




# Response to the Reviewers' Comments

## Trustworthy Knowledge Base Embeddings: A Foundational Study of Box Semantics

We thank the reviewers for their detailed and constructive feedback, which helped us improve the submitted manuscript. This document contains our response to the reviewers' comments. To simplify the identification of our [changes/additions to the text](#) for the purpose of reviewing, we have highlighted significant ones (e.g., not the typos or minor rephrasings) [in blue font](#).

We have looked over the whole paper, added some explanations, streamlined some parts and simplified some proofs. Generally, we focused on adding explanations to Definition 2 (especially, but not only, to the motivation of the interpretation of instances). Additionally, we added further discussion regarding Helly-satisfiability (see Definition 35 and Proposition 36) and added an overall discussion section before we conclude the paper in which we also comment on the computational challenges and computational complexity of determining Helly satisfiability. We further discuss the influence of these considerations on related embedding approaches and sketch a number of applications and further lines of research.

### Response to the detailed comments.

1. R.1: "The paper contains no empirical evaluation demonstrating how often real-world ontologies violate HP or how these violations manifest in trained models. "  
 In this paper we are focusing on the theoretical aspects of this discussion and leave practical questions to future work. We, however, added a discussion section where we are discussing the influence of this result to other embedding approaches and other geometric shapes and discuss briefly the complexity question.
2. R.1: "The computational complexity of the Helly-satisfiability procedure is not well discussed, leaving its feasibility unclear. "  
 A thorough discussion of the complexity is left for future work. We added a discussion on how an approximation of the test on Helly's Property could work in practice and give a complexity result for that.
3. R.1: "The discussion of alternative geometric frameworks (e.g., cones, general convex regions) is brief and could better situate the results; "  
 In the discussion section, we have added a detailed paragraph on how these results fit into related results for alternative frameworks. We additionally explain how our results can be adapted and generalized to the benefit of other frameworks.
4. R.1: "Some parts of the manuscript felt a bit redundant. Overall, this is a solid and relevant theoretical paper that meaningfully analyzes the foundational limitations of box-based embeddings. While it would benefit from tightening some parts and adding more context on practical impact and alternatives, the core results are sound and significant. "

- ✓ We have kept where it seemed appropriate some repetition and redundancy as a means to increase the overall understandability. In other parts we tightened the paper a bit and removed inessential repetition. In particular, we added some more discussion, e.g., for Definition 2.
5. R.2: “In the preliminaries section, role inclusions and role hierarchies are included in the TBox; however, they belong to the RBox. I would suggest revising the definition of ontology presented here. ”
- ✓ We introduced the notion of an Rbox in the preliminaries and throughout the paper. Therefore, we added additional definitions, e.g., Rbox-entailment closure. This is, however, only a notational change without changing the meaning or requiring new technical work, since the old notion of, e.g., TBox-entailment closure was covering already both TBox and Rbox.
6. R.2: “In the section ‘Box Embeddings’, can you elaborate on the rationale for defining individuals as subsets of  $R^n$ ? In classical semantics, given a domain  $D$ , a concept is a subset of  $D$  and an individual is an element of  $D$ . Following this notion, in the embedding space it seems natural to represent the concept as a box and an individual as a point. I think it would be good to have a rationale of why you decide to represent individual embeddings as subsets. ”
- ✓ We’ve added a detailed discussion of this point in the context of Definition 2. In particular, it is important to note that the classical case is a special case of this definition. Modeling individuals as subsets allows, e.g., for modeling uncertainty and allows for modeling individuals in the same way (as the same kind of object) as concepts, thus for instance providing a simplified framework in the training process.
7. R.2: “The proof for proposition 12 [now 13] seems just a (very clear) example but not a general demonstration of the proposition. I would suggest to elaborate a general proof and maybe show the current proof as an example. ”
- ✓ In most cases, a shallow model is both weakly and globally faithful. One prominent example is the classical model that can be interpreted as a shallow model and is clearly both weakly and globally faithful. The interesting aspect is the case where both do not coincide. The main aim of this proposition is thus to demonstrate the necessity of discriminating between global and weak faithfulness. Thus, we show that there is an ontology and an embedding method such that the shallow model is weakly but not globally faithful. This result shows that testing for weak faithfulness is not sufficient. The formulation of the proposition was however misleading. We therefore kept the proof but reformulated the proposition appropriately.
8. R.2: “The paper is quite large and the analysis is very deep. Before jumping to the conclusion, I would suggest adding a discussion section to sum up the key messages of the paper. Furthermore, the conclusion seems to mainly focus on the Helly’s property, the this extended version contains additional results that can be included as well. ”
- ✓ We added a full discussion section where we summarize and discuss the whole paper with a special focus on the influence of the results on real-world embedding approaches. We also slightly extended the conclusion.
9. R.2: “Please replace the Arxiv citation of Query2Box to the published version. ”
- ✓ Done.