Expressive and Geometrically Interpretable Knowledge Graph Embedding (Extended Abstract)

Aleksandar Pavlović **Emanuel Sallinger**
Research Unit of Databases and Artificial Intelligence
TU Wien
Vienna, Austria
{aleksandar.pavlovic,emanuel.sallinger}@tuwien.ac.at

Motivation. A pressing challenge in artificial intelligence (AI) is connecting the interpretability of logic-based AI with the promising prediction performance of machine-learning-based AI. Knowledge graphs (KGs) are prominently subject to this challenge, as real-world KGs are highly incomplete [10], employing logic- or machine-learning-based completion approaches. KG embedding models (KGEs) that embed KGs into vector spaces show promising performance for knowledge graph completion (KGC) [9], i.e., predicting missing links. In the spirit of neuro-symbolic AI, geometric KGEs that embed entities and relations as geometric shapes in vector spaces allow for an intuitive geometric interpretation of their learned rules. However, significant challenges remain that need to be overcome. On the one hand, *composition* rules are ubiquitous for reasoning in KGs, as they infer links based on paths through the KG. However, existing KGEs capture an extremely limited notion of composition [12, 1, 5, 3], precisely *compositional definition* and not *general composition* (formally defined in Table 1). On the other hand, while most KGEs can learn compositional definition [2, 7, 12, 5] or hierarchy rules [11, 4, 8, 1] individually, they cannot capture both rules jointly (see Table 1).

Table 1: This table summarizes the inference capabilities of several KGEs. In particular, a ✓ indicates that a KGE can learn the specific logical rule, and an ✗ displays that it cannot learn the rule.

Logical Rule	ExpressivE	BoxE	RotatE	TransE	DistMult	ComplEx
Symmetry: $r_1(X,Y) \Rightarrow r_1(Y,X)$	✓	✓	✓	✗	✓	✓
Anti-symmetry: $r_1(X,Y) \Rightarrow \neg r_1(Y,X)$	✓	✓	✓	✓	✗	✓
Inversion: $r_1(X,Y) \Leftrightarrow r_2(Y,X)$	✓	✓	✓	✓	✗	✓
Comp. def.: $r_1(X,Y) \wedge r_2(Y,Z) \Leftrightarrow r_3(X,Z)$	✓	✗	✓	✓	✗	✗
Gen. comp.: $r_1(X,Y) \wedge r_2(Y,Z) \Rightarrow r_3(X,Z)$	✓	✗	✗	✗	✗	✗
Hierarchy: $r_1(X,Y) \Rightarrow r_2(X,Y)$	✓	✓	✗	✗	✓	✓
Intersection: $r_1(X,Y) \wedge r_2(X,Y) \Rightarrow r_3(X,Y)$	✓	✓	✓	✓	✗	✗
Mutual exclusion: $r_1(X,Y) \wedge r_2(X,Y) \Rightarrow \bot$	✓	✓	✓	✓	✓	✓

Problem and Contribution. The vast amount of research on KGEs that can learn composition [2, 7, 12, 5] and hierarchy [11, 8, 4, 1] rules emphasizes their importance for KGC. However, any KGE so far cannot (1) learn general composition rules, (2) *jointly* learn composition and hierarchy rules, and (3) provide a *geometric interpretation* of its learned rules. Facing these challenges, we present the results of our top 25% paper accepted at ICLR 2023 [6], of which this paper is an extended abstract. Specifically, we (i) present the *ExpressivE* model and its *virtual triple space*, which allows us to geometrically interpret rules learned by an ExpressivE embedding, (ii) prove that ExpressivE can learn any rule of Table 1, the first such KGE, and (iii) benchmark ExpressivE, revealing that it is competitive with contemporary KGEs, even significantly outperforming them on some datasets.

Background. KGs can be depicted as extensive collections of triples $r_i(e_h, e_t)$ drawn from a finite set of relations $r_i \in \mathbf{R}$ and entities $e_h, e_t \in \mathbf{E}$. Given a triple $r_i(e_h, e_t)$, e_h is referred to as its *head* and e_t as its *tail*. In the rest of the paper, we adhere to the standard definition for learning logical rules (termed *capturing*), as outlined in [7, 1, 6]. Essentially, this implies that a KGE captures a rule if a set of parameters exists such that the logical rule is learned *exactly* (i.e., the KGE predicts any triple inferrable by the rule) and *exclusively* (i.e., the KGE's predictions do not support any unwanted rule).

The First Austrian Symposium on AI, Robotics, and Vision (AIROV24).

