

Generating \mathcal{SROI}^- Ontologies via Knowledge Graph Query Embedding Learning

Yunjie He ^{ib,*,}, Daniel Hernandez ^{ib,a,}, Mojtaba Nayyeri ^{ib,a,}, Bo Xiong ^{ib,a,}, Yuqicheng Zhu ^{ib,ab,}, Evgeny Kharlamov ^{ib,bd} and Steffen Staab ^{ib,ac}

^aUniversity of Stuttgart

^bBosch Center for Artificial Intelligence

^cUniversity of Southampton

^dUniversity of Oslo

Abstract. Query embedding approaches answer complex logical queries over incomplete knowledge graphs (KGs) by computing and operating on low-dimensional vector representations of entities, relations, and queries. However, current query embedding models heavily rely on excessively parameterized neural networks and cannot explain the knowledge learned from the graph. We propose a novel query embedding method, AConE, which explains the knowledge learned from the graph in the form of \mathcal{SROI}^- description logic axioms while being more parameter-efficient than most existing approaches. AConE associates queries to \mathcal{SROI}^- description logic concepts. Every \mathcal{SROI}^- concept is embedded as a cone in complex vector space, and each \mathcal{SROI}^- relation is embedded as a transformation that rotates and scales cones. We show theoretically that AConE can learn \mathcal{SROI}^- axioms, and defines an algebra whose operations correspond one-to-one to \mathcal{SROI}^- description logic concept constructs. Our empirical study on multiple query datasets shows that AConE achieves superior results over previous baselines with fewer parameters. Notably on the WN18RR dataset, AConE achieves significant improvement over baseline models. We provide comprehensive analyses showing that the capability to represent axioms positively impacts the results of query answering.

1 Introduction

Knowledge Graphs (KGs) such as Wikidata [23], Freebase [4], and YAGO [20] represent real-world facts as sets of *triples* of the form (s, p, o) which encode atomic assertions as $p(s, o)$. Graph database engines can store and query KGs efficiently using query languages such as SPARQL [1] that can express a variety of queries that result of combining atomic queries called *triple patterns*. The first step of querying KGs is answering triple patterns with the stored triples. However, when querying incomplete Knowledge Graphs (KGs), some triples are not explicitly available in the triple store. As a result, these triples are neither included in the answers nor in the intermediate results. To provide plausible answers beyond what is known, these missing triples must be inferred. Figure 1 shows an example, where the available triples $(L.Messi, playsFor, ArgentinaNFT)$, and $(ArgentinaNFT, teamWon, WorldCup)$ should suggest that triple $(L.Messi, athleteWon, WorldCup)$ is a missing triple. KG embed-

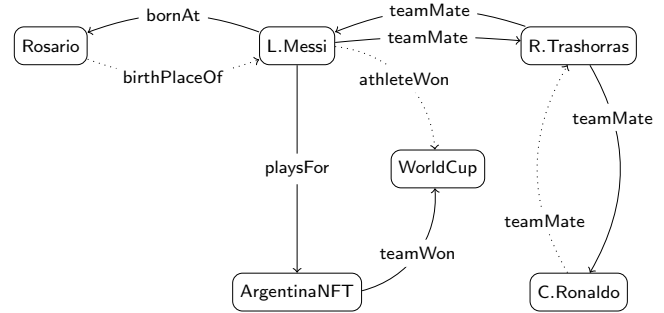


Figure 1: An example of an incomplete KG. The nodes represent entities, the edges with solid lines represent known atomic statements, whereas the edges with dotted lines represent missing atomic statements that must be inferred.

ding methods can predict these missing triples [5, 28, 15] by learning how to embed entities and relations into vector representations, which can be points or more complex geometric objects. These methods use these embeddings to answer triple patterns by computing plausibility scores by applying geometric operations.

Query embedding methods [18, 29, 17] go beyond querying triple patterns. They provide plausible answers to queries that combine triple patterns into first-order logic queries with logical connectives (e.g., negation (\neg), conjunction (\wedge), and disjunction (\vee)). However, current query embedding approaches are restricted to queries with a single unquantified variable and are called *tree-form queries* because their *computation graph* is a tree [19]. The tree-form queries correspond to the \mathcal{SROI}^- description logic concepts that do not include the symbols \top (the concept for all entities) nor \perp (the concept for no elements), nor concept names (e.g., Athlete or Team), but nominals (i.e., concepts with a unique element like $\{C.Ronaldo\}$).

Example 1. The query seeking the birthplaces of the athletes who have won either the World Cup or the Europe Cup but do not play in the same team as C.Ronaldo can be expressed as the \mathcal{SROI}^- concept

$$C \equiv \exists birthPlaceOf. (\neg (\exists teamMate. \{C.Ronaldo\}) \sqcap (\exists athleteWon. \{WorldCup\} \sqcup \exists athleteWon. \{EuroCup\})),$$

whose computation graph is depicted in Figure 2.

* Corresponding Author. Email: yunjie.he@ki.uni-stuttgart.de




Figure 2: Computation graph for the concept C in Example 1.

Query embedding methods [18, 29, 17] learn a model M to predict the answers to queries by embedding entities and relations as geometrical objects of a vector space. These geometrical objects are used to recursively compute geometrical objects for all nodes in the computation graph of the query. Finally, the query embeddings are compared with the entity embeddings of the candidate answers to the query by using a similarity function that represents the plausibility of an answer.

The quality of query embedding methods depends on their ability to represent the *logical patterns* satisfied by a knowledge graph. For example, some relations are symmetric (e.g., *teamMate*) while others are antisymmetric (e.g., *playsFor*); some relations are the inverse of other relations (e.g., *bornAt* and *birthPlaceOf*) and some relations may be composed by others (e.g., if an athlete x playsFor a team y and y teamWon a cup z , then x athleteWon z). As these logical patterns influence the interplay of entities and relations, several works have studied their effect on KG embeddings [21, 14, 26] and demonstrated that an embedding's ability to support them improves its link prediction quality. For query embeddings, similar support of logical patterns in the embedding space is still lacking.

Given a model M of the query embedding method, we can distinguish two ways in which M can exploit a logical pattern P . The first is the ability to *predict* answers that are entailed by P . The second is the ability of model M to *explain* such predictions by inferring P as an axiom that can be obtained from the geometrical relations between the objects in the embedding space.

Although query embedding methods have achieved great success in predicting query results over incomplete data, they fail to explain the learned knowledge because, to achieve better results, they endow geometric operations with neural network operations that impede explainability. For example, the method BetaE [17] represents queries q with multidimensional beta distributions \mathbf{q} , and the set of answers to q in M , denoted $\llbracket q \rrbracket_M$, consists of the entities e such that $e \in \mathbf{q}$. A query $q(y) = \exists x(q_1(x) \wedge r(x, y))$, extending a given query $q_1(x)$ with a relation r , is embedded as a beta distribution \mathbf{q} that results from applying a neural network over the distribution \mathbf{q}_1 for query q_1 . The use of neural networks to represent relations hinders the inference of axioms like $\text{playsFor} \circ \text{teamWon} \sqsubseteq \text{athleteWon}$. On the other hand, BoxEL [25] and Box²EL [11] embed description logic concepts, but are limited to reduced description logics (\mathcal{EL} and \mathcal{EL}^{++}).

