

Inductive Learning on Commonsense Knowledge Graph Completion

Bin Wang

Electrical and Computer Engineering
University of Southern California
Los Angeles, CA, USA
bwang28c@gmail.com

Guangtao Wang

AI Research
JD.com
Mountain View, CA, USA
guangtao.wang@jd.com

Jing Huang

AI Research
JD.com
Mountain View, CA, USA
jing.huang@jd.com

Jiaxuan You

Computer Science
Stanford University
Stanford, CA, USA
jiaxuan@stanford.edu

Jure Leskovec

Computer Science
Stanford University
Stanford, CA, USA
jure@stanford.edu

C.-C. Jay Kuo

Electrical and Computer Engineering
University of Southern California
Los Angeles, CA, USA
cckuo@sipi.usc.edu

Abstract—Commonsense knowledge graph (CKG) is a special type of knowledge graph (KG), where entities are composed of free-form text. Existing CKG completion methods focus on transductive learning setting, where all the entities are present during training. Here, we propose the first inductive learning setting for CKG completion, where unseen entities may appear at test time. We emphasize that the inductive learning setting is crucial for CKGs, because unseen entities are frequently introduced due to the fact that CKGs are dynamic and highly sparse. We propose InductiveE as the first framework targeted at the inductive CKG completion task. InductiveE first ensures the inductive learning capability by directly computing entity embeddings from raw entity attributes. Second, a graph neural network with novel densification process is proposed to further enhance unseen entity representation with neighboring structural information. Experimental results show that InductiveE performs especially well on inductive scenarios where it achieves above 48% improvement over previous methods while also outperforms state-of-the-art baselines in transductive settings.

Index Terms—Commonsense Knowledge Graph, Inductive Learning, Graph Learning, Knowledge Graph Completion

I. INTRODUCTION

Knowledge graphs (KGs) are represented as triplets where entities (nodes) are connected by relationships (edges). It is a structured knowledge base with various applications such as recommendation systems, question answering, and natural language understanding [1]–[3]. In practice, most KGs are far from complete. Therefore, predicting missing facts is one of the most fundamental problems in this field. A lot of embedding-based methods have been proposed and shown to be effective on the KG completion task [4]–[10]. However, relatively little work targets at commonsense knowledge graph (CKG) completion. There are unique challenges encountered by applying existing KG embedding methods to CKGs (e.g. ConceptNet [11] and ATOMIC [12]).

First, many real-world CKGs are dynamic in nature, and entities with unseen text/names are introduced from time to

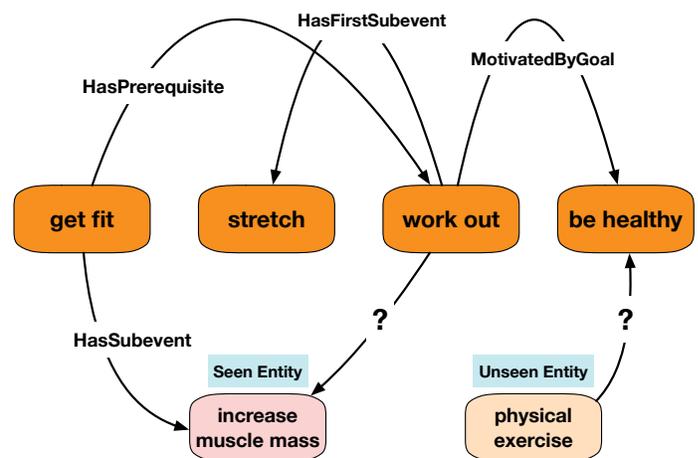


Fig. 1: Illustration of a fragment of ConceptNet commonsense knowledge graph: Link prediction for seen entities (transductive) and unseen entities (inductive).

time. We call these entities unseen entities because they are not involved in training and only appear in testing.

Second, entity attributes of CKGs are composed of free-form texts which are not present in non-attributed KG datasets. As shown in Figure 1, entity description has rich semantic meaning and commonsense knowledge can be largely inferred from their implicit semantic relations. However, we notice that often entities refer to the same concept are stored as distinct ones, resulting in the graph to be larger and sparser. As shown in Table 1, the average in-degree of ConceptNet and ATOMIC is only 1/15 and 1/8 comparing with that of FB15K-237, a popular KG dataset. Since CKGs are highly sparse and can be disconnected, a portion of entities are isolated from the main graph structure. These entities are also unseen entities and how to obtain embeddings for these isolated entities remains

challenging.

Therefore, the inductive learning problem on commonsense knowledge graph completion is particularly important with practical necessities. An example for CKG completion is shown in Figure 1. Transductive setting targets on predicting missing links for seen entities. The predictions can be made from two perspectives: 1) entity attributes and 2) existing links for seen entities. In contrast, inductive setting works on purely unseen entities where only entity attributes can be leveraged in the first place. Different from previous inductive setting on non-attributed KGs [13], [14], inductive learning on CKG assumes unseen entities are purely isolated and does not have any existing links. Therefore, this problem is unique and remains unexplored.

Many existing KG embedding models are focusing on non-attributed graphs with entity embeddings obtained during training [4], [7], [10] or using attribute embedding solely for initialization [15]. Therefore, all entities are required to present in training. Otherwise, the system will not have their embeddings and hence transductive as originally proposed.

