Formal Foundations Of Ontological Metamodeling*

Zekeri Adams*

Comenius University in Bratislava, Faculty of Mathematics, Physics, and Informatics, Mlynská dolina, 842 48 Bratislava, Slovakia

Abstract

Ontological metamodeling offers a promising approach for representing complex domains through structured, hierarchical models that support multiple levels of abstraction. However, many existing metamodeling approaches lack rigorous formal foundations and well-defined semantics, limiting their reliability, scalability, and applicability in tasks such as validation, reasoning, and consistency checking. In this research, we formalize the PURO metamodeling language and verify the correctness of its formalization through automated theorem prover and knowledge-base validation. We further investigate the decidability of metamodelled description logics with fixed interpretations of the subClassOf and instanceOf relations under full set-theoretic semantics of subClassOf. Finally, we examine the capability of large language models (LLMs) to generate ontologies that correctly preserve and represent multi-level semantic structures.

Keywords

Description logics, Metamodelling, Ontologies, Large Language Models

1. Introduction

Ontologies play a central role in knowledge representation by providing formal and machine-interpretable models of a domain [1, 2]. Their success in areas such as the Semantic Web, artificial intelligence, life sciences, cybersecurity, and digital engineering stems from their ability to represent structured knowledge and support automated reasoning. As knowledge graphs and ontology-based systems continue to grow in size and complexity, there is an increasing need to represent not only domain entities but also the conceptual structures used to organize and classify them.

Ontological metamodeling addresses this challenge by enabling concepts, relations, and other universal entities to be treated as modeling objects in their own right. Through metamodeling, entities can participate simultaneously at multiple levels of abstraction, allowing the representation of taxonomies, conceptual hierarchies, modeling languages, and other multi-level structures within a unified framework. Such capabilities are essential in domains where concepts themselves must be described, classified, and related, including biological taxonomies, enterprise architectures, knowledge graphs, and ontology engineering frameworks.

Despite its importance, ontological metamodeling remains a challenging problem. Existing approaches often involve trade-offs between expressivity, semantic precision, and computational tractability. Description-logic-based [3, 4, 5, 6, 7, 8] approaches typically restrict the forms of metamodeling that can be represented, while conceptual metamodelling approaches formalized in first-order logic often lack formal verification, and proper verification and validation of practical applications. Furthermore, the growing adoption of large language models (LLMs) for ontology engineering introduces new challenges regarding the preservation of semantic correctness across multiple abstraction levels.

This research investigates the theoretical foundations, formal semantics, and practical applications of ontological metamodeling. The work is organized around four complementary research directions. First, we formalize the PURO metamodeling language and provide a first-order logic semantics that supports consistency checking and validation of metamodelled knowledge bases. Second, we investigate

Proceedings of FOIS 2026 Satellite events co-located with the 16th International Conference on Formal Ontology in Information Systems (FOIS 2026), September 21-25, 2026, Vitória, Brazil.

*Corresponding author.

zekeri.adams@fmph.uniba.sk (Z. Adams)

0009-0006-3413-2409 (Z. Adams)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-11-27

CEUR-WS.org/Vol-4342/paper13.pdf

the decidability boundaries of higher-order description logics with metamodelled instantiation and subsumption relations, focusing on the logic $\mathcal{HIRS}_{SA}(\mathcal{SROIQ})$. Third, we examine the capability of large language models to generate ontologies that correctly preserve multi-level semantic structures and metamodeling constraints. Fourth, we apply these foundations in the cybersecurity domain through the development of MAECO-Lite, a UFO-grounded malware-analysis ontology designed to provide a semantically precise representation of malware behavior, execution processes, and analysis artifacts while promoting interoperability with cybersecurity standards such as MAEC (Malware Attribute Enumeration and Characteristics) [9, 10], STIX (Structured Threat Information Expression)[11], and MITRE ATT&CK[12].

