DensE: An Enhanced Non-commutative Representation for Knowledge Graph Embedding with Adaptive Semantic Hierarchy

Haonan Lu[*]
OPPO
Guangdong Mobile
Telecommunications Co., Ltd.
`luhaonan@oppo.com`

Hailin Hu
Huawei
Technologies Co., Ltd.
`huhailin2@huawei.com`

Xiaodong Lin
Department of
Management Science
and Information Systems,
Rutgers University
`lin@business.rutgers.edu`

Abstract

Capturing the composition patterns of relations is a vital task in knowledge graph completion. It also serves as a fundamental step towards multi-hop reasoning over learned knowledge. Previously, several rotation-based translational methods have been developed to model composite relations using the product of a series of complex-valued diagonal matrices. However, these methods tend to make several oversimplified assumptions on the composite relations, e.g., forcing them to be commutative, independent from entities and lacking semantic hierarchy. To systematically tackle these problems, we have developed a novel knowledge graph embedding method, named DensE, to provide an improved modeling scheme for the complex composition patterns of relations. In particular, our method decomposes each relation into an SO(3) group-based rotation operator and a scaling operator in the three dimensional (3-D) Euclidean space. This design principle leads to several advantages of our method: (1) For composite relations, the corresponding diagonal relation matrices can be non-commutative, reflecting a predominant scenario in real world applications; (2) Our model preserves the natural interaction between relational operations and entity embeddings; (3) The scaling operation provides the modeling power for the intrinsic semantic hierarchical structure of entities; (4) The enhanced expressiveness of DensE is achieved with high computational efficiency in terms of both parameter size and training time; and (5) Modeling entities in Euclidean space instead of quaternion space keeps the direct geometrical interpretations of relational patterns. Experimental results on multiple benchmark knowledge graphs show that DensE is comparable to the current state-of-the-art models for missing link prediction, especially on composite relations. In addition,
the interpretations generated by DensE also reveal how relations with distinct patterns (i.e., symmetry/anti-symmetry, inversion and composition) are modeled, which suggests several important directions of future studies.

1 Introduction

Knowledge graphs (KGs) are a vital component of a wide range of downstream applications, such as machine reasoning, information retrieval and knowledge-guided natural language processing (Ji et al., 2020; Zhang et al., 2019b; Yang et al., 2019; Lin et al., 2019). Especially, learning how to hop over a variety of concepts or instances stored in a knowledge graph represents a value path towards artificial general intelligence.

Knowledge graphs are defined as a collection of triplets. Each triplet, denoted by (h, r, t), indicates a relation r pointing from the head entity h to tail entity t. Currently, numerous research efforts have been devoted to developing knowledge graph embedding (KGE) methods. These methods aim to learn a set of low-dimensional representations of entities and relations (Ji et al., 2020; Nguyen et al., 2017), which is usually coupled with a score function to enable the knowledge graph completion process, i.e., predicting missing links between entities, for real-world KGs (Nickel et al., 2016; Lacroix et al., 2018; Bordes et al., 2013; Sun et al., 2019; Zhang et al., 2019a). Sometimes, neural networks can be inserted into the process (Dettmers et al., 2018; Schlichtkrull et al., 2018; Nathani et al., 2019), though this requires additional computation costs.

In principle, the desired KGE method should be able to accommodate various relation patterns and to learn representations that are approximately able to reason over the given patterns (expressiveness property of a KGE model (Sun et al., 2019; Wang et al., 2017)). For example, in a relation pattern such as symmetry (e.g., friend), asymmetry

---

(e.g., uncle), inversion relations (e.g., hypernym and hyponym) and compositional relations (e.g., my father's mother is my grandmother), these patterns should be hold in the vector space. While the former three patterns are readily covered by the current methods (Trouillon et al., 2016a; Sun et al., 2019), it still lacks an effective modeling strategy for composite relations due to the complexity of composition patterns. In particular, we find three predominant challenges in this modeling problem. First, the composition of relations can be non-commutative (e.g., my father's mother is my grandmother, while my mother's father is my grandfather, the reasoning result can be different by changing the orders of relations in a path of knowledge graph), which is opposite to the assumption of most KGE methods (Bordes et al., 2013; Sun et al., 2019). Second, the expressiveness of KGE methods is often limited by the counterintuitive lack of interaction between entity and relation embeddings (Yang et al., 2014). Last but not least, while the semantic hierarchy of entities in a knowledge graph is a ubiquitous property (Zhang et al., 2020) (e.g., a triplet in WordNet (Miller, 1995) $(palm, hypernym, tree)$ indicates "tree" is at a higher level than "palm" in the hierarchy), most methods do not pay attention to this and therefore fail to capture the semantic features at different semantic hierarchical levels.

In this work, to address these limitations, we develop a comprehensive solution to provide a highly expressive, efficient and interpretable modeling method for knowledge graph embedding. More specifically, we propose DensE (**D**istance-based **E**mbedding with **N**on-commutative Rotation and **S**caling in 3-D **E**uclidean Space), which decomposes the relation into an SO(3) group-based rotation operator and a scaling operator and in the 3-D Euclidean space. Intuitively, non-Abelian group (here we use SO(3) rotation group) is applied to introduce non-commutative nature to our model, and the scaling operation offers another important dimension to accommodate each triplet in the Euclidean space, which is barely explored in previous research. Our main contributions are summarized as the following:

(1) By integrating infinite non-Abelian group-based relational rotation and scaling operations in the 3-D Euclidean space within a unified framework, we effectively accommodate various relation patterns including non-commutative compositions,

semantic hierarchy, as well as interactions between entities and relations;

(2) Extensive experiments show that DensE achieves comparable to the current state-of-the-art models in link prediction with high computational efficiency, offering a useful tool for knowledge graph completion;

(3) We systematically consider three important scenarios of composition patterns that shall be considered by KGE methods (**Section 3**). Then, we show our method can provide an up-to-date most comprehensive while straightforward geometric interpretation for the modeling process of each relation type in the 3-D Euclidean space.

2 Related Work

In this section, we will discuss two different categories of KGE methods, especially how they evolve in terms of model expressiveness and interpretability.

2.1 Transnational Distance Model

Transnational distance models, represented by TransE (Bordes et al., 2013) and RotatE (Sun et al., 2019) use Eucleadian distance as the score function. In particular, the embedding of relations and entity are fit so that the tail entity can be obtained by from the head entity using the operation defined by the relation. Given this intuition, these methods usually reflect some particular geometric interpretations. However, within the various relation patterns, their modeling capacity for composite relations (i.e., a relation path composed of a series of relations) tend to be insufficient because most methods assume a commutative pattern on the relation path and do not consider entity information in inferring composition patterns.

Specifically, TransE models each relation as a pure translational transformation, so it assumes a fixed addition composition pattern between relations, i.e., $r_3 = r_1 + r_2$, which is commutative and irrelevant to entity embeddings. RotatE made significant progress by modeling relations as rotational operator (rotation matrix) in 2-D Euclidean space. When modeling a relation path composed of multiple relations, RotatE uses Hadamard product to combine the rotation matrices of the relations on the path, i.e., $r_3 = r_1 \circ r_2$. In this model, all relations in the composite relation have the same rotation axis. Thus, the compositions in RotatE are also mandatorily commutative. Also, interactions
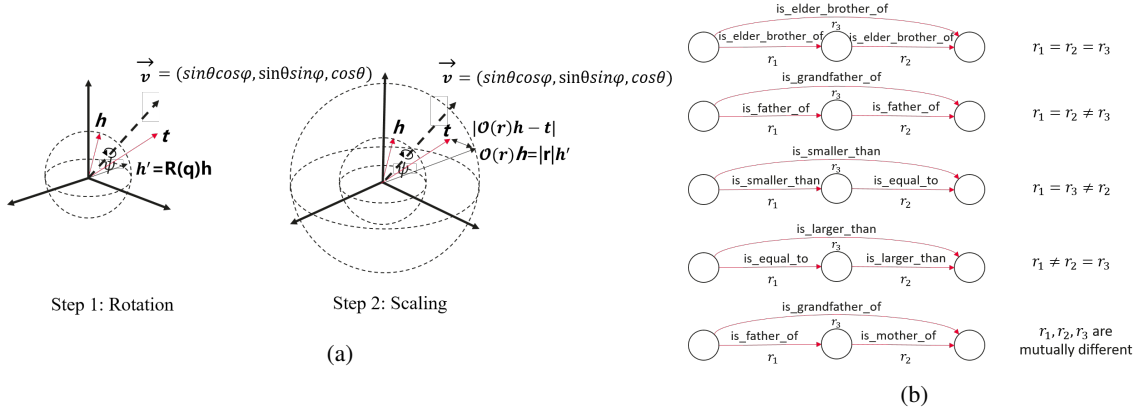
Figure 1: (a) DensE decomposes a relation into a rotation operator and a scaling operator on the head entity h in 3-D Euclidean space. (b) Examples of composition patterns.

between relation and entity embeddings are precluded as the rotation axis is always perpendicular to entity embeddings.

Following the effort of RotatE, several methods have been proposed to enhance the expressiveness of rotation-based translation KGE model. For instance, (Yang et al., 2020) proposes a group-theoretic analysis for KGE methods. Their method, named NagE, represents a preliminary attempt in applying non-Abelian group in modeling relational rotations. RotatE3D (Gao et al., 2020), on the other hand, extends the rotation of RotatE into the 3-D Euclidean space. However, although these methods represent certain conceptual advances, their empirical results show limited performance advance over previous methods, probably challenged by the fitting power of pure rotation-based operations and lacking of ability to model semantic hierarchies in knowledge graphs.

2.2 Semantic Matching Model

In constrast to translational distance model, methods in this category evaluate the matching of latent representation of relations and entities using bilinear model, e.g., RESCAL (Nickel et al., 2011), DistMult (Yang et al., 2014), and ComplEX (Trouillon et al., 2016b). Recently, as a generalization of DistMult and ComplEX, QuatE proposes a transformation on the entity representations by quaternion multiplication with the relation representation (Zhang et al., 2019a), leading to a significant advance in expressiveness. For composite relations, this method does not assume any fixed composition pattern and preserve the non-commutative nature to some extent. However, QuatE requires normalization of relation to unit quaternion, indicating it

is incapable of integrating scale information. In addition, as both entities and relations are embedded in quaternion hyperplanes, QuatE cannot provide a straightforward geometric interpretation in the space, which hinders the understanding of the learned embeddings.

2.3 Methodological Advance of DensE

From the perspective of score function, DensE also belongs to translation distance model. In contrast to the previous works, our model leverages both rotation and scaling operations for relation modeling. The key idea is that we can transform any non-zero vector in the Euclidean space to another arbitrary vector through decoupled rotation and scaling transformations. In addition, our model provides a clear geometric picture to demonstrate the transformation of entity representation in various relation composition patterns.

Conceptually, some modules of our method is also related to other recent KGE methods. For instance, a recent work HAKE (Zhang et al., 2020) has explored the combination of rotation and scaling operations, which are used to model entities at same and different levels of hierarchy, respectively. In HAKE, rotation is defined following the protocol in RotatE, which leverages U(1) Abelian group and thus incapable of handling non-commutative relations. In addition, its rotation axis is vertical to the 2-D representation space of entities, again omitting the entity-relation interaction as in RotatE. On the other hand, while using the non-Abelian group in KGE model has been explored by (Xu and Li, 2019), we argue this model has lower expressiveness than us in principle since it uses finite (non-)Abelian group (using non-Abelian group is

optional) in 2-D space while we consider an infinite non-Abelian group in 3-D space. Also, its rotation operation is combined with reflection, which constitutes a special case of our scaling operation (i.e., scaling with a factor of -1).

