

Accelerating Knowledge Graph and Ontology Engineering with Large Language Models

Cogan Shimizu^a, Pascal Hitzler^b

^aDepartment of Computer Science and Engineering, Wright State University, USA

^bDepartment of Computer Science, Kansas State University, USA

Abstract

Large Language Models bear the promise of significant acceleration of key Knowledge Graph and Ontology Engineering tasks, including ontology modeling, extension, modification, population, alignment, as well as entity disambiguation. We lay out LLM-based Knowledge Graph and Ontology Engineering as a new and coming area of research, and argue that modular approaches to ontologies will be of central importance.

Keywords: Knowledge Graph Engineering, Ontology Engineering, Large Language Models, Modular Ontologies, Ontology Modeling, Ontology Population, Ontology Alignment, Entity Disambiguation

1. Introduction

Knowledge Graph and Ontology Engineering (KGOE, in short) refers to a (vaguely defined) set of tasks that are of central relevance to the life cycle of knowledge graphs, and of ontologies,¹ as data artifacts used in data management and applications.² These include, for example, ontology modeling (i.e., construction), ontology population (i.e., creating a knowledge graph with the given ontology as schema), ontology extension and modification, ontology alignment, entity disambiguation (sometimes called co-reference resolution).

All of the just mentioned tasks have in common that they are *hard*, in the sense that even after more than a quarter century of Semantic Web research, they still defy attempts to automate them at reasonable quality levels and scale. The state of the art on all of these is that they require significant human expert labor, at times (such as for entity disambiguation) together with detailed scripting of algorithms that solve the problem at scale but only for a specific problem instance, i.e., for a very specific knowledge graph and/or ontology.

At the same time, knowledge graphs and ontologies are ever more important for applications in data integration and data management, and more recently also as ground truth to escape from Large Language Model (LLM) hallucinations [29] and as components of other neurosymbolic approaches [21, 23, 24]. As a consequence, improved processes and methods for automating or even semi-automating core KGOE tasks remains a key challenge for the research community.

LLMs enter the scene, and the public perception of Artificial Intelligence (AI), with force in 2022 at the launch of OpenAI's ChatGPT,³ with rapid developments since then. Their human-style conversation capabilities, which include a profound mastery of expression in written language, as well as solid production of more structured information, are as stunning as their sometimes wildly confabulated responses (usually termed *hallucinations*). Perhaps most important for our discussion herein, LLMs appear to capture, and to have the ability to recall in different formats and contexts, a wide swath of human knowledge, both commonsense and specialized, provided it is reflected well enough in the training data. While at this point in time, their reliability in terms of accuracy of content in their responses remains problematic, it is quite apparent and widely reported that working with an LLM can save significant time and effort provided there is a (human) topic expert available as a check on factual accuracy.

The promise of LLMs for KGOE is thus: by using LLMs as approximate natural language knowledge bases that can be approximately queried, plus LLM capabilities to understand and produce information in ways structured for KGOE use, it should be possible to design semi-automatic methods, or human-LLM interactive methods, that can produce at least *draft solutions* for key KGOE problems at a level of quality that will significantly reduce human expert time and effort. Work on this line of research has of course already been started by the Semantic Web community. With this paper, we intend to begin to consolidate the discussion, and – in particular – contribute observations and discussions related to our own research on modular ontologies, which we believe to be highly relevant, for reasons that we will lay out below.

The plan of this paper is as follows. In Section 2 we motivate the need for (a certain type of) modularity for LLM-based

Email addresses: cogan.shimizu@wright.edu (Cogan Shimizu), hitzler@ksu.edu (Pascal Hitzler)

¹We mostly understand ontologies as a type of *schema* for knowledge graphs in the sense of [45], and as such the ontology can in fact also be understood to be *part* of the knowledge graph. We acknowledge other uses, to which our discussion also applies.

²For background on this and the more general Semantic Web field, see, e.g., [20] and the references given therein.

³<https://openai.com/index/hello-gpt-4o/>

KGOE. In Section 3 we discuss our notion of modularity and its research context. In Section 4 we briefly look, in turn, at the key KGOE tasks identified in the introduction. In Section 5 we will list some concrete research challenges that can drive LLM-based KGOE forward, and in Section 6 we will conclude.

2. Divide and Conquer

Useful ontologies and knowledge graphs often tend to be large, or even very large. The ENVO ontology [6] has over 2,100 classes; the Gene Ontology [7] has over 40,000 classes; a Wikipedia-derived class hierarchy used for explainable deep learning analysis has over $2 \cdot 10^6$ classes [41]; DBpedia Core⁴ has about $9 \cdot 10^8$ triples with an ontology of about 800 classes, KnowWhereGraph [48] has over $28 \cdot 10^9$ triples with an ontology of about 300 classes. The size – in particular of the ontology underlying a knowledge graph – is one of the obstacles facing humans in KGOE tasks, as even an ontology with a “mere” 300 classes (and corresponding comprehensive OWL axiomatization) is simply too big for a human to keep an overview of as a whole, or to understand well from scratch within a reasonable amount of time.