In this paper, we propose a novel query embedding model, AConE, that can explain several logical patterns expressed as $SRQI^2$ axioms learned for query answering tasks. Being more parameter-efficient than most existing approaches [29, 6, 31, 18], AConE reduces the dependency on neural networks by translating logical operators to a simpler algebraic structure. To achieve this explainability, AConE embeds each $SRQI^2$ concept as a multidimensional cone in the complex vector space, and relations as scaling and rotations of cones. Then, each logical operator is translated into geometric operations in embedding space. By doing so, our method can generate $SRQI^2$ ontologies by learning query embeddings.

In summary, this paper makes the following contributions:

1. We formalize the notions of tree-form logical query in terms of $SRQI^2$ concepts, and logical pattern in terms of $SRQI^2$ axioms. We propose an algebra of cones in the complex plane, and theoretically identify the subset of the $SRQI^2$ axioms that can be represented with the cone algebra models (Section 4).
2. We present criteria where a multicone embedding expresses six different logical patterns, namely role containment, composition, transitivity, inverse, symmetry, and asymmetry (Section 5).
3. We propose a novel method, AConE, which embeds $SRQI^2$ concepts as cones in the complex plane (Section 6). This technique allows us to leverage the rotation operator through Euler's formula while maintaining the geometric representation of concepts and concept operators to allow the explanation of logical patterns as $SRQI^2$ axioms that the embedding explicitly encodes.
4. We show that modeling cones in the complex plane, coupled with the rotation operator as a complex product, is more parameter-efficient than modeling cones in the 2D real space with multi-layer neural operators (Section 8).
5. We create new datasets and dataset splits (Section 7.3) to conduct a more detailed analysis of the influence of patterns in complex query answering, providing a finer-grained evaluation, and more in-depth insights into the often-overlooked problem of pattern inference in complex query answering.
6. Our experiments show that AConE outperforms state-of-the-art baselines that represent query regions using vectors, geometries, or distributions, using geometric operations to model logical operations one-to-one (Section 7.2).

2 Related Work

Logical patterns in knowledge graphs. KG embedding methods aim to learn KG representations that capture latent structural and logical patterns [5, 28, 22, 21, 8, 24]. In particular, RotatE [21] captures a broad range of logical patterns, such as symmetry, inversion, and composition, among others. While these KG embedding methods excel at predicting links, they cannot answer first-order logical queries.

BoxEL [25, 30] and Box²EL [11] embed description logic concepts, but are limited to reduced description logics (\mathcal{EL} and \mathcal{EL}^{++}). [16] proposes a method to embed \mathcal{ALC} concepts using al-cones, which differs from our proposal in the geometrical representation of concepts. However, they do not propose a concrete geometry to embed relations (which is required for embedding tree-form queries).

Query answering. Path-based [27, 13], neural [6, 18, 29, 17, 12], and neural-symbolic [2, 32, 31] methods have been developed to answer (subsets of) queries. Among these methods, geometric and probabilistic query embedding approaches [6, 18, 29, 17] provide an effective way to answer tree-form queries over incomplete and noisy KGs. This is done by representing entity sets as geometric objects or probability distributions, such as boxes [18], cones [29], or Beta distribution [17], and performing neural logical operations directly on them. The Graph Query Embedding (GQEs) [6] was first proposed to answer only conjunctive queries via modeling the query q as single vector \mathbf{q} through neural translational operators. However, modeling a query as a single vector limits the model's expressiveness in modeling multiple entities. Query2Box [18] remedies this flaw by modeling entities as points within boxes. This allows Query2Box to predict the intersection of entity sets as the intersection of boxes in vector space. ConE [29] was proposed as the first geometry-based query embedding method that can handle negation via embedding the set of entities (query embedding) as cones in Euclidean space.

All of the above query embedding methods commonly apply multi-layer perceptron networks for selecting answer entities of atomic queries by relation and performing logical operations. Such methods suffer from two problems. Firstly, their ability to capture logical patterns in KGs remains unclear due to the limited explainability of neural networks. Secondly, a large number of parameters need to be trained for an outstanding model performance. Our proposed method, AConE, overcomes these issues by providing a one-to-one mapping between logical and geometrical operators.

3 Preliminaries

The Description Logic SROI⁻. We next present the standard SROI⁻ syntax and semantics and assume standard semantics as defined in Baader et al. [3]. For these definitions, we assume three pairwise disjoint sets **C**, **R**, and **I**, whose elements are respectively called *concept names*, *relation names*, and *individual names*.

Definition 1 (SROI⁻ Concept Descriptions). SROI⁻ concept descriptions C and relation descriptions R are defined by the following grammar

$$\begin{aligned} C &::= \top \mid A \mid \{a\} \mid \neg C \mid C \sqcap C \mid \exists R.C \\ R &::= r \mid r^- \end{aligned}$$

where the symbol \top is a special concept name, and symbols A , a , and r stand for concept names, individual names, and relation names, respectively. Concept descriptions $\{a\}$ are called *nominals*.

Given two concept descriptions C and D , the expression $C \sqsubseteq D$ is a concept-axiom. Given the relation descriptions R, S, R_1, \dots, R_n (with $n \geq 1$), the expressions $\rho_1 \circ \dots \circ \rho_{n-1} \sqsubseteq \rho_n$, $\text{Disj}(R, S)$, $\text{Trans}(R)$, $\text{Ref}(R)$, $\text{Irref}(R)$, $\text{Sym}(R)$, and $\text{Asym}(R)$ are relation-axioms. Given two individual names $a, b \in \mathbf{I}$, a concept description C and a relation description ρ , $a(C)$ is a concept-assertion and $\rho(a, b)$ is a relation-assertion.

We write $C \equiv D$ as an abbreviation for two axioms $C \sqsubseteq D$ and $D \sqsubseteq C$, and likewise for $\rho_1 \equiv \rho_2$. We write \perp , $C \sqcup D$, $\forall \rho.C$ as abbreviations for $\neg \top$, $\neg(\neg C \sqcap \neg D)$ and $\neg \exists \rho. \neg C$, respectively.

An SROI⁻ knowledge base (or ontology) \mathcal{K} is a triple $(\mathcal{R}, \mathcal{T}, \mathcal{A})$ where \mathcal{R} is a finite set of relation-axioms called the *RBox*, \mathcal{T} is a finite set of concept-axioms called the *TBox*, and \mathcal{A} is a finite set of assertions called the *ABox*.

An *interpretation* \mathcal{I} is a pair $(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consisting of a set $\Delta^{\mathcal{I}}$, called the *domain*, and a function $\cdot^{\mathcal{I}}$ such that we have for each individual name $a \in \mathbf{I}$, an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$; for each concept name $A \in \mathbf{C}$, a subset $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$; and for each relation name $r \in \mathbf{R}$, a relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. An interpretation \mathcal{I} is a *model* of a knowledge base \mathcal{K} if and only if all the axioms and assertions in \mathcal{K} are satisfied according to the standard semantics defined in Baader et al. [3]. Given two knowledge bases \mathcal{K}_1 and \mathcal{K}_2 , \mathcal{K}_1 *entails* \mathcal{K}_2 , denoted $\mathcal{K}_1 \models \mathcal{K}_2$, if and only if every model \mathcal{I} of \mathcal{K}_1 is also a model of \mathcal{K}_2 .

Knowledge graphs and queries. We next define knowledge graphs in terms of SROI⁻ knowledge bases, and tree-form queries in terms of SROI⁻ concepts.

Definition 2 (Knowledge Graph). A knowledge graph G is a SROI⁻ knowledge base \mathcal{K} whose *RBox* is empty, and its \mathcal{T} contains a unique axiom $\top \sqsubseteq \{a_1\} \sqcup \dots \sqcup \{a_n\}$, where $\{a_1, \dots, a_n\}$ is the set of all individual names occurring in the *ABox*. This axiom is called *domain-closure assumption*.

