Grazia Serratore, University of Calabria (Rende, Italy)/Informatics and Telematics Institute of the National Research Council (Rende, Italy)

Julaine Clunis, Old Dominion University (Norfolk, VA, USA)

The Integration of Artificial Intelligence and Ontologies: Transformations in Knowledge Representation and Application

Abstract

Artificial Intelligence (AI) is reshaping the landscape of knowledge representation. There is an increasingly strong bidirectional relationship, between AI techniques and ontologies. AI techniques revolutionized traditional, manual ontology development and contribute to automated ontology construction, while ontologies enhance the performance of AI systems and their semantic accuracy. Through a comprehensive review of current literature, this paper aims to examine: i) how Machine Learning (ML) techniques contribute to the automated construction, refinement, and validation of ontologies; ii) the most widely used and effective ML approaches for ontology construction; iii) how domain-specific requirements influence the selection and adaptation of AI techniques for building and applying ontologies; iv) how ontologies enhance the interpretability, explainability, and reliability of AI. This overview highlights the integration between AI and ontology engineering across different domains and indicates that so far successful AI-ontology integration typically follows a collaborative model, in which AI acts as an intelligent assistant to human experts, combining computational efficiency with critical domain knowledge. Ethical concerns, such as bias and hallucinations, remain pressing challenges that require standardized frameworks and careful considerations. However, this reciprocal relationship between AI and ontologies points to the development of more dynamic, adaptive, and complete ontologies.

1. Introduction

Artificial intelligence (AI) has become widespread in every domain of knowledge, transforming the way information is processed, analyzed, and utilized. While AI systems have traditionally relied on ontologies for structuring knowledge, recent advancements have demonstrated that AI techniques are also being used to automate the construction of Knowledge Organization Systems (KOSs). In fact, a significant aspect of AI's influence lies in the development and application of KOSs, especially ontologies. This bidirectional relationship has reshaped the landscape of knowledge representation, enabling more efficient and scalable systems.

Considering this dynamic interplay between AI and ontologies, this study will explore the main AI techniques used in ontology building and how ontologies are leveraged for Machine Learning (ML), Natural Language Processing (NLP), and other AI applications. Specifically, this work will try to answer the following questions:
RQ1 How do ML techniques contribute to the automated construction, refinement, and validation of ontologies?
RQ2 What are the most used and effective ML techniques for ontology construction?
RQ3 How do domain-specific requirements influence the selection and adaptation of AI techniques for building and leveraging ontologies?
RQ4 How do ontologies enhance the interpretability, explainability and reliability of AI?

By analyzing these aspects, this study aims to provide a comprehensive overview of how AI shapes ontology development and use across diverse fields, reinforcing the intersection of ML and knowledge representation.

2. Ontologies and AI

Ontologies are structured representations of shared conceptualizations, capturing domain-specific knowledge through the identification and mapping of concepts and their relationships (Gruber 1993). By formally specifying the intended meaning of vocabularies within specific domains, they have significantly transformed various fields of human endeavor, enhancing understanding and interoperability (Naqvi et al. 2021). The maintenance of ontologies requires ongoing enrichment, that involves discovering new concepts and placing them within the appropriate positions in the existing conceptual structure. Traditionally, ontology development and maintenance have been manual, time-consuming, and often error-prone processes. This is particularly critical in the case of domain ontologies, which must guarantee high coverage of the analyzed fields and are thus more susceptible to resource limitations - often resulting in missing elements and challenges in updating (K. Liu, Hogan, and Crowley 2011).

AI has transformed this paradigm by introducing methodologies that support both semi-automatic and automatic ontology construction. The goal is to convert new knowledge into ontological form, thereby enabling processing techniques such as semantic search and retrieval. By minimizing the need for manual intervention, these approaches promise a substantial reduction of time and labor costs required to build ontologies. However, a central challenge lies in converting unstructured data into meaningful ontological representations.

