

On the Benefits of OWL-based Knowledge Graphs for Neural-Symbolic Systems

A position paper

David Herron^{1,*}, Ernesto Jiménez-Ruiz^{1,2} and Tillman Weyde¹

¹City, University of London, London, United Kingdom

²SIRIUS, University of Oslo, Norway

Abstract

Knowledge graphs, as understood within the Semantic Web and Knowledge Representation communities, are more than just graph data. OWL-based knowledge graphs offer the benefits of being based on an ecosystem of open W3C standards that are implemented in a range of reusable existing resources (e.g. curated ontologies, software tools, web-wide linked data) and that also permit researchers to tailor resources for their unique needs (e.g. custom ontologies). Additionally, OWL-based knowledge graphs offer the benefits of formal, logical symbolic reasoning (e.g. reliable inference of new knowledge based on Description Logics, semantic consistency checking, extensions via user-defined Datalog rules). These capabilities allow OWL-based knowledge graphs to be leveraged in the form of active reasoning agents to guide deep learning during training and to participate in refining neural inference. We enumerate a host of such benefits to using OWL-based knowledge graphs in neural-symbolic systems. We illustrate several of these by drawing upon examples from our research in visual relationship detection within images, and we point to promising research directions and challenging opportunities.

Keywords

neural-symbolic, AI, deep learning, Semantic Web, OWL, ontologies, knowledge graphs, reasoning

1. Introduction

OWL-based KGs are exemplars of the explicit symbolic knowledge representation and symbol manipulation and reasoning machinery that prominent voices like those of Chollet [1], Marcus [2, 3, 4] and Kautz [5] have argued over the last few years, should be combined with deep learning in hybrid, neural-symbolic (NeSy) systems. Deep learning continues to advance AI. OpenAI's GPT-4 [6] impresses even more than ChatGPT (GPT-3.5), improving factual correctness and arithmetic consistency to some extent.¹ OpenAI says that GPT-4 is the latest milestone in their effort in scaling up deep learning with bigger models, data and computing power [6]. But last year, Marcus reprised his earlier critiques of deep learning by enumerating the limitations

NeSy 2023: 17th International Workshop on Neural-Symbolic Learning and Reasoning, July 3–5, 2023, Siena, Italy

*Corresponding author.

✉ david.herron@city.ac.uk (D. Herron); ernesto.jimenez-ruiz@city.ac.uk (E. Jiménez-Ruiz); t.e.veyde@city.ac.uk (T. Weyde)

🌐 <https://djherron.github.io/> (D. Herron)

🆔 0009-0008-2736-6789 (D. Herron); 0000-0002-9083-4599 (E. Jiménez-Ruiz); 0000-0001-8028-9905 (T. Weyde)

© 2023 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)

¹E.g., whereas GPT-3.5 explains (with confident eloquence) how it is that 1kg of lead weighs the same as 2kg of feathers, GPT-4 gets it right, saying 2kg is more than 1kg.

of GPT-3 [7], and just recently argued that, despite GPT-4’s advances, those fundamental limitations remain [8]. Scaling up puts a better mask over deep learning’s dependence upon statistical pattern matching, but the limitations of that dependency persist.

Our position is that OWL-based KGs have valuable contributions to make to NeSy systems, particularly due to their deductive reasoning capabilities, but are underutilised. A recent survey of 476 papers exploring approaches to combining machine learning (ML) with SW technologies [9], reports that only 20 (4%) mention using reasoning capabilities to infer new knowledge. We argue that OWL-based KGs merit a larger seat at the NeSy AI research table by highlighting a host of benefits of OWL-based KGs, drawing upon illustrative examples from our own research, and by pointing to research directions both promising and challenging. By doing so, we hope to inspire more NeSy research using OWL-based KGs.

2. Benefits of OWL-based KGs

Here we give our perspective on attractions and benefits of OWL-based KGs, and some examples illustrating how and why they can be usefully applied in NeSy systems.

