

A Non-commutative Bilinear Model for Answering Path Queries in Knowledge Graphs

Katsuhiko Hayashi^{†,*}

katsuhiko-h@sanken.osaka-u.ac.jp

Masashi Shimbo^{‡,*}

shimbo@is.naist.jp

[†]Osaka University
Suita, Osaka, Japan

[‡]NAIST
Ikoma, Nara, Japan

*Riken AIP
Chuo-ku, Tokyo, Japan

Abstract

Bilinear diagonal models for knowledge graph embedding (KGE), such as DistMult and ComplEx, balance expressiveness and computational efficiency by representing relations as diagonal matrices. Although they perform well in predicting atomic relations, composite relations (relation paths) cannot be modeled naturally by the product of relation matrices, as the product of diagonal matrices is commutative and hence invariant with the order of relations. In this paper, we propose a new bilinear KGE model, called BlockHolE, based on block circulant matrices. In BlockHolE, relation matrices can be non-commutative, allowing composite relations to be modeled by matrix product. The model is parameterized in a way that covers a spectrum ranging from diagonal to full relation matrices. A fast computation technique is developed on the basis of the duality of the Fourier transform of circulant matrices.

1 Introduction

Large-scale knowledge graphs (Nickel et al., 2016a) are indispensable resources for knowledge-intensive applications such as question answering, dialog systems, and distantly supervised relation extraction. A knowledge graph is a collection of *triplets* (s, r, o) representing the fact that (binary) relation r holds between subject entity s and object entity o . Although efforts continue to enrich existing knowledge graphs with more facts, many facts are still missing (Nickel et al., 2016a). Knowledge graph completion (KGC) aims to automatically detect missing facts in an *incomplete* knowledge graph, and has become an active field of research in recent years.

Knowledge graph embedding (KGE) is a promising approach to KGC. It embeds entities and relations in vector space, and defines a *scoring function*

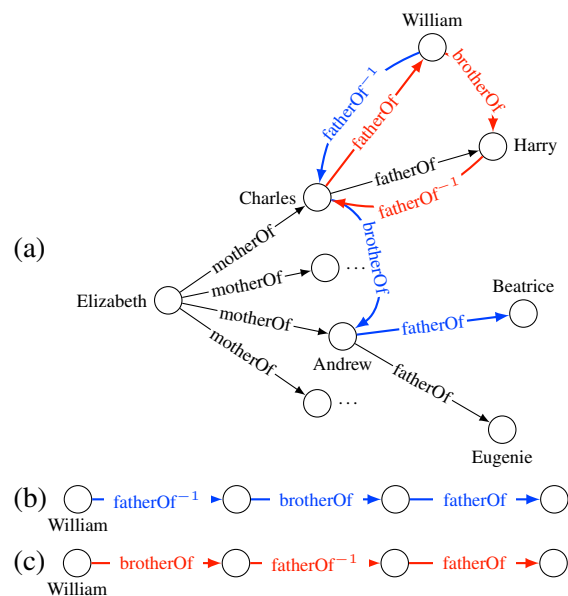


Figure 1: (a) A knowledge graph and (b,c) two relation paths starting from William.

$\phi(s, r, o)$ to evaluate the degree of factuality of a given triplet (s, r, o) in terms of vector operations.

Bilinear KGE models are a popular choice for a scoring function, along with those based on translation and neural networks. RESCAL (Nickel et al., 2011) adopts a generic bilinear form as the scoring function, given by $\phi_{\text{RESCAL}}(s, r, o) = \mathbf{e}_s^T \mathbf{R}_r \mathbf{e}_o$. In this formula, $\mathbf{e}_s, \mathbf{e}_o$ are the n -dimensional vector embeddings of entities s and o , respectively, and \mathbf{R}_r is the $n \times n$ matrix embedding of relation r . Some of the more recent models have constrained the relation matrices to be diagonal. DistMult (Yang et al., 2015) and ComplEx (Trouillon et al., 2016) are two such diagonal models. HolE (Nickel et al., 2016b) does not use diagonal relation matrices, but has been shown (Hayashi and Shimbo, 2017) to be isomorphic to ComplEx. These models have a smaller number of parameters than RESCAL, making them less prone to overfit-

ting, and the performance is usually better.

While all these models were designed with a specific task of KGC in mind, i.e., computing the factuality of triplets, another important task on knowledge graphs was pursued by Guu et al. (2015) and Lin et al. (2015). This latter task, called *path query answering (path QA)*, is to answer composite queries that consist of a cascade of relations, as opposed to an atomic relation. See Figure 1 for instance. A query “Is Beatrice a child of a paternal uncle of William?” can be answered by predicting the truth value of the triplet (William, fatherOf⁻¹/brotherOf/fatherOf, Beatrice) where fatherOf⁻¹/brotherOf/fatherOf is a binary relation not present in the knowledge graph as a relation (edge) label but is composed of a cascade of three atomic relations.¹ Composite queries are also called *path queries*, as they can be represented as paths in a knowledge graph; see, e.g., the blue line in Figure 1(a). Notice however that some of the edges in the path may be missing due to the incompleteness of the knowledge graph; even in such circumstances, the model must ideally be able to answer path queries correctly.

Guu et al. (2015) extended the existing KGE approaches to path QA. For example, to answer a general path query $(s, r_1 / \dots / r_k, o)$ with RESCAL, a composite relation $r_1 / \dots / r_k$ is modeled by matrix product $\mathbf{R}_{r_1} \cdots \mathbf{R}_{r_k}$, and the score for the given query is modeled by $\mathbf{e}_s^T \mathbf{R}_{r_1} \cdots \mathbf{R}_{r_k} \mathbf{e}_o$. This formulation is also applicable to DistMult and ComplEx, which use diagonal relation matrices. In diagonalized models, however, relation matrices are *commutative*, in the sense that $\mathbf{R}_r \mathbf{R}_{r'} = \mathbf{R}_{r'} \mathbf{R}_r$ for any pair of relations r, r' .