In contrast with QuatE that models both entities and relations in the quaternion space and does transformation using quaternion multiplication, DensE is based on 3-D Euclidean space rather than the space of quaternions. The continuous rotation transformation in the n-D Euclidean space ($n > 2$) is modeled by a special orthogonal group (SO(n) group). Compared with the vanilla U(1) abelian group based KGE models (e.g. RotatE/HAKE) that perform rotation transformation in the 2-D Euclidean space, continuous rotation transformation in the 3-D space modeled by SO(3) group is the minimum non-abelian extension with geometric interpretability. Composite relations (relation paths, details can be found in **Section 3**) are usually modeled by the product/summation of relation matrices. The violation of the commutative law of multiplication in the non-abelian case makes modeling the complex composition patterns of relations possible (non-commutative). The quaternion system is related to the SO(3) group and corresponds to the rotation transformation in 3-D Euclidean space, it provides a mathematical way to model continuous rotation transformation in 3D space. In our model, to guarantee geometric interpretability, entities are represented by 3D vectors, relations are modeled by quaternions that perform rotation and scaling transformation in 3D space. The rigidness of the quaternion system corresponds to the mathematical properties of SO(3) group theory since the rotation transformation in 3D space must satisfy several constraints such as: non-commutative (non-abelian nature), orthogonality, invertibility, etc. QuatE also studied that increasing spatial dimensions such as to Octonion does not increase performance compared to modeling relation and entities in the space of quaternion. The reason behind this is that the octonion system is more rigid than the quaternion system, the associative law of multiplication is also violated. However, this property is even harmful to modeling patterns of relations in a KG since there's no relation pattern in a real-world KG that needs to be modeled by violating the associative law of multiplication.

Therefore, we argue that among these concurrent works, DensE is the most comprehensive solution with the geometric interpretability that accounts for all the three *desiderata* for modeling composite relations, i.e., covering the non-commutative relations, preserving interaction between entity and relation, and capturing the entities' semantic hierarchy.

3 Problem Formulation

We denote a directed knowledge graph as $G(\mathcal{E}, \mathcal{R}, \mathcal{F})$, where $\mathcal{E}$, $\mathcal{R}$ and $\mathcal{F}$ are sets of entities, relations and facts, respectively. A fact stored in a KG can be expressed as a triplet $(h, r, t) \in \mathcal{F}$, where $h, t \in \mathcal{E}$ and $r \in \mathcal{R}$. Herein, we focus on the knowledge completion task, which aims to predict missing links based on the observed facts. To fulfill this goal, a score function is used to measure the plausibility of proposed fact candidates, and the goal of model optimization is to give higher scores to true triplets (h, r, t) than the false triplets $(h, r, \bar{t})$ or $(\bar{h}, r, t)$, where $\bar{t}$ and $\bar{h}$ are randomly sampled tail and head entities, respectively. Mathematically, the entity and relation embeddings are usually represented by tensors, and the score function can thus be written into the form of $f_{\mathbf{r}}(\mathbf{h}, \mathbf{t})$.

In principle, KGE models should be designed to accommodate various relation patterns existing in real world KGs, such as symmetry, anti-symmetry, inversion and composition, which are formally defined as follows.

Let x, y, z be the entities in a given KG, and $r(\cdot, \cdot)$ maps the relation between the two entities, we have:

Definition 1. A relation r is **symmetric** if $\forall x, y$,

$$r(x, y) \Rightarrow r(y, x). \tag{1}$$

On the other hand, a relation is said to be **anti-symmetric** if $\forall x, y$,

$$r(x, y) \Rightarrow \neg r(y, x). \tag{2}$$

$Friend$ is a typical example of symmetric relation, which means if we know x is friend of y, we can infer y is also friend of x. $Filiation$ is an example of anti-symmetric relation.

Definition 2. Relation r_1 is **inverse** to relation r_2 if $\forall x, y$

$$r_1(x, y) \Rightarrow r_2(y, x) \tag{3}$$

For instance, has_part and $part\_of$ fit into the scope of inverse relations, which means if we know x is a part of y, we can infer that y has part x. Note

that both **symmetric/antisymmetric** and **inverse** relation patterns can be inferred in one hop, so they are also called **atomic** relation.

In contrast to the above atomic relation patterns (inferable within one hop), the complex composition patterns pose a particular challenge to modeling, as discussed below.

Definition 3. Relation r_3 is **composed** of relation r_1 and relation r_2 if $\forall x, y, z$

$$r_1(x, y) \Lambda r_2(y, z) \Rightarrow r_3(x, z) \qquad (4)$$

Here r_3 is also referred to as a **composite relation** and possesses certain **composition pattern**. In particular, our model design takes the following properties into account:

Property 1. The two relations in the composition are not always **commutative**. For example, given $r_1 = is\_father\_of$, $r_2 = is\_mother\_of$, based on the **Definition 3**, we will get $r_3 = is\_grandmother\_of$. However, when we change the order, i.e., $r_1 = is\_mother\_of$, $r_2 = is\_father\_of$, we will get $r_3 = is\_grandfather\_of$. Recent KGE methods usually model composite relations (relation paths) by the product (e.g., QuatE, HAKE) /summation (e.g., TransE) of relation matrices. However, it is non-trivial to model composition relation patterns since the product/summation of diagonal real-valued/complex-valued matrices is usually commutative and hence invariant with the order of relations. For instance, ComplEx (Trouillon et al., 2016a) models relation path mentioned above $r_1(x, y) \Lambda r_2(y, z)$ by using the product of two complex-valued diagonal matrices: $R_{r_2} R_{r_1}$. However, the product of relation matrices in the diagonalized framework are commutative since that $R_{r_2} R_{r_1} = R_{r_1} R_{r_2}$.

Property 2. The composition patterns are not always inferable by the relations alone. For example, given that y is x's younger sister and z is y's elder brother, we can not answer whether z is elder or younger than x from the given information. Actually, to answer this question, we need to know more about $x/y/z$ from their own attributes and their other relationships.

Property 3. In a composition, the relations involved are not necessarily different. Given the two-hop example above, besides the situation that r_1, r_2 and r_3 are mutually different, there are also four different cases that satisfy the definition of composition, i.e., $r_1 = r_2 = r_3$, $r_1 = r_2 \neq r_3$, $r_1 = r_3 \neq r_2$ and $r_1 \neq r_2 = r_3$ (Figure 1(b)).

4 Method

In this section, we will first discuss the limitation of previous method such as RotatE, which is based on the 2-D Euclidean space. Then we will introduce each module of our method. In particular, we model a relation by a combination of an SO(3) group-based rotation (introducing the non-commutative nature) and a scaling operation (introducing the semantic hierarchy).

4.1 Limitations of Modeling Relational Rotation in the 2-D Euclidean Space

The motivation of RotatE is from Euler's identity $e^{i\theta} = \cos\theta + i\sin\theta$, which applies rotation in the 2D complex plane by using a unitary complex number. The RotatE model maps the entities and relations to the complex vector space and defines each relation as a rotation operator that transforms the source entity to the target entity. However, as shown in the Figure 2, composite relations are assumed to be commutative. Changing the order of relational rotation of r_1 and r_2 gives the same composition r_3. Also, the unit rotation transformation makes it difficult to model the semantic hierarchy which is a ubiquitous property in knowledge graphs. The rotation axis (perpendicular to the paper) of 2-D rotation transformation is orthogonal to entity embeddings, which hinders the method to model interactions between relational operations and entity embeddings.

$$r_1(\boldsymbol{h}, \boldsymbol{h'_1}) \wedge r_2(\boldsymbol{h'_1}, \boldsymbol{t}) \Rightarrow r_3(\boldsymbol{h}, \boldsymbol{t})$$
$$r_2(\boldsymbol{h}, \boldsymbol{h'_2}) \wedge r_1(\boldsymbol{h'_2}, \boldsymbol{t}) \Rightarrow r_3(\boldsymbol{h}, \boldsymbol{t})$$

Figure 2: RotatE models relations as a unit rotation operator in the 2-D Euclidean space.

Continuous rotation transformation in 3D space modeled by SO(3) group is the minimum non-abelian extension with geometric interpretability. By modeling relations and entities as rotation operators and vectors in 3D space, transformation in the 3-D Euclidean space can be either non-commutative or commutative. And since the ro-

tation axis of 3-D transformation is not enforced to be perpendicular to the vectors, interactions between relations and entities can also be considered. A simple example for how rotation in the 3-D Euclidean space can model non-commutative relations is shown in **Supplementary Note 1**. We will formally introduce the mathematical method for modeling rotations in 3D space in **Section 4.2**. In addition, to model the semantic hierarchies of knowledge graphs, the rotation operation is then followed by a scaling transformation. The modulus parts of 3D vectors aims to model the entities in a KG at different levels of the semantic hierarchy. The detail of integrating the scaling transformation in our model is discussed in **Section 4.3**.

4.2 Modeling Relational Rotation Using SO(3) Rotation Group

One of the ways to model a rotation operation in the 3-D space is called axis-angle representation, which parameterizes a rotation by two quantities: 1) A unit vector $\overrightarrow{v}$ indicating the direction of the axis of rotation, i.e., $\overrightarrow{v} = (v_x, v_y, v_z) = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$, where $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$; and 2) An angle ψ describing the magnitude of the rotation about the rotation axis, where $\psi \in [0, 2\pi)$. Given an entity vector $\overrightarrow{w}$ in the 3-D space with the coordinate (x, y, z), its rotation about axis $\overrightarrow{v}$ with an angle of ψ can be modeled using the SO(3) group theory (Figure 1(a), Step 1). More specifically, we can use a unit quaternion to encode the rotation using three degrees of freedom (i.e., θ, ϕ and ψ). Actually, it can be viewed as a group structure on a 3-sphere (i.e., S3) which gives the group Spin(3). Note that this group structure is isomorphic to SU(2) group and also to the universal cover of SO(3) group. Formally, the unit quaternion $\mathbf{q}$ to model a rotation through an angle of ψ around the aforementioned axis $\overrightarrow{v}$ can be derived using an extension of Euler's formula:

$$\mathbf{q} = e^{\frac{\psi}{2}(v_x\mathbf{i}+v_y\mathbf{j}+v_z\mathbf{k})} = \cos\frac{\psi}{2} + \sin\frac{\psi}{2}*(v_x\mathbf{i}+v_y\mathbf{j}+v_z\mathbf{k}) \quad (5)$$

where $\mathbf{i}, \mathbf{j}, \mathbf{k}$ are imaginary units of the quaternion representation, which satisfies the condition $\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1$. Unlike real/complex numbers, the multiplication of quaternions (Hamilton product) is sensitive to the orders as we have: $\mathbf{ij} = \mathbf{k}, \mathbf{ji} = -\mathbf{k}, \mathbf{jk} = \mathbf{i}, \mathbf{kj} = -\mathbf{i}, \mathbf{ki} = \mathbf{j}, \mathbf{ik} = -\mathbf{j}$. For $\mathbf{Q}_1 = a_1 + b_1\mathbf{i} + c_1\mathbf{j} + d_1\mathbf{k}$ and $\mathbf{Q}_2 = a_2 + b_2\mathbf{i} + c_2\mathbf{j} + d_2\mathbf{k}$, their Hamilton product is:

$$\begin{aligned} \mathbf{Q}_1 \otimes \mathbf{Q}_2 = \ & a_1a_2 - b_1b_2 - c_1c_2 - d_1d_2 \\ & +(a_1b_2 + b_1a_2 + c_1d_2 - d_1c_2)\mathbf{i} \\ & +(a_1c_2 - b_1d_2 + c_1a_2 + d_1b_2)\mathbf{j} \\ & +(a_1d_2 + b_1c_2 - c_1b_2 + d_1a_2)\mathbf{k} \end{aligned} \quad (6)$$

A 3-D Euclidean vector $\overrightarrow{w}$ with the coordinate (x, y, z) can be expressed as a pure quaternion (meaning the real part of quaternion is zero), i.e., $\mathbf{W} = x\mathbf{i} + y\mathbf{j} + z\mathbf{k}$, giving the following theorem (Jia, 2019): **Theorem 1** Given a 3-D Euclidean vector $\overrightarrow{w}$ and its counterpart in the quaternion space $\mathbf{W}$, the desired rotation axis $\overrightarrow{v}$, the magnitude of the rotation ψ, the destination coordinate of the vector after the rotation, i.e., $\mathbf{W}' = x'\mathbf{i} + y'\mathbf{j} + z'\mathbf{k}$, can be calculated by the Hamilton product of quaternions:

$$\mathbf{W}' = \mathbf{q}\mathbf{W}\mathbf{q}^{-1} \quad (7)$$

where $\mathbf{q}^{-1}$ is the inverse of $\mathbf{q}$, i.e., $\mathbf{q}^{-1} = e^{-\frac{\psi}{2}(v_x\mathbf{i}+v_y\mathbf{j}+v_z\mathbf{k})} = \cos\frac{\psi}{2} - \sin\frac{\psi}{2}*(v_x\mathbf{i} + v_y\mathbf{j} + v_z\mathbf{k})$.

