

Lifelong Embedding Learning and Transfer for Growing Knowledge Graphs

Yuanning Cui¹, Yuxin Wang¹, Zequn Sun¹, Wenqiang Liu³, Yiqiao Jiang³, Kexin Han³, Wei Hu^{1,2,*}

¹ State Key Laboratory for Novel Software Technology, Nanjing University, China

² National Institute of Healthcare Data Science, Nanjing University, China

³ Interactive Entertainment Group, Tencent Inc, China

{yncui, zqsun, yxwangcs}.nju@gmail.com, {masonqliu, gennyjiang, casseyhan}@tencent.com, whu@nju.edu.cn

Abstract

Existing knowledge graph (KG) embedding models have primarily focused on static KGs. However, real-world KGs do not remain static, but rather evolve and grow in tandem with the development of KG applications. Consequently, new facts and previously unseen entities and relations continually emerge, necessitating an embedding model that can quickly learn and transfer new knowledge through growth. Motivated by this, we delve into an expanding field of KG embedding in this paper, i.e., lifelong KG embedding. We consider knowledge transfer and retention of the learning on growing snapshots of a KG without having to learn embeddings from scratch. The proposed model includes a masked KG autoencoder for embedding learning and update, with an embedding transfer strategy to inject the learned knowledge into the new entity and relation embeddings, and an embedding regularization method to avoid catastrophic forgetting. To investigate the impacts of different aspects of KG growth, we construct four datasets to evaluate the performance of lifelong KG embedding. Experimental results show that the proposed model outperforms the state-of-the-art inductive and lifelong embedding baselines.

1 Introduction

Many knowledge-driven applications are built on knowledge graphs (KGs), which store massive amounts of structured facts about the real world (Ji et al. 2022). Throughout the life-cycle of KG construction and application, new facts, unseen entities, and unseen relations continually emerge into the KG on a regular basis. As a result, the real-world KG is rarely a static graph, but rather evolves and grows alongside the development. Figure 1 illustrates an example excerpted from Wikidata (Vrandečić and Krötzsch 2014), which shows the growth of the KG along with the continuous knowledge extraction. However, KG embedding, a critical task for downstream applications, has primarily focused on static KGs over the years (Wang et al. 2017). Learning from scratch every time is inefficient and wastes previously acquired knowledge, and simply fine-tuning new facts would quickly disrupt previously acquired knowledge. Hence, this paper proposes to investigate lifelong embedding learning and transfer for growing KGs, with the goal of learning new facts while retaining old knowledge without re-training from scratch.

*Corresponding author

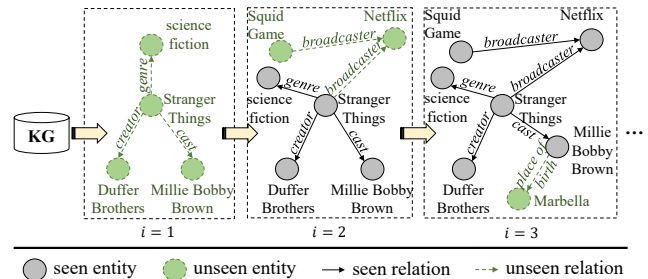


Figure 1: An example of a growing KG. In each snapshot i , new facts are added into the KG, and previously unseen entities and relations emerge with the new facts.

The key idea of this paper comes from the human learning process. Humans are typical lifelong learners, with knowledge transfer and retention being the most important aspects of lifelong learning. Humans, in particular, can continually learn new knowledge given new facts and use previously learned knowledge to help new knowledge learning (*knowledge learning and transfer*), as well as update old knowledge while retaining useful knowledge (*knowledge update and retention*). Motivated by this, we seek to build a lifelong KG embedding model, namely **LKGE**, which is capable of learning, transferring and retaining knowledge for growing KGs efficiently. Existing related work, such as inductive KG embedding (Hamaguchi et al. 2017; Wang et al. 2019b), mainly focuses on knowledge transfer, ignoring new knowledge learning and old knowledge update.

The proposed lifelong KG embedding task faces two major challenges. First, how to strike a balance between new knowledge learning and old knowledge transfer? Learning embeddings for new entities and relations from scratch cannot leverage previously learned knowledge, and inductively generating embeddings for them ignores new knowledge in the new snapshot. Second, how to update old knowledge while retaining useful knowledge? Learning new facts about an old entity usually requires updating the previously learned embeddings, which can be harmful to the old model. This is because updating an old entity embedding would affect many other old embeddings of related entities. This would cause the catastrophic forgetting issue, and therefore affect the applications built on the old KG snapshot.

To resolve the above challenges, we propose three solu-

tions in our LKGE model. First, as the base embedding model for new knowledge learning and old knowledge update, we design a masked KG autoencoder that masks and reconstructs the entities or relations in new facts. It builds connections between locally-related old and new entities, acting as a bridge for knowledge transfer. Second, to aid in the learning of new knowledge, we propose a knowledge embedding transfer strategy that uses previously learned knowledge to initialize the embeddings of new entities and relations. These embeddings are used by the KG autoencoder to learn new facts. Third, to avoid catastrophic forgetting in the old knowledge update, we propose embedding regularization to balance the learning of new facts and the update of old embeddings.

We build four datasets to assess the lifelong KG embedding performance, including the entity-centric, relation-centric, fact-centric, and hybrid growth. Each dataset examines a different aspect of KG growth. By contrast, existing datasets (Hamaguchi et al. 2017; Kou et al. 2020; Daruna et al. 2021) all assume that a KG grows in an ideal way, with balanced new entities or facts in each new snapshot. In our experiments, we compare the link prediction accuracy, knowledge transfer ability, and learning efficiency of the proposed model against baselines on the four datasets. The results show that the proposed LKGE not only achieves the best performance on the four datasets, but also has the best forward knowledge transfer ability and learning efficiency. The main contributions of this paper are summarized as follows:

- We study *lifelong KG embedding learning and transfer*. It is a practical task since the real-world KGs continually evolve and grow, which requires the embedding model to be capable of handling the knowledge growth.
- We propose *a novel lifelong learning model*, LKGE. It includes a masked KG autoencoder as the basis of embedding learning and update, an embedding transfer strategy for knowledge transfer, and an embedding regularization to prevent catastrophic forgetting in knowledge update.
- We conduct *extensive experiments on four new datasets*. The results demonstrate the effectiveness and efficiency of LKGE against a variety of state-of-the-art models.

