computer sciences and mathematics forum

MDPI

Proceeding Paper

Modelling SWRL Enabled OWL Ontologies for Neurosymbolic Computing in the Legal Domain †

Eiman Almami [1], Radmila Juric [2,*] and Anna Maria Cerenko [3]

[1] Independent Researcher, London NW8 8EN, UK; eimanalmami@pm.me
[2] Almais Consultancy, London E10 7ER, UK
[3] Independent Researcher, Atlanta, GA 30305, USA; amcerenko@oglethorpe.edu
* Correspondence: radmila@almais.net
† Presented at the 1st International Conference on Emerging Tech & Innovation (ICETI), Buraydah, Saudi Arabia, 10 February 2026.

Abstract

The rise of neurosymbolic AI promises to bridge the gap between learning technologies and symbolic computing and create models where predictive inference is complemented with logic reasoning and vice versa. Consequently, ontologies and knowledge graphs play important roles in defining neurosymbolic computing and adding to knowledge manipulation accompanied by GenAI. However, traditional ontologies may not be an exclusive driving force behind logic reasoning in neurosymbolic computing. This paper defines factors which define neurosymbolic-ready ontologies. The example is taken from the legal domain which illustrates that neurosymbolic computing empowers legal knowledge dissemination and exceeds the results obtained by generative AI.

Keywords: SWRL enabled OWL ontologies; logic reasoning; neurosymbolic computing; legal domain

1. Introduction

Neurosymbolic computing has come of age as a driving force in narrowing the gap between computing based on Neural Network (NN) powered Large Language Models (LLMs), prevalent in GenAI, and our traditional computing based on the symbolic presentation of the semantic of problem domains [1]. It is also well known that however impressive predictive inference is, it cannot carry the precision of symbolic computing. It has also not secured our trust in GenAI, because we believe in computing 'with facts' and in inference which is not probabilistic [2]. These are really two different worlds: symbolic and non-symbolic, which have their own merits and use, and, in an ideal world, they should co-exist and complement each other. We should not separate them.

Hopefully, neurosymbolic computing can add to our efforts to keep both these worlds together [3]. We can have symbolic and logic reasoning, within the world of AI and NN powered LLMs, 'handy' whenever we need computing precision and logic reasoning, and vice versa [4]. This would mean that the success of knowledge graph (KG) [5,6] in GenAI could make our progress towards neurosymbolic computing faster and consequently put ontological engineering, based on the Semantic Web Technologies (SWT) [7–9], in the spotlight again, more than 20 years after we standardized the SWT. KG in current AI is not successful by accident because its logical structures which can accommodate ontologies can significantly enhance results of predictive inference, particularly within GenAI technologies [10]. This has been considered good news for AI, but it has also opened another

Check for updates

Academic Editor: Jamal Nasser Alotaibi

Published: 12 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

set of questions. Are the ontologies and knowledge structures in KG ready to assist in AI? We have had a generation of successful ontological engineering efforts which started with SWT and resulted in numerous knowledge-based repositories based on OWL ontologies. They are static vocabularies with enormous power of describing the semantics of our outside world within computing environments which can be manipulated and retrieved. Unfortunately, they are not software application-friendly because of the huge computational overload that full-scale ontologies carry if we incorporate them into our everyday computing. We need lightweight computing within current AI if symbolic computing with reasoning will even be a part of neurosymbolic computing [11]. Therefore we need to re-think ontologies and KG if we wish to realize the ideas of the neurosymbolic paradigm.

This paper addresses the issue of ontology design for its long term deployment in neurosymbolic computing. This does not mean that we discard all our ontological static vocabularies created in the last 20 years across many problem domains. One important aspect of this research is to understand where traditional ontologies can be used in neurosymbolic computing and why. The paper also looks at the computational power of using SWRL [8] enabled OWL [9] ontologies which, in addition to ontological vocabularies, can make a difference to bringing ad-hoc reasoning to current AI.

