

# TeRo: A Time-aware Knowledge Graph Embedding via Temporal Rotation

Chengjin Xu and Mojtaba Nayyeri and Fouad Alkhoury and Hamed Shariat Yazdi

Smart Data Analytics Group, University of Bonn / Germany

{xuc, nayyeri}@iai.uni-bonn.de

s6foalkh@uni-bonn.de

shariatyazdi@gmail.com

Jens Lehmann

University of Bonn / Germany

Fraunhofer IAIS/ Germany

jens.lehmann@iais.fraunhofer.de

## Abstract

In the last few years, there has been a surge of interest in learning representations of entities and relations in knowledge graph (KG). However, the recent availability of temporal knowledge graphs (TKGs) that contain time information for each fact created the need for reasoning over time in such TKGs. In this regard, we present a new approach of TKG embedding, TeRo, which defines the temporal evolution of entity embedding as a rotation from the initial time to the current time in the complex vector space. Specially, for facts involving time intervals, each relation is represented as a pair of dual complex embeddings to handle the beginning and the end of the relation, respectively. We show our proposed model overcomes the limitations of the existing KG embedding models and TKG embedding models and has the ability of learning and inferring various relation patterns over time. Experimental results on four different TKGs show that TeRo significantly outperforms existing state-of-the-art models for link prediction. In addition, we analyze the effect of time granularity on link prediction over TKGs, which as far as we know has not been investigated in previous literature.

## 1 Introduction

In recent years, a number of sizable Knowledge Graphs (KGs) have been constructed, including DBpedia (Auer et al., 2007), YAGO (Suchanek et al., 2007), Nell (Carlson et al., 2010) and Freebase (Bollacker et al., 2008). In these KGs, a fact is represented as a triple  $(s, r, o)$ , where  $s$  (subject) and  $o$  (object) are entities (nodes), and  $r$  (relation) is the relation (edge) between them.

Several KG embedding (KGE) models are developed to perform learning and inference over these KGs (Bordes et al., 2013; Yang et al., 2014; Trouillon et al., 2016; Sun et al., 2019; Zhang et al., 2019). The most common learning task for these models is link prediction, which is to complete a fact with the missing entity. For instance, one can use a KGE model to perform an object query like (*Barack Obama, visits, ?*). In this case, there are several valid answers to this question, regardless of the time factor. Obviously, the inclusion of time information can make this query more specific, e.g., (*Barack Obama, visits, ?, 2014-07-08*).

Some temporal KGs (TKGs) including ICEWS (Lautenschlager et al., 2015), GDELT (Leetaru and Schrod, 2013), YAGO3 (Mahdisoltani et al., 2013) and Wikidata (Erxleben et al., 2014) store billions of time-aware facts as quadruples  $(s, r, o, t)$  where  $t$  is the time annotation, e.g., (*Barack Obama, visits, Ukraine, 2014-07-08*). The availability of these TKGs that exhibits complex temporal dynamics in addition to its multi-relational nature has created the need for approaches that can characterize and reason over them. Traditional KGE models disregard time information, leading to an ineffectiveness of performing link prediction on TKGs involving temporary relations (e.g., *visits, live in*, etc.).

---

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

To tackle this problem, TKG embedding (TKGE) models encode time information in their embeddings. Such TKGE models (Trivedi et al., 2017; Leblay and Chekol, 2018; García-Durán et al., 2018; Ma et al., 2018; Dasgupta et al., 2018; Jin et al., 2019) were shown to have better performances on link prediction over TKGs than traditional KGE models. However, most of the existing TKGE models are the extensions of TransE (Bordes et al., 2013) and DistMult (Yang et al., 2014), and thus are not fully expressive for some relation patterns (Sun et al., 2019).

In this paper, we propose a novel approach for TKGEs, TeRo, which defines the temporal evolution of an entity embedding as a rotation from the initial time to the current time in the complex vector space. We show the limitation of the existing TKGE models and the advantage of our proposed model on learning various relation patterns over time.

Specially, for facts involving time intervals, each relation is represented as a pair of dual complex embeddings which are used to handle the beginning and the end of the relation, respectively. In this way, TeRo can adapt well to datasets where time annotations are represented in the various forms: time points, beginning or end time, time intervals.

Most of previous TKGE-related works as far as we know use specific time granularities for various TKGs. For example, the time granularity of the ICEWS datasets is fixed as 24 hours in (Dasgupta et al., 2018; Xu et al., 2019). In this work, we adopt various time-division approaches for different TKG datasets and investigate the effect of the length of the time steps on the performance of our model.

To verify our approach, we compare the performance of our proposed models on link prediction and time prediction tasks over four different TKGs with the state-of-the-art KGE models and the existing TKGE models. The experimental results demonstrate that our proposed model outperforms other baseline models significantly by inferring various relation patterns and encoding time information.

## 2 Related Work

KGE models can be roughly classified into distance-based models and semantic matching models.

Distance-based models measure the plausibility of a fact as the distance between the two entities, usually after a translation or rotation carried out by the relation. A typical example of distance-based models is TransE (Bordes et al., 2013). TransE exhibits deficiencies when learning 1-n relations. Thus, various extensions of TransE (Wang et al., 2014; Ji et al., 2015; Lin et al., 2015; Nayyeri et al., 2019), were proposed to tackle this problem. They use different mapping methods to project entities from entity space to relation space. Specially, RotatE (Sun et al., 2019) defines each relation as a rotation from the subject to the object. Nevertheless, these distance-based distance models are still unable to capture reflexive relations which can hold, i.e. for a particular relation  $r$  each entity is related to itself via  $r$ . In distance-based models, the values or the phases of vectors for all reflexive relations are enforced to be 0, which does not allow to fully express the semantic characteristics of these relations.

