

LinkNBed: Multi-Graph Representation Learning with Entity Linkage

Rakshit Trivedi *
College of Computing
Georgia Tech

Bunyamin Sisman
Amazon.com

Jun Ma
Amazon.com

Christos Faloutsos
SCS, CMU
and Amazon.com

Hongyuan Zha
College of Computing
Georgia Tech

Xin Luna Dong
Amazon.com

Abstract

Knowledge graphs have emerged as an important model for studying complex multi-relational data. This has given rise to the construction of numerous large scale but incomplete knowledge graphs encoding information extracted from various resources. An effective and scalable approach to jointly learn over multiple graphs and eventually construct a unified graph is a crucial next step for the success of knowledge-based inference for many downstream applications. To this end, we propose **LinkNBed**, a deep relational learning framework that learns entity and relationship representations across multiple graphs. We identify **entity linkage** across graphs as a vital component to achieve our goal. We design a novel objective that leverage entity linkage and build an efficient multi-task training procedure. Experiments on link prediction and entity linkage demonstrate substantial improvements over the state-of-the-art relational learning approaches.

1 Introduction

Reasoning over multi-relational data is a key concept in Artificial Intelligence and knowledge graphs have appeared at the forefront as an effective tool to model such multi-relational data. Knowledge graphs have found increasing importance due to its wider range of important applications such as information retrieval (Dalton et al., 2014), natural language processing (Gabrilovich and Markovitch, 2009), recommender systems (Catherine and Cohen, 2016), question-answering (Cui et al., 2017)

and many more. This has led to the increased efforts in constructing numerous large-scale Knowledge Bases (e.g. Freebase (Bollacker et al., 2008), DBpedia (Auer et al., 2007), Google’s Knowledge graph (Dong et al., 2014), Yago (Suchanek et al., 2007) and NELL (Carlson et al., 2010)), that can cater to these applications, by representing information available on the web in relational format.

All knowledge graphs share common drawback of *incompleteness* and *sparsity* and hence most existing relational learning techniques focus on using observed triplets in an incomplete graph to infer unobserved triplets for that graph (Nickel et al., 2016a). Neural embedding techniques that learn vector space representations of entities and relationships have achieved remarkable success in this task. However, these techniques only focus on learning from a single graph. In addition to incompleteness property, these knowledge graphs also share a set of overlapping entities and relationships with varying information about them. This makes a compelling case to design a technique that can learn over multiple graphs and eventually aid in constructing a unified giant graph out of them. While research on learning representations over single graph has progressed rapidly in recent years (Nickel et al., 2011; Dong et al., 2014; Trouillon et al., 2016; Bordes et al., 2013; Xiao et al., 2016; Yang et al., 2015), there is a conspicuous lack of principled approach to tackle the unique challenges involved in learning across multiple graphs.

One approach to multi-graph representation learning could be to first solve graph alignment problem to merge the graphs and then use existing relational learning methods on merged graph. Unfortunately, graph alignment is an important but still unsolved problem and there exist several techniques addressing its challenges (Liu and Yang, 2016; Pershina et al., 2015; Koutra et al., 2013; Buneman and Staworko, 2016) in limited settings.

*Correspondence: rstrivedi@gatech.edu and bunyamis@amazon.com. Work done when Rakshit Trivedi interned at Amazon.

The key challenges for the graph alignment problem emanate from the fact that the real world data are noisy and intricate in nature. The noisy or sparse data make it difficult to learn robust alignment features, and data abundance leads to computational challenges due to the combinatorial permutations needed for alignment. These challenges are compounded in multi-relational settings due to heterogeneous nodes and edges in such graphs.

Recently, deep learning has shown significant impact in learning useful information over noisy, large-scale and heterogeneous graph data (Rossi et al., 2017). We, therefore, posit that combining graph alignment task with deep representation learning across multi-relational graphs has potential to induce a synergistic effect on both tasks. Specifically, we identify that a key component of graph alignment process—entity linkage—also plays a vital role in learning across graphs. For instance, the embeddings learned over two knowledge graphs for an actor should be closer to one another compared to the embeddings of all the other entities. Similarly, the entities that are already aligned together across the two graphs should produce better embeddings due to the shared context and data. To model this phenomenon, we propose **LinkNBed**, a novel deep learning framework that jointly performs representation learning and graph linkage task. To achieve this, we identify key challenges involved in the learning process and make the following contributions to address them:

- We propose novel and principled approach towards jointly learning entity representations and entity linkage. The novelty of our framework stems from its ability to support linkage task across heterogeneous types of entities.
- We devise a graph-independent inductive framework that learns functions to capture contextual information for entities and relations. It combines the structural and semantic information in individual graphs for joint inference in a principled manner.
- Labeled instances (specifically positive instances for linkage task) are typically very sparse and hence we design a novel multi-task loss function where entity linkage task is tackled in robust manner across various learning scenarios such as learning only with unlabeled instances or only with negative instances.

- We design an efficient training procedure to perform joint training in linear time in the number of triples. We demonstrate superior performance of our method on two datasets curated from Freebase and IMDB against state-of-the-art neural embedding methods.

2 Preliminaries

2.1 Knowledge Graph Representation

A knowledge graph \mathcal{G} comprises of set of facts represented as triplets (e^s, r, e^o) denoting the relationship r between subject entity e^s and object entity e^o . Associated to this knowledge graph, we have a set of attributes that describe observed characteristics of an entity. Attributes are represented as set of key-value pairs for each entity and an attribute can have null (missing) value for an entity. We follow *Open World Assumption - triplets not observed in knowledge graph are considered to be missing but not false*. We assume that there are no duplicate triplets or self-loops.

