered by random walks independently of the query relation.

Recent approaches have instead adopted policy gradient based reinforcement learning for a more focused exploration of reasoning paths. Policy gradient based models such as DeepPath (Xiong et al., 2017), MINERVA (Das et al., 2017), MINERVA with Reward Shaping and Action Dropout (Lin et al., 2018), and M-Walk (Yelong Shen and Gao, 2018) formulate the KG reasoning task as a Markov Decision Process and learn a policy conditioned on the query relation. Although the inference paths are explainable in these models (if reward shaping is omitted), there exists a substantial performance gap with the embedding-based models.

3 Model

Our model consists of two modules that are subject to joint end-to-end training. The encoder module learns inductive entity embeddings while accounting for the query relation and the local neighborhood of an entity. The decoder module operates on this learned embedding space of entities and relations. By leveraging the embeddings of the source entity and the query relation, the decoder module infers a reasoning path to the target entity using policy gradient-based reinforcement learning.

3.1 Problem Statement

Formally, a knowledge graph $\mathcal{G}(\mathcal{E}, \mathcal{R}, \mathcal{F})$ is a directed multi-graph such that each node $e \in \mathcal{E}$ represents an entity, each $r \in \mathcal{R}$ represents a unique relation, and each directed edge $(e_s, r, e_o) \in \mathcal{F}$ represents a fact about the subject entity e_s .

Given a structured relational query $(e_s, r_q, ?)$, where e_s is the source entity, r_q is the query relation, and $(e_s, r_q, e_o) \notin \mathcal{F}$, the goal is find a set of plausible answer entities $\{e_o\}$ by navigating paths through the existing entities and relations in \mathcal{G} leading to answer entities. Note that, unlike previous methods that consider transductive settings with a static snapshot of the knowledge graph, we allow for dynamic knowledge graphs, where e_s may be an emerging entity, and therefore, previously unseen. Moreover, while embedding-based methods only deliver candidate answer entities, we here seek the actual paths, i.e., sequences of nodes and edges for better interpretability.¹

¹From here onwards, we will use the terms *node* and *entity*, as well as *edge* and *relation(ship)* interchangeably.

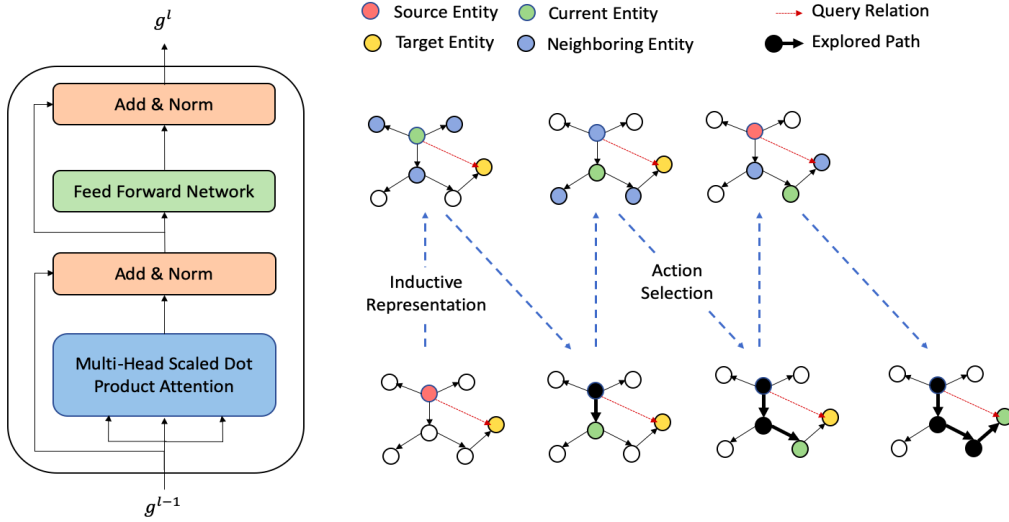


Figure 2: A schematic diagram of a Graph Transformer block, along with an illustration of the workflow of our model, demonstrating successive applications of inductive node representation learning and action selection to find a reasoning path.