Semantic matching models measure plausibility of facts by matching latent semantics of entities and relations embodied in their embedding representations. A few examples of such models include RESCAL (Nickel et al., 2011), DistMult (Yang et al., 2014), ComplEx (Trouillon et al., 2016), QuatE (Zhang et al., 2019) and GeomE (Xu et al., 2020). RESCAL and DistMult cannot capture asymmetric relations since the score of the triple  $(s, r, o)$  is always equal to the score of its symmetric triple  $(o, r, s)$ . ComplEx, QuatE and GeomE have been proven to be able to capture various relation patterns for static KGs, but cannot model temporary relations in TKGs due to their ignorance of time information.

Recent research illustrated that the performances of KGE models can be further improved by incorporating time information in TKGs. Some TKGE models are extended from TransE, e.g., TTransE (Leblay and Chekol, 2018), TA-TransE (García-Durán et al., 2018), HyTE (Dasgupta et al., 2018) and ATiSE (Xu et al., 2019). Another part of TKGE models are temporal extensions of DistMult, e.g., Know-Evolve (Trivedi et al., 2017), TDistMult (Ma et al., 2018) and TA-DistMult (García-Durán et al., 2018). Similar to TransE and DistMult, these TKGE models have issues with capturing reflexive relations or asymmetric relations. Specially, DE-Simple (Goel et al., 2020) incorporates time information into diachronic entity embeddings and has capability of modeling various relation patterns. However, this approach only focuses on event-based TKG datasets, and cannot model facts involving time intervals

shaped like [2003-##-##, 2005-##-##].

### 3 A Novel TKGE Approach based on Temporal Rotation

Although various KGE models have been developed to learn multi-relational interactions between entities, all of them have problems with inferring temporary relations which are only valid for a certain time point or a last for a certain time period. To illustrate this by example, assume we are given a quadruple (*Barack Obama, visits, France, 2014-02-12*) as a training sample, where the relation *visits* is a temporary relation. If we query (*Barack Obama, visits, ?, 2014-07-08*), a trained static KGE model probably returns the incorrect answer *France* due to the validness of the triple *Barack Obama, visits, France*, while the correct answer is *Ukraine* considering the given time constraint. On the other hand, most of the existing TKGE models, which were extended from TransE (Bordes et al., 2013) and DistMult (Yang et al., 2014), incorporate time information in the embedding space, but have limitations on learning transitive relations or asymmetric relations as discussed in Section 2.

To overcome the limitations of these existing KGE and TKGE models on learning and inferring over TKGs, we propose a new TKGE model, TeRo, which defines the temporal evolution of an entity embedding as a rotation in the complex vector space. Let  $\mathcal{E}$  denote the set of entities,  $\mathcal{R}$  denote the set of relations, and  $\mathcal{T}$  denote the set of time steps. Then a TKG is a collection of factual quadruples  $(s, r, o, t)$ , where  $s, o \in \mathcal{E}$  are the subject and object entities,  $r \in \mathcal{R}$  is the relation,  $t$  denotes the actual time when the fact occurs. For any time  $t$ , we have a time step  $\tau \in \mathcal{T}$  representing this actual time. We map  $s, r, o$  to their complex embeddings, i.e.,  $\mathbf{s}, \mathbf{r}, \mathbf{o} \in \mathbb{C}^k$ ; then we define the functional mapping induced by each time step  $\tau$  as an element-wise rotation from the time-independent entity embeddings  $\mathbf{s}$  and  $\mathbf{o}$  to the time-specific entity embeddings  $\mathbf{s}_t$  and  $\mathbf{o}_t$ . The mapping function is defined as follows:

$$\mathbf{s}_t = \mathbf{s} \circ \boldsymbol{\tau}, \quad \mathbf{o}_t = \mathbf{o} \circ \boldsymbol{\tau} \quad (1)$$

where  $\circ$  denotes the Hermitian dot product between complex vectors. Here, we constrain the modulus of each element of  $\boldsymbol{\tau} \in \mathbb{C}^k$ , i.e.,  $\tau_j \in \mathbb{C}$ , to be  $|\tau_j| = 1$ . By doing this,  $\tau_j$  is of the form  $e^{i\theta_{\tau,j}}$ , which corresponds to a counter-clockwise rotation by  $\theta_{\tau,j}$  radians around the origin of the complex plane, and only affects the phases of the entity embeddings in the complex vector space. This idea is motivated by Euler's identity  $e^{i\theta} = \cos\theta + i\sin\theta$ , which indicates that a unitary complex number can be regarded as a rotation in the complex plane.

We regard the relation embedding  $\mathbf{r}$  as translation from the time-specific subject embedding  $\mathbf{s}_t$  to the conjugate of the time-specific object embedding  $\bar{\mathbf{o}}_t$  for a single quadruple  $(s, r, o, t) \in \mathcal{Q}^+$ , where  $r \in \mathcal{R}_b \cup \mathcal{R}_e$  and  $\mathcal{Q}^+$  denotes the set of all positive quadruples. The score function is defined as:

$$f_{\text{TeRo}}(s, r, o, t) = \|\mathbf{s}_t + \mathbf{r} - \bar{\mathbf{o}}_t\| \quad (2)$$

For a fact  $(s, r, o, t)$  occurring in a certain time interval, i.e.,  $t = [t_b, t_e]$  where  $t_b, t_e$  denote the beginning time and the end time of the fact, we separate this fact into two quadruples, namely,  $(s, r_b, o, t_b)$  and  $(s, r_e, o, t_e)$ . Here, we extend the relation set  $\mathcal{R}$  in a TKG which involves time intervals to a pair of dual relation sets,  $\mathcal{R}_b$  and  $\mathcal{R}_e$ . A relation  $r_b \in \mathcal{R}_b$  is used to handle the beginning of relation  $r$ , meanwhile a relation  $r_e \in \mathcal{R}_e$  is used to handle the end of relation  $r$ . By doing this, we score a fact  $(s, r, o, [t_r, t_e])$  as the mean value of scores of two quadruples,  $(s, r_b, o, t_b)$  and  $(s, r_e, o, t_e)$  which represent the beginning and the end of this fact respectively.

$$f_{\text{TeRo}}(s, r, o, [t_b, t_e]) = \frac{1}{2} (\|\mathbf{s}_{t_b} + \mathbf{r}_b - \bar{\mathbf{o}}_{t_b}\| + \|\mathbf{s}_{t_e} + \mathbf{r}_e - \bar{\mathbf{o}}_{t_e}\|) \quad (3)$$

Specially, for a fact missing either the beginning time or the end time, e.g.,  $(s, r, o, [t_b, -])$  or  $(s, r, o, [-, t_e])$ , the score of this fact is equal to the score of the quadruple involving the known time, i.e.,  $f_{\text{TeRo}}(s, r, o, [t_b, -]) = f_{\text{TeRo}}(s, r_b, o, t_b)$ ,  $f_{\text{TeRo}}(s, r, o, [-, t_e]) = f_{\text{TeRo}}(s, r_e, o, t_e)$ .