Definition 3 (Tree-form query). Given a knowledge graph G , and an individual name x that does not occur in G , a tree-form query q is an assertion $C(x)$ where C is a SROI⁻ concept description. The answers to query q , are all individuals a occurring in G such that $G \models C(a)$.

Informally, the query answering task over incomplete data consists of predicting answers to a query q over a knowledge graph pattern G given only a subset $G' \subset G$. Knowledge graph embeddings to generalize the knowledge on G' to predict answers on G .

4 SROI⁻ concepts and Cone Algebra

In this section, we present an algebra of cones in the complex plane, called the Cone Algebra, we show the correspondence between this algebra and SROI⁻ concepts, and we identify the subset of the SROI⁻ axioms that are expressible with this algebra.

Definition 4 (Cone). A cone $\mathcal{C}(\alpha, \beta)$ is a region in the complex plane \mathbb{C} determined by a pair of angles $(\alpha, \beta) \in \mathbb{R}^2$ as follows:

$$\mathcal{C}(\alpha, \beta) = \{e^{i\theta} : \alpha \leq \theta \text{ and } \theta \leq \beta\}.$$

The empty cone, denoted \mathcal{C}_{\perp} , is the cone such that $\alpha > \beta$. A singleton cone with angle α , denoted \mathcal{C}_{α} , is a cone such that $\alpha = \beta$. A proper cone, denoted $\mathcal{C}_{\alpha \rightarrow \beta}$, is a cone where $\alpha + 2\pi > \beta$. A full cone, denoted \mathcal{C}_{\top} , is the cone such that $\alpha + 2\pi \leq \beta$.

Notice that $e^{i\theta} = e^{i\theta + 2k\pi}$, for every natural number k . Thus, the same cone can be determined with multiple combinations of angles.

Notice that the intersection is not closed on the set of cones. Indeed, the region $\mathcal{C}(0, \frac{3}{2}\pi) \cap \mathcal{C}(\pi, \frac{5}{2}\pi)$ is not a cone. This hinders the use of cones to represent SROI⁻ concepts. To overcome this limitation, we can embed concepts on sets of cones.

Definition 5 (Multicone algebra). A multicone $\text{MC}(\mathcal{C})$ is a region determined by a set of cones \mathcal{C} as follows:

$$\text{MC}(\mathcal{C}) = \bigcup_{\mathcal{C}(\alpha, \beta) \in \mathcal{C}} \mathcal{C}(\alpha, \beta)$$

The multicone algebra is the algebra over the set of multicones defined by the binary operations \cap and \cup that are defined as the set operations over the multicon regions.

Proposition 1. The operations \cup and \cap , are commutative, associative, and mutually distributive, and their identity elements are \mathcal{C}_{\perp} and \mathcal{C}_{\top} , respectively.

Definition 6 (Rotation algebra). Given a triple $(\theta, \gamma, \delta) \in \mathbb{R}^3$ where $0 \leq \gamma \leq 2\pi$, a rotation $\text{R}(\theta, \gamma, \delta)$ is a function that maps every singleton and proper cone $\mathcal{C}(\alpha, \beta)$ to the cone $\mathcal{C}(\alpha', \beta')$ such that

$$\beta' + \alpha' = \theta(\beta + \alpha), \quad \beta' - \alpha' = \gamma(\beta - \alpha) + \delta,$$

and that maps the empty and the full cone to themselves:

$$\text{R}(\theta, \gamma, \delta)(\mathcal{C}_{\perp}) = \mathcal{C}_{\perp}, \quad \text{R}(\theta, \gamma, \delta)(\mathcal{C}_{\top}) = \mathcal{C}_{\top}.$$

The rotation parameters θ , γ , and δ are called the rotation angle and the aperture factor, and the aperture adding. We call aperture-multiplicative and aperture-additive rotations to rotations of the respective forms $\text{R}(\theta, \gamma, 0)$ and $\text{R}(\theta, 1, \delta)$.

We call rotation algebra to be the algebraic structure whose ground set is the set of rotations over cones, and has a binary operation \circ denoting function composition (i.e., $(f \circ g)(x) = g(f(x))$).

Table 1: Multicone Embedding Semantics.

Concept	Semantics
\top	MC_{\top}
$\{a\}$	$\text{MC}(\{C_{\theta}\})$ where $a^{\mathcal{E}} = e^{i\theta}$
$C \sqcap D$	$C^{\mathcal{E}} \cap D^{\mathcal{E}}$
$\neg C$	$\text{MC}_{\top} \setminus C^{\mathcal{E}}$
$\exists R.C$	$R^{\mathcal{E}}(C^{\mathcal{E}})$
r^{-}	$(r^{\mathcal{E}})^{-}$
$R \circ S$	$R^{\mathcal{E}} \circ S^{\mathcal{E}}$
$A(a)$	$a^{\mathcal{E}} \in A^{\mathcal{E}}$
$R(a, b)$	$(a^{\mathcal{E}}, b^{\mathcal{E}}) \in R^{\mathcal{E}}$
$C \sqsubseteq D$	$C^{\mathcal{E}} \subseteq D^{\mathcal{E}}$
$R_1 \circ \dots \circ R_n \sqsubseteq S$	$(R_1 \circ \dots \circ R_n)^{\mathcal{E}} \subseteq S^{\mathcal{E}}$
$\text{Disj}(R, S)$	$R^{\mathcal{E}} \cap S^{\mathcal{E}} = \emptyset$
$\text{Trans}(R)$	$(R \circ R)^{\mathcal{E}} \subseteq R^{\mathcal{E}}$
$\text{Ref}(R)$	$\{(a, a) \mid a \in \mathbf{I}\} \subseteq R^{\mathcal{E}}$
$\text{Irref}(R)$	$\{(a, a) \mid a \in \mathbf{I}\} \cap R^{\mathcal{E}} = \emptyset$
$\text{Sym}(R)$	for all $a, b \in \mathbf{I}$, if $(a, b) \in R^{\mathcal{E}}$ then $(b, a) \in R^{\mathcal{E}}$
$\text{Asym}(R)$	for all $a, b \in \mathbf{I}$, if $(a, b) \in R^{\mathcal{E}}$ then $(b, a) \notin R^{\mathcal{E}}$

We write $R\langle\theta, \gamma, \delta\rangle^{-} = R\langle-\theta, \frac{1}{\gamma}, -\delta\rangle$, and with a slight abuse of notation, given a multicone $\text{MC}(\mathcal{C})$, we write

$$R\langle\theta, \gamma, \delta\rangle(\text{MC}(\mathcal{C})) = \text{MC}(\{R\langle\theta, \gamma, \delta\rangle(C(\alpha, \beta)) : C(\alpha, \beta) \in \mathcal{C}\}).$$

Definition 7 (Multicone Embedding). A multicone embedding is a function \mathcal{E} that maps each individual name $a \in \mathbf{I}$ to a complex number $a^{\mathcal{E}} = e^{i\theta} \in \mathbf{C}_{\top}$, each concept name $A \in \mathbf{C}$ to a multicone $A^{\mathcal{E}}$, and each relation name $r \in \mathbf{R}$ to a rotation $r^{\mathcal{E}} = R\langle\theta, \gamma, \delta\rangle$. The multicone embedding of a knowledge graph G is a multicone embedding restricted to the individual names and relation names occurring in knowledge graph G .