3.2 Graph Transformer for Inductive Representation Learning

The state-of-the-art embedding based models either focus on learning entity embeddings by using only the query relations, ignoring the subject entity’s neighborhood, or use message passing neural networks to learn embeddings conditioned on neighboring entities and relations while being oblivious of the query relation. However, we observe that in many cases a new fact can be inferred by using another existing fact. For example, the fact $(PersonX, Place\ of\ Birth, Y)$ can often help to answer to the query $(PersonX, Nationality, ?)$. Motivated by this observation, we propose a Graph Transformer architecture that learns the embedding of the source entity by iterative aggregation of neighborhood information (messages) that are weighted by their relevance to the query relation. To learn the relevance weights, our Graph Transformer model deploys *multi-head scaled dot product attention*, also known as *self-attention* (Vaswani et al., 2017).

Formally, we denote the local neighborhood for each entity $e_i \in \mathcal{E}$ as \mathcal{N}_i such that $\mathcal{N}_i = \{e_j | e_j \in \mathcal{E} \wedge (e_i, r, e_j) \in \mathcal{F} \wedge r \in \mathcal{R}_{ij}\}$. \mathcal{R}_{ij} is the set of relations between entities e_i and e_j .

Each neighboring entity $e_j \in \mathcal{N}_i$ connected to e_i by a relation r sends in a message to entity e_i . The message \mathbf{m}_{ijr} is a linear transformation of the fact (e_i, r, e_j) followed by the application of a non-linear function, i.e., LeakyReLU. Formally,

$$\mathbf{m}_{ijr} = \text{LeakyReLU}(\mathbf{W}_f[\mathbf{e}_i; \mathbf{r}; \mathbf{e}_j]), \quad (1)$$

where $\mathbf{W}_f \in \mathbb{R}^{d \times 3d}$ is a shared parameter for the linear transformation and $[\cdot]$ is the concatenation operator.

To compute an attention head, our model performs linear projections of the query relation r_q , the neighborhood relations $r \in \mathcal{R}_{ij}$, and the neighborhood messages \mathbf{m}_{ijr} to construct queries Q , keys K , and values V , respectively, such that $Q = \mathbf{W}_Q \mathbf{r}_q$, $K = \mathbf{W}_K \mathbf{r}$, and $V = \mathbf{W}_V \mathbf{m}_{ijr}$, where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d' \times d}$ are learnable parameters.

Next, we use the queries Q to perform a dot-product attention over the keys K . Formally,

$$\alpha_{ijr} = \frac{\exp((\mathbf{W}_Q \mathbf{r}_q)^\top (\mathbf{W}_K \mathbf{r}))}{\sum_{z \in \mathcal{N}_i} \sum_{r' \in \mathcal{R}_{ij}} \exp((\mathbf{W}_Q \mathbf{r}_q)^\top (\mathbf{W}_K \mathbf{r}'))} \quad (2)$$

Following Vaswani et al. (2017), we scale the dot products of Q and K by a factor of $\frac{1}{\sqrt{d}}$.

The attention weights are then used to aggregate the neighborhood messages. Note that *self-attention* deploys multiple attention heads, each having its own query, key, and value projectors. The aggregated messages from N attention heads are concatenated and added to the initial embedding \mathbf{e}_i through a residual connection to obtain new intermediate representation

$$\hat{\mathbf{e}}_i = \mathbf{e}_i + \parallel_{n=1}^N \left(\sum_{j \in \mathcal{N}_i} \sum_{r \in \mathcal{R}_{ij}} \alpha_{ijr}^n \mathbf{W}_V^n \mathbf{m}_{ijr} \right), \quad (3)$$

where \parallel is the concatenation operator.

Layer normalization (LN) is applied to the intermediate representation $\hat{\mathbf{e}}_i$, followed by a fully connected two-layer feed forward network (FFN) with a non-linear activation (ReLU) in between. Finally, the output of the feed forward network is added to the intermediate representation through another residual connection. The resulting embedding is again layer normalized to obtain the new representation \mathbf{g}_i^L for e_i . Formally,

$$\mathbf{g}_i^L = \text{LN}(\text{FFN}(\text{LN}(\hat{\mathbf{e}}_i)) + \text{LN}(\hat{\mathbf{e}}_i)) \quad (4)$$

This pipeline is called a *transformer block*. Figure 2 represents a schematic diagram of a transformer block in Graph Transformers. We stack L -layers of transformer blocks to obtain the final embedding \mathbf{g}_i^L for e_i .

