Differentially Private Federated Knowledge Graphs Embedding

Hao Peng[1,4], Haoran Li[2,5], Yangqiu Song[2,5], Vincent Zheng[3], Jianxin Li[1,6]
[1]Beijing Advanced Innovation Center for Big Data and Brain Computing, Beihang University, Beijing 100191, China;
[2]Department of Computer Science and Engineering, HKUST, Hongkong, China; [3]AI Group, Webank Co., Ltd;
[4]School of Cyber Science and Technology, Beihang University, Beijing 100191, China;
[5]Peng Cheng Laboratory, Shenzhen 518066, China; [6]SKLSDE, Beihang University, Beijing 100191, China;
{penghao,lijx}@act.buaa.edu.cn,hlibt@connect.ust.hk,yqsong@cse.ust.hk,vincentz@webank.com.

ABSTRACT

Knowledge graph embedding plays an important role in knowledge representation, reasoning, and data mining applications. However, for multiple cross-domain knowledge graphs, state-of-the-art embedding models cannot make full use of the data from different knowledge domains while preserving the privacy of exchanged data. In addition, the centralized embedding model may not scale to the extensive real-world knowledge graphs. Therefore, we propose a novel decentralized scalable learning framework, *Federated Knowledge Graphs Embedding* (FKGE), where embeddings from different knowledge graphs can be learnt in an asynchronous and peer-to-peer manner while being privacy-preserving. FKGE exploits adversarial generation between pairs of knowledge graphs to translate identical entities and relations of different domains into near embedding spaces. In order to protect the privacy of the training data, FKGE further implements a privacy-preserving neural network structure to guarantee no raw data leakage. We conduct extensive experiments to evaluate FKGE on 11 knowledge graphs, demonstrating a significant and consistent improvement in model quality with at most 17.85% and 7.90% increases in performance on triple classification and link prediction tasks.

CCS CONCEPTS

• **Computing methodologies** → **Knowledge representation and reasoning**; *Unsupervised learning*;

KEYWORDS

Federated Learning; Knowledge Graph Embedding; Differential Privacy; GAN

ACM Reference Format:
Hao Peng[1,4], Haoran Li[2,5], Yangqiu Song[2,5], Vincent Zheng[3], Jianxin Li[1,6]. 2021. Differentially Private Federated Knowledge Graphs Embedding. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21), November 1–5, 2021, Virtual Event, QLD, Australia.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3459637.3482252

Hao Peng and Haoran Li contribute equally.

CIKM '21, November 1–5, 2021, Virtual Event, QLD, Australia

ACM ISBN 978-1-4503-8446-9/21/11…$15.00
https://doi.org/10.1145/3459637.3482252

1 INTRODUCTION

Knowledge graphs (KGs) have been built to benefit many applications, e.g., semantic search, question answering, recommendation systems, etc. [14, 20, 26, 52]. There have been several big and general-purpose KGs such as Freebase [3] (or later Wikidata [43]) and Yago [37], and numerous domain-specific KGs of various sizes such as GeoNames [1] in geography and Lexvo [11] in linguistics. To our best knowledge, most companies build their own commercial KGs, which usually require laborious human annotation and high computational cost*. However, there are multiple reasons that companies would not want to share their KGs. First, each company has its private part of data, which cannot be disclosed to others. Second, even if privacy is not a concern, they would not expose their knowledge to other companies except they can also benefit from others. Third, integrating knowledge itself is not trivial or easy. On the other hand, in many cases, companies indeed have the motivation to exchange knowledge to improve their own data quality and service. For example, a drug discovery company may benefit from a patient social network and a gene bank, and so as the other two interested in a drug network and a disease network owned by the drug discovery company. Sometimes, the data cannot even be bought due to the privacy concern. For example, under EU's GDPR† , companies cannot use or share a user's data without his/her consent. Therefore, a more loosely coupled and principled way to share their KGs that can benefit multiple parties should be considered.

Traditionally, federated database systems [36] were proposed to support unified query language over heterogeneous databases without doing actual data integration. Such systems do not help improve individual KG's quality or service with private data preserved. Recently, machine learning, in particular KG representation learning or KG embedding, has been shown to be powerful for knowledge representation, reasoning, and many downstream applications [20, 28, 31, 44]. When representing entities and their relations in a vector space, different KGs may share information if their embedding spaces are (partially) aligned [8, 40, 49, 55]. However, revealing vector representations to other parties can also leak private information [6]. Currently, for multiple cross-domain KGs, especially for the ones shared by different companies, state-of-the-art KG embedding models cannot make full use of the data from different domains while preserving the privacy. More recently, *federated machine learning* [51] has been widely considered when heterogeneous devices or multiple parties contribute non-IID data where privacy of different sources should be preserved. However,

*See https://www.maana.io/ for a general pipeline of components.
†https://gdpr-info.eu

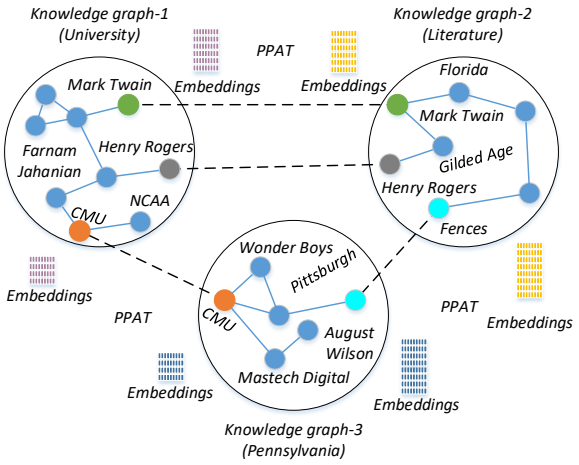

Figure 1: Overview of FKGE framework. Different Knowledge graphs may use different embedding models.

the natures of distributed federated learning [21, 22] (which requires centralized averaging platform) and vertical or horizontal federated learning (which do not consider graph structures) make them non-trivial to be adapted to KG embeddings. There are some existing federated learning mechanisms for graphs [23] but they treat each node to be a computing cell and cannot collaboratively train multiple KGs together.

To improve the quality of individual KGs for multiple cross-domain KGs, we propose a new framework called Federated Knowledge Graphs Embedding (FKGE), where embeddings from different KGs can be learnt jointly in an asynchronous and pairwise manner while being privacy-preserving. In FKGE, we design a privacy-preserving adversarial translation (PPAT) network to mutually enhance embeddings of each pair of KGs based on aligned entities and relations. PPAT network's mechanism guarantees the differential privacy for the paired embeddings of aligned entities and relations not leaked to each other. For example, in Fig. 1, *Mark Twain* and *Henry Rogers* are two of aligned entities that belong to both *University* and *Literature* KGs, they have different representations in two embedding spaces of two KGs. We use PPAT network to separate their embeddings from both sides while we are still able to exploit their embeddings to improve the embedding quality for both *University* and *Literature* KGs.

In summary, we highlight the following characteristics of our FKGE [‡] framework. 1) FKGE framework is asynchronous and decentralized. Different from centralized client-based models, FKGE pairs up KGs from different domains with an adversarial network. 2) FKGE is scalable and compatible with many base embedding models. The asynchronous and decentralized setting leads to parallel computation between pairs of collaborators. Moreover, FKGE can serve as a meta-algorithm for existing KG embedding methods through a handshake protocol. 3) FKGE is privacy-preserving and guarantees no raw data leakage. FKGE's design requires no raw data transmission between collaborators, and transmitted generated embeddings are differentially private.

[‡]Code is available at https://github.com/HKUST-KnowComp/FKGE

2 RELATED WORK

As our work is closely related to federated learning and knowledge graph embedding, the related works are two-fold.

2.1 Federated Learning

Federated learning allows multiple data owners to collaborate on building models without compromising on data privacy. Google first proposed a federated learning framework between the cloud server and edge devices to train on edges and update global model on the cloud [21, 22, 29]. Yang et al. [51] further extended the definition from edge devices to general data owners, so that the collaborations between separate databases were taken into scenarios of federated learning. For graph structured data, Lalitha et al. [23] considered using federated learning in a peer to peer manner. FedE [9] exploited federated learning over a KG through centralized aggregation for the link prediction task. However, both of them handled one single graph by either treating each node to be a computing cell or distributing triplets in a KG into different servers and performed the Federated Averaging algorithm introduced by Google [29], and thus cannot collaboratively train multiple KGs together.