In this work, we first propose and define the inductive learning problem on CKG completion. Then, an inductive learning framework called InductiveE is introduced to address the above-mentioned challenges and several components are specially designed to enhance its inductive learning capabilities. First, the inductive capability of InductiveE is guaranteed by directly building representations from entity descriptions, not merely using entity textual representation as training initialization. Second, a novel graph encoder with densification is proposed to leverage structural information for unseen entities and entities with limited neighboring connections. Overall, InductiveE follows an encoder-decoder framework to learn from both semantic representations and updated graph structures.

The main contributions can be summarized as follows.¹

1. A formal definition of inductive learning on CKGs is presented. We propose the first benchmark for inductive CKG completion task, including new data splits and testing schema, to facilitate future research.
2. InductiveE is the first model that is dedicated to inductive learning on commonsense knowledge graphs. It leverages entity attributes based on transfer learning from word embedding and graph structures based on a novel graph neural network with densification.
3. Comprehensive experiments are conducted on ConceptNet and ATOMIC datasets. The improvements are demonstrated in both transductive and inductive settings. InductiveE performs especially well on inductive learning scenarios with an over 48% improvement on MRR comparing with previous methods.

II. PROBLEM DEFINITION

Definition 1: Commonsense Knowledge Graph (CKG) is represented by $G = (V, E, R)$, where V is the set of nodes/entities, E is the set of edges and R is the set of

relations. Edges consist of triplets (h, r, t) where head entity h and tail entity t are connected by relation r : $E = \{(h, r, t) | h \in V, t \in V, r \in R\}$, and each node comes with a free-text description.

Definition 2: CKG Completion Given a commonsense knowledge graph $G = (V, E, R)$, CKG completion is defined as the task of predicting missing triplets $E' = \{(h, r, t) | (h, r, t) \notin E\}$. It includes both transductive and inductive settings. Transductive CKG completion is defined as predicting missing triplets $E'' = \{(h, r, t) | (h, r, t) \notin E, h \in V, t \in V, r \in R\}$.

Definition 3: Inductive CKG completion is defined as predicting missing triplets $E''' = \{(h, r, t) | (h, r, t) \notin E, h \in V' \text{ or } t \in V', r \in R\}$, where $V' \cap V = \emptyset$ and $V' \neq \emptyset$.

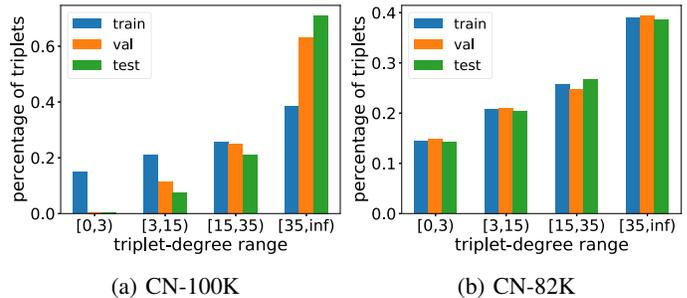


Fig. 2: The triplet-degree distribution for the ConceptNet dataset, where the triplet-degree is the average of head and tail degrees. Triplets with high degrees are easier to predict in general. CN-100K split is clearly unbalanced (comparing with CN-82K).

III. DATASET PREPARATION

Three CKG datasets are used to evaluate the link prediction (i.e., CKG completion) task. Their statistics are shown in Table I. CN-82K is our newly proposed split for better evaluation on link prediction task. Besides the standard split, we create an inductive split for CN-82K and ATOMIC called CN-82K-Ind and ATOMIC-Ind to specifically evaluate model’s generalizability for unseen entities.

A. Standard Split: CN-100K, CN-82K, ATOMIC

CN-100K was first introduced by [16]. It contains Open Mind Common Sense (OMCS) entries in the ConceptNet 5 dataset [11]. “100K” indicates the number of samples in the training data. In the ConceptNet 5 dataset, each triplet is associated with a confidence score that indicates the degree of trust. In the original split of CN-100K, the most confident 1,200 triplets are selected for testing and the next 1,200 most confident triplets are used for validation. Entities have a text description with an average of 2.9 words.

CN-100K was originally proposed to separate true and false triplets. It is not ideal for link prediction for two reasons. First, its data split ratio is biased. For 100,000 training samples, CN-100K contains only 2,400 triplets (2.4%) for validation (1.2%) and testing (1.2%). With such limited testing and

¹Our code is now publicly available: here.

TABLE I: Statistics of CKG datasets. Unseen Entity % indicates the percentages of unseen entities in all test entities.

Dataset	Entities	Relations	Train Edges	Valid Edges	Test Edges	Avg. In-Degree	Unseen Entity %
CN-100K	78,334	34	100,000	1,200	1,200	1.31	6.7%
CN-82K	78,334	34	81,920	10,240	10,240	1.31	52.3%
ATOMIC	304,388	9	610,536	87,700	87,701	2.58	37.6%

* For comparison, a popular KG dataset: FB-15K-237 has 14,541 entities, 310,115 edges and 18.76 Avg. In-Degree.

