Stable Tuple Embeddings for Dynamic Databases

1st Jan Toenshoff
RWTH Aachen University
Aachen, Germany
toenshoff@informatik.rwth-aachen.de

2nd Neta Friedman
Technion
Haifa, Israel
netafr@cs.technion.ac.il

3rd Martin Grohe
RWTH Aachen University
Aachen, Germany
grohe@informatik.rwth-aachen.de

4th Benny Kimelfeld
Technion
Haifa, Israel
bennyk@cs.technion.ac.il

Abstract—We study the problem of computing an embedding of the tuples of a relational database in a manner that is extensible to dynamic changes of the database. In this problem, the embedding should be stable in the sense that it should not change on the existing tuples due to the embedding of newly inserted tuples (as database applications might already rely on existing embeddings); at the same time, the embedding of all tuples, old and new, should retain high quality. This task is challenging since inter-dependencies among the embeddings of different entities are inherent in state-of-the-art embedding techniques for structured data.

We study two approaches to solving the problem. The first is an adaptation of Node2Vec to dynamic databases. The second is the FoRWaRD algorithm (Foreign Key Random Walk Embeddings for Relational Databases) that draws from embedding techniques for general graphs and knowledge graphs, and is inherently utilizing the schema and its key and foreign-key constraints. We evaluate the embedding algorithms using a collection of downstream tasks of column prediction over geographical and biological domains. We find that in the traditional static setting, our two embedding methods achieve comparable results that are compatible with the state-of-the-art for the specific applications. In the dynamic setting, we find that the FoRWaRD algorithm generally outperforms and runs faster than the alternatives, and moreover, it features only a mild reduction of quality even when the database consists of more than half newly inserted tuples after the initial training of the embedding.

Index Terms—Database Embedding, Node2Vec

I. Introduction

Standard machine learning algorithms assume representations of their input data as numerical vectors. Applying these algorithms for the analysis of non-numerical data requires *embeddings* of these data into a (typically) finite dimensional Euclidean vector space. The embedding needs to be "faithful" to the semantics. In particular, similar entities should be mapped to vectors that are close geometrically, and vice versa. In some modalities, the input comes with a useful embedding to begin with; for example, an image can be represented by the RGB intensities of its pixels. In others, semantic-aware embeddings have to be devised, and indeed have been devised, such as Word2Vec [34] and RoBERTa [30] for natural language [27], Node2Vec [15] and GraphSAGE [19] for the nodes of a graph, TransE [10] and RotatE [46] for the entities of a knowledge graph, and Mol2Vec [23] for molecule structures. An approach to evaluating the *quality* of a generic embedding technique is via different downstream tasks: solve a collection of machine-learning tasks by utilizing

machine-learning models that operate over the embedding (see, e.g., [26], [29]), as illustrated in Figure 1.

Generalizing graph embeddings, generic embeddings have also been devised for relational databases. Such embeddings have enabled the deployment of machine-learning architectures to traditional database tasks such as record similarity [7]–[9], [16], [17], record linking [14], [35] and other integration tasks such as schema, token and record matching (entity resolution) [11]. The embedded entities are typically either tuples or attribute values. In this work, we focus on tuple embeddings. Various approaches have been proposed for obtaining embeddings in databases. One is to concatenate predefined embeddings of the attribute values (which are, e.g., words or quantities) [14], [35] or permutations over the list of values [11]. Another approach views the tuples as text documents and applies word and document embeddings [8]. Cappuzzo, Papotti and Thirumuruganathan [11] studied the approach of transforming the database into a graph and applying a node embedding over this graph.

Arguably, an important advantage of the graph approach to database embedding is that it utilizes information that should be highly relevant to the semantics of the data and is freely available in databases: the *structure* of the data. This structure includes more than just the tabular form. An important difference between embedding a single relation and embedding a full relational database is that the latter entails considerable semantics via dependencies between relations: typically, columns in a database store foreign keys, which have no semantic meaning as atomic values but may have rich semantic meaning as references to tuples from other relations. For instance, we can infer considerable information about a tuple with apparently meaningless values, such as internal codes, by looking at the tuples that are referenced by this tuple. Hence, to appropriately capture the semantics of the data, embedding methods need to incorporate such dependencies. Yet, we are not aware of any embedding technique that directly uses the most common way of referencing tuples, namely foreign-key references. Yet, it has been shown useful in the general context of machine learning [42], [44]. As we explain later on, we show here how these can be gracefully utilized for the sake of high-quality tuple embedding.

Incorporating the database structure, and particularly references among tuples, means that the embeddings of tuples depend upon each other. This leads to new challenges since the database is often not a static object but rather serves a

dynamic organization (e.g., with arrival of new customers and purchases, new patients and patient records, etc.). When new data arrives, we are in a situation where we have an embedding for the old tuples but not the new ones. A straightforward solution to this problem is to reapply the embedding algorithm from scratch over the new database. This approach, however, suffers from two main drawbacks. First, it might be computationally too expensive to compute the embedding over the entire database upon every tuple arrival. More fundamentally, reapplying the embedding algorithm is likely to change the embedding of the old tuples due to the inherent randomness in most embedding algorithms.

It has been empirically demonstrated [43], [48] that for many standard node embedding algorithms for graphs, the embeddings change considerably even when applied repeatedly to identical graphs with identical parameter settings. Even worse, standard downstream classifiers change their predictions for individual data points when applied to these embeddings [43]. A similar behaviour has been observed earlier for word embeddings [3], [21], [28]. Clearly, such an instability, in particular when it even changes downstream classification results, is hard to tolerate, even more so if the database changes often and hence the embedding needs to be updated frequently. However, we cannot just ignore changes in the database, and we need to insert new tuples into the embedding as they arrive to keep the downstream tasks up to date.

The most pragmatic and feasible solution to this dilemma is to freeze the embedding of the existing tuples in the database and only compute embeddings for new tuples as they arrive dynamically. Hence, we address the *stable* variant of the problem where the goal is to infer embeddings of new tuples *without changing the embedding of old ones*. Of course, the challenge is to do so while retaining high quality of the embedding. One may suspect that the performance of downstream tasks will suffer significantly when based on a stable embedding of a dynamically changing database. Yet we demonstrate that with the right embedding algorithms, this is not the case.

To be more precise, we study the following task that we refer to as the *stable database embedding* problem. We need to devise two algorithms. The first algorithm, applied in the *static phase*, takes as input a database D over a schema σ and learns a tuple embedding γ that maps every fact of D (i.e., occurrence of a tuple in a certain relation in D) to the vector space $\mathbb{R}^k$ (for some hyperparameter k). Note that this phase solves the task of *static database embedding*, that is, the embedding problem in its traditional sense. The second algorithm, applied in the *dynamic phase*, has access to D and the tuple embedding γ and takes as input a newly arrived tuple t that is not in D; the goal of this algorithm is to extend γ to $D \cup \{t\}$ by determining the value $\gamma(t)$. Importantly, the schema σ specifies the key and foreign-key constraints that can be used for understanding the actual foreign-key references that exist inside the database D and its future evolution.

Since we wish to keep the embedding of the existing tuples stable (for the aforementioned reasons), deleting tuples from
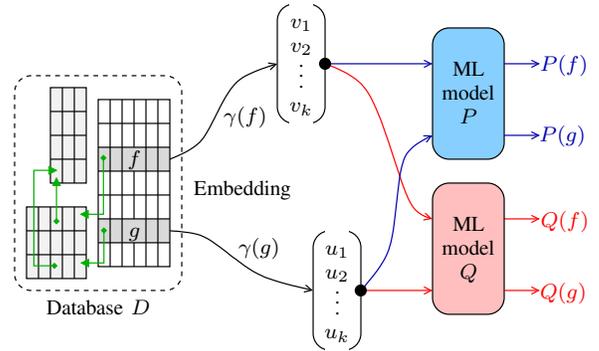
Fig. 1. Embedding tuples in a database. Different machine-learning models for various downstream tasks operate over the same vectors that the tuples are mapped to by the embedding. The embedding algorithm can utilize the database structure, particularly the foreign-key references.

a database is not an issue—we simply delete the tuples and their images under the embedding. This is why we focus on tuple insertion in our framework and experiments. Of course, eventually there may be a point where the database has changed so much that a completely new embedding has to be computed.

Contribution: We design two main solutions to the stable database embedding problem. The first is an adaptation of NODE2VEC, and the second is what we call the FORWARD algorithm.

a) NODE2VEC adaptation.: The first solution that we design is based on our adaption of NODE2VEC to relational database embeddings. We found empirically that this adaptation performs very well for the static embedding problem. For the stable database embedding problem, we devise a dynamic version of NODE2VEC that is based on the following idea. When we extend an existing embedding to new nodes, we sample relevant paths and continue the training of NODE2VEC from where it stopped while performing gradient descent *only on the embeddings of new nodes*.

b) The FORWARD algorithm.: The second solution we devise is a new algorithm, FORWARD (Foreign Key Random Walk Embeddings for Relational Databases), that is inherently built to accommodate the structure of relational databases and, importantly, to be extensible to dynamic databases. We demonstrate experimentally that it performs very well for the stable database embedding problem, which it was designed for, as well as the static database embedding problem. FORWARD draws from node embedding techniques based on random walks [15], [40] as well as the knowledge graph embedding algorithm RESCAL [37]. From each tuple in the database we start random walks in the database by repeatedly following foreign-key references. We embed tuples of the database depending on the similarity of the distributions of these walks, where we measure similarity in terms of a predefined similarity measures on the attribute values. Crucially, while learning the embedding of tuples, we also learn a separate similarity measure for each type (relational scheme) of a walk. This

similarity is encoded as a matrix describing an inner product. Once we have learned an initial static embedding, we can dynamically extend it to each new tuple by essentially solving a system of linear equations that constrain the distances to the existing tuples (or a random subset of these).