Commutativity of relation matrices was not recognized as an issue in the past research because the main focus was on predicting the truth value of atomic triplets. However, when path queries are concerned, commutativity poses a problem. Consider, for example, a relation sequence

fatherOf⁻¹/brotherOf/fatherOf

and its permutation

brotherOf/fatherOf⁻¹/fatherOf.

Although these are two distinct paths (cf. Figure 1(b, c)), in bilinear models with commutative relation matrices, they are represented by the same

¹We regard inverse relations (e.g., fatherOf⁻¹) also as atomic relations.

Symbol	Description
\mathbb{R}, \mathbb{C}	sets of real/complex numbers
$[\mathbf{v}]_j$	j th component of vector \mathbf{v}
$[\mathbf{M}]_{jk}$	(j, k) -component of matrix \mathbf{M}
\mathbf{M}^T	transpose of \mathbf{M}
$\overline{\mathbf{M}}$	conjugate of \mathbf{M}
\odot	componentwise (Hadamard) product
$*$	circular convolution
\star	circular correlation
$\text{Re}(x)$	real part of complex number x
$\text{diag}(\mathbf{v})$	diagonal matrix with main diagonal \mathbf{v}
$\text{circ}(\mathbf{v})$	circulant matrix determined by \mathbf{v}
$\langle \mathbf{x}, \mathbf{y}, \mathbf{z} \rangle$	sum of the componentwise products of $\mathbf{x}, \mathbf{y}, \mathbf{z}$
\mathbf{F}	discrete Fourier matrix
\mathcal{E}	set of entities
\mathcal{R}	set of relations
\mathcal{F}	set of observed facts (triplets)
\mathcal{F}^*	set of ground truth facts
$\mathcal{G}(\mathcal{F})$	Knowledge graph induced by facts \mathcal{F}

Table 1: List of symbols. See Secs. 2 and 3 for detail.

product of relation matrices, which thereby makes the truth values of these permuted queries indistinguishable by their scores.

Drawing on the observation above, this paper proposes a new KGE model called BlockHolE, wherein relations are represented by block circulant matrices. This makes relation matrices non-commutative, and thus it does not suffer from the issues arising from commutativity, yet in general manages to reduce the number of parameters compared with RESCAL. It can be interpreted as a generalization of HolE and ComplEx, and also subsumes RESCAL as an extreme case. We report experimental results in both path and atomic QA tasks.

2 Notation and preliminaries

We first introduce symbols and notation used in this paper, followed by some preliminaries on circulant matrices, circular convolution, correlation, and Fourier transform. The summary of symbols can be found in Table 1.

Let \mathbb{R} be the set of reals, and \mathbb{C} be the set of complex numbers. Let $[\mathbf{v}]_j$ denote the j th component of vector \mathbf{v} , and let $[\mathbf{M}]_{jk}$ the (j, k) element of matrix \mathbf{M} . For a complex number z , vector \mathbf{z} , and matrix \mathbf{Z} , let \bar{z} , $\bar{\mathbf{z}}$, and $\overline{\mathbf{Z}}$ denote their complex conjugate, respectively.

Let \mathbf{x}, \mathbf{y} , and \mathbf{z} be n -dimensional (real or complex) vectors. Let $\text{diag}(\mathbf{x})$ denote an $n \times n$ diagonal matrix with the main diagonal components given by \mathbf{x} . We write $\mathbf{x} \odot \mathbf{y}$ to denote the componentwise product of \mathbf{x} and \mathbf{y} ; i.e., $\mathbf{x} \odot \mathbf{y} = \text{diag}(\mathbf{x})\mathbf{y}$, or

$[\mathbf{x} \odot \mathbf{y}]_i = [\mathbf{x}]_i [\mathbf{y}]_i$, $i = 1, \dots, n$. We also write $\langle \mathbf{x}, \mathbf{y}, \mathbf{z} \rangle = \mathbf{x}^T \text{diag}(\mathbf{y}) \mathbf{z} = \sum_{i=1}^n [\mathbf{x}]_i [\mathbf{y}]_i [\mathbf{z}]_i$.

For n -dimensional real vectors² $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, $\mathbf{x} * \mathbf{y}$ and $\mathbf{x} \star \mathbf{y}$ denote *circular convolution* and *circular correlation*, respectively defined by

$$[\mathbf{x} * \mathbf{y}]_i = \sum_{j=1}^n [\mathbf{x}]_{i-j+1} [\mathbf{y}]_j,$$

$$[\mathbf{x} \star \mathbf{y}]_i = \sum_{j=1}^n [\mathbf{x}]_{j-i+1} [\mathbf{y}]_j, \quad i = 1, \dots, n,$$

where vector indices that do not fall in the range $1, \dots, n$ must be interpreted by $0 = n, -1 = n - 1, \dots, -n + 1 = 1$.

For n -dimensional real vector $\mathbf{v} \in \mathbb{R}^n$, let

$$\text{circ}(\mathbf{v}) = \begin{bmatrix} [\mathbf{v}]_1 & [\mathbf{v}]_n & \dots & [\mathbf{v}]_3 & [\mathbf{v}]_2 \\ [\mathbf{v}]_2 & [\mathbf{v}]_1 & [\mathbf{v}]_n & & [\mathbf{v}]_3 \\ \vdots & [\mathbf{v}]_2 & [\mathbf{v}]_1 & \ddots & \vdots \\ [\mathbf{v}]_{n-1} & & \ddots & \ddots & [\mathbf{v}]_n \\ [\mathbf{v}]_n & [\mathbf{v}]_{n-1} & \dots & [\mathbf{v}]_2 & [\mathbf{v}]_1 \end{bmatrix}$$

be an operation that converts a vector \mathbf{v} to a circulant matrix of size $n \times n$.

