The Joint Knowledge Graph Labs: Neuro-symbolic Reasoning in Action

Luigi Bellomarini[1], Livia Blasi[2,1], Andrea Gentili[1], Rosario Laurendi[1], Eleonora Laurenza[2,1] and Emanuel Sallinger[2,*]

[1]*Bank of Italy, Italy*
[2]*TU Wien, Austria*

Abstract

In this work, we present the Joint Knowledge Graph Labs (Joint KG Labs), a success story across different countries and between academic and industrial research labs. Initially founded at the University of Oxford, it now includes a number of international institutions. Yet, our main interest shall not be on organisation, but on research foci. We shall cover three of its main areas: (1) Knowledge Graphs and reasoning, (2) neuro-symbolic AI, and (3) applications, i.e., seeing these topics in action, in particular in the domain of finance and beyond.

Keywords

Knowledge Graphs, Neuro-symbolic AI, Reasoning

1. Introduction

The Joint Knowledge Graph Labs research focus is on neuro-symbolic data management and AI systems in the broad, with a particular specialisation in Knowledge Graph-based systems and reasoning. We therefore first discuss the labs' work on Knowledge Graphs (KGs) and Knowledge Graph Management Systems (KGMSs). After that, we consider symbolic and sub-symbolic combinations of reasoning, including neuro-symbolic reasoning. Finally, we concentrate on applications, i.e., seeing these topics in action. Our goal is twofold: (a) showing topics and thus collaboration opportunities with the labs and (b) giving an effective entry point through numerous references to the actual scientific works.

An overall picture. Following the layered view of KGs (Figure 1), the philosophy of the labs is to consider the area both from the side of the fundamental representations of knowledge in KGs (i.e., bottom up) as well as through applications (i.e., top down), with a major focus on the middle part, systems.

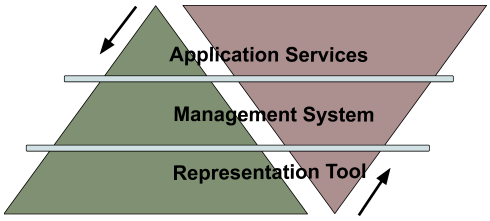

Figure 1: Layered View of KGs [1].

To make this come to life, let us consider a core family of contributions of the labs to this, Vadalog. Conceived from the need to understand which *representations* are required at the foundation of KGMSs [2], and how, through a comprehensive architecture (Figure 2), they enable enterprise AI *applications* [3], the labs developed the Vadalog *system* [4, 5]. Yet, importantly, not standalone but as part of a larger project [6] focusing very much on how this enables data science [7] and real-world applications [8].

Systems built on strong principles. The Vadalog system as well as recent versions such as Vadalog Parallel [9] are built on strong theoretical foundations that make scalability possible. This requires understanding effective joins in the context of KGs [10] and heuristics [11] as well as benchmarks such as iWarded [12] to understand the actual performance implications in KGs. At the core of such

RuleML+RR'25: Companion Proceedings of the 9th International Joint Conference on Rules and Reasoning, September 22–24, 2025, Istanbul, Turkiye
*Corresponding author.

principled foundations are the questions of space efficiency [13] and the key concepts underlying reasoning and query optimization as well as scalability [14, 15].

From a principles perspective, especially interesting are the connections to long-standing theoretical questions in the area on dependencies [16] and schema mappings, especially those designed for tree and graph data [17, 18]. This has far-reaching connections back to the foundations of reasoning in and about these [19, 20] as well as management tasks such as equivalence [21] and limits [22].

Beyond systems – design and interoperability. One critical element beyond the systems themselves are principled methods to design KGs, in particular approaches independent of the particular KG model [23, 24]. Bridging the gap between different KG models (whether RDF-based, property graph-based, relational or other) is the labs work on efficiently evaluating SPARQL together with Datalog-based languages [25]. Of similar importance for interoperability are efforts of bridging different Datalog-based languages such as Shy and Warded Datalog$^{\pm}$ [26, 27].