We describe our experiments over several downstream tasks of (binary and multi-label) column prediction from multiple benchmark databases over geographical and biological domains [13], [31], [36], [47]. We show that FORWARD performs very well even for the static database embedding problem. On the majority of the benchmarks, both FORWARD and our static adaptation of NODE2VEC clearly outperform the state-of-the-art baselines (see Table III).

Regarding the stable database embedding problem, both dynamic NODE2VEC and FORWARD perform well, though overall FORWARD shows a superior performance (as we show in Figure 5 and Table IV). It is also faster than NODE2VEC. Quite remarkably, the performance of FORWARD turns out to be rather stable as we add more and more tuples to the database. Even with 50%, sometimes even 80%, newly added tuples the drop in accuracy for column prediction is small (see Figure 5). Combined with the strong performance of FORWARD for the static embedding problem, this convincingly shows that FORWARD is a viable solution to the stable embedding problem.

In summary, our contributions are as follows. First, we define the problem of stable database embedding. Second, we devise an adaptation of NODE2VEC to dynamic databases. Third, we devise the FORWARD algorithm in both its static and dynamic versions. Fourth, we conduct a thorough experimental evaluation of our algorithms over a collection of tuple prediction benchmarks.

Organization: The remainder of the paper is organized as follows. After giving preliminary definitions and notation in Section II, we define the problem of stable database embedding in Section III. In Section IV we describe our NODE2VEC adaption, and in Section V the FORWARD algorithm. Finally, we present our experimental evaluation in Section VI and conclude in Section VII.

II. PRELIMINARIES

We focus on databases over schemas with key and foreign-key constraints. More precisely, a *database schema* σ consists of a finite collection of *relation schemas* $R(A_1, \ldots, A_k)$ where R is a distinct *relation name* and each A_i is a distinct *attribute name*. For simplicity, we assume that the attribute sets of distinct relations are disjoint. Each attribute A is associated with a *domain*, denoted $\mathsf{dom}(A)$. Each relation schema $R(A_1, \ldots, A_k)$ has a unique *key*, denoted $\mathsf{key}(R)$, such that $\mathsf{key}(R) \subseteq \{A_1, \ldots, A_k\}$.

A *foreign-key constraint* (FK) is an inclusion dependency of the form $R[\boldsymbol{B}] \subseteq S[\boldsymbol{C}]$ where R and S are relation names, $\boldsymbol{B} = B_1, \ldots, B_\ell$ and $\boldsymbol{C} = C_1, \ldots, C_\ell$ are sequences of distinct attributes of R and S, respectively, and $\mathsf{key}(S) = \{C_1, \ldots, C_\ell\}$.

A *database* D over the schema σ is a finite set of *facts* $R(a_1, \ldots, a_k)$ over the relation schemas $R(A_1, \ldots, A_k)$ of σ, so that $a_i \in \mathsf{dom}(A_i)$ for all $i = 1, \ldots, k$. In addition, such an a_i can be missing, in which case we assume that it is a distinguished *null* value (that belongs to none of the attributes domains) denoted by $\bot$. The fact $R(a_1, \ldots, a_k)$ is also called an *R-fact* and a σ-*fact*. We denote by $R(D)$ the restriction of D to its R-facts. For a fact $f = R(a_1, \ldots, a_k)$ over $R(A_1, \ldots, A_k)$, we denote by $f[A_i]$ the value a_i, and by $f[B_1, \ldots, B_\ell]$ the tuple $(f[B_1], \ldots, f[B_\ell])$. For a database D, the *active domain* of an attribute A (*w.r.t. to* D), denoted $\mathsf{adom}_D(A)$ or just $\mathsf{adom}(A)$ if D is clear from the context, is the set of values that occur in D for the attribute A, that is, $\mathsf{adom}(A) = \{f[A] \mid f \in R(D)\}$.

We require that the database D over σ satisfies the constraints of σ. In particular, for the key constraints we require every two distinct R-facts f_1 and f_2 must satisfy $f_1[C] \neq f_2[C]$ for at least one attribute $C \in \mathsf{key}(R)$, and both $f_1[C]$ and $f_2[C]$ are nonnull. Moreover, for every FK φ of the form $R[\boldsymbol{B}] \subseteq S[\boldsymbol{C}]$ and R-fact $f \in D$, if $f[\boldsymbol{B}]$ has no nulls[1] then there exists an S-fact $g \in D$ such that $f[\boldsymbol{B}] = g[\boldsymbol{C}]$; in this case, we say that f *references* g *via* φ (note that f references precisely one fact via φ).

Example 2.1: Figure 2 depicts an example of a movie database over a schema. The leftmost column of each relation includes tuple names for later reference, and is not considered part of the database itself. As conventional, keys are marked by underlining the key attributes. The FKs of each relation schema are given under the corresponding relation. For example, The MOVIES relation has the key constraint $\mathsf{key}(\text{MOVIES}) = \{\mathsf{mid}\}$ and the FK MOVIES[studio] $\subseteq$ STUDIO[sid]. The reader can verify that the FK is indeed satisfied by the database of the figure; for example, s03 is indeed the sid attribute of a fact of STUDIO, namely s_3. Finally, observe that the genre attribute of m_3 is missing, that is, $m_3[\mathsf{genre}] = \bot$. ∎

III. PROBLEM DEFINITION

As explained in the Introduction, the problem of stable database embedding consists of two tasks.

- *Static phase*: The goal is to derive an embedding of the tuples in the traditional sense. Formally, we are given a database D over a schema σ, and we wish to compute an embedding function $\gamma : D \to \mathbb{R}^k$ for some hyperparameter $k > 0$.
- *Dynamic phase*: Here the goal is to extend the embedding γ to a new fact f. Formally, we are given a database D over a schema σ, a precomputed embedding $\gamma : D \to \mathbb{R}^k$ and a new fact $f \notin D$. Our goal is to compute a new embedding $\gamma' : D \cup \{f\} \to \mathbb{R}^k$ such that $\gamma'(f') = \gamma(f')$ for all $f' \in D$. Hence, we only need to compute $\gamma'(f)$.

In the obvious manner, we can generalize the dynamic phase to a set (*batch*) $\{f_1, \ldots, f_\ell\}$ of new facts rather than just a

[1]Note that we adopt the convention that an FK is ignored in a fact that includes nulls in one or more of the referencing attributes.

MOVIES

	mid	studio	title	genre	budget
m_1	m01	s03	Titanic	Drama	200M
m_2	m02	s01	Inception	SciFi	160M
m_3	m03	s01	Godzilla	$\bot$	150M
m_4	m04	s03	Interstellar	SciFi	160M
m_5	m05	s02	Tropic Thunder	Action	90M
m_6	m06	s01	Wolf of Wall St.	Bio	100M

MOVIES[studio] $\subseteq$ STUDIOS[sid]

STUDIOS

	sid	name	loc
s_1	s01	Warner Bros.	LA
s_2	s02	Universal	LA
s_3	s03	Paramount	LA

COLLABORATIONS

	actor1	actor2	movie
c_1	a01	a02	m03
c_2	a04	a05	m04
c_3	a04	a03	m05
c_4	a01	a04	m06

COLLABORATIONS[actor1] $\subseteq$ ACTORS[aid]
COLLABORATIONS[actor2] $\subseteq$ ACTORS[aid]
COLLABORATIONS[movie] $\subseteq$ MOVIES[mid]

ACTORS

	aid	name	worth
a_1	a01	DiCaprio	230M
a_2	a02	Watanabe	40M
a_3	a03	Cruise	600M
a_4	a04	McConaughey	140M
a_5	a05	Damon	170M

Fig. 2. Database example. Key constraints are marked with underline and foreign-key references are specified under the corresponding relations.

single f, and the goal is to extend the embedding to all of the facts of this set.

The objective is to compute an embedding γ that represents the data in a way that makes it accessible for data analysis and machine learning algorithms. To evaluate an embedding algorithm, we empirically test it against various downstream learning tasks on the relational data. That is, we take the embedded data as input for a machine learning algorithm (e.g., an artificial neural network) and measure how well it performs. To evaluate the performance of the dynamic embedding algorithm, we measure the performance not only over the original database, on which the machine learning algorithm was trained, but also on the tuples newly added as the database changes dynamically.

We need to devise two algorithms: one for the static phase and one for the dynamic phase. Note that these algorithms might (and actually should) depend on each other: the algorithm for the static phase not only needs to perform well as a static embedding algorithm, but also allow and enhance the effectiveness of the algorithm for the dynamic phase.