A circulant matrix $\text{circ}(\mathbf{v}) \in \mathbb{R}^{n \times n}$ can be diagonalized as $\text{circ}(\mathbf{v}) = \mathbf{F}^{-1} \text{diag}(\mathbf{F}\mathbf{v})\mathbf{F}$, where $\mathbf{F} \in \mathbb{C}^{n \times n}$ is the discrete Fourier matrix of order n . Also, circular convolution and correlation can be written in terms of $\text{circ}(\cdot)$: $\mathbf{x} * \mathbf{y} = \text{circ}(\mathbf{x})\mathbf{y}$, and $\mathbf{x} \star \mathbf{y} = \text{circ}(\mathbf{x})^T \mathbf{y}$. It follows that

$$\mathbf{x} * \mathbf{y} = \mathbf{F}^{-1}(\mathbf{F}\mathbf{x} \odot \mathbf{F}\mathbf{y}), \quad (1)$$

$$\mathbf{x} \star \mathbf{y} = \mathbf{F}^{-1}(\overline{\mathbf{F}\mathbf{x}} \odot \mathbf{F}\mathbf{y}). \quad (2)$$

These equations imply that circular convolution and correlation can be computed in time $O(n \log n)$ using the fast Fourier transform (FFT).

3 Knowledge graph embedding using bilinear maps

A *knowledge graph* is a labeled multigraph $(\mathcal{E}, \mathcal{R}, \mathcal{F})$, where \mathcal{E} is the set of entities (or vertices), \mathcal{R} is the set of relation labels (or edge labels), and $\mathcal{F} \subset \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ defines the observed instances of binary relations over entities (or labeled edges). An item $(s, r, o) \in \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ is called a *triplet*, with s and o called its *subject* and *object*, respectively. For every entity in $e \in \mathcal{E}$, it is

²Generally, circular convolution, circular correlation, and circulant matrices are defined over \mathbb{C}^n . However, in this paper, it suffices to define them over \mathbb{R}^n .

assumed that \mathcal{F} contains at least one triplet (s, r, o) with $s = e$ or $o = e$; likewise, for every relation in $r \in \mathcal{R}$, \mathcal{F} is assumed to contain at least one triplet (s, r, o) . Because \mathcal{F} determines the sets \mathcal{E} and \mathcal{R} of entities and relations, we write $\mathcal{G}(\mathcal{F})$ to denote the knowledge graph $(\mathcal{E}, \mathcal{R}, \mathcal{F})$ determined by \mathcal{F} .

Aside from observed triplets \mathcal{F} , we also assume the presence of a set $\mathcal{F}^* \subset \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ of (*ground truth*) facts, which is a strict superset of \mathcal{F} , i.e., $\mathcal{F} \subset \mathcal{F}^*$. Thus, \mathcal{F}^* is not fully observable.

3.1 Knowledge graph completion

Knowledge graph completion (KGC) is the task of identifying the set of ground truth facts \mathcal{F}^* from observed facts $\mathcal{F} \subset \mathcal{F}^*$ (or equivalently, $\mathcal{G}(\mathcal{F}^*)$ from $\mathcal{G}(\mathcal{F})$).

A popular approach to KGC is to design a *scoring function* $\phi(s, r, o)$ quantifying how likely a triplet (s, r, o) is true. This scoring function is learned from the observed triplets \mathcal{F} , in a way that it generalizes well to unobserved triplets $\mathcal{F}^* \setminus \mathcal{F}$; i.e., the score must be high for both observed and unobserved facts, and it must be low for nonfactual triplets.

In *knowledge graph embedding* (KGE)-based approaches to KGC, the scoring function $\phi(s, r, o)$ is defined in terms of the embeddings of entities and relations; i.e., s , r , and o are embedded as objects in a vector space, and ϕ is defined in terms of some operations over these objects.

3.2 Bilinear models for knowledge graph embedding

Below, we describe some of the popular KGE models that use bilinear maps to define scoring functions.

3.2.1 RESCAL

RESCAL (Nickel et al., 2011) provides the most general form of bilinear scoring function.

$$\phi_{\text{RESCAL}}(s, r, o) = \mathbf{e}_s^T \mathbf{R}_r \mathbf{e}_o, \quad (3)$$

where $\mathbf{e}_s, \mathbf{e}_o \in \mathbb{R}^n$ are the vector embeddings of entities s and o , respectively, and $\mathbf{R}_r \in \mathbb{R}^{n \times n}$ is the matrix representing relation r . Thus, n^2 parameters are required per relation, which is not only a computational burden but also the cause of overfitting during training (Kazemi and Poole, 2018).

3.2.2 DistMult

DistMult (Yang et al., 2015) is a model obtained by restricting the relation matrices \mathbf{R}_r of RESCAL

to diagonal; i.e., $\mathbf{R}_r = \text{diag}(\mathbf{w}_r)$, $\mathbf{w}_r \in \mathbb{R}^n$. The scoring function is thus

$$\begin{aligned}\phi_{\text{DistMult}}(s, r, o) &= \mathbf{e}_s^\top \text{diag}(\mathbf{w}_r) \mathbf{e}_o \\ &= \langle \mathbf{w}_r, \mathbf{e}_s, \mathbf{e}_o \rangle.\end{aligned}\quad (4)$$

Although the number of parameters is reduced considerably, the scoring function (4) is symmetric with respect to the entities, i.e., $\phi_{\text{DistMult}}(s, r, o) = \phi_{\text{DistMult}}(o, r, s)$. This is a severe limitation because most real-world relations are non-symmetric.

3.2.3 ComplEx: Complex embedding

The *complex embedding* (ComplEx) (Trouillon et al., 2016) represents entities and relations as n -dimensional vectors as in DistMult, but their components are complex-valued.

The scoring function of ComplEx is given by

$$\begin{aligned}\phi_{\text{ComplEx}}(s, r, o) &= \text{Re}(\mathbf{e}_s^\top \text{diag}(\mathbf{w}_r) \overline{\mathbf{e}_o}) \\ &= \text{Re}(\langle \mathbf{w}_r, \mathbf{e}_s, \overline{\mathbf{e}_o} \rangle),\end{aligned}$$

where $\mathbf{e}_s, \mathbf{e}_o, \mathbf{w}_r \in \mathbb{C}^n$ are the embeddings of s, o , and r , respectively. The number of parameters in ComplEx is $2n|\mathcal{E}| + 2n|\mathcal{R}|$, and the score is computable in time linear in the dimension of vector space. Unlike DistMult, ComplEx can model non-symmetric relations, since $\phi(s, r, o) \neq \phi(o, r, s)$ in general.

