Time in a Box: Advancing Knowledge Graph Completion with Temporal Scopes

Ling Cai
lingcai@ucsb.edu
UC Santa Barbara

Krzysztof Janowicz
UC Santa Barbara

Bo Yan
UC Santa Barbara

Rui Zhu
UC Santa Barbara

Gengchen Mai
Stanford University

ABSTRACT

Almost all statements in knowledge bases have a temporal scope during which they are valid. Hence, knowledge base completion (KBC) on temporal knowledge bases (TKB), where each statement *may* be associated with a temporal scope, has attracted growing attention. Prior works assume that each statement in a TKB *must* be associated with a temporal scope. This ignores the fact that the scoping information is commonly missing in a KB. Thus prior work is typically incapable of handling generic use cases where a TKB is composed of temporal statements with/without a known temporal scope. In order to address this issue, we establish a new knowledge base embedding framework, called TIME2BOX, that can deal with atemporal and temporal statements of different types simultaneously. Our main insight is that answers to a temporal query always belong to a subset of answers to a time-agnostic counterpart. Put differently, time is a filter that helps pick out answers to be correct during certain periods. We introduce boxes to represent a set of answer entities to a time-agnostic query. The filtering functionality of time is modeled by intersections over these boxes. In addition, we generalize current evaluation protocols on time interval prediction. We describe experiments on two datasets and show that the proposed method outperforms state-of-the-art (SOTA) methods on both link prediction and time prediction.

CCS CONCEPTS

• Computing methodologies → Knowledge representation and reasoning; Temporal reasoning.

KEYWORDS

Temporal Knowledge Base, Link Prediction, Time Prediction

ACM Reference Format:

Ling Cai, Krzysztof Janowicz, Bo Yan, Rui Zhu, and Gengchen Mai. 2021. Time in a Box: Advancing Knowledge Graph Completion with Temporal Scopes. In *Proceedings of the 11th Knowledge Capture Conference (K-CAP '21)*, December 2–3, 2021, Virtual Event, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3460210.3493566>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

K-CAP '21, December 2–3, 2021, Virtual Event, USA.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8457-5/21/12...\$15.00

<https://doi.org/10.1145/3460210.3493566>

1 INTRODUCTION

A knowledge base (KB) such as Wikidata and DBpedia stores statements about the world around us. A KB is typically represented as a set of triples in the form of (s, r, o) – short for (*subject, relation, object*), encoding the association between entities and relations among them. A statement is often temporally scoped, which indicates during which time period it is valid. Two examples are (*Albert Einstein, educatedAt, ETH Zurich, 1896 - 1900*) and (*Albert Einstein, academicDegree, Doctor of Philosophy in Physics, 1906*). The former specifies the time period during which Albert Einstein studied at ETH, and the latter points out the specific date when he obtained his degree. Graphs that contain a substantial amount of such time-aware statements are often called temporal knowledge base (TKB) in the machine learning literature. Each statement in a TKB is associated with a validity time as (s, r, o, t^*) ¹.

Due to the ever-changing state of the world and missing data, TKBs usually contain inaccurate and incomplete information similar to KBs. The sparsity of TKBs necessitates temporal knowledge base completion (TKBC), namely inferring missing statements from known statements. Temporal link prediction task is proposed to evaluate a TKBC model by testing its performance on answering incomplete temporal queries of the form $(s, r, \exists, t^*)$ or $(\exists, r, o, t^*)$.

Despite recent success stories on time-agnostic KBC, research on TKBC is still in its early age and is facing new challenges. The validity time period of a statement is often missing in a KB. As a result, it is difficult to distinguish whether statements in a KB are atemporal (e.g., (*Albert Einstein, instanceOf, Human*)) or time-dependent (e.g., (*United States of America, instanceOf, Historical Unrecognized State*²)). This leads to the question of which statements should be part of a TKB in the first place. Prior works restrict TKBs to a collection of statements where the validity time period for each statement *must* be available. However, in WIKIDATA114k, a dataset from Wikidata, for instance, 85.1% of all statements are temporal, 56.2% of the temporal statements are missing their validity time information and are excluded in previous studies while only 247,393 out of 1,660,824 statements (i.e., 14.9%) are truly atemporal³. As the number of temporal statements with missing validity information is substantial, excluding them from a TKB will significantly reduce the amount of information that could be useful in TKB studies.

¹ t^* could be a time instant or time interval

²According to Wikidata that statement holds true during 1776–1784.

³For all the statements, we first categorize predicates into two groups – atemporal predicates and temporal predicates. If a predicate has ever been involved in a statement that has temporal scoping, it belongs to temporal predicates; otherwise, it is an atemporal predicate. Atemporal statements are those associated with atemporal predicates.

Retaining these temporally scoped statements leads to several challenges that need to be addressed. For instance, how to design a TKBC model to handle statements with and without known temporal scoping from the data representation perspective and model design perspective? Clearly, the conventional representation in prior TKBC in the form of $(s, r, o, t)^4$ falls short. An ideal TKBC model should be more *flexible* to address cases when the validity information of different types (i.e., point in time, right-open interval (known start time), left-open interval (known end time), closed interval) is presented in a TKB or even no validity information is available for a statement.

The second challenge is how to predict the temporal scope of a statement as it is often missing in TKBs. This task is referred to as time interval prediction, which amounts to answering incomplete queries of the form $(s, r, o, ?I)$. How to generate a predicted time interval and evaluate it require further investigation. This problem has only been addressed very recently by Jain et al. [5]. However, at times their evaluation protocols fail to distinguish one predicted interval from another since they do not consider the gap between the predicted and the gold interval in case of no overlap. For instance, the same metric scores are assigned to two predictions [1998, 1999] and [1998, 2010] when a gold interval [2011, 2020] is considered.