3. Reasoning and symbolic logic

Reasoning plays a pivotal role in extracting relevant data and formalizing unstructured data, as symbolic logic supports the construction of meaningful and accurate ontological representations by enabling machines to interpret, infer, and manage complex within ontologies. Chain-of-thought (CoT) reasoning, used as demonstrations, can enhance reasoning performance (Fu et al. 2023) and CoT prompting significantly improves large language models (LLMs) reasoning abilities, particularly in complex scenarios, by breaking down intricate problems into smaller, manageable steps (Chu et al. 2024). The reasoning capabilities of LLMs can significantly enhance their contribution to ontology construction and, by anchoring unstructured texts to domain-specific vocabularies, LLMs apply domain-specific rules more effectively while maintaining scalability across multiple fields.

By leveraging their deduction ability, LLMs can assist in identifying logical relationships, refining conceptual hierarchies, and suggesting coherent extensions to existing ontologies. When integrated with symbolic reasoning frameworks, LLMs not only support the dynamic evolution of ontologies but also contribute to maintaining semantic consistency and domain relevance, resulting in more adaptable KOSs. For instance, H. Liu, Perl, and Geller (2020) introduce a BERT-based model trained to automatically predict IS-A relationships between new and existing concepts. By modeling sentence pairs where one concept logically follows from another, the approach successfully predicts subsumption relationships, thereby reducing errors in concept placement and improving the hierarchical integration of new terms. This model can not only identify potential parents of a new concept, but also filter out irrelevant concepts, reducing the number of improper placement choices for a concept.

Moreover, symbolic logic also supports the verification of the consistency and coherence of ontological models. Logical reasoning also allows for the detection of contradictions or inconsistencies in the ontologies, ensuring that the data and relationships represented are valid and reliable (Ma, Molnár, and Benczúr 2021). Beyond verification, reasoning mechanisms can drive the automated evolution of ontologies by dynamically integrating new information, refining conceptual structures, and refining concepts based on logical rules and inferences. In this regard, Jabla et al. (2021) propose an automatic, ontology-based model evolution approach designed to function in dynamic runtime environments. Their method analyzes heterogeneous, semi-structured data to learn and extend ontological models, allowing the ontology to evolve alongside changes in its application context.

LLMs applied to ontology learning (OL) and ontology engineering have the transformative potential of combining symbolic logic, enabling systems to handle complexity, ensure reliability, and adapt to evolving needs.

4. AI for the construction and populating of ontologies

The integration of AI in ontology construction and maintenance represents a significant paradigm shift in knowledge representation and management.

By definition, OL encompasses the derivation of relevant knowledge - including concepts, properties, relations, and axioms - from input data to either construct new ontologies or enrich existing ones. Current OL approaches leverage NLP and ML techniques to extract and structure information from unstructured or semi-structured data sources, employing methods such as word embeddings, clustering and text mining algorithms, that identify patterns and hierarchical structures through semi-automatic or fully automatic construction and refinement procedures. Beyond basic extraction, NLP integration enrich and refine the linguistic realization of ontologies, where computational linguistics methods have proven instrumental in partially or fully automating semantic knowledge extraction processes (K. Liu, Hogan, and Crowley 2011). This comprehensive approach streamlines the ontology development process while offering promising pathways for extending existing ontologies with newly discovered knowledge.

Additionally, LLMs have shown considerable potential in addressing the various interconnected subtasks inherent in OL (Asim et al. 2018). A notable example is the OLLM framework, which proposes a general and scalable method for building the taxonomic backbone of an ontology from scratch through an end-to-end framework. Rather than focusing on individual relations between concepts, OLLM employs a finetuned LLM to model entire subcomponents of target ontologies (Lo et al. 2024). LLMs4OL also employs LLMs, but focuses specifically on extracting relations among ontology classes and instances; its scope remains limited to entity relationships rather than comprehensive ontology generation (Giglou, D'Souza, and Auer 2023).

Despite the advancements, ontology construction still demands significant human expertise, as shown with DRAGON-AI. It is an LLM-backed method which specializes in ontology term completion, transforming partially completed ontology terms into comprehensive objects containing all requisite components including textual descriptions, logical definitions, and inter-term relationships. The system operates by constructing contextually informed prompts that guide LLM processing, subsequently parsing results to generate complete term objects. The AI-generated term definitions were found to be decent but not as good as human-generated definitions (Toro et al. 2024).