The ideas of designing ontologies for neurosymbolic computing is illustrated in the legal domain by using the example of finding ***court case precedent*** [12]. This example was chosen for two reasons. The problem domain is text-rich and not very well supported by current GenAI tools due to the complexity of its semantics. Therefore neurosymbolic computing might be an answer. The second reason is that finding precedent court cases is a very good example for illustrating neurosymbolic-ready ontology design. This means that in this problem domain we can debate our main worries:

(a) Are well known and used ontologies in the legal domains directly deployable in neurosymbolic computing?
(b) Which ontologies, and reasoning upon their elements, enable a symbiosis of symbolic (FOL and logic reasoning) with the reasoning using NN powered LLMs of GenAI in the legal domain?

At the time of writing this paper, our answer to both questions is negative.

Finally neurosymbolic-ready ontologies have one important characteristic. They support the bi-directional relationship with the GenAI world. They can learn from GenAI question/answer systems by becoming learning ontologies and, at the same time, they are a validation mechanism, with logic reasoning, for evaluating results generated by predictive inference [13,14]. They might have a slightly different purpose within the same problem domain. Static ontological vocabulary is an excellent candidate for learning ontologies, and the ontological reasoning with SWRL can be a part of a wider computational model, i.e., a driving force to evaluate the results of NN-powered LLMs and increase our trust in AI [15,16].

This paper is organized as follows.

In Section 2 we outline the proposal by elaborating on the nature and purpose of having two types of ontologies in the same problem domain which will be neurosymbolic computing ready. Section 3 outlines the design of creating ontology as a static vocabulary and using SWRL to strengthen it. If we plan that this ontology becomes a learning ontology, then its semantic structure should correspond to taxonomies available in real life, in the domain of court case precedent documentation. Section 4 outlines another ontological design which secures computations upon its elements using first-order logic and reasoning with SWRL. Section 5 debates our proposal, assesses challenges of AI tools in the legal domain and outlines future research steps.

https://doi.org/10.3390/cmsf2026013015

2. The Proposal

We propose two interrelated ontologies which address the problems identified in the previous section.

Ontology A is a powerful controlled vocabulary for storing facts, knowledge and discoveries in the legal domain and possibly addressing a particular problem which can strengthen the semantics of the ontology. From this perspective our proposal might not be far from traditional ontologies derived in the last 3 decades. In other words, these OWL ontologies would enable knowledge extraction with languages like SPARQL and thus facilitate retrievals of the semantics stored in them.

Ontology B must be related to ***Ontology A*** because we deal with the same problem domain (legal facts and knowledge), but it will not be used for retrievals. ***Ontology B*** would be a core part of a wider computational model which would help us to answer numerous questions (called competency questions). They are special questions which are not generic. They are either related to a specific problem/person or derived in a certain context—hence the need for logic reasoning upon selective ontological concepts to infer answers. Therefore, the essence of ***Ontology B*** is to enable a reasoning process with a pathway for inference mechanisms based on first order logic (FOL) for a single purpose: answering a competency question [17,18].

The difference between these two ontologies is very subtle; i.e., it might not be visible immediately. However, ***Ontology B*** addresses deficiencies of controlled vocabularies (***Ontology A***) [19] and sits very well as a part of a lightweight data intensive software application which does not have to be computationally demanding [19,20]. As mentioned earlier, teaming them up together within one software application can create a powerful operating environment for neurosymbolic computing because they will have a different purpose: one (A) will be used as a learning ontology through its synergy with learning and GenAI technologies in neurosymbolic computing and the other (B) would be used for either answering questions in knowledge discovery or evaluating results of GenAI when answering the same question, using a specially designed reasoning process.

In order to illustrate the proposal, we use one problem from the legal domain which could be resolved using neurosymbolic computing. We chose to work on the problem of discovering ***court case precedents*** as one of the most time-consuming tasks for legal professionals.

We could not find any computational environment which promotes neurosymbolic computing and addresses both:

(a) An automated process of finding ***court case precedents*** through software (including GenAI tools);
(b) Verification or evaluation of these findings, particularly if we use GenAI (i.e., addressing our trust in GenAI answers).

In other words, there is no commercially available software which uses ontologies in the legal domain, as part of a computational model, expected to be incorporated into neurosymbolic solutions where SWRL enabled OWL ontologies play a crucial role.

