BoxLitE: A Faithful Knowledge Base Embedding Based on Convex Optimization (Extended Abstract)

Bruno F. Lourenço[1], Hesham Morgan[2], Ana Ozaki[3], Aleksandar Pavlović[4] and Emanuel Sallinger[2]

[1] *Institute of Statistical Mathematics, Japan*
[2] *TU Wien, Austria*
[3] *University of Oslo, Norway*
[4] *University of Applied Sciences Campus Vienna, Austria*

Abstract

We introduce BoxLitE, a knowledge base (KB) embedding model for DL-Lite^{H} that allows for convex optimization. We show that for any satisfiable DL-Lite^{H} KB, there is a BoxLitE embedding that is a weakly faithful model. As a proof of concept, we show how to formulate the KB embedding task as a convex optimization problem and how to obtain embeddings with such desirable faithfulness properties. This is an extended abstract of our paper accepted at KR 2026 (Track: KR meets Machine Learning and Explanation).

Keywords

Knowledge Base Embeddings, Description Logic, Convex Optimization, Faithfulness

1. Introduction

Knowledge base (KB) embeddings combine the capability of knowledge graph embeddings to perform inductive reasoning for link prediction with deductive reasoning, using logic expressions present in an ontology [1]. Several authors have recently explored the idea of mapping concepts in a KB to regions in a vector space [2, 3, 4, 5, 6]. Region-based embeddings are important for KBs as they can naturally represent hierarchies present in KBs: more general concepts can be mapped to larger regions, containing those regions associated with more specific concepts [7, 8, 9, 10]. In this work, we introduce BoxLitE, a KB embedding model for DL-Lite^{H} [11] KBs that allows for convex optimization. In particular, we study the notion of *weakly faithful models* [12, 1] in the optimization task. This notion states that axioms that hold in the embedding are consistent with the KB and entailments of the KB are satisfied in the embedding. Our main contribution lies on theoretical grounds, establishing the existence of a faithful model for DL-Lite^{H} and the proposal of a novel approach that computes a faithful embedding. BoxLitE's implementation and the empirical results serve as a proof of concept that our theoretical results translate into practice. One of the motivations for this work is to try to perform link prediction using theoretically sounds techniques from a convex optimization perspective. We would like to check how much performance can we get from a purely convex approach.

DL 2026: 39th International Workshop on Description Logics, July 17–19, 2026, Lisbon, Portugal

bruno@ism.ac.jp (B. F. Lourenço); hesham.morgan@tuwien.ac.at (H. Morgan); anaoz@uio.no (A. Ozaki); aleksandar.pavlovic@hcw.ac.at (A. Pavlović); sallinger@dbai.tuwien.ac.at (E. Sallinger)
https://bflourenco.github.io/ (B. F. Lourenço); https://informatics.tuwien.ac.at/people/hesham-morgan (H. Morgan); https://www.mn.uio.no/ifi/english/people/aca/anaoz/index.html (A. Ozaki); https://personen.hcw.ac.at/aleksandar-pavlovic/ (A. Pavlović)
0000-0003-2369-0107 (B. F. Lourenço); 0009-0006-2403-4581 (H. Morgan); 0000-0002-3889-6207 (A. Ozaki); 0000-0001-6887-9515 (A. Pavlović)

2. BoxLitE's Semantics

Here we introduce a semantics for DL-Lite^H inspired by box-based geometric models [7] and suitable for defining a KB embedding model that can be optimized using convex optimization. We define a geometric model that uses axis-aligned hyper-rectangles, called *boxes* (Figure 1).

Let ϵ and s_Ω be fixed but arbitrary [1] positive constants with $\epsilon \leq s_\Omega$. Let $\boldsymbol{\epsilon}$ and $\mathbf{s}$ denote the d-dimensional vectors ($d > 0$) whose value in each dimension is equal to ϵ and s_Ω respectively. Henceforth, we denote elementwise comparison operators with $\leq_d$ and $\geq_d$. Using $\boldsymbol{\epsilon}$ and $\mathbf{s}$, we define an axis-aligned hyper-rectangle, called the *universe box*: $\Omega = \{\mathbf{x} \in \mathbb{R}^d \mid -\mathbf{s} \leq_d \mathbf{x} \leq_d \mathbf{s}\}$.