Another notable contribution is OntoGenix, which offers an LLM-powered pipeline for ontology development using GPT-4. Experiments on e-commerce datasets show that OntoGenix modeling is consistent, but some OntoGenix ontologies include hallucinations, such as entities modeled simultaneously as both data and object properties, and inappropriate equivalence relations. The qualitative evaluation of the ontologies revealed that, although OntoGenix ontologies are rich in annotations and often produce a similar number of classes and properties, they are still behind human reasoning when modeling complex contexts. Compared to human-generated ontologies, OntoGenix ontologies present axioms with lower expressivity in class restrictions and properties, less accurate linking with entities from external ontologies, and a less nuanced understanding of OWL modeling (Val-Calvo et al. 2025).

Recent studies have also explored how prompt engineering can help in navigating the complexities of knowledge modeling. In contrast to finetuning, prompt engineering stands out for its simplicity in implementation and adaptability, offering an approach for enhancing knowledge engineering processes while avoiding the need for large, labelled datasets and dedicated finetuned models. For instance, Lippolis, Saeedizade, et al. (2025) and Lippolis, Ceriani, et al. (2025) describe a framework to assess LLM-pipelines with prompting techniques. These methods reveal recurrent issues in automatically generated ontologies, such as lack of annotations, inverse relations, and inconsistent domain or range assignments.

In contrast, one interesting study proposed by Mateiu and Groza (2023) explores the potential of a fine-tuned GPT-3 model to translate natural language into description logic, specifically into OWL functional syntax. The tool is constructed as a Protégé plugin that supports the development of an ontology from scratch or the enrichment of an existing ontology. This plugin aims to be a support tool that saves ontology engineering time, reducing efforts needed for ontology development by automating repetitive tasks. However, challenges remain in achieving structurally and syntactically correct outputs.

Building on the assumption that LLMs encapsulate a substantial amount of domain-specific knowledge during the comprehensive pre-training phase, KGFiller is a framework for semi-automatic ontology population and its validity is assessed through a case study on the food domain. Beginning with an initial schema consisting of interrelated classes and properties and a set of query templates, this method repeatedly queries an LLM to generate instances for both classes and properties from the responses. As a result, an ontology is quickly and automatically enriched with instances, which experts may consider keeping, adjusting, discard, or using according to their own needs and expertise (Ciatto et al. 2025).

To further enhance the capabilities of LLMs in specific domain, such as medical, financial, and science, the research and industry community have begun to focus on developing domain-specific LLMs, such as BloombergGPT, BioMistral and LawGPT. In parallel, OntoTune aims to align LLMs with domain ontology through in-context learning and generate responses guided by the ontology. OntoTune enhances LLM performance through a three-phase workflow integrating structured ontology. First, instruction texts are generated using three ontology-aware prompts focused on diversity, conceptual clarity, and professionalism, encouraging the model to refine responses through in-context learning. Second, inconsistencies are identified by comparing outputs with and without ontology input - significant differences indicate the model has not internalized the ontological structure, making these instances valuable training data. Finally, the model undergoes self-training on this dataset to produce domain-aligned versions of the LLM. This approach demonstrates that even a small but well-constructed ontology can significantly enhance the capabilities of LLMs in specialized domains (Z. Liu et al. 2025).

Interestingly, although ontologies do not directly supply the practical knowledge needed to answer specific queries, the structured nature of ontological input appears to help models better organize and internalize domain-specific knowledge. Altogether, these studies illustrate the evolving role of LLMs in ontology engineering, showing how they can reduce modeling inconsistencies, produce robust and scalable ontologies. Besides increasing the automation of the development of ontologies, LLMs may help to reduce inaccuracies and inconsistencies inherent to manual methods, ensuring the creation of more reliable and robust ontologies.

5. Ontologies and Machine Learning

Although AI-driven techniques can streamline ontology construction, ontologies also play a crucial role in improving the interpretability and transparency of ML models. Ontologies enable ML models to incorporate domain-specific constraints and rules, introduce hierarchical and domain-specific reasoning, and enhancing explainability, thus allowing experts and users to better understand the reasoning behind model outputs (Zmiivskyi, Danova, and Feoktystova 2023; Kulmanov et al. 2020). These ontological approaches integrate diverse data sources, optimize model parameters, and make AI systems more interpretable, reliable, and adaptable (Zmiivskyi, Danova, and Feoktystova 2023).