In this paper, we use the same loss function as the negative sampling loss proposed in (Sun et al., 2019) for optimizing our model. This loss function has been proved to be very effective on optimizing

some distance-based KGE models, e.g., TransE, RotatE (Sun et al., 2019) and ATiSE (Xu et al., 2019).

$$L(\xi) = -\log \sigma(\gamma - f_{\text{TeRo}}(\xi)) - \sum_{i=1}^{\eta} \frac{1}{\eta} \log \sigma(f_{\text{TeRo}}(\xi'_i) - \gamma) \quad (4)$$

where  $\xi \in \mathcal{Q}^+$  is a positive training quadruple,  $\xi'_i$  is the  $i$ th negative sample corresponding to  $\xi$  generated by randomly corrupting the subject or the objects of  $\xi$  such as  $(s', p, o, t)$  and  $(s, p, o', t)$ ,  $\sigma(\cdot)$  denotes the sigmoid function,  $\gamma$  is a fixed margin and  $\eta$  is the ratio of negatives over positive training samples.

### 3.1 Learning Various Relation Patterns

Static KGE models and some existing TKGE models which are the temporal extensions of TransE or DistMult have limitations on capturing some key relation patterns which are defined as follows.

**Definition 1.** A relation  $r$  is a **temporary** relation if  $\exists s, o, t_1, t_2 \ r(s, o, t_1) \wedge \neg r(s, o, t_2)$  holds True

**Definition 2.** A relation  $r$  is **asymmetric** if  $\exists s, o, t \ r(s, o, t) \wedge \neg r(o, s, t)$  holds True.

**Definition 3.** A relation  $r$  is a **reflexive** relation if  $\exists s, t \ r(s, s, t)$  holds True.

As mentioned in Section 2, static KGE models can not model temporary relations, e.g., 'visits', since  $f_{\text{SKGE}}(s, r, o, t_1) \equiv f_{\text{SKGE}}(s, r, o, t_2)$ . Temporal extensions of DistMult (denoted as T-DistMult) including TDistMult, TA-DistMult and Know-Evolve can not model asymmetric relations, e.g., 'parentOf', since  $f_{\text{T-DistMult}}(s, r, o, t) = \langle \mathbf{s}_t, \mathbf{r}_t, \mathbf{o}_t \rangle = f_{\text{T-DistMult}}(o, r, s, t)$ , where  $\mathbf{s}_t, \mathbf{o}_t, \mathbf{r}_t$  are time-specific entity/relation embeddings corresponding to different T-DistMult models. Temporal extensions of TransE (denoted as T-TransE) including HyTE, TTransE, TA-TransE have difficulties of modeling multiple reflexive relations, e.g., 'equalTo' and 'subsetOf', since  $f_{\text{T-TransE}}(s_1, r_1, s_1, t) = f_{\text{T-TransE}}(s_2, r_2, s_2, t) = 0 \Rightarrow \mathbf{r}_1 = \mathbf{r}_2 = 0$ .

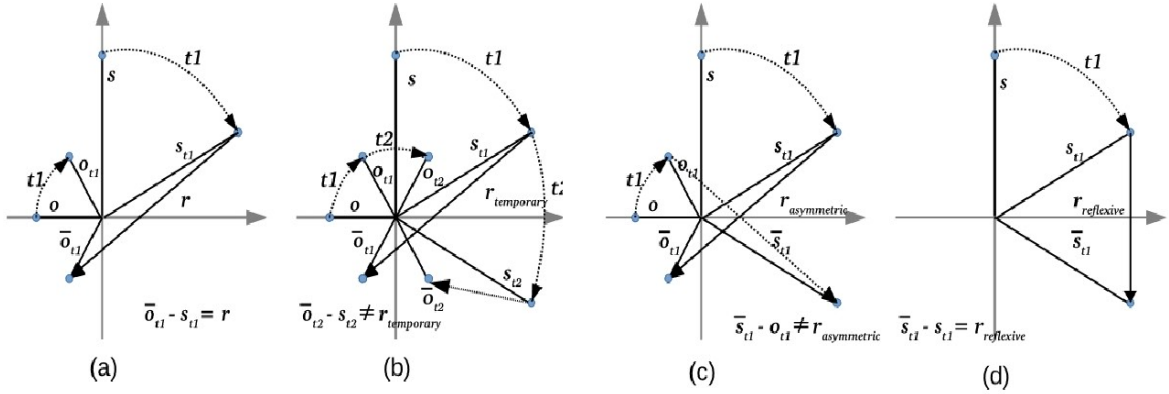


Figure 1: (a) Illustration of TeRo with only one embedding dimension; (b) an example of modeling a temporary relation; (c) an example of modeling an asymmetric relation; (d) an example of modeling a reflexive relation.