Together, these research directions aim to strengthen the formal foundations of ontological metamodeling while demonstrating their practical value in ontology engineering and knowledge graph development. By combining formal semantics, decidable reasoning, ontology validation, and domain-specific ontology engineering, this research seeks to improve the reliability, explainability, and interoperability of ontology-based knowledge representation systems.

2. Research Plan and Objectives

The overall goal of this research is to strengthen the theoretical and practical foundations of ontological metamodeling by combining formal semantics, automated reasoning, ontology validation, and ontology engineering. The research seeks to address limitations in existing metamodeling approaches regarding semantic precision, formal verification, decidability, and applicability to real-world ontology development. The research is guided by the following research questions:

RQ1 How can ontological metamodeling languages be formally specified and verified to ensure semantic correctness and consistency?

RQ2 Under what conditions can metamodelled description logics with instantiation and subsumption relations remain decidable while preserving expressive multi-level semantics?

RQ3 To what extent can large language models generate ontologies that correctly preserve multi-level semantic structures and metamodeling constraints?

RQ4 How can foundational ontology principles improve the representation of complex cybersecurity knowledge, particularly malware behavior and analysis processes?

To achieve our set goals, we organized our research around the following four complementary objectives:

Objective 1: Formalization and validation of metamodeling languages. The first objective is to formalize the PURO metamodeling language using first-order logic and to verify the correctness of its semantics through theorem proving and knowledge-base validation. This work aims to provide a rigorous semantic foundation for PURO and enable the detection of ontological inconsistencies in metamodelled knowledge bases.

Objective 2: Decidability of metamodelled description logics. The second objective is to investigate the decidability of higher-order description logics supporting metamodelled instantiation and subsumption. In particular, we study the logic $\mathcal{HIRS}_{SA}(\mathcal{SROIQ})$, which extends previous work by providing full set-theoretical semantics for the *subClassOf* relation over all concepts. The research focuses on developing reductions to classical description logics and decidable fragments of first-order logic, as well as identifying the expressive and computational boundaries of the framework.

Objective 3: Evaluation of LLMs for ontological metamodeling. The third objective is to investigate the capability of large language models (LLMs) to generate ontologies that correctly preserve multi-level semantic structures. The study will evaluate the extent to which LLM-generated ontologies satisfy metamodeling constraints, maintain ontological distinctions across abstraction levels, and align with formal ontology principles.

Objective 4: Application to cybersecurity ontology engineering. The fourth objective is to apply the developed foundations to the cybersecurity domain through the design of MAECO-Lite, a UFO-grounded ontology for malware analysis. This work investigates how foundational ontologies can improve the representation of malware behavior, execution processes, and analysis artifacts while ensuring interoperability with cybersecurity standards such as MAEC, STIX, and MITRE ATT&CK.

Expected Outcomes

The research is expected to yield the following outcomes:

- A formally verified first-order logic semantics for the PURO metamodeling language and methods for validating metamodelled knowledge bases.
- New theoretical results on the expressiveness and decidability of higher-order description logics with metamodelled instantiation and subsumption relations, particularly within the $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$ framework.
- A benchmark and evaluation methodology for assessing the ability of Large Language Models to generate ontologies that preserve multi-level semantic structures and metamodeling constraints.
- To develop a semantically grounded ontology for dynamic malware analysis that supports interoperability with cybersecurity standards such as MAEC, STIX, and MITRE ATT&CK.
- Scientific publications advancing the state of the art in ontological metamodeling, description logics, ontology engineering, and AI-assisted ontology generation.

Collectively, these outcomes are expected to contribute to the development of more reliable, explainable, and semantically rigorous ontology-based knowledge representation systems capable of supporting both theoretical research and real-world applications.

3. Research Progress

3.1. Formalization and Verification of the PURO Language

A central component of this research is the formalization and verification of the PURO metamodeling language. Although PURO provides a lightweight framework for representing multi-level knowledge structures, its semantics are largely informal, limiting the possibility of rigorous verification and automated consistency checking. To address this limitation, we developed a first-order logic formalization of PURO and investigated its semantic properties.