Ontologies are also widely used in both supervised and unsupervised learning. In unsupervised learning, they are applied for semantic similarity analysis (Kulmanov et al. 2021), embedding techniques for explaining trained neural networks (Sarker et al. 2017), but they also facilitate clustering and data pattern analysis using linked data (Tiddi, d'Aquin, and Motta 2015). In supervised learning, ontologies contribute to explainable recommendations (Ai et al. 2018) and transfer learning explanation (Geng et al. 2019).

Ontologies are particularly employed in describing two types of ML models: supervised classification tasks with neural networks and unsupervised embedding tasks. When combined with deep neural networks, knowledge graphs have been applied as embedding techniques. In this context, applying a knowledge graph as an embedding technique involves using the structured data from the knowledge graph to create vector representations (embeddings) of entities or concepts within the graph. This allows the incorporation of the rich semantic relationships captured by the graph into the machine learning models. It could potentially also improve explanation capabilities and enhance the performance of models applied in various domains. Knowledge graph embeddings transform entities and relationships from a graph into continuous vector spaces, which allows machine learning algorithms to utilize the semantic structure and information encoded in the KG. This enhances the model's ability to understand, represent and reason about complex relationships. Multiple approaches including TransE, TransH, and DistMult exist which result in downstream tasks, such as clustering, classification, or enhancement of deep learning models through providing contextually meaningful input data (Dadoun 2023; Sellami et al. 2025).

6. Different domain applications

Different domains can benefit from different ontology construction techniques.

In smart farming ontologies have been integrated with knowledge graphs to enhance decision-making processes, enabling accurate crop yield predictions (Choudhary et al. 2020). Similarly, in healthcare, ontology-based methods have been using NLP techniques to generate high-quality training data for weakly supervised ML models; enhance Named Entity Recognition (NER), event detection, and semantic querying; and integrate semantic constraints (Magumba, Nabende, and Mwebaze 2018; Batbaatar and Ryu 2019; Kulmanov et al. 2020; Deng et al. 2021; Zmiivskyi, Danova, and Feoktystova 2023). Particularly prominent in the healthcare field, with applications like disease prediction (Kulmanov et al. 2020; Carvalho, Oliveira, and Pesquita 2023), drug-drug interaction analysis (Herrero-Zazo et al. 2015), and EHR mining (Robinson and Haendel 2020).

In other domains it is possible to find some implementations in (a) Business and Management, (b) Engineering and Construction, (c) Sustainable Operations, (d) Tourism and Hospitality Studies (Thayyib et al. 2023), (e) Education (Silva, Mutz, and Ruy 2022), and (f) Enterprise applications (Erekhinskaya, Strebkov, et al. 2020).

Ontologies enable multi-domain customization in tasks such as semantic search and question answering (Erekhinskaya, Strebkov, et al. 2020; Erekhinskaya, Tatu, et al. 2020). Given the complexity of information retrieval due to the growing use of data mining, ontology-based retrieval systems have been applied to enhance interfaces between data and search queries, giving back result sets closer to users' research needs (Munir and Sheraz Anjum 2018). Semantic-based approaches, incorporating domain ontologies, have been adapted for data modeling and information retrieval, supporting visual or interactive query formulation, ontology-based information linking (keyword search), and ontology-based query refinement. Despite these advancements, challenges remain, including the lack of robust interdisciplinary evaluation (Ronzano and Nanavati 2024; Petrenko et al. 2023) and insufficient systematic metrics to evaluate ontology contributions to overall AI system performance (Zmiivskyi, Danova, and Feoktystova 2023; Chen et al. 2025). Addressing these gaps will be crucial in advancing AI-driven ontology development and ensuring their effective application across diverse fields. Nevertheless, the convergence of ontologies and ML shows promise for creating more dynamic and adaptive knowledge representation systems.

7. Ethical considerations

The integration of AI into ontology construction and maintenance marks a significant paradigm shifts in knowledge representation. AI systems can suggest new terms, identify hierarchical relationships, and detect inconsistencies within existing ontological structures. However, this transformation brings challenges and concerns. One of the most critical concerns is represented by hallucinations, which are plausible, but factually incorrect or misleading information generated by AI models. When applied to ontology construction, such outputs can introduce inaccuracies and compromise the semantic integrity and reliability of the resulting KOS.