2.2 Multi-Graph Relational Learning

Definition. Given a collection of knowledge graphs \mathcal{G} , Multi-Graph Relational Learning refers to the task of learning information rich representations of entities and relationships across graphs. The learned embeddings can further be used to infer new knowledge in the form of link prediction or learn new labels in the form of entity linkage. We motivate our work with the setting of two knowledge graphs where given two graphs $G_1, G_2 \in \mathcal{G}$, the task is to match an entity $e_{G_1} \in G_1$ to an entity $e_{G_2} \in G_2$ if they represent the same real-world entity. We discuss a straightforward extension of this setting to more than two graphs in Section 7.

Notations. Let X and Y represent realization of two such knowledge graphs extracted from two different sources. Let n_e^X and n_e^Y represent number of entities in X and Y respectively. Similarly, n_r^X and n_r^Y represent number of relations in X and Y . We combine triplets from both X and Y to obtain set of all observed triplets $\mathcal{D} = \{(e^s, r, e^o)_p\}_{p=1}^P$ where P is total number of available records across from both graphs. Let \mathcal{E} and \mathcal{R} be the set of all entities and all relations in \mathcal{D} respectively. Let $|\mathcal{E}| = n$ and $|\mathcal{R}| = m$. In addition to \mathcal{D} , we also have set of linkage labels \mathcal{L} for entities between X and Y . Each record in \mathcal{L} is represented as triplet $(e^X \in X, e^Y \in Y, l \in \{0, 1\})$ where $l = 1$ when the entities are matched and $l = 0$ otherwise.

3 Proposed Method: LinkNBed

We present a novel *inductive multi-graph* relational learning framework that learns a set of aggregator functions capable of ingesting various contextual information for both entities and relationships in multi-relational graph. These functions encode the ingested structural and semantic information into low-dimensional entity and relation embeddings. Further, we use these representations to learn a relational score function that computes how two entities are likely to be connected in a particular relationship. The key idea behind this formulation is that when a triplet is observed, the relationship between the two entities can be explained using various contextual information such as local *neighborhood* features of both entities, *attribute* features of both entities and *type* information of the entities which participate in that relationship.

We outline two key insights for establishing the relationships between embeddings of the entities over multiple graphs in our framework:

Insight 1 (Embedding Similarity): If the two entities $e^X \in X$ and $e^Y \in Y$ represent the same real-world entity then their embeddings e^X and e^Y will be close to each other.

Insight 2 (Semantic Replacement): For a given triplet $t = (e^s, r, e^o) \in X$, denote $g(t)$ as the function that computes a relational score for t using entity and relation embeddings. If there exists a matching entity $e^{s'} \in Y$ for $e^s \in X$, denote $t' = (e^{s'}, r, e^o)$ obtained after replacing e^s with $e^{s'}$. In this case, $g(t) \sim g(t')$ i.e. score of triplets t and t' will be similar.

For a triplet $(e^s, r, e^o) \in \mathcal{D}$, we describe encoding mechanism of LinkNBed as three-layered architecture that computes the final output representations of $\mathbf{z}^r, \mathbf{z}^{e^s}, \mathbf{z}^{e^o}$ for the given triplet. Figure 1 provides an overview of LinkNBed architecture and we describe the three steps below:

3.1 Atomic Layer

Entities, Relations, Types and Attributes are first encoded in its basic vector representations. We use these basic representations to derive more complex contextual embeddings further.

Entities, Relations and Types. The embedding vectors corresponding to these three components are learned as follows:

$$\mathbf{v}^{e^s} = f(\mathbf{W}^E \mathbf{e}^s) \quad \mathbf{v}^{e^o} = f(\mathbf{W}^E \mathbf{e}^o) \quad (1)$$

$$\mathbf{v}^r = f(\mathbf{W}^R \mathbf{r}) \quad \mathbf{v}^t = f(\mathbf{W}^T \mathbf{t}) \quad (2)$$

where $\mathbf{v}^{e^s}, \mathbf{v}^{e^o} \in \mathbb{R}^d$. $\mathbf{e}^s, \mathbf{e}^o \in \mathbb{R}^n$ are ‘‘one-hot’’ representations of e^s and e^o respectively. $\mathbf{v}^r \in \mathbb{R}^k$ and $\mathbf{r} \in \mathbb{R}^m$ is ‘‘one-hot’’ representation of r . $\mathbf{v}^t \in \mathbb{R}^q$ and $\mathbf{t} \in \mathbb{R}^z$ is ‘‘one-hot’’ representation of t . $\mathbf{W}^E \in \mathbb{R}^{d \times n}$, $\mathbf{W}^R \in \mathbb{R}^{k \times m}$ and $\mathbf{W}^T \in \mathbb{R}^{q \times z}$ are the entity, relation and type embedding matrices respectively. f is a nonlinear activation function (Relu in our case). $\mathbf{W}^E, \mathbf{W}^R$ and \mathbf{W}^T can be initialized randomly or using pre-trained word embeddings or vector compositions based on name phrases of components (Socher et al., 2013). **Attributes.** For a given attribute a represented as key-value pair, we use paragraph2vec (Le and Mikolov, 2014) type of embedding network to learn attribute embedding. Specifically, we represent attribute embedding vector as:

$$\mathbf{a} = f(\mathbf{W}^{\text{key}} \mathbf{a}_{\text{key}} + \mathbf{W}^{\text{val}} \mathbf{a}_{\text{val}}) \quad (3)$$

where $\mathbf{a} \in \mathbb{R}^y$, $\mathbf{a}_{\text{key}} \in \mathbb{R}^u$ and $\mathbf{a}_{\text{val}} \in \mathbb{R}^v$. $\mathbf{W}^{\text{key}} \in \mathbb{R}^{y \times u}$ and $\mathbf{W}^{\text{val}} \in \mathbb{R}^{y \times v}$. \mathbf{a}_{key} will be ‘‘one-hot’’ vector and \mathbf{a}_{val} will be feature vector. Note that the dimensions of the embedding vectors do not necessarily need to be the same.