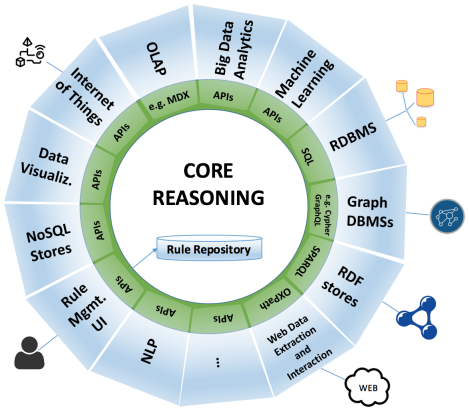

Figure 2: KGMS Reference Architecture [2].

Important extensions – temporal and more. Critical to the success of KGMSs are core properties like full recursion and existential quantification as well as extensions such as arithmetic and aggregation. One particularly important extension is temporal reasoning, as in the Temporal Vadalog system [28, 29]. An interesting challenge is supporting existential quantification [30] and aggregation [31]. Critical for supporting the community are benchmarks such as the iTemporal benchmark generation suite [32].

2. Neuro-symbolic AI – Reasoning and KGEs, GNNs and LLMs

The Joint KG Labs have as one of its main research foci the area typically called neuro-symbolic AI, that is, bringing together symbolic – or logic-based – reasoning, and subsymbolic – or Machine Learning (ML)-based – reasoning [33, 34]. In the area of KGs, key ML methods are Knowledge Graph Embeddings (KGEs), Graph Neural Networks (GNNs) and of course Large Language Models (LLMs) and related (graph) transformer-based architectures. A stylized but evocative representation of the labs' agenda is shown in Figure 3.

KGEs and GNNs. Let us start with an example of how such a neuro-symbolic combinations can work. By combining logic-based and (Knowledge Graph) embedding-based reasoning [35] one can effectively find solutions for real-world problems such as in the domain of finance, concretely companies [36, 37] . More fundamentally, it is critical to design KGE methods that can capture logical constraints [38, 39], that is, build ML models that respect domain logic [40, 41]. Similar considerations can be made for GNNs [42]. It is important to make such ML-methods resilient to noise [43, 44].

Probabilistic reasoning and rule learning. For areas where precise understanding and control of (ML)-based reasoning is necessary, we show two of the labs' foci. Where precise understanding of probabilities is necessary, probabilistic reasoning such as Markov Chain Monte Carlo-based methods are of particular interest [45, 46]. Where full understanding is necessary of the explicit underlying knowledge, rule learning methods can fulfill this role [47, 48].

LLMs. While many form of ML are interesting and relevant, of particular current importance are of course Large Language Models and neuro-symbolic uses of these. One particularly interesting one is the lab's approach to semantic aware query answering with LLMs Semantic-aware query answering with Large Language Models [49].

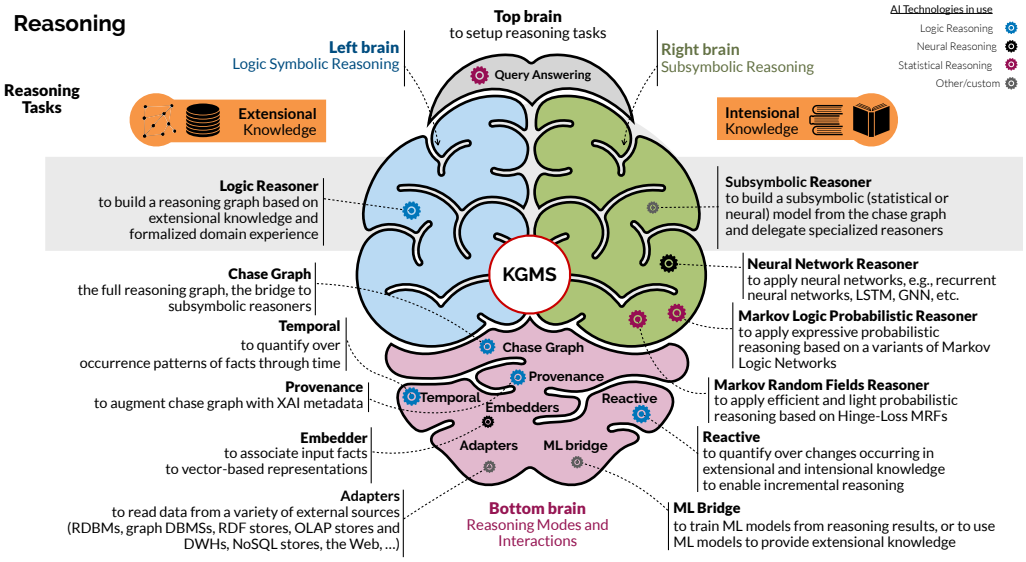

Figure 3: Neuro-symbolic AI Research Overview of the Joint KG Labs.