Even before the advent of LLMs, therefore, Semantic Web researchers have looked into the use of ontology modules in order to segment large ontologies into pieces of a size that are more workable for the human brain. These efforts include this paper’s authors and their collaborators, in a line of work that has developed out of Ontology Design Patterns (ODPs [3, 14]) research and practice, using an approach that we believe to be particularly suited for LLM-based KGOE, and that we will focus on further below: If an ontology is split into (even possibly overlapping) pieces that make conceptual sense (as a type of mini-ontology) for the domain experts, then KGOE tasks can often focus on one or a few coherent pieces (modules) of the ontology, thus simplifying the task significantly. And given the frequent difficulties (and expense) of LLMs to deal with substantial size prompts and with more open-ended and less common scenarios, it is a reasonable expectation that similar benefits will apply to LLM-based KGOE.

Let us provide a case in point from the context of ontology alignment, reported in [1]. The setting is the creation of complex ontology mapping rules between two ontologies in the OAEI GeoLink complex ontology alignment benchmark [52]. The benchmark is of moderate size, with 40 classes, 149 object properties and 49 data properties in one of the ontology, and 156 classes, 124 object properties and 46 data properties in the other ontology. It is a natural rather than synthetic benchmark in that the ontologies were originally developed for an application purpose, and were only cast into a benchmark later. The benchmark (like all complex ontology alignment benchmarks) had defied automation for years, the only approaches that were able to create reasonably good results assumed a shared ABox

(i.e. a shared data graph) [39, 38], which is of course a very unrealistic assumption for data management practice.

As reported in [1], an LLM was prompted to produce a body of an alignment rule, given a rule head together with the two ontologies as part of the prompt. This essentially completely failed, i.e., the LLM produced output that was essentially unusable. But then we made use of the modular structure of the body-side ontology, for which 20 named modules (such as “Organization” or “Physical Sample”) had already been provided at original deployment of the ontology [27]. We thus prompted in two stages: given an alignment rule head and the list of 20 module names, we first asked for the module(s) that would be required to create the rule. Then we prompted for the rule body by providing the modules previously identified. The results we obtained were really good in terms of high precision and recall. The availability and principled use of modules made the difference between an almost complete failure and a reasonably accurate system response.

3. Modules

The term *module* has been used for different things in the Semantic Web context.⁵ Our approach, as in [45], comes out of the tradition of ODPs [3, 14], which are partial ontologies that address commonly occurring, domain invariant ontology design issues. Later, related notions with similar underlying ideas, but generally different objectives (and implementations), include [35, 50].

For us, and this is laid out in [25, 45], a module is a part of an ontology that consists of the classes, properties, and axioms within that ontology that are relevant to a key notion, as considered by a domain expert, for the ontology use case. For example, the above mentioned GeoLink ontologies that focus on oceanographic cruises and data, such key notions (and corresponding modules) include *Trajectory*, *Cruise*, *Physical Sample*, *Organization*. The *Enslaved.org* ontology about historic person and events data on the transatlantic slave trade [47] has modules such as *Event*, *Place*, *Provenance*, *PersonRecord*, *AgeRecord*, etc. Modules can overlap or even be nested, and there are, on purpose, no precise rules which ontology pieces should be thought of as being part of a specific module. Rather, it has to *make sense* from several perspectives, including the domain experts’ perspectives, the use case perspectives, and the perspectives of the available data relevant for the use case and ontology. As such, the modules provide *conceptual bridges* between human expert conceptualization and data reality [25].

Our approach called MOMo for *Modular Ontology Modeling*, also includes a step-by-step ontology design methodology geared towards humans, a module (and pattern) description language OPLa expressed in OWL [22], ontology design pattern libraries [46], and a Protégé plug-in for ontology development [43, 44] – we mention these just in passing, and details can be

⁴<https://www.dbpedia.org/resources/latest-core/>

⁵See, e.g., the proceedings of the WoMO – Workshops on Modular Ontologies, details at <https://iaoa.org/womo/history.html>.

found in [45]; for this paper, the modular structure of the resulting ontologies themselves plus the fact that modules in OPLa annotated modular ontologies be easily be identified programmatically, are really the most important.

Knowledge graphs whose schema is given as a MOMo modular ontology then naturally inherit the modular structure from the ontology: The corresponding graph (ABox) module consists of all those ABox statements that are fully within a given ontology module.

MOMo was designed before the advent of LLMs, geared primarily towards human ontology engineers. However it has also always been on our minds that it should also be helpful for working towards (semi-)automation of hard KGOE tasks. LLMs may now provide the opportunity for the Semantic Web community to close the remaining gaps.

4. Modular LLM-based KGOE

We will now look, in turn, at the key KGOE tasks identified in the introduction and discuss each of them from the perspective of modularity-driven LLM-based KGOE. Specifically, we address the different tasks in essentially in order of abstraction, i.e., from knowledge to data. We note of course that inroads made in any tasks easily have the capacity to impact and improve outcomes in the others.