Definition 1. *A* box *is an element of the set* Box*, defined as:*

$$\mathsf{Box} \coloneqq \{\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{L} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U} - \boldsymbol{\epsilon}\} \mid \ \mathbf{0} \leq_d (\mathbf{U} - \mathbf{L}) \leq_d 2\mathbf{s}, \ \mathbf{L}, \mathbf{U} \in \mathbb{R}^d\}.$$

For a non-empty box, the lower *and* upper *bounds are denoted* $\mathbf{L}$ *and* $\mathbf{U}$ *(resp.) of* $\mathbf{X}$*. If* $\mathbf{X}$ *is the empty box then we set* $\mathbf{L} := \mathbf{U} := \mathbf{0}$*. We denote with* $\mathbf{x}[i]$ *the i-th dimension of a vector* $\mathbf{x}$*. If* $\mathbf{X}$ *is a box with lower and upper bounds* $\mathbf{L}$ *and* $\mathbf{U}$*, resp., then* $\mathbf{X}[i]$ *is the pair* $(\mathbf{L}[i], \mathbf{U}[i])$.

Definition 2. *A* box interpretation η *is a function that maps:*

- *each individual name* $a \in \mathsf{N_I}$ *to two vectors* $\eta(a) = (\mathsf{pos}(a), \mathsf{bump}(a))$*, namely, a position* $\mathsf{pos}(e) \in \Omega$ *and a bump* $\mathsf{bump}(e) \in \Omega$*;*
- *each concept name* $A \in \mathsf{N_C}$ *to a box* $\eta(A) \in \mathsf{Box}$*;*
- *each role name* $R \in \mathsf{N_R}$ *to three boxes* $\eta(R) = (\mathsf{Head}(R), \mathsf{Tail}(R), \mathsf{Bump}(R))$*, which we call R's head* $\mathsf{Head}(R)$*, tail* $\mathsf{Tail}(R)$ *and bump box* $\mathsf{Bump}(R)$*.*

We extend the mapping function η *to arbitrary DL-Lite*^{H} *concept and role expressions as follows:*

$$\eta(\neg C) \coloneqq \overline{\eta(C)}, \quad \eta(R^-) \coloneqq (\mathsf{Tail}(R), \mathsf{Head}(R), \mathsf{Bump}(R)),$$
$$\eta(\exists R) \coloneqq \{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{L_R^H} - \mathbf{U_R^B} + \boldsymbol{\epsilon} \leq_d \mathbf{x} \leq_d \mathbf{U_R^H} - \mathbf{L_R^B} - \boldsymbol{\epsilon}\}$$

where $\mathbf{L_R^{\mathit{X}}}, \mathbf{U_R^{\mathit{X}}}$ *with* $X \in \{\mathbf{H}, \mathbf{T}, \mathbf{B}\}$ *are the lower and upper bounds of* $\mathsf{Head}(R)$, $\mathsf{Tail}(R)$*, and* $\mathsf{Bump}(R)$*; and where* $\overline{\eta(C)}$ *represents* $\eta(C)$*'s complement box. Furthermore, for any* $C \in \mathsf{N_C^\exists}$*, the complement box* $\overline{\eta(C)}$ *is defined as follows:*

$$\overline{\eta(C)} \coloneqq \{\mathbf{x} \in \mathbb{R}^d \mid (-\mathbf{s} - \mathbf{L_C} + \boldsymbol{\epsilon}) \leq_d \mathbf{x} \leq_d (\mathbf{s} - \mathbf{U_C} - \boldsymbol{\epsilon})\}$$

with $\mathbf{L_C}$ *and* $\mathbf{U_C}$ *being* $\eta(C)$*'s lower and upper bounds.*

A box interpretation is *box consistent* if for all $C \in \mathsf{N_C^\exists}$ we have that $\eta(C) \cap \overline{\eta(C)} = \emptyset$.