Meaningful combinations of LLMs and logic-based reasoning are possible at many levels, and can go both ways: KG-based logical reasoning supporting LLMs and the other way around. One such approach is integrating LLMs within the very reasoning mechanism of KGMSs [50] (LLMs supporting KGs), another one the use of logic-based reasoning for fine-tuning LLMs [51] (KGs supporting LLMs).

Explainability – LLMs and more. A final key area is explainability, where the lab has put a major focus on, including both traditional, interactive [52] and LLM-based [53, 54] explanations. This includes visual demonstrators. Combining many of the aforementioned topics, we can see approaches such as combining LLMs, temporal logic-based reasoning and financial use-cases [55, 56].

3. Applications – Finance and Beyond

The key application area of the Joint KG Labs is finance and economics. Core applications include a manifold of using the Vadalog system for financial scenarios [57], including interesting meetings points on topics such as company takeovers at the interface between computer science and economics [58].

Fundamental topics include distributed computation in financial KG scenarios [59], mining of financial knowledge for KGs [60], and KG-based anti-money laundering [61]. Of special interest are two areas, the emergency reaction of the labs to the COVID-19 crisis [62, 63] and a major focus on KGs and neuro-symbolic AI for blockchains and smart contracts [64, 65, 66, 67, 68, 69].

Further key applications. Further key topics of the labs are *privacy and confidentiality* [70, 71, 72], *legal* [73] and critical fields like *healthcare* [74, 75, 76, 77], *enterprise architecture*, in particular KG-based modeling [78, 79, 80, 81], and *education* [82, 83].

Acknowledgements. This work was supported by the Vienna Science and Technology Fund (WWTF) [10.47379/VRG18013, 10.47379/NXT22018, 10.47379/ICT2201] and the Austrian Science Fund [10.55776/COE12].

Declaration on Generative AI

The authors have not employed any Generative AI tools for writing this paper.

References

[1] L. Bellomarini, E. Sallinger, S. Vahdati, Knowledge graphs: The layered perspective, in: Knowledge Graphs and Big Data Processing, volume 12072 of *LNCS*, Springer, 2020, pp. 20–34.

[2] L. Bellomarini, G. Gottlob, A. Pieris, E. Sallinger, Swift logic for big data and knowledge graphs, in: IJCAI, ijcai.org, 2017, pp. 2–10.

[3] L. Bellomarini, D. Fakhoury, G. Gottlob, E. Sallinger, Knowledge graphs and enterprise AI: the promise of an enabling technology, in: ICDE, IEEE, 2019, pp. 26–37.

[4] L. Bellomarini, E. Sallinger, G. Gottlob, The vadalog system: Datalog-based reasoning for knowledge graphs, Proc. VLDB Endow. 11 (2018) 975–987.

[5] L. Bellomarini, D. Benedetto, G. Gottlob, E. Sallinger, Vadalog: A modern architecture for automated reasoning with large knowledge graphs, Inf. Syst. 105 (2022) 101528.

[6] N. Konstantinou, E. Abel, L. Bellomarini, A. Bogatu, C. Civili, E. Irfanie, M. Koehler, L. Mazilu, E. Sallinger, A. A. A. Fernandes, G. Gottlob, J. A. Keane, N. W. Paton, VADA: an architecture for end user informed data preparation, J. Big Data 6 (2019) 74.

[7] L. Bellomarini, R. R. Fayzrakhmanov, G. Gottlob, A. Kravchenko, E. Laurenza, Y. Nenov, S. Reissfelder, E. Sallinger, E. Sherkhonov, S. Vahdati, L. Wu, Data science with vadalog: Knowledge graphs with machine learning and reasoning in practice, FGCS 129 (2022) 407–422.

[8] T. Baldazzi, L. Bellomarini, M. Gerschberger, A. Jami, D. Magnanimi, M. Nissl, A. Pavlovic, E. Sallinger, Vadalog: Overview, extensions and business applications, in: RW, 2022.