4.1. Ontology Modeling, Extension and Modification

Automated ontology design – often referred to as Ontology Learning – has been investigated primarily from a traditional (i.e., pre-LLM) Natural Language Processing (NLP) perspective, as a possible way to address the knowledge acquisition bottleneck with much of the methods established in the early aughts; see, e.g., [30, 31, 10, 12, 5, 11]. However, as the proliferation of early ODP-based methodologies (e.g., eXtreme Design [40]) grew, newer attempts to accomplish ontology learning emerged, leveraging identification of candidate patterns and subsequent human intervention [4].

In remaining pre-LLM years, ontology learning incrementally advanced from these early achievements and established techniques, generally through incremental evolution, inductive logic programming and both linguistic- and statistical-focused NLP [2, 51].⁶ However, results to date have been rather mixed, with resulting ontologies far from being able to compare with carefully crafted ontologies by domain and ontology engineering experts. The promise of LLMs now in this context comes from the fact that they are very powerful NLP tools that approximately capture a wide swath of human (expert) knowledge and can be coaxed by good prompting to even provide this knowledge in a structured form, such as expressed in a formal language over a given vocabulary.

⁶Indeed, the capacity of modern hardware allowed some techniques to newly shine.

Indeed, some attempts have already been made for LLMs, with some middling success [42, 15]. At this stage, LLMs are too limited in commonsense and reasoning power to in a fully automated fashion. Yet, we can thus expect that LLMs should enable a human ontology engineers and domain experts to first draft and then finalize suitable ontologies much more quickly than before.

It appears to us that specific aspects of the MOMo methodology, and generally the modularity idea, should further an LLM-based approach even more:

1. MOMo is based on a principled use of (high-quality) ontology design patterns (ODPs). Developing ODPs is arguably easier than developing full-fledged ontologies, i.e. in a first step LLMs can be used to develop ODP libraries, which then in turn can be made use of in LLM-based ontology design. As a first step, we have used an LLM to generate hundreds of simple “commonsense” patterns for common concepts [13]. These micropatterns – so called due to their shallow semantics and, even for ODPs, simple structure – have been organized into a design library [46], which allows for programmatic access, such as through a RAG [29] system. Due to their simplicity, they are easily instantiated into modules [17] and connected together in a modular fashion.
2. MOMo provides a step-by-step ontology design process that breaks down the complex ontology modeling task into clearly delineated pieces, each of which should be easier to automate than going one-shot from base data (or texts) to an ontology.

The situation for extension and modification of a MOMo ontology is rather similar – in the MOMo approach an ontology is designed module-by-module (with possible modification of earlier developed modules while progressing), i.e. extension and modification are already part of the modeling process.

4.2. Ontology Alignment

Our experience with LLM-based modular ontology alignment was already conveyed above, in Section 2. Ontology Alignment is a core task for ontology-based data integration with a long-standing corresponding research community, benchmarks and annual performance competitions.⁷ The community has mostly focused on full automation (as opposed to human-in-the-loop semi-automation) and on so-called *simple* alignment tasks, i.e. the creation of one-to-one class mappings (and, with much less emphasis, one-to-one property mappings). The need for *complex* ontology alignment – i.e., the creation of mapping *rules* that go beyond one-to-one mapping – has long been recognized, but it has only come into more focus in the community about 10 years ago, with very limited results (as pointed out in Section 2, see also [39, 38]. Regretfully (we think), since the advent of LLMs focus has mostly shifted back to simple alignments, just now with the support of LLMs. Indeed, LLM-based

⁷See <https://om.ontologymatching.org/> for pointers.

simple ontology alignment has already been investigated with good results [33, 19].

We would be negligent to not point out that we also observe issues with some of the established current benchmarks and methods. For example, in [9] a simple alignment benchmark was re-evaluated, finding that even humans disagree on a significant part of the benchmark, essentially demonstrating that the benchmark ontologies simply do not carry enough information such that even humans can solve it perfectly. In [8], it was demonstrated that most of the functionality of simple alignment systems can be obtained by using only string similarity metrics. At the same time, when submitting papers for publication that described complex alignment benchmarks that came out of real data integration work [53, 54], some reviewers dismissively pointed out that they would not be good as benchmarks because they would be far beyond anything that could currently be done.⁸

For some years, progress on the complex alignment task stalled [34], even as new benchmarks were released [55]. However, some new techniques, even incorporating LLMs have shown some promise. For example, the composition of (geometric) language embeddings cast from a KG or sourced from an LLM, can be shown to have correspondences in the latent space [49]. Yet, the performance is still relatively low, ranging from 0.49–0.69 for Semantic F1-score, depending on the ontologies. As described above in Section 2, it was demonstrated in [1] that modularity provides significant performance improvements to LLM-based complex alignment, in this case with 104 out of 109 (i.e., 95%) target alignment mappings correctly identified on the GeoLink benchmark when taking modularity into account.