Definition 3. *A box interpretation* η *satisfies*

- $R(a, b)$ *with* $R \in \mathsf{N_R}$ *and* $a, b \in \mathsf{N_I}$ *iff* (i) $\mathsf{pos}(a) + \mathsf{bump}(b) \in \mathsf{Head}(R)$, ($ii$) $\mathsf{pos}(b) + \mathsf{bump}(a) \in \mathsf{Tail}(R)$, *and* ($iii$) $\mathsf{bump}(a), \mathsf{bump}(b) \in \mathsf{Bump}(R)$;
- $C \sqsubseteq D$*,* $C \in \mathsf{N_C^\exists}$ *and* $D \in \mathsf{N_{C\neg}^\exists}$ *iff* $\eta(C) \subseteq \eta(D)$*;*
- $R \sqsubseteq S$ *with* $R, S \in \mathsf{N_R^-}$ *iff* $\mathsf{Bump}(R) \subseteq \mathsf{Bump}(S)$, $\mathsf{Head}(R) \subseteq \mathsf{Head}(S)$, $\mathsf{Tail}(R) \subseteq \mathsf{Tail}(S)$.

Regarding concept assertions with $C \in \mathsf{N_C^\exists}$ *and* $a \in \mathsf{N_I}$*,* η *satisfies* $C(a)$ *if* $\mathsf{pos}(a) \in \eta(C)$ *and* η *falsifies* $C(a)$ *if* $\mathsf{pos}(a) \in \overline{\eta(C)}$*. Otherwise the status of* $C(a)$ *is* unknown[2].

---

[1]In our work and, in particular, in the implementation, we take $0 < \epsilon < \epsilon_{max}$, where $\epsilon_{max} = 0.5$ and $\epsilon_{max} \leq s_\Omega/8$. Making ϵ and s_Ω learnable parameters would not enhance the representation capabilities of our model but allow for infinitely many equivalent solutions of our learned embeddings which only differ in scale.

[2]Geometric models for languages that have negation normally need to allow for the 'unknown' truth status. This happens in the Cone Semantics [12], tailored for $\mathcal{ALC}$, which has negation. Since DL-Lite^H allows for $B \sqsubseteq \neg C$ we use complement boxes and the 'unknown' truth state.

We write $\eta \models \alpha$ when η satisfies an axiom α, and $\eta \not\models \alpha$ otherwise. Also, if $\eta \models \alpha$ for every axiom α in a KB K then we say that η satisfies K or, equivalently, that η is a model of K, that is, $\eta \models K$.