Consequently, in this paper we create ontologies which are ***court case precedent*** specific. They are modelled according to the semantics stored in the ***court case precedent*** documentation and mostly applicable to the US and UK legal environments. They are generic but could be transferred into any other jurisdiction with minor adjustments. The structure of the documentation on US ***court case precedents*** is publicly available on the Internet and may include all information we used for modeling our OWL ontologies, defining constraints and populating the ontological elements with individuals. Finally, these ontologies prove the concept: we can use them for creating a neurosymbolic computational model which finds relevant ***court case precedents*** for a given request.

https://doi.org/10.3390/cmsf2026013015

3. Ontology A: Controlled Vocabulary

The ontological model from Figure 1 shows basic ontological classes which are taken from the semantics of US Supreme Court case structures, publicly available on the court webpages (US Supreme Court Documents) [21]. We initially focused on US Supreme Court cases because they have been immaculately documented for at least 150+ years. It could be argued that our ontological concepts are very specific and suitable solely for the US environment. However, we have discovered that the semantic of the classes from Figure 1 might be easily interpreted and applied in different legal domains/countries. Furthermore, we deliberately chose class names as they appear in the terminology of US Supreme Court documents, in order to allow readers to inspect the source of our ontological elements (individuals in particular) and evaluate the level of overlapping between the structure of the ontology with taxonomies of information stored on the US Supreme Court webpages [22].

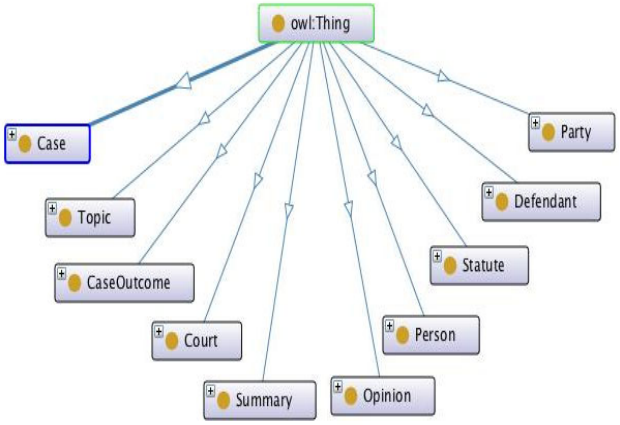

Figure 1. Base ontological model for court cases.

Finally, all ***court case precedents*** available on the US Supreme Court webpages are systematically structured and thus providing specific taxonomy, which helped to create a base ontological model in Figure 1.

Therefore Figure 1 collates base ontological classes: Case, Court, Defendant, Judge, Opinion, Party, Person, Statute, Summary, and Topic. However, Figure 2 explores their semantics further and highlights their potential class hierarchies. These subclasses fully represent the semantic of court cases: TariffCase (as an example of a court case), Petitioner-WinCase, RespondentWinCase, ConcurringOpinion, DissentingOpinion, MajorityOpinion, and Justice, to mention just a few. Due to space restriction, we cannot give detailed justification of our class subhierarchies, but we would encourage readers to inspect all ***precedent court cases*** available on the US website. This can also help to compare the contents of Figures 1 and 2 with their own local/country specific court case specifications.

We would like to draw readers' attention to the fact that this is data intensive software development and as such it requires careful thinking when presenting the semantics of real life court cases in environments as restricted as computational modelling is. The SWT and its languages have interesting provisions for defining and manipulating the semantic of court cases, and Figure 3 illustrates a selection of data type properties which capture data such as legal cases identification, title, summary, decision date, citation, docket number and argument date.

Finally Figure 4 is a snapshot of object properties which define constraints upon the individuals from the selection of the classes from Figures 1 and 2 (e.g., Case, Opinion, Party, Person, and Statute).