In terms of privacy preserving mechanisms, most works used homomorphic encryption (HE) [33], secure multi-party computation (MPC) [30], and differential privacy (DP) [12, 15, 17] to improve security. There are also serval extensions which combine or improve the above frameworks. For example, Lyu et al. [27] used DP, HE, and Blockchain technologies to build a decentralized fair and privacy-preserving deep learning framework. Xu et al. [48] designed a verifiable federated learning framework based on the homomorphic hash function and pseudorandom. Our proposed FKGE framework is targeted for multiple KGs with millions of entities and HE is inefficient for our task. DP addresses individuals' privacy concerns while aggregating different databases so that specific users' information can not be inferred through federation. Therefore, DP algorithm is implemented for our FKGE.

2.2 Knowledge Graph Embedding

KG embedding plays an important role in knowledge base inferences and downstream applications [2, 24, 31, 44]. Popular models are mostly based on a translational model, where a head is translated to a tail through a relational embedding [4], or a bilinear model, where a bilinear matrix is used to combine the head and tail to form a loss function [32]. Recently, there have been many extensions based on these two models [19, 25, 45, 47]. More recently, these models are extended to the complex space instead of using the Euclidean space [38, 42], which can model many important properties of relations such as symmetry, inversion, and composition.

When there are multiple KGs, embedding based entity alignment can be performed. Such entity alignment task usually assumes that different KGs are partially aligned and tries to predict more aligned entities. For example, the entity alignment can be based on cross-lingual KGs [8, 35, 46, 56], KGs with multi-view entity-related information [5, 7, 54], and KGs in similar domains with significant entity overlaps [39, 41, 50, 55]. After joint learning, embeddings for entity alignment are usually aligned in a unified space so the vectors can be used to find nearest entities in other KGs. Note that our FKGE framework is different from above KG alignment

problems. Instead of predicting the potential aligned entities using given ones, we aim to improve individual KG embeddings based on provided aligned entities. In FKGE, after joint training, each KG still does not know other KG's embedding space, but embeddings in each KG are all improved for better downstream tasks such as node classification or link prediction. This is guaranteed by the differential privacy mechanism that we introduced in our PPAT network: when training each pair of embedding sets for the aligned entities, they cannot leak a single embedding since inclusion and exclusion of a particular embedding will not affect the out-come distribution very much. This also allows us to use different base KG embedding models for different KGs.

3 FEDERATED KNOWLEDGE GRAPHS REPRESENTATION LEARNING

In this section, we present detailed descriptions of FKGE. We first give the problem definition and sketch intuitive solution in Section 3.1. Then we introduce more details of our privacy-preserving model in Section 3.2. Finally, we give a comprehensive explanation of our federated training mechanism in Section 3.3.

3.1 Problem Formulation and Proposed Solution

We define the set of knowledge graphs (KGs) from separate owners as $\mathcal{KG} = \{g_1, g_2, g_3, ..., g_N\}$, where N is the total number of KGs. Every element in $\mathcal{KG}$ locates in different databases and cannot access other KGs' databases. Let $g_k = \{\mathcal{E}_k, \mathcal{R}_k, \mathcal{T}_k\}(1 \le k \le N)$ denote the k-th KG in $\mathcal{KG}$. Each triple $(h, r, t) \in \mathcal{T}_k$ (the set of triples of g_k) is composed of a head entity $h \in \mathcal{E}_k$ (the set of entities of g_k), a tail entity $t \in \mathcal{E}_k$ and a relation $r \in \mathcal{R}_k$ (the set of relations of g_k). For any pair of KGs (g_i, g_j) in $\mathcal{KG}$, we assume that both aligned entities $\mathcal{E}_i \cap \mathcal{E}_j$ and relations $\mathcal{R}_i \cap \mathcal{R}_j$ are given which can be done by a secure hash[§]. Our goal is to exploit aligned entities and relations to further improve all embeddings of any individual KG. To aid discussion, Tab. 1 depicts the notations used throughout the paper. Each KG owner g_i trains its own embeddings of entities $\mathcal{E}_i$ and relations $\mathcal{R}_i$ locally. Based on the trained embeddings, FKGE aggregates the embeddings of both aligned entities and relations from paired KGs, and then updates embeddings in a federated manner. For aligned entities and relations from any pair of KGs, e.g., (g_i, g_j), FKGE includes a secure pipeline that can refine the embeddings of $\mathcal{E}_i \cap \mathcal{E}_j$ and $\mathcal{R}_i \cap \mathcal{R}_j$ and further improve embeddings of $\mathcal{E}_i \cup \mathcal{R}_i$ and $\mathcal{E}_j \cup \mathcal{R}_j$ individually. Moreover, FKGE proposes a federated training mechanism to improve all the parties jointly via broadcasting. If g_i or g_j gains improvement from the refined embeddings, it will broadcast signals to other KGs to further improve overall results. Otherwise, it will backtrack to original embeddings before federation. As an example shown in Fig. 2, in the beginning, g_1, g_2, g_3 train their embedding locally. During first federation, they form 3 pairs of KGs: (g_1, g_3), (g_2, g_1) and (g_3, g_2). After first federation, g_1 and g_2 gain improvement for overall embeddings. g_3's training takes longer time and fails to improve its embedding, therefore g_3 backtracks to initial embedding. During second federation, g_1 and g_2 pair up as (g_2, g_1) and (g_1, g_2) and only g_1 gains improvements. g_2 backtracks to previous embedding. Since g_3 is still on the training process, it will not join second federation and will go to sleep state if no available KG exists. For third federation, g_1 finishes its training and broadcasts g_3 to wake up. Then they form (g_1, g_3), (g_1, g_2) and (g_4, g_1) pairs for federation. The whole training procedure continues until no more improvement for all the KGs.

[§]https://csrc.nist.gov/publications/detail/fips/180/4/final

Table 1: Glossary of Notations.

Symbol	Definition
$\mathcal{KG}$, N, g_k	Knowledge graph set, its size, k-th knowledge graph
$\mathcal{E}_k$, $\mathcal{R}_k$, $\mathcal{T}_k$	The set of entities, relations, triples of g_k
X, Y	Embedding of aligned entities and relations of client, host
n	Total number of training samples
d	Embedding dimension of entity and relation
$\mathcal{N}(X)$	Raw embeddings of neighbor entities and joining relations of X
S, T_i	Student discriminator, i-th teacher discriminator
θ_S , θ_T^i	Parameters of student, i-th teacher discriminator
G, θ_G	Generator, parameters of generator that includes the mapping matrix W
$\|T\|$	Total number of teacher discriminators
λ	Noise (scale) of Laplace random variable
n_0, n_1	Total vote number for 0,1 of teacher discriminators
l	A chosen positive integer for moment

3.2 Privacy-Preserving Adversarial Model

Given two KGs (g_i, g_j) with aligned entities $\mathcal{E}_i \cap \mathcal{E}_j$ and relations $\mathcal{R}_i \cap \mathcal{R}_j$, FKGE exploits a GAN [16] structure to unify the embeddings of aligned entities and relations, where we borrow the idea of MUSE [10] which used a translational mapping matrix W in the generator to align two manifolds in the GAN's intermediate vector space. The generator and the discriminator can locate in different sides for a pair of KGs. More specifically, we may put the generator in g_i and the discriminator in g_j. The generator translates aligned entities' embeddings from g_i to g_j and the discriminator distinguishes between synthesized embeddings of the generator and ground truth embeddings in g_j. After GAN training, the synthesized embeddings are able to learn features from both KGs and therefore can replace original embeddings of $\mathcal{E}_i \cap \mathcal{E}_j$ and $\mathcal{R}_i \cap \mathcal{R}_j$ in g_i and g_j as refined and unified embeddings. It is sufficient for GAN training that only the generated outputs and gradients are transmitted between g_i and g_j without revealing raw data. However, even for generated embeddings, there are still privacy concerns for reconstruction attacks. It is possible that neural models may memorize inputs and reconstruct inputs from corresponding outputs [6]. To further address the privacy issue, we introduce differential privacy to privatize generated embeddings. Differential privacy provides a strong guarantee for protecting any single embedding in the generator outputs since inclusion and exclusion of a particular embedding will not affect the outcome distribution very much. Definitions 3.1 and 3.2 give the formal definition of differential privacy.