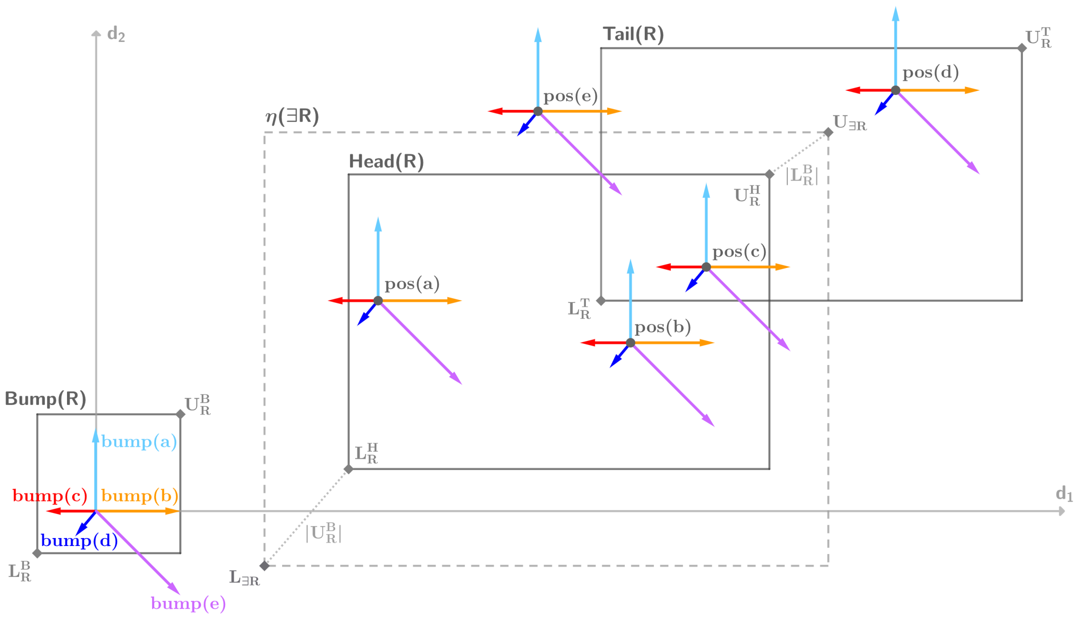

Figure 1: A two-dimensional box interpretation η is shown. It maps a role $R \in \mathsf{N_R}$ to the Head(R), Tail(R), and Bump(R) boxes. Additionally $\eta(\exists R)$ is visualized. The depicted box interpretation satisfies the assertions $\{ R(a, b)$, $R(b, d)$, $R(b, c)$, $R(c, d)$, $R(c, c)$, $\exists R(a)$, $\exists R(b)$, $\exists R(c) \}$ over $\{a, b, c, d, e\} \in \mathsf{N_I}$.

3. Model Faithfulness

Here we study model faithfulness [12], which is a property that is useful to show that geometric models correctly represent the knowledge present in KBs. For KB completion, one is particularly interested in preserving the TBox axioms, while allowing new assertions consistent with the KB to hold.

Definition 4. *(Adapted [12, 1]) Let $\mathcal{L}$ be an ontology language and K a satisfiable KB in $\mathcal{L}$. A box interpretation η is*

- weakly KB faithful *for $\mathcal{L}$ and K if for every KB axiom α in $\mathcal{L}$, $\eta \models \alpha$ implies $K \cup \{\alpha\}$ is consistent;*
- KB entailed *for $\mathcal{L}$ and K if for every KB axiom α in $\mathcal{L}$ that is entailed by K, $\eta \models \alpha$.*

*We may omit "weakly" and just say "faithful". We may also omit $\mathcal{L}$ and/or K if they are clear from the context. If a box η interpretation satisfies both of the properties we say that η is a KB **faithful model**. These notions can be adapted for the case $\mathcal{L}$ is a description logic with TBox and ABox axioms and for the case where we consider only TBox axioms (or only ABox axioms) for TBox (or ABox) faithful models.*

Proposition 1 and Proposition 2 give basic conditions for faithfulness in DL-Lite^{H}.

Proposition 1. *Let $\mathcal{T}$ be a DL-Lite^{H} TBox. If $\eta \models \mathcal{T}$ then η is (weakly) TBox faithful.*

Proposition 1 follows from the fact that, in DL-Lite^{H}, TBoxes are always consistent, that is, inconsistencies only happen when considering a TBox and an ABox.

Proposition 2. *Let K be a DL-Lite^{H} KB and η a box interpretation that is box consistent. If $\eta \models K$ then η is (weakly) KB faithful.*

We are now ready to state Theorem 1, which is the main result of this section.

Theorem 1. *There exists a suitable s_Ω such that every satisfiable DL-Lite^{H} KB K has a box interpretation η_{I_K} that is a KB faithful model and box consistent.*

Table 1
Test Results for BoxLitE on F_v1-4. Standard deviation nearly 0 in all cases.

Dataset	MRR	H@1	H@3	H@10
F_v1	.720	.545	.870	.979
F_v2	.549	.352	.685	.905
F_v3	.433	.288	.509	.708
F_v4	.444	.249	.571	.763

Corollary 1. *Let K be a DL-Lite^H KB with an empty ABox. For any $d \geq d_{\min}$ there is a box interpretation η with dimensionality d such that $d_{\min} = |\mathsf{N_C}| + 3|\mathsf{N_R}|$ and η is a TBox faithful model of K.*

Corollary 2. *Let K be a satisfiable DL-Lite^H KB. For any $d \geq d_{\min}$ there is a box interpretation η with dimensionality d such that: $d_{\min} = |\mathsf{N_C}| + |\mathsf{N_R}|(2 + |\mathsf{N_I}| + 2|\mathsf{N_R}|)$ and η is a KB faithful model of K.*

We show in full version of the paper [13] that one can encode TBox axioms as linear inequalities that are used as convex constraints in an optimization problem that guarantees any BoxLitE embedding solution z to be KB faithful for DL-Lite^H. Let $\mathscr{C}_K \subseteq \mathbb{R}^n$ be the set of z's such that the constraints are satisfied. That is, $z \in \mathscr{C}_K$ if and only if: (a) for each TBox axiom the corresponding inequalities are satisfied; and (b) the box consistency and universe constraints are satisfied, see the supplemental material for details. Also, let $f_{\boldsymbol{\lambda}} : \mathbb{R}^n \to \mathbb{R}$ be the function that maps z to the objective value for a given choice of nonnegative hyperparameters $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3_+$. The next theorem establishes the conditions in which the optimizer is guaranteed to find an optimal solution and associated properties.