The formalization was encoded in the TPTP (Thousands of Problems for Theorem Provers) language and verified using the Vampire theorem prover [13]. The satisfiability results confirmed that the proposed axiomatization is logically consistent and admits valid interpretations, providing evidence for its semantic soundness. Beyond validating the axioms themselves, we applied the formalization to a collection of PURO models and verified their satisfiability through automated reasoning. The successful verification of all evaluated models demonstrates that the formalization can support the construction of logically coherent metamodelled knowledge bases.

To assess the practical applicability of the approach, we extended the analysis to Wikidata, a large-scale collaborative knowledge graph that exhibits multi-level classification structures. Using SPARQL-based analysis informed by the PURO formalization, we identified several structural inconsistencies and metamodeling anti-patterns. These results demonstrate the potential of formal metamodeling frameworks not only for ontology verification but also for the detection and analysis of semantic anomalies in real-world knowledge graphs. This work provides a formally verified semantic foundation for PURO and establishes a basis for future research on automated validation, consistency checking, and reasoning over metamodelled ontologies [14].

3.2. The Description Logic $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$

A major focus of this research is the investigation of expressive and decidable metamodelled description logics. Building upon the work of Kubincová et al. [3], we introduced the description logic $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$, which extends higher-order description logics with full set-theoretical semantics for the subClassOf relation. Unlike previous approaches, the proposed framework provides a uniform treatment of both instanceOf and subClassOf across multiple meta-levels, enabling richer forms of ontological metamodeling.

The current research has focused on the formal definition of the syntax and semantics of $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$ and on investigating its computational properties. In particular, we have analyzed the expressibility of the semantic constraints underlying the subClassOf relation in several decidable fragments of first-order logic, including Maslov's class $\overline{\mathcal{K}}$ [15], FO^3_- [16], the Guarded Fragment with Universal role (GFU) [17], and the Tri-Guarded Fragment (TGF) [17]. Our preliminary results show that the backward implication of the subClassOf semantics can be represented within Maslov's class $\overline{\mathcal{K}}$ and FO^3_-, while an approximation can be obtained in GFU and TGF through the introduction of a counterexample predicate. However, the expressibility of the converse implication remains unresolved in all investigated fragments. Consequently, the decidability status of the full $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$ semantics remains an open research question [18].

Current work focuses on establishing the decidability of $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$ by extending existing tableau-based reasoning procedures with unravelling techniques adapted to the metamodeling setting. The investigation explores restrictions on the interaction between the instanceOf and subClassOf relations and other expressive features of the logic, with the goal of controlling cyclic dependencies and preserving termination. Through this approach, we aim to identify some restricted form of $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$ that retain expressive metamodeling capabilities while achieving decidability that preserves the essential semantics of full set-theoretical metamodeling.

3.3. Evaluating Metamodeling Awareness in LLM-Generated Ontologies

As part of the research on ontological metamodeling, we investigated the ability of Large Language Models (LLMs) to generate ontologies that preserve multi-level semantic structures. Existing evaluations of LLM-based ontology generation primarily focus on structural correctness, consistency, and concept extraction, while largely overlooking metamodeling constraints and abstraction-level semantics. To address this gap, we developed a benchmark grounded in Multi-Level Theory (MLT) and introduced a taxonomy of metamodeling anti-patterns, including class-instance confusion, abstraction-level mixing, incorrect instantiation chains, and incorrect cross-level relations. Using this benchmark, we evaluated several state-of-the-art LLMs across multiple domains and prompting strategies.

The results show that although current LLMs can generate syntactically valid OWL ontologies, they consistently struggle to preserve multi-level semantic structures. The most frequent errors involve abstraction-level mixing and incorrect cross-level relations, indicating that the models rely primarily on surface-level pattern learning rather than principled metamodeling reasoning. While structured semantic templates improved performance in some domains, they did not eliminate these semantic errors. These findings highlight the need for metamodeling-aware evaluation frameworks and suggest that future ontology generation systems may require stronger integration of formal semantic constraints and neuro-symbolic reasoning techniques.

