

Integrating Semantics with Symbolic AI: The Path to Interpretable Hybrid AI Systems

Maria-Esther Vidal^{1,2,3}

¹Leibniz University of Hannover, Germany

²TIB-Leibniz Information Center for Science and Technology, Germany

³L3S Research Center, Germany

Abstract

The increasing complexity of AI-driven decision-making necessitates more interpretable and explainable models, particularly in critical domains like medicine. This talk explores the integration of semantics with symbolic AI to develop hybrid AI systems that combine the strengths of machine learning and formal reasoning. By leveraging knowledge graphs (KGs), ontologies, and logical constraints, these systems enable effective knowledge representation, improve transparency, and support reasoning over heterogeneous data.

This talk discusses the evolution of symbolic AI from early rule-based systems to modern neuro-symbolic approaches and analyzes emerging research trends in semantic data management, knowledge augmentation, and hybrid AI architectures. A key focus is the application of hybrid AI in oncology, demonstrating how integrating symbolic reasoning with machine learning enhances diagnostic accuracy, supports counterfactual reasoning, and aids in treatment decision-making.

Using structured frameworks like a boxology of design patterns¹ and hybrid AI frameworks such as Semantic Web Machine Learning (SWeML)², we outline patterns for model generation, inference, and validation, showcasing their effectiveness in optimizing predictive performance and ensuring compliance with medical protocols. Experimental results from real-world medical datasets highlight improvements in link prediction³, causal graph discovery⁴, and counterfactual reasoning⁵.

This presentation underscores the necessity of fusing semantics with AI to enhance interpretability, usability, and trustworthiness. By bridging the gap between data-driven learning and symbolic reasoning, hybrid AI systems provide a powerful framework for advancing AI-driven decision support across multiple domains. Future directions include optimizing computational efficiency, improving usability, and developing scalable, user-centric hybrid AI solutions.

¹Michael van Bekkum, Maaïke de Boer, Frank van Harmelen, Andre Meyer-Vitali, Annette Ten Teije. Modular design patterns for hybrid learning and reasoning systems a taxonomy, patterns and use cases. *Applied Intelligence* (2021) 51:6528–6546 <https://link.springer.com/article/10.1007/s10489-021-02394-3>

²A. Breit, L. Waltersdorfer, F. J. Ekaputra, M. Sabou, A. Ekelhart, A. Iana, H. Paulheim, J. Portisch, artem Revenko, A. Ten Teije, F. van Harmelen. Combining Machine Learning and Semantic Web: A Systematic Mapping Study. *ACM Computing Surveys*, Vol. 55, No. 14s. 2023 <https://dl.acm.org/doi/10.1145/3586163>

³Disha Purohit, Yashrajsinh Chudasama, Maria Torrente, Maria-Esther Vidal. VISE: Validated and Invalidated Symbolic Explanations for Knowledge Graph Integrity. *EXPLIMED@ECAI 2024*. <https://ceur-ws.org/Vol-3831/paper5.pdf>

⁴Hao Huang, Maria-Esther Vidal. HyKG-CF: A Hybrid Approach for Counterfactual Prediction using Domain Knowledge. *WSDM 2025*: 1104–1105 <https://dl.acm.org/doi/10.1145/3701551.3708813>

⁵Chudasama Yashrajsinh, Huang Hao, Disha Purohit, Maria-Esther Vidal. HyKG-CF: A Hybrid Approach for Counterfactual Prediction using Domain Knowledge. *IEEE Access* 10.1109/ACCESS.2025.3529133

AMW 2024: 16th Alberto Mendelzon International Workshop on Foundations of Data Management, September 30th–October 4th, 2024, Mexico City, Mexico

✉ vidal@l3s.de (M. Vidal)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).