[9] L. Bellomarini, D. Benedetto, M. Brandetti, E. Sallinger, A. Vlad, The vadalog parallel system: Distributed reasoning with datalog+/-, Proc. VLDB Endow. 17 (2024) 4614–4626.

[10] P. Atzeni, L. Bellomarini, D. Benedetto, E. Sallinger, Traversing knowledge graphs with good old (and new) joins, in: KR4L@ECAI, CEUR-WS.org, 2020, pp. 3–9.

[11] T. Baldazzi, D. Benedetto, M. Brandetti, A. Vlad, L. Bellomarini, E. Sallinger, Datalog-based reasoning with heuristics over knowledge graphs, in: Datalog, CEUR-WS.org, 2022, pp. 114–126.

[12] P. Atzeni, T. Baldazzi, L. Bellomarini, E. Sallinger, iwarded: A versatile generator to benchmark warded datalog+/- reasoning, in: RuleML+RR, volume 13752 of *LNCS*, Springer, 2022, pp. 113–129.

[13] G. Berger, G. Gottlob, A. Pieris, E. Sallinger, The space-efficient core of vadalog, ACM Trans. Database Syst. 47 (2022) 1:1–1:46.

[14] T. Baldazzi, L. Bellomarini, E. Sallinger, P. Atzeni, Eliminating harmful joins in warded datalog+/-, in: RuleML+RR, volume 12851 of *Lecture Notes in Computer Science*, Springer, 2021, pp. 267–275.

[15] T. Baldazzi, L. Bellomarini, E. Sallinger, P. Atzeni, Reasoning in warded datalog+/- with harmful joins, in: SEBD, volume 3194, CEUR-WS.org, 2022, pp. 292–299.

[16] L. Bellomarini, D. Benedetto, M. Brandetti, E. Sallinger, Exploiting the power of equality-generating dependencies in ontological reasoning, Proc. VLDB Endow. 15 (2022) 3976–3988.

[17] P. G. Kolaitis, R. Pichler, E. Sallinger, V. Savenkov, On the language of nested tuple generating dependencies, ACM Trans. Database Syst. 45 (2020) 8:1–8:59.

[18] P. G. Kolaitis, R. Pichler, E. Sallinger, V. Savenkov, Nested dependencies: structure and reasoning, in: PODS, ACM, 2014, pp. 176–187.

[19] E. Sallinger, Reasoning about schema mappings, in: Data Exchange, Information, and Streams, volume 5 of *Dagstuhl Follow-Ups*, 2013, pp. 97–127.

[20] I. Feinerer, R. Pichler, E. Sallinger, V. Savenkov, On the undecidability of the equivalence of second-order tuple generating dependencies, Inf. Syst. 48 (2015) 113–129.

[21] R. Pichler, E. Sallinger, V. Savenkov, Relaxed notions of schema mapping equivalence revisited, in: ICDT, ACM, 2011, pp. 90–101.

[22] P. G. Kolaitis, R. Pichler, E. Sallinger, V. Savenkov, Limits of schema mappings, Theory Comput. Syst. 62 (2018) 899–940.

[23] L. Bellomarini, A. Gentili, E. Laurenza, E. Sallinger, Model-independent design of knowledge graphs (short paper), in: AMW, volume 3409 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023.

[24] L. Bellomarini, A. Gentili, E. Laurenza, E. Sallinger, Model-independent design of knowledge graphs - lessons learnt from complex financial graphs, in: EDBT, 2022, pp. 2:524–2:526.

[25] R. Angles, G. Gottlob, A. Pavlovic, R. Pichler, E. Sallinger, Sparqlog: A system for efficient evaluation of SPARQL 1.1 queries via datalog, Proc. VLDB Endow. 16 (2023) 4240–4253.

[26] T. Baldazzi, L. Bellomarini, M. Favorito, E. Sallinger, Ontological reasoning over shy and warded datalog+/- for streaming-based architectures, in: PADL, Springer, 2024, pp. 169–185.

[27] T. Baldazzi, L. Bellomarini, M. Favorito, E. Sallinger, On the relationship between shy and warded datalog+/-, in: KR, 2022.