In this paper, we present a novel TKBC embedding framework, called TIME2BOX, which relies on the intuition that the answer set in a temporal query $(s, r, ?o, t^*)$ is always a subset of answers of its time-agnostic counterpart $(s, r, ?o)$. As illustrated in Fig. 1, there are four correct answer entities to a query (*Albert Einstein, employer, ?o*). However, when time information is specified (e.g., *Albert Einstein, employer, ?o, 1933*) as shown in Fig. 1b, the number of positive answers becomes three. With more temporal information being available (e.g., *Albert Einstein, employer, ?o, [1933, 1955]*), the answer set shrinks further (see Fig. 1c). Therefore, we propose to model a statement in a TKB by imitating the process of answering its corresponding temporal query $(s, r, ?o, t^*)$, which can be achieved in two steps – finding answer entities to its atemporal counterpart $(s, r, ?o)$ by using KBC methods and then picking out entities to be true to the temporal query from preceding answers by including time. We implement this idea by using box embeddings, especially inspired by QUERY2BOX [8], which is originally used for answering conjunctive queries. Boxes, as containers, can naturally model a set of answers they enclose. The filtering functionality of time can be naturally modeled as intersections over boxes similarly to Venn diagrams. Meanwhile, performing an intersection operation over boxes would still result in boxes, thus making it possible to design a unified framework to deal with statements of different types.

Our main research contributions are listed as follows:

- We propose a box-based KG embedding TKBC framework (TIME2BOX) that can represent and model statements with different types of validity information (i.e., unknown temporal scoping, left/right-open intervals and closed intervals).
- We introduce a new evaluation metrics *gaeIOU* for interval evaluation by taking the gap between a gold interval and a predicted interval into consideration if no overlap exists.

- Extensive experiments on two datasets - WIKIDATA12k and WIKIDATA114k - show that TIME2BOX yields the state-of-the-art (SOTA) results in link prediction and outperforms SOTAs in time prediction by significant margins. TIME2BOX code is available at Github⁵.

2 RELATED WORK

Knowledge Base Completion. The core insight of KBC is to embed entities and relations in a KB into low-dimensional vectors, which can be utilized in downstream tasks, such as link prediction. These methods can be roughly classified into two groups: transformation-based models [2, 9] and semantic matching energy based models [11, 12]. All KBC models ignore the temporal scoping of statements, and thus are unable to address temporal statements. However, these models are the foundations for TKBC.

Temporal Knowledge Base Completion. There are two lines of works on temporal link prediction. The first assumes that knowledge in KBs evolves over time and historical statements/events drive the occurrence of new events [6, 10]. The other line is to fill in missing components in TKGs with/without explicitly modeling temporal dependencies between statements. We are on the second line. However, our idea is to deal with cases when time could be an instant, a (left/right-open) interval, closed interval or even missing, while prior works can only handle timestamped statements.

3 PRELIMINARIES

3.1 Temporal Knowledge Bases

Prior TKBC methods typically work on TKBs in which each statement has to be associated with validity information. Thereby, for statements that do not have known temporal scopes, they either exclude them from a TKB in the beginning or assume that these statements hold all the time [7]. However, there are limitations in both ways. As discussed in Section 1, excluding them from a TKB will significantly reduce the amount of information that could be beneficial in TKBC studies as the number of such statements is substantial. For the latter, their assumption would be problematic since a lot of them may only hold for a certain time period. For instance, the statement (*Warsaw, country, Russian Empire*) holds during the time interval [1815-07-09, 1916-11-04]. Following the open-world assumption (OWA), we argue that TKBs are an extension to KBs insofar as the lack of temporal scoping for any given statement does not imply it holding indefinitely.

In the following, we use t and $I_{st}^{et} = [st, et]$ to denote a time point and a time interval, respectively. The symbol $-$ will stand for unknown temporal validity. There are five types of statements in such a TKB: (1) (s, r, o) for a statement without a known temporal scope; (2) (s, r, o, t) for a timestamped statement which holds at a point in time t ; (3) (s, r, o, I_{st}^-) for a right-open interval-based statement, in which only the time when the statement starts to hold is known; (4) (s, r, o, I_{st}^+) for a left-open interval-based statement, in which only the time when the statement ceases to hold is known; and (5) (s, r, o, I_{st}^{et}) for a statement which is temporally scoped by a closed interval I_{st}^{et} . Then a TKB is denoted as $\mathcal{G} = \bigcup_{(s,r,o,t^*)}$, namely the union of statements of the five types, where $s, o \in E$ represent

⁴ t denotes a time point

⁵<https://github.com/ling-cai/Time2Box>

Figure 1: Illustration of TIME2BOX reasoning process. In each figure, the upper part shows entities and relations in the KB space and the latter illustrates their correspondences in the embedding space. In all figures, the final boxes are shaded regions in orange and answer entities are in the boxes. Note that we omit the edges between *Albert Einstein* and associated entities in Fig.1b and 1c for simplicity. Fig.1a shows that for an atemporal query, the reasoning process picks out all possible answers from the whole entity space and encloses them into a time-agnostic box. In Fig.1b a time-aware box is added to enclose entities that are relevant to *Albert Einstein* in 1933. Then the intersection between time-agnostic and time-aware boxes consists of a new box, which contains entities that satisfy both requirements. When more validity information is available, Fig.1c shows that more time-aware boxes can be added and the intersection box that contains correct answers would shrink further. By doing so, TIME2BOX is flexible enough to handle different types of queries.

entities, $r \in R$ denotes a relation and $t^* \in \{t, I_{st}^-, I_{st}^+, I_{st}^t, None\}$ denotes different types of valid time or no valid time available.

3.2 The TKBC Problem

Link prediction and time prediction are two main tasks used to evaluate a TKBC model. Statements in TKBs are split into training, validation, and test sets, used for model training, parameter tuning and model evaluation, respectively.

Link prediction. Queries used in this task are of the form $(s, r, ?o, t^*)$. Performance is evaluated on the rank of a given golden, i.e., ground truth, answer in the list of all the entities sorted by scores in a descending order. Then MRR (mean reciprocal rank), MR (mean rank), HITS@1, HITS@3 and HITS@10 are computed from the ranks over all queries in the test set. However a query may be satisfied by multiple answer entities. Thus another correct answer may be ranked over the given golden answer. In such cases, a KBC/TKBC model should not be penalized. A traditional strategy used in KBC is to filter out those correct answers that are already in the training and validation sets before calculating metrics. This strategy can be directly applied to queries of the form $(s, r, ?o)$ or $(s, r, ?o, t)$. However, it may not be sufficient for queries of the form $(s, r, ?o, I)$, as there may exist other answers that are true during a time period within the interval I . For example, suppose two statements – $(Albert Einstein, employer, Princeton University,$

$[1933, 1955])$ and $(Albert Einstein, employer, Leiden University, [1920, 1946])$, both Princeton University and Leiden University are correct answers during the period $[1933, 1946]$. One naive way to solve this problem is to discretize the interval I to a sequence of time points ts and then to convert $(s, r, ?o, I)$ into timestamped queries of the form $(s, r, ?o, t)$ so that the same filtering process can be performed on each timestamped query. Finally, the ranks over them are averaged to be the rank for a time interval-based query. This idea is well-aligned with the proposal by [5].