Definition 3.1 (Neighboring Dataset)**.** Two datasets D, D' are neighboring if

$$\exists x \in D \; s.t. \; D - \{x\} = D'. \tag{1}$$

Definition 3.2 (Differential Privacy)**.** A randomized *algorithm mechanism* M with domain D and range R satisfies (ϵ, δ)-*differential*

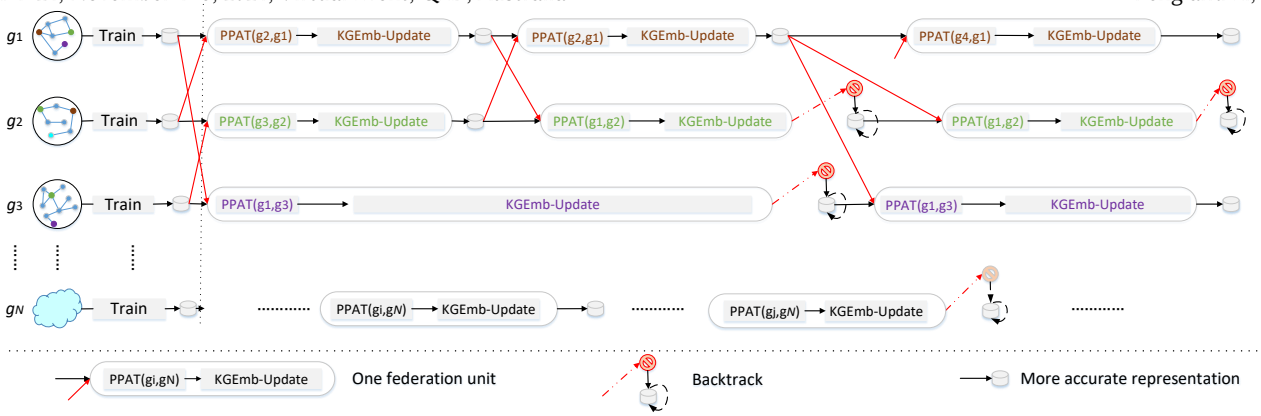

Figure 2: An example of whole training procedure. The order of training is not fixed and depends on individual KG owner's computation power and willingness to cooperate. "Train" indicates training KGE model locally. "PPAT(g_i, g_j)" denotes the PPAT embeddings between g_i and g_j where the generator locates in client g_i and the discriminator lies in host g_j. "KGEmb-Update" updates aligned embeddings with generated output of PPAT network and retrains all the embeddings as "Train". "Backtrack" happens if evaluation result after "KGEmb-Update" is not improved, then newly trained embeddings are abandoned and previous embeddings are kept. Otherwise the result is backtracked.

privacy if for any two neighboring datasets D, D' and for any subsets of output $O \subseteq R$:

$$Pr[M(D) \in O] \leq e^{\epsilon} Pr[M(D') \in O] + \delta. \quad (2)$$

The ϵ corresponds to privacy budget. Smaller ϵ asserts better privacy protection and lower model utility since algorithm outputs of neighboring datasets are similar. δ is the probability of information accidentally being leaked. Relying on above definitions, PATE-GAN [53] proposed a revised GAN structure to generate differentially private generator outputs by applying PATE mechanism [34] with teacher and student discriminators. Previously, PATE-GAN [53] was tested on simple image classification tasks, where all images were in one domain. Putting the translational mapping matrix W in the generator, we are able to apply PATE-GAN structure for the complicated cross-domain KGs.

Based on above intuitions, we implement our privacy-preserving adversarial translation (PPAT) network.

3.2.1 *Model Architecture and Loss Formulation*

For any pair of (g_i, g_j), the architecture of our PPAT network is illustrated in Fig. 3. PPAT network exploits GAN structure to generate differentially private synthetic embedding with high utility. It replaces the original GAN discriminator with multiple teacher discriminators and one student discriminator to achieve differential privacy of generated embeddings. The generator G with parameters θ_G (which also is the translational mapping matrix W in Fig. 3) locates in g_i's database while the student discriminator S with parameters θ_S and multiple teacher discriminators $T = \{T_1, T_2, ..., T_{|T|}\}$ with parameters $\theta_T^1, \theta_T^2, ..., \theta_T^{|T|}$ lie in g_j's database. We denote g_j as host and g_i as client, since the host is responsible for the generator's and all discriminators' loss back-propagation calculations, while the client only transmits its generated embeddings and receives gradients to update its generator parameters. We use $X = \{x_1, x_2, ..., x_n\}$ to denote embeddings for $\mathcal{E}_i \cap \mathcal{E}_j$ and $\mathcal{R}_i \cap \mathcal{R}_j$ in g_i, and $Y = \{y_1, y_2, ..., y_n\}$ to denote embeddings for $\mathcal{E}_i \cap \mathcal{E}_j$ and $\mathcal{R}_i \cap \mathcal{R}_j$ in g_j.

The objective of the generator G is to generate adversarial samples by making $G(X)$ and Y similar so that the student discriminator

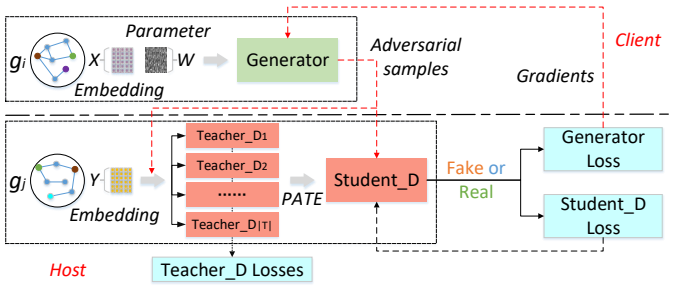

Figure 3: The architecture of PPAT network. The host and client are separated: only adversarial samples and gradients from the host are exchanged, all raw data and embeddings are private. Here "D" in host is short for discriminator.

S cannot distinguish them. Eq. (3) presents the objective function of the generator loss:

$$L_G(\theta_G; S) = \frac{1}{n} \sum_{m=1}^{n} \log(1 - S(G(x_m); \theta_G)), \quad (3)$$

where $G(X) = WX$, and S is the student discriminator parameterized by θ_S, which takes embeddings of both $G(X)$ and Y as an input under the CSLS metric used by MUSE [10].

The learning objective of teacher discriminators is the same as the original discriminator that distinguishes between fake samples $G(X)$ and real samples Y. The only difference is that teacher discriminators are trained on disjointly partitioned data. Teachers' losses are formulated as Eq. (4):

$$L_T^i(\theta_T^i; G) = -[\sum_{m=1}^{n} \log(1 - T_i(G(x_m); \theta_T^i)) + \sum_{y_k \in D_i} \log(T_i(y_k; \theta_T^i))], \quad (4)$$

where D_i is partitioned subset consisted of the dataset X and Y for T_i such that $|D_i| = \frac{n}{|T|}$ and there is no overlap between different subsets.

The learning objective of the student discriminator S is to classify generated samples given aggregated noisy labels. More specifically, teacher discriminators' predictions together with randomly injected

Laplace noises will determine the labels for the student discriminator S. Eqs. (5) and (6) illustrate the PATE mechanism:

$$PATE_\lambda(x) = \arg\max_{j\in\{0,1\}}(n_j(x) + V_j), \tag{5}$$

where V_0, V_1 are i.i.d. $Lap(\frac{1}{\lambda})$ random variables that introduce noises to teachers' votes. $n_j(x)$ denotes the number of teachers that predict class j for input x:

$$n_j(x) = |\{T_i : T_i(x) = j\}| \quad \text{for } j = 0, 1. \tag{6}$$

Then the student discriminator exploits generated samples with voted noisy labels (shown in Eq. (5)) to train itself in the host database. The student loss function is formulated as:

$$L_S(\theta_S; T, G) = \frac{1}{n}\sum_{i=1}^{n}[\gamma_i \log S(G(x_i); \theta_S) + (1-\gamma_i)\log(1 - S(G(x_i); \theta_S))], \tag{7}$$

where $\gamma_i = PATE_\lambda(x_i)$, the noisily aggregated label voted by teacher discriminators.