3.2 Contextual Layer

While the entity and relationship embeddings described above help to capture very generic latent features, embeddings can be further enriched to capture structural information, attribute information and type information to better explain the existence of a fact. Such information can be modeled as context of nodes and edges in the graph. To this end, we design the following canonical aggregator function that learns various contextual information by aggregating over relevant embedding vectors:

$$\mathbf{c}(z) = \text{AGG}(\{\mathbf{z}', \forall z' \in C(z)\}) \quad (4)$$

where $\mathbf{c}(z)$ is the vector representation of the aggregated contextual information for component z . Here, component z can be either an entity or a relation. $C(z)$ is the set of components in the context of z and \mathbf{z}' correspond to the vector embeddings of those components. AGG is the aggregator function which can take many forms such Mean, Max, Pooling or more complex LSTM based aggregators. It is plausible that different components in a context may have varied impact on the component for which the embedding is being learned. To account for this, we employ a soft attention mechanism where we learn attention coefficients

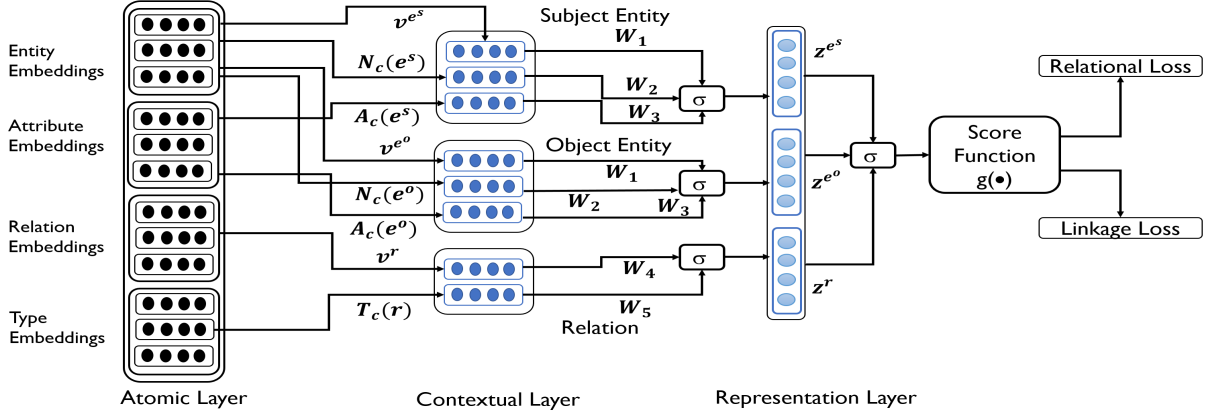


Figure 1: LinkNBed Architecture Overview - one step score computation for a given triplet (e^s, r, e^o) . The Attribute embeddings are not simple lookups but they are learned as shown in Eq 3

to weight components based on their impact before aggregating them. We modify Eq. 4 as:

$$\mathbf{c}(z) = \text{AGG}(\mathbf{q}(z) * \{\mathbf{z}', \forall z' \in C(z)\}) \quad (5)$$

where

$$\mathbf{q}(z) = \frac{\exp(\theta_z)}{\sum_{z' \in C(z)} \exp(\theta_{z'})} \quad (6)$$

and θ_z 's are the parameters of attention model.

Following contextual information is modeled in our framework:

Entity Neighborhood Context $\mathbf{N}_c(e) \in \mathbb{R}^d$. Given a triplet (e^s, r, e^o) , the neighborhood context for an entity e^s will be the nodes located near e^s other than the node e^o . This will capture the effect of local neighborhood in the graph surrounding e^s that drives e^s to participate in fact (e^s, r, e^o) . We use Mean as aggregator function. As there can be large number of neighbors, we collect the neighborhood set for each entity as a pre-processing step using a random walk method. Specifically, given a node e , we run k rounds of random-walks of length l following (Hamilton et al., 2017) and create set $\mathcal{N}(e)$ by adding all unique nodes visited across these walks. This context can be similarly computed for object entity.

Entity Attribute Context $\mathbf{A}_c(e) \in \mathbb{R}^y$. For an entity e , we collect all attribute embeddings for e obtained from Atomic Layer and learn aggregated information over them using Max operator given in Eq. 4.

Relation Type Context $\mathbf{T}_c(r) \in \mathbb{R}^q$. We use type context for relation embedding i.e. for a given relationship r , this context aims at capturing the effect of type of entities that have participated in this relationship. For a given triplet (e^s, r, e^o) , type context for relationship r is computed by aggregation with mean over type embeddings corresponding to the context of r . Appendix C provides specific forms of contextual information.