* In transductive settings, all entities are assumed visible during training including unseen entities. In contrast, unseen entities can only be used during testing stage for inductive settings.

validation sets, we see a large variance in evaluation. Second, the testing and the validation sets are selected as the most confident samples which are the least challenging ones. As shown in Figure 2(a), we see an unbalanced distribution in triplet-degree among train, validation and test sets. In order to test the performance of the link prediction task, we create and release a new data split called CN-82K.

CN-82K is a uniformly sampled version of the CN-100K dataset. 80% of the triplets are used for the training, 10% for validation, and the remaining 10% for testing. As shown in Figure 2(b), CN-82K is more balanced among train, validation and test sets w.r.t the triplet-degree distribution.

ATOMIC contains everyday commonsense knowledge entities organized as *if-then* relations [12]. It contains over 300K entities in total and entities are composed of text descriptions with an average of 4.4 words. Here, we use the same data split as done in [15]. Different from the ATOMIC split in [12], [17], the data split is a random split with ratio 80%/10%/10% for train/valid/test. Therefore, there is still a reasonable percentage of entities are unseen during training.

B. Inductive Split: CN-82K-Ind, ATOMIC-Ind

To evaluate a model’s generalizability to unseen entities, we create new validation and test sets for CN-82K and ATOMIC, while keeping the training set unchanged. The inductive validation and testing sets are subsets of the original validation and test sets, which contain only triplets with at least one unseen entities. After filtering, for CN-82K-Ind, the validation and test sets contain 5,715 and 5,655 triplets, respectively. For ATOMIC-Ind, the validation and test sets contain 24,355 and 24,486 triplets, respectively. We show by experiments that previous models fail in such scenarios (please refer to Table III) and propose a new benchmark called InductivE for inductive CKG completion.

IV. INDUCTIVE

In this section, we introduce InductivE, the first benchmark of inductive learning on CKGs, which includes a free-text encoder, a graph encoder with densification and a convolutional decoder. The overall architecture is illustrated in Figure 3.

A. Free-Text Encoder

Our free-text encoder embeds text attributes with a pre-trained language model and word-embedding. For pre-trained language model, BERT [18] is applied and further fine-tuned on entity textual attributes, which allows BERT to better

model our domain-specific data. For word embedding, fastText model is used and mean-pooling is applied when the text sequence contains more than one word [19], [20]. Finally, two representations are concatenated together as the final entity representation. For inductive purpose, the free-text encoder is viewed as a feature extractor and entity embeddings are fixed during training.

BERT shows superior performance in many NLP tasks. However, for CKG datasets, fastText embedding exhibits comparable performance. We infer that fastText can be close to BERT feature when handling short text sequences. More discussion is provided in Sec. V-C.

B. Graph Encoder

In addition to semantic representation, it is also desired to investigate the possibility to enhance inductive ability with entity neighboring structures. However, it is challenging to leverage structural information for unseen entities because, as isolated ones, unseen entities do not contain any existing links at the beginning. Therefore, we first densify the graph by add similarity links using a novel graph densifier. Then, a gated-relational graph convolutional network is applied to learn from the graph structural information.

1) *Graph densifier*: We observe that some entities often share the same ontological concept in CKGs. For example, “work out” and “physical exercise” in Figure 1 are similar concepts but presented as two distinct entities. Similarity edges are added to densify the graph to provide more interactions between similar concepts.

Here, we propose a novel *graph densifier* to generate high-quality links among semantic-similar entities. We first compute the pairwise cosine similarity using the output entity feature of our graph encoder rather than original entity features. Then, for each node i , we identify k_i nearest neighbors and add directed similarity edges from them to node i to densify the graph. To balance the resulting node degree across the graph, the number of added synthetic edges is node-dependent. More synthetic edges will be added to nodes with fewer connections. Specifically, the number of added edges k_i for each node i is determined by

$$k_i = \begin{cases} 0, & \text{if } m \leq \text{degree}(i), \\ m - \text{degree}(i) & \text{otherwise,} \end{cases} \quad (1)$$

where m is a hyperparameter to determine the number of similarity edges added for each node, and $\text{degree}(i)$ represents the