2 Related Work

2.1 Knowledge Graph Embedding

KG embedding seeks to encode the symbolic representations of KGs into vector space to foster the application of KGs in downstream tasks. Most existing KG embedding models (Bordes et al. 2013; Wang et al. 2014; Dettmers et al. 2018; Schlichtkrull et al. 2018; Guo, Sun, and Hu 2019; Vashishth et al. 2020a,b) focus on static graphs and cannot continually learn new knowledge on the growing KGs.

To embed unseen entities, inductive KG embedding models learn to represent an entity by aggregating its existing neighbors in the previous KG snapshot. MEAN (Hamaguchi et al. 2017) uses a graph convolutional network (GCN) (Scarselli et al. 2009) for neighborhood aggregation. When an unseen entity emerges, the GCN would aggregate its previously seen neighboring entities to generate an embedding. LAN (Wang et al. 2019b) adopts an attention mechanism

to attentively aggregate different neighbors. As MEAN and LAN rely on the entity neighborhood for embedding learning, they cannot handle the new entities that have no neighbors in the previous snapshot. Furthermore, inductive KG embedding disregards learning the facts about new entities.

Our work is also relevant to dynamic KG embedding. pu-TransE (Tay, Luu, and Hui 2017) trains several new models when facts are added. DKGE (Wu et al. 2022) learns contextual embeddings for entities and relations, which can be automatically updated as the KG grows. They both need partial re-training on old facts, but our model does not. In addition, some subgraph-based models, such as GraIL (Teru, Denis, and Hamilton 2020), INDIGO (Liu et al. 2021), and TACT (Chen et al. 2021), can also represent unseen entities using the entity-independent features and subgraph aggregation. Their subgraph-building process is time-consuming, making them only applicable to small KGs. In order to run on large-scale KGs, NBFNet (Zhu et al. 2021) proposes a fast node pair embedding model based on the Bellman-Ford algorithm, and NodePiece (Galkin et al. 2022) uses tokenized anchor nodes and relational paths to represent new entities. However, they do not consider learning new knowledge and cannot support new relations.

2.2 Lifelong Learning

Lifelong learning seeks to solve new problems quickly without catastrophically forgetting previously acquired knowledge. Lifelong learning models are broadly classified into three categories. (i) Dynamic architecture models (Rusu et al. 2016; Lomonaco and Maltoni 2017) extend the network to learn new tasks and avoid forgetting acquired knowledge. (ii) Regularization-based models (Kirkpatrick et al. 2017; Zenke, Poole, and Ganguli 2017) capture the importance of model parameters for old tasks and limit the update of important parameters. (iii) Rehearsal-based models (Lopez-Paz and Ranzato 2017; Wang et al. 2019a) memorize some data from old tasks and replay them when learning new knowledge.

Few lifelong learning models focus on KG embedding. DiCGRL (Kou et al. 2020) is a disentangle-based lifelong graph embedding model. It splits node embeddings into different components and replays related historical facts to avoid catastrophic forgetting. The work (Daruna et al. 2021) combines class-incremental learning models with TransE (Bordes et al. 2013) for continual KG embedding. However, it does not propose a specific lifelong KG embedding model.

3 Lifelong Knowledge Graph Embedding

3.1 Preliminaries

Growing KG. The growth process of a KG yields a snapshot sequence, i.e., $\mathcal{G} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_t\}$. Each snapshot \mathcal{S}_i is defined as a triplet $(\mathcal{T}_i, \mathcal{E}_i, \mathcal{R}_i)$, where $\mathcal{T}_i, \mathcal{E}_i$ and \mathcal{R}_i denote the fact, entity and relation sets, respectively. We have $\mathcal{T}_i \subseteq \mathcal{T}_{i+1}$, $\mathcal{E}_i \subseteq \mathcal{E}_{i+1}$ and $\mathcal{R}_i \subseteq \mathcal{R}_{i+1}$. We use $\mathcal{T}_{\Delta i} = \mathcal{T}_i - \mathcal{T}_{i-1}$, $\mathcal{E}_{\Delta i} = \mathcal{E}_i - \mathcal{E}_{i-1}$, and $\mathcal{R}_{\Delta i} = \mathcal{R}_i - \mathcal{R}_{i-1}$ to denote the new facts, entities and relations, respectively. Each fact is in the form of $(s, r, o) \in \mathcal{T}_i$, where $s, o \in \mathcal{E}_i$ are the subject and object entities, respectively, and $r \in \mathcal{R}_i$ is their relation.

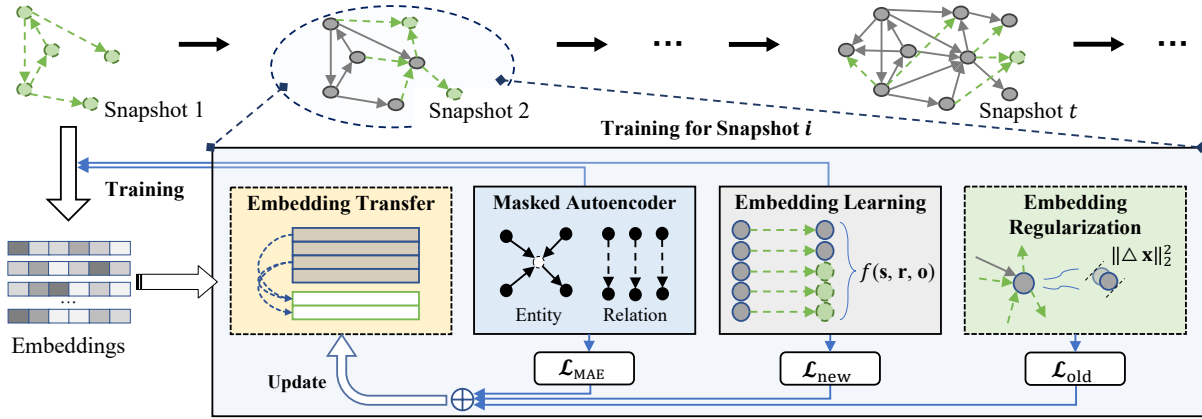


Figure 2: Overview of the proposed model for lifelong KG embedding.