Besides the generated embedding $G(X)$, after the training process of PPAT network is stable, the client generates embeddings for the adjacent entities of the aligned entities in g_i and the joining relations as types of virtual entities and relations that are added into the host. We use $\mathcal{N}(X)$ to denote the raw embeddings of adjacent entities and joining relations, and the generated embeddings of virtual entities and relations can be denoted as $G(\mathcal{N}(X))$. These virtual entities and relations are only used for KGE, and will be removed before responding to other hosts.

3.2.2 Privacy Analysis and Parameter Estimation

For the PPAT network, our dataset X to feed the generator is embedding of $\mathcal{E}_i \cap \mathcal{E}_j$ and $\mathcal{R}_i \cap \mathcal{R}_j$ in g_i, the neighboring dataset is defined by excluding a particular embedding from X. We use X' to denote the neighboring dataset such that there is an embedding $x_i \in X$ and $X - \{x_i\} = X'$. The algorithm to perform parameter estimation is shown in Algo. 2. The raw embeddings of aligned entities and relations in X are fed to the generator to generate adversarial samples $G(X)$ which are later on transmitted to all teacher discriminators T of the host. By adding Laplace noise to voted results of teacher discriminators, differential privacy requirement is satisfied [34]. The student discriminator S is then trained by the synthesized embeddings with aggregated labels that contains 0 or 1 voted by teacher discriminators. The Post-Processing Theorem [13] which states the composition of a data-independent mapping f with an (ϵ, δ) - differentially private algorithm M is also (ϵ, δ) - *differentially private*. By Post-Processing Theorem, the student discriminator S is also differentially private since it is trained by differentially private labels. Moreover, the generator G is differentially private since G is trained by student discriminator S. Hence, the transmitted embeddings are synthesized and differentially private since they are outputs of the generator G.

During the training process, the host calculates the generator's and all discriminators' loss functions locally: gradients of student discriminator loss and teacher discriminator losses are used to update discriminators' parameters locally, while gradients of generator loss are sent back to the generator to update its parameters. Thus, neither of g_i and g_j is able to access the embeddings or raw data of the other's. Therefore the raw data's privacy are protected for any participant of knowledge graph owners.

Algorithm 1: KGProcessor

Input : The Host KG g_j
Output: Best embedding E_b

// Start self iterative training to get the best knowledge graph embedding and the best score S_b
$E_b \leftarrow g_j$.train();
$S_b \leftarrow g_j$.test(E_b);
g_j.state ← Ready;
// Start federated learning, and Q is handshake signal Queue
Q.receive_signal(); // Receiving signal as Host
while $g_j.state$ = *Ready* **do**
  $U_p \leftarrow$ False; // Set the improvement identifier U_p to false
  while Q *is not empty* **do**
    g_j.state = Busy; // Start PPAT, and set the state to busy
    $E_t \leftarrow$ ActiveHandshake(g_j, Q.poll());
    g_j.aggregation(E_t);
    // Start g_j KGEmb-Update, after aggregating embedding
    $T_e \leftarrow g_j$.train();
    $T_s \leftarrow g_j$.test(T_e);
    // Backtrack function is the embodiment of our backtrack mechanism
    $U_p,\ S_b,\ E_b \leftarrow g_j$.backtrack($T_s, S_b, T_e, E_b$);
    g_j.state = Ready;
  if U_p = *False* ***and*** Q *is empty* **then**
    // Both following states are configured by KG owner or activated by handshake signal
    if $g_j.state$ = *Sleep* **then**
      g_j.sleep();
    if $g_j.state$ = *Ready* **then**
      // g_j begins self iterative training
      $T_e \leftarrow g_j$.train();
      $T_s \leftarrow g_j$.test(T_e);
      $U_p,\ S_b,\ E_b \leftarrow g_j$.backtrack($T_s, S_b, T_e, E_b$);
  if U_p = *True* ***and*** Q *is empty* **then**
    // Broadcast handshake signal to other KGs having aligned entities and relations
    g_j.send_handshake_signal();
return Best embedding E_b ;

As we incorporate the differential privacy mechanism in our PPAT network, following PATE-GAN's proof, we can also estimate ϵ in Def. 3.2 for PPAT. We formulate the upper bound of $\hat{\epsilon}$ in PPAT in Eq. (8):

$$\hat{\epsilon} = \min_l \frac{\alpha(l) + \log(\frac{1}{\delta})}{l}, \tag{8}$$

where $\alpha(l)$ is the moments accountant for l-th moment. Its upper bound is derived by Theorem 2 and Theorem 3 in PATE mechanism [34] as shown in Eq. (9):

$$\alpha(l) = \alpha(l) + \min\left\{2\lambda^2 l(l+1), \log\left((1-q)\left(\frac{1-q}{1-e^{2\lambda}q}\right)^l + qe^{2\lambda l}\right)\right\}, \tag{9}$$

where q is an intermediary value which is formulated in Eq.(10):

$$q = \frac{2 + \lambda|n_0 - n_1|}{4\exp(\lambda|n_0 - n_1|)}, \tag{10}$$

where n_0, n_1 denote the number of teachers' votes for 0 or 1 separately. More details about updating α and $\hat{\epsilon}$ are shown in the Algo. 2. By choosing δ and λ, the $\hat{\epsilon}$ can be calculated.

3.3 Federated Training

For multiple KGs, we construct PPAT networks between any pair of $(g_i, g_j) \in \mathcal{KG}$ where $\mathcal{E}_j \cap \mathcal{E}_i \neq \emptyset$ or $\mathcal{R}_j \cap \mathcal{R}_i \neq \emptyset$, and produce $2 \times \binom{N}{2}$ PPAT networks at most at the same time. For any pair of (g_i, g_j), at least one client and one host are required separately. Our asynchronous and decentralized setting allows individual KG owner to decide whether it should collaborate with other KGs. The collaboration process can be described as a handshake protocol. Any g_i has three states: *Ready*, *Busy*, and *Sleep*. *Ready* state indicates g_i having available computational resources and being active to pair up with other KGs. *Busy* state indicates that g_i does not have enough resources and will not respond to any collaboration request at the moment. Instead, collaborators will be put in a queue till g_i finishes its work and is ready for collaborations. *Sleep* state indicates that though g_i has the computational resources, it has not received any collaboration request yet. That is, If *Ready* state cannot find a partner, it will switch to *Sleep* state and wake up to *Ready* state after a certain time period or being notified by a collaboration request. A successful handshake process between g_i and g_j implies $state(g_i) \neq Busy, state(g_j) \neq Busy$, and at least one of them has *Ready* state. Algo. 1 describes how a KG refines its embeddings. The whole handshake mechanism is explained in the Algo. 2.

Algorithm 2: ActiveHandshake

Input : The caller or Host KG g_j, Client KG g_i, Parameter δ, Noise λ
Output : Translated embedding E_t

// Initialization
PPAT.initialize(g_j, g_i); $\alpha = 0$;
repeat
  $AdvS \leftarrow g_i$.generate(X); // Generate adversarial samples
  g_j.receive(g_i.send($AdvS$)); // Communication between processes
  // Train all teacher discriminators with partitioned $AdvS$
  $Probs \leftarrow g_j$.teacher_Ds(Y, $AdvS$);
  $vote \leftarrow \text{PATE}_\lambda(Probs)$;
  // Train student discriminator with noisily voted labels
  $Proba \leftarrow g_j$.student_D($vote$, $AdvS$);
  $L_G, L_S \leftarrow g_j$.loss_calculation($Proba$);
  $L_T \leftarrow g_j$.loss_calculation($Probs$);
  $grad\_G, grad\_T\_Ds, grad\_S\_D \leftarrow$ backpropagation(L_G, L_T, L_S);
  g_i.receive(g_j.send($grad\_G$));
  g_j.update_parameters(Θ_T, $grad\_T\_Ds$);
  g_j.update_parameters(Θ_S, $grad\_S\_D$);
  g_i.update_parameters(Θ_G, $grad\_G$);
  // Update moments
  $q \leftarrow \frac{2+\lambda|n_0-n_1|}{4\exp(\lambda|n_0-n_1|)}$;
  $\alpha(l) \leftarrow \alpha(l) + \min\left\{2\lambda^2 l(l+1), \log\left((1-q)\left(\frac{1-q}{1-e^{2\lambda}q}\right)^l + qe^{2\lambda l}\right)\right\}$;
until PPAT *training converged*;
$\hat{\epsilon} \leftarrow \min_l \frac{\alpha(l)+\log(\frac{1}{\delta})}{l}$;
$E_t \leftarrow g_j$.receive(g_i.generate(X), g_i.generate($N(X)$));
return Translated embedding E_t;

4 EXPERIMENTS

In this section, we present extensive experiments to evaluate the performance of the proposed FKGE framework. In Section 4.1, we will introduce basic experimental setup and metrics. Then we will use FKGE to conduct both triple classification and link prediction experiments in Section 4.2. Moreover, we will conduct ablation study to show the effectiveness of FKGE in Section 4.3. Finally, we will analyze time cost for our experiment in Section 4.4.