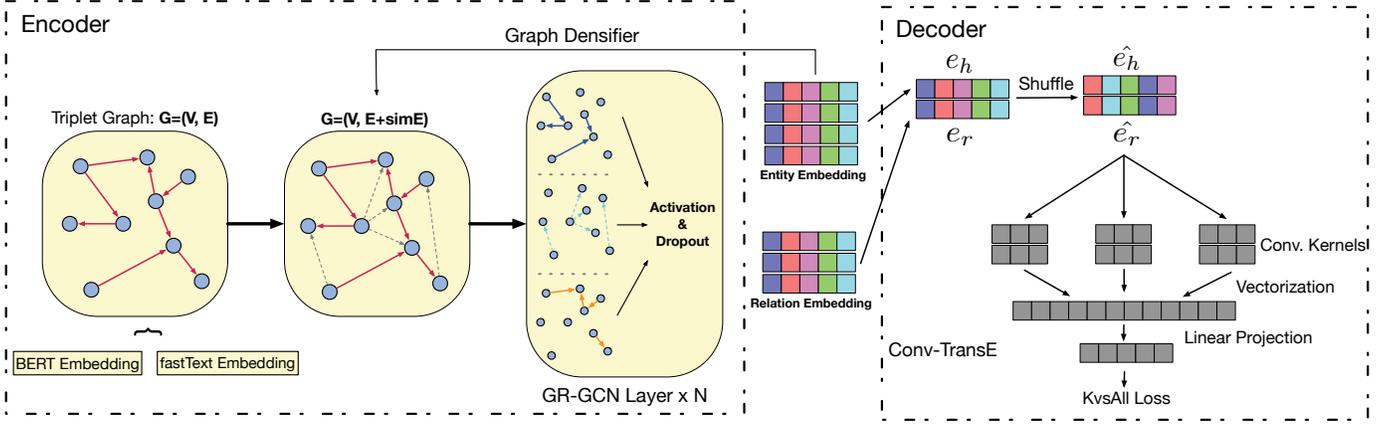


Fig. 3: Illustration of the proposed InductivE model. The encoder contains multiple GR-GCN layers applied to the graph to aggregate node’s local information. The decoder uses an enhanced Conv-TransE model.

degree of node i . The densified graph is updated periodically during training based on the above-mentioned scheme.

In the testing stage, to infer the embedding for unseen entities, we first get entity representations by applying our trained graph encoder on the CKG graph, in which the unseen entities are isolated; then similarity edges are added according to these representations to densify the graph; finally, the densified graph is encoded again using our trained graph encoder to get final entity embeddings, which serve as the input to the decoder.

The logic behind our densification design is as follows: first, we did not use raw feature with threshold because a lower threshold leads to noisy edges and a higher threshold provides little extra information as the feature already serves as the input; second, more edges are added to unseen entities to ensure enough structural information are incorporated for inductive purposes.

2) *Gated-Relational GCN*: R-GCN [21] is effective in learning the node representation over graph with relational information between nodes (e.g., Knowledge graph). Our model is an extension of the R-GCN [21] model. First, in CKGs, the neighboring conditions can vary a lot from node to node. It is desired to adaptively control the amount of information fused to the center node from their neighboring connections. Therefore, a gating function is added to R-GCN for this purpose based on the interaction of the center and neighboring nodes. Second, to increase the efficiency of R-GCN model, instead of using relation-type-specific transformation matrices, one unified transformation matrix W_1 is adopted for all neighboring relation types. As a result, our graph encoder is composed of multiple newly proposed gated-relational graph convolutional (GR-GCN) layers.

The first convolutional layer takes the output from free-text encoder as the input: $h_i^{(0)} = x_i$. At each layer, the update message is a weighted sum of a transformation of center node u_c and a transformation of its neighbors u_n in form of

$$h_i^{(l+1)} = \sigma(u_c^{(l)} \odot \beta_i^{(l)} + u_n^{(l)} \odot (1 - \beta_i^{(l)})), \quad (2)$$

where β_i denotes a gating function, \odot denotes an element-wise multiplication, $\sigma(\cdot)$ is a nonlinear activation function. The two transformations in Eq. (2) are defined as

$$u_c^{(l)} = W_o^{(l)} h_i^{(l)}, \quad (3)$$

$$u_n^{(l)} = \sum_{r \in R} \sum_{j \in N_i^r} \frac{1}{|N_i^r|} \alpha_r^{(l)} W_1^{(l)} h_j^{(l)}, \quad (4)$$

where N_i^r denotes neighbors of node i with relation type r , R denotes all relation types, $\alpha_r^{(l)}$ is the relation weight at layer l . For all neighboring nodes, we use one unified transformation $W_1^{(l)}$, which differs from the transformation of the self-loop message denoted by $W_o^{(l)}$, to account for the relation gap between neighborhood information and self-connection.

A gating mechanism controls the amount of the neighborhood message flowing into the center node. In this work, a learnable gating function is used, which takes both $u_c^{(l)}$ and $u_n^{(l)}$ as the input. It can be written as

$$\beta_i^{(l)} = \text{sigmoid}(f([u_c^{(l)}, u_n^{(l)}])), \quad (5)$$

where $[,]$ is the concatenation of self-loop and neighboring messages and f is a linear transformation. Finally, a sigmoid function is used to ensure $0 < \beta_i^{(l)} < 1$.

C. Decoder - Conv-TransE

Convolutional decoder is effective in scoring triplets in KGs with high parameter efficiency [6]. Conv-TransE [22] is a simplified version of ConvE [6]. It removes the reshaping process before the convolution operation and use 1D convolution operation instead of 2D. Yet, the 2D convolution increases the expressive power of the ConvE model since it allows more interactions between embeddings as discussed in [6]. InteractE was proposed recently to allow even more interactions [23].