Given a multicone embedding \mathcal{E} , Table 1 defines the multicone $C^{\mathcal{E}}$ corresponding to a SROI^{-} concept description C . The definition of the multicone semantics for relations, assertions and axioms is straightforward. For example, the embedding of $R \circ S$ is the relation $\{(a, c) : (a, b) \in R^{\mathcal{E}}, (b, c) \in S^{\mathcal{E}}\}$, and an $\mathcal{E} \models C(a)$ if and only if $a^{\mathcal{E}} \in C^{\mathcal{E}}$, $\mathcal{E} \models R(a, b)$ if and only if $(a^{\mathcal{E}}, b^{\mathcal{E}}) \in R^{\mathcal{E}}$, and $C \sqsubseteq D$ if and only if $C^{\mathcal{E}} \subseteq D^{\mathcal{E}}$.

The full cone. Since \mathbf{C}_{\top} is an absorbing element for the rotation algebra, multicone embeddings infer wrong axioms $\top \sqsubseteq \exists R.\top$.

The rotation commutativity. In certain cases, the order in which rotations are applied does not affect the resulting cone.

Proposition 2. Given two aperture-multiplicative rotations $R\langle\theta_1, \gamma_1, 0\rangle$ and $R\langle\theta_2, \gamma_2, 0\rangle$ such that $\gamma_1 \leq 1$ and $\gamma_2 \leq 1$, then

$$R\langle\theta_1, \gamma_1, 0\rangle \circ R\langle\theta_2, \gamma_2, 0\rangle = R\langle\theta_2, \gamma_2, 0\rangle \circ R\langle\theta_1, \gamma_1, 0\rangle.$$

The issue with the commutativity is that if two relation names r and s are embedded with aperture-multiplicative rotations $R\langle\theta_r, \gamma_r, 0\rangle$ and $R\langle\theta_s, \gamma_s, 0\rangle$ with $\gamma_r \leq 1$ and $\gamma_s \leq 1$, then we can infer that $r \circ s \equiv s \circ r$. So, the commutativity of these rotations may introduce a bias into the models that can lead to wrong predictions.

This bias favoring commutativity does not hold if $\gamma > 1$, as the following counter example shows.

$$\begin{aligned} (R\langle 0, 2, 0\rangle \circ R\langle 0, \frac{1}{2}, 0\rangle)(C(0, \pi)) &= \mathbf{C}_{\top}, \\ (R\langle 0, \frac{1}{2}, 0\rangle \circ R\langle 0, 2, 0\rangle)(C(0, \pi)) &= C(0, \pi). \end{aligned}$$

However, rotations with $\gamma > 1$ generate larger cones (sometimes full cones), and thus can lead to a reduced prediction accuracy.

Aperture-additive rotations are not commutativity because, for every $\delta \neq 0$, we can find a cone C such that $R\langle\theta, 1, \delta\rangle(C)$ is either \mathbf{C}_{\top} or \mathbf{C}_{\perp} (which are absorbing elements). However, under certain conditions, commutativity holds. For example, if u, v , and w are three positive real numbers with $u + v \leq w$, and $R_1 = R\langle\theta_1, 1, u\rangle$, $R_2 = R\langle\theta_2, 1, v\rangle$, $R_3 = R\langle\theta_1, 1, -u\rangle$, $R_4 = R\langle\theta_2, 1, -v\rangle$, then

$$\begin{aligned} (R_1 \circ R_2)(C(0, 2\pi - w)) &= (R_2 \circ R_1)(C(0, 2\pi - w)), \\ (R_3 \circ R_4)(C(0, w)) &= (R_4 \circ R_3)(C(0, w)). \end{aligned}$$

Under these circumstances, there is a bias that favors the inference of axioms $\exists R_1(\exists R_2.C) \equiv \exists R_2(\exists R_1.C)$. Hence, both aperture-multiplicative and aperture-additive rotations are biased.

Distributivity of the existential over the disjunction. To guarantee that the operations are closed, we needed to extend the embeddings to multicones. However, multicones do not satisfy a basic property. The full multicone MC_{\top} can be understood as two possible multicones, namely $\text{MC}(\{\mathbf{C}_{\top}\})$ and $\text{MC}(\{C(0, \pi), C(\pi, 2\pi i)\})$. However, the application of the rotation $R^{\mathcal{E}} = R\langle\theta, \frac{1}{2}, 0\rangle$ returns different multicones for both representations of the multicone. Since a multicone $\text{MC}(\{C_1, C_2\})$ is equal to $C_1 \cup C_2$, the tautology $\exists R.(C \sqcup D) \equiv \exists R.C \sqcup \exists R.D$ does not hold in the embedding space. Indeed, it only holds for positive aperture-additive rotations.

Proposition 3. The axiom $\exists R.(C \sqcup D) \equiv \exists R.C \sqcup \exists R.D$ holds in multicone embeddings for rotations $R\langle\theta, \gamma, \delta\rangle$ if $\gamma = 1$ and $\delta \geq 0$.

5 Expressing Logical Patterns

In this section, we present how logical patterns are expressed in the multicone embedding space. Given the relation names $r_1, r_2, r_3 \in \mathbf{R}$, correspond to the following SROI^{-} axioms:

$$\begin{array}{ll} r_1 \sqsubseteq r_2 & (\text{role containment}), \quad r_1 \circ r_2 \sqsubseteq r_3 \quad (\text{composition}), \\ \text{Trans}(r_1) & (\text{transitivity}), \quad r_1 \equiv r_2^{-} \quad (\text{inverse}), \\ \text{Sym}(r_1) & (\text{symmetry}), \quad \text{Asym}(r_1) \quad (\text{asymmetry}). \end{array}$$

Role containment pattern. Given two rotations $R^{\mathcal{E}} = R\langle\theta_1, \gamma_1, 0\rangle$ and $S^{\mathcal{E}} = R\langle\theta_2, \gamma_2, 0\rangle$ with $\theta_1 \neq \theta_2$ and $\gamma_1 \neq \gamma_2$, it never hold that $R^{\mathcal{E}} \sqsubseteq S^{\mathcal{E}}$. Indeed, we can always consider a cone with a sufficiently small aperture angle (e.g., a singleton cone) to show that this inclusion does not hold. However, for rotations $R^{\mathcal{E}} = R\langle\theta_1, 1, \delta_1\rangle$ and $S^{\mathcal{E}} = R\langle\theta_2, 1, \delta_2\rangle$ the inclusion can hold, even with distinct values for θ_1 and θ_2 .

Proposition 4. Given two rotations $R^{\mathcal{E}} = R\langle\theta_1, 1, \delta_1\rangle$ and $S^{\mathcal{E}} = R\langle\theta_2, 1, \delta_2\rangle$, the axiom $R \sqsubseteq S$ holds if $\theta_2 + \frac{\delta_2}{2} \geq \theta_1 + \frac{\delta_1}{2}$ and $\theta_2 - \frac{\delta_2}{2} \leq \theta_1 - \frac{\delta_1}{2}$.

Proposition 5. Given two rotations $R^{\mathcal{E}} = R\langle\theta, \gamma_1, \delta_1\rangle$ and $S^{\mathcal{E}} = R\langle\theta, \gamma_2, \delta_2\rangle$, the axiom $R \sqsubseteq S$ holds if $\gamma_1 \leq \gamma_2$ and $\delta_1 \leq \delta_2$.

Composition and transitivity patterns. The rules for composition and transitivity are a corollary of Proposition 4 and Proposition 5. For space limitations, we omit the conditions for the Transitivity pattern. The conditions can be obtained from the equivalence of axioms $\text{Trans}(R)$ and $R \circ R \sqsubseteq R$.