Open standards and reusable resources The Web Ontology Language (OWL) [10, 11] and OWL-based knowledge graphs (KGs) [12, 13] are key components of the W3C open standards ecosystem of the Semantic Web (SW) [14, 15, 16, 17]. Open standards facilitate interoperability and promote development of reusable, often free, software resources that make it easy to work with OWL-based KGs. Amongst the many such resources are: (i) public SW KGs like DBpedia [18], Wikidata [19] and Yago [20]; (ii) public repositories of curated OWL ontologies like BioPortal [21] and OBO Foundry [22] in the biomedical domain; (iii) RDF stores like GraphDB (not open, but has free version) [23] and RDFox (not open, but has free academic license) [24]; and (iv) efficient OWL reasoners like Hermit [25], Pellet [26], RDFox and ELK [27].

Custom ontologies and custom KGs Reusing state-of-the-art ontologies and/or public KGs is a good practice option. But researchers can also design custom, domain-specific OWL ontologies tailored to their unique needs and use them to construct custom KGs. Custom ontologies can subsequently be aligned with publicly available ontologies to enhance interoperability [28].

This is the approach we took for our own work on visual relationship detection in images. We designed an ontology to describe the domain of common object classes and relationships (predicates) referred to in the human-supplied visual relationship annotations for the everyday images of the VRD dataset [29]. As depicted in Figure 1, our custom OWL ontology, called VRD-World [30], drives a custom KG in the hybrid systems with which we explore combining neural learning with symbolic reasoning. We used the free ontology editor Protégé [31] to engineer our VRD-World ontology and took guidance from the vast literature on ontology engineering [e.g. 32, 33, 34, 35]. Additionally, many ML-based tools have been developed to support aspects of ontology development (see [36]), such as for *concept learning*.

Knowledge completion and infusion KGs have inspired a vast amount of research into encoding their symbolic background knowledge into vectors — *KG embeddings* — that preserve

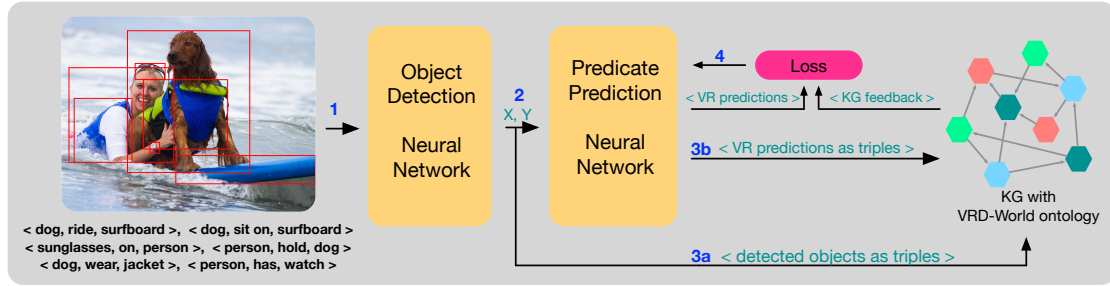


Figure 1: A representative illustration of our hybrid, NeSy systems for detecting visual relationships within images by leveraging OWL-based KG reasoning to guide neural learning.

semantic similarity and reflect similarity by proximity within the (low-dimensional) embedding space [36, 37, 38, 39, 40, 41, 42]. The primary application of KG embeddings has been the task of *KG completion*: *link prediction* (relating individuals in the KG) or *type prediction* (classifying individuals in the KG). These problems are cast as neural classification problems, and the embedded KG knowledge is used to guide neural learning using tactics that exploit the proximity principle. *Link inference* and *type inference* are the bread and butter of OWL reasoners. They materialise (complete) a KG commensurate with the richness of the governing ontology. And the logical soundness of their inferences is guaranteed, whereas NeSy KG completion can be approximate and error prone. One can also combine these approaches: use OWL materialisation reasoning to complete a KG as far as possible, and NeSy KG completion (emulated reasoning) for special cases OWL cannot address.

A secondary (but growing) application of KG embeddings is *knowledge infused learning* [43, 44], where embedded knowledge is infused into NNs internally to guide neural learning. OWL-based KGs can help here too: a materialised KG (where everything implicit is made explicit) will contain more knowledge to embed and deliver richer embeddings for infusion.

Deductive reasoning and agency ChatGPT and similar models are notorious for their lack of reliability in reasoning. The reliability (logical soundness) of OWL’s deductive reasoning is guaranteed because it is grounded in formal Description Logics (DLs) that are decidable fragments of first-order logic [45, 46, 47, 48]. The highly expressive DL *SROIQ* is used in the latest OWL 2 [49]. Given an OWL ontology, OWL reasoners infer new knowledge for KGs and enforce the logical consistency of KGs. Both of these capabilities are commonly used to debug ontology models during ontology development [13, 50].