Example 3.1: Let D' be the database of Figure 2 and $D = D' \setminus c_4$. Hence, D' is obtained from D by inserting the fact

$$c_4 = \text{COLLABORATIONS}(\texttt{a01}, \texttt{a04}, \texttt{m06}).$$

In static phase, for the input D we need to compute a mapping $\gamma : D \to \mathbb{R}^k$. In the dynamic phase, when inserting c_4 into D we wish to extend γ to c_4 by computing $\gamma(c_4)$ without changing it on the facts of D; for example, $\gamma(c_1)$ and $\gamma(m_1)$ should remain intact. To determine $\gamma(c_4)$ we can utilize the semantic knowledge that the new c_4 references the existing a_1, a_4 and m_6. ∎

IV. DYNAMIC DATABASE EMBEDDINGS BASED ON NODE2VEC

The first algorithm we propose for the stable database embedding problem is an adaptation of NODE2VEC, a well-known algorithm for node embeddings of graphs [15]. In addition to the original NODE2VEC, we incorporate ideas from [11], where NODE2VEC was adapted to embedding static relational databases, and ideas from [32], where NODE2VEC was adapted for node embeddings of dynamic graphs. A crucial novelty in our approach is that we incorporate foreign key constraints by identifying certain nodes in the graph model of the relational database to which we apply NODE2VEC.

We start with a description of the static embedding algorithm. For the rest of the section, let D be a database of schema σ. We build a bipartite graph G_D such that one side represents the facts and the other side represents the attribute values that occur in the facts, as illustrated in Figure 3 for a fragment of the database of Figure 2. The graph is related to, but slightly different from the graph in [11]. For each relation schema $R(A_1, \ldots, A_k)$ in σ, each attribute A_i, and each value a of A_i that occurs in $R(D)$, we add a node $u(R, A_i, a)$ to G_D. For each fact $f = R(a_1, \ldots, a_k)$ in $R(D)$ we add a node $v(f)$ and, for $i = 1, \ldots, k$ we add edges between $v(f)$ and $u(R, A_i, a_i)$. So far, the components of the graph for different relations R, S in σ are disconnected and completely independent. In a second step, we use the foreign key constraints to connect these subgraphs. Suppose we have an FK $R[B_1, \ldots, B_\ell] \subseteq S[C_1, \ldots, C_\ell]$. Then for all $i \in [\ell]$ and all values $a \neq \bot$, we identify the two nodes $u(R, B_i, a)$ and $u(S, C_i, a)$ if they both exist. Figure 3 illustrates the
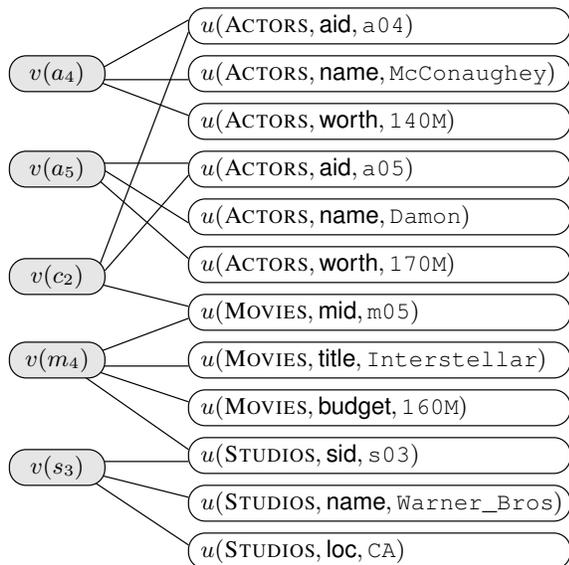
Fig. 3. A partial graph of the database in Figure 2.

construction.

Note that we are careful in the way we introduce connections in our graph: if the same value occurs in different columns or different relations, there is no connection between the two occurrences, except if they are linked by a foreign key. To understand why this is the correct way of modeling the database by a graph, just imagine that in our movie database of Figure 2 we had a tuple m_7 representing a movie called "Universal". Then the occurrence of `Universal` in the title column of the MOVIES relation would be completely independent of the occurrence of `Universal` in the name column of the STUDIOS relation. So it is appropriate to represent these two by distinct nodes in the graph. However, the occurrence of `s01` in the studio column of the MOVIES relation is referring to the same object as the occurrence of `s01` in the sid column of the STUDIOS relation, and therefore it is justified to identify the two nodes.

A. Extension to the Dynamic Setting

We now describe our extension of NODE2VEC to the dynamic setting. Beres et al. [5] suggested an online extension of NODE2VEC on graphs in the setting where the graph arrives as a stream of edges, and Mahdavi et al. [32] suggested a method based on evolving walks generation, in the setting where the input is a series of graphs that arrive at discrete timestamps. Neither of these applies directly to the stable embedding problem, because is does not allow to freeze of the old embeddings (i.e., in these methods the old embeddings change when new nodes arrive). Nevertheless, we use similar ideas as [32].

Our method is based on sampling new walks and continuing the gradient descent. Suppose we have already trained a node embedding for the graph G_D associated with our current database D. As a new fact f arrives, we update the graph to $G_{D'}$ for the new database $D \cup \{f\}$. Note that the new

nodes in $V(G_{D'}) \setminus V(G_D)$ are $v(f)$ and possibly so nodes $u(R, A, a)$ for values a in f that have not been present before. We sample new random walks starting at the new nodes. Then we train a new NODE2VEC model for $G_{D'}$ taking the old model for G_D and a random initialization for the new nodes as the intialization. As we train the new model by standard gradient descent techniques, we freeze the old nodes and only update the embedding on the new nodes.

V. THE FORWARD ALGORITHM

We now present our second algorithm for the stable database embedding problem. Our embedding algorithm will use *random walks* along foreign-key constraints to incorporate the structure of the database into the embedding, and it will use *kernelized domains* to incorporate the semantics of the values appearing in the facts. Before we describe the algorithm, we establish the background on random walks and on kernelized domains in Sections V-A and V-B.

Throughout this section, we fix a database schema σ, and we always assume that D is a database of schema σ.

A. Random Walks over Database Facts

We consider random walks over database facts, where the transition from one fact to another follows a pattern (or scheme) of FK (forward or backward) references. Formally, a *walk scheme* is a sequence s of the form

$$R_0[\mathbf{A}^0] \!-\! R_1[\mathbf{B}^1], R_1[\mathbf{A}^1] \!-\! R_2[\mathbf{B}^2], \dots \quad (1)$$
$$, R_{\ell-1}[\mathbf{A}^{\ell-1}] \!-\! R_\ell[\mathbf{B}^\ell]$$

such that for all $k = 1, \dots, \ell$, either $R_{k-1}[\mathbf{A}^{k-1}] \subseteq R_k[\mathbf{B}^k]$ is an FK or $R_k[\mathbf{B}^k] \subseteq R_{k-1}[\mathbf{A}^{k-1}]$ is an FK. We say that s has *length* ℓ, that it *starts from* R_0 and that it *ends with* R_ℓ.

Example 5.1: Figure 4 depicts nine walk schemes, $s_1, \dots, s_9$ that start from the ACTOR relation. The reader can verify that these are all of the walk schemes of length at most three that start from ACTOR. For illustration, let us consider the scheme s_5 from the figure. In our notation, this walk scheme is written as

ACTORS[aid]—COLLABORATIONS[actor2],
COLLABORATIONS[movie]—MOVIES[mid].

Note that s_1 ends with COLLABORATIONS while s_5 ends with with MOVIES. ∎

A *walk* with the scheme s is a sequence $(f_0, \dots, f_\ell)$ of facts such that f_k is an R_k-fact and $f_{k-1}[\mathbf{A}^{k-1}] = f_k[\mathbf{B}^k]$ for all $k = 1, \dots, \ell$. We say that $(f_0, \dots, f_\ell)$ *starts from*, or has the *source*, f_0, and that it *ends with*, or has the *destination*, f_ℓ.

Example 5.2: Continuing our example, consider again the walk scheme s_5 of Figure 4. Starting at the fact a_1 of Figure 2, there are two walks that follow the scheme s_5, namely (a_1, c_1, m_3) and (a_1, c_4, m_6). ∎

Note that we allow walk schemes and walks of length zero. For each relation R there is a scheme of length zero that starts and ends in R. The walks of this scheme have the form (f_0) and simply end directly at the start fact f_0 in R.

Let s be a walk scheme as written in (1). By a *random walk* with the scheme s we refer to the walk obtained by uniformly selecting the next valid fact in the walk. More formally, let $f_0 = f$ be an R_0-fact. We denote by $\mathcal{W}(f, s)$ the distribution over the walks with the walk scheme s where each walk is sampled by starting from f_0 and then iteratively selecting f_k, for $k = 1, \ldots, \ell$, randomly and uniformly from the set $\{f \in R_k \mid f[\boldsymbol{B}^k] = f_{k-1}[\boldsymbol{A}^{k-1}]\}$. We denote by $d_{f,s}$ the random variable/element that maps each walk in $\mathcal{W}(f, s)$ to its destination, that is, the last fact in the walk. Then for a fact $g \in R^k(D)$, the probability that a walk sampled from $\mathcal{W}(f, s)$ ends with g is $\Pr(d_{f,s} = g)$. Observe that for every attribute A of R_k we get the random variable $d_{f,s}[A]$ that forms the value of the random walk's destination in the attribute A. Given a start fact f_0 in the start relation R_0 of walk scheme s, one can compute the distribution $d_{f,s}$ through a simple breadth first search along the sequence of foreign keys specified by s.

