Strong Faithfulness for ELH Ontology Embeddings (Extended Abstract)

Victor Lacerda[1], Ana Ozaki[1,2] and Ricardo Guimarães[3]

[1]University of Bergen, Norway
[2]University of Oslo, Norway
[3]Zivid

Abstract

We present a region-based geometric model for embedding normalized $\mathcal{ELH}$ ontologies into a continuous vector space and formally prove that normalized $\mathcal{ELH}$ has the strong faithfulness property on convex geometric models. This means that there is an embedding that precisely captures the original ontology. We first prove the statement assuming (possibly) non-convex regions, and then impose convexity on the regions, showing that the property still holds.

Keywords

Geometric Models, Description Logics, Knowledge Graph Embedding, Faithfulness

1. Introduction

Knowledge Graphs (KGs) are a popular method for representing knowledge using triples of the form (subject, predicate, object), called *facts*, corresponding to *role assertions* in DL ABoxes. Although KGs contain a large number of facts, they are incomplete. This has sparked interest in using machine learning methods to suggest plausible facts to add to the KG based on patterns found in the data using KG embedding techniques, which aim to create representations of KGs in low-dimensional finite vector spaces [1, 2, 3, 4].

Embeddings that consider TBoxes formulated in description logic are a more recent phenomenon, here referred as *ontology embeddings*, where the *ontology* can have both facts and a TBox [5, 6, 7, 8, 9, 10, 11]. Ontology embeddings offer advantages over traditional KGEs as they exploit the semantic relationships between concepts and roles. One question that arises is how similar to the source ontology these embeddings are, and, more strictly, whether the generated embeddings are *guaranteed* to precisely represent the meaning of the source ontology and its entailments (of particular interest, the TBox entailments), such that one can freely move from its classical set-theoretic interpretation to its geometric one. This property is called the *strong* model faithfulness property [7]. So far, no previous work for ontology embeddings for fragments of $\mathcal{EL}^{++}$ has attempted to prove this property holds for their embedding method, nor has its existence been formally proven for the $\mathcal{ELH}$ language.

We investigate whether $\mathcal{ELH}$ has the strong faithfulness property over convex geometric models. We first prove the statement for embeddings in low dimensions, considering a region-based representation for (possibly) non-convex regions. We investigate strong faithfulness on convex geometric models with a number of dimensions connected to an ontology's signature and domain. We do so including embeddings for role inclusions, a problem that has not been well studied in the $\mathcal{ELH}$ ontology embedding literature. Additionally, we consider model checking in convex geometric models.

2. Preliminary Definitions

Geometric models We go from the traditional model-theoretic interpretation of the $\mathcal{ELH}$ description logic to geometric interpretations, adapting definitions by Gutiérrez-Basulto and Schockaert [4] and

DL 2024: 37th International Workshop on Description Logics, June 18–21, 2024, Bergen, Norway
victor.botelho@uib.no (V. Lacerda); anaoz@uio.no (A. Ozaki); ricardo.guimaraes@uib.no (R. Guimarães)

Bourgaux et al. [12]. Here we present the basic ingredients of the notion of a geometric interpretation. The main idea of geometric models is that the sets on which we define our interpretation functions are *continuous* vector spaces, and that the interpretation of concepts and roles are defined over geometric *regions*.

Definition 1 (Isomorphism Preserving Linear Maps). *Let m be a natural number and $f\colon \mathbb{R}^m \times \mathbb{R}^m \mapsto \mathbb{R}^{2\cdot m}$ a fixed but arbitrary linear map satisfying the following:*

1. *the restriction of f to $\mathbb{R}^m \times \{0\}^m$ is injective;*
2. *the restriction of f to $\{0\}^m \times \mathbb{R}^m$ is injective;*
3. *$f(\mathbb{R}^m \times \{0\}^m) \cap f(\{0\}^m \times \mathbb{R}^m) = \{0^{2\cdot m}\}$;*

where 0^m denotes the vector $(0, ..., 0)$ with m zeros. For instance, the concatenation function is a linear map that satisfies Points 1, 2, and 3.

Definition 2 (Geometric Interpretation). *Let f be an isomorphism preserving linear map and m a natural number. An m-dimensional f-geometric interpretation η of (N_C, N_R, N_I) assigns to each $A \in N_C$ a region $\eta(A) \subseteq \mathbb{R}^m$, to each $r \in N_R$ a region $\eta(r) \subseteq \mathbb{R}^{2\cdot m}$, and to each $a \in N_I$ a vector $\eta(a) \in \mathbb{R}^m$. We have that $\eta(\bot) := \emptyset$, $\eta(\top) := \mathbb{R}^m$, and for arbitrary $\mathcal{ELH}$ concepts:*

$$\eta(C \sqcap D) := \eta(C) \cap \eta(D) \text{ and } \eta(\exists r.C) := \{v \in \mathbb{R}^m \mid \exists u \in \eta(C) \text{ with } f(v, u) \in \eta(r)\}.$$