As shown in Figure 3, we deploy an improved Conv-TransE model that adds a shuffling operation before convolution to enable more interactions across dimensions inspired by Inter-

actE [23]. Formally, let ϕ_s represent the horizontal shuffling process, the score function in our decoder is defined as:

$$\text{score}(e_h, e_r, e_t) = f(\text{vec}(f(\phi_s([e_h; e_r]) * w))W))e_t \quad (6)$$

where $e_h/e_t \in \mathbb{R}^d$ is the head/tail embedding and $e_r \in \mathbb{R}^d$ is the relation embedding. These entity embeddings come from the output of the graph encoder. The relation embedding is randomly initialized and jointly trained with the end-to-end framework. f denotes a non-linear activation function. In a feed-forward pass, e_h and e_r are first stacked as a $\{2 \times d\}$ matrix and shuffled horizontally. It is used as the input to a 1D convolutional layer with filters w . Each filter is with 1D convolutional kernel with size $\{2 \times n\}$. The output is further reshaped as a vector and projected back to the original d dimensions using a linear transformation. Finally, an inner product with the tail embedding e_t is performed in the d -dimensional space as the final triplet score.

To train our model parameter, we use KvsAll training schema [6] by considering all entities simultaneously. Instead of scoring each triplet (h, r, t) , we take each (h, r) pair and score it with all entities (positive or negative) as tail. Some pairs could have more than one positive tail entities. Thus, it is a multi-label problem, for which we adopt the binary cross-entropy loss:

$$L(p, t) = -\frac{1}{N} \sum_i (t_i \cdot \log(p_i) + (1 - t_i) \cdot \log(1 - p_i)), \quad (7)$$

where t is the true label vector with dimension \mathbb{R}^N . We get the probability of (h, r, t) being positive as $p_i = \text{sigmoid}(\text{score}(e_h, e_r, e_i))$.

V. EXPERIMENTS

A. Experimental Setup

1) *Evaluation protocol*: We use link prediction task with standard evaluation metrics including Mean Reciprocal Rank (MRR) and Hits@10, to evaluate CKG completion models. We report the average results in percentage with five runs. Following [6], [7], each triplet (h, r, t) is measured in two directions: $(h, r, ?)$ and $(t, r^{-1}, ?)$. Inverse relations r^{-1} are added as new relation types and the filtered setting is used to filter out all valid triplets before ranking.

2) *Hyper-parameter settings*: Our graph encoder consists of two GR-GCN layers with hidden dimension 500. The hyperparameter m used in graph densifier is set to 5 for ConceptNet and 3 for ATOMIC. The graph is updated every 100 epochs for ConceptNet and 500 epochs for ATOMIC. For convolutional decoder, we use 300 kernels of size $\{2 \times 5\}$. For model training, we set the initial learning rate to $3e-4$ and $1e-4$ for ConceptNet and ATOMIC, respectively, and halve the learning rate when validation metric stops increasing for three times. The checkpoint with the highest Mean Reciprocal Rank (MRR) on validation set is used for testing.²

²More details are presented in the our released project page.

Training GCN on large graphs demands large memory space to keep the entire graph parameters. Here, for efficiency purposes, we perform uniform random sampling on nodes at each training epoch with sampling size 50k in all experiments.

3) *Baselines*: We compare with several representative models, including DistMult [24], ComplEx [5], RotatE [7], ConvE [6], Malaviya [15] and COMET [17]. The first four models are competitive KG embedding models without using entity textual attributes. The last two models are focusing on CKGs and also take entity textual attributes into consideration. For non-inductive models, we allow the presence of unseen entities during training in order to perform evaluation. We use their respective official implementations to obtain the baseline results and tune several hyperparameters including batch size, learning rate and embedding dimensions. Finally, our obtained baseline results are compared with the best results reported in existing literature and a higher one is reported.

B. Result and Analysis

1) *Transductive link prediction*: Table II summarizes results on CN-100K, CN-82K and ATOMIC. For CN-100K, our model outperforms previous state-of-the-art by 9.8% on MRR and 6.1% on Hits@10. In contrast, the performance for CN-82K is much lower since CN-82K has more unseen entities in the testing as shown in Table I. Over 50% of all entities in testing are unseen. This is very challenging for all existing methods. InductivE can learn high-quality embedding for all entities and outperforms the previous best model by over 20% across all evaluation metrics on CN-82K. For ConceptNet datasets, without using the textual information, ConvE and RotatE perform better than ComplEx and DistMult. [15] outperforms ConvE by a large margin with BERT features as initialization. This indicates that semantic information plays an important role for ConceptNet entities.

For ATOMIC, ComplEx and InductivE provide the best performance among all methods. Text attribute feature is less effective than that for the ConceptNet dataset. We observe that the relation types from ATOMIC (e.g. xAttr, xIntent, oReact) are more complex and require more high-level reasoning. Thus, it is more difficult to infer directly from semantic embeddings. As compared with ComplEx, InductivE is good at Hits@10 score, which means more entities are ranked higher. This is attributed to synthetic edges that help entities with limited connections to get reasonable performance as more synthetic connections are added.

2) *Inductive link prediction*: Table III summaries results on CN-82K-Ind and ATOMIC-Ind. From this table, we observe that:

- The conventional KG embedding models, ConvE and RotatE, perform badly in our proposed inductive settings. ConvE and RotatE learn entity embeddings via learning to score positive/negative links between entities. After training, unseen entities remain to be randomly initialized, because no links over unseen entities can be observed during training. These random embeddings for

TABLE II: Comparison of CKG completion results on CN-100K, CN-82K and ATOMIC datasets. Improvement is computed by comparing with [15].