ExpressivE represents entities $e_j \in \boldsymbol{E}$ as vectors $\boldsymbol{e_j} \in \mathbb{R}^d$ in the embedding space $\mathbb{R}^d$ and relations $r_i \in \boldsymbol{R}$ as hyper-parallelograms in the virtual triple space $\mathbb{R}^{2d}$ (see Figure 1a). For each arity position $p \in \{h, t\}$ of a relation r_i, ExpressivE assigns a *slope vector* $\boldsymbol{r_i^p} \in \mathbb{R}^d$, a *center vector* $\boldsymbol{c_i^p} \in \mathbb{R}^d$, and a *width vector* $\boldsymbol{d_i^p} \in (\mathbb{R}_{\geq 0})^d$, defining the hyper-parallelogram's boundaries. A triple $r_i(e_h, e_t)$ is true in an ExpressivE model if the embeddings of r_i, e_h, and e_t satisfy the inequalities below:

$$(\boldsymbol{e_h} - \boldsymbol{c_i^h} - \boldsymbol{r_i^t} \odot \boldsymbol{e_t})^{|.|} \preceq \boldsymbol{d_i^h}\,, \qquad (\boldsymbol{e_t} - \boldsymbol{c_i^t} - \boldsymbol{r_i^h} \odot \boldsymbol{e_h})^{|.|} \preceq \boldsymbol{d_i^t}$$

Notation. The notation $\boldsymbol{x}^{|.|}$ denotes the element-wise absolute value of a vector $\boldsymbol{x}$, $\odot$ denotes the Hadamard product, and $\preceq$ denotes the element-wise less-than-or-equal-to operator. Due to the intricate nature of interpreting this model in the embedding space $\mathbb{R}^d$, we introduce the *virtual triple space* $\mathbb{R}^{2d}$ to facilitate the interpretation of ExpressivE's parameters and inference capabilities.

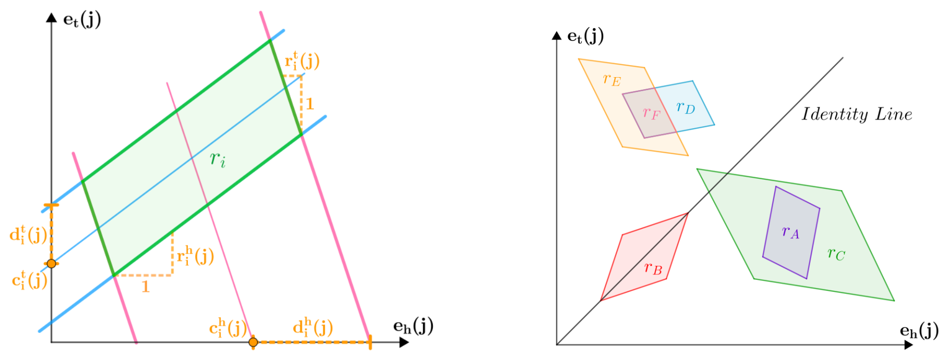

(a) (b)

Figure 1: The figures show (a) the interpretation of ExpressivE's relation parameters (orange dashed) as a parallelogram (green solid) in the j-th correlation subspace; (b) multiple relation parallelograms that satisfy the following rules in the j-th dimension: r_B (Symmetry); r_A, r_D, r_E, and r_F (Anti-Symmetry); $r_D = r_A^{-1}$ (Inversion); $r_A(X,Y) \Rightarrow r_C(X,Y)$ (Hierarchy); $r_D(X,Y) \wedge r_E(X,Y) \Rightarrow r_F(X,Y)$ (Intersection); and e.g., $r_A(X,Y) \wedge r_B(X,Y) \Rightarrow \bot$ (Mutual Exclusion).

Virtual Triple Space. To construct the virtual triple space $\mathbb{R}^{2d}$, the head $\boldsymbol{e_h}$ and tail embeddings $\boldsymbol{e_t}$ are concatenated. We refer to the 2-dimensional sub-space of $\mathbb{R}^{2d}$, formed from the j-th dimension of $\boldsymbol{e_h}$ and $\boldsymbol{e_t}$, as the *j-th correlation subspace* since it visualizes this dimension's captured logical rules. As depicted in Figure 1a, the relation parameters can be interpreted as a hyper-parallelogram in $\mathbb{R}^{2d}$. With these concepts in place, we proceed to analyze ExpressivE's theoretical capabilities.