3.2.4 HolE: Holographic embedding

The *holographic embedding* (HolE) (Nickel et al., 2016b) uses circular correlation to define a scoring function

$$\phi_{\text{HolE}}(s, r, o) = \mathbf{w}_r^\top (\mathbf{e}_s \star \mathbf{e}_o), \quad (5)$$

where $\mathbf{w}_r, \mathbf{e}_s, \mathbf{e}_o \in \mathbb{R}^n$ are n -dimensional real vectors representing relation r , and entities s and o , respectively. HolE has only n parameters per relation, and it can model non-symmetric relations since $\phi_{\text{HolE}}(s, r, o) \neq \phi_{\text{HolE}}(o, r, s)$ in general. Computing circular correlation requires $O(n \log n)$ time if FFT is employed. Eq. (5) is not a bilinear form, but it has been shown (Hayashi and Shimbo, 2017) that HolE is isomorphic to ComplEx, and thus any model in HolE can be converted to an equivalent model in ComplEx, and vice versa.

4 Path question answering over a knowledge graph

4.1 Path query answering

Let \mathcal{F}^* be the set of ground truth facts, and let $\mathcal{G}(\mathcal{F}^*) = (\mathcal{E}, \mathcal{R}, \mathcal{F}^*)$ be its induced knowledge

graph. For k relations $r_1, \dots, r_k \in \mathcal{R}$, we call $r_1 / \dots / r_k$ a *relation path* of length k . When $k = 1$, the relation path is *atomic*; otherwise, it is *composite*. Let $s, o \in \mathcal{E}$. We say a *path query* $(s, r_1 / \dots / r_k, o)$ holds (or “is true”) in $\mathcal{G}(\mathcal{F}^*)$ (or with respect to \mathcal{F}^*) if

$$\begin{aligned}\exists e_1, \dots, e_{k-1} \in \mathcal{E} \quad \forall j = 1, \dots, k \\ (e_{j-1}, r_j, e_j) \in \mathcal{F}^*,\end{aligned}$$

where $e_0 = s$ and $e_k = o$. *Path query answering* (path QA) is the task of predicting the truth value of path queries with respect to the unobserved set \mathcal{F}^* of ground truth facts, when its incomplete subset $\mathcal{F} \subset \mathcal{F}^*$ is only available. In other words, we want to predict that $(s, r_1 / \dots / r_k, o)$ is true if a path from s to o exists in $\mathcal{G}(\mathcal{F}^*)$, although some of the edges that constitute the path may be missing in the observed graph $\mathcal{G}(\mathcal{F})$.

For atomic path queries (i.e., those with length $k = 1$), path QA reduces to that of knowledge graph completion introduced in Section 3.1. Thus, it is natural to address general path QA by extending the scoring function $\phi(s, r, o)$ of KGC methods so that composite relation $r_1 / \dots / r_k$ is allowed in place of atomic relation r ; i.e., by defining $\phi(s, r_1 / \dots / r_k, o)$. Previous work (Guu et al., 2015) explored this direction, which is also pursued in the rest of this paper.

4.2 Issues in existing KGE models applied to path QA

We now discuss the extension of existing bilinear KGE models to path QA. We begin with RESCAL, which is the most general among existing bilinear models. In RESCAL, if we assume $\mathbf{R}_r^\top \mathbf{e}_s \approx \mathbf{e}_o$ for true triplets (s, r, o) , we can model path QA as computing

$$\begin{aligned}\phi_{\text{RESCAL}}(s, r_1 / \dots / r_k, o) \\ = \mathbf{e}_s^\top \mathbf{R}_{r_1} \cdots \mathbf{R}_{r_k} \mathbf{e}_o.\end{aligned}\quad (6)$$

As seen in this formula, a composite relation is represented by the product of the matrices for atomic relations (Guu et al., 2015).

Likewise, DistMult and ComplEx can also be used for path QA, by computing

$$\begin{aligned}\phi_{\text{DistMult}}(s, r_1 / \dots / r_k, o) \\ = \mathbf{e}_s^\top \text{diag}(\mathbf{w}_{r_1}) \cdots \text{diag}(\mathbf{w}_{r_k}) \mathbf{e}_o,\end{aligned}$$

and

$$\begin{aligned} \phi_{\text{CompEx}}(s, r_1 / \dots / r_k, o) \\ = \text{Re}(\mathbf{e}_s^T \text{diag}(\mathbf{w}_{r_1}) \cdots \text{diag}(\mathbf{w}_{r_k}) \overline{\mathbf{e}_o}), \end{aligned}$$

respectively. However, because diagonal matrices are commutative, the score of $(s, r_1 / \dots / r_k, o)$ is equal to any path query in which r_1, \dots, r_k are permuted, such as $(s, r_k / r_{k-1} / \dots / r_1, o)$. That is, because $\phi(s, r_1 / \dots / r_k, o) = \phi(s, r_k / r_{k-1} / \dots / r_1, o)$, their truth values cannot be distinguished by the magnitude of scores. More recent bilinear models such as ANALOGY³ (Liu et al., 2017) and Simple (Kazemi and Poole, 2018) also represent relations by diagonal matrices, and thus they can only model commutative relation paths. Moreover, for Simple, which represents subject and object entities in different vector spaces, it is not clear how it can be applied to path QA.

In the translation-based model TransE (Bordes et al., 2013), the scoring function is given by⁴

$$\phi_{\text{TransE}}(s, r, o) = -\|\mathbf{e}_s + \mathbf{w}_r - \mathbf{e}_o\|_2^2. \quad (7)$$

Guu et al. (2015) extended this function for a path query by

$$\begin{aligned} \phi_{\text{TransE}}(s, r_1 / \dots / r_k, o) \\ = -\|\mathbf{e}_s + \mathbf{w}_{r_1} + \dots + \mathbf{w}_{r_k} - \mathbf{e}_o\|_2^2. \quad (8) \end{aligned}$$

Thus, a composite relation is represented as the sum of the embedding vectors for its constituent atomic relations. Unfortunately, Eq. (8) is also invariant with the permutation of relations r_1, \dots, r_k , and their order is not respected.