[28] L. Bellomarini, L. Blasi, M. Nissl, E. Sallinger, The temporal vadalog system: Temporal datalog-based reasoning, Theory Pract. Log. Program. 25 (2025) 168–196.
[29] L. Bellomarini, L. Blasi, M. Nissl, E. Sallinger, The temporal vadalog system, in: RuleML+RR, volume 13752 of *LNCS*, Springer, 2022, pp. 130–145.
[30] M. Lanzinger, M. Nissl, E. Sallinger, P. A. Walega, Temporal datalog with existential quantification, in: IJCAI, ijcai.org, 2023, pp. 3277–3285.
[31] L. Bellomarini, M. Nissl, E. Sallinger, Monotonic aggregation for temporal datalog, in: RuleML+RR (Supplement), volume 2956, CEUR-WS.org, 2021.
[32] L. Bellomarini, M. Nissl, E. Sallinger, itemporal: An extensible generator of temporal benchmarks, in: ICDE, IEEE, 2022, pp. 2021–2033.
[33] L. Bellomarini, E. Sallinger, S. Vahdati, Reasoning in knowledge graphs: An embeddings spotlight, in: Knowledge Graphs and Big Data Processing, volume 12072, Springer, 2020, pp. 87–101.
[34] E. Sallinger, Knowledge graphs: Meeting point of knowledge representation and representation learning, in: KR4L@ECAI, volume 3020, CEUR-WS.org, 2020, pp. 1–2.
[35] A. Vlad, S. Vahdati, M. Nayyeri, L. Bellomarini, E. Sallinger, Towards hybrid logic-based and embedding-based reasoning on financial knowledge graphs, in: EDBT/ICDT Workshops, 2022.
[36] P. Atzeni, L. Bellomarini, M. Iezzi, E. Sallinger, A. Vlad, Augmenting logic-based knowledge graphs: The case of company graphs, in: KR4L@ECAI, volume 3020, CEUR-WS.org, 2020, pp. 22–27.
[37] P. Atzeni, L. Bellomarini, M. Iezzi, E. Sallinger, A. Vlad, Weaving enterprise knowledge graphs: The case of company ownership graphs, in: EDBT, 2020, pp. 555–566.
[38] A. Pavlovic, E. Sallinger, Expressive: A spatio-functional embedding for knowledge graph completion, in: ICLR, OpenReview.net, 2023.
[39] A. Pavlovic, E. Sallinger, Speede: Euclidean geometric knowledge graph embedding strikes back, in: NAACL-HLT (Findings), Association for Computational Linguistics, 2024, pp. 69–92.
[40] A. Pavlovic, E. Sallinger, Building bridges: Knowledge graph embeddings respecting logical rules (short paper), in: AMW, volume 3409 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023.
[41] A. Pavlovic, E. Sallinger, Raising the efficiency of knowledge graph embeddings while respecting logical rules, in: AMW, volume 3954 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024.
[42] A. Pavlovic, E. Sallinger, S. Schockaert, Differentiable reasoning about knowledge graphs with region-based graph neural networks, CoRR abs/2406.09529 (2024).