Table 2: Statistics of the Knowledge Graphs.

KGs	#Relation	#Entity	#Triple
Dbpedia	14,085	49,1078	1,373,644
Geonames	6	300,000	1,163,878
Yago	37	286,389	1,824,322
Geospecies	38	41,943	782,120
Poképédia	28	238,008	548,883
Sandrart	20	14,765	18,243
Hellenic	4	11,145	33,296
Lexvo	6	9,810	147,211
Tharawat	12	4,693	31,130
Whisky	11	642	1,339
World lift	10	357	1,192
Summation	14,257	1,398,830	5,915,596

Table 3: Statistics of Aligned Entities (AEs).

KGs	# AEs	KGs	# AEs	KGs
Geonames	118,939	Dbpedia	27	Poképédia
Yago	123,853	Dbpedia	133	Geospecies
Yago	53,553	Geonames	89	Geospecies
Sandrart	379	Dbpedia	41	Hellenic
Dbpedia	507	Lexvo	245	Geonames
Dbpedia	403	Tharawat	90	Geonames
Dbpedia	70	Whisky	39	Geonames
Dbpedia	25	World lift	18	Yago
Lexvo	77	Yago	266	Tharawat
Whisky	49	Yago	-	-

4.1 Experimental Settings

Dataset. We select 11 KGs at different scales from the Linked Data community[¶]. For the KGE, OpenKE framework [18] is used so that FKGE is compatible with various KGE models. For each KG, we count the numbers of relation, entity and triple, and divide the triples into train, valid, and test sets with ratio 90 : 5 : 5 according to the default setting in the OpenKE. Note that in order to reduce computational time of training and testing in KGE, we cut out some sparse entities and triples that are not relevant to aligned entities and triples from the original KGs. A summary statistics of these KGs is shown in Tab. 2. The Linked Data community provides aligned entities (AE) between different KGs in RDF files, statistics as recorded in Tab. 3.

Hyperparameter Setting. To simulate real-world asynchronous training without data leakage, we set each KG to one process and implement all comparative experiments on 11 independent processes using the same type of GPU devices with the same configurations. During a handshake process, we use pipeline communication between processes to transmit generated adversarial samples from client to host and gradients from host to client. In consideration of computational time and testing results, we set default dimension d of the embedding vector to $d = 100$, and step of testing to 1,000 epochs. For other essential parameters of KGE, we set the learning rate to 0.5 and batch size to 100 following the default setting of the OpenKE [18]. Negative samples are generated by corrupting either head or tail entities and ratios of negative samples and positive samples are 1:1. For the essential parameters of the PPAT network,

[¶]https://lod-cloud.net/

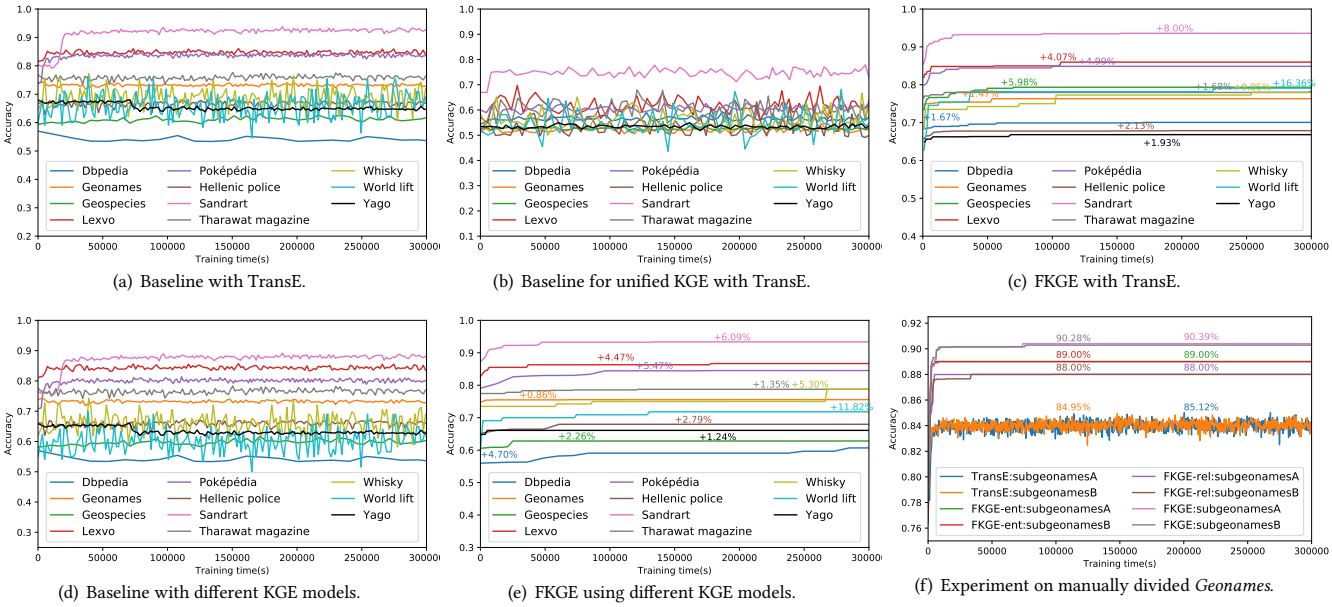

(a) Baseline with TransE. (b) Baseline for unified KGE with TransE. (c) FKGE with TransE.

(d) Baseline with different KGE models. (e) FKGE using different KGE models. (f) Experiment on manually divided *Geonames*.

Figure 4: Performance of Triple Classification.

we set batch size, teacher number, learning rate, momentum to 32, 4, 0.02, 0.9. We train for 1,000 epochs for each KG to get the initial best score and embeddings, and then activate the federal unit.

Privacy Setting. We set λ to 0.05 and δ to 10^{-5}. For each pair of KGs in federated training, there is an PPAT network with varied ϵ according to the input dataset X since total the number of queries of paired embeddings may change according to number of aligned entities. We estimate the largest upper bound $\hat{\epsilon}$ for all the ϵ and the bound is $\hat{\epsilon} = 2.73$. For each round of federated training, the max $\alpha(l)$ is 0.29 among all the ActiveHandshake in the Alg. 2, putting $\ln \frac{1}{\delta} = 11.5$ and $l = 9$ together, we obtain upper bound 2.73 for all ϵ following Eqs. (8 - 10).

Evaluation Metrics. The proposed FKGE framework is compatible with different KGE methods. We choose popular and simple translation-based models including TransE [4], TransH [45], TransR [25], and TransD [19] from the OpenKE [18] to evaluate the quality of the embeddings trained by different methods under two classical testing tasks of KGE: triple classification and link prediction. For triple classification, we apply the **Accuracy** as evaluation metric. For link prediction, we apply the widely used proportion of correct entities in top-1, 3 and 10 ranked entities (**Hit@1, 3 and 10**) and **Mean Rank** as evaluation metrics.

4.2 Evaluation

Here, we demonstrate the advantages of the FKGE based on experimental results on triple classification and link prediction tasks.

4.2.1 Triple Classification

We give the baseline accuracies of triple classification of the 11 KGs using TransE in Fig. 4(a). The performances of baseline are unstable: accuracies of *Yago* and *Dbpedia* are even reduced. To verify the performance of one unified structure of multiple KGs, we integrate 11 KGs into a united KG by merging aligned entities, and then test the performance of TransE on each KG independently. The accuracy of triple classification of the unified KG is shown in Fig. 4(b). Compared with the independent KG embedding in Fig. 4(a), the unified KG embedding even has generally decreased by 6.82% - 17.63%. Hence, integrating the embeddings of multiple KGs into one unified vector space does not help to obtain effective representation learning of KGs.