Lifelong KG embedding. KG embedding seeks to encode the symbolic representations of entities and relations into vector space and capture KG semantics using vector operations. For a growing KG, a lifelong KG embedding model learns to represent the snapshot sequence $\mathcal{G} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_t\}$ continually. When a new fact set $\mathcal{T}_{\Delta i}$ emerges, the current KG embedding model \mathcal{M}_{i-1} needs update to fit the new facts and learn embeddings for the new entities $\mathcal{E}_{\Delta i}$ and new relations $\mathcal{R}_{\Delta i}$. The resulting model is denoted by \mathcal{M}_i .

Lifelong link prediction. The link prediction task asks the KG embedding model to predict the missing subject or object entity in an incomplete fact like $(s, r, ?)$ or $(?, r, o)$. For each snapshot \mathcal{S}_i , the new fact set $\mathcal{T}_{\Delta i}$ is divided into a training set \mathcal{D}_i , a validation set \mathcal{V}_i and a test set \mathcal{Q}_i . In the lifelong setting, the model is required to learn the training data sets, $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_t$, in turn. After finishing the learning on \mathcal{D}_i , the model is evaluated on the accumulated test data, which is $\cup_{j=1}^i \mathcal{Q}_j$, to assess the overall learning performance. After learning \mathcal{D}_i , then $\mathcal{D}_i, \mathcal{V}_i$ would no longer be available for the following learning. Note that the goal of lifelong KG embedding is to improve the overall performance on all snapshots, which requires the KG embedding model to continually learn new knowledge and retain the learned knowledge.

3.2 Model Overview

The overview of our model, LKGE, is shown in Figure 2. It continually learns knowledge over a sequence of KG snapshots without re-training on previously seen data. The foundation is a masked KG autoencoder that can reconstruct the entity and relation embeddings from the masked related subgraph. To enable knowledge transfer from old snapshots to the new one, we propose embedding transfer to inject learned knowledge into unseen entities and relations and then iteratively optimize the model to fit the new facts. We also propose a lightweight regularization method to retain old knowledge.

3.3 Masked Knowledge Graph Autoencoder

In lifelong KG embedding learning, the new facts, e.g., $\mathcal{T}_{\Delta i}$, bring new entities $\mathcal{E}_{\Delta i}$, and new relations $\mathcal{R}_{\Delta i}$. It would also involve some old entities from \mathcal{E}_{i-1} , and old relations from

\mathcal{R}_{i-1} . Thus, a new entity may be connected with both new and old entities (referred to as new knowledge). The involved old entities also receive more facts, and therefore their previously learned embeddings (referred to as old knowledge) need to be updated to fit the new facts. Hence, the base KG embedding model for lifelong learning should be capable of capturing new knowledge as well as updating old knowledge. To this end, we propose a masked KG autoencoder motivated by the recent success of self-supervised learning (Hou et al. 2022). The key idea is to reconstruct the embedding for an entity or relation based on its masked subgraph, which may include both other new entities and some old entities. Specifically, we use the first-order subgraph of an entity or relation to reconstruct its embedding \bar{x}_i :

$$\bar{x}_i = \text{MAE}(\cup_{j=1}^i \mathcal{N}_j(x)), \quad (1)$$

where x denotes either an entity or a relation, and $\mathcal{N}_j \subseteq \mathcal{D}_j$ denotes the involved facts of x in the j -th snapshot. $\text{MAE}()$ is an encoder to represent the input subgraph. The objective of our KG encoder is to align the entity or relation embedding with the reconstructed representation as follows:

$$\mathcal{L}_{\text{MAE}} = \sum_{e \in \mathcal{E}_i} \|\mathbf{e}_i - \bar{\mathbf{e}}_i\|_2^2 + \sum_{r \in \mathcal{R}_i} \|\mathbf{r}_i - \bar{\mathbf{r}}_i\|_2^2. \quad (2)$$

The key then becomes how to design an effective and efficient encoder for lifelong learning. GCN (Kipf and Welling 2017) and Transformer (Vaswani et al. 2017) are two common choices for the encoder. The two encoders both introduce additional model parameters (e.g., the weight matrices). In our lifelong learning setting, the encoder needs to be updated to fit new facts or subgraphs. In this case, once the GCN or Transformer is updated, the changed model parameters would affect the embedding generation of *all* old entities (not just the involved old entities in new facts), increasing the risk of catastrophically forgetting previous snapshots. To avoid this issue, we use the entity and relation embedding transition functions as encoders, which do not introduce additional parameters. We borrow the idea of TransE (Bordes et al. 2013) and interpret a relation embedding as the translation vector between the subject and object entity embeddings,

i.e., $\mathbf{s} + \mathbf{r} \approx \mathbf{o}$, where $\mathbf{s}, \mathbf{r}, \mathbf{o}$ denote the embeddings of subject entity, relation and object entity, respectively. Based on this, we can deduce two transition functions for entity and relation embeddings. The subject entity of (s, r, o) can be represented by $f_{\text{sub}}(\mathbf{r}, \mathbf{o}) = \mathbf{o} - \mathbf{r}$, and the relation embedding is $f_{\text{rel}}(\mathbf{s}, \mathbf{o}) = \mathbf{o} - \mathbf{s}$. We can define the encoders as

$$\bar{\mathbf{e}}_i = \frac{\sum_{j=1}^i \sum_{(s,r,o) \in \mathcal{N}_j(e)} f_{\text{sub}}(\mathbf{r}_i, \mathbf{o}_i)}{\sum_{j=1}^i |\mathcal{N}_j(e)|}, \quad (3)$$

$$\bar{\mathbf{r}}_i = \frac{\sum_{j=1}^i \sum_{(s,r,o) \in \mathcal{N}_j(r)} f_{\text{rel}}(\mathbf{s}_i, \mathbf{o}_i)}{\sum_{j=1}^i |\mathcal{N}_j(r)|}, \quad (4)$$

where $\mathbf{s}_i, \mathbf{r}_i, \mathbf{o}_i$ are the embeddings of s, r, o during the training on \mathcal{D}_i . $\mathcal{N}_j(x) \subseteq \mathcal{D}_j$ is the set of facts containing x .