An m-dimensional f-geometric interpretation η satisfies

- *a concept assertion $A(a)$, if $\eta(a) \in \eta(A)$, a role assertion $r(a, b)$, if $f(\eta(a), \eta(b)) \in \eta(r)$,*
- *an $\mathcal{ELH}$ instance query (IQ) $C(a)$, if $\eta(a) \in \eta(C)$,*
- *an $\mathcal{ELH}$ CI $C \sqsubseteq D$, if $\eta(C) \subseteq \eta(D)$, and an RI $r \sqsubseteq s$, if $\eta(r) \subseteq \eta(s)$.*

We write $\eta \models \alpha$ if η satisfies an $\mathcal{ELH}$ axiom α. A geometric interpretation *satisfies* an ontology $\mathcal{O}$, in symbols $\eta \models \mathcal{O}$, if it satisfies all axioms in $\mathcal{O}$. We define model faithfulness based on the work by [7].

Definition 3 (Faithfulness). *Let $\mathcal{O}$ be a satisfiable ontology (or any other representation allowing the distinction between IQs and TBox axioms). Given an m-dimensional f-geometric interpretation η, we say that:*

- *η is a* strongly concept-faithful model *of $\mathcal{O}$ iff, for every concept C and individual name b, if $\eta(b) \in \eta(C)$ then $\mathcal{O} \models C(b)$;*
- *η is a* strongly IQ-faithful model *of $\mathcal{O}$ iff it is strongly concept-faithful and for each role r and individual names a, b: if $f(\eta(a), \eta(b)) \in \eta(r)$, then $\mathcal{O} \models r(a, b)$;*
- *η is a* strongly TBox-faithful model *of $\mathcal{O}$ iff for all TBox axioms τ: if $\eta \models \tau$, then $\mathcal{O} \models \tau$.*

We say that an ontology language *has the strong faithfulness property* over a class of geometric interpretations $\mathcal{C}$ if for every satisfiable ontology $\mathcal{O}$ in this language there is a geometric interpretation in $\mathcal{C}$ that is both a strongly IQ-faithful and a strongly TBox-faithful model of $\mathcal{O}$.

3. Strong Faithfulness on Convex Models

We prove that normalized $\mathcal{ELH}$ has the strong faithfulness property over a class of *convex* geometric models. We introduce a new mapping μ from the domain of a classical interpretation $\mathcal{I}$ to a vector space and a new geometric interpretation $\eta_{\mathcal{I}}$ based on this mapping. Our proofs now require us to fix the isomorphism preserving linear map f used in the definition of geometric interpretations (Definition 2). We choose the concatenation function, denoted $\oplus$, as done in Gutiérrez-Basulto and Schockaert [4]. The strategy for proving strong faithfulness for normalized $\mathcal{ELH}$ requires us to (a) find a suitable non-convex geometric interpretation for concepts and roles, and (b) show that the convex hull of the region maintains the property.

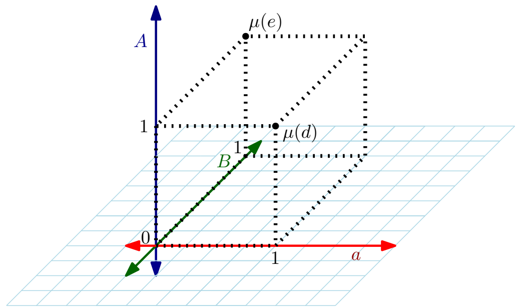

Figure 1: Axes colored in red, blue, and green correspond to the dimensions for a, A, and B, respectively.

Definition 4. *Let $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ be a classical $\mathcal{ELH}$ interpretation, and $\mathcal{O}$ an $\mathcal{ELH}$ ontology. We start by defining a new map $\mu\colon \Delta^{\mathcal{I}} \mapsto \mathbb{R}^{\mathsf{d}}$, where d corresponds to $|N_I(\mathcal{O})| + |N_C(\mathcal{O})| + |N_R(\mathcal{O})| \cdot |\Delta^{\mathcal{I}}|$. We assume, without loss of generality, a fixed ordering in our indexing system for positions in vectors, where indices 0 to $|N_I(\mathcal{O})| - 1$ correspond to the indices for individual names; $|N_I(\mathcal{O})|$ to $k = |N_I(\mathcal{O})| + |N_C(\mathcal{O})| - 1$ correspond to the indices for concept names; and k to $k + (|N_R(\mathcal{O})| \cdot |\Delta^{\mathcal{I}}|) - 1$ correspond to the indices for role names together with an element of $\Delta^{\mathcal{I}}$. We write $v[a]$, $v[A]$, and $v[r, d]$ to refer to the position in a vector v corresponding to a, A, and r together with an element d, respectively (according to our indexing system). For example, $v[a] = 0$ means the value at the index corresponding to the individual name a is 0. A vector is* binary *iff $v \in \{0, 1\}^{\mathsf{d}}$. We define μ using binary vectors. For all $d \in \Delta^{\mathcal{I}}$, $a \in N_I$, $A \in N_C$ and $r \in N_R$:*