Corollary 6. Given three rotations $R_1^{\mathcal{E}} = R\langle\theta_1, 1, \delta_1\rangle$, $R_2^{\mathcal{E}} = R\langle\theta_2, 1, \delta_2\rangle$, and $R_3^{\mathcal{E}} = R\langle\theta_3, 1, \delta_3\rangle$, the axiom $R_1 \circ R_2 \sqsubseteq R_3$ holds if $|\theta_2 - (\theta_1 + \theta_2)| \leq \delta_3 - (\delta_1 + \delta_2)$, $\delta_1 \geq 0$, $\delta_2 \geq 0$, and $\delta_3 \geq 0$.

Corollary 7. Given $R_1^{\mathcal{E}} = R\langle\theta, \gamma_1, \delta_1\rangle$, $R_2^{\mathcal{E}} = R\langle\theta, \gamma_2, \delta_2\rangle$, and $R_3^{\mathcal{E}} = R\langle\theta, \gamma_3, \delta_3\rangle$, the axiom $R_1 \circ R_2 \sqsubseteq R_3$ holds if $\gamma_1 \gamma_2 \leq \gamma_3$, $\delta_1 + \delta_2 \leq \delta_3$, $\delta_1 \geq 0$, and $\delta_2 \geq 0$.

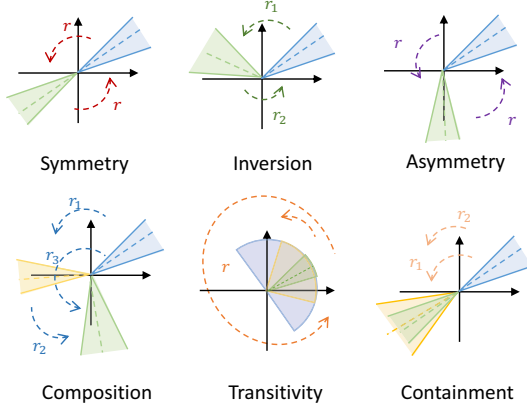


Figure 3: Logical patterns captured by AConE in a single dimension

Inverse pattern. Because inverses cannot generally be found, we cannot define inverse patterns for relations. However, inverse patterns still can be inferred under the circumstances where the commutativity is held, which is discussed in Section 4.

Symmetry and asymmetry patterns. The symmetry pattern remains consistent when a single cone is rotated twice, resulting in the same cone. This happens when two rotations complete the circle. Similarly, the asymmetry pattern holds when two rotations do not coincide with a circle.

Proposition 8. Given a rotation $R^\varepsilon = R\langle\theta, \gamma, \delta\rangle$ with $\gamma \geq 1$ and $\rho \geq 0$, if there is a natural number k such that $2\theta + \frac{\gamma\rho + \rho}{2} \geq 2k\pi$ and $2\theta - \frac{\gamma\rho + \rho}{2} \leq 2k\pi$, then the axiom $\text{Sym}(r)$ holds.

6 Tree-form Query Answering with AConE

To accommodate the learning of logical patterns in query answering over KGs, we propose a new model, AConE, which embeds KGs according to the multicone embedding (Definition 7). Our approach distinguishes itself from ConE [29], which represents cone embeddings using two-dimensional vectors in real space. We reframe these embeddings in the complex plane. By doing so, we introduce an inductive bias that facilitates the learning of logical patterns. This bias is achieved by combining the embeddings with relational rotations on cones, enhancing the model's ability to capture complex relationships and patterns. In AConE, the embedding of a tree-form query $q = C(x)$ is parameterized by a pair $\mathbf{q} = (\mathbf{h}_U, \mathbf{h}_L) \in \mathbb{C}^d$, where d is the embedding dimension, $|\mathbf{h}_U|_2 = 1$, $|\mathbf{h}_L|_2 = 1$ with $|\cdot|_2$ being the L2 norm. The vectors \mathbf{h}_U and \mathbf{h}_L represent the counterclockwise upper and lower boundaries of the cone, such that

$$\mathbf{h}_U = e^{i\theta_U} = e^{i(\theta_{ax} + \frac{\theta_{ap}}{2})}, \quad \mathbf{h}_L = e^{i\theta_L} = e^{i(\theta_{ax} - \frac{\theta_{ap}}{2})}, \quad (1)$$

where $\theta_{ax} \in [-\pi, \pi)^d$ represents the angle of the symmetry axis of the cone and $\theta_{ap} \in [0, 2\pi]^d$ represents the cone aperture.

For each index j with $1 \leq j \leq d$, we write $[v]_j$ for the value of the j -th component of a vector \mathbf{v} . The i -th component of \mathbf{q} represents the proper cone $C([\theta_L]_j, [\theta_U]_j)$ if $[\theta_U]_j > 0$ and $[\theta_{ap}]_j < 2\pi$, the full cone C_\top if $[\theta_{ap}]_j = 2\pi$, and the singleton cone $C_{\{\theta_L\}_j}$ if $\theta_{ap} = 0$. Notice that we do not provide an encoding for the empty cone because interesting queries are satisfiable. The queries are thus modeled as multidimensional cones according to Definition 7, and each entity answering the query is modeled as a vector $\mathbf{h}^* = e^{i\theta^*}$ in the cone parametrized with the query embedding.

6.1 Geometric Operators

To answer tree-form queries, AConE translates the concept constructors \exists , \sqcap , \sqcup , and \neg into corresponding geometric operators, i.e. *relational rotation* \mathcal{P}_\exists , *intersection* \mathcal{P}_\sqcap , *union* \mathcal{P}_\sqcup , and *negation* \mathcal{P}_\neg , in the complex plane. AConE derives the final query embedding by executing these geometric operators along the computation graph of the query. We next describe these geometric operators.

Relational Rotation. Given a set of entities $\mathcal{S} \subset \mathbf{I}$ and a relation $r \in \mathbf{R}$, the transformation operator for r selects the entities $\mathcal{S}' = \{e' \in \mathbf{I} : G \models r(e, e'), e \in \mathcal{S}\}$. To this end, we model r with a vector $\mathbf{r} = (\mathbf{r}_U, \mathbf{r}_L) \in \mathbb{C}^{2 \times d}$ encoding a counterclockwise rotation on query embeddings about the complex plane origin such that $|\mathbf{r}_U| = 1$ and $|\mathbf{r}_L| = 1$. The rotation \mathbf{r} transforms a query embedding $\mathbf{q} = (\mathbf{h}_U, \mathbf{h}_L)$ into a query embedding $\mathcal{P}_\exists(\mathbf{q}, \mathbf{r}) = \mathbf{q}' = (\mathbf{h}_U', \mathbf{h}_L')$ where

$$\mathbf{h}_U' = \mathbf{h}_U \circ \mathbf{r}_U, \quad \mathbf{h}_L' = \mathbf{h}_L \circ \mathbf{r}_L, \quad (2)$$

where \circ is the Hadamard (element-wise) product (i.e., for each component j , $[\mathbf{h}_U']_j = [\mathbf{h}_U]_j \cdot [\mathbf{r}_U]_j$ and $[\mathbf{h}_L']_j = [\mathbf{h}_L]_j \cdot [\mathbf{r}_L]_j$). If a pair $([\mathbf{r}_U]_j, [\mathbf{r}_L]_j)$ is seen as a cone, with axis angle $[\theta_{ax}, \mathbf{r}]_j$ and aperture angle $[\theta_{ap}, \mathbf{r}]_j$, then the rotation corresponds to the aperture-additive rotation $R\langle[\theta_{ax}, \mathbf{r}]_j, 1, [\theta_{ap}, \mathbf{r}]_j\rangle$ (see Definition 6). That is, the axis and aperture angles of \mathbf{q}' are $\theta_{ax}' = \theta_{ax} + \theta_{ax, \mathbf{r}}$, and $\theta_{ap}' = \theta_{ap} + \theta_{ap, \mathbf{r}}$. Figure 3 illustrates how AConE models the logical patterns described in Section 5, using the example on a single dimension.