Example 5.3: Recall from Example 5.2 that the walks (a_1, c_1, m_3) and (a_1, c_4, m_6) are the only two walks that follow the scheme s_5 (Figure 4) and start from a_1 (Figure 2). Therefore, for the random variable d_{a_1,s_5} it holds that $\Pr(d_{a_1,s_5} = m_3) = 0.5$ and $\Pr(d_{a_1,s_5} = m_6) = 0.5$. Moreover, we have the following.

$$\Pr(d_{a_1,s_5}[\mathsf{budget}] = 150\mathrm{M}) = 0.5$$
$$\Pr(d_{a_1,s_5}[\mathsf{budget}] = 100\mathrm{M}) = 0.5$$
$$\Pr(d_{a_1,s_5}[\mathsf{genre}] = \mathrm{Bio}) = 1.0$$

Hence, each of $d_{a_1,s_5}[\mathsf{budget}]$ and $d_{a_1,s_5}[\mathsf{genre}]$ indeed defines a distribution over attribute values. ∎

Recall that databases (in real life and in our formal framework) may have missing values. A random walk starting at f might end at a fact R_ℓ-fact g which has no known value for an attribute A of R_ℓ. Therefore, the random variable $d_{f,s}[A]$ can in theory assume the value $\perp \notin dom(A)$. As a convention, we define the probability distribution of $d_{f,s}[A]$ as the posterior distribution after $d_{f,s}[A] \neq \perp$ is known. With this modification we enforce $d_{f,s}[A] \in dom(A)$. This will be crucial in Section V-B, where we define similarity measures for $d_{f,s}[A]$ based on $dom(A)$. If all walks from f with scheme s end at facts g with $g[A] = \perp$, then $d_{f,s}[A]$ does not exist and is not considered by FORWARD. This also includes the case where no walks with scheme s exist from start f.

B. Kernelized Domains

In addition to the structural information carried by the foreign-key random walks described in the previous section, our embeddings are based on similarities between values occurring in the tuples. Formally, we assume that these similarities are given by *kernels*. For every attribute A occurring in the database schema σ we assume that we have a symmetric binary function κ_A mapping pairs of elements from $\mathsf{dom}(A)$ to the nonnegative reals. Intuitively, $\kappa_A(a, b)$ measures the similarity between elements $a, b \in \mathsf{dom}(A)$. Formally, κ_A needs to satisfy certain properties that turn it into a kernel function. This means that there is an embedding $\alpha_A : \mathsf{dom}(A) \to \mathcal{H}_A$,

where $\mathcal{H}_A$ is a Hilbert space, that is, a (possibly infinite) vector space with an inner product, which we denote by $\langle \cdot, \cdot \rangle$, and κ_A is defined by $\kappa_A(a, b) = \langle \alpha_A(a), \alpha_A(b) \rangle$. Importantly, the mapping α_A is only implicit, our algorithms only need access to the kernel function κ_A. For details on kernels, we refer the reader to [45, Chapter 16], [22].

For many natural domains, such a kernel function can be obtained by standard embedding techniques. For examples, we can use word embeddings, possibly tailored towards specific applications, for natural language domains [27], or molecule embeddings in the case of a domain of molecules [23]. For numerical domains, we can use a Gaussian kernel $\kappa(a, b) = \exp\left(-(a - b)^2/2\upsilon\right)$ for some "variance" $\upsilon > 0$. As a fallback, we can always use the *equality kernel* defined by $\kappa(a, a) = 1$ and $\kappa(a, b) \neq 0$ for $a \neq b$. We would typically use the equality kernel for finite categorical domains and domains consisting of identifiers, which have no semantic meaning.

Kernel functions offer a straightforward way of encoding domain knowledge by modeling the similarity of the domain values. Kernels are also helpful when dealing with noisy data. For example, on text, kernels based on the edit distance can be used to smooth out random typos. In the following we use these kernel functions to define similarity measures for the random variables $d_{f,s}[A]$.

Let s be a walk scheme of length ℓ from R to R'. Let A be an attribute of R' and let f and f' be two distinct R-facts. Under our assumption, $d_{f,s}[A]$ and $d_{f',s}[A]$ are random variables over a shared kernelized domain $dom(A)$. We utilize this to quantify the similarity between $d_{f,s}[A]$ and $d_{f',s}[A]$ with respect to the underlying kernel κ_A. To this end we define the *Expected Kernel Distance* KD is the expected distance between two random values selected independently at random:

$$\mathrm{KD}(d_{s,f}[A], d_{s,f'}[A]) = \underset{\mathcal{W}(f,s) \times \mathcal{W}(f',s)}{\mathbb{E}}[\kappa_A(d_{s,f}[A], d_{s,f'}[A])] \quad (2)$$

C. Embedding

Let us now describe our static embedding algorithm. Recall that D denotes a database of schema σ. We describe how to embed the tuples of a single relation $R(D)$; of course we can then apply the method to embed all relations in D. So our goal is to compute a vector embedding $\varphi : R(D) \mapsto \mathbb{R}^d$ of dimension $d \in \mathbb{N}$. Intuitively, FORWARD embeddings aim to model the similarity of the random walk destinations $d_{s,f}$ for all R-facts f and all walk schemes s starting from R of length up to a certain walk length $\ell_{\max}$. Formally, we define $\mathcal{T}(R, \ell_{\max})$ as the set of all pairs (s, A) such that s is a walk scheme of length at most $\ell_{\max}$ starting from R, and A is an attribute of the destination R' of s that is not involved in any foreign-key constraints. Together with φ we compute an auxiliary embedding $\psi : \mathcal{T}(R, \ell_{\max}) \to \mathbb{R}^{d \times d}$ that maps each pair (s, A) to a symmetric matrix $\psi(s, A)$.

Ideally, our goal is to find φ, ψ satisfying

$$\varphi(f)^\top \psi(s, A)\varphi(f') = \mathrm{KD}(d_{s,f}[A], d_{s,f'}[A]) \quad (3)$$
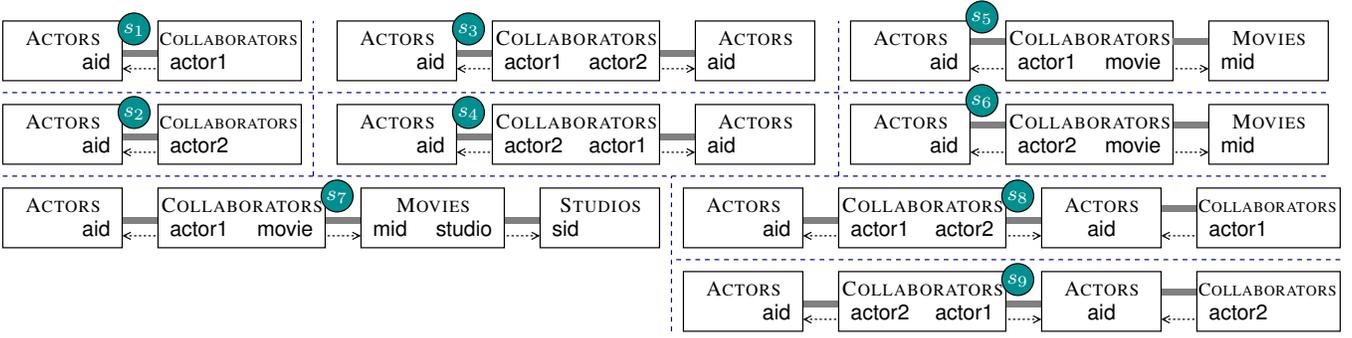
Fig. 4. All walk schemes of length at most three, for the database schema of Figure 2, that start from the Actor relation.

for all $f, f' \in R(D)$ and $(s, A) \in \mathcal{T}(R, \ell_{\max})$. Usually, we will not be able to achieve this. Instead, we aim at jointly minimising

$$\left| \varphi(f)^\top \psi(s, A) \varphi(f') - \mathrm{KD}(d_{s,f}[A], d_{s,f'}[A]) \right| \quad (4)$$

for all f, f', s, A. Then the embedding φ is the primary output of FORWARD.

Essentially, what we do here is learn an inner product $\langle \cdot, \cdot \rangle_{S,A}$ on the latent space of the embedding φ defined by $\langle x, y \rangle_{S,A} = x^\top \psi(s, A) y$ for all s and A, and we try to find a φ such that the similarity of facts f and f' with respect to this inner product matches the similarity between the random variables $d_{f,s}[A]$ and $d_{f',s}[A]$ with respect to the underlying kernel $\kappa_{R\ell.A}$. The idea of learning auxiliary inner products (or their matrices) jointly with the actual embedding goes back to the knowledge graph embedding algorithm RESCAL [37].

It remains to describe the optimisation procedure that we use for minimising (4).

D. Optimization

We utilize gradient descent to optimize FORWARD embeddings. That is, the embeddings φ and ψ are initialized randomly and then stochastic gradient descend is used to minimize the ℓ_2-loss of the objective in Equation 4.

During training, we sample a large number of tuples of the form (f, f', s, A, g, g'). Here, f and f' are R facts from the database and $(s, A) \in \mathcal{T}(R, \ell_{\max})$. The R_ℓ facts g and g' are the destinations of random walks with scheme s sampled for f and f', respectively. To this end, we specify a hyperparameter $n_{\text{samples}} \in \mathbb{N}$. For each R-fact f and each $(s, A) \in \mathcal{T}(R, \ell_{\max})$ for which $d_{s,f}[A]$ exists we uniformly sample n_{samples} of the form (f, f', s, A, g, g') with $f' \neq f$. If the total number of such samples is less than n_{samples}, then we just use all samples without duplicates. Using these samples, we minimize the following term with stochastic gradient descent:

$$\mathcal{L} = \frac{1}{2} |\varphi(f)^\top \psi(s, A) \varphi(f') - \kappa_{R_\ell.A}(g[A], g'[A])|^2. \quad (5)$$

This objective uses the value $\kappa_{R\ell.A}(g[A], g'[A])$ as a (stochastic) estimate of $\mathrm{KD}(d_{s,f}[A], d_{s,f'}[A])$. This procedure avoids computing $\mathrm{KD}(d_{s,f}[A], d_{s,f'}[A])$ explicitly, which would be prohibitive in large databases.