- *$\mu(d)[a] = 1$ if $d = a^{\mathcal{I}}$, otherwise $\mu(d)[a] = 0$,*
- *$\mu(d)[A] = 1$ if $d \in A^{\mathcal{I}}$, otherwise $\mu(d)[A] = 0$, and*
- *$\mu(d)[r, e] = 1$ if $(d, e) \in r^{\mathcal{I}}$, otherwise $\mu(d)[r, e] = 0$.*

We introduce a definition for (possibly) non-convex geometric interpretations, using μ.

Definition 5. *Let $\mathcal{I}$ be a classical $\mathcal{ELH}$ interpretation. The* geometric interpretation *of $\mathcal{I}$, denoted $\eta_{\mathcal{I}}$, is defined as:*

$$\eta_{\mathcal{I}}(a) := \mu(a^{\mathcal{I}}), \text{ for all } a \in N_I,$$
$$\eta_{\mathcal{I}}(A) := \{\mu(d) \mid \mu(d)[A] = 1, d \in \Delta^{\mathcal{I}}\}, \text{ for all } A \in N_C,$$
$$\eta_{\mathcal{I}}(r) := \{\mu(d) \oplus \mu(e) \mid \mu(d)[r, e] = 1, d, e \in \Delta^{\mathcal{I}}\}, \text{ for all } r \in N_R.$$

An intuitive way of thinking about our definition μ is that it maps domain elements to a subset of the vertex set of the d-dimensional unit hypercube, as in Fig. 1.

Example 1. *Consider $A, B \in N_C$ and $a \in N_I$. Let $\mathcal{I}$ be an interpretation with $d, e \in \Delta^{\mathcal{I}}$ such that $d = a^{\mathcal{I}}$, $d \in A^{\mathcal{I}}$, and $e \in A^{\mathcal{I}} \cap B^{\mathcal{I}}$. We illustrate $\mu(d)$ and $\mu(e)$ in Fig. 1. In symbols, $\mu(d)[a] = 1$, $\mu(d)[A] = 1$, and $\mu(d)[B] = 0$, while $\mu(e)[a] = 0$, $\mu(e)[A] = 1$, and $\mu(e)[B] = 1$.*

Theorem 1. *For all $\mathcal{ELH}$ axioms α, $\mathcal{I} \models \alpha$ iff $\eta_{\mathcal{I}_{\mathcal{O}}} \models \alpha$.*

The canonical models need to be finite because the definition of $\eta_{\mathcal{I}}$ uses vectors in a finitely-dimensional space whose dimension depends on the size of Δ^{I} and $\mathcal{O}$. We then employ *finite* canonical models for normalized $\mathcal{ELH}$ since canonical models for arbitrary $\mathcal{ELH}$ CIs are not guaranteed to be finite. The following holds for the canonical model just defined.

Theorem 2. *Let $\mathcal{O}$ be a normalized $\mathcal{ELH}$ ontology. The following holds*

- *for all $\mathcal{ELH}$ IQs and CIs α in normal form over $\mathsf{sig}(\mathcal{O})$, $\mathcal{I}_{\mathcal{O}} \models \alpha$ iff $\mathcal{O} \models \alpha$; and*

- *for all RIs α over $\mathsf{sig}(\mathcal{O})$, $\mathcal{I}_{\mathcal{O}} \models \alpha$ iff $\mathcal{O} \models \alpha$.*

Theorem 3. *Let $\mathcal{O}$ be an $\mathcal{ELH}$ ontology and let $\mathcal{I}_{\mathcal{O}}$ be the canonical model of $\mathcal{O}$. The d-dimensional (possibly non-convex) $\oplus$-geometric interpretation $\eta_{\mathcal{I}_{\mathcal{O}}}$ of $\mathcal{I}_{\mathcal{O}}$ is a strongly and IQ and TBox faithful model of $\mathcal{O}$.*

Here, the vectors mapped by μ and the regions given by the non-convex geometric interpretation $\eta_{\mathcal{I}}$ are the anchor points for the convex closure of these sets. The geometric interpretation obtained by taking the convex hull is denoted $\eta^*_{\mathcal{I}}$.