Intersection. The intersection of a set of query embeddings $Q = \{q_1, \dots, q_n\}$, denoted $\mathcal{P}_\sqcap(Q) = \mathbf{q}'$ is defined with the permutation-invariant functions SemanticAverage and CardMin [29], which calculate the center θ_{ax}' and the aperture θ_{ap}' of the resulting cone encoded by \mathbf{q}' as follows:

$$\begin{aligned} \theta_{ax}' &= \text{SemanticAverage}(\{q_k\}_{k=1}^n), \\ \theta_{ap}' &= \text{CardMin}(\{q_k\}_{k=1}^n). \end{aligned} \quad (3)$$

SemanticAverage is expected to approximate the axis of the cone resulting from the intersection of the input cones. CardMin predicts the aperture θ_{ap}' of the intersection set such that $[\theta_{ap}']_j$ should be no larger than the aperture of any cone $q_i \in Q$, since the intersection set is the subset of all input entity sets. Both functions, SemanticAverage and CardMin, use a neural network to improve the results. We extend the details of them in Appendix D in the supplementary material [9].

Union. Given a set of query embeddings $Q = \{q_1, \dots, q_n\}$, the union operator $\mathcal{P}_\sqcup(Q)$ returns the set Q . Intuitively, for each k where $1 \leq k \leq d$, the set Q encodes a multicone embedding for the query $(q_1 \sqcup \dots \sqcup q_n)(x)$, which represents the multicone $\text{MC}(C_1, \dots, C_n)$ where C_i is the cone encoded in $[q_i]_k$.

Notice that multicones cannot be further used as input of other operations because all AConE operations are applied on cones. To compute queries, we thus follow [18], translating queries into disjunctive normal form, so we only perform the disjunction operator in the last step in the computation graph.

Negation. The negation of a cone $\mathbf{q} = (\mathbf{h}_U, \mathbf{h}_L)$, denoted $\mathcal{P}_\neg(\mathbf{q})$, is the cone \mathbf{q}' that contains the entities in the complement of \mathbf{q} , that is, $\mathbf{q}' = (\mathbf{h}_L, \mathbf{h}_U)$.

6.2 Optimization

Learning Objective. Given a set of training samples, our goal is to minimize the distance between the query embedding q and the answer entity vector h^* , while maximizing the distance between q and its negative samples. Thus, we define our training objective as

$$L = -\log \sigma(\gamma - d_{\text{comb}}(h^*, q)) - \frac{1}{k} \sum_{i=1}^k \log \sigma(d(h'_i, q) - \gamma), \quad (4)$$

where d_{comb} is the *combined distance* defined below, γ is a margin, h^* is a positive entity, and h'_i is the i -th negative entity, k is the number of negative samples, and σ represents the sigmoid function.

Combined Distance. Inspired by [29, 18], the distance between q and h^* is defined as a combination of inside and outside distances, $d_{\text{comb}}(q, h^*) = d_o(q, h^*) + \lambda d_i(q, h^*)$:

$$d_o(q, h^*) = \min\{\|h_U - h^*\|_1, \|h_L - h^*\|_1\}, \quad (5)$$

$$d_i(q, h^*) = \min\{\|h_{ax} - h^*\|_1, \|h_U - h_{ax}\|_1\}, \quad (6)$$

where $\|\cdot\|_1$ is the L_1 norm, h_{ax} represents the cone center, and $\lambda \in (0, 1)$. Note that d_{comb} can only be used for measuring the distance between a single query embedding and an answer entity vector. Since we represent the disjunctive queries in Disjunctive Normal Form as a set of query embeddings Q , the distance between the answer vector and such set of embeddings is the minimum distance:

$$d_{\text{comb}}(Q, h^*) = \min\{d_{\text{comb}}(q, h^*) : q \in Q\}. \quad (7)$$

7 Experiments and Analysis

In this section, we answer the following research questions with experimental statistics and corresponding case analyses. **RQ1:** how well does AConE improve query answering on incomplete KGs? **RQ2:** How is the improvement of results related to the capturing of logical patterns?

7.1 Experimental Setup

Dataset and Query Structure. For a fair comparison with the baseline models, we use the same query structures and logical query datasets from NELL-QA [27] and WN18RR-QA [10], and the open-sourced framework created by [17] for logical query answering tasks. Figure 4 illustrates all query structures in experiments. We trained our model using 10 specific query structures and tested it on all 14 query structures to evaluate the generalization ability of our model regarding unseen query structures. More experimental details are in the supplementary material [9]. Our code is available at [7].

Evaluation Metrics. We use Mean Reciprocal Rank (MRR) as the evaluation metric. Given a sample of queries Q , MRR represents the average of the reciprocal ranks of results, $MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{\text{rank}_i}$.

7.2 RQ1: how well does AConE improve query answering on incomplete KGs?

We evaluate AConE on two logical query-answering benchmark datasets that include a variety of complex logical patterns: NELL-QA created by [17] and WN18RR-QA by [10]. AConE is compared with various query embedding models that also define query regions as vectors, geometries, or distributions, including GQE [6], Query2Box (Q2B) [18], BetaE [17], ConE [29], and LinE [10]. To ensure a fair

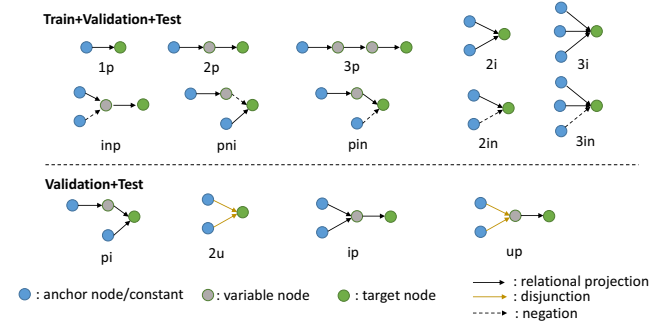


Figure 4: Fourteen types of queries used in the experiments. p represents an edge to another entity labeled with a relation name r (i.e., an operation $\exists r$), i represents the intersection (\cap), u represents the disjunction (\cup), and n represents the negation (\neg).

comparison, we indeed referenced the results of all baseline models from corresponding papers. In the case of ConE's performance on the WN18RR dataset, we conducted our own experiments due to the absence of previously reported results. This involved rerunning ConE with an optimized search for hyperparameters to generate the relevant data.

Main Results. Table 2 summarizes the performance of all methods on answering various query types. Compared with baselines that can only model queries without negation. Overall, AConE (aperture-additive) outperforms baseline methods on the majority of query types in Figure 4 while achieving competitive results on the others. Specifically, AConE consistently achieves improvements on non-negation queries on all datasets. Furthermore, our model brings notable improvement over baseline models on the WN18RR-QA dataset, where the average accuracy of AConE is 18.35% higher than that of the baseline models.

On the other hand, we observe that AConE performs closely to the baseline models on answering query types involving negation on dataset NELL-QA, in spite of its consistently good performance on WN18RR-QA. This may be due to two factors: firstly, handling negation is a challenging research question for complex logical query-answering tasks. When it comes to negation queries, all the current models show inferior performance compared to their performance on non-neg queries. Secondly, modeling negation as the complement leads to bias in prediction and high uncertainty in a large number of answers for negation queries. Thus, we leave the task of further improving our model on negation queries for future research.