The form of Eq.7 and a factor of $\frac{1}{2}$ for the angle ψ in Eq.5 indicate that there is a $2:1$ homomorphism from quaternions of unit norm to SO(3). Considering each 3-D Euclidean vector can also be expressed as a pure quaternion, we can now represent the rotation using a matrix $\mathbf{R}(\mathbf{q})$ by expanding Eq.7 and letting $C = \cos\psi$ and $S = \sin\psi$:

$$\overrightarrow{w}' = \mathbf{R}(\mathbf{q})\overrightarrow{w} =$$

$$\begin{bmatrix} C + v_x^2(1-C) & v_xv_y(1-C)+v_zS & v_xv_z(1-C)-v_yS \\ v_xv_y(1-C)-v_zS & C + v_y^2(1-C) & v_yv_z(1-C)+v_xS \\ v_xv_z(1-C)+v_yS & v_yv_z(1-C)-v_xS & C + v_z^2(1-C) \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} \quad (8)$$

In our framework, two rotations can be combined into one equivalent rotation operation (this is also consistent with the closure property of group theory). In other words, we can define $\mathbf{q} = \mathbf{q}_2\mathbf{q}_1$, where $\mathbf{q}$ corresponds to the rotation $\mathbf{q}_1$ followed by the rotation $\mathbf{q}_2$. Therefore, a series of rotations can be **composed** together and then applied as a single rotation. Note that quaternion multiplication is not commutative unless $\mathbf{q}_1$ and $\mathbf{q}_2$ share the same rotation axes (i.e., $\overrightarrow{v}_1 = \overrightarrow{v}_2$), which can be seen from Eq.6. This makes it possible to model both commutative and non-commutative relation patterns.

4.3 Integrating the Scaling Operation

In a knowledge graph, different entities may have different level of semantic hierarchy given a particular relation. For example, in WN18RR, *trade*

is a *hypernym* of *transaction*, and *man* is recorded to *has_part* to be *arm*. In these cases, the head entity and tail entity show different abstraction levels or showing inclusion relationships. Intuitively, we argue that the difference of semantic hierarchy can be reflected by the scale of entity, as the entities possessing same level of abstraction tend to be achieved through rotation operations.

To define this intuition mathematically, we first obtain of norm of quaternions. Following Eq.5 and letting $\mathbf{q}$ to be the unit quaternion, an arbitrary quaternion with non-unit norm can be written as: $\mathbf{Q} = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k} = |Q|\mathbf{q}$, with the norm given by

$$|Q| = \sqrt{a^2 + b^2 + c^2 + d^2} \qquad (9)$$

where

$$a = |Q|\cos\frac{\psi}{2},$$
$$b = |Q|\sin\frac{\psi}{2}\sin\theta\cos\phi,$$
$$c = |Q|\sin\frac{\psi}{2}\sin\theta\sin\phi, \qquad (10)$$
$$d = |Q|\sin\frac{\psi}{2}\cos\theta.$$

By multiplying a scalar $|Q|$ in the Eq.8, we can further introduce length as another degree of freedom to better match the ground-truth tail embedding vector (Figure 1(a), Step 2). Formally, we have:

$$\overrightarrow{\boldsymbol{w}}' = |Q|\mathbf{R}(\mathbf{q})\overrightarrow{\boldsymbol{w}} = \mathcal{O}(\mathbf{Q})\overrightarrow{\boldsymbol{w}}, \text{ where } \mathbf{Q} \in \mathbb{H}, \overrightarrow{\boldsymbol{w}}, \overrightarrow{\boldsymbol{w}}' \in \mathbb{R}^3.$$
$$(11)$$

where $\mathcal{O}(\mathbf{Q}) = |Q|\mathbf{R}(\mathbf{q})$ is the combined operator of rotation and scaling transformations, $\mathbb{H}$ denotes the quaternion algebra, and $\mathbb{R}^3$ represents the 3-D Euclidean algebra. Here we call $|Q|$ the *scaling factor*. Therefore, we now have a uniform framework with interpretable geometric meaning, i.e., $(|Q|, \theta, \phi, \psi)$ to describe the transformation corresponding to a specific relation type. We can also define the reverse operation $\mathcal{O}(\mathbf{Q}^{-1}) = |Q|^{-1}\mathbf{R}(\mathbf{q}^{-1})$, which describes the reverse process: rotate a vector about the axis $\overrightarrow{\boldsymbol{v}}$ with angle $-\psi$ (from another direction), and then scale the vector with a factor of $|Q|^{-1}$. Combining the Eq.8 and Eq.11, we can always find a operator $\mathcal{O}(\mathbf{Q}_3) = \mathcal{O}(\mathbf{Q}_2)\mathcal{O}(\mathbf{Q}_1)$, which corresponds to the application of $\mathcal{O}(\mathbf{Q}_1)$ followed by the application of $\mathcal{O}(\mathbf{Q}_2)$, where we have $|Q_3| = |Q_1| * |Q_2|$ and $\mathbf{R}(\mathbf{q}_3) = \mathbf{R}(\mathbf{q}_2)\mathbf{R}(\mathbf{q}_1)$.

4.4 Score Function and Optimization

A score function aims to correctly measure the plausibility of a triple of interest. Formally, as a distance-based model, our scoring function is defined as

$$f_{\mathbf{r}}(\mathbf{h}, \mathbf{t}) = -\frac{1}{2}(|\mathcal{O}(\mathbf{r})\mathbf{h} - \mathbf{t}| + |\mathcal{O}(\mathbf{r}^{-1})\mathbf{t} - \mathbf{h}|). \quad (12)$$

Here, $|\cdot|$ denotes the Euclidean distance and $\mathcal{O}(\cdot)$ stands for the transformation conducted on each element of the entity embeddings. That is to say, for the i-th embedding unit of $\mathbf{h}$, the optimization target is to minimize the Euclidean distance between $\boldsymbol{t}_i$ and $\mathcal{O}(\mathbf{r}_i)\boldsymbol{h}_i$, as well as the Euclidean distance between $\boldsymbol{h}_i$ and $\mathcal{O}(\mathbf{r}_i^{-1})\boldsymbol{t}_i$, where $\mathbf{r}_i \in \mathbb{H}, \boldsymbol{h}_i, \boldsymbol{t}_i \in \mathbb{R}^3$. $\mathbb{H}$ and $\mathbb{R}^3$ stand for the quaterrion and 3-D Euclidean algebra, respectively. The arrow of $\boldsymbol{h}_i$ and $\boldsymbol{t}_j$ are omitted for clarity. To properly train the model parameters, here we use a loss function similar to the self-adversarial negative sampling loss proposed in (Sun et al., 2019):

$$L = -\log\sigma(\gamma + f_{\mathbf{r}}(\mathbf{h}, \mathbf{t}))$$
$$-\sum_{j=1}^{n} p(\bar{h}_j, r, \bar{t}_j)\log\sigma(-(\gamma + f_{\mathbf{r}}(\bar{\mathbf{h}}^{(j)}, \bar{\mathbf{t}}^{(j)}))),$$
$$(13)$$

where γ is a fixed margin, n is the number of negative sampling size, $(\bar{h}_j, r, \bar{t}_j)$ is the j-th negative triplet of the fact (h, r, t), and σ is the sigmoid function. $\bar{\mathbf{h}}^{(j)}$ and $\bar{\mathbf{t}}^{(j)}$ are the embeddings corresponding to the negative triplet $(\bar{h}_j, r, \bar{t}_j)$. $p(\bar{h}_j, r, \bar{t}_j)$ is the weight of the negative sample, which gives the higher scored negative samples with larger weight during training. The details about self-adversarial negative sampling technique can be found in (Sun et al., 2019).

5 Experimental Settings

Datasets and evaluation metrics The experiments are conducted mainly on three commonly used benchmark datasets, including WN18RR, FB15k-237 and YAGO3-10. As pointed out by (Toutanova and Chen, 2015a; Dettmers et al., 2018), WN18 and FB15k suffer from the test set leakage problem. One can predict missing links and attain the state-of-the-art results even using a simple rule-based model. To avoid this issue, two much more challenging datasets (WN18RR (Dettmers et al., 2018) and FB15k-237 (Toutanova and Chen,

Table 1: Statistics of datasets used in this study.

Dataset	# Entities	# Relations	#Training	#Validation	#Test
WN18RR	40943	11	86835	3034	3134
FB15k-237	14541	237	272115	17535	20466
YAGO3-10	123182	37	1079040	5000	5000

Table 2: Performance comparison on benchmark datasets. Best results are labeled in bold and the second best are underlined. The reporting scheme generally follows that in (Ruffinelli et al., 2019). *First* indicates the originally reported performance of each method. *Enhanced* records the improved performance with tuned training techniques and hyperparameters by (Ruffinelli et al., 2019). *Recent* shows the best results of more selected recent models. *Adv+Recip* reports the model performance using the same training scheme of DensE, i.e., using self-adversarial negative sampling and reciprocal learning. *Ours* reports the performance of DensE as well as its ablation counterparts.