Theorem 4. *Let $\eta_{\mathcal{I}}$ be a geometric interpretation as in Definition 5. If α is an $\mathcal{ELH}$ CI, an $\mathcal{ELH}$ RI, or an $\mathcal{ELH}$ IQ in normal form then $\eta_{\mathcal{I}} \models \alpha$ iff $\eta^*_{\mathcal{I}} \models \alpha$.*

Since normalized $\mathcal{ELH}$ has a finite canonical model, we can take advantage of the previous theorem to generalize it to arbitrary interpretations by applying it to the canonical $\mathcal{I}_{\mathcal{O}}$ of $\mathcal{O}$.

Theorem 5. *Let $\mathcal{O}$ be a normalized $\mathcal{ELH}$ ontology and let $\mathcal{I}_{\mathcal{O}}$ be the canonical model of $\mathcal{O}$. The d-dimensional convex $\oplus$-geometric interpretation of $\mathcal{I}_{\mathcal{O}}$ is a strongly IQ and TBox faithful model of $\mathcal{O}$.*

Corollary 6. *Normalized $\mathcal{ELH}$ has the strong faithful property over $\oplus$-geometric interpretations.*

Acknowledgments

Lacerda is supported by the Research Council of Norway, project number 316022, led by Ozaki.

References

[1] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: C. J. Burges, L. Bottou, M. Welling, Z. Ghahramani, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems, volume 26, Curran Associates, Inc., ., 2013. URL: https://proceedings.neurips.cc/paper_files/paper/2013/file/1cecc7a77928ca8133fa24680a88d2f9-Paper.pdf.

[2] R. Abboud, I. Ceylan, T. Lukasiewicz, T. Salvatori, BoxE: A box embedding model for knowledge base completion, in: H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, volume 33, Curran Associates, Inc., ., 2020, p. 9649–9661. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/6dbbe6abe5f14af882ff977fc3f35501-Paper.pdf.

[3] A. Pavlovic, E. Sallinger, ExpressivE: A spatio-functional embedding for knowledge graph completion, in: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023, OpenReview.net, 2023. URL: https://openreview.net/pdf?id=xkev3_np08z.

[4] V. Gutiérrez-Basulto, S. Schockaert, From knowledge graph embedding to ontology embedding? an analysis of the compatibility between vector space representations and rules, in: M. Thielscher, F. Toni, F. Wolter (Eds.), KR, AAAI Press, ., 2018, pp. 379–388. URL: https://aaai.org/ocs/index.php/KR/KR18/paper/view/18013.

[5] B. Xiong, N. Potyka, T.-K. Tran, M. Nayyeri, S. Staab, Faithful embeddings for EL++ knowledge bases, in: The Semantic Web – ISWC 2022, Springer International Publishing, ., 2022, pp. 22–38. doi:10.1007/978-3-031-19433-7_2.

[6] M. Jackermeier, J. Chen, I. Horrocks, Box^{2}el: Concept and role box embeddings for the description logic EL++, CoRR abs/2301.11118 (2023). URL: https://doi.org/10.48550/arXiv.2301.11118.

[7] Ö. L. Özçep, M. Leemhuis, D. Wolter, Cone semantics for logics with negation, in: C. Bessiere (Ed.), IJCAI, ijcai.org, ., 2020, pp. 1820–1826.

[8] X. Peng, Z. Tang, M. Kulmanov, K. Niu, R. Hoehndorf, Description logic EL++ embeddings with intersectional closure, CoRR abs/2202.14018 (2022). URL: https://arxiv.org/abs/2202.14018. arXiv:2202.14018.

[9] X. Peng, Z. Tang, M. Kulmanov, K. Niu, R. Hoehndorf, Description logic EL++ embeddings with intersectional closure, CoRR abs/2202.14018 (2022). URL: https://arxiv.org/abs/2202.14018.

[10] S. Mondal, S. Bhatia, R. Mutharaju, Emel++: Embeddings for EL++ description logic, in: A. Martin, K. Hinkelmann, H. Fill, A. Gerber, D. Lenat, R. Stolle, F. van Harmelen (Eds.), AAAI-MAKE, volume 2846 of *CEUR Workshop Proceedings*, CEUR-WS.org, ., 2021. URL: https://ceur-ws.org/Vol-2846/paper19.pdf.

[11] C. d'Amato, N. F. Quatraro, N. Fanizzi, Injecting background knowledge into embedding models for predictive tasks on knowledge graphs, in: R. Verborgh, K. Hose, H. Paulheim, P.-A. Champin, M. Maleshkova, O. Corcho, P. Ristoski, M. Alam (Eds.), The Semantic Web, Springer International Publishing, Cham, 2021, p. 441–457.

[12] C. Bourgaux, A. Ozaki, J. Z. Pan, Geometric models for (temporally) attributed description logics, in: M. Homola, V. Ryzhikov, R. A. Schmidt (Eds.), DL, volume 2954 of *CEUR Workshop Proceedings*, CEUR-WS.org, ., 2021. URL: https://ceur-ws.org/Vol-2954/paper-7.pdf.