E. Extending Embeddings to New Tuples

We consider the situation where a new R-fact f_{new} is inserted into the database D, and our goal is to extend the existing embedding φ over D to incorporate f_{new}. Hence, our goal is to determine the vector $\varphi(f_{\text{new}}) \in \mathbb{R}^d$.

Let f_{old} be an R-fact such that the embedding $\varphi(f_{\text{old}}) \in \mathbb{R}^d$ is already known and let $(s, A) \in \mathcal{T}(R, \ell_{\max})$. Then $\psi(s, A) \in \mathbb{R}^{d \times d}$ is a known matrix (that we referred to in Equation (3)). We wish for our new embedding $\varphi(f_{\text{new}})$ to satisfy the objective in Equation (3) with respect to $f = f_{\text{new}}$ and $f' = f_{\text{old}}$:

$$\varphi(f_{\text{new}})^\top \cdot \psi(s, A) \cdot \varphi(f_{\text{old}}) \overset{!}{=} \mathrm{KD}(d_{s,f_{\text{old}}}[A], d_{s,f_{\text{new}}}[A]) \quad (6)$$

Hence, we obtain a linear equation $\varphi(f_{\text{new}})^\top \cdot c = y$ where $c \in \mathbb{R}^d$ and $y \in \mathbb{R}$ are known. If we stack these linear equations for many choices of (f', s, A), then we obtain an overdetermined system of linear equations that we can (approximately) solve for $\varphi(f_{\text{new}})$.

We randomly sample a sufficiently large number of such triples. In particular, we sample $n_{\text{samples}}^{\text{new}} \in \mathbb{N}$ distinct samples for each $(s, A) \in \mathcal{T}(R, \ell_{\max})$ where $n_{\text{samples}}^{\text{new}}$ is a hyperparameter. Let k be the total number of drawn samples and let (f_i, s_i, A_i) be the i-th sample with $i \in [k]$. We define the matrix $C \in \mathbb{R}^{k \times d}$ such that

$$C_i = \psi(s_i, A_i) \cdot \varphi(f_i). \quad (7)$$

Define $b \in \mathbb{R}^k$ to be the vector with

$$b_i = \mathrm{KD}(d_{s_i, f_i}[A_i], d_{s_i, f_{\text{new}}}[A_i]). \quad (8)$$

To obtain a new embedding for f_{new} we solve for

$$C \cdot \varphi(f_{\text{new}}) = b. \quad (9)$$

Thus, we can infer embeddings for novel data simply by solving systems of linear equations. Note that we aim to find an approximate solution, since there are no exact solution for overdetermined linear systems in general. Any standard method for solving such systems can be applied. In our case, we use the pseudoinverse C^+ of C to obtain a solution that is optimal in the Euclidean norm:

$$\varphi(f_{\text{new}}) = C^+ \cdot b \quad (10)$$

F. Hyperparameters

The main hyperparameters are the embedding dimension d, the maximum walk length $\ell_{\max}$ and the number of samples n_{samples}, as described in Section V-D. When extending a FORWARD embedding to new tuples $n_{\text{samples}}^{\text{new}}$ is an additional hyperparameter. The batch size, learning rate and number of epochs of the gradient descent training are additional parameters. The domain kernels $\kappa_{R.A}$ for each domain can also be viewed as hyperparameters. However, most common data types allow for simple default choices.

VI. EXPERIMENTAL EVALUATION

We now describe our experimental study. The goal of this study is to evaluate the quality of the embeddings produced, in both the static phase and (more importantly) the dynamic phase. Our quality evaluation is via a collection of downstream tasks of tuple classification. (We can also view the tuple-classification task as column prediction, where the class of the tuple can be seen as a new column.) We focus on databases that involve multiple relations, and particularly relations that are not the target of the downstream tasks; yet these relations can be used as context for inferring embeddings, as we do in our proposed embedding FORWARD. We first describe the downstream tasks (Section VI-A), the compared methods (Section VI-B), and the general runtime setup (Section VI-C). We then describe our experiments on the static setting (Section VI-D) and the dynamic setting (Section VI-E). Finally we discuss the execution times of the algorithms (Section VI-F) and the main conclusions of the experimental section (Section VI-G).

A. Datasets and Tasks

We now describe the datasets (tasks) of our experiments. The information on the structure of the datasets is summarized in Table I. Each dataset is a database of multiple relations, where one relation contains an attribute that we wish to predict. Hereafter, we refer to this relation as the *prediction relation*. Note that neither FORWARD nor NODE2VEC see the predicted attribute during training.

1) Hepatitis: This database is from the 2002 ECML/PKDD Discovery Challenge.[2] We use the modified version of Neville et al. [36]. The goal in this task is to predict the type column, which is either *Hepatitis B* or *Hepatitis C* based on medical examinations. There are in total 206 instances of the former and 484 cases of the latter. The relation with the predicted column contains, in addition to the type classification, the age, sex and identifier of the patient. The other relations contain the rest of the medical data. The dataset contains seven relations with a total of 26 attributes and 12,927 tuples.

2) Mondial: This dataset contains information from multiple geographical resources [33]. We predict the religion of a country. There are 114 countries classified as *Christian* and 71 as *non-Christian*. The prediction is based on a variety of fields such as the language, population, geography, and government

of the country. The target relation, where we predict a column, is binary—it contains only the name of the country and the (predicted) classification. The dataset contains 40 different relations with a total of 167 attributes and 21,497 tuples. We use the whole database and use the TARGET relation as the prediction relation as previously done by Bina et al. [6].

3) Genes: This dataset is from the KDD 2001 competition [12], and contains data from genomic and drug-design applications. We predict the localization of the gene, based on biological data, with 15 different labels. The prediction relation contains only the class and an identifier for the gene, while the rest contain the biological data such as the function, gene type, cellular location and the expression correlation between different genes. The dataset contains 3 relations with a total of 15 attributes and 6,063 tuples. We remove two tuples which have a unique class to prevent split in-balances during cross-validation.

4) Mutagenesis: This dataset contains data on the mutagenicity of molecules on Salmonella typhimurium [13]. We predict the mutagenicity (the mutagenic attribute) of the molecules, based on chemical properties of the molecule, with 122 positive samples and 63 negative samples. The prediction relation contains the binary class, molecule ID, and some of the chemical data, while the other relations contain more chemical data and information about the relations between the molecules. The dataset contains 3 relations with a total of 14 attributes and 10,324 tuples. Note that we do not use any external features, in contrast to some past methods for this dataset [31] that we revisit later on.

5) World: This dataset contains data on states and their cities. We predict the continent of a country with 7 different labels. The prediction is based on general data on the country such as population, GNP, Capital city and information on the spoken languages and cities. The dataset contains 3 relations, with a total of 24 attributes and 5,411 tuples.

B. Compared Methods

We compare between the following alternatives.

- S.o.A.: These are state-of-the-art methods that solve the multi-relational classification problem without using an embedding, such as multi-relational decision trees and forests [4], [6], multi-relational Bayes nets [42] and Inductive Logic Programming (ILP) [31], [50]. This applies only to the static experiment.
- N2V: In the static phase, this is the NODE2VEC method, with our own implementation that is based on the original paper. In the dynamic phase, we use our adaption of NODE2VEC described in Section IV-A
- FWD: FORWARD embeddings, as defined in Section V.

In the next section, we discuss the implementation of the N2V and FORWARD variants.

In all datasets and experiments, we have a full separation between the embedding process and the downstream task. This means that we generate the embedding independently from the task (as opposed to training for the task), and then use these embeddings as the input to a downstream classifier (that

Dataset	Prediction Rel.	Prediction Attr.	#Samples	#Relations	#Tuples	#Attributes
Hepatitis	Dispat	type	500	7	12927	26
Genes	Classification	localization	862	3	6063	15
Mutagenesis	Molecule	mutagenic	188	3	10324	14
World	Country	continent	239	3	5411	24
Mondial	Target	target	206	40	21497	167

sees only the embeddings and none of the other database information).

Specifically, we train and apply an SVM classifier (Scikit-learn's SVC implementation) as the downstream machine-learning architecture, in both experiments. Performance assessment is conducted via k-fold cross validation with $k = 10$ folds.

C. Experimental Setup

1) Implementation: We implemented NODE2VEC on our own based on its original publication [15]. For the state-of-the-art results (S.o.A.) we used the reported numbers. Both NODE2VEC and FORWARD (in both static and dynamic versions) are implemented in Python using PyTorch [38], Numpy [20] for the numerical operations, Scikit-learn [39] for the downstream classifiers and validation, and NetworkX [18] for the graph implementations. The code is publicly available on GitHub[3].

Note that in the FORWARD implementation we embed only the relation that contains the tuples that we wish to classify. We use the default kernels in all of our experiments: Gaussian distance for numbers, and equality for all other data types.

2) Hyperparameters: The hyperparameters of the FORWARD and NODE2VEC implementations are listed in Table II. Note that we use different hyperparameters for the Genes dataset, since the distribution of values between the relations is different compared to the rest of the datasets. In the dynamic experiment the number of epochs for the second

training phase of NODE2VEC was set to 5 and $n_{\text{samples}}^{\text{new}}$ was set to 2500 for all FORWARD runs.

3) Hardware: We run the experiments on a server with 1 Intel i7 processor, 64 GB RAM, and an Nvidia RTX 2070 GPU. We utilize the GPU in all methods to achieve better runtimes.