Closely related to this is the pervasive issue of biases, that can emerge at every stage of the AI development lifecycle, from data collection and model training to deployment and evaluation, leading to multiple forms of bias - computational, historical, human, institutional, societal, and systemic (Schwartz et al. 2022). These biases often stem from underlying assumptions and may be either explicit or implicit (Rovetto 2024). When embedded in AI systems, such assumptions risk distorting the ontologies they help to build, particularly in sensitive or underrepresented domains.

Therefore, the standardization of AI concepts, including the identification and mitigation of bias, is essential for developing AI technologies that are effective and equitable (Joachimiak et al. 2024). While many general approaches try to ensure technological safety and accountability, the challenges posed by AI demand new perspectives. Current efforts to mitigate AI bias have primarily focused on computational aspects, such as ensuring dataset representativeness and algorithmic fairness. These are important first steps, but more comprehensive strategies are needed, especially in the context of knowledge organization, where the consequences of conceptual distortion can be harmful, as this can reinforce existing inequalities.

A crucial aspect of these strategies is represented by the terminological consistency. Promoting shared definitions and standardized conceptual frameworks helps reduce ambiguity and miscommunication, especially as AI methods are introduced into new domains. In this direction, the AI Ontology (AIO) plays a pivotal role by advancing the standardization and understanding of AI concepts and methodologies. Through its comprehensive terminological framework, AIO facilitates clearer communication, collaboration, and result sharing within the AI community (Joachimiak et al. 2024).

In conclusion, while AI brings transformative capabilities to KOSs construction, its integration must be approached with critical awareness. Addressing the risks of hallucinations and biases, and promoting clarity and transparency through standardization, is essential to ensure that AI-enhanced ontologies remain accurate, equitable, and trustworthy.

8. Conclusion

The relationship between ontologies and AI is increasingly recognized as bidirectional. While AI accelerates the ontology construction process, ontologies enable AI systems to enhance their performance, particularly from a semantic perspective.

Additionally, the current state of art suggests that the most successful applications of AI in ontology workflow follow a collaborative model, where AI serves as an intelligent assistant to human experts. This hybrid approach leverages the computational power of AI while maintaining the critical thinking, domain expertise and judgment that only human experts can provide so far.

9. References

Ai, Qingyao, Vahid Azizi, Xu Chen, and Yongfeng Zhang. 2018. ‘Learning Heterogeneous Knowledge Base Embeddings for Explainable Recommendation’. *Algorithms* 11 (9): 137. https://doi.org/10.3390/a11090137.

Asim, Muhammad Nabeel, Muhammad Wasim, Muhammad Usman Ghani Khan, Waqar Mahmood, and Hafiza Mahnoor Abbasi. 2018. ‘A Survey of Ontology Learning Techniques and Applications’. *Database* 2018 (January):bay101. https://doi.org/10.1093/database/bay101.

Batbaatar, Erdenebileg, and Keun Ho Ryu. 2019. ‘Ontology-Based Healthcare Named Entity Recognition from Twitter Messages Using a Recurrent Neural Network Approach’. *International Journal of Environmental Research and Public Health* 16 (19): 3628. https://doi.org/10.3390/ijerph16193628.

Carvalho, Ricardo M. S., Daniela Oliveira, and Catia Pesquita. 2023. ‘Knowledge Graph Embeddings for ICU Readmission Prediction’. *BMC Medical Informatics and Decision Making* 23 (1): 12. https://doi.org/10.1186/s12911-022-02070-7.

Chen, Jiaoyan, Olga Mashkova, Fernando Zhapa-Camacho, Robert Hoehndorf, Yuan He, and Ian Horrocks. 2025. ‘Ontology Embedding: A Survey of Methods, Applications and Resources’. arXiv. https://doi.org/10.48550/arXiv.2406.10964.

Choudhary, Nitu Kedarmal, Sai Sree Laya Chukkapalli, Sudip Mittal, Maanak Gupta, Mahmoud Abdelsalam, and Anupam Joshi. 2020. ‘YieldPredict: A Crop Yield Prediction Framework for Smart Farms’. In *2020 IEEE International Conference on Big Data (Big Data)*, 2340–49. https://doi.org/10.1109/BigData50022.2020.9377832.