By defining each time step as a rotation in the complex vector spaces, TeRo can capture all of the above three relation patterns. Given an observed fact  $(s, r, o, t_1)$  where  $\mathbf{s}_{t_1} + \mathbf{r} = \bar{\mathbf{o}}_{t_1}$ :

- as shown in Figure 1(b), if  $r$  is a temporary relation, we can have  $\mathbf{s}_{t_2} + \mathbf{r} \neq \bar{\mathbf{o}}_{t_2}$  for TeRo to make  $r(s, o, t_1) \wedge \neg r(s, o, t_2)$  hold true.
- as shown in Figure 1(c), if  $r$  is an asymmetric relation, we can have  $\mathbf{o}_{t_1} + \mathbf{r} \neq \bar{\mathbf{s}}_{t_1}$  for TeRo to make  $r(s, o, t_1) \wedge \neg r(o, s, t_1)$  hold true.
- as shown in Figure 1(d), if  $r$  is a reflexive relation, we have  $\text{Im}(\mathbf{r}) = 2\text{Im}(\mathbf{s}_{t_1})$  for TeRo. Thus, TeRo can represent multiple reflexive relations as different embeddings due to the conjugate operations of object embeddings.

### 3.2 Complexity

In Table 1, we summarize the scoring functions and the space complexities of several state-of-the-art TKGE approaches and our model as well as TransE.  $n_e$ ,  $n_r$ ,  $n_\tau$  and  $n_{token}$  are numbers of entities, relations, time steps and temporal tokens used in (García-Durán et al., 2018);  $d$  is the dimensionality of embeddings.  $\mathcal{P}_t(\cdot)$  denotes the temporal projection for embeddings (Dasgupta et al., 2018).  $\text{LSTM}(\cdot)$  denotes an LSTM neural network;  $[\mathbf{r}; \mathbf{t}_{seq}]$  denotes the concatenation of the relation embedding and the sequence of temporal tokens (García-Durán et al., 2018).  $\rightarrow$  and  $\leftarrow$  denote the temporal part and untemporal part of a time-specific diachronic entity embedding (Goel et al., 2020);  $\mathbf{r}^{-1}$  denotes the inverse relation embedding of  $r$ , i.e.,  $(s, r, o, t) \leftrightarrow (o, r^{-1}, s, t)$ .  $\mathcal{D}_{\mathcal{KL}}(\cdot)$  denotes the KL divergence between two Gaussian distributions;  $\mathbf{P}_{s,t}$ ,  $\mathbf{P}_{o,t}$ ,  $\mathbf{P}_{r,t}$  denote the Gaussian embeddings of  $s$ ,  $r$  and  $o$  at time  $t$  (Xu et al., 2019).

As shown in Table 1, the space complexity of TeRo and TransE will be close if  $n_\tau < n_e$ . In practice, we can achieve this condition by tuning the time granularity.

Model	Scoring Function	Space Complexity
TransE	$\ \mathbf{s} + \mathbf{r} - \mathbf{o}\ $	$\mathcal{O}(n_e d + n_r d)$
TTransE	$\ \mathbf{s} + \mathbf{r} + \boldsymbol{\tau} - \mathbf{o}\ $	$\mathcal{O}(n_e d + n_r d + n_\tau d)$
HyTE	$\ P_t(\mathbf{s}) + P_t(\mathbf{r}) - P_t(\mathbf{o})\ $	$\mathcal{O}(n_e d + n_r d + n_\tau d)$
TA-TransE	$\ \mathbf{s} + \text{LSTM}([\mathbf{r}; \mathbf{t}_{seq}]) - \mathbf{o}\ $	$\mathcal{O}(n_e d + n_r d + n_{token} d)$
TA-DistMult	$\langle \mathbf{s}, \text{LSTM}([\mathbf{r}; \mathbf{t}_{seq}]), \mathbf{o} \rangle$	$\mathcal{O}(n_e d + n_r d + n_{token} d)$
DE-SimplE	$\frac{1}{2}(\langle \vec{\mathbf{s}}^t, \mathbf{r}, \overleftarrow{\mathbf{o}}^t \rangle + \langle \vec{\mathbf{o}}^t, \mathbf{r}^{-1}, \overleftarrow{\mathbf{s}}^t \rangle)$	$\mathcal{O}(n_e d + n_r d)$
ATiSE	$\mathcal{D}_{\mathcal{KL}}(\mathbf{P}_{s,t} - \mathbf{P}_{o,t}, \mathbf{P}_{r,t})$	$\mathcal{O}(n_e d + n_r d)$
TeRo	$\ \mathbf{s}_t + \mathbf{r} - \overline{\mathbf{o}}_t\ $	$\mathcal{O}(n_e d + n_r d + n_\tau d)$

Table 1: Comparison of our models with several baseline models for space complexity.

## 4 Experiments

### 4.1 Temporal Knowledge Graph Datasets

Common TKG benchmarks include GDEL (Goel et al., 2020), ICEWS14, ICEWS05-15, YAGO15k, Wikidata11k (García-Durán et al., 2018), YAGO11k and Wikidata12k (Dasgupta et al., 2018). In this work, we choose ICEWS14, ICEWS05-15, YAGO11k and Wikidata12k as datasets for the following reasons: 1. ICEWS14 and ICEWS05-15 are two well-established event-based datasets which are commonly used in previous literature (García-Durán et al., 2018; Goel et al., 2020; Xu et al., 2019), these two datasets are subsets of ICEWS (Lautenschlager et al., 2015) corresponding to facts in 2014 and facts between 2005 and 2015, where all time annotations are time points; 2. YAGO15k, Wikidata11k, YAGO11k and Wikidata12k are subsets of YAGO3 (Mahdisoltani et al., 2013) and Wikidata (Erleben et al., 2014) where a part of time annotations are time intervals. In YAGO15k and Wikidata11k, each time interval only contains either beginning dates or end dates, shaped like 'occurSince 2003' or 'occurUntill 2005' and a part of facts in YAGO15k exclude time information. Thus we prefer to using YAGO11k and Wikidata12k where each fact includes time information and time annotations are represented in the various forms, i.e., time points like [2003-01-01, 2003-01-01], beginning or end time like [2003, ##], and time intervals like [2003, 2005]. We list the statistics of the four datasets we use in Table 2.