Crucially, these same capabilities also enable OWL-based KGs to be leveraged as *active reasoning agents* in hybrid, NeSy systems. This is the hypothesis that informs much of our research, as depicted in Figure 1. We have identified scenarios where KG reasoning agents can guide neural learning and participate in neural inference. During training, our KG can reason over *ground-truth* and/or *predicted* visual relationships, and its reasoned judgements used to guide the learning of a Predicate Prediction NN by modifying the loss. During inference, it can reason over visual relationship predictions produced by a Predicate Prediction NN to *filter out* implausible candidates and help ensure the predictions submitted by the hybrid system for performance evaluation are optimal.

For example, *link inference* in KGs is driven by reasoning over KG data with respect to an ontology’s object property hierarchy. The 70 predicates of the VRD dataset (mostly common spatial relations and verbs) have corresponding object properties in our VRD-World ontology that allow a rich web of characteristics (e.g. symmetry, transitivity) and relationships (e.g. inverses, subPropertyOf, equivalentPropertyOf) to be defined. The human-annotated visual relationships of the VRD images are sparse, arbitrary, and semantically noisy. We leverage KG *link inference* for a form of *data augmentation: annotation augmentation*. One version of VRD-World’s property hierarchy increases the annotated visual relationships per image by a factor of 2.5 (on average), resulting in denser, more consistent, semantically de-noised supervision for (hypothesised) faster and better neural learning.

Yet another benefit of OWL-based KGs is that, in certain cases, OWL inference semantics can be extended via Datalog rules that capture nuanced inference cases beyond OWL’s reach. Part of our research will explore this opportunity, e.g. with RDFox, which implements a fast engine that seamlessly blends reasoning over the OWL 2 RL profile and Datalog rules. Many of our planned Datalog rules contain goals that rely on KG *type inference*. For example, a rule for determining when it is plausible (or implausible) to predict that two detected objects, X and Y, be related with predicate wear can be represented as

$$\text{wear}(X, Y) := \text{WearCapableThing}(X), \text{WearableThing}(Y), \text{ir}(Y, X) > 0.8$$

where function $\text{ir}()$ measures an *inclusion ratio* (the extent to which the bounding box for Y is included within the bounding box for X). Ontology classes `WearCapableThing` and `WearableThing` are generalisations for multiple, low-level classes tied to the dataset (e.g. `Person`, `Dog`, `Teddy Bear`, or `Jacket`, `Sunglasses`, `Hat`, respectively). KG *type inference* thus makes it feasible to define a single rule that captures a vast multiplicity of cases.

3. Promising and Challenging Research Directions

Here we briefly point to several areas of NeSy research where the capabilities of OWL-based KGs might usefully be explored.

OWL-based KG plausibility reasoning In the VRD dataset, the visual relationship annotations are sparse and somewhat arbitrary, so the supervision they provide is incomplete and many conditions occur for few-shot and zero-shot learning. This has revealed several scenarios where an OWL-based KG’s reasoned judgements as to the plausibility (or implausibility) of visual relationships can be leveraged to guide neural learning. OWL-based KG plausibility judgements could also be applied to other non-exhaustively annotated and k-shot supervised learning problems (within vision or other domains), to semi-supervised learning problems (where some examples are labelled, others not), and potentially to unsupervised learning problems.

Transferring KG subsumption reasoning capability to neural networks As part of our research, we are developing a technique for representing an OWL class hierarchy with an extension to the architecture of a classification NN, similar to [51]. We equip the NN with the ability to perfectly emulate the subsumption reasoning of an OWL-based KG, using OWL

reasoning as part of the solution strategy. A further possibility is to place this technique within a neural network, so that the subsequent layers can benefit from the class generalisation. Another direction is to explore transferring other aspects of OWL-based KG background knowledge and reasoning capability to structural strong priors within NN architectures.

Using KG reasoning as logical constraints Much NeSy research explores using background knowledge expressed in first-order or propositional logic axioms as constraints to guide neural learning, often by manipulating loss to encourage constraint satisfaction. Examples are Logic Tensor Networks (LTN) [52, 53], the ROAD-R dataset [54], and [55]. The ability of OWL reasoners to check and enforce logical consistency of a KG means that aspects of OWL ontologies can be used as direct counterparts of logical constraints.