3.3 Policy Gradient for Explainable Reasoning

To infer the answer entity, we could leverage the entity representations obtained by the Graph Transformers. However, our goal is not only to infer the answer entity, but to find a symbolic reasoning path to support the inference. Following Das et al. (2017) and Lin et al. (2018), we formulate the reasoning task as a finite horizon, deterministic partially observable Markov Decision Process (POMDP). As mentioned in Das et al. (2017), a knowledge graph can be seen as a partially observable environment with out-going relations at each entity node corresponding to a set of discrete actions that an agent can explore to reach the target answer from the source entity.

Knowledge Graph Environment Formally, a Markov Decision Process is defined by a 4-tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R})$, where \mathcal{S} is a finite set of states, \mathcal{A} is a finite set of actions, \mathcal{P} captures state transition probabilities, and \mathcal{R} is the reward function. In a knowledge graph environment, the state space is defined as a set of tuples $s_t = (e_t, r_q) \in \mathcal{S}$, where e_t is an entity node in the knowledge graph, and r_q is the query relation. The action space $A_t \in \mathcal{A}$ for a state s_t is defined as the set of outgoing edges from entity node e_t in the knowledge graph. Formally, $A_t = \{(r_{t+1}, s_{t+1}) \mid (e_t, r_{t+1}, s_{t+1}) \in \mathcal{G}\}$. Since state transitions in a KG environment are deterministic, the transition probabilities $P(s_{t+1} \mid s_t, a_t) = 1 \forall P \in \mathcal{P}$. The agent receives a terminal reward of 1 if it arrives at the correct answer entity at the end.

Graph Search Policy To find a plausible path to the answer entity, the model must have a policy to choose the most promising action at each state. Note that in the KG environment, the decision of choosing the next action is not only dependent on the current state, but also on the sequence of observations and actions taken so far in the path. We use a multi-layer LSTM as a sequence encoder to encode the path history.

Formally, each state s_t is represented by a vector $\mathbf{s}_t = [\mathbf{e}_t; \mathbf{r}_q] \in \mathbb{R}^{2d}$ and each possible action $a_t \in A_t$ is represented by $\mathbf{a}_t = [\mathbf{e}_{t+1}; \mathbf{r}_{t+1}] \in \mathbb{R}^{2d}$, where $\mathbf{e}_t, \mathbf{e}_{t+1} \in \mathbb{R}^d$ are the embeddings of the entity nodes at timestep t and $t+1$, respectively, that are obtained from Graph Transformer encoders. $\mathbf{r}_{t+1} \in \mathbb{R}^d$ is the embedding of an out-going relation from entity e_t , and $\mathbf{r}_q \in \mathbb{R}^d$ corresponds to the embedding of the query relation r_q . Each of these embeddings is also obtained from the graph transformer encoder. The path history is encoded as $\mathbf{h}_t = \text{LSTM}(\mathbf{h}_{t-1}, \mathbf{a}_{t-1})$. Given the embedded action space $\mathbf{A}_t \in \mathbb{R}^{2|A_t|}$, i.e., the stacked embeddings of actions $a_t \in A_t$, and the path history \mathbf{h}_t , we define the parameterized policy as:

$$\pi_\theta(a_t \mid s_t) = \text{Softmax}(\mathbf{A}_t(\mathbf{W}_2 \text{ReLU}(\mathbf{W}_1[\mathbf{h}_t; \mathbf{e}_t; \mathbf{r}_q])))$$

Policy Optimization The policy network is trained to maximize the expected reward for all (e_s, r_q, e_o) triples in the training sub-graph. The agent learns an optimal policy π_θ by exploring a state space of all possible actions. The objective of the agent is to take actions to maximize the expected end reward. Formally,

$$J(\theta) = \mathbb{E}_{(e_s, r_q, e_o)} [\mathbb{E}_{a_1, \dots, a_{T-1} \sim \pi_\theta} [R(s_T \mid e_s, r_q)]] \quad (5)$$

Since policy gradient uses gradient-based optimization techniques, the estimated gradient of the objective function can be derived as follows:

$$\nabla_\theta J(\theta) = \mathbb{E}_{a_{1:T} \sim \pi_\theta} [\nabla_\theta \log \pi_\theta(a_{1:T} \mid e_s, r_q) R(s_T \mid e_s, r_q)] \quad (6)$$

$$\approx \frac{1}{N} \sum_{n=1}^N \nabla_\theta \log \pi_\theta(a_{1:T}^n \mid e_s, r_q) R \quad (7)$$

Here, N is the number of policy rollouts.