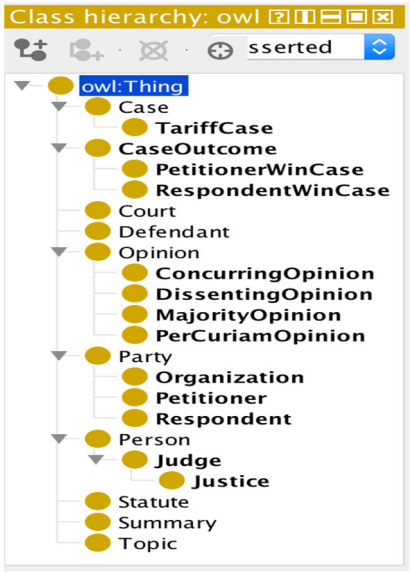

Figure 2. Proposed base class hierarchies.

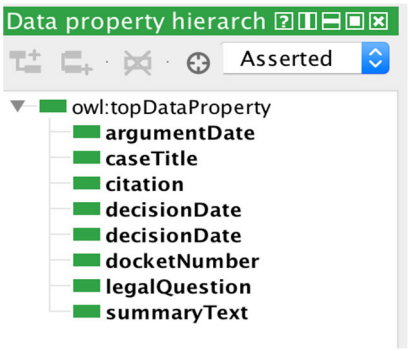

Figure 3. Proposed data properties.



Figure 4. Proposed object properties.

3.1. Strengthening Ontology A

The semantics of ***Ontology A*** can be enhanced by running a selection of SWRL rules which can add more to understanding the knowledge of the problem domain. In other words, SWRL rules may infer either classes or individuals or object properties into the existing ontology. This inference is not necessarily needed in advance, because it would depend on the facts we would like to retrieve from the ontology. If we are using SPARQL, we might be able to prepare this inference in advance for SPARQL queries.

We give two examples.

Let us assume that we wish to have a few more facts ready for retrieval. One fact would be the answer to this question:

*Who won a given case X (in our example the case is **Roe versus Wade**).*

We can answer this question using the 'winner' object property alongside the 'petitioner' and 'respondent' information. We needed to run three SWRL rules given in Figures 5–7. The rules create a new inference, i.e., the name of the winner in this court case, ***'Jane Roe' (Petitioner)***, available in the screenshot from Figure 8. In this example we infer an individual 'Jane Roe' within the class WINNER.

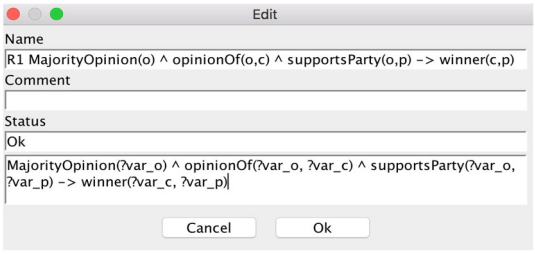

Figure 5. Inferring Winner (majority of opinions).

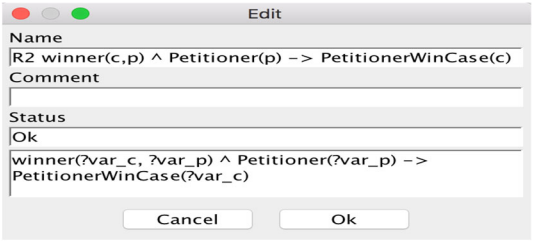

Figure 6. Inferring PetitionerWinCase.

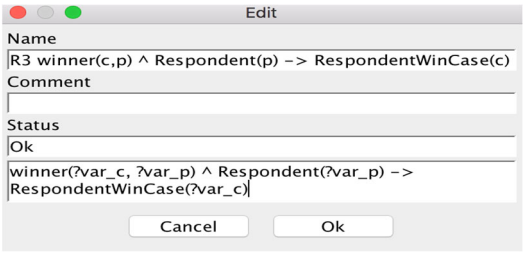

Figure 7. Inferring RespondentWinCase.