Chu, Zheng, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2024. ‘Navigate through Enigmatic Labyrinth A Survey of Chain of Thought Reasoning: Advances, Frontiers and Future’. arXiv. https://doi.org/10.48550/arXiv.2309.15402.

Ciatto, Giovanni, Andrea Agiollo, Matteo Magnini, and Andrea Omicini. 2025. ‘Large Language Models as Oracles for Instantiating Ontologies with Domain-Specific Knowledge’. *Knowledge-Based Systems* 310 (February):112940. https://doi.org/10.1016/j.knosys.2024.112940.

Dadoun, Amine. 2023. ‘Knowledge Graph Embeddings 101’. *Towards Data Science* (blog). 9 April 2023. https://towardsdatascience.com/knowledge-graph-embeddings-101-2cc1ca5db44f/.

Deng, Shumin, Ningyu Zhang, Luoqiu Li, Chen Hui, Tou Huaixiao, Mosha Chen, Fei Huang, and Huajun Chen. 2021. ‘OntoED: Low-Resource Event Detection with Ontology Embedding’. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2828–39. Online: Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.acl-long.220.

Erekhinskaya, Tatiana, Dmitriy Strebkov, Sujal Patel, Mithun Balakrishna, Marta Tatu, and Dan Moldovan. 2020. ‘Ten Ways of Leveraging Ontologies for Natural Language Processing and Its Enterprise Applications’. In *Proceedings of The International Workshop on Semantic Big Data*, 1–6. SBD ’20. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3391274.3393639.

Erekhinskaya, Tatiana, Marta Tatu, Mithun Balakrishna, Sujal Patel, Dmitry Strebkov, and Dan Moldovan. 2020. ‘Ten Ways of Leveraging Ontologies for Rapid Natural Language Processing Customization for Multiple Use Cases in Disjoint Domains’ 7 (1).

Fu, Yao, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023. 'COMPLEXITY-BASED PROMPTING FOR MULTI-STEP REASONING'.

Geng, Yuxia, Jiaoyan Chen, Ernesto Jimenez-Ruiz, and Huajun Chen. 2019. 'Human-Centric Transfer Learning Explanation via Knowledge Graph [Extended Abstract]'. arXiv. https://doi.org/10.48550/arXiv.1901.08547.

Giglou, Hamed Babaei, Jennifer D'Souza, and Sören Auer. 2023. 'LLMs4OL: Large Language Models for Ontology Learning'. arXiv. https://doi.org/10.48550/arXiv.2307.16648.

Gruber, Thomas R. 1993. 'A Translation Approach to Portable Ontology Specifications'. *Knowledge Acquisition* 5 (2): 199–220. https://doi.org/10.1006/knac.1993.1008.

Herrero-Zazo, María, Isabel Segura-Bedmar, Janna Hastings, and Paloma Martínez. 2015. 'Application of Domain Ontologies to Natural Language Processing: A Case Study for Drug-Drug Interactions'. *International Journal of Information Retrieval Research (IJIRR)* 5 (3): 19–38. https://doi.org/10.4018/IJIRR.2015070102.

Jabla, Roua, Maha Khemaja, Félix Buendia, and Sami Faiz. 2021. 'Automatic Ontology-Based Model Evolution for Learning Changes in Dynamic Environments'. *Applied Sciences* 11 (22): 10770. https://doi.org/10.3390/app112210770.

Joachimiak, Marcin P, Mark A Miller, J Harry Caufield, Ryan Ly, Nomi L Harris, Andrew Tritt, Christopher J Mungall, and Kristofer E Bouchard. 2024. 'The Artificial Intelligence Ontology: LLM-Assisted Construction of AI Concept Hierarchies'. *Applied Ontology* 19 (4): 408–18. https://doi.org/10.1177/15705838241304103.

Kulmanov, Maxat, Fatima Zohra Smaili, Xin Gao, and Robert Hoehndorf. 2020. 'Machine Learning with Biomedical Ontologies'. bioRxiv. https://doi.org/10.1101/2020.05.07.082164.

———. 2021. 'Semantic Similarity and Machine Learning with Ontologies'. *Briefings in Bioinformatics* 22 (4): bbaa199. https://doi.org/10.1093/bib/bbaa199.