In lifelong learning, the model learns from a snapshot sequence. Eqs. (3) and (4) require training samples from the first i snapshots, which are not in line with lifelong learning. To reduce the reliance on learned data, we use \mathbf{e}_{i-1} and \mathbf{r}_{i-1} as the approximate average embeddings of e and r in the first $i-1$ snapshots, respectively, and rewrite the encoders as

$$\bar{\mathbf{e}}_i \approx \frac{\sum_{j=1}^{i-1} |\mathcal{N}_j(e)| \mathbf{e}_{i-1} + \sum_{(s,r,o) \in \mathcal{N}_i(e)} f_{\text{sub}}(\mathbf{r}_i, \mathbf{o}_i)}{\sum_{j=1}^{i-1} |\mathcal{N}_j(e)| + |\mathcal{N}_i(e)|}, \quad (5)$$

$$\bar{\mathbf{r}}_i \approx \frac{\sum_{j=1}^{i-1} |\mathcal{N}_j(r)| \mathbf{r}_{i-1} + \sum_{(s,r,o) \in \mathcal{N}_i(r)} f_{\text{rel}}(\mathbf{s}_i, \mathbf{o}_i)}{\sum_{j=1}^{i-1} |\mathcal{N}_j(r)| + |\mathcal{N}_i(r)|}. \quad (6)$$

The encoders use the facts involving both old and new entities and relations for embedding reconstruction and they build a bridge for knowledge transfer.

For each snapshot, to learn the knowledge from the new data and update the learned parameters, we leverage TransE (Bordes et al. 2013) to train the embedding model:

$$\mathcal{L}_{\text{new}} = \sum_{(s,r,o) \in \mathcal{D}_i} \max(0, \gamma + f(\mathbf{s}, \mathbf{r}, \mathbf{o}) - f(\mathbf{s}', \mathbf{r}, \mathbf{o}')), \quad (7)$$

where γ is the margin. $(\mathbf{s}', \mathbf{r}, \mathbf{o}')$ is the embedding of a negative fact. For each positive fact, we randomly replace the subject or object entity with a random entity $e' \in \mathcal{E}_i$.

3.4 Embedding Transfer

During the lifecycle of a growing KG, there are abundant unseen entities and some unseen relations emerge with the new facts. Learning effective embeddings for them is an essential aspect of lifelong KG embedding learning. However, these unseen ones are not included in any learned snapshots, so only inheriting the learned parameters cannot transfer the acquired knowledge to their embeddings. To avoid learning from scratch, we propose embedding transfer that seeks to leverage the learned embeddings to help represent unseen entities and relations. Specifically, we initialize the embeddings of each unseen entity by aggregating its facts:

$$\mathbf{e}_i = \frac{1}{|\mathcal{N}_i(e)|} \sum_{(e,r,o) \in \mathcal{N}_i(e)} f_{\text{sub}}(\mathbf{r}_{i-1}, \mathbf{o}_{i-1}), \quad (8)$$

where $\mathcal{N}_i(e) \subseteq \mathcal{D}_i$ is the set of facts containing e . For the new entities that do not have common facts involving existing

entities, we randomly initialize their embeddings. We also use this strategy to initialize the embeddings of unseen relations:

$$\mathbf{r}_i = \frac{1}{|\mathcal{N}_i(r)|} \sum_{(s,r,o) \in \mathcal{N}_i(r)} f_{\text{rel}}(\mathbf{s}_{i-1}, \mathbf{o}_{i-1}), \quad (9)$$

where $\mathcal{N}_i(r) \subseteq \mathcal{D}_i$ is the set of facts containing r .

3.5 Embedding Regularization

Learning new snapshots is likely to overwrite the learned knowledge from old snapshots. To avoid catastrophic forgetting, some regularization methods (Kirkpatrick et al. 2017; Lopez-Paz and Ranzato 2017) constrain the updates of parameters that are important to old tasks. The loss function of regularization methods is

$$\mathcal{L}_{\text{old}} = \sum_{e \in \mathcal{E}_{i-1}} \omega(e) \|\mathbf{e}_i - \mathbf{e}_{i-1}\|_2^2 + \sum_{r \in \mathcal{R}_{i-1}} \omega(r) \|\mathbf{r}_i - \mathbf{r}_{i-1}\|_2^2, \quad (10)$$

where $\omega(x)$ is the regularization weight for x .

Conventional regularization-based methods for classification tasks such as (Kirkpatrick et al. 2017) model the importance of each parameter at a high cost based on the gradient or parameter change during training. This problem is even more severe for KG embedding models that have a large number of embedding parameters (i.e. entity and relation embeddings). To resolve this problem, we propose a lightweight embedding regularization method, which calculates the regularization weight of each entity or relation by the number of new and old facts containing it:

$$\omega(x) = 1 - \frac{|\mathcal{N}_i(x)|}{\sum_{j=1}^i |\mathcal{N}_j(x)|}. \quad (11)$$

As a lightweight technique, it only keeps the total number of involved trained facts for each entity or relation and only updates regularization weights once per snapshot.

3.6 Overall Learning Objective

To learn new knowledge while retaining acquired knowledge, the overall lifelong learning objective \mathcal{L} is defined as follows:

$$\mathcal{L} = \mathcal{L}_{\text{new}} + \alpha \mathcal{L}_{\text{old}} + \beta \mathcal{L}_{\text{MAE}}, \quad (12)$$

where α, β are hyperparameters for balancing the objectives.

3.7 Complexity Analysis

Compared with fine-tuning, the proposed model requires few extra resources. It does not increase the size of training samples like the rehearsal models (Wang et al. 2019a; Lopez-Paz and Ranzato 2017; Kou et al. 2020). In addition to fine-tuning, the proposed model calculates the loss of masked autoencoder and embedding regularization. The additional time complexity in each iteration is $O(|\mathcal{E}| + |\mathcal{R}| + |\mathcal{D}|)$. In practice, we find that the loss of autoencoder can accelerate learning, and its time consumption is close to that of fine-tuning. The space complexity of fine-tuning is $O((|\mathcal{E}| + |\mathcal{R}|) \times d)$, and the space complexity of the proposed model is $O((|\mathcal{E}| + |\mathcal{R}|) \times (d+1))$, where d is the dimension of embeddings.