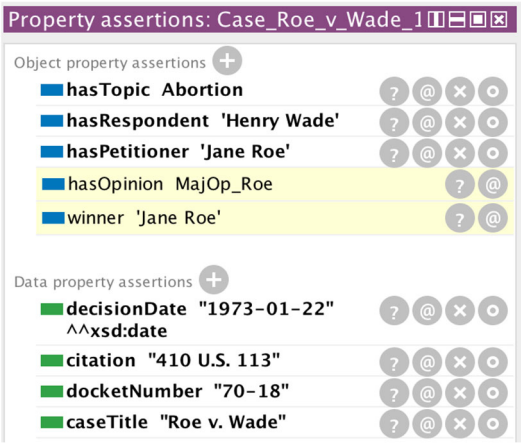

Figure 8. RESULT: Winner from MajorityOpinion for case Rose versus Wade.

The second example illustrates a situation when we might want to know:

Which opinions exist and who wrote them for court case ***Marbury v. Madison.***

It is easy to find out from Figures 9 and 10 that if we wish to retrieve facts from the ontology which would give the answer to this question, we should infer object properties 'participatesAsJudge' between correct individuals and use 'wroteOpinion' and 'opinionOf' as earlier defined object properties.

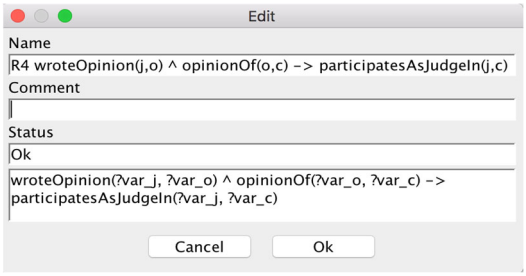

Figure 9. Inferring opinion authorship.

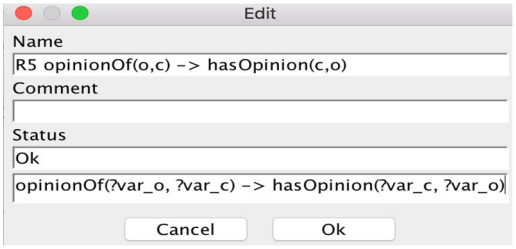

Figure 10. Inferring hasOpinion property.

Figure 11 shows the result for case ***Marbury v. Madison***: Judge John Marshall authored Majority Opinion.

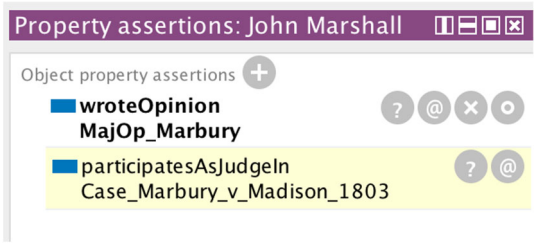

Figure 11. RESULT: 'John Marshal' has been asserted.

4. Ontology B: The Reasoning

In order to show the process of performing logic reasoning upon SWRL enabled OWL ontologies, we must define a competency question first, because the purpose of the process is to give an answer to it.

> *I am a barrister defending client Bob. He is accused of manslaughter of another person (XX), but he acted in defense of a third person who was attacked by XX. I can also prove that Bob acted with reasonable force. I need a precedent court case where manslaughter happened in defense of another person with reasonable force and the accused in that case was acquitted of manslaughter on the grounds of self-defense.*

Figure 12 below is a model of our reasoning process. The model might not be easy to follow, but it is divided into a set of reasoning steps. They use semantic overlapping between the individuals of participating classes and then move these individuals from one class to another. Therefore, if we place this reasoning process within a computational model,

it will not be a black box: the steps of the reasoning are well defined. However, the model in Figure 12 follows the SWT and the syntax of its languages and consequently should be read according to the SWT terminology. Details on SWRL enabled OWL ontologies and their modeling elements are available on the SWT webpages.