D. Results for Static Database Embeddings

Table III summarizes the results for the static case. Also note that for each of the ten folds, we train a new embedding. Hence, to account for the randomness of both the folds and the embeddings, we report the standard deviation ($\pm x$) next to each number. Note that the embedding methods always see the full database, and the downstream classifier uses the different splits. For the state-of-the-art (S.o.T.), we use the reported results from the publications that we mention in the table.

As can be seen in Table III, both our methods perform better than the state-of-the-art methods on most datasets and are competitive with dataset-specific methods even without using external knowledge. Furthermore, the experiment shows that the embeddings produced with FORWARD are as good as the embeddings produced by the NODE2VEC based algorithm for static classification.

We also offer some interesting insights about the algorithms: First, NODE2VEC preforms better on Hepatitis and World - two datasets where most of the information lays in *categorical* data. This is in-fact expected as the graph that is used by NODE2VEC encodes this data very well. Second, the standard

[3]https://github.com/toenshoff/DynamicDBEmbedding

Alg.	Param.	Value
FORWARD	embedding dim. (k)	100
	#samples (n_{samples})	5,000
	batch size	50,000
	max walk len. (ℓ_{max})	1–3
	#epochs	5–10
NODE2VEC	embedding dim.	100
	#walks per node	40
	#steps per walk	30
	context window	5
	#neg/#pos samples	20
	batch size	40,000
	#epochs	10

Task	FORWARD	N2V	S.o.A.
Hepatitis	84.20% ± 4.94	**93.60%** ± 2.5	84.00%[(*)] [6]
Genes	**97.91%** ± 0.87	97.19% ± 1.25	85.00% [4]
Mutagenesis	**90.00%** ± 7.96	88.23% ± 4.56	**91.00%**[(**)] [50]
World	85.83% ± 5.34	**94.00%** ± 4.4	77.00% [50]
Mondial	80.95% ± 6.73	77.62% ± 5.24	**85%** [42]

[(*)] 95% achieved in a method specific for the Hepatitis dataset [2].

[(**)] 96% achieved in a method specific for the Mutagenesis dataset [31].

(a) Genes

(b) Hepatitis

(c) World

(d) Mondial

(e) Mutagenesis
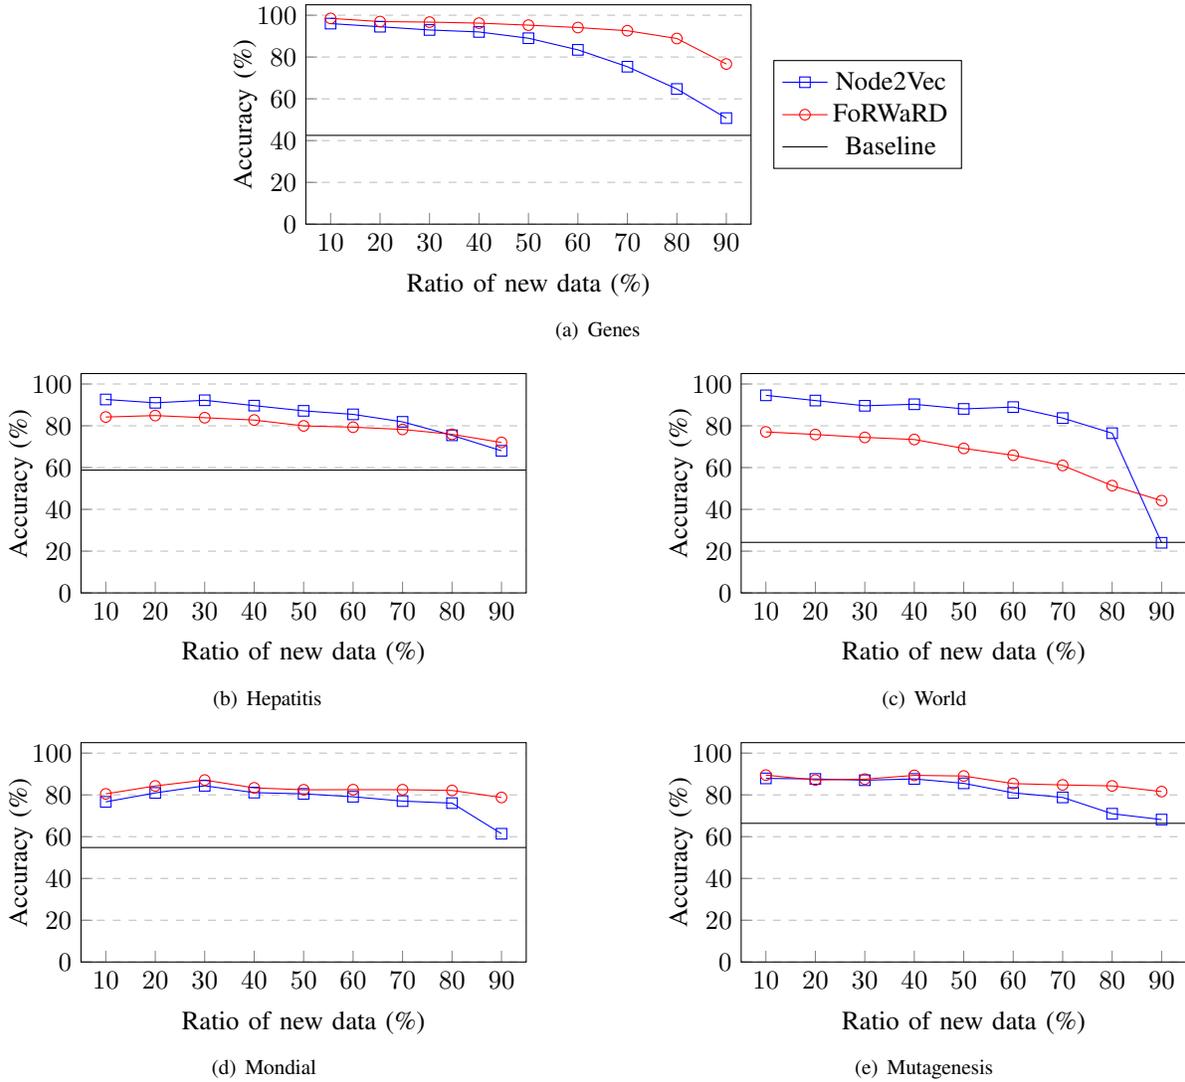
Fig. 5. Results for the dynamic experiment - baseline is the accuracy of guessing the most common class. We report the accuracy as a function of the percentage of new tuples.

deviations of all methods are on the same scale in all datasets. Third, both NODE2VEC and FORWARD excel at the Genes dataset, managing to capture the database's structure almost perfectly. Generally, both methods are competitive with the state-of-the-art in all of our datasets.

All in all, we can conclude that both NODE2VEC and FORWARD achieve very good results on these datasets, and that the structure of the database affects the quality of the embeddings - and both manage to capture the structure very well.

E. Results for Dynamic Database Embeddings

1) Experiment Description: In this experiment we aim to measure how well we can produce new embeddings when new data arrives. We test that by splitting the database into two parts and treat the first part as the static database, and the other part as the new data that arrives. The experiment is comprised of five steps:

1) Partition all facts of the database into two sets $\mathcal{F}_{old}$ and $\mathcal{F}_{new}$. The facts in $\mathcal{F}_{new}$ are then removed from the database.
2) Train an embedding *only* on the static part of the split.
3) Label the vectors in this embedding using the correct class labels and train a downstream classifier on these labelled data.
4) Add the facts from from $\mathcal{F}_{new}$ back to the database to simulate the arrival of new data. Generate new embeddings for these new facts.
5) Evaluate the trained classifier only on the embeddings of "new" data.

It is important to note that, by (5), the accuracy results we obtain refer exclusively to new data added after the embedding was trained.

The first question that arises is how to split the data in a way that simulates a real-life scenario. In real-world scenarios,

	All at Once		One by One	
Task	NODE2VEC	FORWARD	NODE2VEC	FORWARD
Hepatitis	**93.34**% $\pm$ **2.70**	82.20% $\pm$ 4.94	**92.60**% $\pm$ **2.37**	84.20% $\pm$ 5.02
Genes	94.50% $\pm$ 1.89	**97.91**% $\pm$ **0.87**	96.20% $\pm$ 0.78	**98.49**% $\pm$ **0.53**
Mutagenesis	87.58% $\pm$ 7.80	**90.00**% $\pm$ **6.84**	87.89% $\pm$ 7.82	**89.47**% $\pm$ **6.66**
World	**91.25**% $\pm$ **4.95**	87.50% $\pm$ 3.73	**94.58**% $\pm$ **4.58**	77.08% $\pm$ 5.67
Mondial	77.62% $\pm$ 6.75	**80.00**% $\pm$ **7.32**	76.67% $\pm$ 7.20	**80.47**% $\pm$ **7.20**

Task	NODE2VEC	FORWARD
Hepatitis	**189**	540
Genes	**78**	204
Mutagenesis	**166**	230
World	**219**	440
Mondial	**462**	810

new facts tend to arrive in batches that are distributed across all relations. For our example database from Figure 2, a new fact that is added to the MOVIES relation may be accompanied by new actors or new collaborations that are also added to the database. In our experiments we would like to simulate this behavior, where the new facts that arrive are are semantically related. Note that such a setup is more challenging than a simple random partition, since the new facts are less connected to the old data.