Datasets	Snapshot 1			Snapshot 2			Snapshot 3			Snapshot 4			Snapshot 5		
	$ \mathcal{T}_{\Delta 1} $	$ \mathcal{E}_1 $	$ \mathcal{R}_1 $	$ \mathcal{T}_{\Delta 2} $	$ \mathcal{E}_2 $	$ \mathcal{R}_2 $	$ \mathcal{T}_{\Delta 3} $	$ \mathcal{E}_3 $	$ \mathcal{R}_3 $	$ \mathcal{T}_{\Delta 4} $	$ \mathcal{E}_4 $	$ \mathcal{R}_4 $	$ \mathcal{T}_{\Delta 5} $	$ \mathcal{E}_5 $	$ \mathcal{R}_5 $
ENTITY	46,388	2,909	233	72,111	5,817	236	73,785	8,275	236	70,506	11,633	237	47,326	14,541	237
RELATION	98,819	11,560	48	93,535	13,343	96	66,136	13,754	143	30,032	14,387	190	21,594	14,541	237
FACT	62,024	10,513	237	62,023	12,779	237	62,023	13,586	237	62,023	13,894	237	62,023	14,541	237
HYBRID	57,561	8,628	86	20,873	10,040	102	88,017	12,779	151	103,339	14,393	209	40,326	14,541	237

Table 1: Statistical data of the four constructed growing KG datasets. For the i -th snapshot, $\mathcal{T}_{\Delta i}$ denotes the set of new facts in this snapshot, and $\mathcal{E}_i, \mathcal{R}_i$ denote the sets of cumulative entities and relations in the first i snapshots, respectively.

4 Dataset Construction

To simulate a variety of aspects of KG growth, we create four datasets based on FB15K-237 (Toutanova and Chen 2015), which are entity-centric, relation-centric, fact-centric, and hybrid. We denote them by ENTITY, RELATION, FACT and HYBRID, respectively. Given a KG $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{T}\}$, we construct five snapshots with the following steps:

- Seeding.** We randomly sample 10 facts from \mathcal{T} and add them into \mathcal{T}_1 for initialization. The entities and relations in the 10 facts form the initial \mathcal{E}_1 and \mathcal{R}_1 , respectively.
- Expanding.** To build ENTITY, RELATION and FACT, we iteratively sample a fact containing at least one seen entity in \mathcal{E}_i , add it into \mathcal{T}_i , and extract the unseen entity and relation from it to expand \mathcal{E}_i and \mathcal{R}_i . For ENTITY, once $|\mathcal{E}_i| \geq \frac{i+1}{5}|\mathcal{E}|$, we add all new facts $\{(s, r, o) \mid s \in \mathcal{E}_i \wedge o \in \mathcal{E}_i\}$ into \mathcal{T}_i and start building the next snapshot. In the same way, we construct RELATION and FACT. As for HYBRID, we uniformly sample an entity, relation or fact without replacement from $\mathcal{U} = \mathcal{E} \cup \mathcal{R} \cup \mathcal{T}$ to join $\mathcal{E}_i, \mathcal{R}_i$ and \mathcal{T}_i . Note that when the sampled fact contains an unseen entity or relation, we re-sample a fact that only contains seen entities and relations to replace it. After each iteration, we terminate the expansion of this snapshot with a probability $\frac{5}{|\mathcal{U}|}$. Consequently, the expansion of HYBRID is uneven, making it more realistic and challenging. For all datasets, we take the whole KG as the last snapshot, i.e., $\mathcal{T}_5 = \mathcal{T}$, and $\mathcal{E}_5 = \mathcal{E}, \mathcal{R}_5 = \mathcal{R}$.
- Dividing.** For each snapshot, we randomly divide the new fact set $\mathcal{T}_{\Delta i}$ into a training set \mathcal{D}_i , a validation set \mathcal{V}_i and a test set \mathcal{Q}_i by a split ratio of 3:1:1.

The statistics of the four datasets are presented in Table 1.

5 Experimental Results

We conduct experiments regarding link prediction accuracy, knowledge transfer capability, and learning efficiency to validate the proposed model, LKGE. The datasets and source code are available at <https://github.com/nju-websoft/LKGE>.

5.1 Experiment Settings

Competitors. We compare our model with 12 competitors, including (i) three baseline models: snapshot only, re-training, and fine-tuning; (ii) two inductive models: MEAN (Hamaguchi et al. 2017), and LAN (Wang et al. 2019b); (iii) two dynamic architecture models: PNN (Rusu et al. 2016), and CWR (Lomonaco and Maltoni 2017); (iv) two regularization-based models: SI (Zenke, Poole, and Ganguli 2017), and

EWC (Kirkpatrick et al. 2017); and (v) three rehearsal-based models: GEM (Lopez-Paz and Ranzato 2017), EMR (Wang et al. 2019a), and DiCGRL (Kou et al. 2020).

Evaluation metrics. Following the convention, we conduct the experiments on link prediction. Given a snapshot \mathcal{S}_i , for each test fact $(s, r, o) \in \mathcal{Q}_i$, we construct two queries $(s, r, ?)$ and $(?, r, t)$. When evaluating on \mathcal{Q}_i , we set all seen entities in \mathcal{E}_i as candidate entities. We select seven metrics to evaluate all models, including (i) Four metrics on link prediction accuracy: mean reciprocal rank (MRR) and Hits@ k ($k = 1, 3, 10$, and H@ k for shot). We conduct the model \mathcal{M}_5 trained on the last snapshot to evaluate on the union of the test sets in all snapshots. (ii) Two metrics on knowledge transfer capability: forward transfer (FWT) and backward transfer (BWT) (Lopez-Paz and Ranzato 2017). FWT is the influence of learning a task to the performance on the future tasks, while BWT is the influence of learning to the previous tasks:

$$\text{FWT} = \frac{1}{n-1} \sum_{i=2}^n h_{i-1,i}, \text{BWT} = \frac{1}{n-1} \sum_{i=1}^{n-1} (h_{n,i} - h_{i,i}), \quad (13)$$

where n is the number of snapshots, $h_{i,j}$ is the MRR scores on \mathcal{Q}_j after training the model \mathcal{M}_i on the i -th snapshot. Higher scores indicate better performance. (iii) Time cost: The cumulative time cost of the learning on each snapshot.