Lippolis, Anna Sofia, Miguel Ceriani, Sara Zuppiroli, and Andrea Giovanni Nuzzolese. 2025. 'Ontogenia: Ontology Generation with Metacognitive Prompting in Large Language Models'. In *The Semantic Web: ESWC 2024 Satellite Events*, edited by Albert Meroño Peñuela, Oscar Corcho, Paul Groth, Elena Simperl, Valentina Tamma, Andrea Giovanni Nuzzolese, Maria Poveda-Villalón, et al., 15344:259–65. Lecture Notes in Computer Science. Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-78952-6_38.

Lippolis, Anna Sofia, Mohammad Javad Saeedizade, Robin Keskisärkkä, Sara Zuppiroli, Miguel Ceriani, Aldo Gangemi, Eva Blomqvist, and Andrea Giovanni Nuzzolese. 2025. 'Ontology Generation Using Large Language Models'. arXiv. https://doi.org/10.48550/arXiv.2503.05388.

Liu, Hao, Yehoshua Perl, and James Geller. 2020. 'Concept Placement Using BERT Trained by Transforming and Summarizing Biomedical Ontology Structure'. *Journal of Biomedical Informatics* 112 (December):103607. https://doi.org/10.1016/j.jbi.2020.103607.

Liu, Kaihong, William R Hogan, and Rebecca S Crowley. 2011. 'Natural Language Processing Methods and Systems for Biomedical Ontology Learning'. *Journal of Biomedical Informatics* 44 (1): 163–79. https://doi.org/10.1016/j.jbi.2010.07.006.

Liu, Zhiqiang, Chengtao Gan, Junjie Wang, Yichi Zhang, Zhongpu Bo, Mengshu Sun, Huajun Chen, and Wen Zhang. 2025. 'OntoTune: Ontology-Driven Self-Training for

Aligning Large Language Models'. arXiv. https://doi.org/10.48550/arXiv.2502.05478.

Lo, Andy, Albert Q. Jiang, Wenda Li, and Mateja Jamnik. 2024. 'End-to-End Ontology Learning with Large Language Models'. arXiv. https://doi.org/10.48550/arXiv.2410.23584.

Ma, Chuangtao, Bálint Molnár, and András Benczúr. 2021. 'A Semi-Automatic Semantic Consistency-Checking Method for Learning Ontology from Relational Database'. *Information* 12 (5): 188. https://doi.org/10.3390/info12050188.

Magumba, Mark Abraham, Peter Nabende, and Ernest Mwebaze. 2018. 'Ontology Boosted Deep Learning for Disease Name Extraction from Twitter Messages'. *Journal of Big Data* 5 (1): 31. https://doi.org/10.1186/s40537-018-0139-2.

Mateiu, Patricia, and Adrian Groza. 2023. 'Ontology Engineering with Large Language Models'. In *2023 25th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC)*, 226–29. https://doi.org/10.1109/SYNASC61333.2023.00038.

Munir, Kamran, and M. Sheraz Anjum. 2018. 'The Use of Ontologies for Effective Knowledge Modelling and Information Retrieval'. *Applied Computing and Informatics* 14 (2): 116–26. https://doi.org/10.1016/j.aci.2017.07.003.

Naqvi, Muhammad Raza, Muhammad Waseem Iqbal, Muhammad Usman Ashraf, Shafiq Ahmad, Ahmed T. Soliman, Shahzada Khurram, Muhammad Shafiq, and Jin-Ghoo Choi. 2021. 'Ontology Driven Testing Strategies for IoT Applications'. *Computers, Materials & Continua* 70 (3): 5855–69. https://doi.org/10.32604/cmc.2022.019188.

Petrenko, Mykola, Ellen Cohn, Oleksandr Shchurov, and Kyrylo Malakhov. 2023. 'Ontology-Driven Computer Systems: Elementary Senses in Domain Knowledge Processing'. *South African Computer Journal* 35 (2). https://doi.org/10.18489/sacj.v35i2.17445.

Robinson, Peter N., and Melissa A. Haendel. 2020. 'Ontologies, Knowledge Representation, and Machine Learning for Translational Research: Recent Contributions'. *Yearbook of Medical Informatics* 29 (01): 159–62. https://doi.org/10.1055/s-0040-1701991.