3.3 Representation Layer

Having computed the atomic and contextual embeddings for a triplet (e^s, r, e^o) , we obtain the final embedded representations of entities and relation in the triplet using the following formulation:

$$\mathbf{z}^{e^s} = \sigma \left(\underbrace{\mathbf{W}_1 \mathbf{v}^{e^s}}_{\text{Subject Entity Embedding}} + \underbrace{\mathbf{W}_2 \mathbf{N}_c(e^s)}_{\text{Neighborhood Context}} + \underbrace{\mathbf{W}_3 \mathbf{A}_c(e^s)}_{\text{Subject Entity Attributes}} \right) \quad (7)$$

$$\mathbf{z}^{e^o} = \sigma \left(\underbrace{\mathbf{W}_1 \mathbf{v}^{e^o}}_{\text{Object Entity Embedding}} + \underbrace{\mathbf{W}_2 \mathbf{N}_c(e^o)}_{\text{Neighborhood Context}} + \underbrace{\mathbf{W}_3 \mathbf{A}_c(e^o)}_{\text{Object Entity Attributes}} \right) \quad (8)$$

$$\mathbf{z}^r = \sigma \left(\underbrace{\mathbf{W}_4 \mathbf{v}^r}_{\text{Relation Embedding}} + \underbrace{\mathbf{W}_5 \mathbf{T}_c(r)}_{\text{Entity Type Context}} \right) \quad (9)$$

where $\mathbf{W}_1, \mathbf{W}_2 \in \mathbb{R}^{d \times d}$, $\mathbf{W}_3 \in \mathbb{R}^{d \times y}$, $\mathbf{W}_4 \in \mathbb{R}^{d \times k}$ and $\mathbf{W}_5 \in \mathbb{R}^{d \times q}$. σ is nonlinear activation function – generally Tanh or Relu.

Following is the rationale for our formulation: An entity’s representation can be enriched by encoding information about the local neighborhood features and attribute information associated with the entity in addition to its own latent features. Parameters $\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3$ learn to capture these different aspects and map them into the entity embedding space. Similarly, a relation’s representation can be enriched by encoding information about entity types that participate in that relationship in addition to its own latent features. Parameters $\mathbf{W}_4, \mathbf{W}_5$ learn to capture these aspects and map them into the relation embedding space. Further, as the ultimate goal is to jointly learn over multiple graphs, shared parameterization in our model facilitate the propagation of information across graphs thereby making it a graph-independent inductive model. The flexibility of the model stems from the ability to shrink it (to a very simple model considering atomic entity and relation embeddings only) or expand it (to a complex model by adding different contextual information) without affecting any other step in the learning procedure.

3.4 Relational Score Function

Having observed a triplet (e^s, r, e^o) , we first use Eq. 7, 8 and 9 to compute entity and relation representations. We then use these embeddings to capture relational interaction between two entities using the following score function $g(\cdot)$:

$$g(e^s, r, e^o) = \sigma(\mathbf{z}^{r^T} \cdot (\mathbf{z}^{e^s} \odot \mathbf{z}^{e^o})) \quad (10)$$

where $\mathbf{z}^r, \mathbf{z}^{e^s}, \mathbf{z}^{e^o} \in \mathbb{R}^d$ are d -dimensional representations of entity and relationships as described below. σ is the nonlinear activation function and \odot represent element-wise product.

4 Efficient Learning Procedure

4.1 Objective Function

The complete parameter space of the model can be given by: $\Omega = \{\{\mathbf{W}_i\}_{i=1}^5, \mathbf{W}^E, \mathbf{W}^R, \mathbf{W}^{\text{key}}, \mathbf{W}^{\text{val}}, \mathbf{W}^t, \Theta\}$. To learn these parameters, we design a novel multi-task objective function that jointly trains over two graphs. As identified earlier, the goal of our model is to leverage the available linkage information across graphs for optimizing the entity and relation embeddings such that they can explain the observed triplets across the graphs. Further, we want to leverage these optimized embeddings to match

entities across graphs and expand the available linkage information. To achieve this goal, we define following two different loss functions catering to each learning task and jointly optimize over them as a multi-task objective to learn model parameters:

Relational Learning Loss. This is conventional loss function used to learn knowledge graph embeddings. Specifically, given a p -th triplet $(e^s, r, e^o)_p$ from training set \mathcal{D} , we sample C negative samples by replacing either head or tail entity and define a contrastive max margin function as shown in (Socher et al., 2013):

$$L_{rel} = \sum_{c=1}^C \max(0, \gamma - g(e_p^s, r_p, e_p^o) + g'(e_c^s, r_p, e_p^o)) \quad (11)$$

where, γ is margin, e_c^s represent corrupted entity and $g'(e_c^s, r_p, e_p^o)$ represent corrupted triplet score.

Linkage Learning Loss: We design a novel loss function to leverage pairwise label set \mathcal{L} . Given a triplet (e_X^s, r_X, e_X^o) from knowledge graph X , we first find the entity e_Y^+ from graph Y that represent the same real-world entity as e_X^s . We then replace e_X^s with e_Y^+ and compute score $g(e_Y^+, r_X, e_X^o)$. Next, we find set of all entities E_Y^- from graph Y that has a negative label with entity e_X^s . We consider them analogous to the negative samples we generated for Eq. 11. We then propose the label learning loss function as:

$$L_{lab} = \sum_{z=1}^Z \max(0, \gamma - g(e_Y^+, r_X, e_X^o) + (g'(e_Y^-, r_X, e_X^o)_z)) \quad (12)$$

where, Z is the total number of negative labels for e_X . γ is margin which is usually set to 1 and $e_Y^- \in E_Y^-$ represent entity from graph Y with which entity e_X^s had a negative label. Please note that this applies symmetrically for the triplets that originate from graph Y in the overall dataset. Note that if both entities of a triplet have labels, we will include both cases when computing the loss. Eq. 12 is inspired by *Insight 1* and *Insight 2* defined earlier in Section 2. Given a set \mathcal{D} of N observed triplets across two graphs, we define complete multi-task objective as:

$$\mathbf{L}(\Omega) = \sum_{i=1}^N [b \cdot L_{rel} + (1-b) \cdot L_{lab}] + \lambda \|\Omega\|_2^2 \quad (13)$$

Algorithm 1 LinkNBed mini-batch Training

Input: Mini-batch \mathcal{M} , Negative Sample Size C , Negative Label Size Z , Attribute data att_data , Neighborhood data $nhbr_data$, Type data $type_data$, Positive Label Dict pos_dict , Negative Label Dict neg_dict

Output: Mini-batch Loss $\mathcal{L}_{\mathcal{M}}$.