Implementation details. We use TransE (Bordes et al. 2013) as the base model and modify the competitors to do our task:

- Snapshot only.** For the i -th snapshot, we reinitialize and train a model only on the training set \mathcal{D}_i .
- Re-training.** For the i -th snapshot, we reinitialize and train a model on the accumulated training data $\cup_{j=1}^i \mathcal{D}_j$.
- Fine-tuning.** For the i -th snapshot, the model inherits the learned parameters of the model trained on the previous snapshots, and we incrementally train it on \mathcal{D}_i .
- Inductive models.** We train each model on the first snapshot and obtain the embeddings of unseen entities in the following snapshots by neighborhood aggregation.
- Dynamic architecture models.** For PNN, following the implementation of (Daruna et al. 2021), we freeze the parameters learned on previous snapshots and update new parameters. For CWR, after training on \mathcal{D}_1 , we replicate a model as the consolidated model. For the following i -th snapshot, we reinitialize and train a temporal model on \mathcal{D}_i , and merge the temporal model into the consolidated model by copying new parameters or averaging old ones.
- Regularization models.** Since the base model parameters increase with the emergence of unseen entities and

Models	ENTITY				RELATION				FACT				HYBRID			
	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10
Snapshot	0.084	0.028	0.107	0.193	0.021	0.010	0.023	0.043	0.082	0.030	0.095	0.191	0.036	0.015	0.043	0.077
Re-train	0.236	0.137	0.274	0.433	0.219	0.128	0.250	0.403	0.206	0.118	0.232	0.385	0.227	0.134	0.260	0.413
Fine-tune	0.165	0.085	0.188	0.321	0.093	0.039	0.106	0.195	0.172	0.090	0.193	0.339	0.135	0.069	0.151	0.262
MEAN	0.117	0.068	0.123	0.212	0.039	0.024	0.040	0.067	0.084	0.051	0.088	0.146	0.046	0.029	0.049	0.080
LAN	0.141	0.082	0.149	0.256	0.052	0.033	0.052	0.092	0.106	0.056	0.113	0.200	0.059	0.032	0.062	0.113
PNN	0.229	0.130	0.265	0.425	0.167	0.096	0.191	0.305	0.157	0.084	0.188	0.290	0.185	0.101	0.216	0.349
CWR	0.088	0.028	0.114	0.202	0.021	0.010	0.024	0.043	0.083	0.030	0.095	0.192	0.037	0.015	0.044	0.077
SI	0.154	0.072	0.179	0.311	0.113	0.055	0.131	0.224	0.172	0.088	0.194	0.343	0.111	0.049	0.126	0.229
EWC	0.229	0.130	0.264	0.423	0.165	0.093	0.190	0.306	0.201	0.113	0.229	0.382	0.186	0.102	0.214	0.350
GEM	0.165	0.085	0.188	0.321	0.093	0.040	0.106	0.196	0.175	0.092	0.196	0.345	0.136	0.070	0.152	0.263
EMR	0.171	0.090	0.195	0.330	0.111	0.052	0.126	0.225	0.171	0.090	0.191	0.337	0.141	0.073	0.157	0.267
DiCGRL	0.107	0.057	0.110	0.211	0.133	0.079	0.147	0.241	0.162	0.084	0.189	0.320	0.149	0.083	0.168	0.277
LKGE	0.234	0.136	0.269	0.425	0.192	0.106	0.219	0.366	0.210	0.122	0.238	0.387	0.207	0.121	0.235	0.379

Table 2: Result comparison of link prediction on the union of the test sets in all snapshots.

relations, we only use the parameters learned from the previous snapshot to calculate the regularization loss.

- *Rehearsal models.* We store 5,000 training facts from previous snapshots and add them to the current training set of the i -th snapshot. After the learning, we randomly replace half of these facts with those in \mathcal{D}_i .

For a fair comparison, we first tune the hyperparameters of the base model using grid-search: learning rate in $\{0.0005, 0.0001, 0.001\}$, batch size in $\{1024, 2048\}$, embedding dimension in $\{100, 200\}$. Then, we use the same base model for LKGE and all competitors, and tune other hyperparameters. For the regularization models, the α of regularization loss is in $\{0.01, 0.1, 1.0\}$. For our model, the β of MAE loss is in $\{0.01, 0.1, 1.0\}$. For all competitors, we use Adam optimizer and set the patience of early stopping to 3.

5.2 Link Prediction Accuracy

In Table 2, we run 5-seeds experiments for all models on our datasets and report the means. The results show that: (i) our model consistently achieves the best performance across all datasets. Some results of our model in FACT even outperforms re-training. This is because our masked KG autoencoder effectively improves information propagation based on both old and new embeddings, and the embedding regularization avoids catastrophic forgetting. Most competitors only work well on ENTITY, while our model shows stable and promising results on all these datasets. (ii) Re-training is far superior to most baseline models on RELATION and HYBRID, while the gaps on ENTITY and FACT are small. This is because the KG embedding model learns two aspects of knowledge: relational patterns and entity embeddings. In ENTITY and FACT, the relational patterns are stable, while in RELATION and HYBRID, their relational patterns are constantly changing due to unseen relations. These phenomena illustrate that the variation of relational patterns is more challenging for lifelong KG embedding. (iii) The inductive models are only trained on the first snapshot, and cannot transfer knowledge to unseen relations. So, their results are lower than other

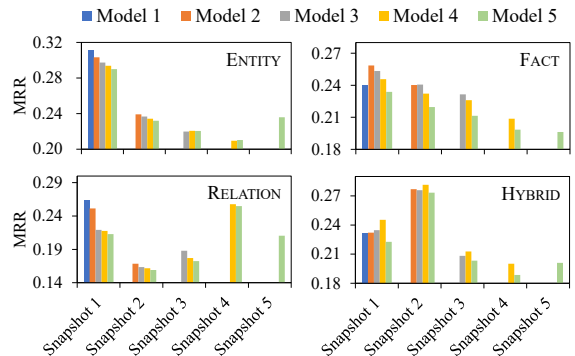


Figure 3: MRR changes. \mathcal{M}_i is trained for the i -th snapshot and evaluated using the test data of previous snapshots 1 to i .

models, especially on RELATION and HYBRID with many unseen relations. (iv) Since the learned parameters are not updated, PNN preserves the learned knowledge well. But on FACT, due to a few unseen entities, it lacks new learnable parameters, and the performance is not well. CWR averages the old and new model parameters, which does not work well on the embedding learning task. (v) EWC performs well because it can model the importance of each parameter using Fisher information matrices to the learned snapshots. (vi) Unlike the classification tasks, most parameters of KG embedding models correspond to specific entities, so the training data cannot be divided into a few types, and we cannot use 5,000 old samples to replay the learned facts for all entities. The performance of GEM, EMR and DiCGRL is limited.