Building on these findings, our ongoing work focuses on extending the benchmark to additional application domains and ontology engineering scenarios. By incorporating a broader collection of ontologies with varying levels of complexity, abstraction depth, and metamodeling requirements, we aim to obtain a more comprehensive understanding of the generalization capabilities of LLMs and the domain-specific characteristics of metamodeling errors. Furthermore, we plan to investigate ontology-guided prompting techniques and the integration of formal reasoning services into the ontology generation process to improve the preservation of multi-level semantic structures.

3.4. MAECO-Lite Ontology

As part of this research, we developed MAECO-Lite, a lightweight ontology for malware analysis grounded in the Unified Foundational Ontology (UFO). The motivation for this work arose from the observation that existing cybersecurity standards and ontologies, such as MAEC, STIX, and MITRE ATT&CK, primarily focus on information exchange and interoperability, while providing limited ontological guidance for representing malware behavior, execution processes, and analysis artifacts.

The research began with a conceptual and ontological analysis of existing malware knowledge representations. Using UFO as a reference framework, we examined core concepts from MAEC, STIX, and MITRE ATT&CK and identified several semantic ambiguities, including the insufficient distinction between enduring and event like entities. This analysis informed the refinement and reinterpretation of key cybersecurity concepts from a foundational ontology perspective.

Based on these findings, we designed a UFO-grounded conceptual model that captures malware samples, processes, actions, techniques, files, registry artifacts, and dynamic analysis features. Particular attention was given to the representation of malware behavior as the manifestation of underlying capabilities and to the separation between static characteristics and runtime phenomena. The resulting conceptual model was subsequently formalized as an OWL ontology and organized into modular components to support extensibility and reuse. The MAECO-Lite ontology has been implemented in OWL and evaluated through consistency checking and ontology inspection using standard Semantic Web tools.. Preliminary results indicate that the ontology provides a semantically coherent representation of malware analysis knowledge while maintaining compatibility with established cybersecurity standards. Ongoing work focuses on extending the ontology with additional malware-analysis concepts, evaluating its applicability to real-world malware datasets, and investigating its use in ontology-based machine learning and symbolic concept learning environments [19].

4. Related Work

We summarized the related work section based on the respective aspect of our research.

4.1. Multi-Level Modelling

Ontological metamodeling enables the representation of concepts, relations, and classification structures across multiple levels of abstraction, supporting the modeling of complex domains and conceptual hierarchies [20, 21]. To provide formal foundations for such representations, several metamodeling approaches have been proposed. Among them, Multi-Level Theory (MLT) [22] offers a first-order logic framework for modeling entities across different classification levels through formally defined instantiation and specialization relations. MLT has been successfully applied to the analysis of large-scale knowledge bases such as Wikidata [23], revealing the presence of multi-level anti-patterns and classification inconsistencies [24, 22]. However, existing work has primarily focused on the identification of modeling anomalies, with limited attention to the formal verification of metamodeling languages themselves and the validation of the logical foundations used to detect such inconsistencies.

PURO [25] is a metamodeling language that combines a graphical modeling notation with a conceptual framework for representing multi-level knowledge structures. While PURO has been applied in several ontology engineering scenarios, its semantics remain largely informal, limiting the ability to perform rigorous verification and automated consistency checking. In contrast to previous work, this research provides a first-order logic formalization of PURO and investigates its semantic correctness through automated theorem prover [13]. Furthermore, we evaluate the applicability of the formalization to real-world knowledge bases, enabling the identification and validation of structural inconsistencies in metamodelled knowledge. This contributes a formally verified semantic foundation for PURO and advances the broader goal of trustworthy and semantically grounded ontological metamodeling.