	Model	WN18RR		FB15K-237		YAGO3-10	
		MRR	H@10	MRR	H@10	MRR	H@10
First	RESCAL (Wang et al., 2018)	0.420	0.447	0.270	0.427	-	-
	TransE (Nguyen et al., 2017)	0.226	0.501	0.294	0.465	-	-
	DistMult (Dettmers et al., 2018)	0.430	0.490	0.241	0.419	0.340	0.540
	ComplEx (Dettmers et al., 2018)	0.440	0.510	0.247	0.428	0.360	0.550
Enhanced	RESCAL (Ruffinelli et al., 2019)	0.467	0.517	<u>0.357</u>	0.541	-	-
	TransE (Ruffinelli et al., 2019)	0.228	0.520	0.313	0.497	-	-
	DistMult (Ruffinelli et al., 2019)	0.452	0.531	0.343	0.531	-	-
	ComplEx (Ruffinelli et al., 2019)	0.475	0.547	0.348	0.536	-	-
Recent	RotatE (Sun et al., 2019)	0.476	0.571	0.338	0.533	0.495	0.670
	NagE (Yang et al., 2020)	0.477	0.574	0.340	0.530	-	-
	QuatE (Jia, 2019)	0.481	0.564	0.311	0.495	-	-
	D4-STE (Xu and Li, 2019)	0.480	0.536	0.320	0.502	0.472	0.643
	TuckER (Balažević et al., 2019)	0.470	0.526	**0.358**	**0.544**	-	-
	Rotate3D (Gao et al., 2020)	0.489	0.579	0.347	<u>0.543</u>	-	-
	HAKE[1] (Zhang et al., 2020)	**0.497**	<u>0.584</u>	0.336	0.533	0.522	<u>0.693</u>
	HAKE[2] (Zhang et al., 2020)	**0.497**	0.582	0.346	0.542	**0.545**	**0.694**
Adv+Recip	TransE	0.230	0.535	0.330	0.525	0.460	0.661
	DistMult	0.444	0.533	0.316	0.497	0.427	0.627
	ComplEx	0.475	0.559	0.334	0.525	0.510	0.681
	RotatE	0.478	0.567	0.337	0.531	0.497	0.676
Ours	DensE	<u>0.492</u>	**0.586**	0.351	**0.544**	<u>0.541</u>	0.678
	-Scaling	0.475	0.562	0.335	0.527	0.486	0.642
	-Reciprocal	0.487	0.572	0.343	0.527	0.530	0.667
	-Adv	0.486	0.572	0.306	0.481	0.452	0.642

2015b)) were released. WN18RR comes from WordNet (Miller, 1995), compared with the previous version WN18, it removes inverse relations to provide a more realistic KGE method benchmark. Similarly, the FB15k-237 dataset is also extracted from the original Freebase dataset FB15K (Bordes et al., 2013) by removing inverse relations.

In addition, we also use the YAGO3-10 (Mahdisoltani et al., 2013) dataset, which consists of a large collection of triplets from multilingual Wikipedia. These three datasets aim to assess the model performance on composition patterns. The main relation patterns of them are symmetry/anti-symmetry and composition. The basic statistics of the datasets

are provided in Table 1. Here, we report mean reciprocal rank (MRR) and Hits at 10 (H@10) for evaluation (the higher, the better), which is consistent with (Ruffinelli et al., 2019). Other performance metrics are provided in the Supplementary Material.

Baselines We mainly compare DensE with top-performing baseline models for KG link prediction, including both translational model and semantic matching model. As the early implementation of the baseline models may lack thorough configuration tuning or advanced learning techniques, direct compassion with these performances (denoted as *First*) may be biased to later methods. Therefore, for early models such as RESCALL (Nickel et al., 2011), TransE (Bordes et al., 2013), DistMult (Yang et al., 2014) and ComplEx (Trouillon et al., 2016a), we also provide two improved versions, including *Enhanced*, which was obtained through a sophisticated hyperparameter tuning procedure by (Ruffinelli et al., 2019) and *Adv+Recip*, which is obtained by us using the same self-adversarial negative sampling and reciprocal learning as for DensE. These two improved versions are provided to prompt the fairness of the comparison.

In addition to these methods, we also compare our model with more recently proposed KGE models (denoted as *Recent*), such as RotatE (Sun et al., 2019), QuatE (Jia, 2019), D4-STE (Xu and Li, 2019), TuckER (Balaževič et al., 2019), Rotate3D (Gao et al., 2020) and HAKE (Zhang et al., 2020). All these recent methods have included some advanced training techniques similar to (Ruffinelli et al., 2019).

Implementation details We use the Adam optimizer and tune the hyperparameters on the validation dataset. During training, we adopt a similar reciprocal learning approach as used in (Lacroix et al., 2018; Zhang et al., 2019a). Early stopping is applied based on the performance on the validation dataset every 1,000 steps. The ranges for hyperparameter grid search and the best hyperparameter settings are listed in **Supplementary Note 2**. All the parameters are randomly initialized from the interval
$[-\frac{1}{\sqrt{2k}}, \frac{1}{\sqrt{2k}}]$, where k is the embedding size.

6 Results and Analysis

6.1 Prediction Performance

We report the link prediction results on the three benchmark datasets in Table 2. On WN18RR, we show that DensE performs on par with HAKE and outperforms most other models on both the metrics, even after the baseline models are improved by hyper-parameter tuning or using advanced learning techniques. On FB15k-237, we show that the performance of DensE is superior to most of the baseline models, including RotatE and QuatE. While on this dataset we notice a particular good performance of a method called TuckER (Balaževič et al., 2019), this method also shows a significantly inferior performance on WN18RR, suggesting a potential drawback in generalizability. On YAGO3-10, DensE also shows a significant margin over RotatE and D4-STE (a KGE method based on 2D finite (non-)Abelian group), and also ComplEx when using the comprehensive metric MRR, which further demonstrates the superiority of DensE on various types of datasets. We also provide additional performance metrics, i.e., MR, MRR, Hits at 1 (H@1), Hits at 3 (H@3), and Hits at 10 (H@10) in Supplementary Note 3 (Supplementary Tables 2, 3, and 4).

Then we carefully compare the performance of DensE to two recent extension of RotatE models, namely RotatE3D (Gao et al., 2020) and HAKE (Zhang et al., 2020). We find that DensE performs better than Rotate3D in most cases, validating the contribution of the scaling operation. On the other hand, HAKE and DensE generally perform comparably. After dissecting into the training details, we find that different from DensE and most other translation distance model, HAKE pays more attention to model the hierarchical nature of knowledge graphs. **Firstly**, unlike the score function of DensE, which directly optimize the Euclidean distance between two vectors in the 3D space, HAKE decomposes the score function into two part: 1. The modulus-distance part corresponds to the hierarchy of the knowledge graph; 2. The phase-distance part corresponds to the rotation operation in the 2D space. HAKE leverages a task-specifically calibrated loss term by tuning the relative contributions of two terms in its score function manually to make it get better performance in the dataset with a clear hierarchical structure. Therefore, for a dataset like WN18RR with a majority of types of relations that link two entities at differ-
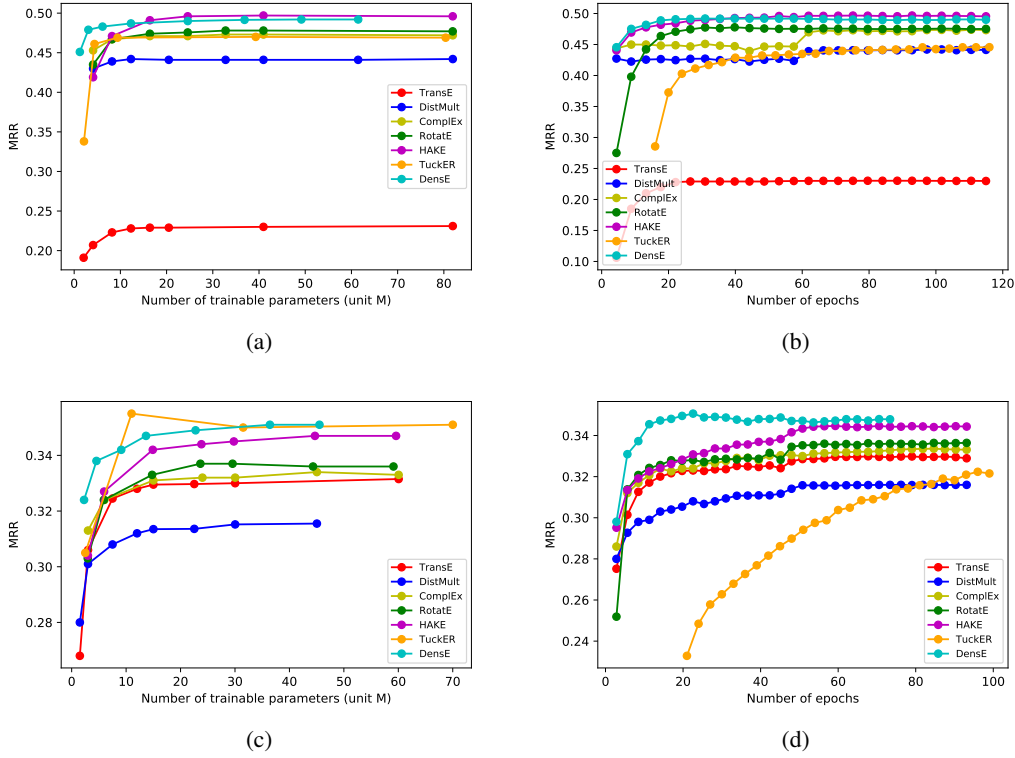
(a)

(b)

(c)

(d)

Figure 3: The effect of trainable parameter size and training epoch number on model performances. (a) and (b) shows the results for WN18RR dataset. The corresponding results for FB15K-237 is shown in (c) and (d). All the results are achieved using the same setting (*Adv+Recip*) as described above.

(a) $\psi_{has\_part} + \psi_{part\_of}$ (b) $|Q_{has\_part}| * |Q_{part\_of}|$ (c) $|Q_{has\_part}|$ (d) $|Q_{part\_of}|$

(e) $\psi(\mathcal{O}(\mathbf{r2})\mathcal{O}(\mathbf{r1})) - \psi(\mathcal{O}(\mathbf{r1}))$ (f) $\psi(\mathcal{O}(\mathbf{r2})\mathcal{O}(\mathbf{r1})) - \psi(\mathcal{O}(\mathbf{r3}))$ (g) $\theta(\mathcal{O}(\mathbf{r1})) - \theta(\mathbf{h})$
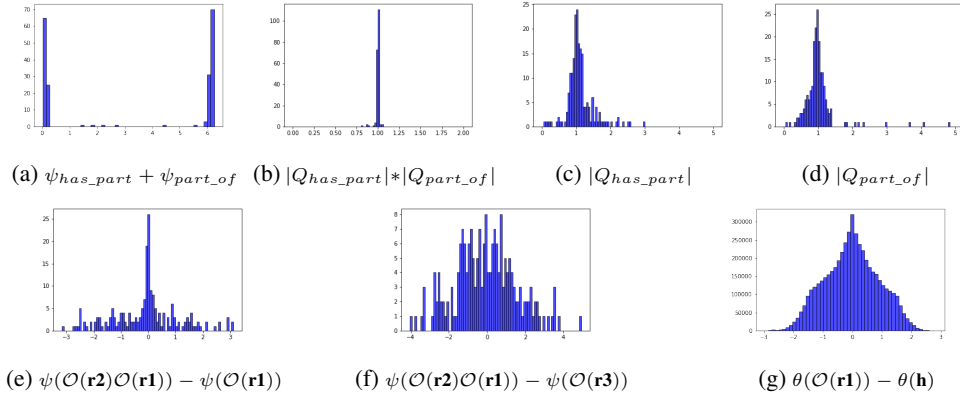
Figure 4: Geometric interpretation provided by DensE. Each histogram shows a distribution of each dimension of the learned embeddings. Angular parameters are in radian units. (a)-(d) A case study for inversion patterns (Other degrees of freedom can be found in the **Supplementary Note 5**). (e)-(f) A case study for composition patterns, reflecting the scenario of $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_1(h, t)$ and $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_3(h, t)$, respectively. $\psi(\cdot)$ denotes the rotation angle about the rotation axis of a relational operator. (g) Collinearity of entity and relation embedding. $\theta(\mathcal{O}(\mathbf{r1}))$ is the θ component of relation r_1. $\theta(\mathbf{h})$ is the θ value of the head entity h's embedding in the spherical coordinate system of the 3-D Euclidean space. All the head entities satisfying $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_1(h, t)$ are included for the analysis. The entities and relations involved can be found in the **Supplementary Note 5**.

ent levels of the hierarchy, HAKE can get better performance than DensE. The FB15k-237 dataset has more complex relation types (237 types of relation) and fewer entities (higher average degree of vertices) than WN18RR and YAGO3-10. The advantage of tuning relative contributions of two terms in the score function manually does not exist anymore, while the information of hierarchy can

be learning automatically in our DensE with adaptive semantic hierarchy. That's why we outperform HAKE in FB15k-237 dataset. **Secondly**, as mentioned in the HAKE, it has two versions of the score function to model the modulus-distance part. We report the results of two versions of HAKE in Table 2, labeled by HAKE[1] and HAKE[2], respectively. Version 1 has a clear and simple mathematical form that models rotation in the 2D space. Compared with version 1, a bias and re-scaling operation on relational embedding are introduced into the model in version 2 (An additional freedom to tune embeddings of relation, thus the element of the embedding of relation is a 3D vector). For the YAGO3-10 dataset which is more complicated than WN18RR and also has a clear semantic hierarchy property, DensE outperforms HAKE[1] and get comparable result with HAKE[2]. The above two techniques proposed by HAKE do improve the ability to model the semantic hierarchy and complement its lack of expressiveness in rotation operation. We believe these techniques will also be important tricks that can be used to improve performance in future studies (just like the self-adversarial negative sampling technique).