Model	CN-100K			CN-82K			ATOMIC		
	MRR	Hits@3	Hits@10	MRR	Hits@3	Hits@10	MRR	Hits@3	Hits@10
DistMult	10.62	10.94	22.54	2.80	2.90	5.60	12.39	15.18	18.30
ComplEx	11.52	12.40	20.31	2.60	2.70	5.00	14.24	14.13	15.96
ConvE	20.88	22.91	34.02	8.01	8.67	13.13	10.07	10.29	13.37
RotatE	24.72	28.20	45.41	5.71	6.00	11.02	11.16	11.54	15.60
COMET	6.07	2.92	21.17	-	-	-	4.91	2.40	21.60
Malaviya et al.	52.25	58.46	73.50	16.26	17.95	27.51	13.88	14.44	18.38
InductiveE	57.35	64.50	78.00	20.35	22.65	33.86	14.21	14.82	20.57
Improvement	9.8%	10.3%	6.1%	25.2%	26.2%	23.1%	2.38%	2.63%	11.92%

TABLE III: Comparison of CKG completion results on unseen entities for CN-82K-Ind and ATOMIC-Ind.

Model	CN-82K-Ind		ATOMIC-Ind	
	MRR	Hits@10	MRR	Hits@10
ConvE	0.21	0.40	0.08	0.09
RotatE	0.32	0.50	0.10	0.12
Malaviya et al.	12.29	19.36	0.02	0.07
InductiveE	18.15	29.37	2.51	5.45

unseen entities explain the poor performance of ConvE and RotatE.

- Comparing with [15], InductiveE achieves the best performance on both inductive datasets. It has an improvement of 48% on MRR when comparing with [15] on CN-82K-Ind dataset. As shown in the transductive result (Table II), text features for ATOMIC dataset are less effective than that for ConceptNet dataset, thus results in much worse performance on ATOMIC-Ind than on CN-82K-Ind.

In [15] the model training adopts a two-branch training process: in the GCN branch, node features are randomly initialized; in the text encoder branch, entities embeddings are *initialized* by BERT features then finetuned during training. For both branches, training *does not* change the embeddings for unseen entities, because no links over unseen entities can be observed during training. Therefore, at test time, embeddings for unseen entities are computed from random node features (in the GCN branch) plus raw BERT features (in the text encoder branch), which are different from those of trained seen entities, thus results in poor performance.

In contrast, InductiveE first ensures inductive learning ability by using fixed text embedding as the input to GR-GCN model and further enhances the neighboring structure for unseen entities via the proposed graph densifier. Both modules contribute to the good performance of InductiveE.

To conclude, InductiveE is the first benchmark in the inductive setting while providing competitive performance in the transductive setting.

C. Ablation Study and Analysis

1) *Ablation study*: To better understand the contribution of different modules in InductiveE to the performance, we present

TABLE IV: Ablation study on CN-82K-Ind dataset.

Model	CN-82K-Ind	
	MRR	Hits@10
InductiveE	18.15	29.37
replace GR-GCN with MLP	-1.70	-2.60
remove graph densifier	-3.03	-4.60
remove gating in GR-GCN	-3.79	-5.37
BERT feature only	16.57	27.68
fastText feature only	16.14	26.87

ablation studies in Table IV.

- **Feature Analysis**: The performance gap between BERT feature and fastText feature is not huge. It is because the text attributes for CKG datasets are short phrases, with an average of 2.9 words for ConceptNet and 4.4 words for ATOMIC. BERT is more powerful for sentence-level sequences and that is why we use a concatenation of features to have multi-view entity representations.
- **Replace GR-GCN with MLP**: The MRR score drops from 18.15 to 16.45. This indicates that learning from the graph structural information can further boost the performance, although our free-text encoder provides a strong baseline and most essential for inductive learning.
- **Remove graph densifier**: This means the model only learns from the original triplet graph structure. The MRR score drops from 18.15 to 15.12. This indicates our graph densifier is helpful for inductive learning, since that the added synthetic edges are particularly helpful for unseen entities since they do not have existing neighboring structure to learn from at the beginning.
- **Remove gating in GR-GCN**: The performance drops from 18.15 to 14.36. This demonstrates the importance of the gating function, which adaptively controls the amount of neighboring information flowing into the center node. Neighboring nodes in CKGs are diverse and connected to the center node with different relations. Without the gating function, different information sources are injected directly to the center node which can cause confusion to the model.

2) *Analysis on graph densifier*: Our iterative graph densifier is a special design to provide more structural information for unseen entities which does not contain any existing links at the beginning. In the meantime, there are other alternatives [15] to build syntactic links such as directly building syntactic links from raw features. Even though constructing synthetic links directly from raw features is easier to implement, it has two disadvantages when targeting on our inductive learning framework: 1) A low threshold leads to noisy synthetic links and a high threshold provide little extra information as the feature has served as the input. 2) Synthetic links can be unbalanced between entities and some unseen entities will still have no connection even after densification.

To verify our graph densifier is superior than its alternatives, we compare it with two alternatives:

- Raw feature with global thresholding (GS) [15]: The BERT feature is taken and compute pairwise cosine similarity between entities. A global threshold (0.95) is set for the whole graph and similarity links are added between entities if their cosine similarity is above the threshold.
- Raw feature with fixed neighboring (FN): To make sure that all entities get a reasonable amount of neighboring information, we add similarity links to entities until they have at least 5 neighboring entities. The candidates are selected by ranking cosine similarity of the BERT feature.