4.3. Ontology Population

In some sense, ontology population and entity co-reference resolution/disambiguation are not substantially different. On one hand there is a *theoretical or unknown* entity which satisfies, or otherwise conforms to, a portion of an ontology (i.e., it should populate the ontology). On the other hand, there is a known entity described, or otherwise named, in the natural language corpus. Now the system needs to determine if those two entities *match* to move forward in populating the ontology.

In this case, we can leverage two aspects of modularity for ontology population with LLMs: conceptual consistency and tight scope. That is, the conceptual consistency of a module – based on our definition, of course – means that the constituent classes and properties of a module somehow belong together, especially from a human perspective. This in turn provides a tighter cluster of “sentiment” (for lack of a better term), priming an LLM in the prompt.

On the other hand, we have seen that LLMs tend to be better at following patterns, rather than instructions, and it correlates as well to the length of the prompt [28]. Thus, we posited that

⁸One of the reasons for creating these benchmarks was precisely to point out that the state of the art is way behind practical needs.

attempting ontology population on a per-module basis would succeed where other attempts had poorer results.

Indeed, in a recent set of experiments, we obtained rather excellent results [32]. Using prompts constructed with a simple schematic representations of modules and a corresponding extraction example, an LLM was able to achieve $\tilde{90}\%$ extraction of related triples from text, as compared to ground truth. While additional experiments are of course required, especially examining appropriate modeling characteristics for modules, we take this as an excellent indicator that modularity is of significant added value to the process.

4.4. Entity Disambiguation

As above, we note that ontology population and entity disambiguation are tightly intertwined. While ontology population focuses on the extraction of an entity and (ostensibly) comparing it to a theoretical candidate, entity disambiguation compares two known, or otherwise named, entities and must produce a value corresponding to the degree to which they are the same entity.

However, we also note that it is difficult to assess the current capabilities of LLMs to produce correct entity disambiguation due to the possibility for data leakage (i.e., the appearance of existing and gold-standard benchmarks in the testing data). For example, some benchmarks (after prompt-tuning) achieve upwards of $F_1 = 91$.

As successful co-reference resolution correlates with available context, per improved ability to resolve matching context (i.e., complex ontology alignment) we see immediate benefit in this parallel task of entity disambiguation. We thus posit that the added context from a tight, conceptual description, i.e., a module, will significantly improve outcomes for this KGOE task.

5. Concrete Research Challenges

Throughout this paper so far, we have largely discussed a specific type of module, which we have been calling a *conceptual module* and distinct from other uses in the literature for both ontologies [37, 26] and KGs [18]. We are not aware of any work on the capabilities of LLMs for integrating for these other definitions of module. We suspect that due to their (in general) rigorous approach to the partitioning of knowledge in a logically consistent way, we may see similar benefits as when using our notion of modularity. That is to say, by somehow limiting the scope, we achieve a more human-like approach – and one more capable of being expressed succinctly in language, and thus more appropriate for LLM-based assistance in the task.

Of course, it is important to not altogether *neglect* further research on non-modular, LLM-based KGOE; we need to understand its limitations, as well as its capabilities. Indeed, additional focus on modular or pattern-based techniques for KGOE can inform these other lines of research.

From a broader perspective, the added value of modularity that seems to become apparent for LLM-based KGOE prompts the

question what other improvements to established ontology and knowledge graph paradigms would be similarly easy to make and could also provide major benefits. It is the task of the Semantic Web, Knowledge Graphs, and Ontologies research communities to look into this very question.

6. Conclusion

Knowledge Graphs and Ontologies have seen new resurgence in the era of LLMs, for example in structuring data for RAG systems [36] or providing guide-rails to limit or prevent confabulation in textual output [16]. Yet, for all their established and growing importance, many of the tasks pertaining to their development, maintenance, extension, and population (to name a few) are still very difficult. It seems somehow appropriate that we incorporate LLMs into improving KGOE tasks. LLMs have shown some initial success in *simple* tasks, such as the generation of simple schemas or topic extraction.

To tackle the *hard* tasks, we have proposed the use of modules – conceptual, human-centric partitions of an ontology or schema – that provide internal structure. These conceptual boundaries assist LLMs in a variety of tasks, seeing already success in ontology construction, (complex) alignment, and population. It seems obvious now that modularity is a *missing link* for bridging human conceptualization and machine interoperability. We thus fully believe that modularity must be incorporated from the start, both structurally and in documentation, so as to further enable the various improvements to KGOE tasks we have outlined above.

Acknowledgment. This work was supported by the National Science Foundation under award 2333532 "Proto-OKN Theme 3: An Education Gateway for the Proto-OKN" (EduGate).