4.2. Subsumption Metamodeling in Description Logics

Metamodeling in description logics has traditionally focused on the representation of instantiation relationships across multiple abstraction levels. Extending this line of research, Kubincová et al. [3] introduced subsumption metamodeling through the description logic family $\mathcal{HIRS}_*(\mathcal{SROIQ})$, where the relations instanceOf and subClassOf are treated as ordinary roles. This approach enables explicit representation of hierarchical relationships at the meta-level while supporting automated reasoning over classification structures. The authors distinguish between non-set-theoretical and set-theoretical interpretations of the subClassOf relation. In the non-set-theoretical setting, subClassOf captures the intended meaning of specialization and is required to satisfy only the necessary condition that instances of a subclass are also instances of its superclass. In contrast, the set-theoretical interpretation identifies subsumption with inclusion between class extensions, thereby satisfying both the necessary and sufficient conditions for subclasshood.

Based on these interpretations, Kubincová et al. [3] established the decidability of several metamodelled description logics. The logics $\mathcal{HIRS}_{NN}(\mathcal{SROIQ})$ and $\mathcal{HIRS}_{NA}(\mathcal{SROIQ})$ support non-set-theoretical subsumption for named concepts and for all concepts, respectively, while $\mathcal{HIRS}_{SN}(\mathcal{SROIQ})$ extends the framework with set-theoretical subsumption for named concepts. These results provide an important foundation for reasoning in metamodelled ontologies. However, the decidability of a framework supporting full set-theoretical semantics for subClassOf over all concepts remains unresolved. Addressing this limitation is one of the central motivations for the development of $\mathcal{HIRS}_{SA}(\mathcal{SROIQ})$ investigated in this research.

4.3. Metamodelling in Large Language Models(LLMs)

Recent advances in Large Language Models (LLMs) have stimulated growing interest in their application to ontology engineering. Existing research has explored the use of LLMs for ontology generation, completion, concept extraction, relation discovery, competency question generation, and ontology-assisted modeling environments [26, 27, 28, 29, 30, 31]. Several studies have also proposed benchmarks and evaluation frameworks for assessing the quality of LLM-generated ontologies [32, 33, 34, 35]. While these approaches demonstrate the potential of LLMs to support ontology development, their evaluations primarily focus on structural correctness, extraction accuracy, consistency, and coverage.

A key limitation of existing work is the lack of evaluation from a metamodeling perspective. Current approaches rarely assess whether generated ontologies preserve semantic relationships across multiple levels of abstraction or satisfy formal metamodeling constraints. Although multi-level modeling frameworks such as MLT [36] provide rigorous foundations for representing cross-level semantic structures, these principles have not been systematically incorporated into the evaluation of LLM-generated ontologies. This research addresses this gap by investigating the ability of LLMs to generate ontologies that preserve multi-level semantics and by developing a benchmark and anti-pattern taxonomy for assessing metamodeling correctness.

4.4. Ontology for Malware domain

Several ontologies [37, 38, 39] and knowledge representation frameworks have been proposed to model malware and cybersecurity knowledge. Standards such as MAEC [9], STIX [11], and MITRE ATT&CK [12] provide rich vocabularies for describing malware samples, cyber observables, adversarial techniques, and threat intelligence. However, these standards were primarily designed for information exchange and interoperability rather than ontologically grounded conceptual modeling. Existing malware ontologies often focus on static malware characteristics, indicators of compromise, and threat intelligence integration, while providing limited support for representing the semantic distinctions between malware entities, runtime behavior, execution processes, and analysis artifacts. Moreover, many approaches lack an explicit foundational ontology, which can lead to ambiguities regarding the nature of malware behaviors, capabilities, and their manifestations during execution. MAECO-Lite addresses these limitations by adopting the Unified Foundational Ontology (UFO) [40, 41, 42, 43] as a

conceptual foundation and by reusing and refining concepts from MAEC, STIX, and MITRE ATT&CK to provide a semantically precise representation of malware behavior and dynamic analysis knowledge.