[43] M. Nayyeri, S. Vahdati, E. Sallinger, M. M. Alam, H. S. Yazdi, J. Lehmann, Pattern-aware and noise-resilient embedding models, in: ECIR (1), volume 12656 of *LNCS*, Springer, 2021, pp. 483–496.
[44] M. Nayyeri, C. Xu, S. Vahdati, N. Vassilyeva, E. Sallinger, H. S. Yazdi, J. Lehmann, Fantastic knowledge graph embeddings and how to find the right space for them, in: ISWC (1), 2020.
[45] L. Bellomarini, E. Laurenza, E. Sallinger, E. Sherkhonov, Swift markov logic for probabilistic reasoning on knowledge graphs, Theory Pract. Log. Program. 23 (2023) 507–534.
[46] L. Bellomarini, E. Laurenza, E. Sallinger, E. Sherkhonov, Reasoning under uncertainty in knowledge graphs, in: RuleML+RR, volume 12173 of *LNCS*, Springer, 2020, pp. 131–139.
[47] L. Wu, E. Sallinger, E. Sherkhonov, S. Vahdati, G. Gottlob, Rule learning over knowledge graphs with genetic logic programming, in: ICDE, IEEE, 2022, pp. 3373–3385.
[48] L. Wu, E. Sallinger, E. Sherkhonov, S. Vahdati, G. Gottlob, An evolutionary algorithm for rule learning over knowledge graphs, in: KR4L@ECAI, CEUR-WS.org, 2020, pp. 52–59.
[49] P. Atzeni, T. Baldazzi, L. Bellomarini, E. Laurenza, E. Sallinger, Semantic-aware query answering with large language models, Data Knowl. Eng. 161 (2026) 102494.
[50] T. Baldazzi, D. Benedetto, L. Bellomarini, E. Sallinger, A. Vlad, Softening ontological reasoning with large language models, in: RuleML+RR (Companion), CEUR-WS.org, 2024.
[51] T. Baldazzi, L. Bellomarini, S. Ceri, A. Colombo, A. Gentili, E. Sallinger, Fine-tuning large enterprise language models via ontological reasoning, in: RuleML+RR, 2023, pp. 86–94.
[52] L. Bellomarini, M. Benedetti, A. Gentili, D. Magnanimi, E. Sallinger, Kg-roar: Interactive datalog-based reasoning on virtual knowledge graphs, Proc. VLDB Endow. 16 (2023) 4014–4017.
[53] T. Baldazzi, L. Bellomarini, S. Ceri, A. Colombo, A. Gentili, E. Sallinger, "please, vadalog, tell me why": Interactive explanation of datalog-based reasoning, in: EDBT, 2024, pp. 834–837.
[54] T. Baldazzi, L. Bellomarini, S. Ceri, A. Colombo, A. Gentili, E. Sallinger, P. Atzeni, Explaining enterprise knowledge graphs with large language models and ontological reasoning, in: Tannen's Festschrift, volume 119 of *OASIcs*, 2024, pp. 1:1–1:20.
[55] A. Colombo, T. Baldazzi, L. Bellomarini, A. Gentili, E. Sallinger, LLM-based DatalogMTL modelling of MiCAR-compliant crypto-assets markets, in: Datalog, volume 3801, 2024, pp. 17–22.
[56] A. Colombo, T. Baldazzi, L. Bellomarini, E. Sallinger, S. Ceri, Template-based explainable inference over