References

- [1] R. Amini, S. S. Norouzi, P. Hitzler, and R. Amini. Towards complex ontology alignment using large language models. In *Proceedings KGSWC 2024*, to appear, 2024.
- [2] M. N. Asim, M. Wasim, M. U. G. Khan, W. Mahmood, and H. M. Abbasi. A survey of ontology learning techniques and applications. *Database J. Biol. Databases Curation*, 2018:bay101, 2018.
- [3] E. Blomqvist. *Semi-automatic Ontology Construction based on Patterns*. PhD thesis, Linköping University, 2009.
- [4] E. Blomqvist. *Semi-automatic Ontology Construction based on Patterns*. PhD thesis, Linköping University, Sweden, 2009.
- [5] P. Buitelaar and P. Cimiano, editors. *Ontology Learning and Population: Bridging the Gap between Text and Knowledge*, volume 167 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2008.
- [6] P. L. Buttigieg, E. Pafilis, S. E. Lewis, M. P. Schildhauer, R. L. Walls, and C. J. Mungall. The environment ontology in 2016: bridging domains with increased scope, semantic density, and interoperation. *J. Biomed. Semant.*, 7:57, 2016.
- [7] S. Carbon, E. Douglass, B. M. Good, D. R. Unni, N. L. Harris, C. J. Mungall, et al. The gene ontology resource: enriching a gold mine. *Nucleic Acids Res.*, 49(Database-Issue):D325–D334, 2021.
- [8] M. Cheatham and P. Hitzler. String similarity metrics for ontology alignment. In H. Alani, L. Kagal, A. Fokoue, P. Groth, C. Biemann, J. X. Parreira, L. Aroyo, N. F. Noy, C. Welty, and K. Janowicz, editors, *The Semantic Web – ISWC 2013 - 12th International Semantic Web Conference, Sydney, NSW, Australia, October 21-25, 2013, Proceedings, Part II*, volume 8219 of *Lecture Notes in Computer Science*, pages 294–309. Springer, 2013.
- [9] M. Cheatham and P. Hitzler. Conference v2.0: An uncertain version of the OAEI conference benchmark. In P. Mika, T. Tudorache, A. Bernstein, C. Welty, C. A. Knoblock, D. Vrandečić, P. Groth, N. F. Noy, K. Janowicz, and C. A. Goble, editors, *The Semantic Web - ISWC 2014 - 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part II*, volume 8797 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2014.
- [10] P. Cimiano, A. Hotho, and S. Staab. Learning concept hierarchies from text corpora using formal concept analysis. *J. Artif. Intell. Res.*, 24:305–339, 2005.
- [11] P. Cimiano, A. Mädche, S. Staab, and J. Völker. Ontology learning. In S. Staab and R. Studer, editors, *Handbook on Ontologies*, International Handbooks on Information Systems, pages 245–267. Springer, 2009.
- [12] P. Cimiano and J. Völker. Text2onto. In A. Montoyo, R. Muñoz, and E. Métais, editors, *Natural Language Processing and Information Systems*, pages 227–238, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [13] A. Eells, B. Dave, P. Hitzler, and C. Shimizu. Commonsense ontology micropatterns. In T. R. Besold, A. d’Avila Garcez, E. Jiménez-Ruiz, R. Confalonieri, P. Madhyastha, and B. Wagner, editors, *Neural-Symbolic Learning and Reasoning - 18th International Conference, NeSy 2024, Barcelona, Spain, September 9-12, 2024, Proceedings, Part II*, volume 14980 of *Lecture Notes in Computer Science*, pages 51–59. Springer, 2024.
- [14] A. Gangemi. Ontology design patterns for semantic web content. In Y. Gil, E. Motta, V. R. Benjamins, and M. A. Musen, editors, *The Semantic Web – ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6-10, 2005, Proceedings*, volume 3729 of *Lecture Notes in Computer Science*, pages 262–276. Springer, 2005.
- [15] H. B. Giglou, J. D’Souza, and S. Auer. Llm4ol: Large language models for ontology learning. In T. R. Payne, V. Presutti, G. Qi, M. Poveda-Villalón, G. Stoilos, L. Hollink, Z. Kaoudi, G. Cheng, and J. Li, editors, *The Semantic Web - ISWC 2023 - 22nd International Semantic Web Conference, Athens, Greece, November 6-10, 2023, Proceedings, Part I*, volume 14265 of *Lecture Notes in Computer Science*, pages 408–427. Springer, 2023.
- [16] L. H. Gilpin, V. Penubarthi, and L. Kagal. Explaining multimodal errors in autonomous vehicles. In *8th IEEE International Conference on Data Science and Advanced Analytics, DSAA 2021, Porto, Portugal, October 6-9, 2021*, pages 1–10. IEEE, 2021.
- [17] K. Hammar and V. Presutti. Template-based content ODP instantiation. In K. Hammar, P. Hitzler, A. Krisnadhi, A. Lawrynowicz, A. G. Nuzzolese, and M. Solanki, editors, *Advances in Ontology Design and Patterns [revised and extended versions of the papers presented at the 7th edition of the Workshop on Ontology and Semantic Web Patterns, WOP@ISWC 2016, Kobe, Japan, 18th October 2016]*, volume 32 of *Studies on the Semantic Web*, pages 1–13. IOS Press, 2016.
- [18] F. Hannou, V. Charpenay, M. Lefrançois, C. Roussey, A. Zimmermann, and F. Gandon. The ACIMOV methodology: Agile and continuous integration for modular ontologies and vocabularies. In F. Toyoshima, M. Katsumi, G. Righetti, S. D. Giorgis, M. M. Hedblom, O. Kutz, B. Brodaric, M. Grüninger, T. Hahmann, D. D. Simon, M. Lange, H. Küçük-McGinty, A. S. Simon, R. Cameron, B. Gajderowicz, D. Rosu, J. Hastings, D. Audrito, L. D. Caro, F. Grasso, R. Nai, E. Sulis, L. Bozzato, C. Shimizu, A. Zimmermann, R. Baratella, S. Borgo, S. de Cesare, and T. P. Sales, editors, *Proceedings of the Joint Ontology Workshops 2023 Episode IX: The Quebec Summer of Ontology co-located with the 13th International Conference on Formal Ontology in Information Systems (FOIS 2023), Sherbrooke, Québec, Canada, July 19-20, 2023*, volume 3637 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2023.
- [19] S. Hertling and H. Paulheim. Olala: Ontology matching with large language models. In K. B. Venable, D. Garijo, and B. Jalaian, editors, *Proceedings of the 12th Knowledge Capture Conference 2023, K-CAP 2023, Pensacola, FL, USA, December 5-7, 2023*, pages 131–139. ACM, 2023.
- [20] P. Hitzler. A review of the semantic web field. *Commun. ACM*, 64(2):76–83, 2021.
- [21] P. Hitzler, A. Dalal, M. S. Mahdavinejad, and S. S. Norouzi, editors. *Neurosymbolic Artificial Intelligence and Knowledge Graphs (tentative title)*. Frontiers in Artificial Intelligence and Applications. IOS Press, 2025. In