Dataset	#Entities	#Relations	Time Span	#Training	#Validation	#Test
ICEWS14	6,869	230	2014	72,826	8,941	8,963
ICEWS05-15	10,094	251	2005-2015	368,962	46,275	46,092
YAGO11k	10,623	10	-453-2844	16,406	2,050	2,051
Wikidata12k	12,554	24	1479-2018	32,497	4,062	4,062

Table 2: Statistics of datasets.

## 4.2 Time Granularity

In some recent work (Dasgupta et al., 2018; Xu et al., 2019), the time span of a TKG dataset was splitted into a number of time steps. For ICEWS14 and ICEWS05-15, the time granularity was fixed as 1 day. For YAGO11k and Wikidata12k, month and day information was dropped, and less frequent year mentioned were clubbed into same time steps but years with high frequency formed individual time steps in order to alleviate the effect of the long-tail property of time data. In other words, the lengths of different time steps were different for the balance of numbers of triples in different time steps. However, it has not been investigated whether the lengths of time steps affect the performances of TKGE models.

In this work, we test our model with different time units, denoted as  $u$ , in a range of  $\{1, 2, 3, 7, 14, 30, 90 \text{ and } 365\}$  days for ICEWS datasets. Dasgupta et al. (2018) and Xu et al. (2019) applied a minimum threshold of 300 triples per interval during construction for YAGO11k and Wikidata12k. We follow their time-division approaches for these two datasets and test different minimum thresholds, denoted as  $thre$ , in a range of  $\{1, 3, 10, 30, 100, 300, 1000, 3000, 10000, 30000\}$ . The change of time granularity will reconstruct the set of time steps  $\mathcal{T}$ . For ICEWS14, when the time unit  $u$  is 1 day, we have totally 365 time steps and the date *2014-01-02* is represented by the second time step, i.e.,  $\tau_1$ . If the time unit is changed as 2 days, the total number of time steps will be 183 and the date *2014-01-02* will be denoted as  $\tau_0$ . For YAGO11k, when the mini threshold  $thre = 1$ , we have 396 time steps since there are totally 396 different years existing as timestamps in YAGO11k. Years like *-453*, *100* and *2008* are all taken as independent time steps. When  $thre$  for YAGO11k rises to 300, the number of time steps drops to 127 and years between *-431* and *100* are clubbed into a same time step.

## 4.3 Evaluation Metrics

We evaluate our model by testing the performances of our model on link prediction task over TKGs under the time-wise filtered setting defined in (Xu et al., 2019; Goel et al., 2020). This task is to complete a time-wise fact with a missing entity. For a test quadruple  $(s, p, o, t)$ , we first generate candidate quadruples  $\mathcal{C} = \{(s, p, o', t) : o' \in \mathcal{E}\} \cup \{(s', p, o, t) : s' \in \mathcal{E}\}$  by replacing  $s$  or  $o$  with all possible entities. Different from the time-unwise filtered setting (Bordes et al., 2013) which filters the triples appearing either in the training, validation or test set from the candidate list, we only filter the quadruples  $\xi \in \mathcal{Q}^+$  existing in the dataset. This ensures that the facts which do not appear at time  $t$  are still considered as candidates for evaluating the given test quadruple. We obtain the final rank of the test quadruple among filtered candidate quadruples  $\bar{\mathcal{C}} = \{\xi : \xi \in \mathcal{C}, \xi \notin \mathcal{Q}^+\}$  by sorting their scores.

Two commonly used evaluation metrics are used here, i.e., Mean Reciprocal Rank and Hits@k. The Mean Reciprocal Rank (MRR) is the means of the reciprocal values of all computed ranks. And the fraction of test quadruples ranking in the top  $k$  is called Hits@k.

## 4.4 Baselines

We compare our approach with several state-of-the-art KGE approaches and existing TKGE approaches, including TransE (Bordes et al., 2013), DistMult (Yang et al., 2014), ComplEx-N3 (Lacroix et al., 2018), RotatE (Sun et al., 2019), QuatE (Zhang et al., 2019), TTransE (Leblay and Chekol, 2018), TA-TransE, TA-DistMult (García-Durán et al., 2018), DE-Simple (Goel et al., 2020) and ATiSE (Xu et al., 2019). The results of most baselines are taken from some recent work (Goel et al., 2020; Xu et al., 2019) which used the same evaluation protocol as ours. DE-Simple which mainly focuses on event-based datasets, cannot model time intervals or time annotations missing month and day information which are common in YAGO and Wikidata. Thus its result on YAGO11k and Wikidata12k are unobtainable. Since the original source code of TA-TransE and TA-DistMult (García-Durán et al., 2018) is not released, we reimplement these models according to the implementation details reported in the original paper, in order to obtain their results on YAGO11k and Wikidata12k.

## 4.5 Experimental Setup

We implement our proposed model in PyTorch. The code is available at <https://github.com/soledad921/ATiSE>.

We select the optimal hyperparameters by early validation stopping according to MRR on the validation set. We restrict the iterations to 5000. Following the setup used in (Xu et al., 2019), the batch size  $b = 512$  is kept for all datasets, the embedding dimensionality  $k$  is tuned in  $\{100, 200, 300, 400, 500\}$ , the ratio of negative over positive training samples  $\eta$  is tuned in  $\{1, 3, 5, 10\}$  and the margin  $\gamma$  is tuned in  $\{1, 2, 3, 5, 10, 20, \dots, 120\}$ . Regarding optimizer, we choose Adagrad for TeRo and tune the learning rate  $r$  in a range of  $\{1, 0.3, 0.1, 0.03, 0.01\}$ . Specially, the time granularity parameters  $u$  and  $thre$  are also regraded as hyperparameters for TeRo as mentioned in Section 4.2.

The default configuration for TeRo is as follows:  $d = 500, \eta = 10$ . Below, we only list the non-default parameters:  $lr = 0.1, \gamma = 110, u = 1$  on ICEWS14;  $lr = 0.1, \gamma = 120, u = 2$  on ICEWS05-15;  $lr = 0.1, \gamma = 50, thre = 100$  on YAGO11k;  $lr = 0.3, \gamma = 20, thre = 300$  on Wikidata12k.