Time prediction. Time prediction queries in TKBs are of the form $(s, r, o, ?I)$. Despite the fact that the validity information could be a point in time or a time interval, a point in time can be viewed as a special time interval, in which start time and end time coincide. Thus, time prediction boils down to time interval prediction. Its performance is evaluated by the overlap between a gold interval and a predicted interval or the closeness between those in case of no overlap. We describe the existing evaluation protocols and propose a generalized evaluation metric in Section 5.

4 METHOD

The key insight of TIME2BOX lies in an intuition that the answer set of a temporal query $(s, r, ?o, t^*)$ is always a subset of answers of its time-agnostic counterpart $(s, r, ?o)$ and set size decreases by adding more temporal constraints. As illustrated in Fig. 1a, four

object entities satisfy the atemporal query (*Albert Einstein, employer, ?o*) while three entities are the correct answers when the query is restricted to the year of 1993 (see Fig. 1b) and only one entity is correct when another time information is further added in the statement, shown in Fig. 1c. Inspired by this observation, we propose to model a temporal statement (s, r, o, t^*) by imitating the process of answering its corresponding temporal query $(s, r, ?o, t^*)$, which can be achieved through two steps: 1) finding a set of answer entities that are true for the corresponding atemporal query by using any KBC model and 2) imposing a filtering operation enforced by time to restrict answers afterwards. In following sections, we take a time instant-based statement as example to formalize our idea in a KB space and a vector space, respectively.

4.1 Formalization in a KB Space

For a statement (s, r, o, t) in a TKB, the first step of TIME2BOX, as shown in Fig. 1a, is to project the subject s to a set of object entities that are true to its corresponding atemporal query in the form of $(s, r, ?o)$ enforced by the relation r . This is a prerequisite for any statement and can theoretically be addressed by any KBC method. Formally, the relation projector is defined as:

Relation Projector – OP_r : Given the subject entity s and the relation r , this operator obtains: $S_r = \{o' \mid (s, r, o') \in \mathcal{G}^N\}$. $\mathcal{G}^N$ is the time-agnostic counterpart of $\mathcal{G}$.

Then time information is used to filter out entities that are incorrect during the time of interest from the answer set S_r . This can be achieved by first *projecting* the subject s to a set of object entities that co-occur with s in statements at a given time point (as shown in blue edges in Fig. 1b) and then finding the *intersection* over them and S_r (see the three entities in red in Fig. 1b). Accordingly, the two involved steps are defined as:

Time Projector – OP_t : Given the subject s and the timestamp t , this operator obtains: $S_t = S_t = \{o' \mid o' \in E \text{ and } (s, r', o', t) \in \mathcal{G} \text{ and } r' \in R\}$.

Intersection Operator – OI : Given S_r and S_t , this operator obtains the intersection $S_{inter} = \{o \mid o \in (S_r \text{ and } S_t)\}$.

In fact, such a modeling process also fits to left/right-open interval-based statements directly. For a left/right-open interval-based statement, we only consider the known endpoint time in such an interval as we follow the open-world assumption. However, for an atemporal statement, we only need one relation projector to obtain S_r , which is the final set consisting of correct answer entities to its query form. For a closed interval-based query, one commonly used approach is to randomly pick one timestamp within the interval and to associate it with (s, r, o) . Then it can be modeled the same way as an instant-based statement. At training, a timestamp is randomly picked from the interval to ensure that all the timestamps in the interval are used. In addition to the common strategy, TIME2BOX allows sampling of a sub-time interval within the given interval so that two constraints (i.e., start time and end time) can be imposed by using two time projectors, as shown in Fig. 1c⁶.

⁶Alternatively, one could also enumerate all the timestamps within the interval and use different OP_t to project the subject to multiple sets of entities, each of which is specific for one timestamp. Subsequently, an intersection operator again is performed over all the sets of entities obtained from OP_r and OP_t in the previous step. However, in spite of its efficiency, this practice is hard to implement in mini-batch training manner since time intervals in different statements usually have varying duration and thus contain different number of timestamps.

4.2 Implementation in a Vector Space

In order to implement this idea in a vector space, two key points are 1) how to model a set of answers returned by a KBC model and 2) how to instantiate two projectors and one intersection operator.

Prior KBC models are incapable of directly representing a set of answer entities in a vector space. Instead, they usually represent entities and relations as single points in the vector space and model point-to-point projections, e.g., TransE. Inspired by QUERY2BOX [8], which is used to deal with complex queries that involve conjunctions, existential quantifiers, and disjunctions, we introduce the idea of boxes in the vector space and thus name the proposed framework TIME2BOX. The reasons for adopting boxes are three-fold. First, boxes are containers that can naturally model a set of answer entities they enclose. Second, finding the intersection set among sets of entities amounts to finding the intersected area over boxes similar to the concept of Venn diagram. Third, the result of performing an intersection operation over boxes is still a box, which makes it possible to deal with statements of different types in a unified framework.

In TIME2BOX, each entity $e \in E$, relation $r \in R$, and timestamp $t \in T$ (T is the set of all discrete timestamps in a TKB) are initialized as vector embeddings $\mathbf{e} \in \mathbb{R}^d$, $\mathbf{r} \in \mathbb{R}^d$, and $\mathbf{t} \in \mathbb{R}^d$. S_r , S_t , and S_{inter} refer to sets of entities and thus are modeled by boxes, represented as box embeddings in the vector space. In the following, we first introduce the definition of box embeddings and then introduce main components of modeling and reasoning.