5. Summary of Contributions and Future Work

This research advances the formal and practical foundations of ontological metamodeling along four directions. First, it provides a first-order formalization of the PURO metamodeling language and verifies its semantic consistency using automated reasoning. Second, it introduces the description logic $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$, extending prior metamodeling frameworks with full set-theoretical semantics for subClassOf and examining its relationship to decidable fragments of first-order logic. Third, it evaluates the ability of large language models to generate ontologies that preserve multi-level semantics, supported by a benchmark and a taxonomy of metamodeling anti-patterns. Fourth, it develops MAECO-Lite, a UFO-grounded ontology for dynamic malware analysis that enables semantically precise and interoperable cybersecurity representations.

These results underscore the importance of rigorous semantic foundations for metamodeling and reveal remaining challenges, particularly regarding the decidability of expressive metamodelled description logics and the integration of formal constraints into AI-assisted ontology engineering.

Future work will focus on establishing decidability results for $\mathcal{HIRS}_{\mathcal{SA}}(\mathcal{SROIQ})$, extending LLM-based ontology evaluation to broader domains, and further developing MAECO-Lite through additional concepts and empirical validation on real-world malware datasets.

Acknowledgments

The author gratefully acknowledges the supervision of Dr. Martin Homola and the guidance of Ján Kľuka at the Department of Applied Informatics, Faculty of Mathematics, Physics and Informatics, Comenius University Bratislava, Slovakia.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT-4 for grammar and spelling checks. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] S. Staab, R. Studer (Eds.), Handbook on Ontologies, 2nd ed., Springer Publishing Company, Incorporated, 2009.

[2] S. Rudolph, Foundations of description logics, in: Reasoning Web. Semantic Technologies for the Web of Data - 7th International Summer School 2011, Galway, Ireland, August 23-27, 2011, Tutorial Lectures, volume 6848 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 76–136. doi:10.1007/978-3-642-23032-5_2.

[3] P. Kubincová, Higher-order description logics for metamodelling, Master's thesis, Comenius University Bratislava, 2016.

[4] G. D. Giacomo, M. Lenzerini, R. Rosati, Higher-order description logics for domain metamodeling, in: Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI-11), AAAI Press, 2011.

[5] B. Glimm, S. Rudolph, J. Völker, Integrated metamodeling and diagnosis in OWL 2, in: Proceedings of the International Semantic Web Conference (ISWC), 2010.

[6] M. Homola, J. Kľuka, V. Svátek, M. Vacura, Typed higher-order variant of SROIQ – why not?, in: Proceedings of the 2014 International Workshop on Description Logics (DL), 2014.