## 5 Results and Analysis

### 5.1 Comparative Study

Datasets	ICEWS14				ICEWS05-15			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
TransE*	.280	.094	-	.637	.294	.090	-	.663
DistMult*	.439	.323	-	.672	.456	.337	-	.691
ComplEx-N3 <sup>†</sup>	.467	.347	.527	.716	.481	.362	.535	.729
RotatE <sup>†</sup>	.418	.291	.478	.690	.304	.164	.355	.595
QuatE <sup>2 †</sup>	.471	.353	.530	.712	.482	.370	.529	.727
TTransE <sup>◇</sup>	.255	.074	-	.601	.271	.084	-	.616
HyTE <sup>◇</sup>	.297	.108	.416	.655	.316	.116	.445	.681
TA-TransE*	.275	.095	-	.625	.299	.096	-	.668
TA-DistMult*	.477	.363	-	.686	.474	.346	-	.728
DE-Simple <sup>◇</sup>	.526	.418	.592	.725	.513	.392	.578	.748
ATiSE <sup>†</sup>	.550	.436	<b>.629</b>	<b>.750</b>	.519	.378	.606	.794
TeRo	<b>.562</b>	<b>.468</b>	.621	.732	<b>.586</b>	<b>.469</b>	<b>.668</b>	<b>.795</b>

Table 3: Link prediction results on ICEWS14 and ICEWS05-15. \*: results are taken from (García-Durán et al., 2018). <sup>◇</sup>: results are taken from (Goel et al., 2020). <sup>†</sup>: results are taken from (Xu et al., 2019). Dashes: results are unobtainable. The best results among all models are written bold.

Datasets	YAGO11k				Wikidata12k			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
TransE <sup>†</sup>	.100	.015	.138	.244	.178	.100	.192	.339
DistMult <sup>†</sup>	.158	.107	.161	.268	.222	.119	.238	.460
ComplEx-N3 <sup>†</sup>	.167	.106	.154	.282	.233	.123	.253	.436
RotatE <sup>†</sup>	.167	.103	.167	.305	.221	.116	.236	.461
QuatE <sup>2 †</sup>	.164	.107	.148	.270	.230	.125	.243	.416
TTransE <sup>†</sup>	.108	.020	.150	.251	.172	.096	.184	.329
HyTE <sup>†</sup>	.105	.015	.143	.272	.180	.098	.197	.333
TA-TransE	.127	.027	.160	.326	.178	.030	.267	.429
TA-DistMult	.161	.103	.171	.292	.218	.122	.232	.447
ATiSE <sup>†</sup>	.170	.110	.171	.288	.280	.175	.317	.481
TeRo	<b>.187</b>	<b>.121</b>	<b>.197</b>	<b>.319</b>	<b>.299</b>	<b>.198</b>	<b>.329</b>	<b>.507</b>

Table 4: Link prediction results on YAGO11k and Wikidata12k. <sup>†</sup>: results are taken from (Xu et al., 2019). The best results among all models are written bold.

Table 3 and 4 list all link prediction results of our proposed model and baseline models on four datasets. TeRo surpassed all baseline embedding models regarding all metrics on all datasets except that the ATiSE got the better Hits@3 and Hits@10 than TeRo on ICEWS14. Compared to ATiSE, TeRo achieved the

improvement of 1.2 MRR points, 6.7 MRR points, 1.7 MRR points and 1.9 MRR points on ICEWS14, ICEWS05-15, YAGO11k and Wikidata12k respectively.

## 5.2 Ablation Study

In this work, we analyze the effect of the change of the time granularity on the performance of our model. As mentioned in Section 4.2, we adopt two different time-division approaches for event-based datasets, i.e., ICEWS datasets, and time-wise KGs involving time intervals, i.e., YAGO11k as well as Wikidata12k. For ICEWS14 and ICEWS05-15, we use time steps with fixed length since the the distribution of numbers of facts in ICEWS datasets over time are relatively uniform as shown in Figure 2. The time granularities of ICEWS datasets are equal to the lengths of time units  $u$ . On the other hand, the time distributions of numbers of facts in YAGO15k and Wikidata12k are long-tailed. Thus we divide the time steps in YAGO15k and Wikidata12k by setting a mini threshold for the numbers of facts in each time step. Time granularities of these two datasets can be changed by setting different thresholds  $thre$ .

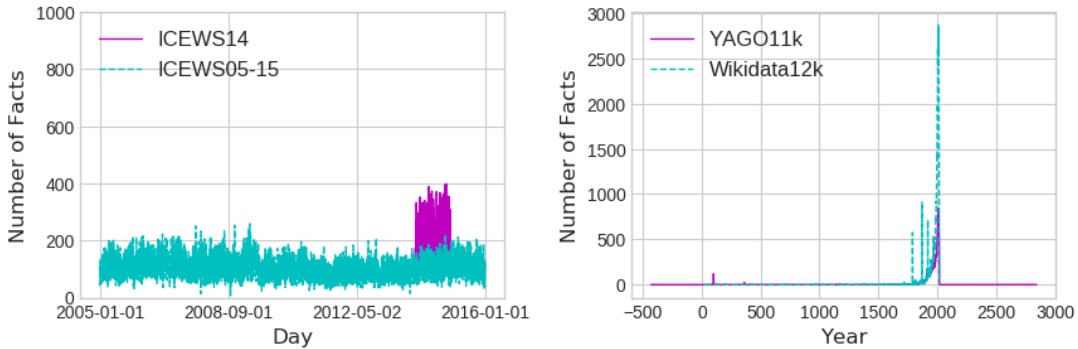


Figure 2: Time distribution of numbers of facts.