Box embeddings: Mathematically, they are axis-aligned hyper-rectangle in a vector space, which can be determined by the position of the box (i.e., a center point) and its length (i.e., offsets). Formally, in a vector space $\mathbb{R}^d$, a box can be represented by $\mathbf{b} = (\text{Cen}(\mathbf{b}), \text{Off}(\mathbf{b}))$, where $\text{Cen}(\mathbf{b}) \in \mathbb{R}^d$ is its center point and $\text{Off}(\mathbf{b}) \in \mathbb{R}_{\geq 0}^d$ specifies the length/2 of the box in each dimension. If an entity belongs to a set, its entity embedding is modeled as a point inside the box of the set. The interior of a box in the vector space can be specified by points inside it:

$$\text{box}_{\mathbf{b}} = \{\mathbf{e} \in \mathbb{R}^d \mid \text{Cen}(\mathbf{b}) - \text{Off}(\mathbf{b}) \leq \mathbf{e} \leq \text{Cen}(\mathbf{b}) + \text{Off}(\mathbf{b})\} \quad (1)$$

where $\leq$ denotes element-wise inequality.

Projection operators in a vector space. In previous work, relations are commonly assumed to be projectors that transform a subject embedding to an object embedding in terms of *points* in a vector space, e.g., TransE [2] and RotatE [9]. Here we adopt a similar idea but take both relations and timestamps as projectors (OP_r and OP_t) to project a subject to a set of entities in S_r – represented as a time-agnostic box $\mathbf{b}_{S_r}$ and to a set of entities in S_t – represented as a time-aware box $\mathbf{b}_{S_t}$, respectively, which are illustrated in Fig. 1.

The center of a box can be defined as the resulting embedding after applying a projection operator (OP_r or OP_t) on the subject embedding. The centers of $\mathbf{b}_{S_r}$ and $\mathbf{b}_{S_t}$ can be formulated as below:

$$\text{Cen}(\mathbf{b}_{S_r}) = \mathbf{e} \odot \mathbf{r}; \text{Cen}(\mathbf{b}_{S_t}) = \mathbf{e} \otimes \mathbf{t} \quad (2)$$

where $\odot \mathbf{r}$ and $\otimes \mathbf{t}$ are projectors OP_r and OP_t , respectively. Theoretically, projection operators could be instantiated by any projector in existing KBC models, such as element-wise addition in TransE, element-wise product in DistMult, and Hadamard product in RotatE. Even though OP_r and OP_t can be different, we choose the

same projector for both and implement two TIME2BOX models by taking element-wise addition and element-wise product as operators by following TransE and DistMult, respectively. Accordingly, these two models are named as TIME2BOX-TE and TIME2BOX-DM.

Ideally, the size of the box $\mathbf{b}_{S_r}$ should be determined by both the subject entity and the relation, since the box contains all object entities that satisfy a query in the form of $(s, r, ?o)$. The same applies to $\mathbf{b}_{S_t}$. However, as the entity space is usually large in a KB, introducing entity-specific parameters would result in high computational cost. Therefore, in practice, $\text{Off}(\mathbf{b}_{S_r})$ and $\text{Off}(\mathbf{b}_{S_t})$ are only determined by the relation $r \in R$ and the timestamp $t \in T$, respectively. Put differently, the size of $\mathbf{b}_{S_r}$ and $\mathbf{b}_{S_t}$ are initialized based on r and t , which are learned through training.

Intersection Operators in a vector space. An intersection operator aims to find the intersection box $\mathbf{b}_{inter} = (\text{Cen}(\mathbf{b}_{inter}), \text{Off}(\mathbf{b}_{inter}))$ of a set of box embeddings $\mathbf{B} = \{\mathbf{b}_{S_r}, \mathbf{b}_{S_{t1}}, \dots, \mathbf{b}_{S_{tn}}\}$ obtained from the previous step. The intersection operator should be able to deal with $\mathbf{B}$ of different sizes, as required in Fig. 1. Thus, both $\text{Cen}(\mathbf{b}_{inter})$ and $\text{Off}(\mathbf{b}_{inter})$ are implemented by using attention mechanisms. Following the idea in [1], the center point $\text{Cen}(\mathbf{b}_{inter})$ is calculated by performing element-wise attention over the centers of boxes in $\mathbf{B}$. This can be formulated as follows:

$$\text{Cen}(\mathbf{b}_{inter}) = \sum_i \text{softmax}(\text{NN}(\text{Cen}(\mathbf{b}_i)) \odot \text{Cen}(\mathbf{b}_i)) \quad (3)$$

where NN is a one-layer neural network and $\mathbf{b}_i \in \mathbf{B}$.

Since the intersection box $\mathbf{b}_{inter}$ must be smaller than any of the box in $\mathbf{B}$, we use element-wise min-pooling to make sure the new box must be shrunk and perform DeepSets over all the $\text{Off}(\mathbf{b}_i)$ ($\mathbf{b}_i \in \mathbf{B}$) to downscale $\mathbf{b}_{inter}$ [8]. This can be written as below:

$$\text{Off}(\mathbf{b}_{inter}) = \text{Min}(\mathbf{Off}) \odot \sigma(\text{DeepSets}(\mathbf{Off})) \quad (4)$$

where $\text{DeepSets}(\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}) = \text{MLP}(1/n) \cdot \sum_i^n \text{MLP}(\mathbf{x}_i)$, σ denotes the sigmoid function, and $\mathbf{Off} = \{\text{Off}(\mathbf{b}_i) : \mathbf{b}_i \in \mathbf{B}\}$.

4.3 Optimization Objective

For a query, TIME2BOX aims to pull correct entity embedding into the final box $\mathbf{b}_{inter}$ while pushing incorrect entity embedding far away from it. The distance-based loss proposed by Sun et al. [9] satisfies this need:

$$\text{Loss} = -\log \sigma(\gamma - D(\mathbf{o}, \mathbf{b}_{inter})) - \frac{1}{k} \sum_{i=1}^k \log \sigma(D(\mathbf{o}', \mathbf{b}_{inter}) - \gamma) \quad (5)$$

where σ is the sigmoid function, γ is a fixed margin, $\mathbf{o}$ is the embedding of a positive entity to the given query, and k is the number of negative samples $\mathbf{o}'$. $D(\mathbf{o}, \mathbf{b}_{inter})$ measures the distance between entity $\mathbf{o}$ and the final box $\mathbf{b}_{inter}$. With the size of a box being considered, the distance is divided into two parts: outside distance $D_{outside}(\mathbf{o}, \mathbf{b}_{inter})$ and inside distance $D_{inside}(\mathbf{o}, \mathbf{b}_{inter})$. For cases when $\mathbf{o}$ is outside of $\mathbf{b}_{inter}$, the former refers to the distance of an entity embedding $\mathbf{o}$ to the boundary of the box $\mathbf{b}_{inter}$, and the latter calculates the distance between the box's center $\text{Cen}(\mathbf{b}_{inter})$ and its boundary. This can be formalized as below:

$$D(\mathbf{o}, \mathbf{b}_{inter}) = \alpha \cdot D_{inside}(\mathbf{o}, \mathbf{b}_{inter}) + D_{outside}(\mathbf{o}, \mathbf{b}_{inter}) \quad (6)$$

where $\alpha \in [0, 1]$. When $\alpha = 0$, it means that a positive entity is required to be in a $\mathbf{b}_{inter}$, but its distance to the center is not as important. $D_{inside}(\mathbf{o}, \mathbf{b}_{inter})$ and $D_{outside}(\mathbf{o}, \mathbf{b}_{inter})$ are written as:

$$D_{inside}(\mathbf{o}, \mathbf{b}_{inter}) = \|\text{Cen}(\mathbf{b}_{inter}) - \text{Min}(\mathbf{b}_{max}, \text{Max}(\mathbf{b}_{min}, \mathbf{o}))\|_1$$

$$D_{outside}(\mathbf{o}, \mathbf{b}_{inter}) = \|\text{Max}(\mathbf{o} - \mathbf{b}_{max}, \mathbf{0}) + \text{Max}(\mathbf{b}_{min} - \mathbf{o}, \mathbf{0})\|_1$$

where $\mathbf{b}_{min} = \text{Cen}(\mathbf{b}_{inter}) - \text{Off}(\mathbf{b}_{inter})$ and $\mathbf{b}_{max} = \text{Cen}(\mathbf{b}_{inter}) + \text{Off}(\mathbf{b}_{inter})$ are embeddings of the bottom left corner and the top right corner of $\mathbf{b}_{inter}$, respectively.

Compared to answering atemporal queries, finding correct answers to temporal ones is more challenging. Therefore, the loss function should reward more in the optimization direction that is capable of correctly answering temporal queries. For a given query q_i , we use $\frac{1}{n_{q_i}}$, where n_{q_i} is the number of correct answers to q_i that appear in training as a weight to adjust the loss. The core idea here is that time-aware queries often are satisfied with fewer answers, and, thus, are harder to answer compared to atemporal queries.

4.4 Time Negative Sampling

Entity negative sampling is widely used in KBC. For a positive sample (s, r, o) , negative samples are constructed by replacing o with other entities o' , ensuring that (s, r, o') must not appear in training set. In this paper, we adopt this strategy so that the model is able to learn the association between entities, relations, and time occurring in a positive sample by distinguishing the correct answers from the negative samples. Moreover, for time-aware statements, we perform time negative sampling, which corrupts a statement (s, r, o, t) by replacing t with a number of timestamps t' . This is important for statements where only start time or end time is available. As shown in Fig. 1, the proposed architecture cannot distinguish those statements from time instant-based statements. But time negative sampling can mitigate this issue to some degree. The following is used for time negative sampling concerning different types of statements (st and et are short for start and end time):

$$T' = \begin{cases} \{t' \in T : (s, r, o, t') \notin \mathcal{G}\} & (s, r, o, t) \\ \{t' \in T : (s, r, o, t') \notin \mathcal{G}, t' < st\} & (s, r, o, I_{st}^-) \\ \{t' \in T : (s, r, o, t') \notin \mathcal{G}, t' > et\} & (s, r, o, I_{et}^+) \\ \{t' \in T : (s, r, o, t') \notin \mathcal{G}, t' \notin T_{st}^{et}\} & (s, r, o, I_{st}^{et}) \end{cases} \quad (7)$$

where T_{st}^{et} denotes a set of time points within the interval I_{st}^{et} .

4.5 Time Smoothness Regularizer

Time is continuous. We may expect that neighboring timestamps would have similar representations in the vector space. Following Lacroix et al. [7], we penalize time difference between embeddings of two consecutive timestamps by using L_2 :

$$\Lambda(T) = \frac{1}{|T| - 1} \sum_{i=1}^{|T|-1} \|\mathbf{t}_{i+1} - \mathbf{t}_i\|_2^2 \quad (8)$$

During the training step, for batches with temporal statements, we add this regularizer with a weight scalar β to the loss function in Eq. 5, where β specifies the degree of penalization.

5 EVALUATION METRICS

Time Interval Evaluation. $gIOU$ and $aeIOU$ [5] are two evaluation metrics recently adopted in time interval prediction. Both are

Figure 2: Evaluation Comparison between aeIOU and gaeIOU on different predicted intervals. Suppose a gold interval is [2011, 2016], seven possible predicted intervals are represented as rectangles in black. Intersections between the predicted and the gold are in pink and gaps are in orange if no overlap exists. Notably, gaeIOU is able to distinguish these predictions while aeIOU fails to do so.

built on *Intersection Over Union* that is commonly used for bounding box evaluation in Computer Vision.

The idea of *gIOU* is to compare the intersection between a predicted interval and a gold interval against the maximal extent that the two intervals may expand. It can be formulated as below:

$$gIOU(I^{gold}, I^{pred}) = \frac{D(I^{gold} \cap I^{pred})}{D(I^{gold} \cup I^{pred})} \quad (9)$$

$$\frac{D(I^{gold} \cup I^{pred} \setminus (I^{gold} \cap I^{pred}))}{D(I^{gold} \cup I^{pred})} \in (-1, 1]$$

where $I^{gold} \cap I^{pred}$ is the overlapping part of two intervals, $I^{gold} \cup I^{pred}$ denotes the shortest contiguous interval (hull) that contains both I^{gold} and I^{pred} . As shown in Fig. 2, if $I^{gold} = [2011, 2016]$ and $I^{pred} = [2009, 2013]$, then $I^{gold} \cap I^{pred} = [2011, 2013]$ and $I^{gold} \cup I^{pred} = [2009, 2016]$. $D(I) = I_{max} - I_{min} + 1$ is the number of time points at a certain granularity (e.g., year in this paper) during the time interval I .