7.3 RQ2: How is the improvement of results related to the capturing of logical patterns?

To investigate the influence of logical patterns learning on the query-answering model, we provide an ablation analysis considering both the model and data:

Model Perspective. We compared the performance of three AConE's variants: AConE (Base), AConE (aperture-multiplicative) and AConE (aperture-additive). They utilize alternative relational rotating transformation strategies for capturing the logical patterns while keeping other modules unchanged. Note that AConE (Base) represents AConE model with any neural relational transformation module. This configuration operates similarly to ConE in terms of its foundational approach. Consequently, we used the results of ConE as a direct representation of "AConE (Base)." This decision was based on their operational similarity, ensuring a fair and coherent presentation of our findings.

Table 2: MRR results (%) of AConE, LinE (the results of LinE are only available on WN18RR-QA), ConE, BETAE, Q2B, and GQE on answering tree-form ($\exists, \wedge, \vee, \neg$) queries on datasets NELL-QA, and WN18RR-QA. The best statistic is highlighted in bold, while the second best is highlighted in underline.

Dataset	Model	1p	2p	3p	2i	3i	pi	ip	2u	up	2in	3in	inp	pin	pni
WN18RR-QA	GQE	18.0	4.5	2.7	19.3	23.9	9.9	10.6	2.3	3.5	-	-	-	-	-
	Q2B	22.4	4.6	2.3	25.6	41.2	13.2	11.0	2.9	3.4	-	-	-	-	-
	BetaE	44.1	9.8	3.8	57.2	76.2	32.6	17.9	7.5	5.3	12.7	59.9	5.1	4.0	7.4
	LinE	45.1	12.3	6.7	47.1	67.1	24.8	14.7	8.4	6.9	12.5	60.8	7.3	5.2	<u>7.7</u>
	ConE	<u>46.8</u>	<u>14.5</u>	<u>9.3</u>	<u>59.0</u>	<u>83.9</u>	<u>33.6</u>	<u>18.7</u>	<u>10.0</u>	<u>9.8</u>	<u>13.9</u>	<u>61.8</u>	<u>10.6</u>	<u>7.3</u>	<u>7.6</u>
	AConE	50.9	17.6	9.9	70.5	89.0	38.9	29.6	18.4	14.0	18.3	70.1	13.4	7.6	10.4
NELL-QA	GQE	33.1	12.1	9.9	27.3	35.1	18.5	14.5	8.5	9.0	-	-	-	-	-
	Q2B	42.7	14.5	11.7	34.7	45.8	23.2	17.4	12.0	10.7	-	-	-	-	-
	BetaE	53.0	13.0	11.4	37.6	47.5	24.1	14.3	12.2	8.5	5.1	7.8	10.0	3.1	3.5
	ConE	<u>53.1</u>	<u>16.1</u>	<u>13.9</u>	<u>40.0</u>	<u>50.8</u>	26.3	<u>17.5</u>	<u>15.3</u>	<u>11.3</u>	5.7	8.1	10.8	3.5	3.9
	AConE	54.5	17.7	14.4	41.9	53.0	<u>26.1</u>	20.7	16.5	12.8	<u>5.2</u>	7.7	9.4	<u>3.2</u>	<u>3.7</u>

Table 3: Ablation study of AConE on NELL dataset.

Model	1p	2p	3p	2i	3i	pi	ip	2u	up
AConE (Base)	53.1	16.1	13.9	40.0	50.8	26.3	17.5	15.3	11.3
AConE (Ap-Mul)	51.3	16.6	13.8	38.4	48.4	18.9	19.5	14.7	12.1
AConE (Ap-Add)	54.5	17.7	14.4	41.9	53.0	26.1	20.7	16.5	12.8

As Table 3 shows, the performance variation among AConE’s variants highlights the impact of relation transformation on the model’s ability to capture logical patterns. Moreover, the outperformance of AConE confirms the efficiency of logical patterns in query reasoning tasks.

Data Perspective. To more effectively examine the specific impact of AConE on queries involving logical patterns, a detailed analysis was conducted on the NELL query answering dataset. We categorize the test dataset into five categories based on the relations involved in the queries. For each query in subgroups ¹ *Inverse*, *Symmetry*, *Composition*, *Containment* and *Transitivity*, there is at least one relation encompassing the corresponding logical pattern. Table 7 and Appendix B in the supplementary material [9] provide more details about these subgroups of queries and elaborate the classification process. The category *Others* corresponds to queries that do not involve any of these logical patterns. Figure 5 shows the average performances of AConE and neural baseline model (ConE) on these subgroups. It is observed that AConE outperforms the neural baseline model on queries that have logical patterns, especially inverse relations. However, AConE does not generalize as well to queries that were not influenced by logical patterns compared to the baseline model. This analysis supports our hypothesis that AConE is better suited for capturing logical patterns in tree-form query answering tasks than purely neural models, with fewer model parameters and a simpler model structure, which is elaborated in section 8.

8 Analysis on Model Parameters

The acquisition of logical patterns not only enhances the model performance but also alleviates the reliance on excessively parameterized neural networks in existing methods. Table 4 summarizes the number of parameters and the average performance of our method AConE and other baseline models on non-negation queries, which all methods can handle. It shows that AConE has the second-fewest number of parameters among these models, though it achieves much better performances than the baseline models.

¹ Note that *Asymmetry* is not included here because such type of relations is missing from the selected rules mined from NELL.

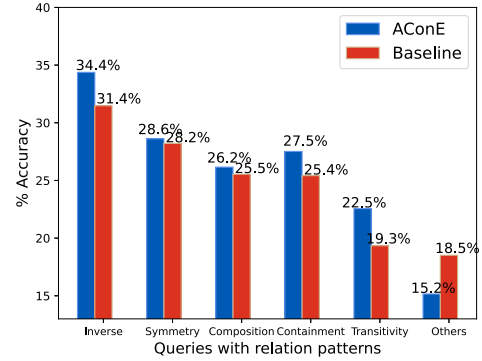


Figure 5: Average performances of AConE and Baseline model over query subgroups with different logical patterns.

Table 4: Number of parameters in AConE and other baseline models.

Model	AVG MRR	Parameters
BetaE[17]	28.26	57,574,000
GQE[6]	10.52	52,290,400
ConE[29]	31.73	44,010,401
Q2B[18]	14.06	26,306,000
Ours	37.64	36,325,601

9 Conclusion and Future Work

In this work, we investigate the important yet understudied impact of logical pattern inference on the query-answering task. We propose a new embedding method, the multicone embedding, and study its algebraic structure and capability to express SROI² ontologies. We show some limitations of the method, like breaking the tautology $\exists R.(C \sqcap D) \equiv \exists R.C \sqcup \exists R.D$, the bias favoring the commutativity, and the (general) non-existence of inverse rotations.

Our practical method, AConE, is motivated by the multicone algebra, but it has differences. First, multicones consisting of multiple cones are difficult to manage due to the unconstrained number of parameters they require. Hence, we limit the representation of queries to single multicones (i.e., cones). Since this simplification makes the cone intersection no longer closed, we define it as a neural operation returning a single cone. We limit rotations to be aperture-additive because they outperform aperture-multiplicative rotations (see Table 3).

Although the experiment datasets had no concept names, multicone embeddings support concept names. We use these datasets to make a fair comparison with the existing methods, which do not support concept names. Since most real datasets include concept names and concept names are a key component of knowledge representation, we plan to extend our research to datasets with concept names. Furthermore, we think that AConE can be extended for knowledge graphs including also non-empty RBoxes and TBoxes.