TABLE V: Analysis on our graph densifier

Model	CN-82K-Ind	
	MRR	Hits@10
Our graph densifier	18.15	29.37
Raw feature with GS	13.25	22.98
Raw feature with FN	10.12	17.04

The results are shown in Table V. Our graph densifier achieves the best performance. In comparison, *Raw feature with GS* is not providing satisfactory result because the raw feature can be noisy and imprecise links can be added. Even though some high-quality similarity links are constructed, it does not provide much extra information. *Raw feature with FN* is the worst choice because adding fixed number of links directly from raw feature will allow more noisy links injected to the graph structure which leads more confusion in general. Since our graph densifier incorporates an iterative approach, the features we use are generally more fine-tuned with current datasets and keep improving during training and thus providing synthetic links with a higher quality.

3) *Case study on graph densifier*: By examining our graph densifier carefully, we notice some interesting phenomena associated with synthetic similarity links. In Table VI, we list top-3 nearest neighbors of unseen entities. As expected, some similarity relations can be discovered with our graph densifier. For example, “*pay for subway*” shares almost the same semantic meaning with “*pay subway fare*”. Therefore, when computing embedding for “*pay for subway*”, the embedding

TABLE VI: Top-3 nearest neighbors of unseen entities

Unseen entity	Ours	Malaviya et al.
pay for subway	pay subway fare pay traffic ticket pay fare	pay for drink pay on billet pay for package
perform experiment	try experiment do experiment test his hypothesis	do experiment conduct experiment run experiment
item fill with air	inflate thing with air balloon fly because they balloon go up because they	open container wled like item microscopic thing
wait for you airplane	get prepare to wait run short of fly transport passenger	wait for while wait for windy day wait for blue bird

for “*pay subway fare*” can be a reliable reference. Unexpectedly, other complex relation types can also be discovered using our densifier. Entities marked as bold in Table VI indicates the candidate entity has more complex relation with the unseen entity. For example, “*perform experiment*” is the “*Goal*” to “*test his hypothesis*”. The “*Reason*” for “*wait for you airplane*” is “*run short of fly*”. With rich local connectivity, unseen entities can perform multi-hop reasoning over graphs and obtain high quality embeddings. Comparing with previous method [15], our proposed graph densifier can construct higher quality synthetic edges. This could be especially helpful for unseen entities whose neighboring structures are built using our graph densifier.

VI. RELATED WORK

1) *Knowledge Graph Embedding*: Knowledge graph completion by predicting missing links has been intensively investigated in recent years. Most methods are embedding-based. TransE [4] models the relationship between entities with nice translation property ($e_h + e_r \approx e_t$) in the embedding space. ComplEx [5] and RotatE [7] represent embeddings in a complex space to model more complicated relation interactions. Instead of using a simple score function, ConvE [6] applied convolution to embedding so as to allow more interactions among triplet features. To exploit the structural information, GNNs are applied to multi-relational graphs as done in R-GCN [21] and SCAN [22]. All above-mentioned methods learn embedding based on a fixed set of entities and are transductive as originally proposed.

2) *Inductive Learning on Graphs*: Inductive learning is investigated in the last several years for both graphs and knowledge graphs. GraphSage [25] relies on node features and learns a local aggregation function that is generalizable to newly observed subgraphs. Most KGs embedding models are focusing non-attributed graph so that inductive learning is mostly relying on local connections for unseen entities [14]. [26] generates embedding for unseen nodes by aggregating the information from surrounding known nodes. [13] models relation prediction as a subgraph reasoning problem. However, because unseen entities in CKGs have no existing link, structural-based methods cannot be applied [13], [14]. As an alternative, there exists work that incorporates entity description in the embedding process and can be inductive by nature. For example, [27] learns a joint embedding space for

conventional entity embedding and description-based embedding. Our work exploits both structure information and textual description for inductive purpose on CKGs.

3) *Language model on CKGs*: Recently, researchers attempt to link commonsense knowledge with pre-trained language models [18], [28], [29]. COMET [17] is a generative model that transfers knowledge from pre-trained language models and generates new facts in CKGs. It could achieve performance close to human beings. However, COMET always introduces novel/unseen entities to an existing graph, leading to an even sparser graph. With unseen entities constantly introduced by generative models, an inductive learning method, such as InductivE, is important in practice.

VII. CONCLUSION

In this work, we propose to study the inductive learning problem on CKG completion, where unseen entities are involved in link prediction. To better evaluate CKG completion task in both transductive and inductive settings, we release one new ConceptNet dataset and two inductive data splits for future research and development purposes. Dedicate to this task, a new embedding-based framework InductivE is proposed as the first benchmarking on inductive CKG completion. InductivE leverages entity attributes with transfer learning and considers structural information with GNNs. Experiments on both transductive and inductive settings show that InductivE outperformed the state-of-the-art method considerably.

Inductive learning on CKG completion is still at its infancy. There are many promising directions for future work. For example, large pre-trained language models (LMs) have shown effective in capturing implicit commonsense knowledge from large corpus. How to effectively merge the knowledge in large pre-trained LMs with structured CKG could be an interesting direction to explore. Another direction is to explore inductive learning on unseen/new relations that are truly useful in real CKG expansion task.