To confirm the source of performance gains, we conduct a further analysis that compares the MRR performance of DensE to RotatE on each relation type of WN18RR (Table 3). Besides the taxonomy mentioned in **Section 3**, relations in the WN18RR dataset can be also divided into two categories: (a) relations that link two entities in the same semantic hierarchy (e.g., "similar_to"); (b) relations that link two entities at different levels of the hierarchy (e.g., "has_part"). One can see that most of the relations that fall into category (b) are also overlap with composite relation patterns. Intriguingly, we notice a large performance increase in composite relations, as exemplified by $hypernym$, the most abundant composite relation in test data. We show that DensE improves MRR on this relation by as much as 3.3%. These results indicates particular advantages of DensE in modeling composition relation patterns and semantic hierarchies of knowledge graphs.

6.2 Ablation Study

To examine the effectiveness of each module in our model, we perform a series of ablation experiments (Table 2). On WN18RR dataset, the most significant performance decrease occurs when we cancel the scaling operation, i.e., only model the relation as rotations. This confirms the contribution from scaling to the whole model. On FB15K-237 and YAGO3-10 datasets, we also observe a large drop in performance when removing the scaling operation. Also, self-adversarial negative sampling (adv) shows significant contribution, indicating the necessity to incorporate proper training techniques. In **Supplementary Note 4**, we also compared DensE and RotatE models without self-adversarial negative sampling and confirmed the superiority of DensE in this setting. Note that HAKE also uses the self adversarial technique in training. However, the ablation results are not provided.

6.3 Computational Complexity

We show that compared with high-performance models such as RotatE and HAKE, DensE is generally more computationally efficient in terms of parameter number and training epochs. For TuckRE, although it can get better performance on FB15k237 with a relatively small model size, it needs much larger training epochs than other models on both WN18RR and FB15k237 datasets. We find that TuckRE needs roughly 500 epochs to converge to its best results. We plot the results of the first 100 epochs in Figure 3(b) and Figure 3(d) here. As is shown in Figure 3, when compared with other baseline models, DensE achieves significantly higher performance with the same parameter size or epoch number on both WN18RR and FB15K-237 dataset. When comparing the time efficiency, as different models have different training time per epoch, here we also report the training time of DensE and other baseline models. We show that although the more complex math formulation may cause longer training time, the resulting better-designed model can lead to a much faster convergence speed that significantly shortens the total training time under the same machine condition (Table 4). On the other hand, if we let models have similar performance, e.g., only make DensE reach the final performance of RotatE, we can see it only needs 10 epochs and uses 27% of RotatE's training time. These results indicate a clear advance of DensE in computation efficiency. We reason that this is mainly achieved by introducing a decoupled scaling operation, thus lowering the embedding dimension required in rotation-only modeling.

Table 3: MRR comparison on each relation type of WN18-RR. Performance increases are in parentheses.

Relation type	Relation Name	% in test data	RotatE	DensE
Atomic	derivationally_related_form	34%	0.947	0.955 (+0.008)
	also_see	1.8%	0.585	0.647 (+0.062)
	verb_group	1.3%	0.943	0.955 (+0.012)
	similar_to	0.2%	1	1 (+0)
Composite	hypernym	39.5%	0.148	0.181 (+0.033)
	instance_hypernym	4%	0.318	0.349 (+0.031)
	member_meronym	8.1%	0.232	0.249 (+0.017)
	synset_domain_topic_of	3.8%	0.341	0.412 (+0.071)
	has_part	5.5%	0.184	0.205 (+0.021)
	member_of_domain_usage	0.8%	0.318	0.326 (+0.008)
	member_of_domain_region	1%	0.2	0.407 (+0.207)

Table 4: The comparison of training time of each method on WN18RR. A lower total time results in a higher efficiency. To ensure a fair comparison, here we unify the hyperparateters of each method so that all the models have a similar parameter size around 36M.

Model	Training time per epoch (s)	# of epochs	Total training time (s)	MRR
DensE	92	21	1932	0.492
RotatE	83	40	3320	0.478
ComplEx	75	62	4650	0.475
DisMult	52	68	3536	0.444

7 Geometric Interpretation

7.1 Theoretical Analysis

In this section, we first discuss mathematically how DensE provides geometric interpretation of relation patterns including symmetry, antisymmetry, inversion and composition. According to the convention of axis-angle representation described in **Section 4.2**, all angle-related parameters (θ, ϕ, and ψ) are restricted to be in the half-closed intervals as we mentioned before. The positive direction of rotation is based on the right-handed coordinate system, the rotation angle with the minus value indicates a rotation opposite to the positive direction. To keep the values of angle to be within the above interval, we relocate angles outside intervals into the desired regions by leveraging the periodic property of the rotation system. To begin with, a relation r is **symmetric** in DensE if and only if each dimension of its embedding $\mathbf{r}_i$ satisfies $|\mathbf{r}_i| = 1$ and the rotation angle satistifies $\psi_{r_i} = 0$ or π. For **anti-symmetry** relation pattern, the embedding $\mathbf{r}_i$ satisfies $|\mathbf{r}_i| = 1$, but the rotation angle ψ_{r_i} should be neither 0 nor π. Also, two relations r_1 and r_2 are in **inverse** pattern, if and only if they sat-

isfy: $|\mathbf{r}_{1i}| * |\mathbf{r}_{2i}| = 1, \theta_{r_{1i}} = \theta_{r_{2i}}, \phi_{r_{1i}} = \phi_{r_{2i}}$ and $\psi_{r_{1i}} + \psi_{r_{2i}} = 2\pi$, meaning the embeddings of these two relations share the same rotation axes, but rotate in two opposite directions.

As discussed in **Section 4.2**, the **commutative** and **non-commutative composition** patterns can be naturally modeled by the guarantee of the property of group theory, which covers **Property 1** of composite relations. Also, following the intuition of **Property 2**, our model does not enforce a uniform mode of each element in relation representations. Instead, it learns to model the interaction between relations and entities as well as the ambiguity in composition pattern inference, leading to a disperse distribution in the relation embedding space. Last but not least, our model can smoothly deal with constraints posed by relation types in inferring composition patterns, as stated in **Property 3**. For instance, when modeling the pattern $r_1(x, y) \Lambda r_2(y, z) \Rightarrow r_2(x, z)$, the representation from RotatE tends to degenerate to a trivial case where the rotation angle of r_1 and r_2 both set to be 0 or 2π or $r_1 = 2\pi, r_2 = \pi$. In DensE, since the entity embeddings are not required to be perpendicular to the rotation axis, it can also place the

embedding of entity x to be collinear with the rotation axis of r_1. In this way, the rotation axis of r_1 and r_2 are not required to be the same, making the model to more expressive. In another example, as for the pattern $r_1(x,y)\Lambda r_1(y,z) \Rightarrow r_2(x,z)$, besides capturing the relationship of the two rotation angles (i.e., $\psi_{r_{2i}} = 2\psi_{r_{1i}}$), the scaling transformation offers an additional degree of freedom, where our model tends to give $|\mathbf{r}_{2i}| = |\mathbf{r}_{1i}|^2$. Again, we point out that these "rules" are not constant solutions, as the information entities will further guide the model to deviate from the statistical mode for better accommodation of each triplet. Other composition patterns presented in the **Property 3** can be analyzed in a similar way (see **Supplementary Note 5.3**).

7.2   Case Studies

Here, we show several examples to illustrate the geometric insight given by DensE, which basically reflects the geometric intuition discussed above.

We start our analysis with **inverse** relations, which comes from the original WN18 dataset (We have also confirmed the good prediction performance of DensE on the WN18 dataset in **Supplementary Note 6**). In Figure 4(a), we show the distribution of element-wise addition of embeddings from two **inverse** relations (has_part and $part\_of$) of ψ, one representative degree of freedom in modeling relational rotation (Other degrees of freedom can be found in the **Supplementary Note 5**). In this way, we can visualize how the two embeddings agree with each other. As the two relations are fully inferable by each other, we do observe a clear conjugation as expected in **Section 7.1**. This is also reflected in the representation of scaling from these two relations, where the element-wise products tend to be one (Figure 4(b)). Interestingly, we do observe that two complementary embeddings are learned these two relations (Figure 4(c)-(d)). In particular, the relation whose head entity has a higher semantic hierarchy (i.e., has_part) tend to show a scaling norm $|Q|$ larger than one, while the relation whose head entity has a lower semantic hierarchy (i.e., $part\_of$) generally shows a scaling norm $|Q|$ smaller than one. This clearly verifies the intuition of introducing the scaling operation do capture relation-specific semantic hierarchy of entities.

For **composition patterns**, we slightly change the experiment protocol, with each histogram show-

ing the element-wise difference between the embeddings of a composite relation and the embeddings calculated by multiplying each relation in the relation path. In a case from WN18RR, we demonstrate how DensE models a composition pattern for

$$r_1(h,h')\Lambda r_2(h',t) \Rightarrow r_1(h,t),$$

where

$$r_1 = derivationally\_related\_form,$$
$$r_2 = hypernym.$$

This is a typical case where the composite relation equals the first relation in the relation path. As shown in Figure 2(e), while most embedding dimensions still agree well between the actual composite relation and the calculated relation path, the distribution tends to disperse to a large range, indicating the existence of ambiguity and interaction between entities and relations. To further explore the ambiguity issue in the above case, we perform the same analysis on another small portion of triplets that actually give

$$r_1(h,h')\Lambda r_2(h',t) \Rightarrow r_3(h,t)$$

where

$$r_3 = synset\_domain\_topic\_of.$$

Interestingly, we also observe that part of embedding dimensions of $\mathcal{O}(\mathbf{r2})\mathcal{O}(\mathbf{r1})$ are aligned with $\mathcal{O}(\mathbf{r3})$ (embedding difference close to zero), demonstrating the flexibility of our model to capture potentially ambiguous relation compositions (Figure 4(f)). On the other hand, the model can also learn to put the rotation axis of r_1 collinear with the embedding of head entities h in the composition mode expressed as

$$r_1(h,h')\Lambda r_2(h',t) \Rightarrow r_1(h,t)$$

reflecting the interaction of entities and relations (Figure 4(g)). This example clearly demonstrates the interpretability of DensE in modeling complex composition patterns. The geometric patterns for other relation patterns can be found in **Supplementary Note 5**. Together with the properties discussed in **Section 3**, here we clearly demonstrate the pros and cons of the current rotation-based translational KGE method in modeling composition relation patterns.