Compared to *gIOU*, affinity enhanced IOU, denoted as *aeIOU*, provides a better evaluation in case of non-overlapping intervals and outputs scores in $[0, 1]$. It can be written as follow:

$$aeIOU(I^{gold}, I^{pred}) = \begin{cases} \frac{D(I^{gold} \cap I^{pred})}{D(I^{gold} \cup I^{pred})} & D(I^{gold} \cap I^{pred}) > 0 \\ \frac{1}{D(I^{gold} \cup I^{pred})} & otherwise \end{cases} \quad (10)$$

We notice that *aeIOU* cannot distinguish some cases. As shown in Fig. 2, *aeIOU* results in the same scores for ⑤, ⑥, and ⑦ when compared to the gold interval [2011, 2016]. Intuitively one would assume that ⑦ is better than the others and ⑥ is the least desirable. The former has a one-year intersection between ⑦ and the gold. For the latter, the gap between ⑤ and the gold is smaller than that between ⑥ and the gold, despite the fact that neither ⑤ and ⑥ overlaps with the gold. Its failure lies in that it doesn't consider the gap between the gold and predicted interval in case of no overlap.

In the following, we take both the hull and the intersection/gap between a gold interval and a predicted interval into the design of the metric. The intuition is when the size of the hull remains the same, the score of a predicted interval *decreases* with a larger

gap to the gold in case of no overlap and *increases* with a larger intersection. *aeIOU* is therefore generalized to *gaeIOU* as below:

$$gaeIOU(I^{gold}, I^{pred}) = \begin{cases} \frac{D(I^{gold} \cap I^{pred})}{D(I^{gold} \cup I^{pred})} & D(I^{gold} \cap I^{pred}) > 0 \\ \frac{D'(I^{gold}, I^{pred}) - 1}{D(I^{gold} \cup I^{pred})} & otherwise \end{cases} \quad (11)$$

where $D'(I^{gold}, I^{pred}) = \max(I_{min}^{gold}, I_{min}^{pred}) - \min(I_{max}^{gold}, I_{max}^{pred}) + 1$ is the length of the gap.

Accordingly, the Property (P) that a good evaluation metric must satisfy can be rewritten as: if predicted intervals (partially) overlap with the gold interval with the same size, then the prediction having a smaller hull with the gold interval should be awarded more by M ; if there is no overlap, the prediction that has a smaller hull and a narrower gap with the gold should be scored higher by M . It follows that *gaeIOU* satisfies Property P, whereas *aeIOU* does not satisfy it; see Fig. 2. It can be formalized as below:

Property P: In case of $D(I^{gold} \cap I^{pred1}) = D(I^{gold} \cap I^{pred2}) \neq 0$, $M(I^{gold}, I^{pred1}) > M(I^{gold}, I^{pred2})$ if and only if $D(I^{gold} \cup I^{pred1}) < D(I^{gold} \cup I^{pred2})$.

In case of non-overlapping, $M(I^{gold}, I^{pred1}) > M(I^{gold}, I^{pred2})$ if and only if $D(I^{gold} \cup I^{pred1}) \cdot D'(I^{gold} \cap I^{pred1}) > D(I^{gold} \cup I^{pred2}) \cdot D'(I^{gold} \cap I^{pred2})$.

6 EXPERIMENT

Our goal here is to evaluate TIME2BOX in both link prediction and time prediction tasks. For a test sample $(s, r, ?o, t^*)$, we replace $?o$ with each entity $o' \in E$ and use $\log\sigma(\gamma - D(o', b_{inter}))$, a variation of the inverse distance used in Eq. 5, as scores for link prediction. Entities that have higher scores are more likely to form new links. Likewise, in terms of time prediction, for a query $(s, r, o, ?I)$, we first replace I with each timestamp $t \in T$ and calculate its score. Then we use the greedily coalescing method proposed in [5] to generate time intervals as predictions.

6.1 Datasets

We report experimental results using two TKBC datasets, which both are rooted in Wikidata. WIKIDATA12k is a widely used benchmark dataset in TKBC where each statement is associated with a time “interval” [3]. Such an “interval”, in fact, could be a time instant, where start time and end time are the same, a left/right-open interval, or a closed interval. Note that this dataset excludes statements that do not have known temporal scopes in Wikidata, although they may be time-dependent and useful in TKBC, as discussed in Section 1. The other dataset is a subset of WIKIDATA432k proposed by Lacroix et al. [7], which is the only TKB dataset where the start time, end time, or both of a statement can remain unspecified. Although this dataset is more appropriate for a TKBC problem, there are two limitations. First, it poses a computational burden as it contains 432k entities and 407 relations, consisting of 7M tuples in the training set. Second, there are several mistakes in the temporal information. For instance, 2014 was written as 2401. We extract a subgraph, named as WIKIDATA114k, and correct temporal information by checking it against Wikidata. More details about data pre-processing and statistics are in Appendix A (All

Appendices are available online⁷). Since our focus is on generic knowledge bases, we do not consider event-based datasets, such as ICEWS14 and ICEWS05-15, in which each statement is associated with a timestamp.

6.2 Baselines and Model Variants

In the following experiments, we regard TIME2BOX-TE as our main model, in which both the relation operator and the time projector are instantiated as an element-wise addition. It is denoted as TIME2BOX in resulting tables. We compare it against two SOTAs in TKBC: TNT-Complex and TIMEPLEX base model by using the implementation in Jain et al. [5], both of which are based on the time-agnostic KBC model: ComplEx [11].

In addition to comparison with existing SOTAs, we also conduct an ablation study, in which several variants of the proposed model are compared: (1) TIME2BOX-SI, short for Sample Interval: for a closed interval-based statement, this variant randomly samples a sub time interval from a given interval at each training step and train it as shown in Fig. 1c. (2) TIME2BOX-TR: previous works in TKBC often explicitly fused relations with time information to obtain time-aware relations and empirically demonstrated its effectiveness [4, 5, 7]. We also explicitly model the association between relations and time as a new point $p_{rt} = r + t$ in the vector space and incorporate it into Eq. 3 to help locate the intersection box. (3) TIME2BOX-DM: this variant implements the relation and time projectors as an element-wise product in real space as DistMult does. (4) TIME2BOX-TNS: this variant tests the effect of time negative samples, in which we replace a number of entity negative samples with time negative samples, as introduced in Section 4.4.