- preparation.
- [22] P. Hitzler, A. Gangemi, K. Janowicz, A. A. Krisnadhi, and V. Presutti. Towards a simple but useful ontology design pattern representation language. In E. Blomqvist, Ó. Corcho, M. Horridge, D. Carral, and R. Hoekstra, editors, *Proceedings of the 8th Workshop on Ontology Design and Patterns (WOP 2017) co-located with the 16th International Semantic Web Conference (ISWC 2017)*, Vienna, Austria, October 21, 2017., volume 2043 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2017.
- [23] P. Hitzler and M. K. Sarker, editors. *Neuro-Symbolic Artificial Intelligence: The State of the Art*, volume 342 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2021.
- [24] P. Hitzler, M. K. Sarker, and A. Eberhart, editors. *Compendium of Neurosymbolic Artificial Intelligence*, volume 369 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2023.
- [25] P. Hitzler and C. Shimizu. Modular ontologies as a bridge between human conceptualization and data. In P. Chapman, D. Endres, and N. Pernelle, editors, *Graph-Based Representation and Reasoning - 23rd International Conference on Conceptual Structures, ICCS 2018, Edinburgh, UK, June 20-22, 2018, Proceedings*, volume 10872 of *Lecture Notes in Computer Science*, pages 3–6. Springer, 2018.
- [26] P. Koopmann and J. Chen. Deductive module extraction for expressive description logics. In C. Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 1636–1643. International Joint Conferences on Artificial Intelligence Organization, 7 2020. Main track.
- [27] A. Krisnadhi, Y. Hu, K. Janowicz, P. Hitzler, R. A. Arko, S. Carbotte, C. Chandler, M. Cheatham, D. Fils, T. W. Finin, P. Ji, M. B. Jones, N. Karima, K. A. Lehnert, A. Mickle, T. W. Narock, M. O’Brien, L. Raymond, A. Shepherd, M. Schildhauer, and P. Wiebe. The geolink modular oceanography ontology. In M. Arenas, Ó. Corcho, E. Simperl, M. Strohmaier, M. d’Aquin, K. Srinivas, P. T. Groth, M. Dumontier, J. Heflin, K. Thirunarayan, and S. Staab, editors, *The Semantic Web – ISWC 2015 – 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II*, volume 9367 of *Lecture Notes in Computer Science*, pages 301–309. Springer, 2015.
- [28] M. Levy, A. Jacoby, and Y. Goldberg. Same task, more tokens: the impact of input length on the reasoning performance of large language models, 2024.
- [29] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.
- [30] A. Maedche and S. Staab. Discovering conceptual relations from text. In *Proceedings of the 14th European Conference on Artificial Intelligence, ECAI’00*, page 321–325, NLD, 2000. IOS Press.
- [31] A. Maedche and S. Staab. Ontology learning for the semantic web. *IEEE Intell. Syst.*, 16(2):72–79, 2001.
- [32] S. S. Norouzi, A. Barua, A. Christou, N. Gautam, A. Eells, P. Hitzler, and C. Shimizu. Ontology population using LLMs, 2024.
- [33] S. S. Norouzi, M. S. Mahdavejad, and P. Hitzler. Conversational ontology alignment with chatgpt. In P. Shvaiko, J. Euzenat, E. Jiménez-Ruiz, O. Hassanzadeh, and C. Trojahn, editors, *Proceedings of the 18th International Workshop on Ontology Matching co-located with the 22nd International Semantic Web Conference (ISWC 2023)*, Athens, Greece, November 7, 2023, volume 3591 of *CEUR Workshop Proceedings*, pages 61–66. CEUR-WS.org, 2023.
- [34] OAEI 2020 complex track results. <http://oaei.ontologymatching.org/2020/results/complex/index.html>. Retrieved November 18, 2020.
- [35] D. Osumi-Sutherland, M. Courtot, J. P. Balhoff, and C. J. Mungall. Dead simple OWL design patterns. *J. Biomed. Semant.*, 8(1):18:1–18:7, 2017.
- [36] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*, 36(7):3580–3599, July 2024.
- [37] C. Parent and S. Spaccapietra. An overview of modularity. In H. Stuckenschmidt, C. Parent, and S. Spaccapietra, editors, *Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization*, pages 5–23. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [38] M. A. N. Pour, A. Algergawy, F. Amardeilh, R. Amini, O. Fallatah, D. Faria, I. Fundulaki, I. Harrow, S. Hertling, P. Hitzler, M. Huschka, L. Ibanescu, E. Jiménez-Ruiz, N. Karam, A. Laadhar, P. Lambrix, H. Li, Y. Li, F. Michel, E. Nasr, H. Paulheim, C. Pesquita, J. Portisch, C. Roussey, T. Saveta, P. Shvaiko, A. Splendiani, C. Trojahn, J. Vataschinová, B. Yaman, O. Zamazal, and L. Zhou. Results of the ontology alignment evaluation initiative 2021. In P. Shvaiko, J. Euzenat, E. Jiménez-Ruiz, O. Hassanzadeh, and C. Trojahn, editors, *Proceedings of the 16th International Workshop on Ontology Matching co-located with the 20th International Semantic Web Conference (ISWC 2021)*, Virtual conference, October 25, 2021, volume 3063 of *CEUR Workshop Proceedings*, pages 62–108. CEUR-WS.org, 2021.