$\mathcal{L}_{\mathcal{M}} = 0$

score_pos = []; score_neg = []; score_pos_lab = [];
score_neg_lab = []

for $i = 0$ **to** $size(\mathcal{M})$ **do**

input_tuple = $\mathcal{M}[i] = (e^s, r, e^o)$

sc = compute_triplet_score(e^s, r, e^o) (Eq. 10)

score_pos.append(sc)

for $j = 0$ **to** C **do**

Select e_c^s from entity list such that $e_c^s \neq e^s$
and $e_c^s \neq e^o$ and $(e_c^s, r, e^o) \notin \mathcal{D}$

sc_neg = compute_triplet_score(e_c^s, r, e^o)

score_neg.append(sc_neg)

end for

if e^s in pos_dict **then**

e^{s+} = positive label for e^s

sc_pos_l = compute_triplet_score(e^{s+}, r, e^o)

score_pos_lab.append(sc_pos_l)

end if

for $k = 0$ **to** Z **do**

Select e^{s-} from neg_dict

sc_neg_l = compute_triplet_score(e^{s-}, r, e^o)

score_neg_lab.append(sc_neg_l)

end for

end for

$\mathcal{L}_{\mathcal{M}} +=$ compute_minibatch_loss(score_pos,
score_neg, score_pos_lab, score_neg_lab)
(Eq. 13)

Back-propagate errors and update parameters Ω

return $\mathcal{L}_{\mathcal{M}}$

where Ω is set of all model parameters and λ is regularization hyper-parameter. b is weight hyper-parameter used to attribute importance to each task. We train with mini-batch SGD procedure (Algorithm 1) using Adam Optimizer.

Missing Positive Labels. It is expensive to obtain positive labels across multiple graphs and hence it is highly likely that many entities will not have positive labels available. For those entities, we will modify Eq. 12 to use the original triplet (e_X^s, r_X, e_X^o) in place of perturbed triplet $g(e_Y^+, r_X, e_X^o)$ for the positive label. The rationale here again arises from *Insight 2* wherein embeddings of two duplicate entities should be able to

replace each other without affecting the score.

Training Time Complexity. Most contextual information is pre-computed and available to all training steps which leads to constant time embedding lookup for those context. But for attribute network, embedding needs to be computed for each attribute separately and hence the complexity to compute score for one triplet is $\mathcal{O}(2a)$ where a is number of attributes. Also for training, we generate C negative samples for relational loss function and use Z negative labels for label loss function. Let $k = C + Z$. Hence, the training time complexity for a set of n triplets will be $\mathcal{O}(2ak * n)$ which is linear in number of triplets with a constant factor as $ak \ll n$ for real world knowledge graphs. This is desirable as the number of triplets tend to be very large per graph in multi-relational settings.

Memory Complexity. We borrow notations from (Nickel et al., 2016a) and describe the parameter complexity of our model in terms of the number of each component and corresponding embedding dimension requirements. Let $H_a = 2 * N_e H_e + N_r H_r + N_t H_t + N_k H_k + N_v H_v$. The parameter complexity of our model is: $H_a * (H_b + 1)$. Here, N_e, N_r, N_t, N_k, N_v signify number of entities, relations, types, attribute keys and vocab size of attribute values across both datasets. Here H_b is the output dimension of the hidden layer.

5 Experiments

5.1 Datasets

We evaluate LinkNBed and baselines on two real world knowledge graphs: D-IMDB (derived from large scale IMDB data snapshot) and D-FB (derived from large scale Freebase data snapshot). Table 5.1 provides statistics for our final dataset used in the experiments. Appendix B.1 provides complete details about dataset processing.

Dataset Name	# Entities	# Relations	# Attributes	# Entity Types	# Available Triples
D-IMDB	378207	38	23	41	143928582
D-FB	39667	146	69	324	22140475

Table 1: Statistics for Datasets: D-IMDB and D-FB

5.2 Baselines

We compare the performance of our method against state-of-the-art representation learning baselines that use neural embedding techniques to learn entity and relation representation. Specifically, we consider compositional methods of

RESCAL (Nickel et al., 2011) as basic matrix factorization method, DISTMULT (Yang et al., 2015) as simple multiplicative model good for capturing symmetric relationships, and Complex (Trouillon et al., 2016), an upgrade over DISTMULT that can capture asymmetric relationships using complex valued embeddings. We also compare against translational model of STransE that combined original structured embedding with TransE and has shown state-of-art performance in benchmark testing (Kadlec et al., 2017). Finally, we compare with GAKE (Feng et al., 2016), a model that captures context in entity and relationship representations.

In addition to the above state-of-art models, we analyze the effectiveness of different components of our model by comparing with various versions that use partial information. Specifically, we report results on following variants:

LinkNBed - Embed Only. Only use entity embeddings, **LinkNBed - Attr Only.** Only use Attribute Context, **LinkNBed - Nhbr Only.** Only use Neighborhood Context, **LinkNBed - Embed + Attr.** Use both Entity embeddings and Attribute Context, **LinkNBed - Embed + Nhbr.** Use both Entity embeddings and Neighbor Context and **LinkNBed - Embed All.** Use all three Contexts.