All models are trained on statements in training set and evaluated by answering queries where either the object or the time information is missing. Hyper-parameter settings are introduced in Appendix B and a comparison of parameters used in different models in Appendix F. There are several limitations in current experimental setups of SOTAs and we detail them in Appendix C.

6.3 Main Results

Datasets	WIKIDATA12k				WIKIDATA114k			
Metrics	MRR	MR	HITS@1	HITS@10	MRR	MR	HITS@1	HITS@10
TNT-Complex	31.77	415	19.24	51.74	49.25	638	41.02	66.99
TIMEPLEX base	34.55	302	21.91	53.25	49.99	337	41.25	66.10
TIME2BOX-TR	34.99	102	24.79	56.32	50.25	85	41.73	67.13
TIME2BOX-DM	35.90	139	25.52	56.74	48.84	284	41.09	64.33
TIME2BOX-SI	36.79	100	27.16	56.43	50.42	139	41.65	67.58
TIME2BOX-TNS	37.25	100	27.41	57.31	50.55	185	41.77	67.78
TIME2BOX	37.30	101	27.38	57.36	50.49	168	41.69	67.91

Table 1: Link prediction evaluation across two datasets.

Link Prediction Task. We report main results of link prediction in Table 1. TIME2BOX and all its variants consistently outperform or are on a par with the performance on SOTAs in terms of MRR, MR, HITS@1 and HITS@10. On WIKIDATA12k, TIME2BOX outperforms TIMEPLEX base by around 3 points in terms of MRR and over 5 points in HITS@1. On WIKIDATA114k, TIME2BOX is slightly better than two SOTAs in general for MRR, HITS@1 and HITS@10. In addition, we notice that TIME2BOX beats SOTAs by large margins in time interval-based link prediction, as shown in

⁷Link to online Appendices.

Table 8 in Appendix D. Our method improves around 20 and 7 HITS@1 points in terms of half-open interval-based link prediction and closed interval-based link prediction on WIKIDATA12k, respectively. On WIKIDATA114k TIME2BOX improves around 6 and 4 HITS@1 points, respectively.

Another critical observation in Table 1 is the substantial improvements of using TIME2BOX in terms of MR on both datasets. TIME2BOX returns an MR of 100 and 139 on WIKIDATA12k and WIKIDATA114k, respectively and TIMEPLEX base obtains 302 and 337 for MR on both datasets. It indicates that TIME2BOX is capable of giving a fair rank for a gold answer to any test query on average. This is likely because of the idea of using boxes to constraint the potential answer set. As a time-agnostic box is optimized towards embracing entities satisfying atemporal queries of the form $(s, r, ?o)$ in the learning process, boxes implicitly manage to learn common characteristics of the satisfied $?o$. Therefore, TIME2BOX is less likely to output extremely bad predictions. Examples in Section 6.4 exemplify this hypothesis.

Time Prediction Task. Table 2 and Table 3 summarizes the results for two datasets. On both datasets, TIME2BOX and its variants consistently outperform SOTAs by significant margins. Specifically, TIME2BOX improves over TIMEPLEX by about 5.56, 7.25, and 4.87 points with respect to gIOU@1, aeIOU@1, and gaeIOU@1, respectively, on WIKIDATA12k. As for WIKIDATA114k, despite subtle improvements in link prediction, the advancement of TIME2BOX is more pronounced in time prediction, which shows that it gains 8.7, 5.87, and 4.66 points on gIOU@1, aeIOU@1, and gaeIOU@1, respectively. Furthermore, the improvements on gaeIOU@10 are much more notable with gains of 15.81 and 11.07 points on the two datasets, respectively.

Datasets	WIKIDATA12k					
	gIOU@1	gIOU@10	aeIOU@1	aeIOU@10	gaeIOU@1	gaeIOU@10
TNT-Complex	31.44	55.18	18.86	40.94	11.01	29.51
TIMEPLEX base	35.63	60.86	18.60	37.75	12.61	32.63
TIME2BOX-TR	39.63	67.83	23.47	44.64	15.87	41.53
TIME2BOX-DM	38.78	62.44	21.91	41.55	14.94	37.14
TIME2BOX-SI	39.68	65.30	23.66	42.16	16.09	38.54
TIME2BOX-TNS	42.30	70.16	25.78	50.04	17.41	47.54
TIME2BOX	41.20	68.53	24.70	46.05	16.98	43.08

Table 2: Time prediction evaluation on WIKIDATA12k.

Datasets	WIKIDATA114k					
	gIOU@1	gIOU@10	aeIOU@1	aeIOU@10	gaeIOU@1	gaeIOU@10
TNT-Complex	27.94	48.31	16.18	35.32	7.31	23.68
TIMEPLEX base	29.31	57.68	18.56	36.70	12.53	32.47
TIME2BOX-TR	37.49	67.95	25.05	49.02	15.41	45.72
TIME2BOX-DM	35.88	66.62	24.33	48.03	14.89	44.48
TIME2BOX-SI	34.02	62.89	23.10	44.74	14.07	40.05
TIME2BOX-TNS	37.31	66.91	25.07	48.18	15.57	44.66
TIME2BOX	38.01	71.29	24.42	50.07	15.88	47.77

Table 3: Time prediction evaluation on WIKIDATA114k.

6.4 Qualitative Study

Table 4 showcases examples of timestamp-based link prediction on WIKIDATA12k. The comparison between TIMEPLEX base and TIME2BOX reveals that TIME2BOX is able to learn common characteristics of entities by adopting boxes. For instance, the predicted top 10 returned by TIME2BOX are possible affiliations (e.g., institutes, colleges, universities) in the first query and are countries in the second query. By contrast, TIMEPLEX base returns a mixture of entities with distinct classes for both queries. Furthermore, Table 5 shows an example of time interval-based link prediction, in which