Formally, we compute the partition as follows. We first partition the prediction table according to a specified ratio of *old* and *new* tuples. This split is randomly chosen and stratified, i.e. the classes of the downstream task are (roughly) equally distributed in both partitions. We then iteratively remove the *new* tuples from the prediction relation in a random order. We remove each tuple with an "On Delete Cascade" deletion, which will automatically fix the foreign-key constraints throughout the database. In particular, data that is only referenced by the tuple that is being deleted is also removed from the database. The tuples that remain after these deletions form $\mathcal{F}_{\text{old}}$ while the deleted tuples form $\mathcal{F}_{\text{new}}$.

Example 6.1: Consider again the database of Figure 2 and suppose that our prediction relation is COLLABORATIONS. If we remove the fact c_1, then we will also remove the fact m_4 (Interstellar) from the MOVIES relation, and the tuple a_2 (Watanabe) from the ACTORS relation. Note that we will *not* remove a_1 (DiCaprio) as it is still connected to non-removed facts in other relations (specifically, it is connected to c_4 in the COLLABORATIONS relation). ∎

After computing the initial embedding of $\mathcal{F}_{\text{old}}$ we add the facts of $\mathcal{F}_{\text{new}}$ back to the database and extend the embedding to this new data. By default, we add the deleted facts from the prediction table *one-by-one* in the inverse order of their

deletion. With each new tuple in the prediction table we also add all the referenced tuples that were removed in the "On Delete Cascade" deletion. We then extend the embedding to these new tuples before the next tuple is added to the prediction table. Note that in this setting we do not recompute the paths starting at the old tuples, neither for NODE2VEC nor for FORWARD. The distribution of paths originating at the old tuples is changing as new data is added. However, recomputing these paths for each new tuple would slow down both methods substantially.

Additionally, we also study the case where all new tuples are added at once. We refer to this as a *all-at-once* embedding extension. In this setting we do recompute paths in the old data, since this only needs to be done once.

In our experiment we vary the ratio of old and new facts to understand how the relative amount of data effects the quality of the induced embeddings for the new data. For each tested ratio, we run the experiment 10 times with different partitions of the database. We measure the evaluation accuracy of the downstream classifier on the new data and the corresponding standard deviation across the 10 runs.

2) Results: Figure 5 provides the accuracies achieved in the one-by-one extension of the embeddings. We omit the results of the all-at-once extension in this figure, as the accuracy results are very similar in both versions. The term "baseline" here means the accuracy obtained by always predicting the most common class. Both FORWARD and NODE2VEC have similar performance and are able to produce useful embeddings for new tuples. As the percentage of new data increases the performance of both methods tends to decrease. However, this decrease is fairly slow and only becomes more significant when over 50% of the data is removed.

Table IV compares the two methods in both experimental setups (all-at-once vs one-by-one). Here, we look at the specific scenario where 10% of tuples are new data. Recall that in the all-at-once setup we do recompute paths for all tuples prior to extending the embedding to the new data. This is not done in the one-by-one experiment. One would expect the performance of both methods to be better in the all-at-once setup, since it takes into account paths which start at old tuples and traverse new data. This information is otherwise neglected. Surprisingly, the results are very similar in both setups. Only on the World dataset FORWARD achieves a significantly lower accuracy in the one-by-one setting when

Task	All at Once		One by One	
	NODE2VEC	FORWARD	NODE2VEC	FORWARD
Hepatitis	**0.265**	0.620	0.679	**0.111**
Genes	**0.062**	0.176	0.173	**0.079**
Mutagenesis	0.650	**0.280**	0.764	**0.134**
World	**0.640**	0.733	0.283	**0.149**
Mondial	1.550	**1.090**	1.710	**0.385**

compared to the all-at-once embedding extension. Otherwise, the results are only marginally worse or even slightly better when adding tuples one-by-one without recomputing paths for old data.

Overall, we observe that both FORWARD and NODE2VEC empirically achieve the goal of producing viable embeddings for new data, while keeping the old embeddings intact. In most of the datasets, the accuracy is almost as good as the static case, even when up to 50% of the data is new.

F. Execution Times

Here, we compare FORWARD's execution times to that of NODE2VEC in both the static and dynamic phases. Table V reports the runtime for the static classification. We observe that NODE2VEC is faster than FORWARD on all of the datasets.

Finally, we are also interested in the average time for generating an embedding for a newly arrived tuple, looking at both versions of the dynamic experiment. The results are shown in Table VI. The reported numbers are the average time (in seconds) that it takes to embed a newly-arrived tuple - we already have the embedding for the old tuples, and we measure only how much time it takes to train and infer a new embedding for the new tuple. Note that these numbers change dramatically from one dataset to another, as according to the experiment each new tuple in the classification relation is also accompanied by new tuples from other relations, thus the structure of the database affects the execution time. In the all-at-once setting, neither method is consistently faster than the other. In the 'one-by-one setting, where new tuples arrive iteratively, FORWARD is significantly faster than NODE2VEC across all datasets. Note that this is expected, because for NODE2VEC we have to re-train the embedding using gradient descent for every new tuple that arrives, where as for FOR-WARD we only have to solve a system of linear equations. This insight was essential in the design of FORWARD.

G. Conclusions

We offer some final conclusions regarding the experiments.

First, the tested methods all provide good tuple embeddings in the static setting, across all datasets, competing with the state-of-the-art methods.

Second, in the dynamic phase both NODE2VEC and FOR-WARD give good embeddings even when more than half of the data is new.

Finally, which method performs best strongly depends on the database. Both NODE2VEC and FORWARD can significantly outperform the other method respectively.

VII. CONCLUDING REMARKS

We studied the stable database embedding problem, that is, the problem of computing vector embeddings of dynamically changing databases where the embedding of tuples remains stable once they have been added to the database. Past techniques for the static database embedding problem are not designed for dynamic database embedding. We propose two algorithmic solutions to this problem, an adaptation of the node embedding algorithm NODE2VEC for graphs and a new algorithm forward FORWARD. Both of our embedding methods incorporate the structure of the entire databases with relations linked by foreign key constraints, while previous embedding methods focus on embedding single relations. Our experiments with a variety of downstream tasks demonstrate that both methods perform surprisingly well even as large parts of the database have changes in the dynamic process, with a slight edge for the FORWARD method in the dynamic setting.

In future work, it will be important to study the performance of our embedding algorithm with respect to downstream machine-learning tasks other than column prediction: record linking [14], [35], entity resolution [11], [24], data imputation [25], [49], data cleaning [1], [41], [49] and so on.

We argued in the Introduction that tuple deletion is a trivial operation in our dynamic setup where we preserve the embedding of existing tuples. This is why we focused on tuple insertions. However, there is a subtle issue about deletions that poses interesting questions somewhat orthogonal to what we study here. When deleting a tuple t^- and the corresponding point from the embedding, we do not delete all information about t^-, since the existence of the tuple had impact on how the other tuples were embedded. This has consequences for privacy considerations. For instance, if a user wants to be deleted from a database, then all information about the user must be deleted; this may even be a legal requirement. For this, the embedding of the remaining tuples needs to be adapted. Rather than re-computing the whole embedding, we may try to find a minimal set of changes that removes all information about the deleted tuple, but keeps the overall embedding intact.

REFERENCES

[1] ABEDJAN, Z., CHU, X., DENG, D., FERNANDEZ, R. C., ILYAS, I. F., OUZZANI, M., PAPOTTI, P., STONEBRAKER, M., AND TANG, N. Detecting data errors: Where are we and what needs to be done? *Proc. VLDB Endow. 9*, 12 (2016), 993–1004.

[2] AFIF, M. H., HEDAR, A.-R., HAMID, T. H. A., AND MAHDY, Y. B. Ss-svm (3svm): a new classification method for hepatitis disease diagnosis. *Int. J. Adv. Comput. Sci. Appl 4* (2013).

[3] ANTONIAK, M., AND MIMNO, D. Evaluating the stability of embedding-based word similarities. *Transactions of the Association for Computational Linguistics 6* (2018), 107–119.

[4] ATRAMENTOV, A., LEIVA, H., AND HONAVAR, V. A multi-relational decision tree learning algorithm – implementation and experiments. In *Inductive Logic Programming* (Berlin, Heidelberg, 2003), T. Horváth and A. Yamamoto, Eds., Springer Berlin Heidelberg, pp. 38–56.

[5] BÉRES, F., KELEN, D. M., PÁLOVICS, R., AND BENCZÚR, A. A. Node embeddings in dynamic graphs. *Applied Network Science 4*, 1 (2019), 1–25.

[6] BINA, B., SCHULTE, O., CRAWFORD, B., QIAN, Z., AND XIONG, Y. Simple decision forests for multi-relational classification. *Decis. Support Syst. 54*, 3 (2013), 1269–1279.

[7] BORDAWEKAR, R., BANDYOPADHYAY, B., AND SHMUELI, O. Cognitive database: A step towards endowing relational databases with artificial intelligence capabilities. *CoRR abs/1712.07199* (2017).

[8] BORDAWEKAR, R., AND SHMUELI, O. Using word embedding to enable semantic queries in relational databases. In *Proceedings of the 1st Workshop on Data Management for End-to-End Machine Learning* (New York, NY, USA, 2017), DEEM'17, ACM.

[9] BORDAWEKAR, R., AND SHMUELI, O. Exploiting latent information in relational databases via word embedding and application to degrees of disclosure. In *CIDR* (2019), www.cidrdb.org.