5.3 Evaluation Scheme

We evaluate our model using two inference tasks: **Link Prediction.** Given a test triplet (e^s, r, e^o) , we first score this triplet using Eq. 10. We then replace e^o with all other entities in the dataset and filter the resulting set of triplets as shown in (Bordes et al., 2013). We score the remaining set of perturbed triplets using Eq. 10. All the scored triplets are sorted based on the scores and then the rank of the ground truth triplet is used for the evaluation. We use this ranking mechanism to compute HITS@10 (predicted rank ≤ 10) and reciprocal rank ($\frac{1}{rank}$) of each test triplet. We report the mean over all test samples.

Entity Linkage. In alignment with *Insight 2*, we pose a novel evaluation scheme to perform entity linkage. Let there be two ground truth test sample triplets: $(e_X, e_Y^+, 1)$ representing a positive duplicate label and $(e_X, e_Y^-, 0)$ representing a negative duplicate label. Algorithm 2 outlines the procedure to compute linkage probability or score q ($\in [0, 1]$) for the pair (e_X, e_Y) . We use L1 distance between the two vectors analogous

Algorithm 2 Entity Linkage Score Computation

Input: Test pair $-(e_X \in X, e_Y \in Y)$.

Output: Linkage Score $- q$.

1. Collect all triplets involving e_X from graph X and all triplets involving e_Y from graph Y into a combined set \mathcal{O} . Let $|\mathcal{O}| = k$.
 2. Construct $S_{orig} \in \mathbb{R}^k$.
For each triplet $o \in \mathcal{O}$, compute score $g(o)$ using Eq. 10 and store the score in S_{orig} .
 3. Create triplet set \mathcal{O}' as following:
if $o \in \mathcal{O}$ contain $e^X \in X$ **then**
 Replace e^X with e^Y to create perturbed triplet o' and store it in \mathcal{O}'
end if
if $o \in \mathcal{O}$ contain $e^Y \in Y$ **then**
 Replace e^Y with e^X to create perturbed triplet o' and store it in \mathcal{O}'
end if
 4. Construct $S_{repl} \in \mathbb{R}^k$.
For each triplet $o' \in \mathcal{O}'$, compute score $g(o')$ using Eq. 10 and store the score in S_{repl} .
 5. Compute q .
Elements in S_{orig} and S_{repl} have one-one correspondence so take the mean absolute difference:
 $q = |S_{orig} - S_{repl}|_1$
return q
-

to Mean Absolute Error (MAE). In lieu of hard-labeling test pairs, we use score q to compute Area Under the Precision-Recall Curve (AUPRC).

For the baselines and the unsupervised version (with no labels for entity linkage) of our model, we use second stage multilayer Neural Network as classifier for evaluating entity linkage. Appendix B.2 provides training configuration details.

5.4 Predictive Analysis

Link Prediction Results. We train LinkNBed model jointly across two knowledge graphs and then perform inference over individual graphs to report link prediction reports. For baselines, we train each baseline on individual graphs and use parameters specific to the graph to perform link prediction inference over each individual graph. Table 5.4 shows link prediction performance for all methods. Our model variant with attention mechanism outperforms all the baselines with 4.15% improvement over single graph state-of-the-art Complex model on D-IMDB and 8.23% improvement on D-FB dataset. D-FB is more challenging dataset to

Method	D-IMDB-HITS10	D-IMDB-MRR	D-FB-HITS10	D-FB-MRR
RESCAL	75.3	0.592	69.99	0.147
DISTMULT	79.5	0.691	72.34	0.556
Complex	83.2	0.752	75.67	0.629
STransE	80.7	0.421	69.87	0.397
GAKE	69.5	0.114	63.22	0.093
LinkNBed-Embed Only	79.9	0.612	73.2	0.519
LinkNBed-Attr Only	82.2	0.676	74.7	0.588
LinkNBed-Nhbr Only	80.1	0.577	73.4	0.572
LinkNBed-Embed + Attr	84.2	0.673	78.39	0.606
LinkNBed-Embed + Nhbr	81.7	0.544	73.45	0.563
LinkNBed-Embed All	84.3	0.725	80.2	0.632
LinkNBed-Embed All (Attention)	86.8	0.733	81.9	0.677
Improvement (%)	4.15	1.10	8.23	7.09

Table 2: Link Prediction Results on both datasets

learn as it has a large set of sparse relationships, types and attributes and it has an order of magnitude lesser relational evidence (number of triplets) compared to D-IMDB. Hence, LinkNBed’s pronounced improvement on D-FB demonstrates the effectiveness of the model. The simplest version of LinkNBed with only entity embeddings resembles DISTMULT model with different objective function. Hence closer performance of those two models aligns with expected outcome. We observed that the Neighborhood context alone provides only marginal improvements while the model benefits more from the use of attributes. Despite being marginal, attention mechanism also improves accuracy for both datasets. Compared to the baselines which are obtained by trained and evaluated on individual graphs, our superior performance demonstrates the effectiveness of multi-graph learning.

Entity Linkage Results. We report entity linkage results for our method in two settings: a.) Supervised case where we train using both the objective functions. b.) Unsupervised case where we learn with only the relational loss function. The latter case resembles the baseline training where each model is trained separately on two graphs in an unsupervised manner. For performing the entity linkage in unsupervised case for all models, we first train a second stage of simple neural network classifier and then perform inference. In the supervised case, we use Algorithm 2 for performing the inference. Table 5.4 demonstrates the performance of all methods on this task. Our method significantly outperforms all the baselines with 33.86% over second best baseline in supervised case and 17.35% better performance in unsupervised case. The difference in the performance of our method in two cases demonstrate that the two training objectives are helping one another by learning across the graphs. GAKE’s superior performance on this task compared to the other state-of-the-art relational baselines shows the importance of using contex-

Method	AUPRC (Supervised)	AUPRC (Unsupervised)
RESCAL	-	0.327
DISTMULT	-	0.292
Complex	-	0.359
STransE	-	0.231
GAKE	-	0.457
LinkNBed-Embed Only	0.376	0.304
LinkNBed-Attr Only	0.451	0.397
LinkNBed-Nhbr Only	0.388	0.322
LinkNBed-Embed + Attr	0.512	0.414
LinkNBed-Embed + Nhbr	0.429	0.356
LinkNBed-Embed All	0.686	0.512
LinkNBed-Embed All (Attention)	0.691	0.553
Improvement (%)	33.86	17.35

Table 3: Entity Linkage Results - Unsupervised case uses classifier at second step

tual information for entity linkage. Performance of other variants of our model again demonstrate that attribute information is more helpful than neighborhood context and attention provides marginal improvements. We provide further insights with examples and detailed discussion on entity linkage task in Appendix A.