Ronzano, Francesco, and Jay Nanavati. 2024. 'Towards Ontology-Enhanced Representation Learning for Large Language Models'. arXiv. https://doi.org/10.48550/arXiv.2405.20527.

Rovetto, Robert J. 2024. 'The Ethics of Conceptual, Ontological, Semantic and Knowledge Modeling'. *AI & SOCIETY* 39 (4): 1547–68. https://doi.org/10.1007/s00146-022-01563-3.

Sarker, Md Kamruzzaman, Ning Xie, Derek Doran, Michael Raymer, and Pascal Hitzler. 2017. 'Explaining Trained Neural Networks with Semantic Web Technologies: First Steps'. arXiv. https://doi.org/10.48550/arXiv.1710.04324.

Schwartz, Reva, Apostol Vassilev, Kristen Greene, Lori Perine, Andrew Burt, and Patrick Hall. 2022. 'Towards a Standard for Identifying and Managing Bias in Artificial Intelligence'. NIST SP 1270. Gaithersburg, MD: National Institute of Standards and Technology (U.S.). https://doi.org/10.6028/NIST.SP.1270.

Sellami, Dorsaf, Wissem Inoubli, Imed Riadh Farah, and Sabeur Aridhi. 2025. 'Knowledge Graph Representation Learning: A Comprehensive and Experimental

Overview'. *Computer Science Review* 56 (May):100716. https://doi.org/10.1016/j.cosrev.2024.100716.

Silva, Eduardo Moreira da, Filipe Wall Mutz, and F. Ruy. 2022. 'Uso de Ontologias No Suporte a Aplicação de Machine Learning: Um Caso No Domínio de Evasão Escolar'. In . https://www.semanticscholar.org/paper/Uso-de-Ontologias-no-Suporte-a-Aplica%C3%A7%C3%A3o-de-Machine-Silva-Mutz/2a4fdd1675bc4107d0633e6c5dc2d2b6d2eecc2f.

Thayyib, P. V., Rajesh Mamilla, Mohsin Khan, Humaira Fatima, Mohd Asim, Imran Anwar, M. K. Shamsudheen, and Mohd Asif Khan. 2023. 'State-of-the-Art of Artificial Intelligence and Big Data Analytics Reviews in Five Different Domains: A Bibliometric Summary'. *Sustainability* 15 (5): 4026. https://doi.org/10.3390/su15054026.

Tiddi, Ilaria, Mathieu d'Aquin, and Enrico Motta. 2015. 'Data Patterns Explained with Linked Data'. In *Machine Learning and Knowledge Discovery in Databases*, edited by Albert Bifet, Michael May, Bianca Zadrozny, Ricard Gavalda, Dino Pedreschi, Francesco Bonchi, Jaime Cardoso, and Myra Spiliopoulou, 9286:271–75. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-23461-8_28.

Toro, Sabrina, Anna V. Anagnostopoulos, Sue Bello, Kai Blumberg, Rhiannon Cameron, Leigh Carmody, Alexander D. Diehl, et al. 2024. 'Dynamic Retrieval Augmented Generation of Ontologies Using Artificial Intelligence (DRAGON-AI)'. *Journal of Biomedical Semantics* 15 (1): 19. https://doi.org/10.1186/s13326-024-00320-3.

Val-Calvo, Mikel, Mikel Egaña Aranguren, Juan Mulero-Hernández, Ginés Almagro-Hernández, Prashant Deshmukh, José Antonio Bernabé-Díaz, Paola Espinoza-Arias, José Luis Sánchez-Fernández, Juergen Mueller, and Jesualdo Tomás Fernández-Breis. 2025. 'OntoGenix: Leveraging Large Language Models for Enhanced Ontology Engineering from Datasets'. *Information Processing & Management* 62 (3): 104042. https://doi.org/10.1016/j.ipm.2024.104042.

Zmiivskyi, Volodymyr, Mariia Danova, and Olena Feoktystova. 2023. 'APPLICATION OF ONTOLOGICAL ENGINEERING METHODS IN OPTIMIZATION OF COMPLEX ML MODELS AND AI APPLICATIONS'. *Grail of Science*, no. 28 (June), 202–4. https://doi.org/10.36074/grail-of-science.09.06.2023.32.