Query Example 1: (Yury Vasilyevich Malyshev, educatedAt, ?, 1977)	
TIMEPLEX base	TIME2BOX
1. Bauman Moscow State Technical University, 2. Gold Star, 3. Communist Party of the Soviet Union, 4. Order of Lenin, 5. S.P. Korolev Rocket and Space Corporation Energia, 6. Hero of the Soviet Union, 7. Gagarin Air Force Academy , 8. Balashov Higher Military Aviation School of Pilots, 9. Ashok Chakra, 10. Heidelberg University	1. Bauman Moscow State Technical University, 2. Gagarin Air Force Academy , 3. S.P. Korolev Rocket and Space Corporation Energia, 4. Saint Petersburg State Polytechnical University, 5. University of Oxford, 6. Saint Petersburg State University, 7. Steklov Institute of Mathematics, 8. Leipzig University, 9. Heidelberg University, 10. Moscow Conservatory
Query Example 2: (Pedro Pablo Kuczynski, countryOfCitizenship, ?, 2015)	
TIMEPLEX base	TIME2BOX
1. doctor honoris causa, 2. President of Peru, 3. Minister of Economy and Finance of Peru, 4. Grand Cross of the Order of the Sun of Peru, 5. President of the Council of Ministers of Peru, 6. World Bank, 7. Serbia, 8. Royal Spanish Academy, 9. Meurthe-et-Moselle, 10. Norwegian Sportsman of the Year	1. France, 2. Germany, 3. United States of America , 4. Austria, 5. Romania, 6. United Kingdom, 7. Poland, 8. Kingdom of Italy, 9. Russian Soviet Federative Socialist Republic, 10. Russian Empire

Table 4: Examples of timestamp-based link prediction on WIKIDATA12k. Top 10 entities predicted by TIMEPLEX base and TIME2BOX are numbered, where 1 denotes Top One. Correct answers are in bold.

Query Example: (Kyiv, country, ?, [1905, 1919])	
Gold Answers: (1)[1905, 1916]->Russian Empire; (2)[1917, 1919]->Ukrainian People's Republic	
year	TIMEPLEX
1905	Soviet Union
1906	Soviet Union
1907	Russian Empire
1908	Soviet Union
1909	Ukrainian Soviet Socialist Republic
1910	Soviet Union
1911	Russian Empire
1912	Ukrainian Soviet Socialist Republic
1913	Soviet Union
1914	Ukrainian Soviet Socialist Republic
1915	Soviet Union
1916	Ukrainian People's Republic
1917	Ukrainian People's Republic
1918	Ukrainian People's Republic
1919	Ukrainian People's Republic

Table 5: An example of interval-based link prediction on WIKIDATA12k. For time interval-based link prediction, the current strategy is to discretize intervals to timestamps and average ranks for each timestamp-based prediction result as the final evaluation. Only top 1 predictions are shown here. TIME2BOX is able to consistently output correct predictions across time and precisely discern the changes of objects over time (i.e., the correct answer shifts from Russian Empire to Ukrainian People's Republic in 1916), while TIMEPLEX base fails. This can be attributed to the ability of TIME2BOX to capture the order of timestamps and the idea of temporal boxes as a constraint over potential answer entities. Hence, answer entities that are true in two consecutive years can be enclosed in the intersection of temporal boxes.

6.5 Model Variation Study

In this section, we report on observations of results about different model variations, which are shown in Table 1, 2 and 3. Compared to TIME2BOX-DM, which adopts element-wise product as operators, element-wise addition projectors (TIME2BOX) perform better in link prediction and time prediction on both datasets. Moreover, we observe that explicitly modeling association between time and relation (i.e., TIME2BOX-TR) does not significantly improve the performance of TIME2BOX framework, although it speeds up the convergence at training, indicating that intersection operators are good enough to learn the association between time and relations

implicitly. As for different time-aware strategies incorporated in TIME2BOX-SI and TIME2BOX-TNS, we find that on both datasets TIME2BOX-SI does not outperform TIME2BOX, indicating that using one sample strategy (i.e., Fig. 1b) is better at modeling time interval-based statements in TIME2BOX. In addition, we find that by incorporating time negative samples, the performance on time prediction can be further improved on WIKIDATA12k, although TIME2BOX-TNS is not superior to TIME2BOX in link prediction.

7 CONCLUSION

In this work, we presented a box-based temporal knowledge graph (TKBC) completion framework (called TIME2BOX) to represent and model statements with different types of validity information (i.e., no time, known start time, known end time, instant, both start and end time) in a vector space. We argued that a TKBC problem can be solved in two steps. First by solving an atemporal KBC problem and then narrowing down the correct answer sets that are only true at the time of interest. Therefore, we introduced time-agnostic boxes to model sets of answers obtained from KBC models. Time-aware boxes are used as a filter to pick out time-dependent answers. TIME2BOX outperforms existing TKBC methods in both link prediction and time prediction on WIKIDATA12k and WIKIDATA114K. By investigating the model performance on statements with different types of validity information, we found that the improvement of TIME2BOX largely attributes to its better ability to handle statements with interval-based validity information. In the future, we will explore how to include spatial scopes of statements so that KBC can benefit from both spatial and temporal scoping.

ACKNOWLEDGMENTS

This work was partially supported by the NSF award 2033521.

REFERENCES

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [2] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems* 26 (2013), 2787–2795.
- [3] Shib Sankar Dasgupta, Swayambhu Nath Ray, and Partha Talukdar. 2018. Hyte: Hyperplane-based temporally aware knowledge graph embedding. In *EMNLP*.
- [4] Alberto García-Durán, Sebastijan Dumančić, and Mathias Niepert. 2018. Learning Sequence Encoders for Temporal Knowledge Graph Completion. In *EMNLP*. Association for Computational Linguistics, Brussels, Belgium, 4816–4821.
- [5] Prachi Jain, Sushant Rath, Mausam, and Soumen Chakrabarti. 2020. Temporal Knowledge Base Completion: New Algorithms and Evaluation Protocols. In *EMNLP*. Association for Computational Linguistics, Online, 3733–3747.
- [6] Woojeong Jin, Meng Qu, Xisen Jin, and Xiang Ren. 2020. Recurrent Event Network: Autoregressive Structure Inference over Temporal Knowledge Graphs. In *EMNLP*. 6669–6683.
- [7] Timothée Lacroix, Guillaume Obozinski, and Nicolas Usunier. 2019. Tensor Decompositions for Temporal Knowledge Base Completion. In *ICML*.
- [8] Hongyu Ren, Weihua Hu, and Jure Leskovec. 2019. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. In *ICML*.
- [9] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. *arXiv preprint arXiv:1902.10197* (2019).
- [10] Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. 2017. Know-evolve: deep temporal reasoning for dynamic knowledge graphs. In *ICML-Volume 70*.
- [11] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. *ICML*.
- [12] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575* (2014).