8 Conclusion

In this work, we propose an effective method, named DensE, for knowledge graph embedding. DensE decomposes a relation operator into an SO(3) group-based rotation as well as a scaling transformation. Extensive experiments show that DensE possesses good performance in knowledge completion with high computational efficiency. Also, DensE provides a straightforward geometric interpretation for the relations, leading to meaningful insights for the future work for modeling complex relation patterns.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

A   The geometrical interpretation of non-commutative compositions

$$r_1(h, h_1') \wedge r_2(h_1', t) \Rightarrow r_3(h, t) \qquad r_2(h, h_2') \wedge r_1(h_2', t) \Rightarrow r_1(h, t)$$

Figure 5: A simple example for how rotation in the 3-D Euclidean space can model non-commutative relations

Left: A rotation about axis-z followed by a rotation about axis-x, the initial vector **h** is placed along axis-x. It can be seen that the final state is along axis-z, and two rotation operations are equivalent to one operation with the rotation axis to be about axis-y.

Right: The rotation operation sequence is reversed from the left figure. A rotation about axis-x is followed by a rotation about axis-z. The final state is then changed to be along axis-y. Since the initial vector is collinear with the first rotation axis, the two rotation operations are equal to the last rotation (rotation about axis-z).

B   Hyperparameters setting

The ranges of the hyperparameters for the grid search are set as follows: Embedding size $k \in \{100, 200, 500, 1000\}$ (In our model, each entity is represented with a matrix with a size of $3 \times k$, and each relation with a matrix with a size of $4 \times k$), batch size $b \in \{256, 512, 1024\}$, fixed margin $\gamma \in \{3.0, 6.0, 9.0, 12.0, 15.0, 24.0, 30.0\}$, negative sampling size $n \in \{256, 512, 1024\}$, self-adversarial sampling temperature $\alpha \in \{0.3, 0.5, 1.0\}$. The initial learning rate η is set to be 0.1, and it decays with a factor of $1/2$ if the training loss does not decrease in 1000 epochs. We list the best hyperparameters setting of DensE on the benchmark datasets in Supplementary Table 5.

C   Additional performance metrics

For a more complete comparison of each method, for each dataset we list MR, MRR, H@1, H@3, and H@10 in Supplementary Tables 6, 7, and 8.

D   Effect of self-adversarial negative sampling on DensE and RotatE

In the ablation study, we observe a significant contribution of the self-adversarial negative sampling technique on the prediction performance of FB15k-237 and YAGO3-10. Therefore, we compare our model with RotatE in the setting where both models are trained without self-adversarial negative sampling (Supplementary Table 9). These results further confirm the superiority of our model without self-adversarial negative sampling.

E   Capability of DensE in modeling relation patterns

In this Section, we provide a detailed analysis on how our method tend to model each relation pattern in an interpretable way. In our experiment, we calculate the statistical rule of each degree of freedom to reflect the effect of specific relation patterns. In addition, we also sometimes compare the embeddings of two relation types (or a relation

Table 5: Hyperparameters settings of DensE in this study.

Dataset	Embedding size k	Batch size b	Margin γ	Negative sample size n	adv temperature α
WN18	200	512	12.0	1024	0.3
WN18RR	300	512	6.0	512	0.5
FB15k-237	800	1024	9.0	256	1.0
YAGO3-10	200	1024	24.0	512	1.0

Table 6: Performance comparison on WN18RR. Best results are labeled in bold and the second best are underlined. *First* indicates the originally reported performance of each method. *Recent* shows the best results of more selected recent models. *Ours* reports the performance of DensE. For MRR, the lower, the better; for other metrics, the higher, the better.

	Model	MR	MRR	H@1	H@3	H@10
				WN18RR		
First	RESCAL (Wang et al., 2018)	-	0.420	-	-	0.447
	TransE (Nguyen et al., 2017)	3384	0.226	-	-	0.501
	DistMult (Dettmers et al., 2018)	5110	0.430	0.390	0.440	0.490
	ComplEx (Dettmers et al., 2018)	5261	0.440	0.410	0.460	0.510
Recent	RotatE (Sun et al., 2019)	3340	0.476	0.428	0.492	0.571
	NagE (Yang et al., 2020)	-	0.477	0.432	0.493	0.574
	QuatE (Jia, 2019)	3472	0.481	0.436	0.500	0.564
	D4-STE (Xu and Li, 2019)	-	0.480	**0.452**	0.491	0.536
	TuckER (Balažević et al., 2019)	-	0.470	<u>0.443</u>	0.482	0.526
	Rotate3D (Gao et al., 2020)	<u>3328</u>	0.489	0.442	0.505	0.579
	HAKE (Zhang et al., 2020)	-	**0.497**	**0.452**	**0.516**	<u>0.582</u>
Ours	DensE	**2934**	<u>0.492</u>	<u>0.443</u>	<u>0.509</u>	**0.586**

Table 7: Performance comparison on FB15K-237. Best results are labeled in bold and the second best are underlined. *First* indicates the originally reported performance of each method. *Recent* shows the best results of more selected recent models. *Ours* reports the performance of DensE. For MRR, the lower, the better; for other metrics, the higher, the better.

	Model	MR	MRR	H@1	H@3	H@10
				FB15K-237		
First	RESCAL (Wang et al., 2018)	-	0.270	-	-	0.427
	TransE (Nguyen et al., 2017)	357	0.294	-	-	0.465
	DistMult (Dettmers et al., 2018)	254	0.241	0.155	0.263	0.419
	ComplEx (Dettmers et al., 2018)	339	0.247	0.158	0.275	0.428
Recent	RotatE (Sun et al., 2019)	177	0.338	0.241	0.375	0.533
	NagE (Yang et al., 2020)	-	0.340	0.244	0.378	0.530
	QuatE (Jia, 2019)	176	0.311	0.221	0.342	0.495
	D4-STE (Xu and Li, 2019)	-	0.320	0.230	0.353	0.502
	TuckER (Balažević et al., 2019)	-	**0.358**	**0.266**	**0.394**	**0.544**
	Rotate3D (Gao et al., 2020)	<u>165</u>	0.347	0.250	0.385	<u>0.543</u>
	HAKE (Zhang et al., 2020)	-	0.346	0.250	0.381	0.542
Ours	DensE	**161**	<u>0.351</u>	<u>0.256</u>	<u>0.386</u>	**0.544**

Table 8: Performance comparison on YAGO3-10. Best results are labeled in bold and the second best are underlined. *First* indicates the originally reported performance of each method. *Recent* shows the best results of more selected recent models. *Ours* reports the performance of DensE. For MRR, the lower, the better; for other metrics, the higher, the better.

| | Model | YAGO3-10 | | | | |
		MR	MRR	H@1	H@3	H@10
First	DistMult (Dettmers et al., 2018)	5926	0.340	0.240	0.380	0.540
	ComplEx (Dettmers et al., 2018)	6351	0.360	0.260	0.400	0.550
Recent	RotatE (Sun et al., 2019)	<u>1767</u>	0.495	0.402	0.550	0.670
	D4-STE (Xu and Li, 2019)	-	0.472	0.381	0.523	0.643
	HAKE (Zhang et al., 2020)	-	**0.545**	<u>0.462</u>	**0.596**	**0.694**
Ours	DensE	**1450**	<u>0.541</u>	**0.465**	<u>0.585</u>	<u>0.678</u>

Table 9: Results of DensE and RotatE without self-adversarial negative sampling training technique, where "adv" represents "self-adversarial".

Model	WN18	WN18RR	FB15k-237	YAGO3-10
DensE (w/o adv)	0.950	0.486	0.306	0.452
RotatE (w/o adv)	0.947	0.470	0.297	0.439

type and an entity) *per element*, i.e., we perform element-wise addition, subtraction, multiplication on each embedding dimension. Then, we use the distribution of these results to demonstrate how the two compared embeddings agree with each other. Note that below we use an addition subscript i to denote each dimension in the embeddings.

E.1 Symmetry/anti-symmetry pattern

As pointed out in main text **Section 7.1**, for the **symmetry** relation pattern, the scaling factor $|Q|$ of symmetric relation tend to be one, and the rotation angle ψ should be 0 or π in $[0, 2\pi)$. For **anti-symmetry** relation pattern, one can easily check that the scaling factor $|Q|$ should also be one, but the rotation angle ψ should be neither 0 nor π in the range of $[0, 2\pi)$. Here we show the distributions of rotation angle ψ and scaling factor $|Q|$ of four relations with **symmetry** pattern in WN18RR (Supplementary Figure 6 (a)-(h)). We also show distributions of anti-symmetry relation pattern "member_meronym" in Supplementary Figure 6 (i)-(j). As we can see that the scaling factor is roughly around one. For the rotation angle ψ, there are just few elements fall into the bin that contains π. It should be noted that since the embedding size k in our model for WN18RR is set to be 300, the sum of frequency in these distributions also equals to 300.

(a) $|Q_{derivationally\_related\_form}|$

(b) $\psi_{derivationally\_related\_form}$

(c) $|Q_{also\_see}|$

(d) ψ_{also_see}

(e) $|Q_{similar\_to}|$

(f) $\psi_{similar\_to}$

(g) $|Q_{verb\_group}|$

(h) ψ_{verb_group}

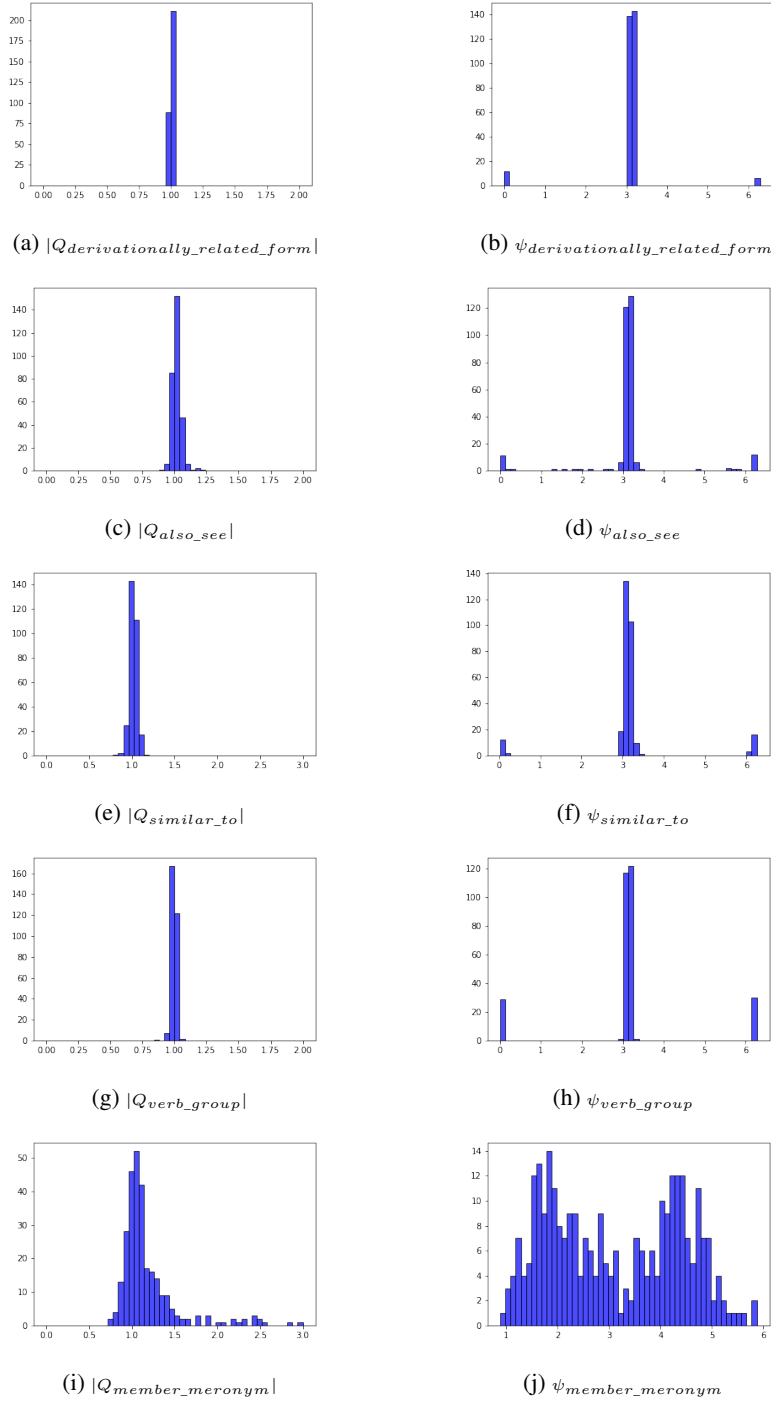(i) $|Q_{member\_meronym}|$

(j) $\psi_{member\_meronym}$

Figure 6: Geometric interpretation of how DensE models symmetry patterns and anti-symmetry patterns. Each row shows the distribution of $|Q|$ and ψ for a given relation, respectively.