- [39] M. A. N. Pour, A. Algergawy, R. Amini, D. Faria, I. Fundulaki, I. Harrow, S. Hertling, E. Jiménez-Ruiz, C. Jonquet, N. Karam, A. Khat, A. Laadhar, P. Lambrix, H. Li, Y. Li, P. Hitzler, H. Paulheim, C. Pesquita, T. Saveta, P. Shvaiko, A. Splendiani, É. Thiéblin, C. Trojahn, J. Vataschinová, B. Yaman, O. Zamazal, and L. Zhou. Results of the ontology alignment evaluation initiative 2020. In P. Shvaiko, J. Euzenat, E. Jiménez-Ruiz, O. Hassanzadeh, and C. Trojahn, editors, *Proceedings of the 15th International Workshop on Ontology Matching co-located with the 19th International Semantic Web Conference (ISWC 2020)*, Virtual conference (originally planned to be in Athens, Greece), November 2, 2020, volume 2788 of *CEUR Workshop Proceedings*, pages 92–138. CEUR-WS.org, 2020.
- [40] V. Presutti, E. Daga, A. Gangemi, and E. Blomqvist. eXtreme Design with content ontology design patterns. In E. Blomqvist, K. Sandkuhl, F. Scharffe, and V. Svátek, editors, *Proceedings of the Workshop on Ontology Patterns (WOP 2009)*, collocated with the 8th International Semantic Web Conference (ISWC-2009), Washington D.C., USA, 25 October, 2009., volume 516 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2009.
- [41] M. K. Sarker, J. Schwartz, P. Hitzler, L. Zhou, S. Nadella, B. S. Minnery, I. Juvina, M. L. Raymer, and W. R. Aue. Wikipedia knowledge graph for explainable AI. In B. Villazón-Terrazas, F. Ortiz-Rodríguez, S. M. Tiwari, and S. K. Shandilya, editors, *Knowledge Graphs and Semantic Web - Second Iberoamerican Conference and First Indo-American Conference, KGSWC 2020, Mérida, Mexico, November 26-27, 2020, Proceedings*, volume 1232 of *Communications in Computer and Information Science*, pages 72–87. Springer, 2020.
- [42] M. Schaeffer, M. SesboÅ¼ÁÁ©, J.-P. Kotowicz, N. Delestre, and C. Zanni-Merk. Olaf: An ontology learning applied framework. *Procedia Computer Science*, 225:2106–2115, 2023. 27th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES 2023).
- [43] C. Shimizu and K. Hammar. CoModIDE – the Comprehensive Modular Ontology Engineering IDE. In M. C. Suárez-Figueroa, G. Cheng, A. L. Gentile, C. Guéret, C. M. Keet, and A. Bernstein, editors, *Proceedings of the ISWC 2019 Satellite Tracks (Posters & Demonstrations, Industry, and Outrageous Ideas) co-located with 18th International Semantic Web Conference (ISWC 2019)*, Auckland, New Zealand, October 26-30, 2019, volume 2456 of *CEUR Workshop Proceedings*, pages 249–252. CEUR-WS.org, 2019.
- [44] C. Shimizu, K. Hammar, and P. Hitzler. Modular graphical ontology engineering evaluated. In A. Harth, S. Kirrane, A. N. Ngomo, H. Paulheim, A. Rula, A. L. Gentile, P. Haase, and M. Cochez, editors, *The Semantic Web - 17th International Conference, ESWC 2020, Heraklion, Crete, Greece, May 31-June 4, 2020, Proceedings*, volume 12123 of *Lecture Notes in Computer Science*, pages 20–35. Springer, 2020.
- [45] C. Shimizu, K. Hammar, and P. Hitzler. Modular ontology modeling. *Semantic Web*, 14(3):459–489, 2023.
- [46] C. Shimizu, Q. Hirt, and P. Hitzler. MODL: A modular ontology design library. In K. Janowicz, A. A. Krisnadhi, M. P. Villalón, K. Hammar, and C. Shimizu, editors, *Proceedings of the 10th Workshop on Ontology Design and Patterns (WOP 2019) co-located with 18th International Semantic Web Conference (ISWC 2019)*, Auckland, New Zealand, October 27, 2019, volume 2459 of *CEUR Workshop Proceedings*, pages 47–58. CEUR-WS.org, 2019.
- [47] C. Shimizu, P. Hitzler, Q. Hirt, D. Rehberger, S. G. Estrecha, C. Foley, A. M. Sheill, W. Hawthorne, J. Mixer, E. Watrall, R. Carty, and D. Tarr. The enslaved ontology: Peoples of the historic slave trade. *J. Web Semant.*, 63:100567, 2020.
- [48] C. Shimizu, S. Stephe, A. Barua, L. Cai, A. Christou, K. Currier, A. Dalal, C. K. Fisher, P. Hitzler, K. Janowicz, W. Li, Z. Liu, M. S. Mahdavejad, G. Mai, D. Rehberger, M. Schildhauer, M. Shi, S. S. Norouzi, Y. Tian,