5 Knowledge graph embedding with block circulant matrices

5.1 BlockHoIE

In this section, we propose a bilinear KGE model suitable for path QA. In this model, the relation matrices are non-commutative. It thus respects the order of relations in a path query. Further, it has a smaller number of parameters than RESCAL in

³ We categorize ANALOGY as a diagonal model because each 2×2 block diagonal element of its relation matrices can be substituted by a single equivalent complex-valued component.

⁴ The original TransE defines a *penalty* function, which gives a smaller value if a triplet is more likely to be true. We thus changed the sign to make it a scoring function in Eq. (7).

general. To be specific, our model constrains the relation matrices to be block circulant.

A matrix is *block circulant* if it can be written in the form

$$\begin{bmatrix} \mathbf{W}^{(11)} & \dots & \mathbf{W}^{(1b)} \\ \vdots & \ddots & \vdots \\ \mathbf{W}^{(b1)} & \dots & \mathbf{W}^{(bb)} \end{bmatrix}, \quad (9)$$

where each $\mathbf{W}^{(ij)} = \text{circ}(\mathbf{w}^{(ij)})$, $i, j = 1, \dots, b$, is a circulant matrix determined by $\mathbf{w}^{(ij)} \in \mathbb{R}^m$. Thus, if the dimension of the matrix in Eq. (9) is $n \times n$, we have $n = bm$. A block circulant matrix is non-commutative when $b \geq 2$; i.e., for two block circulant matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{bm \times bm}$, $\mathbf{A} \neq \mathbf{B}$, $\mathbf{A}\mathbf{B} \neq \mathbf{B}\mathbf{A}$ in general.

Substituting a block circulant matrix of Eq. (9) for matrix \mathbf{R}_r in the bilinear scoring function (Eq. (3)) yields

$$\begin{aligned} \phi_{\text{BlockHoIE}}(s, r, o) &= \mathbf{e}_s^T \mathbf{R}_r \mathbf{e}_o \\ &= \underbrace{[\mathbf{e}_s^{(1)T} \dots \mathbf{e}_s^{(b)T}]}_{\mathbf{e}_s^T} \overbrace{\begin{bmatrix} \mathbf{W}_r^{(11)} & \dots & \mathbf{W}_r^{(1b)} \\ \vdots & \ddots & \vdots \\ \mathbf{W}_r^{(b1)} & \dots & \mathbf{W}_r^{(bb)} \end{bmatrix}}^{\mathbf{R}_r} \overbrace{\begin{bmatrix} \mathbf{e}_o^{(1)} \\ \vdots \\ \mathbf{e}_o^{(b)} \end{bmatrix}}^{\mathbf{e}_o}, \quad (10) \end{aligned}$$

where $\mathbf{e}_s^{(i)}, \mathbf{e}_o^{(i)} \in \mathbb{R}^m$, and $\mathbf{W}_r^{(ij)} = \text{circ}(\mathbf{w}_r^{(ij)}) \in \mathbb{R}^{m \times m}$, $i, j = 1, \dots, b$. Recall that $n = bm$, and thus $\mathbf{e}_s, \mathbf{e}_o \in \mathbb{R}^n$, $\mathbf{R}_r \in \mathbb{R}^{n \times n}$. Using equalities $\mathbf{x}^T(\mathbf{y} * \mathbf{z}) = \mathbf{y}^T(\mathbf{x} \star \mathbf{z})$ (Nickel et al., 2016b) and $\text{circ}(\mathbf{x})\mathbf{y} = \mathbf{x} * \mathbf{y}$ to rewrite Eq. (10), we have

$$\begin{aligned} \phi_{\text{BlockHoIE}}(s, r, o) &= [\mathbf{e}_s^{(1)T} \dots \mathbf{e}_s^{(b)T}] \begin{bmatrix} \sum_{j=1}^b \mathbf{w}_r^{(1j)} * \mathbf{e}_o^{(j)} \\ \vdots \\ \sum_{j=1}^b \mathbf{w}_r^{(bj)} * \mathbf{e}_o^{(j)} \end{bmatrix} \\ &= \sum_{i=1}^b \mathbf{e}_s^{(i)T} \left(\sum_{j=1}^b \mathbf{w}_r^{(ij)} * \mathbf{e}_o^{(j)} \right) \\ &= \sum_{i,j=1}^b \mathbf{w}_r^{(ij)T} \left(\mathbf{e}_o^{(j)} \star \mathbf{e}_s^{(i)} \right). \quad (11) \end{aligned}$$

We call this model *BlockHoIE*, after the fact that it reduces to HoIE when $b = 1$; cf. Eq. (5). Also, BlockHoIE is identical to b -dimensional RESCAL when $m = 1$ (or equivalently $b = n$).

The number of parameters in BlockHoIE is $bm|\mathcal{E}| + b^2m|\mathcal{R}|$ (or $n|\mathcal{E}| + bn|\mathcal{R}|$), and naive

computation of Eq. (11) takes time $O(b^2 m \log m)$ using FFT. However, we can make this computation faster by exploiting the duality of the Fourier transform, as shown below.

5.2 Fast computation in complex space

Using a similar technique used by Hayashi and Shimbo (2017) to show the equivalence of ComplEx and HolE, we can eliminate Fourier transform to speed up the computation of BlockHolE scores. We first rewrite Eq. (11) as follows:

$$\begin{aligned} \phi_{\text{BlockHolE}}(s, r, o) &= \sum_{i,j=1}^b \mathbf{w}_r^{(ij)\text{T}} (\mathbf{e}_o^{(j)} \star \mathbf{e}_s^{(i)}) \\ &= \sum_{i,j=1}^b \mathbf{w}_r^{(ij)\text{T}} \mathbf{F}^{-1} (\overline{\mathbf{F}\mathbf{e}_o^{(j)}} \odot \mathbf{F}\mathbf{e}_s^{(i)}) \\ &= \frac{1}{m} \sum_{i,j=1}^b (\overline{\mathbf{F}\mathbf{w}_r^{(ij)}})^{\text{T}} (\overline{\mathbf{F}\mathbf{e}_o^{(j)}} \odot \mathbf{F}\mathbf{e}_s^{(i)}) \\ &= \frac{1}{m} \sum_{i,j=1}^b \text{Re}(\langle \overline{\mathbf{F}\mathbf{w}_r^{(ij)}}, \overline{\mathbf{F}\mathbf{e}_s^{(i)}} \rangle, \mathbf{F}\mathbf{e}_o^{(j)}), \end{aligned}$$