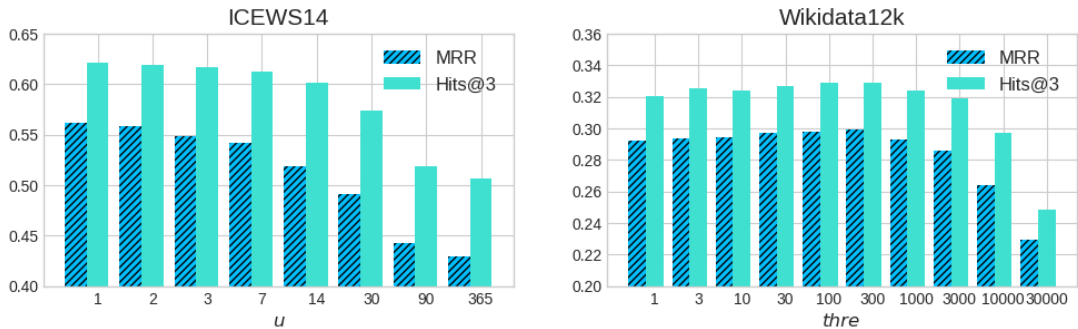


Figure 3: Results of TeRo with different time granularities on ICEWS14 and Wikidata12k.

In ICEWS14, time distribution is relatively uniform and thus representing time with a small time granularity can provide more abundant time information. As shown in Figure 3, TeRo with small time granularities, e.g., 1 day, 2 days and 3 days, had better performance on ICEWS14 compared to TeRo with big time granularities regarding MRR and Hits@3. Likewise, the optimal time unit for TeRo on ICEWS05-15 was proven by our experiments to be 2 days. For Wikidata12k, using a very small time granularity was non-optimal due to the long-tail property of time data. On the other hand, using an overly big time granularity resulted in the invalid incorporation of time information. Figure 3 demonstrates the low performances of TeRo with big time granularities. More concretely, when time unit  $u$  was 1 year, all of time annotations in ICEWS14 were represented by a uniform time embedding, which meant this time embedding was temporally unmeaningful. Table 5 demonstrates a few examples of link prediction results on ICEWS14 of TeRo models with time units  $u$  of two days and one year.

As shown in Table 5, in many cases, TeRo with  $u = 1$  predicated correctly, meanwhile TeRo with  $u = 365$  gave the wrong predictions. We notice that these predictions of TeRo with  $u = 365$  in Table 5 would be valid if we disregarded the time constraint. For instance, (*Colombia, Host a visit, John F. Kelly*) happened on 2014-03-27, (*UN Security Council, Criticize or denounce, Armed Band (South Sudan)*) was



Link Prediction	TeRo with $u = 1$ day	TeRo with $u = 365$ days
Colombia, Host a vist, ?, 2014-06-04	<b>Kyung-wha Kang</b>	John F. Kelly
Head of Government (China), visits, ?, 2014-07-04	<b>South Korea</b>	Serbia
UN Security Council, Criticize or denounce, ?, 2014-08-10	<b>North Korea</b>	Armed Band (South Sudan)
South Korea, Host a vist, ?, 2014-06-20	<b>Kim Jong-Un</b>	National Security Advisor (Japan)
Police (Australia), Accuse, ?, 2014-10-22	<b>Criminal (Australia)</b>	Citizen (Australia)

Table 5: Examples of link prediction results on ICEWS14. The correct predictions are written **bold**.

true on 2014-08-07. As mentioned in Section 3, *Host a visit* and *Criticize or denounce* are temporary relations. The above results prove that using a reasonable time granularity is helpful for TeRo to effectively incorporate time information. And the inclusion of time information enables TeRo to capture temporary relations and improve its performance on link prediction over TKGs.

### 5.3 Efficiency Study

TeRo has the same space complexitiy as TTransE (Leblay and Chekol, 2018) and HyTE (Dasgupta et al., 2018). Since we constrained the numbers of time steps of the four TKG datasets by tuning time granularities (183 time steps in ICEWS14, 1339 time steps in ICEWS05-15, 127 time steps in YAGO11k and 82 time steps in Wikidata12k), the numbers of time steps are much less than the numbers of entities in these datasets, which means that the space complexity of TeRo is close to the space complexity of TransE (Bordes et al., 2013) as mentioned in Section 3.2. Regarding the concrete memory consumption, the recent state-of-the-art TKGE models, ATiSE (Xu et al., 2019) and DE-Simple (Goel et al., 2020) have 1.8 times and 2.2 times as large memory size as TeRo on ICEWS14 with the same embedding dimensionality. The training processes of TeRo with 500-dimensional embeddings on ICEWS14, ICEWS05-15, YAGO11k and Wikidata12k take 4.3 seconds, 25.9 seconds, 1.9 seconds and 4.1 seconds per epoch on a single GeForce RTX 2080 device, respectively.

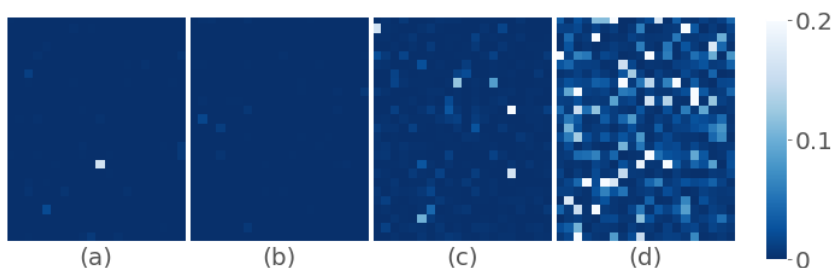


Figure 4: Visualization of the absolute difference vectors between  $\mathbf{r}_b$  and  $\mathbf{r}_e$  for relations *deadIn* and *isMarriedTo* (reshaped into  $25 \times 20$  matrices): (a)  $|Re(\mathbf{r}_b - \mathbf{r}_e)| / |Re(\mathbf{r}_b)|$  for relation *deadIn*; (b)  $|Im(\mathbf{r}_b - \mathbf{r}_e)| / |Im(\mathbf{r}_b)|$  for relation *deadIn*; (c)  $|Re(\mathbf{r}_b - \mathbf{r}_e)| / |Re(\mathbf{r}_b)|$  for relation *isMarriedTo*; (d)  $|Im(\mathbf{r}_b - \mathbf{r}_e)| / |Im(\mathbf{r}_b)|$  for relation *isMarriedTo*.