Each policy rollout explores a sequence of actions $a_{1:T}$. At each timestep $t \in \{1 : T\}$, the agent selects an action a_t conditioned on the current state s_t . Therefore, the gradient of the log-likelihood in

Dataset	\mathcal{E}	\mathcal{R}	\mathcal{U}	\mathcal{F}			
				train	dev	test	aux
FB15k-237-Inductive	13,119	237	1,389	227,266	17,500	32,197	61,330
WN18RR-Inductive	35,928	11	4,029	67,564	3,000	11,015	19,395
NELL-995-Inductive	71,578	200	776	137,221	500	1,679	2,267

Table 1: Evaluation datasets for inductive setting

Eq. 6 can be expressed as

$$\nabla_{\theta} \log \pi_{\theta}(a_{1:T}|e_s, r_q) = \sum_{t=1}^T \nabla_{\theta} \log \pi_{\theta}(a_t|s_t, e_s, r_q) \quad (8)$$

Reward Shaping It is observed by Lin et al. (2018) that a soft reward for the target entities is more beneficial than a binary reward. Following their work, we use pre-trained ConvE (Dettmers et al., 2018) embeddings for the observed entities and relations to shape the reward function. If the agent reaches the correct answer entity, it receives reward 1. Otherwise, the agent receives a reward estimated by the scoring function of the pre-trained ConvE.

4 Evaluation

4.1 Datasets

We evaluate our model based on three standard benchmark knowledge graph completion datasets. (1) FB15k-237, introduced by Toutanova et al. (2016) as a replacement for the FB15k dataset originally introduced by Bordes et al. (2013). The original FB15k dataset suffers from test set leakage such that a significant number of test triples can be determined by simply reversing the relation between the two entities in the training set. In FB15k-237, the reverse relations are removed, rendering the dataset more challenging for inference. (2) WN18RR, introduced by Dettmers et al. (2018), is a subset of the WN18 benchmark dataset. Similar to FB15k-237, the reverse relations are removed for this dataset. (3) NELL-995, introduced by Xiong et al. (2017), is a subset of the 995-th iteration of NELL.

To test the effectiveness of our model for inductive representation learning and reasoning, we create new splits of training, development, and test sets for each of the three benchmark datasets mentioned above. This new split of the dataset is necessary, as in an inductive setting, the subject entities in the test set must not be present anywhere in the

training subgraph. To satisfy this requirement, we first sample 10% of all the entities present in the each of the benchmark datasets. We denote this set as *unseen entities* \mathcal{U} . The remaining entities are denoted as *seen entities* \mathcal{E} . Then, we proceed to split the triples in the datasets into three disjoint sets. The first set contains the triples in which both the head and the tail entities are in \mathcal{E} . The second set consists of the triples with head entities belonging to \mathcal{U} , but tail entities in \mathcal{E} . In the third set, the head entities belong to \mathcal{E} , but the tail entities are in \mathcal{U} . We further split the first set into *train* and *dev* triples. The second set becomes the *test* triples, and the union of the second and the third set is denoted as *auxiliary* data. Auxiliary triples are required to obtain the local neighborhood of a source entity at inference time. We use the suffix "-Inductive" to distinguish these newly derived datasets from their original counterparts. A summary of these datasets is presented in Table 1.

4.2 Baselines

We choose a state-of-the-art graph convolution-based approach SACN (Shang et al., 2019) as a baseline. Our choice is motivated by two factors: (1) SACN performs strongly in the transductive setting by outperforming the other baselines for most of the datasets, and (2) since its encoder module deploys neighborhood integration through Graph Convolution Networks, it has similar characteristics to our model, and therefore, is a good candidate for inductive representation learning. We also compare our model to R-GCN (Schlichtkrull et al., 2018), which also leverages the graph structure by aggregating neighborhood information. We further consider ConvE (Dettmers et al., 2018) as an additional baseline. Although ConvE is particularly unsuitable for the inductive setting, we include it to better demonstrate the challenges of applying such algorithms in an inductive setting.