We apply the FKGE framework to the 11 KGs with TransE. The results of triple classification are shown in Fig. 4(c), where the marked improvements are compared with the results before training (at time 0). After the same training time, it can be observed that the accuracy of each KG increases. Specifically, compared with the baseline method in Fig. 4(a), the accuracy results (KG ordered as in Tab. 2) have been improved by 16.49%, 2.98%, 2.06%, 17.85%, 2.11%, 0.60%, 0.48%, 0.77%, 1.82%, 12.88% and 14.55% on triple classification task, respectively. The improvements in accuracy of the above 11 KGs benefit from the cross-knowledge embedding integration in the FKGE. Moreover, the continuous and steady improvements also show the effectiveness of the backtrack mechanism in the FKGE. Therefore, based on the FKGE framework and TransE, the 11 KGs have achieved consistent improvements in triple classification.

Not only TransE benefits from the PPAT network, but also other mainstream KGE models can be improved through the FKGE framework. We also randomly select KGE method from 4 popular translation family models, including TransR, TransE, TransD and TransH, for each KG. We give the results of the baseline using different translation family models in Fig. 4(d). After 300,000 seconds of training, the accuracy of the FKGE in the triple classification for the 11 KGs is shown in Fig. 4(e), where the marked improvements are also compared with the results before training (at time 0). Specifically, compared with their respective base methods, the 11 KGs have been improved by 7.08% (TransR), 2.23% (TransD), 1.33% (TransE), 1.32% (TransR), 1.77% (TransE), 0.38% (TransD), 0.57% (TransD), 1.92%

Methods	Independent-TransE			FKGE			Random-Independent-KGE			Multi-FKGE		
Metric	Hit@10	Hit@3	Hit@1	Hit@10	Hit@3	Hit@1	Hit@10	Hit@3	Hit@1	Hit@10	Hit@3	Hit@1
Dbpedia	23.29	12.88	5.12	**25.07**	**14.41**	**6.37**	5.46	2.51	1.10	**6.67**	**3.20**	**1.24**
Geonames	8.82	3.69	1.93	**9.65**	**4.88**	**2.12**	8.45	4.53	1.90	**8.85**	**4.97**	**2.14**
Yago	2.05	0.76	0.25	**2.59**	**0.88**	**0.29**	2.03	**0.75**	**0.24**	**2.36**	**0.75**	**0.24**
Geospecies	58.49	45.81	34.01	**60.97**	**46.95**	**35.03**	38.68	26.43	13.12	**40.92**	**28.04**	**14.38**
Poképédia	38.14	29.04	19.31	**45.58**	**35.48**	**24.90**	34.22	25.13	16.43	**42.12**	**32.14**	**22.65**
Sandrart	87.39	83.16	67.18	**88.65**	**84.97**	**72.14**	87.71	83.71	68.91	**87.99**	**84.22**	**69.69**
Hellenic	32.18	21.87	18.96	**33.00**	**22.87**	**19.35**	32.21	22.23	18.59	**32.82**	**22.59**	**19.44**
Lexvo	85.67	76.07	58.29	**87.35**	**77.74**	**62.90**	84.21	75.82	58.09	**85.72**	**76.99**	**59.76**
Tharawat	12.48	4.56	1.67	**13.45**	**5.26**	**2.19**	12.30	4.38	1.39	**12.55**	**5.21**	**1.77**
Whisky	28.78	15.15	9.84	**35.60**	**18.93**	**10.60**	28.78	18.93	12.87	**30.12**	**19.45**	**12.92**
World lift	45.76	24.57	7.62	**51.69**	**28.88**	**11.17**	18.64	8.47	1.69	**18.85**	**9.32**	**2.54**

Table 4: Evaluation results on link prediction (%). We show the best results with boldface.

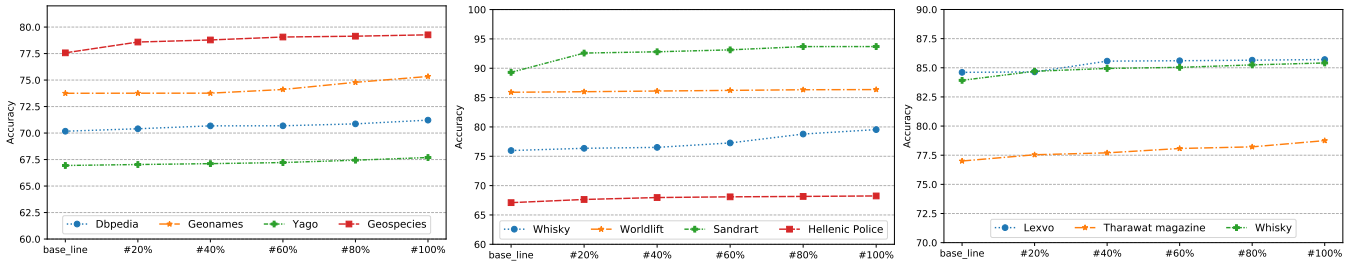

(a) The first part containing 4 knowledge graphs. (b) The second part containing 4 knowledge graphs. (c) The third part containing 3 knowledge graphs.

Figure 5: Evaluation results on triple classification with different sampling rates.

Sampling	#20%				#40%				#60%				#80%			
Metric	Mean Rank	Hit@10	Hit@3	Hit@1	Mean Rank	Hit@10	Hit@3	Hit@1	Mean Rank	Hit@10	Hit@3	Hit@1	Mean Rank	Hit@10	Hit@3	Hit@1
Geonames	63,080.34	8.93	3.86	1.93	62,421.95	9.17	4.47	1.99	61,276.67	9.35	4.49	2.04	59,786.04	9.65	4.57	2.04
Yago	30,638.51	2.05	0.76	0.25	28,838.64	2.235	0.79	0.26	26,563.40	2.35	0.81	0.26	26,421.51	2.55	0.82	0.27
Geospecies	415.47	59.31	45.97	34.81	409.06	59.70	45.99	35.08	400.73	60.17	46.38	35.12	397.69	60.28	46.72	35.30
Poképédia	7,463.42	38.37	30.01	20.01	7,450.85	43.25	30.50	21.61	7,445.54	44.09	33.08	23.39	7,434.55	45.49	33.43	24.19
Sandrart	56.23	87.61	83.38	67.96	52.39	88.38	83.43	69.51	51.71	88.65	84.34	70.38	47.39	88.98	84.37	72.10
Hellenic Police	1,886.24	32.29	22.06	19.03	1,853.55	32.47	22.41	19.08	1,823.84	32.63	22.78	19.11	1,751.37	32.98	22.86	19.27
Lexvo	43.64	86.01	76.59	58.43	44.96	86.14	76.86	59.68	43.43	86.70	77.11	60.47	41.57	86.77	77.62	61.30
Tharawat magazine	232.63	12.50	4.58	1.76	232.62	12.92	4.81	1.96	228.65	13.13	5.09	1.99	226.67	13.19	5.15	2.10
Whisky	49.47	29.39	15.98	9.91	49.85	31.14	16.24	10.15	44.47	31.85	17.06	10.25	43.12	33.35	17.64	10.58
World lift	22.98	46.80	25.47	7.85	22.82	49.71	25.60	7.99	22.74	50.06	28.62	8.02	22.11	50.15	28.87	10.30

Table 5: Evaluation results on link prediction tasks in Filter with different sampling ratio(%) of aligned entities.

(TransD), 2.42% (TransD), 13.64% (TransH), and 7.27% (TransR) on the triple classification, respectively. It confirms that the FKGE framework has the advantage of being compatible with different KGE methods.

4.2.2 *Link Prediction*

We compare the performance of link prediction with type constraint following the OpenKE in multiple scenarios. As shown in Tab. 4, we present the evaluation results of Hit@1, 3, and 10 in the Filter setting. The Filter means that those corrupted triples in test set and validation set are removed for the link prediction. The Independent-TransE means using the traditional TransE based KGE individually. The FKGE means using TransE and federated training to improve every KG collaboratively. Besides TransE, we also use other KGE methods for link prediction. The Random-Independent-KGE means that each KG is applied with a KGE randomly from translation-based models and is training independently. The Multi-FKGE means that each KG randomly chooses a KGE from the translation-based models and also employs the FKGE for further training. Here, we keep the same base KGE methods selected in Fig. 4(e). Compared with the baseline methods, the TransE based FKGE has increased at most by 7.44%, 6.44%, and 5.59% in terms of Hit@10, 3, and 1, and Multi-FKGE has increased at most by 7.90%, 7.01%, and 5.87% in terms of Hit@10, 3, and 1. For example, benefited from the TransE based FKGE, the *World lift* gains 5.93%, 4.31%, and 3.55% improvements in terms of Hit@10, 3, and 1 under the Filter. The above experiments under link prediction also demonstrate the effectiveness and adaptability of the proposed FKGE framework.