where \mathbf{F} is the discrete Fourier matrix. Here we used Eq. (1) to derive the second equation, and $\mathbf{F}^{-1} = (1/m)\overline{\mathbf{F}}^{\text{T}}$ to derive the third. Defining complex vectors $\overline{\mathbf{w}_r^{(ij)}} = (1/m)\overline{\mathbf{F}\mathbf{w}_r^{(ij)}}$, $\mathbf{e}'_s^{(i)} = \overline{\mathbf{F}\mathbf{e}_s^{(i)}}$, and $\mathbf{e}'_o^{(j)} = \overline{\mathbf{F}\mathbf{e}_o^{(j)}}$ yields

$$\phi_{\text{BlockHolE}}(s, r, o) = \sum_{i,j=1}^b \text{Re} \left(\langle \overline{\mathbf{w}_r^{(ij)}}, \mathbf{e}'_s^{(i)}, \overline{\mathbf{e}'_o^{(j)}} \rangle \right). \quad (12)$$

On the basis of Eq. (12), we train $\mathbf{e}'_k^{(i)} \in \mathbb{C}^m$ directly in complex space (i.e., the Fourier domain) instead of $\mathbf{e}_k^{(j)} \in \mathbb{R}^m$ and use it as the vector embedding of entity k , for all $k \in \mathcal{E}$; similarly, $\overline{\mathbf{w}'_r^{(ij)}} \in \mathbb{C}^m$ is directly trained in complex space to represent relation $r \in \mathcal{R}$. The number of parameters in this model is $2n|\mathcal{E}| + 2bn|\mathcal{R}|$, and Eq. (12) can be computed in $O(bn)$ time. Typically, we set $b \ll n$. For instance, in the experiment of Section 6, we set $b = 2$ and $m = 50$, and thus $n = bm = 100$. In this case, factor b is negligible and the computational complexity is linear in n .

		WN11	FB13
Base	Train	112,581	316,232
	Valid	2,609	5,908
	Test	10,544	23,733
Path	Train	2,129,539	6,266,058
	Valid	11,277	27,163
	Test-Deduction	24,749	77,883
	Test-Induction	21,828	31,674

Table 2: Dataset provided by Guu et al. (2015).

5.3 Modeling path QA

BlockHolE can be used in path QA as follows. First, for any $\ell \in \mathcal{E}$ and $r \in \mathcal{R}$, let

$$\begin{aligned} \mathbf{e}'_\ell{}^{\text{T}} &= [\mathbf{e}'_\ell^{(1)\text{T}} \cdots \mathbf{e}'_\ell^{(b)\text{T}}], \\ \mathbf{W}'_r &= \begin{bmatrix} \text{diag}(\mathbf{w}'_r^{(11)}) & \cdots & \text{diag}(\mathbf{w}'_r^{(1b)}) \\ \vdots & \ddots & \vdots \\ \text{diag}(\mathbf{w}'_r^{(b1)}) & \cdots & \text{diag}(\mathbf{w}'_r^{(bb)}) \end{bmatrix}. \end{aligned}$$

Then, Eq. (12) can be rewritten as

$$\phi_{\text{BlockHolE}}(s, r, o) = \text{Re} \left(\mathbf{e}'_s{}^{\text{T}} \mathbf{W}'_r \overline{\mathbf{e}'_o} \right),$$

and we can compute the score of relation paths by

$$\begin{aligned} \phi_{\text{BlockHolE}}(s, r_1 / \cdots / r_k, o) &= \text{Re}(\mathbf{e}'_s{}^{\text{T}} \mathbf{W}'_{r_1} \cdots \mathbf{W}'_{r_k} \overline{\mathbf{e}'_o}). \end{aligned}$$

Since $\mathbf{W}'_{r_i} \mathbf{W}'_{r_j} \neq \mathbf{W}'_{r_j} \mathbf{W}'_{r_i}$ for $b \geq 2$, this scoring function respects the order of relations in $r_1 / \cdots / r_k$.

6 Experiments

In this section, we report the results of empirical evaluation investigating the commutativity property of bilinear KGE models on the path QA task. As expected, the proposed BlockHolE model, which uses non-commutative relation matrices, outperformed commutative bilinear KGE models.

6.1 Dataset and evaluation protocol

The comparison of KGE models was performed in two path QA tasks: (i) ranking and (ii) binary classification tasks.

6.1.1 Path QA ranking

For the path QA ranking task, we adopted the same protocol and dataset used by Guu et al. (2015). Table 2 shows the statistics of their dataset. The dataset consists of two parts, ‘‘Base’’ and ‘‘Path’’.

	WN11						FB13					
	Base		Deduction		Induction		Base		Deduction		Induction	
	P@10	MQ	P@10	MQ	P@10	MQ	P@10	MQ	P@10	MQ	P@10	MQ
DistMult	45.6	83.0	33.5	97.7	29.6	79.8	62.7	91.6	63.6	86.4	59.3	86.5
ComplEx	60.9	83.1	68.7	99.2	46.1	79.7	76.8	93.0	71.5	90.0	70.5	88.9
RESCAL	51.8	74.2	43.2	97.9	51.2	76.8	65.2	91.1	66.9	88.4	69.8	89.0
$b = 2, m = 25$	80.9	83.4	70.2	99.5	54.9	81.0	79.2	93.2	75.0	91.5	71.3	90.0
$b = 4, m = 25$	80.5	75.6	69.3	99.2	54.5	77.4	76.2	92.1	72.1	90.5	70.9	89.5

Table 3: Path QA ranking result: Comparing BlockHole ($b = 2, m = 25$ and $b = 4, m = 25$) to other bilinear models. The dimension of the embedding space for DistMult, ComplEx and RESCAL was set to $n = 50$ as the result of grid search.