One such aspect relates to *domain/range restrictions* defined for OWL object properties. If permitting insertion of a triple would lead to a restriction being violated and the KG’s state becoming inconsistent, insertion is rejected. This response can be used to penalise NN loss. We plan to exploit this to essentially replicate (using a KG) the research done with the VRD dataset in [56] using LTN and negative first-order domain/range LTN Real Logic axioms. That research reveals a limitation of the logical constraint approach to which OWL-based KGs are immune: the combinatorial explosion in the number of logical constraints that may be needed as the number of dataset classes grows even only moderately large.

Another such aspect relates to *disjoint classes and properties*. The propositional constraints in [54], such as $(\neg\text{RedTL} \vee \neg\text{GreenTL})$, meaning “a traffic light cannot be both red and green”, can be expressed in OWL by declaring classes to be disjoint (or not). And OWL can go further.

Integrating KGs with existing NeSy frameworks OWL-based KG knowledge and deductive reasoning can conceivably be integrated with existing logic-based NeSy frameworks such as LTN. So long as (i) there is sufficient contextual information contained in the tensors of NN input data (or otherwise) to permit meaningful SPARQL queries to be constructed, and (ii) the KG’s responses to those SPARQL queries can be mapped to fuzzy truth values in $[0, 1]$, then functions encapsulating interactions with OWL-based KGs can participate in the constraint axioms used by LTN to train NNs.

4. Conclusion

Given the rich ecosystem that exists of free, standards-based resources, ontologies, and support for ontology design, researching the use of OWL-based KGs in NeSy AI is easier than may be suspected. Given the capabilities of OWL ontologies and OWL-based KGs for structured symbolic knowledge representation, query response, sound and scalable deductive reasoning, and agency, the range of possibilities for leveraging KGs in NeSy AI is broader than may be suspected. A recent overview of NeSy systems [57] reports success using an ontology to boost expert user satisfaction with large language model performance, and, like us, advocates KGs for NeSy AI. The aim of this paper is to raise awareness of these matters and inspire more research using OWL-based KGs in NeSy systems in order that their potential for contributing to NeSy AI be better explored.

References

- [1] F. Chollet, Deep Learning: Current Limits and What Lies Beyond Them, Presentation at RAAIS, 2018. URL: <https://raais.co/speakers-2018>.
- [2] G. Marcus, Deep Learning: A Critical Appraisal, CoRR (2018). URL: <http://arxiv.org/abs/1801.00631>.
- [3] G. Marcus, Innateness, AlphaZero, and Artificial Intelligence, CoRR (2018). URL: <http://arxiv.org/abs/1801.05667>.
- [4] G. Marcus, The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence, CoRR (2020). URL: <https://arxiv.org/abs/2002.06177>.
- [5] H. Kautz, The Third AI Summer, AAAI Robert S. Englemore Memorial Lecture, Thirty-fourth AAAI Conference on Artificial Intelligence, New York, NY, 2020. URL: <https://www.cs.rochester.edu/u/kautz/talks/>, presentation slides and video.
- [6] OpenAI, GPT-4, 2023. URL: <https://openai.com/research/gpt-4>, research notes.
- [7] G. Marcus, Deep Learning is Hitting a Wall, Nautilus science magazine, 2022. URL: <https://nautil.us/deep-learning-is-hitting-a-wall-238440/>, online essay.
- [8] G. Marcus, GPT-4's successes, and GPT-4's failures, Substack newsletter: The Road to AI We Can Trust, 2023. URL: <https://garymarcus.substack.com/p/gpt-4s-successes-and-gpt-4s-failures>, blog post.
- [9] A. Breit, L. Waltersdorfer, F. J. Ekaputra, M. Sabou, A. Ekelhart, A. Iana, H. Paulheim, J. Portisch, A. Revenko, A. t. Teije, F. van Harmelen, Combining Machine Learning and Semantic Web: A Systematic Mapping Study, ACM Computing Surveys (2023). URL: <https://doi.org/10.1145/3586163>, just accepted.
- [10] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, S. Rudolph, OWL 2 Web Ontology Language Primer, W3C Recommendation, World Wide Web Consortium, 2012. URL: <http://www.w3.org/TR/owl2-primer/>.
- [11] B. Motik, P. F. Patel-Schneider, B. Parsia, C. Bock, A. Fokoue, P. Haase, R. Hoekstra, I. Horrocks, A. Ruttenberg, U. Sattler, M. Smith, OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax, W3C Recommendation, World Wide Web Consortium, 2012. URL: <http://www.w3.org/2007/OWL/draft/owl2-syntax/>.
- [12] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutiérrez, S. Kirrane, J. E. Labra Gayo, R. Navigli, S. Neumaier, A.-C. Ngonga Ngomo, A. Polleres, S. M. Rashid, A. Rula, L. Schmelzeisen, J. F. Sequeda, S. Staab, A. Zimmermann, Knowledge Graphs, Synthesis Lectures on Data, Semantics, and Knowledge, Springer, 2021. URL: <https://kgbook.org/>.
- [13] M. Kejriwal, C. A. Knoblock, P. Szekely, Knowledge Graphs: Fundamentals, Techniques, and Applications, The MIT Press, 2021.
- [14] T. Berners-Lee, J. Hendler, O. Lassila, The Semantic Web, Scientific American 284 (2001) 34–43. URL: <https://www.jstor.org/stable/pdf/26059207.pdf>.
- [15] The Semantic Web Wiki, World Wide Web Consortium, 2001. URL: https://www.w3.org/2001/sw/wiki/Main_Page.
- [16] P. Hitzler, M. Krötzsch, S. Rudolph, Foundations of Semantic Web Technologies, 3 ed., CRC Press, 2010.
- [17] The Semantic Web Stack, Wikipedia, 2022. URL: https://en.wikipedia.org/wiki/Semantic_Web_Stack.