Theorem 2. *Let K be a satisfiable DL-Lite^H KB. For nonnegative $\boldsymbol{\lambda}$ the following problem is convex.*

$$\min_{z \in \mathbb{R}^n} f_{\boldsymbol{\lambda}}(z), \quad \text{subject to } z \in \mathscr{C}_K. \tag{1}$$

In particular, $f_{\boldsymbol{\lambda}}$ is a convex function, $\mathscr{C}_K$ is a polyhedral set and the following items hold.

i) *For d as in Corollary 1, and s_Ω as in Theorem 1, $\mathscr{C}_K$ is nonempty.*

ii) *Any embedding solution $z \in \mathscr{C}_K$ corresponds to a box consistent interpretation that is TBox faithful.*

iii) *If $\lambda_1 = 0$ and there is an optimal solution z^* such that $f_{\boldsymbol{\lambda}}(z^*) \leq 0$ then the box interpretation corresponding to z^* is KB faithful.*

iv) *Suppose that $\lambda_1 = \lambda_2 = \lambda_3 = 0$. For d as in Corollary 2, and s_Ω as in Theorem 1, there is an optimal solution z^* s.t. $f_{\boldsymbol{\lambda}}(z^*) \leq 0$ holds.*

Informally, Theorem 2 states that we can find a box interpretation for a satisfiable K via convex optimization over a polyhedral set. Any $z \in \mathscr{C}_K$ (whether optimal or not) corresponds to a box interpretation η that is TBox faithful and box consistent. Item *i*) gives a bound on the minimum required d to ensure that $\mathscr{C}_K$ is nonempty, but this estimate seems to be conservative. Items *iii*) and *iv*) imply that if we wish to find a solution that is KB faithful, this can be done by setting the hyperparameters associated to regularization terms to 0, setting d to be sufficiently large and solving (1).

4. Proof of Concept

We evaluate BoxLitE's performance on subsets (F_v1-v4) of the Family dataset [14], providing first results for its scalability and reasoning capabilities. The sizes of the datasets vary from 155 to 895 individuals. We implemented BoxLitE's (convex) second-order cone optimization problem in Python 3.12 using CVXPY [15, 16] for modeling and MOSEK [17] for solving it [3]. The results reveal that MOSEK finds a solution that is KB faithful and satisfies the concept inclusions in the TBox (Table 1).

[3] The code is available at https://github.com/AleksVap/BoxLitE.

Acknowledgments

This work was supported by the "Strategic Research Projects" grant from ROIS (Research Organization of Information and Systems). Bruno F. Lourenço's work was partially supported by the JSPS Grant-in-Aid for Early-Career Scientists 23K16844. Ana Ozaki was supported by the Research Council of Norway, projects (316022, 322480) and Integreat - Norwegian Centre for knowledge-driven machine learning (332645). Emanuel Sallinger's and Hesham Morgan's work on this paper was funded by the Vienna Science and Technology Fund (WWTF) [Grant ID: 10.47379/VRG18013, 10.47379/ICT25032, 10.47379/NXT22018, 10.47379/ ICT2201, 10.47379/DCDH001], and by the Austrian Science Fund (FWF) 10.55776/COE12.

Declaration on Generative AI

The authors have not employed any Generative AI tools in the extended abstract. In the full version of the paper, Generative AI was used for grammar and spelling checks.

References

[1] C. Bourgaux, R. Guimarães, R. Koudijs, V. Lacerda, A. Ozaki, Knowledge base embeddings: Semantics and theoretical properties, in: P. Marquis, M. Ortiz, M. Pagnucco (Eds.), KR, 2024. URL: https://doi.org/10.24963/kr.2024/77. doi:10.24963/KR.2024/77.

[2] 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.

[3] A. Pavlović, E. Sallinger, ExpressivE: A spatio-functional embedding for knowledge graph completion, in: the 11th International Conference on Learning Representations (ICLR 2023), OpenReview.net, 2023. URL: https://openreview.net/pdf?id=xkev3_np08z.