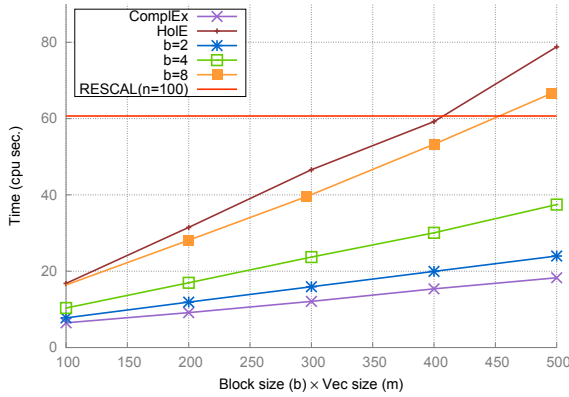


Figure 2: CPU run time per epoch of BlockHole on WN11 Base with single CPU thread.

The Base part only contains facts (i.e., path queries with $k = 1$), and thus it is essentially for evaluating KGC performance. Its training samples constitute the observed facts \mathcal{F} , and the facts in the entire Base part (training/validation/test sets) make the ground truth facts \mathcal{F}^* .

The Path part contains path queries sampled from the same \mathcal{F} and \mathcal{F}^* as the Base part. The test samples in the Path part is divided into “deduction” and “induction” sets. In the “deduction” set, test samples were sampled from the Base training graph $\mathcal{G}(\mathcal{F})$. By contrast, in the “induction” set, the test samples were chosen from the ground truth graph $\mathcal{G}(\mathcal{F}^*)$ such that none of them have a corresponding path in $\mathcal{G}(\mathcal{F})$. Thus, the “induction” set is intended to measure how well a model generalizes to unobserved paths, whereas the “deduction” set is to test its ability to faithfully encode the observed training graph.

At the time of evaluation, for each a test sample $(s, r_1/\dots/r_k, o)$, a *candidate set*

$$\mathcal{T}(r_k) = \{t : \exists h \in \mathcal{E} (h, r_k, t) \in \mathcal{F}^*\},$$

was first computed. In other words, the candidates

are the entities for which r_k (i.e., the last relation in the test query) takes as its object at least once in \mathcal{F}^* . Then, for each compared model, we made the ranking of the candidates entities in $\mathcal{T}(r_k)$ by the score $\phi(s, r_1/\dots/r_k, e)$, where ϕ is learned by the model from the training set.

The quality of the ranking was measured by two evaluation metrics: averaged mean quantile (MQ) and P@10 (percentage of correct answers ranked in the top 10). For $q = s/p$ where $p = r_1/\dots/r_k$, the correct answer set $\llbracket q \rrbracket$ is the set of all entities that can be reached from s by traversing p over $\mathcal{G}(\mathcal{F}^*)$. Formally, let $\llbracket s \rrbracket = \{s\}$, and the answer set can be recursively defined: $\llbracket q/r \rrbracket = \{t : \exists h \in \llbracket q \rrbracket, (h, r, t) \in \mathcal{F}^*\}$. With these definitions, MQ is computed by the following formula:

$$\frac{|\{o' | o' \in \mathcal{N}(q) : \phi(s, p, o') \leq \phi(s, p, o)\}|}{|\mathcal{N}(q)|}, \quad (13)$$

where $\mathcal{N}(q) = \mathcal{T}(r_k) \setminus \llbracket q \rrbracket$ is the set of incorrect answers. Eq. (13) cannot be computed for queries with which $\mathcal{T}(r_k) = \llbracket q \rrbracket$, and these queries were excluded from evaluation. For further details, see the original paper by Guu et al. (2015).

6.1.2 Path QA classification

In the path QA classification task, we simply report classification accuracy. After the scoring function ϕ was trained with logistic regression, a path query $q = (s, r_1/\dots/r_k, o)$ was classified as true if $\phi(q) \geq 0$, or false otherwise.

Since the test and validation sets of Path in Table 2 contain only correct queries, we sampled negative ones by the following procedure: For a correct query $q = (s, r_1/\dots/r_k, o)$ ($k \geq 2$), we generated its reverse relation path query $q' = (s, r_k/r_{k-1}/\dots/r_1, o)$. If q' does not exist in $\mathcal{G}(\mathcal{F}^*)$, we used it as a negative.

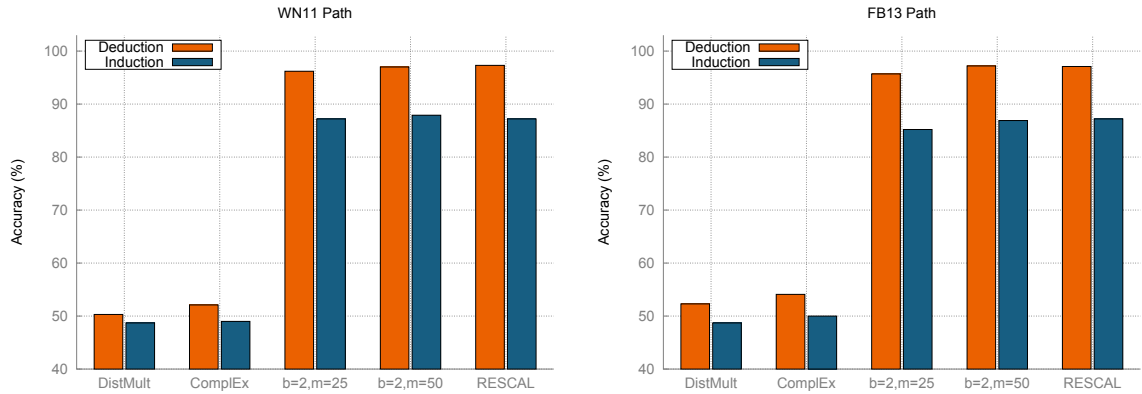


Figure 3: Path QA classification result comparing BlockHole ($b = 2, m = 25$ and $b = 2, m = 50$) to DistMult, ComplEx and RESCAL models (all with $n = 50$ as the result of grid search).