Model	WN18RR-Inductive				FB15k-237-Inductive			
	MRR	Hits@N			MRR	Hits@N		
		@1	@3	@10		@1	@3	@10
ConvE (Dettmers et al., 2018)	1.9	1.1	2.1	3.5	26.3	20.0	28.7	38.8
R-GCN (Schlichtkrull et al., 2018)	14.7	11.4	15.1	20.7	19.1	11.5	20.9	34.3
SACN (Shang et al., 2019)	17.5	9.7	20.3	33.5	29.9	20.5	32.8	50.0
Our Model with Reward Shaping	48.8	42.1	52.2	60.6	39.8	30.7	44.5	57.6

Table 2: Results of our model as compared to the state-of-the-art neighborhood-aware embedding models for the WN18RR and FB15K237 datasets in the inductive setting, where 10% of entities from the original datasets are unseen during training. The Hits@N and MRR metrics are multiplied by 100.

Model	NELL-995-Inductive			
	MRR	Hits@N		
		@1	@3	@10
ConvE (Dettmers et al., 2018)	43.4	32.5	50.3	60.9
R-GCN (Schlichtkrull et al., 2018)	58.4	50.9	62.9	71.6
SACN (Shang et al., 2019)	42.4	37.0	42.9	53.2
Our Model with Reward Shaping	75.2	69.7	79.1	84.4

Table 3: Results of our model as compared to the state-of-the-art neighborhood-aware embedding models for the NELL-995 dataset in the inductive setting, where 10% of entities from the original datasets are unseen during training. The Hits@N and MRR metrics are multiplied by 100.

4.3 Experimental Details

Training Protocol Since the benchmark knowledge graph completion datasets contain only unidirectional edges (e_s, r_q, e_o) , we augment the training sub-graph with the reverse edges (e_o, r_q^{-1}, e_s) . During the Graph Transformer based inductive representation learning, $n\%$ of local neighboring entities are randomly selected and masked. During training, we mask 50%, 50%, and 30% of neighboring nodes, respectively, for the FB15k-237, WN18RR, and NELL-995 datasets. Neighborhood masking helps in learning robust representations and reduces the memory footprint, and has been shown to be effective by Hamilton et al. (2017). Following Das et al. (2017) and Lin et al. (2018), during training of the policy network, we also retain the top- k outgoing edges for each entity that are ranked by the PageRank scores of the neighboring entities. We set the value of k for each dataset following Lin et al. (2018). Such a cut-off threshold is necessary to prevent memory overflow. Finally, we adopt the false-negative masking technique in the final timestep of the policy rollouts to guide the agent to the correct answer entities as described in Das et al. (2017) and Lin et al. (2018) and demonstrated by them to be helpful when multiple answer entities are present in the training graph.

Hyperparameters For a fair comparison to the baselines, we keep the dimensionality of the entity and relation embeddings at 200. For ConvE (Dettmers et al., 2018) and SACN (Shang et al., 2019), we follow the same hyperparameter settings as specified in the original implementations. For our model, we deploy one layer of a Transformer block ($L = 1$) and 4 attention heads ($N = 4$). We choose a minibatch size of 64 during training due to limited GPU memory. We use Adam (Kingma and Ba, 2014) as the stochastic optimizer and keep the learning rate fixed at 0.001 across all training epochs. Additionally, we adopt entropy regularization to improve the learning dynamics of the policy gradient method. The regularizer is weighted by a hyperparameter β set to a value within $[0, 0.1]$. We apply dropout to the entity and relation embeddings, the feedforward networks, and the residual connections. The policy rollout is done for $T = 3$ timesteps for every dataset.