Theorem 1. *ExpressivE is fully expressive, i.e., for any arbitrary graph G over $\boldsymbol{R}$ and $\boldsymbol{E}$, there is an ExpressivE embedding with finite dimensionality d (in particular, d in $O(|\boldsymbol{E}| * |\boldsymbol{R}|)$) capturing G.*

Logical Rules. Theorem 2 shows that ExpressivE captures common logical rules in KGE literature [2, 7, 11, 8, 4, 1]. Figure 1b illustrates how ExpressivE embeddings with $d = 1$ capture various rules.

Theorem 2. *ExpressivE captures (a) symmetry, (b) anti-symmetry, (c) inversion, (d) hierarchy, (e) intersection, (f) mutual exclusion, (g) general composition, and (h) compositional definition.*

ExpressivE's Relatives. As delineated by [6], geometric KGEs can be categorized into three families: *Functional KGEs*, which embed relations as functions; *bilinear KGEs*, which embed relations as bilinear products; and *spatial KGEs*, which embed relations as regions. Notably, ExpressivE is the pioneering KGE that belongs to both the spatial and functional families. Among its closest relatives, BoxE [1] aligns with the spatial family, while RotatE [7] aligns with the functional family.

Space Efficiency. While RotatE and BoxE employ $(2|\boldsymbol{E}| + 2|\boldsymbol{R}|)d$, ExpressivE embeddings employ $(|\boldsymbol{E}| + 6|\boldsymbol{R}|)d$ parameters. Given that in the majority of graphs $|\boldsymbol{R}| \ll |\boldsymbol{E}|$ holds, it is noteworthy that ExpressivE halves the parameter count of BoxE and RotatE for a d-dimensional embedding.

Benchmark Results. In our final evaluation, we benchmarked ExpressivE on the standard KGC benchmarks WN18RR and FB15k-237. The results indicate that ExpressivE, despite having only *half* the number of parameters compared to BoxE and RotatE, excels within its own model family on FB15k-237, reaching competitive results with state-of-the-art KGEs. Notably, ExpressivE significantly *outperforms all* competing KGEs on WN18RR.

Acknowledgments and Disclosure of Funding

Financial support for this research has been provided by the Vienna Science and Technology Fund (WWTF) under grants [10.47379/VRG18013, 10.47379/NXT22018, 10.47379/ICT2201], as well as the Christian Doppler Research Association (CDG) JRC LIVE.

References

[1] R. Abboud, İ. İ. Ceylan, T. Lukasiewicz, and T. Salvatori. Boxe: A box embedding model for knowledge base completion. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Conference on Neural Information Processing Systems*, 2020.

[2] A. Bordes, N. Usunier, A. García-Durán, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In C. J. C. Burges, L. Bottou, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26: 27th Conference on Neural Information Processing Systems 2013.*, pages 2787–2795, 2013.

[3] C. Gao, C. Sun, L. Shan, L. Lin, and M. Wang. Rotate3d: Representing relations as rotations in three-dimensional space for knowledge graph embedding. In M. d'Aquin, S. Dietze, C. Hauff, E. Curry, and P. Cudré-Mauroux, editors, *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management*, pages 385–394. ACM, 2020.

[4] S. M. Kazemi and D. Poole. Simple embedding for link prediction in knowledge graphs. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31: Conference on Neural Information Processing Systems*, pages 4289–4300, 2018.

[5] H. Lu and H. Hu. Dense: An enhanced non-abelian group representation for knowledge graph embedding. *CoRR*, abs/2008.04548, 2020.

[6] A. Pavlovic and E. Sallinger. Expressive: A spatio-functional embedding for knowledge graph completion. In *The Eleventh International Conference on Learning Representations*, 2023.

[7] Z. Sun, Z. Deng, J. Nie, and J. Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *7th International Conference on Learning Representations*, 2019.

[8] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard. Complex embeddings for simple link prediction. In M. Balcan and K. Q. Weinberger, editors, *Proceedings of the 33nd International Conference on Machine Learning*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2071–2080. JMLR.org, 2016.

[9] Q. Wang, Z. Mao, B. Wang, and L. Guo. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, (12):2724–2743, 2017.

[10] R. West, E. Gabrilovich, K. Murphy, S. Sun, R. Gupta, and D. Lin. Knowledge base completion via search-based question answering. In *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, page 515–526. Association for Computing Machinery, 2014.

[11] B. Yang, W. Yih, X. He, J. Gao, and L. Deng. Embedding entities and relations for learning and inference in knowledge bases. In Y. Bengio and Y. LeCun, editors, *3rd International Conference on Learning Representations*, 2015.

[12] S. Zhang, Y. Tay, L. Yao, and Q. Liu. Quaternion knowledge graph embeddings. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Conference on Neural Information Processing Systems*, pages 2731–2741, 2019.