6 Related Work

6.1 Neural Embedding Methods for Relational Learning

Compositional Models learn representations by various composition operators on entity and relational embeddings. These models are multiplicative in nature and highly expressive but often suffer from scalability issues. Initial models include RESCAL (Nickel et al., 2011) that uses a relation specific weight matrix to explain triplets via pairwise interactions of latent features, Neural Tensor Network (Socher et al., 2013), more expressive model that combines a standard NN layer with a bilinear tensor layer and (Dong et al., 2014) that employs a concatenation-projection method to project entities and relations to lower dimensional space. Later, many sophisticated models (Neural Association Model (Liu et al., 2016), HoLE (Nickel et al., 2016b)) have been proposed. Path based composition models (Toutanova et al., 2016) and contextual models GAKE (Feng et al., 2016) have been recently studied to capture more information from graphs. Recently, model like Complex (Trouillon et al., 2016) and Analogy (Liu et al., 2017) have demonstrated state-of-the art performance on relational learning tasks. **Translational Models** ((Bordes et al., 2014), (Bordes et al., 2011), (Bordes et al., 2013), (Wang et al., 2014), (Lin et al., 2015), (Xiao et al., 2016)) learn representation by

employing translational operators on the embeddings and optimizing based on their score. They offer an additive and efficient alternative to expensive multiplicative models. Due to their simplicity, they often lose expressive power. For a comprehensive survey of relational learning methods and empirical comparisons, we refer the readers to (Nickel et al., 2016a), (Kadlec et al., 2017), (Toutanova and Chen, 2015) and (Yang et al., 2015). None of these methods address multi-graph relational learning and cannot be adapted to tasks like entity linkage in straightforward manner.

6.2 Entity Resolution in Relational Data

Entity Resolution refers to resolving entities available in knowledge graphs with entity mentions in text. (Dredze et al., 2010) proposed entity disambiguation method for KB population, (He et al., 2013) learns entity embeddings for resolution, (Huang et al., 2015) propose a sophisticated DNN architecture for resolution, (Campbell et al., 2016) proposes entity resolution across multiple social domains, (Fang et al., 2016) jointly embeds text and knowledge graph to perform resolution while (Globerson et al., 2016) proposes Attention Mechanism for Collective Entity Resolution.

6.3 Learning across multiple graphs

Recently, learning over multiple graphs have gained traction. (Liu and Yang, 2016) divides a multi-relational graph into multiple homogeneous graphs and learns associations across them by employing product operator. Unlike our work, they do not learn across multiple multi-relational graphs. (Pujara and Getoor, 2016) provides logic based insights for cross learning, (Perschina et al., 2015) does pairwise entity matching across multi-relational graphs and is very expensive, (Chen et al., 2017) learns embeddings to support multi-lingual learning and Big-Align (Koutra et al., 2013) tackles graph alignment problem efficiently for bipartite graphs. None of these methods learn latent representations or jointly train graph alignment and learning which is the goal of our work.

7 Concluding Remarks and Future Work

We present a novel relational learning framework that learns entity and relationship embeddings across multiple graphs. The proposed representation learning framework leverage an efficient learning and inference procedure which takes into account the duplicate entities representing the same

real-world entity in a multi-graph setting. We demonstrate superior accuracies on link prediction and entity linkage tasks compared to the existing approaches that are trained only on individual graphs. We believe that this work opens a new research direction in joint representation learning over multiple knowledge graphs.

Many data driven organizations such as Google and Microsoft take the approach of constructing a unified super-graph by integrating data from multiple sources. Such unification has shown to significantly help in various applications, such as search, question answering, and personal assistance. To this end, there exists a rich body of work on linking entities and relations, and conflict resolution (e.g., knowledge fusion (Dong et al., 2014)). Still, the problem remains challenging for large scale knowledge graphs and this paper proposes a deep learning solution that can play a vital role in this construction process. In real-world setting, we envision our method to be integrated in a large scale system that would include various other components for tasks like conflict resolution, active learning and human-in-loop learning to ensure quality of constructed super-graph. However, we point out that our method is not restricted to such use cases—one can readily apply our method to directly make inference over multiple graphs to support applications like question answering and conversations.

For future work, we would like to extend the current evaluation of our work from a two-graph setting to multiple graphs. A straightforward approach is to create a unified dataset out of more than two graphs by combining set of triplets as described in Section 2, and apply learning and inference on the unified graph without any major change in the methodology. Our inductive framework learns functions to encode contextual information and hence is graph independent. Alternatively, one can develop sophisticated approaches with iterative merging and learning over pairs of graphs until exhausting all graphs in an input collection.

Acknowledgments

We would like to give special thanks to Ben London, Tong Zhao, Arash Einolghozati, Andrew Borthwick and many others at Amazon for helpful comments and discussions. We thank the reviewers for their valuable comments and efforts towards improving our manuscript. This project was supported in part by NSF(IIS-1639792, IIS-1717916).