high-stakes financial knowledge graphs, in: EDBT, 2025, pp. 503–515.
[57] T. Baldazzi, L. Bellomarini, E. Sallinger, Reasoning over financial scenarios with the vadalog system, in: EDBT, OpenProceedings.org, 2023, pp. 782–791.
[58] L. Bellomarini, L. Bencivelli, C. Biancotti, L. Blasi, F. P. Conteduca, A. Gentili, R. Laurendi, D. Magnanimi, M. S. Zangrandi, F. Tonelli, S. Ceri, D. Benedetto, M. Nissl, E. Sallinger, Reasoning on company takeovers: From tactic to strategy, Data Knowl. Eng. 141 (2022) 102073.
[59] A. Gulino, S. Ceri, G. Gottlob, E. Sallinger, L. Bellomarini, Distributed company control in company shareholding graphs, in: ICDE, IEEE, 2021, pp. 2637–2648.
[60] L. Bellomarini, D. Magnanimi, M. Nissl, E. Sallinger, Neither in the programs nor in the data, in: MIDAS@PKDD/ECML, volume 12591 of *LNCS*, Springer, 2020, pp. 119–134.
[61] L. Bellomarini, E. Laurenza, E. Sallinger, Rule-based anti-money laundering in financial intelligence units: Experience and vision, in: RuleML+RR (Supplement), CEUR-WS.org, 2020, pp. 133–144.
[62] L. Bellomarini, M. Benedetti, A. Gentili, R. Laurendi, D. Magnanimi, A. Muci, E. Sallinger, COVID-19 and company knowledge graphs, CoRR abs/2004.10119 (2020).
[63] L. Bellomarini, M. Benedetti, S. Ceri, A. Gentili, R. Laurendi, D. Magnanimi, M. Nissl, E. Sallinger, Reasoning on company takeovers during the COVID-19 crisis with knowledge graphs, in: RuleML+RR (Supplement), CEUR-WS.org, 2020, pp. 145–156.
[64] L. Bellomarini, M. Favorito, E. Laurenza, M. Nissl, E. Sallinger, Towards fateful smart contracts, in: AMW, volume 3954 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024.
[65] M. Nissl, E. Sallinger, Towards bridging traditional and smart contracts with datalog-based languages, in: Datalog, volume 3203, CEUR-WS.org, 2022, pp. 68–82.
[66] M. Nissl, E. Sallinger, S. Schulte, M. Borkowski, Towards cross-blockchain smart contracts, in: DAPPS, IEEE, 2021, pp. 85–94.
[67] M. Nissl, E. Sallinger, Modelling smart contracts with datalogmtl, in: EDBT/ICDT Workshops, volume 3135, CEUR-WS.org, 2022.
[68] L. Bellomarini, G. Galano, M. Nissl, E. Sallinger, Rule-based blockchain knowledge graphs: Declarative AI for solving industrial blockchain challenges, in: RuleML+RR (Supplement), 2021.
[69] L. Bellomarini, M. Nissl, E. Sallinger, Blockchains as knowledge graphs - blockchains for knowledge graphs (vision paper), in: KR4L@ECAI, volume 3020, CEUR-WS.org, 2020, pp. 43–51.
[70] L. Bellomarini, L. Blasi, R. Laurendi, E. Sallinger, Financial data exchange with statistical confidentiality: A reasoning-based approach, in: EDBT, OpenProceedings.org, 2021, pp. 558–569.
[71] L. Bellomarini, L. Blasi, R. Laurendi, E. Sallinger, A reasoning approach to financial data exchange with statistical confidentiality, in: PKDD/ECML Workshops (2), 2021, pp. 226–231.
[72] L. Bellomarini, C. Catalano, A. Coletta, M. Iezzi, P. Samarati, Chase anonymisation: Privacy-preserving knowledge graphs with logical reasoning, in: ICDE, 2026. Accepted.
[73] L. Bellomarini, M. Brandetti, A. Gentili, R. Laurendi, D. Magnanimi, Knowledge graphs & reasoning in concert: Automating regulation in corporate economics, in: LPLR, 2025.
[74] L. Chammas, O. P. Dwyer, E. Sallinger, J. Davies, E. J. Morris, Care records and healthcare processes: adding context to clinical codes, in: HICSS, ScholarSpace, 2024, pp. 3697–3706.
[75] O. P. Dwyer, L. Chammas, E. Sallinger, J. Davies, Investigating an ontology-informed approach to event log generation in healthcare, in: ICPM Workshops, Springer, 2023, pp. 235–246.
[76] O. P. Dwyer, T. Baldazzi, J. Davies, E. Sallinger, A. Vlad, Reasoning over health records with vadalog: a rule-based approach to patient pathways, in: RuleML+RR (Companion), 2023.
[77] M. Fazeli-Varzaneh, A. Ghorbi, M. Ausloos, E. Sallinger, S. Vahdati, Sleeping beauties of coronavirus research, IEEE Access 9 (2021) 21192–21205.
[78] P. Glaser, S. J. Ali, E. Sallinger, D. Bork, Model-based construction of enterprise architecture knowledge graphs (extended abstract), in: EMISA, volume 3397, CEUR-WS.org, 2023.
[79] P. Glaser, E. Sallinger, D. Bork, EA modelset - A FAIR dataset for machine learning in enterprise modeling, in: PoEM, volume 497 of *LNBIP*, Springer, 2023, pp. 19–36.
[80] P. Glaser, S. J. Ali, E. Sallinger, D. Bork, Model-based construction of enterprise architecture knowledge graphs, in: EDOC, volume 13585 of *LNCS*, Springer, 2022, pp. 57–73.
[81] P. Glaser, S. J. Ali, E. Sallinger, D. Bork, Exploring enterprise architecture knowledge graphs in archi: The EAKG toolkit, in: EDOC Workshops, volume 466 of *LNBIP*, Springer, 2022, pp. 332–338.
[82] D. Graux, V. Janev, H. Jabeen, E. Sallinger, Deploying a strategy to unlock big data research and teaching activities in the west balkan region, in: ITiCSE (1), ACM, 2021, pp. 491–497.
[83] D. Graux, V. Janev, H. Jabeen, E. Sallinger, A big data learning platform for the west balkans and beyond, in: ITiCSE (2), ACM, 2021, pp. 617–618.