[4] A. Pavlović, E. Sallinger, SpeedE: Euclidean geometric knowledge graph embedding strikes back, in: K. Duh, H. Gomez, S. Bethard (Eds.), Findings of the Association for Computational Linguistics: NAACL 2024, Association for Computational Linguistics, Mexico City, Mexico, 2024, pp. 69–92. URL: https://aclanthology.org/2024.findings-naacl.6/. doi:10.18653/v1/2024.findings-naacl.6.

[5] A. Pavlović, E. Sallinger, S. Schockaert, Faithful Differentiable Reasoning with Reshuffled Region-based Embeddings, in: Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning, 2025, pp. 489–499. URL: https://doi.org/10.24963/kr.2025/48. doi:10.24963/kr.2025/48.

[6] H. Morgan, A. Pavlović, A. Ozaki, B. F. Lourenço, E. Sallinger, Convexity for trustworthy knowledge graph embeddings: From interpretability to reliability, in: Proceedings of the 19th Forschungsforum der österreichischen Fachhochschulen (FHK 2026), Graz, Austria, April 15 - April 16, 2026, 2026.

[7] R. Abboud, İ. İ. Ceylan, T. Lukasiewicz, T. Salvatori, BoxE: A box embedding model for knowledge base completion, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), NeurIPS, 2020.

[8] M. Jackermeier, J. Chen, I. Horrocks, Dual box embeddings for the description logic EL++, in: Proceedings of the ACM Web Conference, WWW, Association for Computing Machinery, 2024, p. 2250–2258. URL: https://doi.org/10.1145/3589334.3645648. doi:10.1145/3589334.3645648.

[9] A. Pavlović, E. Sallinger, Building bridges: Knowledge graph embeddings respecting logical rules, in: B. Kimelfeld, M. V. Martinez, R. Angles (Eds.), Proceedings of the 15th Alberto Mendelzon International Workshop on Foundations of Data Management (AMW 2023), Santiago de Chile,

Chile, May 22-26, 2023, volume 3409 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023. URL: https://ceur-ws.org/Vol-3409/paper9.pdf.
[10] A. Pavlović, E. Sallinger, Raising the efficiency of knowledge graph embeddings while respecting logical rules, in: G. Montoya, E. Sallinger, G. Vargas-Solar (Eds.), Proceedings of the 16th Alberto Mendelzon International Workshop on Foundations of Data Management (AMW 2024), Mexico City, Mexico, September 30 - October 4, 2024, CEUR Workshop Proceedings, CEUR-WS.org, 2024.
[11] A. Artale, D. Calvanese, R. Kontchakov, M. Zakharyaschev, The DL-Lite family and relations, J. Artif. Intell. Res. 36 (2009) 1–69.
[12] O. Lütfü Özçep, M. Leemhuis, D. Wolter, Cone semantics for logics with negation, in: C. Bessiere (Ed.), IJCAI, International Joint Conferences on Artificial Intelligence Organization, 2020, pp. 1820–1826. URL: https://doi.org/10.24963/ijcai.2020/252. doi:10.24963/ijcai.2020/252, main track.
[13] B. F. Lourenço, H. Morgan, A. Ozaki, A. Pavlovic, E. Sallinger, Boxlite: A faithful knowledge base embedding based on convex optimization, CoRR abs/2605.23937 (2026). URL: https://doi.org/10.48550/arXiv.2605.23937. doi:10.48550/ARXIV.2605.23937. arXiv:2605.23937.
[14] A. Imenes, R. Guimarães, A. Ozaki, Marrying query rewriting and knowledge graph embeddings, in: RuleML+RR, Springer-Verlag, Berlin, Heidelberg, 2023, p. 126–140. URL: https://doi.org/10.1007/978-3-031-45072-3_9. doi:10.1007/978-3-031-45072-3_9.
[15] S. Diamond, S. Boyd, CVXPY: A Python-embedded modeling language for convex optimization, Journal of Machine Learning Research 17 (2016) 1–5.
[16] A. Agrawal, R. Verschueren, S. Diamond, S. Boyd, A rewriting system for convex optimization problems, Journal of Control and Decision 5 (2018) 42–60.
[17] M. ApS, MOSEK Optimizer API for Python 11.1.3, 2026. URL: https://docs.mosek.com/latest/pythonapi/index.html.