It is also noteworthy that representing each relation as a pair of dual complex embeddings is helpful to save training time on TKGs involving time intervals. Given a fact  $(s, r, o, [t_b, t_e])$ , some TKGE models, e.g., HyTE and ATiSE, discretize this fact into several quadruples involving continuous time points, i.e.,  $[(s, r, o, t_b), (s, r, o, t_b + 1), \dots, (s, r, o, t_e)]$ . When  $thre = 300$ , each fact lasts for averagely around 15 and 8 time steps in YAGO11k and Wikidata12k. In other words, such method that discretizes facts involving time intervals expands the sizes of both datasets by 15 and 8 times. In our model, we propose a more efficient method to handle time intervals by using two different quadruples,  $(s, r_b, o, t_b)$  and  $(s, r_e, o, t_b)$  to represent the beginning and the end of each fact. In this way, we only expand the sizes of datasets as less than twice as their original sizes.

For relations  $r$  in YAGO11k, we analyze the similarities between the embeddings  $\mathbf{r}_b$  and  $\mathbf{r}_e$ . As shown in Figure 4, for short-term relations, e.g., *deadIn*, the real parts of  $\mathbf{r}_b$  and  $\mathbf{r}_e$ , as well as their imaginary parts, have high similarities since  $r_b$  and  $r_e$  always happen at the same time and have the same semantics. By contrast, for long-term relations, e.g., *isMarriedTo*, the real parts of  $\mathbf{r}_b$  and  $\mathbf{r}_e$  show their semantic

similarities and the imaginary parts capture their temporal dissimilarities.

## 6 Conclusion

In this work, we introduce TeRo, a new TKGE model which represents entities or relations as single or dual complex embeddings and temporal changes as rotations of entity embeddings in the complex vector space. Our model is advantageous with its capability in modelling several key relation patterns and handling time annotations in various forms. Experimental results show that TeRo remarkably outperforms the existing state-of-the-art KGE models and TKGE models on link prediction over four well-established TKG datasets. Specially, we adopt two different time-division approaches for various datasets and investigate the effect of the time granularity on the performance of our model.

## Acknowledgements

This work is supported by the CLEOPATRA project (GA no. 812997), the German national funded BmBF project MLwin and the BOOST project.

## References

- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. *The semantic web*, pages 722–735.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250. AcM.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*, pages 2787–2795.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka Jr, and Tom M Mitchell. 2010. Toward an architecture for never-ending language learning. In *AAAI*, volume 5, page 3.
- Shib Sankar Dasgupta, Swayambhu Nath Ray, and Partha Talukdar. 2018. Hyte: Hyperplane-based temporally aware knowledge graph embedding. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2001–2011.
- Fredo Erxleben, Michael Günther, Markus Krötzsch, Julian Mendez, and Denny Vrandečić. 2014. Introducing wikidata to the linked data web. In *International Semantic Web Conference*, pages 50–65. Springer.
- Alberto García-Durán, Sebastijan Dumančić, and Mathias Niepert. 2018. Learning sequence encoders for temporal knowledge graph completion. *arXiv preprint arXiv:1809.03202*.
- Rishab Goel, Seyed Mehran Kazemi, Marcus Brubaker, and Pascal Poupart. 2020. Diachronic embedding for temporal knowledge graph completion. In *AAAI*.
- Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 687–696.
- Woojeong Jin, Changlin Zhang, Pedro Szekely, and Xiang Ren. 2019. Recurrent event network for reasoning over temporal knowledge graphs. *arXiv preprint arXiv:1904.05530*.
- Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. 2018. Canonical tensor decomposition for knowledge base completion. In *International Conference on Machine Learning*, pages 2869–2878.
- Jennifer Lautenschlager, Steve Shellman, and Michael Ward. 2015. Icews event aggregations.
- Julien Leblay and Melisachew Wudage Chekol. 2018. Deriving validity time in knowledge graph. In *Companion of the The Web Conference 2018 on The Web Conference 2018*, pages 1771–1776. International World Wide Web Conferences Steering Committee.

- Kalev Leetaru and Philip A Schrodt. 2013. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, volume 2, pages 1–49. Citeseer.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-ninth AAAI conference on artificial intelligence*.
- Yunpu Ma, Volker Tresp, and Erik A Daxberger. 2018. Embedding models for episodic knowledge graphs. *Journal of Web Semantics*, page 100490.
- Farzaneh Mahdisoltani, Joanna Biega, and Fabian M Suchanek. 2013. Yago3: A knowledge base from multilingual wikipedias. In *CIDR*.
- Mojtaba Nayyeri, Chengjin Xu, Yadollah Yaghoobzadeh, Hamed Shariat Yazdi, and Jens Lehmann. 2019. Toward understanding the effect of loss function on the performance of knowledge graph embedding.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In *ICML*, volume 11, pages 809–816.
- Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706. ACM.
- 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*.
- Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. 2017. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. *arXiv preprint arXiv:1606.06357*.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *AAAI*, pages 1112–1119. Citeseer.
- Chengjin Xu, Mojtaba Nayyeri, Fouad Alkhoury, Jens Lehmann, and Hamed Shariat Yazdi. 2019. Temporal knowledge graph embedding model based on additive time series decomposition. *arXiv preprint arXiv:1911.07893*.
- Chengjin Xu, Mojtaba Nayyeri, Yung-Yu Chen, and Jens Lehmann. 2020. Knowledge graph embeddings in geometric algebras. *arXiv preprint arXiv:2010.00989*.
- 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*.
- Shuai Zhang, Yi Tay, Lina Yao, and Qi Liu. 2019. Quaternion knowledge graph embeddings. In *Advances in Neural Information Processing Systems*, pages 2731–2741.