4.3 Ablation Study

Effectiveness of aligned entities and relations. Here we first consider whether the inclusion of aligned entities and relations is beneficial to FKGE's performance. Because the existing KGs provide no aligned relations, we manually divide *Geonames* into two subsets named *SubgeonamesA* and *SubgeonamesB* of the same size to verify the contribution of aligned entities and relations. We treat relations similarly as entities and simply put them together for model training. Note that the size of divided KGs shrinks on both entity and relation, and are tested on different sets of triples. Therefore the accuracies may differ from baselines in Fig. 4(a). As shown

in Fig. 4(f), the blue and orange lines represent the accuracy of the two subsets with TransE only, namely *TransE: subgeonamesA* and *TransE: subgeonamesB*, in the triple classification; the green and red lines represent the accuracy by using only the aligned entities based on FKGE with TransE, namely *FKGE-ent: subgeonamesA* and *FKGE-ent: subgeonamesB*; the purple and brown lines represent the accuracy by using only the aligned relations based on FKGE with TransE, namely *FKGE-rel: subgeonamesA* and *FKGE-rel: subgeonamesB*; the pink and grey lines represent the accuracy by using both aligned entities and relations based on FKGE with TransE, namely *FKGE: subgeonamesA* and *FKGE: subgeonamesB*, respectively. Obviously, compared with the *TransE: subgeonamesA* and *TransE: subgeonamesB*, *FKGE: subgeonamesA* and *FKGE: subgeonamesB* achieve significant improvements of 5.33% and 5.27% in triple classification, respectively. Besides that, both *FKGE-ent: subgeonamesA* and *FKGE-ent: subgeonamesB* achieve 1.00% improvement over the relation alignment based *FKGE-rel: subgeonamesA* and *FKGE-rel: subgeonamesB*. In general, for the proposed FKGE framework, aligned entities and relations both contribute to the improvement of knowledge graph representations, and the refined embeddings show that PPAT network is also effective.

Size of aligned entities. In order to measure the impact of the scale of aligned entities involved in the FKGE, we compare the performance of different numbers of the aligned entities and different numbers of KGs under triple classification. We randomly sample 20%, 40%, 60% and 80% of the aligned entities for fusing them through the PPAT network, respectively. The results on triple classification with different sampling rates are shown in Fig. 5. Intuitively, the more entities' embeddings that can be fused, the more obvious the improvements are. For example, by sampling 20%, 40%, 60% and 80% and 100% sampling rates, the *World lift* gains 3.52%, 4.05%, 5.19%, 6.44% and 7.16% improvements on the triple classification task, respectively. Besides triple classification, Tab. 5 presents the results of link prediction by using TransE based FKGE, in terms of Mean Rank, Hit@10, Hit@3 and Hit@1, with different sampling ratio in the PPAT network. In the Filter metric, by using the FKGE of 20%, 40%, 60%, 80% and 100% (Presented in Tab. 4) sampling rates, the *Poképédia* gains up to 0.23%, 5.11%, 5.95%, 7.35% and 7.44% improvements in terms of Hit@10 on the link prediction task, respectively. In addition to the metric of Hit@10, most of other KGs have also achieved improvements in Mean Rank, Hit@3 and Hit@1. The above experimental results have once again proved the scalability and effectiveness of the proposed FKGE framework.

4.4 Time Consumption

To analyze the time cost for individual KG and demonstrate scalability of FKGE. During federation, we evaluate the time cost for different ratios of aligned entities between *Geonames* and *Dbpedia*. We keep the exact experimental setup as Section 4.1 except that only *Geonames* and *Dbpedia* are used for the experiment. It's easy to verify that KGEmb-Update and PPAT training pairwisely constitute major time costs for individual KGs. Both of the KGs share similar results, and we show time consumption of *Geonames* in Fig. 6. For each ratio, we run the experiments for 10 times and average the corresponding time consumption. It can be observed that KGEmb-Update usually costs much more time than PPAT network and its cost remains around 4,000s as number of aligned entities increases.

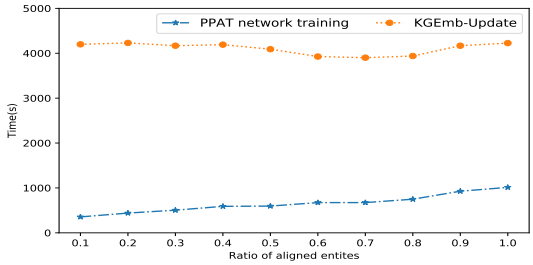

Figure 6: Time cost experiment on *Geonames* (in seconds).

For PPAT training, the cost increases roughly linearly from 350s to 1,000s as number of aligned entities increases, which constitutes around 8% ∼ 20% among overall training time. In general, for each KG in FKGE, overall time cost is acceptable compared with gained improvements. The linear time cost of the PPAT network indicates FKGE is scalable for aligned entities. Moreover, for transmission between client and host, the client sends translated embeddings of size (batch size, d) and host transmits gradients for client of size (d, d). Using our experimental settings batch size = 32, $d = 100$ and 64 bit for double precision, total communication cost for a batch training of the PPAT network is at most 0.845 Mb. Therefore, it is practical for FKGE to train PPAT networks online like a peer-to-peer network.

5 CONCLUSION

In this paper, we present a well-performed and decentralized federated knowledge graph embedding framework. First, we exploit a privacy-preserving adversarial model between pairs of knowledge graphs to merge identical entities or relations of different domains, and guarantee no raw data leakage. Then, we implement an asynchronous and pairwise federated representation learning framework on knowledge graphs. We conduct extensive experiments to evaluate the FKGE on 11 knowledge graphs, demonstrating significant and consistent improvements in model quality through performance on triple classification and link prediction tasks. The ablation study also indicates that merging both aligned entities and relations are beneficial for overall improvement. Besides, the time consumption demonstrates the scalability of FKGE. In the future, we plan to integrate more advanced knowledge graph representation learning models and extend our federated learning framework to other complex knowledge data exchanging scenarios.

ACKNOWLEDGMENT

The authors of this paper were supported by the NSFC through grants U20B2053 and 62002007, the S&T Program of Hebei through grant 20310101D, the Fundamental Research Funds for the Central Universities, the National Key Research and Development Program of China (208AAA0101100), the RIF (R6020-19 and R6021-20) and the GRF (16211520) from RGC of Hong Kong, the MHKJFS (MHP/001/19) from ITC of Hong Kong. We thank Linfeng Du, Qi Teng, and Lichao Sun for useful comments and discussions, and also acknowledge the support from the HKUST-WeBank Joint Lab at HKUST. For any correspondence, please refer to Hao Peng.

REFERENCES

[1] Dirk Ahlers. 2013. Assessment of the accuracy of GeoNames gazetteer data. In *Proceedings of the 7th workshop on geographic information retrieval*. 74–81.

[2] Federico Bianchi, Gaetano Rossiello, Luca Costabello, Matteo Palmonari, and Pasquale Minervini. 2020. Knowledge Graph Embeddings and Explainable AI. *arXiv preprint arXiv:2004.14843* (2020).

[3] Kurt D. Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of SIGMOD*. 1247–1250.

[4] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Proceedings of NIPS*. 2787–2795.

[5] Yixin Cao, Zhiyuan Liu, Chengjiang Li, Zhiyuan Liu, Juanzi Li, and Tat-Seng Chua. 2019. Multi-Channel Graph Neural Network for Entity Alignment. In *ACL (1)*. 1452–1461.

[6] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2020. Extracting Training Data from Large Language Models. *arXiv preprint arXiv:2012.07805* (2020).

[7] Muhao Chen, Yingtao Tian, Kai-Wei Chang, Steven Skiena, and Carlo Zaniolo. 2018. Co-training Embeddings of Knowledge Graphs and Entity Descriptions for Cross-lingual Entity Alignment. In *IJCAI*. 3998–4004.

[8] Muhao Chen, Yingtao Tian, Mohan Yang, and Carlo Zaniolo. 2017. Multilingual Knowledge Graph Embeddings for Cross-lingual Knowledge Alignment. In *Proceedings of IJCAI*. 1511–1517.