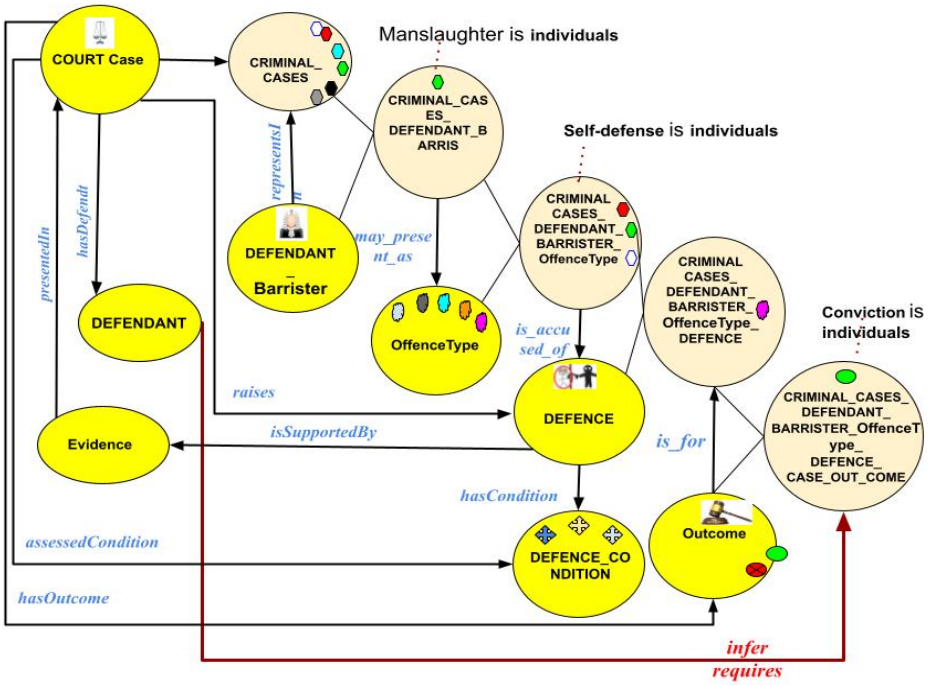

Figure 12. The reasoning process for Ontology B (colored symbols in ontological classes are individuals which can be moved and thus inferred).

Yellow colored OWL classes are sourced within Ontology A, built in the previous section, and they illustrate the semantic richness of Ontology B. Also, the selection of OWL classes in Figure 12 is dictated by the competency question above and therefore they do not need special explanation. Constraints between OWL classes in the format of object properties are labelled and self-explanatory, and some of them are also sourced in Ontology A from the previous section.

However, this diagram shows more classes than anticipated. This is because it has to accommodate the reasoning process. It shows where the reasoning, based on the semantic overlapping [23] between individuals of OWL classes, happens, and which individuals are moved from one class to another because of the semantic overlapping. For example, individuals of the DEFENDANT_BARRISTER class are matched with individuals of the CRIMINAL_CASES class, and only individuals which belong to specific criminal cases (in our example, 'manslaughter') are moved from the CRIMINAL-CASES class to CRIMINAL-CASES DEFENDANT—BARISTER class. This explains the different coloring of classes in the diagram. Light beige colors indicate that we have auxiliary classes which store moved individuals and will in turn, at the end of reasoning, contain its result and answer the competency question.

Individuals of the OFFENCE TYPE class ('self defense') will help further to filter all criminal cases which are 'manslaughter' but committed in 'self-defense'. Considering that all court cases have an outcome ('convicted'/'not guilty'), they influence the result of the reasoning and store all criminal cases suitable for planning DEFENCE (for DEFENDANT) by the BARRISTER.

The class which contains individuals moved from the COURT CASES after all reasoning steps (CRIMINAL_CASES_DEFENDANT_BARRISTER_OfenceTtpe_DEFENCE_CAS

E_OUTCOME) contains a set of court cases which overlap with Bob's case in this particular instance. Therefore, these court cases are specially selected for the DEFENDANT, and they can serve as a guideline for outlining his defense by the barrister. At this moment, the model from Figure 12 infers the object property named 'requires', between individuals of our DEFENDANT class and results of reasoning (final choice of precedent court cases).

It is important to note that the reasoning from Figure 12 is performed for a particular court case (DEFENDANT and BARRISTER) and the result of reasoning might not have long-term significance for knowledge stored in the ontology. We have a result which is created in a particular context and might not be of direct help to anyone else (except for Bob's defense). It is not necessary to keep the results of this reasoning within the ontology. If we need to get the same results of reasoning once again in future, we can run this very computationally efficient reasoning process in a fraction of a second. Running it again might bring benefits because these ontologies are constantly updated with new individuals and object properties. Therefore the results of reasoning for a particular DEFENDANT might change in future, with new ***court case precedents*** added, and thus it is worthwhile to run this reasoning again.