To show the performance evolution of LKGE during the learning process, we evaluate the model \mathcal{M}_i trained for the i -th snapshot using the test data from previous snapshots. The MRR results are in Figure 3. LKGE can maintain the learned knowledge during lifelong learning. On some snapshots like ENTITY Snapshot 3, the knowledge update improves the performance on old test data, which shows that old knowledge update has the potential for backward knowledge transfer.

Variants	ENTITY				RELATION				FACT				HYBRID			
	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10
LKGE (full)	0.234	0.136	0.269	0.425	0.192	0.106	0.219	0.366	0.210	0.122	0.238	0.387	0.207	0.121	0.235	0.379
– fine-tuning	0.123	0.068	0.136	0.225	0.126	0.073	0.146	0.231	0.154	0.091	0.176	0.269	0.109	0.060	0.126	0.201
– autoencoder	0.222	0.124	0.255	0.415	0.185	0.100	0.212	0.355	0.191	0.105	0.215	0.369	0.198	0.111	0.227	0.367
– embedding transfer	0.240	0.141	0.275	0.433	0.174	0.091	0.201	0.339	0.210	0.123	0.237	0.390	0.200	0.112	0.229	0.372
– regularization	0.166	0.089	0.184	0.316	0.040	0.014	0.049	0.089	0.175	0.095	0.195	0.338	0.154	0.079	0.171	0.300

Table 3: Ablation results of link prediction on the union of the test sets in all snapshots.

5.3 Knowledge Transfer Capability

To evaluate the knowledge transfer and retention capability of all models, we report the FWT and BWT of MRR results in Figure 4. Because of the embedding transfer, the FWT of LKGE is higher than all lifelong learning competitors. Even on RELATION and HYBRID where the KG schema changes, LKGE still keeps the FWT capability well. MEAN and LAN are designed to transfer knowledge forward to embed new entities. So, they work well on ENTITY. However, their FWT capability is limited on other datasets since they cannot update the old embeddings to adapt to new snapshots.

BWT is usually negative due to the overwriting of learned knowledge. PNN, MEAN, and LAN do not update old parameters. Their BWT scores are “NA”. The poor scores of CWR show the harmful effects of the average operation. The scores of rehearsal models are also not good as they cannot store enough facts. LKGE gets good BWT scores as the embedding regularization can well maintain the learned knowledge.

5.4 Learning Efficiency

We show the training time on FACT as all snapshots of FACT have the same training set size. Figure 5 shows the results. Unsurprisingly, re-training is most time-consuming. Snapshot is also costly because it cannot inherit knowledge from previous snapshots. By contrast, our model is most efficient, and its advantage is more significant in the final snapshot. This is because the embedding transfer can use the learned knowledge to accelerate the learning of new facts.

5.5 Ablation Study

We conduct an ablation study by designing four variants of LKGE: “w/o fine-tuning”, “w/o autoencoder”, “w/o embedding transfer” and “w/o regularization”. The “w/o fine-tuning” variant is trained on \mathcal{D}_1 and performs the embedding transfer on other \mathcal{D}_i . The latter three variants disable the specific components in LKGE. The results are shown in Table 3. We see that (i) although fine-tuning is disabled, “w/o fine-tuning” can still perform well with only the knowledge from the first snapshot, showing that embedding transfer can effectively transfer learned knowledge to unseen entities and relations. (ii) Both “w/o autoencoder” and “w/o regularization” significantly drop, showing the effects of masked KG autoencoder and knowledge retention. (iii) Embedding transfer enables the model to be trained at a starting point closer to the optimal parameters and stabilizes the embedding space. There are declines when using the embedding transfer on ENTITY. This is because ENTITY contains massive new entities and needs

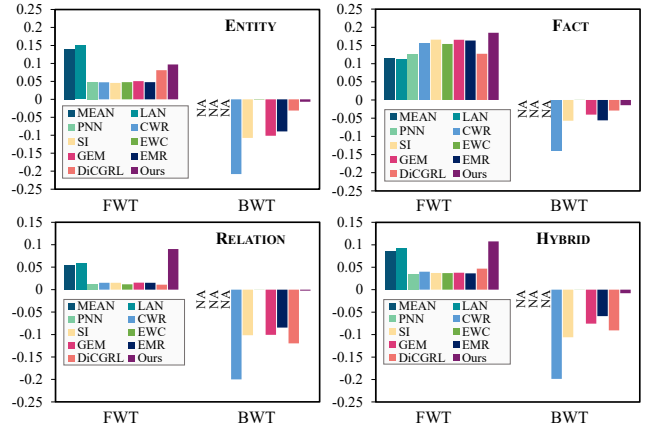


Figure 4: Forward transfer and backward transfer of MRR.

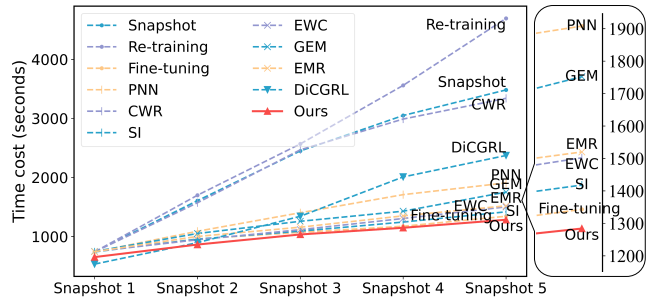


Figure 5: Cumulative time cost on FACT.

more plasticity rather than stability. But on RELATION and HYBRID, the results of “w/o embedding transfer” are lower than the full model, showing that embedding transfer can reduce the interference caused by the KG schema changes. On FACT, the results of “w/o embedding transfer” are similar to the full model. This shows that, even without embedding transfer, the model can still capture the knowledge.

