

Dear Editors,

I am pleased to submit our manuscript titled "*Towards Interpretable Embeddings: Aligning Representations with Semantic Aspects*" for consideration in the Neurosymbolic AI Journal Special Issue on NeSy 2024. This manuscript is an extension of the short paper we presented at NeSy 2024, which was titled "*Bringing Back Semantics to Knowledge Graph Embeddings: An Interpretability Approach*." We believe this extended version provides significant additional contributions to the field.

The new manuscript offers the following key enhancements:

- **Detailed Introduction and Motivation:** We have expanded the introduction to provide a more in-depth discussion of the motivation for this work and its relevance to the community.
- **Preliminaries Section:** This new section introduces key background concepts, ensuring that readers have the necessary foundation to fully understand our approach.
- **Comprehensive Related Work:** The related work section has been considerably expanded to include a broader and detailed review of prior studies, establishing a clearer context for our contributions.
- **Detailed data analysis:** A new section explaining the process of data-driven feature selection for different datasets has been added, with several illustrations and plots.
- **Formalization of the Approach and Algorithm:** We have formalized our approach and included an algorithm to offer a more structured and transparent presentation of our method.
- **Added figure explaining the method:** The figure representing our method has been revised and improved for greater clarity.
- **Significant Expansion of Experiments:** We have greatly expanded the experimental section, providing details of feature extraction, and introducing new experiments on two additional Knowledge Graph datasets, YAGO and Freebase. The results are presented in multiple tables, with supporting vector plots to enhance clarity.
- **Alternative Feature Extraction and Evaluation with LLMs:** An additional section discusses an alternative approach to feature extraction and the evaluation with large language models, along with the reasoning for not pursuing this approach further.

We are confident that these extensions bring substantial value to the original work, making this manuscript a strong contribution to the special issue. We appreciate the opportunity to share our work with the NeSy community and thank you for your consideration.

We look forward to your feedback and are available to provide any further information or clarifications.

Sincerely,

Nitisha Jain