- [18] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, C. Bizer, DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia., *Semantic Web 6* (2015) 167–195. URL: <http://dblp.uni-trier.de/db/journals/semweb/semweb6.html#LehmannIJKMHMK15>.
- [19] D. Vrandečić, M. Krötzsch, Wikidata: A Free Collaborative Knowledgebase, *Communications of the ACM* 57 (2014) 78–85.
- [20] T. P. Tanon, G. Weikum, F. M. Suchanek, YAGO 4: A Reason-able Knowledge Base, in: *The Semantic Web - 17th International Conference, ESWC 2020, Proceedings*, volume 12123 of *Lecture Notes in Computer Science*, Springer, 2020, pp. 583–596. URL: https://doi.org/10.1007/978-3-030-49461-2_34.
- [21] P. L. Whetzel, N. F. Noy, N. H. Shah, P. R. Alexander, C. Nyulas, T. Tudorache, M. A. Musen, BioPortal: Enhanced Functionality via new Web Services from the National Center for Biomedical Ontology to Access and use Ontologies in Software Applications, *Nucleic Acids Research* 39 (2011) W541–W545. <https://bioportal.bioontology.org/>.
- [22] R. Jackson, N. Matentzoglou, J. A. Overton, R. Vita, J. P. Balhoff, P. L. Buttigieg, S. Carbon, M. Courtot, A. D. Diehl, D. M. Dooley, W. D. Duncan, N. L. Harris, M. A. Haendel, S. E. Lewis, D. A. Natale, D. Osumi-Sutherland, A. Ruttenberg, L. M. Schriml, B. Smith, C. J. Stoeckert Jr., N. A. Vasilevsky, R. L. Walls, J. Zheng, C. J. Mungall, B. Peters, OBO Foundry in 2021: Operationalizing Open Data Principles to Evaluate Ontologies, *Database 2021* (2021). URL: <https://doi.org/10.1093/database/baab069>, <http://obofoundry.org/>.
- [23] Ontotext GraphDB, Wikipedia, 2023. URL: https://en.wikipedia.org/wiki/Ontotext_GraphDB, A popular RDF store; <https://www.ontotext.com/products/graphdb/>.
- [24] Y. Nenov, R. Piro, B. Motik, I. Horrocks, Z. Wu, J. Banerjee, RDFox: A Highly-Scalable RDF Store, in: *The Semantic Web - 14th International Conference, ISWC 2015, Proceedings*, volume 9367 of *Lecture Notes in Computer Science*, Springer Verlag, 2015, pp. 3–20. URL: <https://ora.ox.ac.uk/objects/uuid:2a08b023-77be-431a-a08c-89b47381586a>, <https://www.oxfordsemantic.tech/product>.
- [25] B. Glimm, I. Horrocks, B. Motik, G. Stoilos, Z. Wang, HermiT: An OWL 2 Reasoner, *J. Autom. Reason.* 53 (2014) 245–269. URL: <https://doi.org/10.1007/s10817-014-9305-1>.
- [26] E. Sirin, B. Parsia, B. C. Grau, A. Kalyanpur, Y. Katz, Pellet: A Practical OWL-DL Reasoner, *J. Web Semant.* 5 (2007) 51–53. URL: <https://doi.org/10.1016/j.websem.2007.03.004>.
- [27] Y. Kazakov, M. Krötzsch, F. Simančík, Unchain My *EL* Reasoner, in: *Proceedings of the 24th International Workshop on Description Logics (DL'11)*, volume 745 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2011.
- [28] J. Euzenat, P. Shvaiko, *Ontology Matching*, Second Edition, Springer, 2013.
- [29] C. Lu, R. Krishna, M. Bernstein, L. Fei-Fei, Visual Relationship Detection with Language Priors, in: *European Conference on Computer Vision*, Springer, 2016, pp. 852–869. URL: <https://cs.stanford.edu/people/ranjaykrishna/vrd/>.
- [30] D. Herron, E. Jiménez-Ruiz, G. Tarroni, T. Weyde, NeSy4VRD: A Multifaceted Resource for Neurosymbolic AI Research using Knowledge Graphs in Visual Relationship Detection, *CoRR* (2023). URL: <http://arxiv.org/abs/2305.13258>, (Submitted to conference).
- [31] M. A. Musen, The Protégé Project: A Look Back and a Look Forward, *AI Matters* 1 (2015) 4–12. URL: <https://protege.stanford.edu/>. doi:10.1145/2757001.2757003.
- [32] N. Noy, D. McGuinness, *Ontology Development 101: A Guide to Creating Your First*