6 Conclusion and Future Work

This paper studies lifelong embedding learning for growing KGs. For better knowledge transfer and retention, we propose a lifelong KG embedding model consisting of masked KG autoencoder, embedding transfer, and embedding regularization. Experiments on new datasets show better link prediction accuracy, knowledge transfer capability, and learning efficiency of our model. In future, we plan to study lifelong embedding learning in long-tail and low-resource settings.

Acknowledgments

This work was supported by National Natural Science Foundation of China (No. 62272219).

References

- Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *NeurIPS*, 2787–2795.
- Chen, J.; He, H.; Wu, F.; and Wang, J. 2021. Topology-Aware Correlations Between Relations for Inductive Link Prediction in Knowledge Graphs. In *AAAI*, 6271–6278.
- Daruna, A. A.; Gupta, M.; Sridharan, M.; and Chernova, S. 2021. Continual Learning of Knowledge Graph Embeddings. *IEEE Robotics Autom. Lett.*, 6(2): 1128–1135.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2018. Convolutional 2D Knowledge Graph Embeddings. In *AAAI*, 1811–1818.
- Galkin, M.; Denis, E. G.; Wu, J.; and Hamilton, W. L. 2022. NodePiece: Compositional and Parameter-Efficient Representations of Large Knowledge Graphs. In *ICLR*.
- Guo, L.; Sun, Z.; and Hu, W. 2019. Learning to Exploit Long-term Relational Dependencies in Knowledge Graphs. In *ICML*, volume 97, 2505–2514.
- Hamaguchi, T.; Oiwa, H.; Shimbo, M.; and Matsumoto, Y. 2017. Knowledge Transfer for Out-of-Knowledge-Base Entities: A Graph Neural Network Approach. In *IJCAI*, 1802–1808.
- Hou, Z.; Liu, X.; Cen, Y.; Dong, Y.; Yang, H.; Wang, C.; and Tang, J. 2022. GraphMAE: Self-Supervised Masked Graph Autoencoders. *CoRR*, abs/2205.10803: 1–11.
- Ji, S.; Pan, S.; Cambria, E.; Marttinen, P.; and Yu, P. S. 2022. A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *IEEE Trans. Neural Netw. Learn. Syst.*, 33(2): 494–514.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N. C.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; Hassabis, D.; Clopath, C.; Kuzmaran, D.; and Hadsell, R. 2017. Overcoming Catastrophic Forgetting in Neural Networks. *PNAS*, 114(13): 3521–3526.
- Kou, X.; Lin, Y.; Liu, S.; Li, P.; Zhou, J.; and Zhang, Y. 2020. Disentangle-based Continual Graph Representation Learning. In *EMNLP*, 2961–2972.
- Liu, S.; Grau, B. C.; Horrocks, I.; and Kostylev, E. V. 2021. INDIGO: GNN-Based Inductive Knowledge Graph Completion Using Pair-Wise Encoding. In *NeurIPS*, 2034–2045.
- Lomonaco, V.; and Maltoni, D. 2017. CORE50: A New Dataset and Benchmark for Continuous Object Recognition. In *CoRL*, volume 78, 17–26.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient Episodic Memory for Continual Learning. In *NeurIPS*, 6467–6476.
- Rusu, A. A.; Rabinowitz, N. C.; Desjardins, G.; Soyer, H.; Kirkpatrick, J.; Kavukcuoglu, K.; Pascanu, R.; and Hadsell, R. 2016. Progressive Neural Networks. *CoRR*, abs/1606.04671: 1–14.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2009. The Graph Neural Network Model. *IEEE Trans. Neural Netw.*, 20(1): 61–80.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; van den Berg, R.; Titov, I.; and Welling, M. 2018. Modeling Relational Data with Graph Convolutional Networks. In *ESWC*, 593–607.
- Tay, Y.; Luu, A. T.; and Hui, S. C. 2017. Non-Parametric Estimation of Multiple Embeddings for Link Prediction on Dynamic Knowledge Graphs. In *AAAI*, 1243–1249.
- Teru, K.; Denis, E.; and Hamilton, W. 2020. Inductive Relation Prediction by Subgraph Reasoning. In *ICML*, 9448–9457.
- Toutanova, K.; and Chen, D. 2015. Observed Versus Latent Features for Knowledge Base and Text Inference. In *CVSC*, 57–66.
- Vashishth, S.; Sanyal, S.; Nitin, V.; Agrawal, N.; and Talukdar, P. P. 2020a. InteractE: Improving Convolution-based Knowledge Graph Embeddings by Increasing Feature Interactions. In *AAAI*, 3009–3016.
- Vashishth, S.; Sanyal, S.; Nitin, V.; and Talukdar, P. P. 2020b. Composition-based Multi-Relational Graph Convolutional Networks. In *ICLR*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NeurIPS*, 5998–6008.
- Vrandečić, D.; and Krötzsch, M. 2014. Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM*, 57(10): 78–85.
- Wang, H.; Xiong, W.; Yu, M.; Guo, X.; Chang, S.; and Wang, W. Y. 2019a. Sentence Embedding Alignment for Lifelong Relation Extraction. In *NAACL*, 796–806.
- Wang, P.; Han, J.; Li, C.; and Pan, R. 2019b. Logic Attention Based Neighborhood Aggregation for Inductive Knowledge Graph Embedding. In *AAAI*, 7152–7159.
- Wang, Q.; Mao, Z.; Wang, B.; and Guo, L. 2017. Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE Trans. Knowl. Data Eng.*, 29(12): 2724–2743.
- Wang, Z.; Zhang, J.; Feng, J.; and Chen, Z. 2014. Knowledge Graph Embedding by Translating on Hyperplanes. In *AAAI*, 1112–1119.
- Wu, T.; Khan, A.; Yong, M.; Qi, G.; and Wang, M. 2022. Efficiently Embedding Dynamic Knowledge Graphs. *Knowl. Based Syst.*, 250: 109124.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual Learning through Synaptic Intelligence. In *ICML*, 3987–3995.
- Zhu, Z.; Zhang, Z.; Xhonneux, L. A. C.; and Tang, J. 2021. Neural Bellman-Ford Networks: A General Graph Neural Network Framework for Link Prediction. In *NeurIPS*, 29476–29490.