Evaluation Protocol Following Lin et al. (2018), we adopt beam search decoding during inference with a beam width of 512 for NELL-995 and 256 for the other datasets. If more than one path leads to the same target entity, then the path with the maximum log-likelihood is chosen over the others. During evaluation in the inductive setting, the auxiliary graph augments the training graph to construct

Query	(William Green, worksfor, ?)
Answer	Accenture
Explanation	William Green $\xrightarrow{\text{personLeadsOrganization}}$ Accenture
Query	(Florida State, organizationhiredperson, ?)
Answer	Bobby Bowden
Explanation	Florida State $\xleftarrow{\text{worksFor}}$ Bobby Bowden
Query	(Messi, athlethomestadium, ?)
Answer	Camp Nou
Explanation	Messi $\xrightarrow{\text{athletePlaysForTeam}}$ Barcelona $\xrightarrow{\text{teamHomeStadium}}$ Camp Nou
Query	(Adrian Griffin, athlethomestadium, ?)
Answer	United Center
Explanation	Adrian Griffin $\xrightarrow{\text{athletePlaysForTeam}}$ Knicks $\xleftarrow{\text{athletePlaysForTeam}}$ Eddy Curry $\xrightarrow{\text{athleteHomeStadium}}$ United Center
Query	(Bucks, teamplaysinleague, ?)
Answer	NBA
Explanation	Bucks $\xrightarrow{\text{organizationHiredPerson}}$ Scott Stiles $\xleftarrow{\text{organizationHiredPerson}}$ Chicago Bulls $\xrightarrow{\text{teamPlaysInLeague}}$ NBA

Table 4: Example queries from the NELL-995 test set with unseen source entities. The answers are supported by the explainable reasoning paths derived by our model.

the KG environment with unseen entities and their relations to the seen entities. For our model and the baselines, the embeddings of all unseen entities are initialized with Xavier normal initialization (Glorot and Bengio, 2010) at inference time.

Evaluation Metrics We adopt the ranking based metrics *Mean Reciprocal Rank* and *Hits@k* that are also used by prior work for evaluation. We follow the *filtered setting* proposed by Bordes et al. (2013) and adopted by other prior work. In the filtered setting, the scores for the false negative answer entities are masked to facilitate correct ranking of the target entity.

4.4 Results

We present the experimental results of our method and the baselines in Tables 2 and 3. We observe that our model significantly outperforms the baselines for the FB15k-237, WN18RR, and NELL-995 datasets. The substantially inferior results of ConvE across all datasets demonstrates its ineffectiveness for inductive representation learning and reasoning.

Although both R-GCN and SACN use the neighborhood information for learning representations, unlike our method, their neighborhood integration methods do not explicitly consider the query relations.

5 Analysis

In this section, we show our model’s ability to provide reasoning paths as supporting evidence for inference. Additionally, we analyze the effect of different relation types on the inference process.

5.1 Qualitative Analysis of Explainability

Since explainability is one of the key objective of our model, we provide examples of explainable reasoning paths for queries that involve previously unseen source entity at inference time. Table 4 shows examples of 1-hop, 2-hop, and 3-hop reasoning paths. These examples demonstrate our model’s effectiveness in learning inductive representations for the unseen entities, which helps to infer the reasoning paths.

5.2 Effect of Relation Types

Dataset	to-Many		to-1	
	%	MRR	%	MRR
FB15k-237-Inductive	77.4	31.6	22.6	75.5
WN18RR-Inductive	48.1	60.8	51.9	30.1
NELL-995-Inductive	7.6	41.4	92.4	78.5

Table 5: The MRR for the test triples in inductive setting with *to-Many* and *to-1* relation types. The % columns show the percentage of test triples for each relation type.

Following Bordes et al. (2013), we categorize the relations in the seen snapshot of the knowledge graph into Many-to-1 and 1-to-Many relations. The categorization is done based on the ratio of the cardinality of the target answer entities to the source entities. If the ratio is greater than 1.5, we categorize the relation as *to-Many*, otherwise it is categorized as *to-1*. We analysed the results of the test set for these two types of relations. We report the percentage of triples with these two types of relations and the corresponding MRR achieved by our model in Table 5. For FB15k-237 and NELL-995, our model performs better for *to-1* relations than *to-many* relations. On the contrary, we observe a reverse trend for the WN18RR dataset. Note however that *to-many* relations have alternative target entities. In the current evaluation protocol, our model is punished for predicting any alternative target entity other than the ground truth target.

6 Conclusion

The ever-expanding number of entities in knowledge graphs warrants the exploration of link prediction methods that can be applied to emerging entities without retraining the model. While prior approaches assume a static snapshot of the knowledge graph, we introduce a joint framework for inductive representation learning to predict missing links in a dynamic knowledge graph with many emerging entities. Additionally, our method provides explainable reasoning paths for the inferred links as support evidence. Through experiments we demonstrate that our model significantly outperforms the baselines across the new benchmark datasets introduced in this paper.