- Ontology, Technical Report, Knowledge Systems Laboratory, Stanford University, 2001. URL: <http://www.ksl.stanford.edu/people/dlm/papers/ontology-tutorial-noy-mcguinness.pdf>.
- [33] E. F. Kendall, D. L. McGuinness, *Ontology Engineering, Synthesis Lectures on The Semantic Web: Theory and Technology*, Morgan & Claypool Publishers, 2019.
- [34] D. Allemang, J. Hendler, F. Gandon, *Semantic Web for the Working Ontologist*, 3 ed., ACM Books, 2020.
- [35] C. M. Keet, *An Introduction to Ontology Engineering*, 'v1.5' ed., 2020. URL: <https://people.cs.uct.ac.za/~mkeet/OEbook/>.
- [36] C. d'Amato, *Machine Learning for the Semantic Web: Lessons learnt and next research directions*, *Semantic Web* 11 (2020) 195–203. URL: <https://doi.org/10.3233/SW-200388>.
- [37] M. Nickel, K. Murphy, V. Tresp, E. Gabrilovich, *A Review of Relational Machine Learning for Knowledge Graphs*, *CoRR* (2015). URL: <http://arxiv.org/abs/1503.00759>.
- [38] Y. Dai, S. Wang, N. N. Xiong, W. Guo, *A Survey on Knowledge Graph Embedding: Approaches, Applications and Benchmarks*, *Electronics* 9 (2020). URL: <https://www.mdpi.com/2079-9292/9/5/750>. doi:10.3390/electronics9050750.
- [39] M. Wang, L. Qiu, X. Wang, *A Survey on Knowledge Graph Embeddings for Link Prediction*, *Symmetry* 13 (2021). URL: <https://doi.org/10.3390/sym13030485>.
- [40] Z. Chen, Y. Wang, B. Zhao, J. Cheng, X. Zhao, Z. Duan, *Knowledge Graph Completion: A Review*, *IEEE Access* 8 (2020) 192435–192456. URL: <https://doi.org/10.1109/ACCESS.2020.3030076>.
- [41] A. Rossi, D. Barbosa, D. Firmani, A. Matinata, P. Merialdo, *Knowledge Graph Embedding for Link Prediction: A Comparative Analysis*, *ACM Trans. Knowl. Discov. Data* 15 (2021) 14:1–14:49. URL: <https://doi.org/10.1145/3424672>.
- [42] J. Chen, P. Hu, E. Jiménez-Ruiz, O. M. Holter, D. Antonyrajah, I. Horrocks, *OWL2Vec*: Embedding of OWL Ontologies*, *Mach. Learn.* 110 (2021) 1813–1845. URL: <https://doi.org/10.1007/s10994-021-05997-6>.
- [43] U. Kursuncu, M. Gaur, A. P. Sheth, *Knowledge Infused Learning (K-IL): Towards Deep Incorporation of Knowledge in Deep Learning*, in: A. Martin, K. Hinkelmann, H. Fill, A. Gerber, D. Lenat, R. Stolle, F. van Harmelen (Eds.), *Proceedings of the AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice, AAAI-MAKE 2020*, volume 2600 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020. URL: <https://ceur-ws.org/Vol-2600/paper14.pdf>.
- [44] R. Wickramarachchi, C. A. Henson, A. P. Sheth, *Knowledge-infused Learning for Entity Prediction in Driving Scenes*, *Frontiers in Big Data* 4 (2021) 759110. URL: <https://doi.org/10.3389/fdata.2021.759110>.
- [45] R. Brachman, H. Levesque, *Knowledge Representation and Reasoning*, Morgan Kaufman, 2004.
- [46] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. F. Patel-Schneider (Eds.), *The Description Logic Handbook*, 2 ed., Cambridge University Press, 2007.
- [47] F. Baader, I. Horrocks, C. Lutz, U. Sattler, *An Introduction to Description Logic*, Cambridge University Press, 2017.
- [48] D. Nardi, R. J. Brachman, *An Introduction to Description Logics*, in: F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, P. F. Patel-Schneider (Eds.), *The Description Logic Handbook*:

- Theory, Implementation, and Applications, Cambridge University Press, 2003, pp. 1–40.
- [49] I. Horrocks, O. Kutz, U. Sattler, The Even More Irresistible SROIQ, in: P. Doherty, J. Mylopoulos, C. A. Welty (Eds.), Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, AAAI Press, 2006, pp. 57–67. URL: <http://www.aaai.org/Library/KR/2006/kr06-009.php>.
- [50] E. Jiménez-Ruiz, Logic-based Support for Ontology Development in Open Environments, Ph.D. thesis, Jaume I University, Spain, 2010. URL: <http://hdl.handle.net/10803/10493>.
- [51] R. Kopparti, T. Weyde, Weight Priors for Learning Identity Relations, in: Workshop Knowledge Representation & Reasoning Meets Machine Learning at NeurIPS, 2019, pp. 8–14.
- [52] L. Serafini, A. S. d’Avila Garcez, Logic Tensor Networks: Deep Learning and Logical Reasoning from Data and Knowledge, in: Proceedings of the 11th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy’16), volume 1768 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2016. URL: http://ceur-ws.org/Vol-1768/NESY16_paper3.pdf.
- [53] S. Badreddine, A. d’Avila Garcez, L. Serafini, M. Spranger, Logic Tensor Networks, *Artificial Intelligence* 303 (2022). URL: <https://doi.org/10.1016/j.artint.2021.103649>.
- [54] E. Giunchiglia, M. C. Stoian, S. Khan, F. Cuzzolin, T. Lukasiewicz, ROAD-R: The Autonomous Driving Dataset with Logical Requirements, *Machine Learning* (2023). URL: <https://doi.org/10.1007/s10994-023-06322-z>.
- [55] C. Cornelio, J. Stuehmer, S. X. Hu, T. Hospedales, Learning Where and When to Reason in Neuro-Symbolic Inference, in: International Conference on Learning Representations, 2023. URL: <https://openreview.net/forum?id=en9V5F8PR->.
- [56] I. Donadello, L. Serafini, Compensating Supervision Incompleteness with Prior Knowledge in Semantic Image Interpretation, in: IJCNN Hungary, July 14–19, IEEE, 2019, pp. 1–8. URL: <https://doi.org/10.1109/IJCNN.2019.8852413>.
- [57] A. P. Sheth, K. Roy, M. Gaur, Neurosymbolic AI - Why, What, and How, *CoRR* (2023). URL: <https://doi.org/10.48550/arXiv.2305.00813>.