E.2 Inversion pattern

In main text **Section 7.1**, we assert that if two relations r_1 and r_2 satisfy the **inverse** pattern, if and only if they satisfy: $|\mathbf{r}_{1i}| * |\mathbf{r}_{2i}| = 1, \theta_{r_{1i}} = \theta_{r_{2i}}, \phi_{r_{1i}} = \phi_{r_{2i}}$ and $\psi_{r_{1i}} + \psi_{r_{2i}} = 2\pi$. In Supplementary Figure 7, we show a case of paired relations with **inversion** pattern from WN18 dataset, namely $has\_part, part\_of$. We plot the scaling factor $|Q|$, magnitude of the rotation ψ, and (θ, ϕ) that describe the rotation axis for each relation (first two columns), as well as their element-wise alignment results (the last column).

E.3 Composition pattern

E.3.1 Ambiguity in composition pattern

In a real-world KG (here we take a sub-graph from WN18RR as an example), due to the ambiguity issue mentioned in the main text (composition pattern **Property 2**), there exist plenty of examples where a third relation (the composite relation) cannot be inferred given the two participating relations alone (Supplementary Figure 8(a)). For example, given $r_1 = derivationally\_related\_form$ and $r_2 = hypernym$, we have the composition pattern as shown with the blue lines: (Trade(VB), derivationally_related_form, Trade(NN)), (Trade(NN), hypernym, transaction) and (Trade(VB), derivationally_related_form , transaction). From these cases, it seems that one can summarize the composition pattern as: $r_3 = r_1 = derivationally\_related\_form$, i.e., $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_1(h, t)$. However, we also have the triangle with red lines, i.e., (Trade(VB), derivationally_related_form, Selling), (Selling, hypernym, mercantilism) and (Trade(VB), synset_domain_topic_of, mercantilism). In these cases, it looks like the composition pattern has the form that $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_3(h, t)$, where $r_3 = synset\_domain\_topic\_of$. This ambiguity means that the composition mode is not uniform but depends on specific entities and their other neighborhoods. Therefore, in order to give the model sufficient flexibility to learn this, our model does not require all the dimensions in a relation embedding to fit in one single composition mode (e.g., $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_1(h, t)$ or $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_3(h, t)$). In consequence, the learned relation embedding for a composite relation are actually distributed in a disperse manner, with the majority of embedding dimensions following mode $r_1(h, h')\Lambda r_2(h', t) \Rightarrow$ $r_1(h, t)$, and some minor portions following $r_1(h, h')\Lambda r_2(h', t) \Rightarrow r_3(h, t)$, which is consistent with the abundance of each mode in the training data.

E.3.2 Case study: a two-hop composition pattern

As we mentioned in the main text, in a composition pattern, the relations involved are not necessarily different (composition pattern **Property 3**). Here we discuss the simplest case of a two-hop relation composition. Even under the assumption that the composite relation can be inferred from its composition alone (i.e., an **unambiguous** composition), there are still five possible situations, as shown in Supplementary Figure 8(b).

Here we provide a detailed analysis on these five situations, using real cases in WN18RR as examples (Supplementary Figure 8(a)).

Situation 1: r_1, r_2 and r_3 are all the same, e.g., the triangle in yellow color with three triplets: (man, has_part, arm), (arm, has_part, paw) and (man, has_part, paw). To satisfy the composition relation constraint, in principle the model can adopt several approaches to fit this situation. From the perspective of relation embedding, it can learn to fit $|Q_{has\_part}| = 1$ or $\psi_{has\_part} \in \{0, 2\pi\}$. Also, in terms of interaction between entities and relations, it can also learn to align the rotation axis of relation with the entity embedding (see Supplementary Figure 9(c), showing the differences between the θ component of r_1 in the mode $r_1(h, h')\Lambda r_1(h', t) \Rightarrow r_1(h, t)$ and the corresponding head entities h in spherical coordinate system). Note that when the rotation axis aligns with head entity embedding in spherical coordinate system, the composition pattern can be modeled by the scaling factor alone, i.e., $|Q_{has\_part}|^2 = |Q_{has\_part}|$, regardless of the rotation magnitude ψ.

In this case, we see the model mainly adopts the first approach (Supplementary Figure 9(a)), as $|Q_{has\_part}|$ is roughly around 1. In addition, the partial alignment between relation rotation axis and entities is also observed (Supplementary Figure 9(c)). In contrast, for ψ_{has_part}, we see very few dimensions satisfy the condition (0 or 2π), which is also consistent with the dominate role of scaling factor in this case. Noticed the fact that the relation "hypernym" dominates the composite patterns (Refer to Table 3 in the main text), we also find that $r_1 = r_2 = r_3 =$ hypernym pattern exists in the WN18RR dataset. The above similar analy-
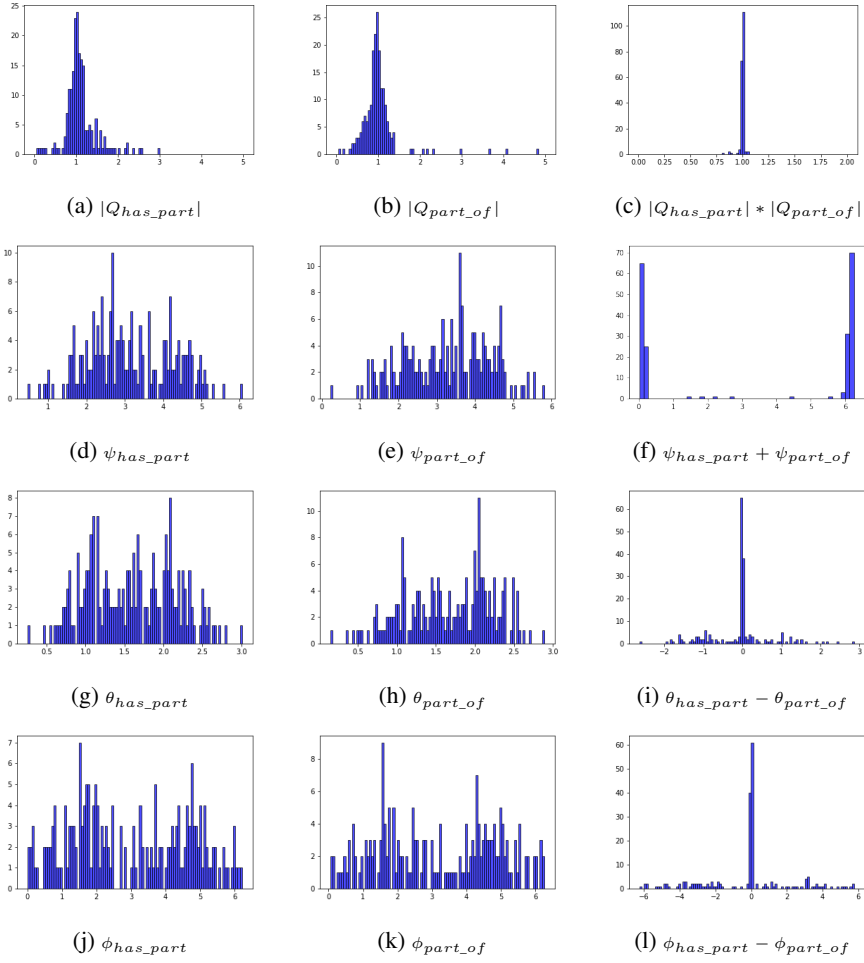
(a) $|Q_{has\_part}|$

(b) $|Q_{part\_of}|$

(c) $|Q_{has\_part}| * |Q_{part\_of}|$

(d) ψ_{has_part}

(e) ψ_{part_of}

(f) $\psi_{has\_part} + \psi_{part\_of}$

(g) θ_{has_part}

(h) θ_{part_of}

(i) $\theta_{has\_part} - \theta_{part\_of}$

(j) ϕ_{has_part}

(k) ϕ_{part_of}

(l) $\phi_{has\_part} - \phi_{part\_of}$

Figure 7: Geometric interpretation of how DensE models inverse pattern, using an example of $(has\_part, part\_of)$ from WN18. At each row, we show the embedding from one degree of freedom of our model. The first two columns show the embeddings of each relation type, and the last column shows the alignment of the two embeddings regarding a specific degree of freedom.
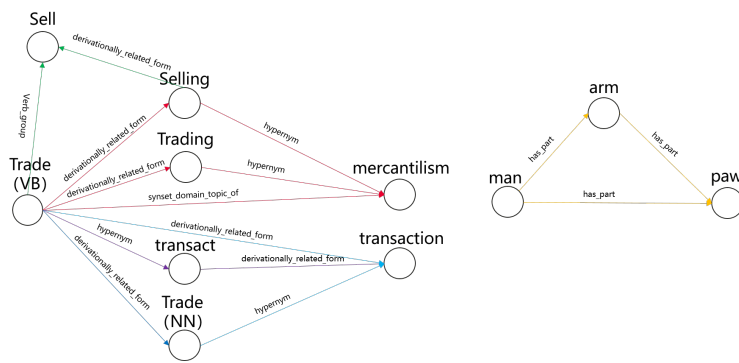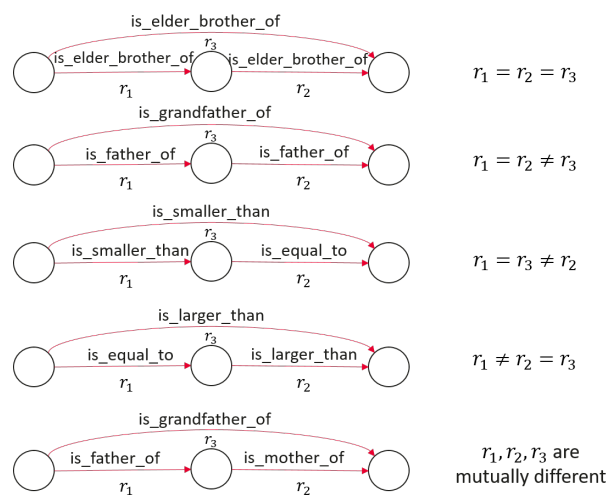
Figure 8: (a) A representative subgraph from WN18RR to show ambiguity. (b) The relations involved in a composition pattern are not necessarily to be mutually different.

sis can be done for this case. The relevant results are shown in Supplementary Figure 9(d), (e), and (f), respectively.