[7] B. Motik, On the properties of metamodeling in OWL, Journal of Logic and Computation 17 (2007) 617–637.
[8] E. Rohrer, P. Severi, R. Motz, A. Díaz, Metamodelling in an ontology network, in: Proceedings of the 8th International Workshop on Modular Ontologies (WOMO), CEUR-WS.org, 2014.
[9] MAEC Project, MAEC™ Vocabularies Specification Version 5.0, Technical Report, MITRE Corporation, 2015. URL: https://maecproject.github.io/releases/5.0/MAEC_Vocabularies_Specification.pdf, accessed: 2025-05-08.
[10] Malware Attribute Enumeration and Characterization (MAEC) Project, 2015. URL: https://maecproject.github.io/.
[11] B. Jordan, R. Piazza, T. Darley, STIX™ Version 2.1, Committee Specification 02, OASIS Cyber Threat Intelligence (CTI) Technical Committee, 2021. URL: https://docs.oasis-open.org/cti/stix/v2.1/cs02/stix-v2.1-cs02.html, approved 25 January 2021.
[12] MITRE Corporation, MITRE ATT&CK, Technical Report, 2025. URL: https://attack.mitre.org/, accessed: 2026-02-18.
[13] L. Kovács, A. Voronkov, First-order theorem proving and vampire, in: N. Sharygina, H. Veith (Eds.), Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings, volume 8044 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 1–35. URL: https://doi.org/10.1007/978-3-642-39799-8_1.
[14] Z. Adams, M. Homola, J. Kluka, V. Svátek, Leveraging meta-modelling language for ontology structuring and validation, in: Joint Proceedings of Posters, Demos, Workshops, and Tutorials of the 24th International Conference on Knowledge Engineering and Knowledge Management (EKAW-PDWT 2024) co-located with 24th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2024), Amsterdam, Netherlands, November 26-28, 2024, volume 3967 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024. URL: https://ceur-ws.org/Vol-3967/PD_paper_180.pdf.
[15] O. Fiuk, E. Kieronski, V. Michielini, On the complexity of Maslov's class K, in: Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2024, Tallinn, Estonia, July 8-11, 2024, ACM, 2024, pp. 35:1–35:14. URL: https://doi.org/10.1145/3661814.3662097. doi:10.1145/3661814.3662097.
[16] O. Fiuk, E. Kieronski, Alternating quantifiers in uniform one-dimensional fragments with an excursion into three-variable logic, Journal of Symbolic Logic (2024).
[17] S. Rudolph, M. Simkus, The triguarded fragment of first-order logic, in: LPAR-22. 22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Awassa, Ethiopia, 16-21 November 2018, volume 57 of *EPiC Series in Computing*, EasyChair, 2018, pp. 604–619. URL: https://doi.org/10.29007/m8ts. doi:10.29007/M8TS.
[18] Z. Adams, M. Homola, J. Kluka, Can full set-theoretical subsumption semantics in meta-modelled description logics be captured within decidable FOL fragments?, in: Proceedings of the 38th International Workshop on Description Logics - DL 2025, Opole, Poland, September 3-6, 2025, volume 4091 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2025. URL: https://ceur-ws.org/Vol-4091/paper36.pdf.
[19] Z. Adams, P. Švec, J. Kl'uka, R. Ploszek, M. Onoja, Š. Balogh, M. Homola, Maeco-lite: Modular ontology for dynamic malware analysis, arXiv preprint arXiv:2605.31199 (2026).
[20] A. Gangemi, Ontology design patterns for semantic web content, in: Proceedings of the International Semantic Web Conference (ISWC), 2005, pp. 262–276.
[21] N. Guarino, D. Oberle, S. Staab, What is an ontology?, in: Handbook on Ontologies, Springer, 2009.
[22] C. M. Fonseca, J. P. Almeida, G. Guizzardi, V. A. Almeida, Multi-level conceptual modeling: Theory, language, and application, 2016.
[23] Wikidata Contributors, Wikidata: A free and open knowledge base, 2026. URL: https://www.wikidata.org/wiki/Wikidata:Main_Page, accessed: 2026-06-14.
[24] V. Almeida, J. P. Almeida, Toward a well-founded theory for multi-level conceptual modeling (2015).

[25] V. Svátek, M. Homola, J. Kluka, M. Vacura, Mapping structural design patterns in OWL to ontological background models, in: Proceedings of the Seventh International Conference on Knowledge Capture (K-CAP '13), ACM, New York, NY, USA, 2013, pp. 117–120. doi:10.1145/2479832.2479847.

[26] P. Mateiu, A. Groza, Ontology engineering with large language models, in: 2023 25th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), IEEE, 2023, pp. 226–229.

[27] V. K. Kommineni, B. König-Ries, S. Samuel, From human experts to machines: An llm supported approach to ontology and knowledge graph construction, arXiv preprint arXiv:2403.08345 (2024).

[28] A. Lo, et al., End-to-end ontology learning with large language models, arXiv preprint (2024). Methods for building taxonomic backbones with LLMs and structural/semantic evaluation metrics.

[29] A. S. Lippolis, M. J. Saeedizade, R. Keskisärkkä, S. Zuppiroli, M. Ceriani, A. Gangemi, E. Blomqvist, A. G. Nuzzolese, Ontology generation using large language models, in: European Semantic Web Conference, Springer, 2025, pp. 321–341.