6.2 Experiment setup

We compared BlockHole with state-of-the-art bilinear KGE models: DistMult, RESCAL and ComplEx. We have implemented BlockHole in Java. BlockHole reduces to ComplEx when $b = 1$, and with the imaginary parts of parameters set to 0, it reduces to RESCAL when $m = 1$ and to DistMult when $b = 1$. For a fair run time comparison, however, we separately implemented RESCAL using jblas-1.2.4 for matrix computation. Through all experiments, we optimized the logistic loss with L2 regularization on the parameters Θ :

$$\min_{\Theta} \sum_{(q,y) \in \mathcal{D}} \log\{1 + \exp(-y\phi(q; \Theta))\} + \lambda \|\Theta\|_2^2$$

where y denotes the truth value of a query q in a training data \mathcal{D} . Given a correct query $q = (s, r_1 / \dots / r_k, o)$, we generated negative samples by replacing o with an entity randomly sampled from \mathcal{E} .

We selected the hyperparameters via grid search such that on the validation set they maximize classification accuracy in the path QA classification task and MQ in the path QA ranking task. For all models except BlockHole, all combinations of $\lambda \in \{0.0001, 0\}$, learning rate $\eta \in \{0.005, 0.01, 0.025, 0.05\}$, and the embedding size $n \in \{50, 100, 150, 200\}$ were tried during grid search. For BlockHole, all combinations of $(b, m) \in \{(2, 25), (2, 50), (2, 100), (4, 25), (4, 50), (8, 25)\}$, $\lambda \in \{0.0001, 0\}$ and $\eta \in \{0.005, 0.01, 0.025, 0.05\}$ were tried. The maximum number of training epochs was set to 500. The number of negatives generated per positive sample was 5 during training.

Label	Relation Path	ComplEx	BlockHole
+	*/parents/religion/*	96.7	100.0
-	*/religion/parents/*	3.3	100.0

Table 4: Classification accuracy on selected queries.

6.3 Results

6.3.1 Path QA ranking

Table 3 shows the results on the path QA ranking data. BlockHole outperforms other bilinear KGE models considerably both on deductive and inductive test settings. These results strongly suggest that BlockHole is more expressive in modeling path QA than DistMult and ComplEx, while effectively reducing redundant parameters in RESCAL which can cause model overfitting. Figure 2 shows the empirical scalability of BlockHole. When b is small, BlockHole scales linearly in the dimension $n = bm$ of the embedding space.

6.3.2 Path QA classification

Figure 3 shows the accuracy of path QA classification. DistMult and ComplEx were considerably worse than BlockHole and RESCAL for both WN11 and FB13. This result confirms our claim: The non-commutativity of relation matrices plays a critical role in modeling path QA. The performance of BlockHole ($b = 2, m = 25$) was comparable to that of RESCAL but the former was 12 times faster.

6.4 Analysis

The accuracies of BlockHole and RESCAL on the path QA classification task were markedly better than those of DistMult and ComplEx. We analyzed the results further. We extracted all queries from

$\mathcal{G}(\mathcal{F}^*)$ of FB13 that consist of an interpretable relation path $*/\text{parents}/\text{religion}/*$ where $*$ denotes “can match any relation path”. For such queries $(s, */\text{parents}/\text{religion}/*, o)$, we also generated meaningless queries $(s, */\text{religion}/\text{parents}/*, o)$ as negatives. Table 4 shows the classification accuracies of ComplEx and BlockHolE ($b = 2, m = 25$). The results clearly show that ComplEx cannot correctly answer the negative queries at all due to the lack of the non-commutative property.

7 Summary

In this paper, we have pointed out the problems of existing bilinear KGE models in path QA, and proposed a new model that overcomes these problems. This model, called BlockHolE, represents relations as block circulant matrices. As a result, it respects the order of relations in path queries, while enjoying linear-time computation of scoring functions when the number b^2 of blocks is sufficiently small. It generalizes HolE/ComplEx, and it can also be interpreted as an interpolation between RESCAL and HolE/ComplEx. Its effectiveness was shown empirically in path QA.

Our proposal can be useful in not only path QA but also many tasks such as associative rule mining (Yang et al., 2015), path regularization (Lin et al., 2015), and more complex QA (Hamilton et al., 2018), in which composite relations need to be embedded as a vector. Other future directions include reducing the increased parameters in the proposed block circulant matrices, such as by using multiplicative L1 regularization for ComplEx (Manabe et al., 2018).

Acknowledgments

We thank anonymous reviewers for helpful comments. This work was partially supported by JSPS Kakenhi Grant Numbers 19H04173, 18K11457, and 18H03288.

References

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 26 (NIPS)*, pages 2787–2795.

Kelvin Guu, John Miller, and Percy Liang. 2015. [Traversing knowledge graphs in vector space](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 318–327.

William L. Hamilton, Payal Bajaj, Marinka Zitnik, Dan Jurafsky, and Jure Leskovec. 2018. [Embedding logical queries on knowledge graphs](#). In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 2030–2041.

Katsuhiko Hayashi and Masashi Shimbo. 2017. [On the equivalence of holographic and complex embeddings for link prediction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 2: Short Papers, pages 554–559.

Seyed Mehran Kazemi and David Poole. 2018. [Simple embedding for link prediction in knowledge graphs](#). In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 4284–4295.

Yankai Lin, Zhiyuan Liu, Huan-Bo Luan, Maosong Sun, Siwei Rao, and Song Liu. 2015. [Modeling relation paths for representation learning of knowledge bases](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 705–714.

Hanxiao Liu, Yuexin Wu, and Yiming Yang. 2017. [Analogical inference for multi-relational embeddings](#). In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 2168–2178.

Hitoshi Manabe, Katsuhiko Hayashi, and Masashi Shimbo. 2018. [Data-dependent learning of symmetric/anti-symmetric relations for knowledge base completion](#). In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 3754–3761.

Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2016a. [A review of relational machine learning for knowledge graphs](#). *Proceedings of the IEEE*, 104(1):11–33.

Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016b. [Holographic embeddings of knowledge graphs](#). In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1955–1961.

Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. [A three-way model for collective learning on multi-relational data](#). In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 809–816.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 2071–2080.

Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. [Embedding entities and relations for learning and inference in knowledge bases](#). In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.