Situation 2: r_1 and r_2 are the same, but not equal to r_3. An example is given by the triangle in green with three triplets: $(Trade(VB)$, $derivationally\_related\_form$, $Selling)$, $(Selling$, $derivationally\_related\_form$, $Sell)$ and $(Trade(VB)$, $verb\_group$, $Sell)$. As mentioned in main text **Section 7.1**, we expect the model to learn $\psi_{verb\_group} = 2\psi_{derivationally\_related\_form}$ and $|Q_{verb\_group}| = |Q_{derivationally\_related\_form}|^2$, both of which are confirmed in this case. Besides this, we also observe that the rotation axes of the two relations are aligned on some dimensions, i.e, $\theta_{verb\_group} = \theta_{derivationally\_related\_form}$ and $\phi_{verb\_group} = \phi_{derivationally\_related\_form}$. The corresponding distributions are shown in Supplementary Figure 10, where we have $r_1 = r_2 = derivationally\_related\_form$ and $r_3 = verb\_group$.

Situation 3: r_1 and r_3 are the same, but not equal to r_2, e.g., the triangle in blue color with three triplets set: $(Trade(VB), derivationally\_related\_form$, $Trade(NN))$, $(Trade(NN)$, $hypernym$, $transaction)$ and $(Trade(VB)$, $derivationally\_related\_form$, $transaction)$. Here we compare the difference between the embedding of a composite relation and the embedding calculated by multiplying each relation in the relation path. We have already shown the distribution of ψ in the main context. Other distributions are shown in Supplementary Figure 11, where we have $r_1 = r_3 = derivationally\_related\_form$ and $r_2 = hypernym$.

Situation 4: r_2 and r_3 are the same, but not equal to r_1, e.g., the triangle in purple color with three triplets set: $(Trade(VB)$, $hypernym$, $transact)$, $(transact$, $derivationally\_related\_form$, $transaction)$ and $(Trade(VB)$, $derivationally\_related\_form$, $transaction)$. Here we compare the difference between the embedding of a composite relation and the embedding calculated by multiplying each relation in the relation path. Related distributions are shown in Supplementary Figure 12, where we have $r_1 = hypernym$ and $r_2 = r_3 = derivationally\_related\_form$. In

Table 10: Performance comparison on WN18. The performances of RotatE are obtained from the original paper.

Model	MR	MRR	H@1	H@3	H@10
RotatE	309	0.949	0.944	0.952	0.959
DensE	285	0.950	0.945	0.954	0.959

this case, we see that the composition pattern given by $\mathcal{O}(\mathbf{r2})\mathcal{O}(\mathbf{r1})$ is learned to have similar scaling factor $|Q|$ and rotation magnitude ψ with r_2, and the rotation axes (θ, ϕ) are also partially aligned.

Situation 5: r_1, r_2 and r_3 are mutually different relations, e.g., the triangle in red color with three triplets set: $(Trade(VB)$, $derivationally\_related\_form$, $Selling)$, $(Selling$, $hypernym$, $mercantilism)$ and $(Trade(VB)$, $synset\_domain\_topic\_of$, $mercantilism)$. Here we compare the difference between the embedding of a composite relation and the embedding calculated by multiplying each relation in the relation path. Related distributions are shown in Supplementary Figure 13, where we have $r_1 = derivationally\_related\_form$, $r_2 = hypernym$ and $r_3 = synset\_domain\_topic\_of$. The reason for relatively large dispersion here is discussed in the main text, i.e., the majority of triplets exemplify $derivationally\_related\_form \Lambda hypernym \Rightarrow derivationally\_related\_form$ (See also the discussion above of **Situation 3**), while only a small portion have the pattern $derivationally\_related\_form \Lambda hypernym \Rightarrow synset\_domain\_topic\_of$.

F   Link prediction results on WN18 dataset

In the main text, we report the MRR and H@10 performance on WN18RR, FB15K237 and YAGO3-10. Here, we also report our model's performance on WN18 (Supplementary Table 10), from which **inverse** relations are extracted by the demonstration of geometric interpretation of DensE (main text Figure 3).

G   Variance of the prediction performance

The mean values and corresponding variance of MRR on WN18, WN18RR, FB15k-237 and YAGO3-10 datasets are shown in Supplementary
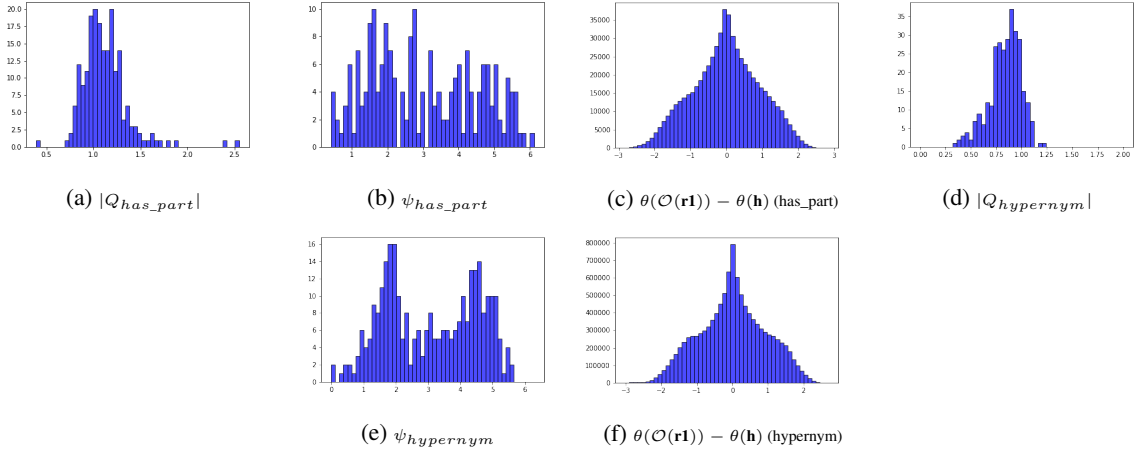
Figure 9: Geometric interpretation of composition patterns in **Situation 1**.


Figure 10: Geometric interpretation of composition patterns in **Situation 2**.

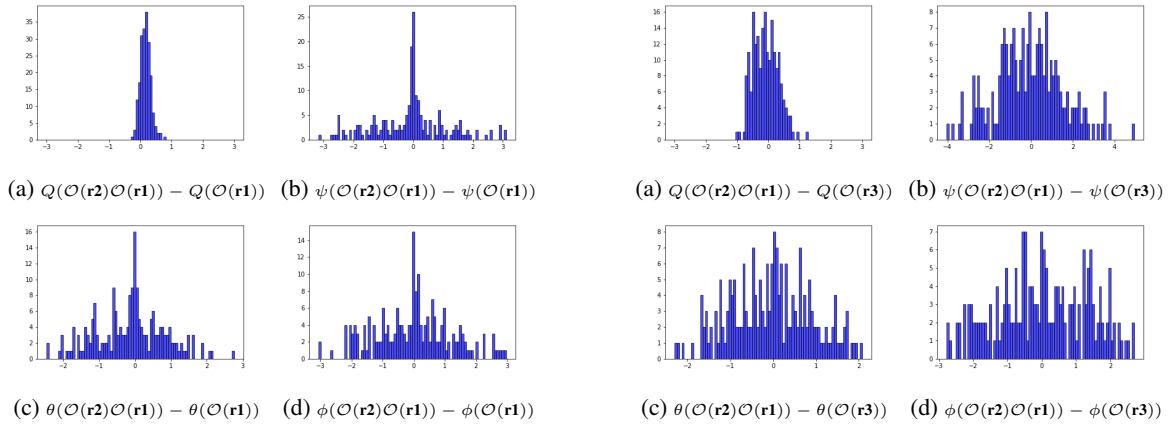Figure 12: Geometric interpretation of composition patterns in **Situation 4**.

Figure 11: Geometric interpretation of composition patterns in **Situation 3**.

Figure 13: Geometric interpretation of composition patterns in **Situation 5**.

Table 11. The results are obtained by training DensE with five different random seeds, showing that the prediction performance of DensE is relatively stable.

References

Ivana Balažević, Carl Allen, and Timothy M Hospedales. 2019. Tucker: Tensor factorization for knowledge graph completion. *arXiv preprint arXiv:1901.09590*.

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.

Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

Chang Gao, Chengjie Sun, Lili Shan, Lei Lin, and Mingjiang Wang. 2020. Rotate3d: Representing relations as rotations in three-dimensional space for knowledge graph embedding. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 385–394.

Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S Yu. 2020. A survey on knowledge graphs: Representation, acquisition and applications. *arXiv preprint arXiv:2002.00388*.

Yan-Bin Jia. 2019. Quaternions. *Com S*, 477:577.

Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. 2018. Canonical tensor decomposition for knowledge base completion. In *ICML*.

Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 2019. Kagnet: Knowledge-aware graph networks for commonsense reasoning. In *EMNLP/IJCNLP*.

Farzaneh Mahdisoltani, Joanna Biega, and Fabian M Suchanek. 2013. Yago3: A knowledge base from multilingual wikipedias.

George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.

Deepak Nathani, Jatin Chauhan, Charu Sharma, and Manohar Kaul. 2019. Learning attention-based embeddings for relation prediction in knowledge graphs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4710–4723.

Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. 2017. A novel embedding model for knowledge base completion based on convolutional neural network. *arXiv preprint arXiv:1712.02121*.

Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016. Holographic embeddings of knowledge graphs. In *Thirtieth Aaai conference on artificial intelligence*.

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.

Daniel Ruffinelli, Samuel Broscheit, and Rainer Gemulla. 2019. You can teach an old dog new tricks! on training knowledge graph embeddings. In *International Conference on Learning Representations*.

Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, pages 593–607. Springer.

Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. In *The Seventh International Conference on Learning Representations*.

Kristina Toutanova and Danqi Chen. 2015a. Observed versus latent features for knowledge base and text inference.

Kristina Toutanova and Danqi Chen. 2015b. Observed versus latent features for knowledge base and text inference. In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality*, pages 57–66.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016a. Complex embeddings for simple link prediction. In *International Conference on Machine Learning*, pages 2071–2080.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016b. Complex embeddings for simple link prediction. *ArXiv*, abs/1606.06357.

Yanjie Wang, Rainer Gemulla, and Hui Li. 2017. On multi-relational link prediction with bilinear models. *CoRR*, abs/1709.04808.

Yanjie Wang, Daniel Ruffinelli, Rainer Gemulla, Samuel Broscheit, and Christian Meilicke. 2018. On evaluating embedding models for knowledge base completion. *arXiv preprint arXiv:1810.07180*.

Canran Xu and Ruijiang Li. 2019. Relation embedding with dihedral group in knowledge graph. *arXiv preprint arXiv:1906.00687*.

Table 11: The mean values and variance of MRR on WN18, WN18RR, FB15k-237 and YAGO3-10 datasets.

	WN18	WN18RR	FB15k-237	YAGO3-10
MRR	0.950±0.001	0.492±0.001	0.351±0.001	0.541±0.001

An Yang, Quan Wang, Jing Liu, Kai Liu, Yajuan Lyu, Hua Wu, Qiaoqiao She, and Sujian Li. 2019. Enhancing pre-trained language representations with rich knowledge for machine reading comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2346–2357, Florence, Italy. Association for Computational Linguistics.

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*.

Tong Yang, Long Sha, and Pengyu Hong. 2020. Nage: Non-abelian group embedding for knowledge graphs. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, CIKM '20, pages 1735–1742, New York, NY, USA. Association for Computing Machinery.

Shuai Zhang, Yi Tay, Lina Yao, and Qi Liu. 2019a. Quaternion knowledge graph embeddings. In *Advances in Neural Information Processing Systems*, pages 2731–2741.

Zhanqiu Zhang, Jianyu Cai, Yongdong Zhang, and Jie Wang. 2020. Learning hierarchy-aware knowledge graph embeddings for link prediction. In *AAAI*, pages 3065–3072.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019b. Ernie: Enhanced language representation with informative entities. *ArXiv*, abs/1905.07129.