[30] M. L. Coutinho, J. P. A. Almeida, G. Guizzardi, Llm-based modeling assistance for textual ontology-driven conceptual modeling, in: Companion Proceedings of the 44th International Conference on Conceptual Modeling: Industrial Track, ER Forum, 8th SCME, Doctoral Consortium, Tutorials, Project Exhibitions, Posters and Demos Co-located with ER 2025, Poitiers, France, October 20-23, 2025, CEUR Workshop Proceedings, CEUR-WS.org, 2025, pp. 338–342. URL: https://ceur-ws.org/Vol-4099/ER25_PAD_Coutinho.pdf.

[31] B. Zhang, V. A. Carriero, K. Schreiberhuber, S. Tsaneva, L. S. González, J. Kim, J. de Berardinis, Ontochat: a framework for conversational ontology engineering using language models, in: European semantic web conference, Springer, 2024, pp. 102–121.

[32] J. Plu, O. M. Escobar, E. Trouillez, A. Gapin, R. Troncy, A comprehensive benchmark for evaluating llm-generated ontologies., in: HGAIS@ ISWC, 2024.

[33] H. Babaei Giglou, J. D'Souza, S. Auer, Llms4ol: Large language models for ontology learning, in: International semantic web conference, Springer, 2023, pp. 408–427.

[34] H. Ghanem, C. Cruz, Fine-tuning vs. prompting: evaluating the knowledge graph construction with llms, in: 3rd International Workshop on Knowledge Graph Generation from Text (Text2KG) Co-located with the Extended Semantic Web Conference (ESWC 2024), volume 3747, 2024, p. 7.

[35] M. J. Saeedizade, R. Alharbi, H. B. Giglou, A. S. Lippolis, E. Blomqvist, V. Tamma, F. Grasso, T. R. Payne, J. D'Souza, S. Auer, et al., A framework for assessing llm consistency in knowledge engineering, Semantic Web–Interoperability, Usability, Applicability (SWJ), under review (2025).

[36] V. A. Carvalho, J. P. A. Almeida, C. M. Fonseca, G. Guizzardi, Multi-level ontology-based conceptual modeling, Data & Knowledge Engineering 109 (2017) 3–24.

[37] I. R. Chowdhury, D. Bhowmik, Capturing malware behaviour with ontology-based knowledge graphs, in: IEEE Conference on Dependable and Secure Computing, DSC 2022, Edinburgh, UK, June 22-24, 2022, IEEE, 2022, pp. 1–7. URL: https://doi.org/10.1109/DSC54232.2022.9888860. doi:10.1109/DSC54232.2022.9888860.

[38] A. Bryniarska, W. Pokuta, Ontology-based knowledge representation in the iot cybersecurity system, Semantic Web Journal 0 (2023) 1–12. URL: https://doi.org/10.1016/j.cose.2019.101574.

[39] H.-S. Chiang, W.-J. Tsaur, et al., Ontology-based mobile malware behavioral analysis, in: IEEE Second International Conference on Social Computing (SocialCOM 2010), volume 10, 2010.

[40] G. Guizzardi, Ontological foundations for structural conceptual models, Ph.D. thesis, University of Twente, 2005.

[41] G. Guizzardi, G. Wagner, J. P. A. Almeida, Towards ontological foundations for the conceptual modeling of events, in: Conceptual Modeling – 32nd International Conference (ER 2013), Lecture Notes in Computer Science, Springer, 2013, pp. 327–341.

[42] G. Guizzardi, A. B. Benevides, C. M. Fonseca, D. Porello, J. P. A. Almeida, Ufo: Unified foundational ontology, Applied Ontology 17 (2022) 1–44.

[43] G. Guizzardi, J. P. A. Almeida, R. S. S. Guizzardi, Ontological Analysis and Design for Engineering Information Systems, Springer, Berlin, Germany, 2015.