Ethics Statement

The authors declare that we have no conflicts of interest. This article does not contain any studies involving business and personal data.

Acknowledgements

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under the DFG Germany's Excellence Strategy – EXC 2120/1 – 390831618, and the DFG Excellence Strategy – EXC 2075 – 390740016. We acknowledge the support by the Stuttgart Center for Simulation Science (SimTech). The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Yunjie He. Mojtaba Nayyeri is funded by the BMBF ATLAS project. This work was partially supported by the European Commission funded projects enRichMyData(101070284), GraphMassivizer(101093202), and Dome 4.0 (953163).

References

- [1] SPARQL 1.1 Query Language. Technical report, W3C, 2013. URL <http://www.w3.org/TR/sparql11-query>.
- [2] E. Arakelyan, D. Daza, P. Minervini, and M. Cochez. Complex query answering with neural link predictors. In *ICLR*. OpenReview.net, 2021.
- [3] F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003. ISBN 0-521-78176-0.
- [4] K. D. Bollacker, C. Evans, P. K. Paritosh, T. Sturge, and J. Taylor. Freebase: a collaboratively created graph database for structuring human knowledge. In J. T. Wang, editor, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, pages 1247–1250. ACM, 2008. doi: 10.1145/1376616.1376746. URL <https://doi.org/10.1145/1376616.1376746>.
- [5] A. Bordes, N. Usunier, A. García-Durán, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In *NIPS*, pages 2787–2795, 2013.
- [6] W. L. Hamilton, P. Bajaj, M. Zitnik, D. Jurafsky, and J. Leskovec. Embedding logical queries on knowledge graphs. In *NeurIPS*, pages 2030–2041, 2018.
- [7] Y. He, D. Hernandez, M. Nayyeri, B. Xiong, Y. Zhu, E. Kharlamov, and S. Staab. Code and data for “generating SROI² ontologies via knowledge graph query embedding learning.”, 2022. Available at <https://github.com/RoyaHe/AConE>.
- [8] Y. He, M. Nayyeri, B. Xiong, Y. Zhu, E. Kharlamov, and S. Staab. Can pattern learning enhance complex logical query answering? In *ISWC (Posters/Demos/Industry)*, volume 3632 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2023.
- [9] Y. He, D. Hernandez, M. Nayyeri, B. Xiong, Y. Zhu, E. Kharlamov, and S. Staab. Generating SROI² ontologies via knowledge graph query embedding learning., 2024. URL <https://arxiv.org/abs/2407.09212>. Full version of this paper.
- [10] Z. Huang, M. Chiang, and W. Lee. Line: Logical query reasoning over hierarchical knowledge graphs. In *KDD*, pages 615–625. ACM, 2022.
- [11] M. Jackermeier, J. Chen, and I. Horrocks. Box²el: Concept and role box embeddings for the description logic EL++. *CoRR*, abs/2301.11118, 2023. doi: 10.48550/ARXIV.2301.11118. URL <https://doi.org/10.48550/arXiv.2301.11118>.
- [12] B. Kotnis, C. Lawrence, and M. Niepert. Answering complex queries in knowledge graphs with bidirectional sequence encoders. In *AAAI*, pages 4968–4977. AAAI Press, 2021.
- [13] X. V. Lin, R. Socher, and C. Xiong. Multi-hop knowledge graph reasoning with reward shaping. In *EMNLP*, pages 3243–3253. Association for Computational Linguistics, 2018.
- [14] M. Nayyeri, C. Xu, Y. Yaghoobzadeh, S. Vahdati, M. M. Alam, H. S. Yazdi, and J. Lehmann. Loss-aware pattern inference: A correction on the wrongly claimed limitations of embedding models. In *PAKDD (3)*, volume 12714 of *Lecture Notes in Computer Science*, pages 77–89. Springer, 2021.
- [15] M. Nickel, V. Tresp, and H. Kriegel. A three-way model for collective learning on multi-relational data. In *ICML*, pages 809–816. Omnipress, 2011.
- [16] Ö. L. Özçep, M. Leemhuis, and D. Wolter. Cone semantics for logics with negation. In C. Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 1820–1826. ijcai.org, 2020. doi: 10.24963/IJCAI.2020/252.
- [17] H. Ren and J. Leskovec. Beta embeddings for multi-hop logical reasoning in knowledge graphs. In *NeurIPS*, 2020.
- [18] H. Ren, W. Hu, and J. Leskovec. Query2box: Reasoning over knowledge graphs in vector space using box embeddings. In *ICLR*. OpenReview.net, 2020.
- [19] H. Ren, M. Galkin, M. Cochez, Z. Zhu, and J. Leskovec. Neural graph reasoning: Complex logical query answering meets graph databases, 2023.
- [20] F. M. Suchanek, G. Kasneci, and G. Weikum. Yago: a core of semantic knowledge. In C. L. Williamson, M. E. Zurko, P. F. Patel-Schneider, and P. J. Shenoy, editors, *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, pages 697–706. ACM, 2007. doi: 10.1145/1242572.1242667. URL <https://doi.org/10.1145/1242572.1242667>.
- [21] Z. Sun, Z. Deng, J. Nie, and J. Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *ICLR (Poster)*. OpenReview.net, 2019.
- [22] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard. Complex embeddings for simple link prediction. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2071–2080. JMLR.org, 2016.
- [23] D. Vrandečić and M. Krötzsch. Wikidata: a free collaborative knowledgebase. *Commun. ACM*, 57(10):78–85, 2014.
- [24] Z. Wang, J. Zhang, J. Feng, and Z. Chen. Knowledge graph embedding by translating on hyperplanes. In *AAAI*, pages 1112–1119. AAAI Press, 2014.
- [25] B. Xiong, N. Potyka, T. Tran, M. Nayyeri, and S. Staab. Faithful embeddings for \mathcal{EL}^{++} knowledge bases. In *ISWC*, volume 13489 of *Lecture Notes in Computer Science*, pages 22–38. Springer, 2022.
- [26] B. Xiong, S. Zhu, M. Nayyeri, C. Xu, S. Pan, C. Zhou, and S. Staab. Ultrahyperbolic knowledge graph embeddings. In *KDD*, pages 2130–2139. ACM, 2022.
- [27] W. Xiong, T. Hoang, and W. Y. Wang. DeepPath: A reinforcement learning method for knowledge graph reasoning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 564–573. Copenhagen, Denmark, Sept. 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1060. URL <https://aclanthology.org/D17-1060>.
- [28] B. Yang, W. Yih, X. He, J. Gao, and L. Deng. Embedding entities and relations for learning and inference in knowledge bases. In *ICLR (Poster)*, 2015.
- [29] Z. Zhang, J. Wang, J. Chen, S. Ji, and F. Wu. Cone: Cone embeddings for multi-hop reasoning over knowledge graphs. In *NeurIPS*, pages 19172–19183, 2021.
- [30] Y. Zhu, N. Potyka, B. Xiong, T.-K. Tran, M. Nayyeri, E. Kharlamov, and S. Staab. Approximating probabilistic inference in statistical el with knowledge graph embeddings, 2024. URL <https://arxiv.org/abs/2407.11821>.
- [31] Z. Zhu, M. Galkin, Z. Zhang, and J. Tang. Neural-symbolic models for logical queries on knowledge graphs. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, pages 27454–27478. PMLR, 2022.
- [32] Z. Zhu, M. Galkin, Z. Zhang, and J. Tang. Neural-symbolic models for logical queries on knowledge graphs. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, pages 27454–27478. PMLR, 2022.