[9] Mingyang Chen, Wen Zhang, Zonggang Yuan, Yantao Jia, and Huajun Chen. 2020. FedE: Embedding Knowledge Graphs in Federated Setting. *arXiv preprint arXiv:2010.12882* (2020).

[10] Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word Translation Without Parallel Data. *Proceedings of ICLR*.

[11] Gerard De Melo. 2015. Lexvo. org: Language-related information for the linguistic linked data cloud. *Semantic Web* 6, 4 (2015), 393–400.

[12] Cynthia Dwork. 2008. Differential Privacy: A Survey of Results. In *TAMC (Lecture Notes in Computer Science)*, Vol. 4978. 1–19.

[13] C. Dwork and A. Roth. 2014. The Algorithmic Foundations of Differential Privacy. In *The Algorithmic Foundations of Differential Privacy*. 19–20.

[14] Lisa Ehrlinger and Wolfram Wöß. 2016. Towards a Definition of Knowledge Graphs. *SEMANTiCS (Posters, Demos, SuCCESS)* 48 (2016), 1–4.

[15] Robin C. Geyer, Tassilo Klein, and Moin Nabi. 2017. Differentially Private Federated Learning: A Client Level Perspective. In *NIPS Workshop on Machine Learning on the Phone and other Consumer Devices*.

[16] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Proceedings of NIPS*. 2672–2680.

[17] Andreas Grammenos, Rodrigo Mendoza Smith, Jon Crowcroft, and Cecilia Mascolo. 2020. Federated Principal Component Analysis. In *Proceedings of NeurIPS*.

[18] Xu Han, Shulin Cao, Xin Lv, Yankai Lin, Zhiyuan Liu, Maosong Sun, and Juanzi Li. 2018. OpenKE: An Open Toolkit for Knowledge Embedding. In *Proceedings of the EMNLP*. 139–144.

[19] Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of ACL-IJNLP*. 687–696.

[20] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and S Yu Philip. 2021. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems* (2021).

[21] Jakub Konecný, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. 2016. Federated Optimization: Distributed Machine Learning for On-Device Intelligence. *arXiv preprint arXiv:1610.02527* (2016).

[22] Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtarik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated Learning: Strategies for Improving Communication Efficiency. In *NIPS Workshop on Private Multi-Party Machine Learning*.

[23] Anusha Lalitha, Osman Cihan Kilinc, Tara Javidi, and Farinaz Koushanfar. 2019. Peer-to-peer Federated Learning on Graphs. *arXiv preprint arXiv:1901.11173* (2019).

[24] Chen Li, Xutan Peng, Shanghang Zhang, Hao Peng, S Yu Philip, Min He, Linfeng Du, and Lihong Wang. 2020. Modeling relation paths for knowledge base completion via joint adversarial training. *Knowledge-Based Systems* 201 (2020).

[25] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Proceedings of AAAI*. 2181–2187.

[26] Ye Liu, Yao Wan, Lifang He, Hao Peng, and S Yu Philip. 2021. KG-BART: Knowledge Graph-Augmented BART for Generative Commonsense Reasoning. In *Proceedings of the AAAI*, Vol. 35. 6418–6425.

[27] Lingjuan Lyu, Jiangshan Yu, Karthik Nandakumar, Yitong Li, Xingjun Ma, Jiong Jin, Han Yu, and Kee Siong Ng. 2020. Towards fair and privacy-preserving federated deep models. *IEEE TPDS* 31, 11 (2020), 2524–2541.

[28] Qianren Mao, Xi Li, Hao Peng, Jianxin Li, Dongxiao He, Shu Guo, Min He, and Lihong Wang. 2021. Event prediction based on evolutionary event ontology knowledge. *Future Generation Computer Systems* 115 (2021), 76–89.

[29] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the AISTATS*. 1273–1282.

[30] P. Mohassel and Y. Zhang. 2017. SecureML: A System for Scalable Privacy-Preserving Machine Learning. In *Proceedings of S&P*. 19–38.

[31] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2016. A Review of Relational Machine Learning for Knowledge Graphs. *Proc. IEEE* 104, 1 (2016), 11–33.

[32] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *ICML*. 809–816.

[33] Pascal Paillier. 1999. Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In *Proceedings of EUROCRYPT '99*, Vol. 1592. 223–238.

[34] Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian J. Goodfellow, and Kunal Talwar. 2017. Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data. In *Proceedings of ICLR*.

[35] Shichao Pei, Lu Yu, Guoxian Yu, and Xiangliang Zhang. 2020. REA: Robust Cross-lingual Entity Alignment Between Knowledge Graphs. In *KDD*. 2175–2184.

[36] Amit P. Sheth and James A. Larson. 1990. Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases. *ACM Comput. Surv.* 22, 3 (1990), 183–236.

[37] Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of WWW*. 697–706.

[38] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. *arXiv preprint arXiv:1902.10197* (2019).

[39] Zequn Sun, Wei Hu, Qingheng Zhang, and Yuzhong Qu. 2018. Bootstrapping Entity Alignment with Knowledge Graph Embedding. In *IJCAI*. 4396–4402.

[40] Zequn Sun, Chengming Wang, Wei Hu, Muhao Chen, Jian Dai, Wei Zhang, and Yuzhong Qu. 2020. Knowledge graph alignment network with gated multi-hop neighborhood aggregation. In *Proceedings of AAAI*. 222–229.

[41] Bayu Distiawan Trisedya, Jianzhong Qi, and Rui Zhang. 2019. Entity Alignment between Knowledge Graphs Using Attribute Embeddings. In *AAAI*. 297–304.

[42] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *Proceedings of ICML*. 2071–2080.

[43] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (2014), 78–85.

[44] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge graph embedding: A survey of approaches and applications. *TKDE* 29, 12 (2017), 2724–2743.

[45] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of AAAI*.

[46] Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, Rui Yan, and Dongyan Zhao. 2019. Relation-Aware Entity Alignment for Heterogeneous Knowledge Graphs. In *IJCAI*. 5278–5284.

[47] Han Xiao, Minlie Huang, and Xiaoyan Zhu. 2016. From one point to a manifold: knowledge graph embedding for precise link prediction. In *Proceedings of IJCAI*. 1315–1321.

[48] Guowen Xu, Hongwei Li, Sen Liu, Kan Yang, and Xiaodong Lin. 2019. Verifynet: Secure and verifiable federated learning. *IEEE TIFS* 15 (2019), 911–926.

[49] Kun Xu, Linfeng Song, Yansong Feng, Yan Song, and Dong Yu. 2020. Coordinated Reasoning for Cross-Lingual Knowledge Graph Alignment. In *Proceedings of AAAI*.

[50] Kai Yang, Shaoqin Liu, Junfeng Zhao, Yasha Wang, and Bing Xie. 2020. COTSAE: CO-Training of Structure and Attribute Embeddings for Entity Alignment. In *AAAI*. 3025–3032.

[51] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated Machine Learning: Concept and Applications. *ACM TIST* 10, 2 (2019), 12:1–12:19.

[52] Yiying Yang, Xi Yin, Haiqin Yang, Xingjian Fei, Hao Peng, Kaijie Zhou, Kunfeng Lai, and Jianping Shen. 2021. KGSynNet: A Novel Entity Synonyms Discovery Framework with Knowledge Graph. In *Proceedings of DASFAA*. Springer International Publishing, Cham, 174–190.

[53] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. 2019. PATE-GAN: Generating Synthetic Data with Differential Privacy Guarantees. In *Proceedings of ICLR*.

[54] Qingheng Zhang, Zequn Sun, Wei Hu, Muhao Chen, Lingbing Guo, and Yuzhong Qu. 2019. Multi-view knowledge graph embedding for entity alignment. In *Proceedings of IJCAI*. AAAI Press, 5429–5435.

[55] Hao Zhu, Ruobing Xie, Zhiyuan Liu, and Maosong Sun. 2017. Iterative Entity Alignment via Joint Knowledge Embeddings. In *Proceedings of IJCAI*. 4258–4264.

[56] Qiannan Zhu, Xiaofei Zhou, Jia Wu, Jianlong Tan, and Li Guo. 2019. Neighborhood-Aware Attentional Representation for Multilingual Knowledge Graphs. In *IJCAI*. 1943–1949.