REFERENCES

- [1] H. Wang, F. Zhang, X. Xie, and M. Guo, "Dkn: Deep knowledge-aware network for news recommendation," in *Proceedings of the 2018 world wide web conference*, 2018, pp. 1835–1844.
- [2] W. Cui, Y. Xiao, H. Wang, Y. Song, S.-w. Hwang, and W. Wang, "Kbqa: learning question answering over qa corpora and knowledge bases," *arXiv preprint arXiv:1903.02419*, 2019.
- [3] W. Liu, P. Zhou, Z. Zhao, Z. Wang, Q. Ju, H. Deng, and P. Wang, "K-bert: Enabling language representation with knowledge graph," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 2901–2908.
- [4] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," in *Advances in neural information processing systems*, 2013, pp. 2787–2795.
- [5] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard, "Complex embeddings for simple link prediction," in *International Conference on Machine Learning (ICML)*, 2016.
- [6] T. Dettmers, P. Minervini, P. Stenetorp, and S. Riedel, "Convolutional 2d knowledge graph embeddings," in *AAAI*, 2018, pp. 1811–1818.
- [7] Z. Sun, Z.-H. Deng, J.-Y. Nie, and J. Tang, "Rotate: Knowledge graph embedding by relational rotation in complex space," *International Conference on Learning Representations (ICLR)*, 2019.
- [8] I. Balazevic, C. Allen, and T. Hospedales, "TuckER: Tensor factorization for knowledge graph completion," in *EMNLP*, 2019.
- [9] T. Bansal, D.-C. Juan, S. Ravi, and A. McCallum, "A2N: Attending to neighbors for knowledge graph inference," in *ACL*, 2019.
- [10] Y. Tang, J. Huang, G. Wang, X. He, and B. Zhou, "Orthogonal relation transforms with graph context modeling for knowledge graph embedding," *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [11] R. Speer and C. Havasi, "Conceptnet 5: A large semantic network for relational knowledge," in *The People's Web Meets NLP*. Springer, 2013, pp. 161–176.
- [12] M. Sap, R. Le Bras, E. Allaway, C. Bhagavatula, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, and Y. Choi, "ATOMIC: An atlas of machine commonsense for if-then reasoning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 3027–3035.
- [13] K. K. Teru, E. Denis, and W. L. Hamilton, "Inductive relation prediction by subgraph reasoning," *International Conference on Machine Learning*, 2020.
- [14] M. Albooyeh, R. Goel, and S. M. Kazemi, "Out-of-sample representation learning for knowledge graphs," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, 2020, pp. 2657–2666.
- [15] C. Malaviya, C. Bhagavatula, A. Bosselut, and Y. Choi, "Commonsense knowledge base completion with structural and semantic context," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 2925–2933.
- [16] X. Li, A. Taheri, L. Tu, and K. Gimpel, "Commonsense knowledge base completion," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 2016, pp. 1445–1455.
- [17] A. Bosselut, H. Rashkin, M. Sap, C. Malaviya, A. Çelikyilmaz, and Y. Choi, "Comet: Commonsense transformers for automatic knowledge graph construction," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [18] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *NAACL*, 2019, pp. 4171–4186.
- [19] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching word vectors with subword information," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, 2017.
- [20] B. Wang, A. Wang, F. Chen, Y. Wang, and C.-C. J. Kuo, "Evaluating word embedding models: Methods and experimental results," *APSIPA transactions on signal and information processing*, vol. 8, 2019.
- [21] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. Van Den Berg, I. Titov, and M. Welling, "Modeling relational data with graph convolutional networks," in *European Semantic Web Conference*. Springer, 2018, pp. 593–607.
- [22] C. Shang, Y. Tang, J. Huang, J. Bi, X. He, and B. Zhou, "End-to-end structure-aware convolutional networks for knowledge base completion," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 3060–3067.
- [23] S. Vashishth, S. Sanyal, V. Nitin, N. Agrawal, and P. P. Talukdar, "Interact: Improving convolution-based knowledge graph embeddings by increasing feature interactions," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 3009–3016.
- [24] B. Yang, W.-t. Yih, X. He, J. Gao, and L. Deng, "Embedding entities and relations for learning and inference in knowledge bases," *International Conference on Learning Representations (ICLR)*, 2014.
- [25] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Advances in neural information processing systems*, 2017, pp. 1024–1034.
- [26] T. Hamaguchi, H. Oiwa, M. Shimbo, and Y. Matsumoto, "Knowledge transfer for out-of-knowledge-base entities: A graph neural network approach," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2017, pp. 1802–1808.
- [27] R. Xie, Z. Liu, J. Jia, H. Luan, and M. Sun, "Representation learning of knowledge graphs with entity descriptions," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [28] F. Petroni, T. Rocktäschel, P. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, and S. Riedel, "Language models as knowledge bases?" *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [29] L. Yao, C. Mao, and Y. Luo, "Kg-bert: Bert for knowledge graph completion," *arXiv preprint arXiv:1909.03193*, 2019.