[10] BORDES, A., USUNIER, N., GARCIA-DURAN, A., WESTON, J., AND YAKHNENKO, O. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems* (2013), pp. 2787–2795.

[11] CAPPUZZO, R., PAPOTTI, P., AND THIRUMURUGANATHAN, S. Creating embeddings of heterogeneous relational datasets for data integration tasks. In *SIGMOD* (2020), D. Maier, R. Pottinger, A. Doan, W. Tan, A. Alawini, and H. Q. Ngo, Eds., ACM, pp. 1335–1349.

[12] CHENG, J., HATZIS, C., HAYASHI, H., KROGEL, M.-A., MORISHITA, S., PAGE, D., AND SESE, J. Kdd cup 2001 report. *SIGKDD Explor. Newsl. 3*, 2 (Jan. 2002), 47–64.

[13] DEBNATH, A. K., LOPEZ DE COMPADRE, R. L., DEBNATH, G., SHUSTERMAN, A. J., AND HANSCH, C. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of Medicinal Chemistry 34*, 2 (1991), 786–797.

[14] EBRAHEEM, M., THIRUMURUGANATHAN, S., JOTY, S. R., OUZZANI, M., AND TANG, N. Distributed representations of tuples for entity resolution. *Proc. VLDB Endow. 11*, 11 (2018), 1454–1467.

[15] GROVER, A., AND LESKOVEC, J. Node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York, NY, USA, 2016), KDD '16, ACM, pp. 855–864.

[16] GÜNTHER, M. FREDDY: fast word embeddings in database systems. In *SIGMOD Conference* (2018), ACM, pp. 1817–1819.

[17] GÜNTHER, M., THIELE, M., NIKULSKI, E., AND LEHNER, W. Retro-Live: Analysis of relational retrofitted word embeddings. In *EDBT* (2020), OpenProceedings.org, pp. 607–610.

[18] HAGBERG, A. A., SCHULT, D. A., AND SWART, P. J. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference* (2008), G. Varoquaux, T. Vaught, and J. Millman, Eds., pp. 11 – 15.

[19] HAMILTON, W., YING, Z., AND LESKOVEC, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems* (2017), I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30, Curran Associates, Inc.

[20] HARRIS, C. R., MILLMAN, K. J., VAN DER WALT, S. J., GOMMERS, R., VIRTANEN, P., COURNAPEAU, D., WIESER, E., TAYLOR, J., BERG, S., SMITH, N. J., KERN, R., PICUS, M., HOYER, S., VAN KERKWIJK, M. H., BRETT, M., HALDANE, A., DEL R'IO, J. F., WIEBE, M., PETERSON, P., G'ERARD-MARCHANT, P., SHEPPARD,

K., REDDY, T., WECKESSER, W., ABBASI, H., GOHLKE, C., AND OLIPHANT, T. E. Array programming with NumPy. *Nature 585*, 7825 (Sept. 2020), 357–362.

[21] HELLRICH, J., AND HAHN, U. Bad company—neighborhoods in neural embedding spaces considered harmful. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers* (2016), pp. 2785–2796.

[22] HOFMANN, T., SCHÖLKOPF, B., AND SMOLA, A. J. Kernel methods in machine learning. *The annals of statistics 36*, 3 (2008), 1171–1220.

[23] JAEGER, S., FULLE, S., AND TURK, S. Mol2vec: Unsupervised machine learning approach with chemical intuition. *J. Chem. Inf. Model. 58*, 1 (2018), 27–35.

[24] KUSHAGRA, S., BEN-DAVID, S., AND ILYAS, I. F. Semi-supervised clustering for de-duplication. In *AISTATS* (2019), vol. 89 of *Proceedings of Machine Learning Research*, PMLR, pp. 1659–1667.

[25] LAKSHMINARAYAN, K., HARP, S. A., AND SAMAD, T. Imputation of missing data in industrial databases. *Appl. Intell. 11*, 3 (1999), 259–275.

[26] LAU, J. H., AND BALDWIN, T. An empirical evaluation of doc2vec with practical insights into document embedding generation. In *Rep4NLP@ACL* (2016), ACL, pp. 78–86.

[27] LE, Q. V., AND MIKOLOV, T. Distributed representations of sentences and documents. In *ICML* (2014), vol. 32 of *JMLR Workshop and Conference Proceedings*, JMLR.org, pp. 1188–1196.

[28] LESZCZYNSKI, M., MAY, A., ZHANG, J., WU, S., ABERGER, C., AND RE, C. Understanding the downstream instability of word embeddings. In *Proceedings of Machine Learning and Systems 2020* (2020), pp. 262–290.

[29] LI, Y., AND YANG, T. *Word Embedding for Understanding Natural Language: A Survey*. Springer International Publishing, Cham, 2018, pp. 83–104.

[30] LIU, Y., OTT, M., GOYAL, N., DU, J., JOSHI, M., CHEN, D., LEVY, O., LEWIS, M., ZETTLEMOYER, L., AND STOYANOV, V. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).

[31] LODHI, H., AND MUGGLETON, S. Is mutagenesis still challenging? In *ILP* (2005), pp. 35–40.

[32] MAHDAVI, S., KHOSHRAFTAR, S., AND AN, A. dynnode2vec: Scalable dynamic network embedding. In *2018 IEEE International Conference on Big Data (Big Data)* (2018), IEEE, pp. 3762–3765.

[33] MAY, W. Information extraction and integration: The mondial case study. Tech. rep., Universität Freiburg, Institut für Informatik, 1999.

[34] MIKOLOV, T., SUTSKEVER, I., CHEN, K., CORRADO, G., AND DEAN, J. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 27th Annual Conference on Neural Information Processing Systems* (2013), pp. 3111–3119.

[35] MUDGAL, S., LI, H., REKATSINAS, T., DOAN, A., PARK, Y., KRISHNAN, G., DEEP, R., ARCAUTE, E., AND RAGHAVENDRA, V. Deep learning for entity matching: A design space exploration. In *SIGMOD Conference* (2018), ACM, pp. 19–34.

[36] NEVILLE, J., JENSEN, D. D., FRIEDLAND, L., AND HAY, M. Learning relational probability trees. In *KDD* (2003), ACM, pp. 625–630.

[37] NICKEL, M., TRESP, V., AND KRIEGEL, H.-P. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on Machine Learning* (2011), pp. 809–816.

[38] PASZKE, A., GROSS, S., MASSA, F., LERER, A., BRADBURY, J., CHANAN, G., KILLEEN, T., LIN, Z., GIMELSHEIN, N., ANTIGA, L., DESMAISON, A., KOPF, A., YANG, E., DEVITO, Z., RAISON, M., TEJANI, A., CHILAMKURTHY, S., STEINER, B., FANG, L., BAI, J., AND CHINTALA, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035.

[39] PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R., DUBOURG, V., VANDERPLAS, J., PASSOS, A., COURNAPEAU, D., BRUCHER, M., PERROT, M., AND DUCHESNAY, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research 12* (2011), 2825–2830.

[40] PEROZZI, B., AL-RFOU, R., AND SKIENA, S. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2014), pp. 710–710.

[41] SA, C. D., ILYAS, I. F., KIMELFELD, B., RÉ, C., AND REKATSINAS, T. A formal framework for probabilistic unclean databases. In *ICDT* (2019), vol. 127 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 6:1–6:18.

[42] SCHULTE, O., BINA, B., CRAWFORD, B., BINGHAM, D., AND XIONG, Y. A hierarchy of independence assumptions for multi-relational bayes net classifiers. In *IEEE Symposium on Comxputational Intelligence and Data Mining, CIDM 2013, Singapore, 16-19 April, 2013* (2013), IEEE, pp. 150–159.

[43] SCHUMACHER, T., WOLF, H., RITZERT, M., LEMMERICH, F., BACHMANN, J., FRANTZEN, F., KLABUNDE, M., GROHE, M., AND STROHMAIER, M. The effects of randomness on the stability of node embeddings. In *Machine Learning and Principles and Practice of Knowledge Discovery in Databases. ECML PKDD 2021* (2021), M. K. et al., Ed., vol. 1524 of *Communications in Computer and Information Science*, Springer.

[44] SHAH, V., KUMAR, A., AND ZHU, X. Are key-foreign key joins safe to avoid when learning high-capacity classifiers? *arXiv preprint arXiv:1704.00485* (2017).

[45] SHALEV-SHWARTZ, S., AND BEN-DAVID, S. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press, 2014.

[46] SUN, Z., DENG, Z.-H., NIE, J.-Y., AND TANG, J. Rotate: Knowledge graph embedding by relational rotation in complex space. *arXiv preprint arXiv:1902.10197* (2019).

[47] TASKAR, B., ABBEEL, P., AND KOLLER, D. Discriminative probabilistic models for relational data. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence* (San Francisco, CA, USA, 2002), UAI'02, Morgan Kaufmann Publishers Inc., pp. 485–492.

[48] WANG, C., RAO, W., GUO, W., WANG, P., LIU, J., AND GUAN, X. Towards understanding the instability of network embedding. *IEEE Transactions on Knowledge and Data Engineering* (2020), 1–1.

[49] WU, R., ZHANG, A., ILYAS, I. F., AND REKATSINAS, T. Attention-based learning for missing data imputation in HoloClean. In *MLSys* (2020), mlsys.org.

[50] YIN, X., HAN, J., YANG, J., AND YU, P. S. Efficient classification across multiple database relations: a crossmine approach. *IEEE Transactions on Knowledge and Data Engineering 18*, 6 (2006), 770–783.