- S. Wang, Z. Wang, J. Zalewski, L. Zhou, and R. Zhu. The KnowWhere-Graph ontology, 2024. arXiv 2410.13948.
- [49] M. C. Silva, D. Faria, and C. Pesquita. Complex multi-ontology alignment through geometric operations on language embeddings. In U. Endriss, F. S. Melo, K. Bach, A. J. B. Diz, J. M. Alonso-Moral, S. Barro, and F. Heintz, editors, *ECAI 2024 - 27th European Conference on Artificial Intelligence, 19-24 October 2024, Santiago de Compostela, Spain - Including 13th Conference on Prestigious Applications of Intelligent Systems (PAIS 2024)*, volume 392 of *Frontiers in Artificial Intelligence and Applications*, pages 1333–1340. IOS Press, 2024.
- [50] M. G. Skjæveland, D. P. Lupp, L. H. Karlsen, and H. Forssell. Practical ontology pattern instantiation, discovery, and maintenance with reasonable ontology templates. In D. Vrandečić, K. Bontcheva, M. C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L. Kaffee, and E. Simperl, editors, *The Semantic Web – ISWC 2018 – 17th International Semantic Web Conference, Monterey, CA, USA, October 8-12, 2018, Proceedings, Part I*, volume 11136 of *Lecture Notes in Computer Science*, pages 477–494. Springer, 2018.
- [51] J. Wotrobski. Ontology learning methods from text - an extensive knowledge-based approach. *Procedia Computer Science*, 176:3356–3368, 2020. Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 24th International Conference KES2020.
- [52] L. Zhou, M. Cheatham, A. Krisnadhi, and P. Hitzler. A complex alignment benchmark: Geolink dataset. In D. Vrandečić, K. Bontcheva, M. C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L. Kaffee, and E. Simperl, editors, *The Semantic Web - ISWC 2018 - 17th International Semantic Web Conference, Monterey, CA, USA, October 8-12, 2018, Proceedings, Part II*, volume 11137 of *Lecture Notes in Computer Science*, pages 273–288. Springer, 2018.
- [53] L. Zhou, M. Cheatham, A. Krisnadhi, and P. Hitzler. Geolink data set: A complex alignment benchmark from real-world ontology. *Data Intell.*, 2(3):353–378, 2020.
- [54] L. Zhou, C. Shimizu, P. Hitzler, A. M. Sheill, S. G. Estrecha, C. Foley, D. Tarr, and D. Rehberger. The enslaved dataset: A real-world complex ontology alignment benchmark using wikibase. In M. d’Aquin, S. Dietze, C. Hauff, E. Curry, and P. Cudré-Mauroux, editors, *CIKM ’20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, pages 3197–3204. ACM, 2020.
- [55] L. Zhou, É. Thiéblin, M. Cheatham, D. Faria, C. Pesquita, C. T. dos Santos, and O. Zamazal. Towards evaluating complex ontology alignments. *Knowl. Eng. Rev.*, 35:e21, 2020.