References

- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 2787–2795.
- Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alexander J. Smola, and Andrew McCallum. 2017. [Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning](#). *CoRR*, abs/1711.05851.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. [Convolutional 2d knowledge graph embeddings](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 1811–1818. AAAI Press.
- Matt Gardner, Partha Pratim Talukdar, Bryan Kisiel, and Tom M. Mitchell. 2013. [Improving learning and inference in a large knowledge-base using latent syntactic cues](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 833–838. ACL.
- Matt Gardner, Partha Pratim Talukdar, Jayant Krishnamurthy, and Tom M. Mitchell. 2014. [Incorporating vector space similarity in random walk inference over knowledge bases](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 397–406. ACL.
- Xavier Glorot and Yoshua Bengio. 2010. [Understanding the difficulty of training deep feedforward neural networks](#). In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy. PMLR.
- Kelvin Guu, John Miller, and Percy Liang. 2015. [Traversing knowledge graphs in vector space](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 318–327. The Association for Computational Linguistics.
- Takuo Hamaguchi, Hidekazu Oiwa, Masashi Shimbo, and Yuji Matsumoto. 2017. [Knowledge transfer for out-of-knowledge-base entities: A graph neural network approach](#). In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI’17*, pages 1802–1808. AAAI Press.
- William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. [Inductive representation learning on large graphs](#). In *NIPS*.
- Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. [Knowledge graph embedding via dynamic mapping matrix](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 687–696. Association for Computational Linguistics.

- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). Cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. [Text Generation from Knowledge Graphs with Graph Transformers](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2284–2293, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ni Lao, Tom M. Mitchell, and William W. Cohen. 2011. [Random walk inference and learning in A large scale knowledge base](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP 2011, 27-31 July 2011, John McIntyre Conference Centre, Edinburgh, UK, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 529–539. ACL.
- Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2018. [Multi-hop knowledge graph reasoning with reward shaping](#). *CoRR*, abs/1808.10568.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. [Learning entity and relation embeddings for knowledge graph completion](#).
- Arvind Neelakantan, Benjamin Roth, and Andrew McCallum. 2015. [Compositional vector space models for knowledge base completion](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 156–166. The Association for Computer Linguistics.
- Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. 2018. [A Novel Embedding Model for Knowledge Base Completion Based on Convolutional Neural Network](#). In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 327–333.
- Dat Quoc Nguyen. 2017. [An overview of embedding models of entities and relationships for knowledge base completion](#). *CoRR*, abs/1703.08098.
- Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. [Modeling relational data with graph convolutional networks](#). In *The Semantic Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings*, volume 10843 of *Lecture Notes in Computer Science*, pages 593–607. Springer.
- Chao Shang, Yun Tang, Jing Huang, Jinbo Bi, Xiaodong He, and Bowen Zhou. 2019. [End-to-end structure-aware convolutional networks for knowledge base completion](#). In *AAAI*.
- Kristina Toutanova, Victoria Lin, Wen-tau Yih, Hoi-fung Poon, and Chris Quirk. 2016. [Compositional learning of embeddings for relation paths in knowledge base and text](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2071–2080. JMLR.org.
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph Attention Networks](#). *International Conference on Learning Representations*. Accepted as poster.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. [Knowledge graph embedding by translating on hyperplanes](#). In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, AAAI 14*, pages 1112–1119. AAAI Press.
- Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Mach. Learn.*, 8(3-4):229–256.
- Wenhan Xiong, Thien Hoang, and William Yang Wang. 2017. [DeepPath: A reinforcement learning method for knowledge graph reasoning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 564–573. Association for Computational Linguistics.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. [Embedding entities and relations for learning and inference in knowledge bases](#). *CoRR*, abs/1412.6575.

Po-Sen Huang Yuqing Gao Yelong Shen, Jianshu Chen and Gianfing Gao. 2018. M-walk: Learning to walk over graphs using monte carlo tree search.

Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. 2019. [Graph trans-](#)

[former networks](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 32, pages 11983–11993. Curran Associates, Inc.