References

- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. DBpedia: A nucleus for a web of open data. In *The Semantic Web*.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *SIGMOD Conference*.
- Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2014. A semantic matching energy function for learning with multi-relational data. *Machine Learning*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795.
- Antoine Bordes, Jason Weston, Ronan Collobert, and Yoshua Bengio. 2011. Learning structured embeddings of knowledge bases. In *AAAI*.
- Peter Buneman and Slawek Staworko. 2016. Rdf graph alignment with bisimulation. *Proc. VLDB Endow*.
- W. M. Campbell, Lin Li, C. Dagli, J. Acevedo-Aviles, K. Geyer, J. P. Campbell, and C. Priebe. 2016. Cross-domain entity resolution in social media. *arXiv:1608.01386v1*.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R. Hruschka Jr., and Tom M. Mitchell. 2010. Toward an architecture for never-ending language learning. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence (AAAI 2010)*.
- Rose Catherine and William Cohen. 2016. Personalized recommendations using knowledge graphs: A probabilistic logic programming approach. In *Proceedings of the 10th ACM Conference on Recommender Systems*.
- Muhao Chen, Yingtao Tian, Mohan Yang, and Carlo Zaniolo. 2017. Multilingual knowledge graph embeddings for cross-lingual knowledge alignment.
- Wanyun Cui, Yanghua Xiao, Haixun Wang, Yangqiu Song, Seung-won Hwang, and Wei Wang. 2017. Kbqa: Learning question answering over qa corpora and knowledge bases. *Proc. VLDB Endow*.
- Jeffrey Dalton, Laura Dietz, and James Allan. 2014. Entity query feature expansion using knowledge base links. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*.
- Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmman, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 601–610.
- Mark Dredze, Paul McNamee, Delip Rao, Adam Gerber, and Tim Finin. 2010. Entity disambiguation for knowledge base population. In *Proceedings of the 23rd International Conference on Computational Linguistics*.
- Wei Fang, Jianwen Zhang, Dilin Wang, Zheng Chen, and Ming Li. 2016. Entity disambiguation by knowledge and text jointly embedding. In *CoNLL*.
- Jun Feng, Minlie Huang, Yang Yang, and Xiaoyan Zhu. 2016. Gake: Graph aware knowledge embedding. In *COLING*.
- Evgeniy Gabrilovich and Shaul Markovitch. 2009. Wikipedia-based semantic interpretation for natural language processing. *J. Artif. Int. Res.*
- Amir Globerson, Nevena Lazic, Soumen Chakrabarti, Amarnag Subramanya, Michael Ringgaard, and Fernando Pereira. 2016. Collective entity resolution with multi-focal attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Representation learning on graphs: Methods and applications. *arXiv:1709.05584*.
- Zhengyan He, Shujie Liu, Mu Li, Ming Zhou, Longkai Zhang, and Houfeng Wang. 2013. Learning entity representation for entity disambiguation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*.
- Hongzhao Huang, Larry Heck, and Heng Ji. 2015. Leveraging deep neural networks and knowledge graphs for entity disambiguation. *arXiv:1504.07678v1*.
- Rudolph Kadlec, Ondrej Bajgar, and Jan Kleindienst. 2017. Knowledge base completion: Baselines strike back. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*.
- Danai Koutra, HangHang Tong, and David Lubensky. 2013. Big-align: Fast bipartite graph alignment. In *2013 IEEE 13th International Conference on Data Mining*.
- Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the 31st International Conference on Machine Learning*.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. *AAAI Conference on Artificial Intelligence*.

- Hanxiao Liu, Yuexin Wu, and Yimin Yang. 2017. Analytical inference for multi-relational embeddings. In *Proceedings of the 34th International Conference on Machine Learning*.
- Hanxiao Liu and Yimin Yang. 2016. Cross-graph learning of multi-relational associations. In *Proceedings of the 33rd International Conference on Machine Learning*.
- Quan Liu, Hui Jiang, Andrew Evdokimov, Zhen-Hua Ling, Xiaodan Zhu, Si Wei, and Yu Hu. 2016. Probabilistic reasoning via deep learning: Neural association models. *arXiv:1603.07704v2*.
- Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2016a. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*.
- Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016b. Holographic embeddings of knowledge graphs.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 809–816.
- Maria Pershina, Mohamed Yakout, and Kaushik Chakrabarti. 2015. Holistic entity matching across knowledge graphs. In *2015 IEEE International Conference on Big Data (Big Data)*.
- Jay Pujara and Lise Getoor. 2016. Generic statistical relational entity resolution in knowledge graphs. In *Sixth International Workshop on Statistical Relational AI*.
- Ryan A. Rossi, Rong Zhou, and Nesreen K. Ahmed. 2017. Deep feature learning for graphs. *arXiv:1704.08829*.
- Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. 2013. Reasoning with neural tensor networks for knowledge base completion. In *Advances in Neural Information Processing Systems*, pages 926–934.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: A core of semantic knowledge. In *Proceedings of the 16th International Conference on World Wide Web*.
- Kristina Toutanova and Danqi Chen. 2015. Observed versus latent features for knowledge base and text inference. In *ACL*.
- Kristina Toutanova, Xi Victoria Lin, Wen-tau Yih, Hoi-fung Poon, and Chris Quirk. 2016. Compositional learning of embeddings for relation paths in knowledge bases and text. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 1434–1444.
- Theo Trouillon, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *Proceedings of the 33rd International Conference on Machine Learning*.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes.
- Han Xiao, Minlie Huang, and Xiaoyan Zhu. 2016. Transg: A generative model for knowledge graph embedding. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding entities and relations for learning and inference in knowledge bases. *arXiv:1412.6575*.