Figure 13 is a screenshot of the result of the reasoning performed by the process in Figure 12. The reasoning is run to answer the competency question, but we see that it finds a ***precedent court case*** for Bob's defense and consequently directly influences the defense. This is displayed in the protégé tool screen, and it shows that the defense for Bob will be supported by one old court case (an individual of the yellow-shaded ontological concepts in the lower part of the screen). The result of that old court case (precedent) has ***Acquital_Outcome*** and the court case shows that this acquittal outcome was based on 'self-defense' with 'reasonable force'. Consequently, the 'requires' object property is now inferred between Bob's current court case and the outcome of a very old court case (***precedent court case***) which is ***acquittal.***

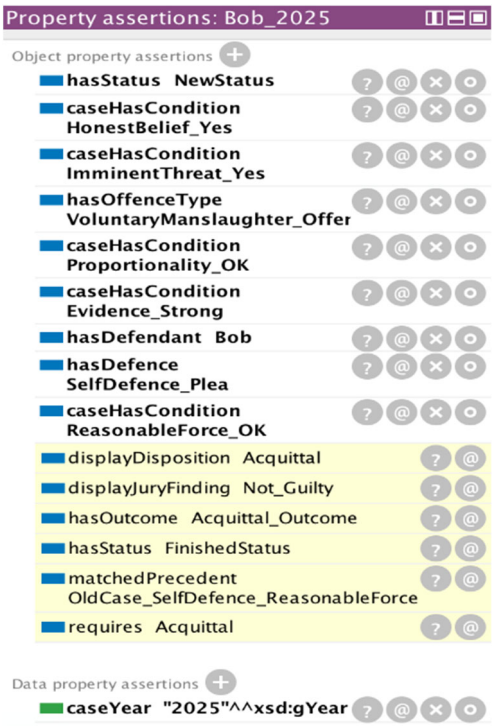

Figure 13. Preparing defense for Bob.

The upper part of the screen from Figure 13 is a full description of Bob's court case, which is formatted according to OWL ontological structures. It allows us to find the

semantic overlapping of Bob's case with an old, finished, ***precedent court case***, which has decisions that could shape Bob's defense.

If we wish to find the details of the ***precedent court case*** which is important for Bob's defense, its ID and title can be displayed through the software application which accommodates the reasoning process. These ***precedent court case*** details are moved as individuals in the reasoning process from the CRIMINAL_CASE class to the class where the result of the reasoning process (CRIMINAL_CASES_DEFENDANT_ BARRISTER_OfenceTtpe) is shown and available for inspection (see also Figure 12).

5. Evaluation and Debate

This paper is one of the rare publications which debate ontological design for creating *neurosymbolic* computing-*ready* SWRL enabled OWL ontology. It also highlights the multiple roles ontologies may play in current AI. The idea of revisiting traditional ontologies for neurosymbolic computing originated in the huge success of KG in AI, but this paper goes much further. It looks at possibilities of bi-directional relationships between logic structures in OWL and reasoning with SWRL, and extremely powerful NN enabled LLMs in modern AI.

The aim of this paper was also to raise awareness that OWL ontologies are not solely static vocabularies from which we can extract knowledge. They can be a part of a viable and software engineering friendly computational model considering the power of SWRL in logic reasoning [16,24,25].

Obviously for the same problem domain ***(court case precedents),*** we may have two different ontological designs, but they will be deployed in different circumstances and according to the need of either AI software application or neuromyotonic computational modeling in general. It is still important to distinguish the OWL static vocabulary in Figure 1 from the computational model with SWRL-enabled OWL ontologies in Figure 12. They are two different aspects of computing and can be mixed if needed, particularly in neurosymbolic computing.

The example of ***court case precedents*** used in this paper is deliberately simple and easily explainable. The aim was to use the most common legal case example for all readers across disciplines, including readers without any legal background. This is important for understanding the power of ontological design and the complexity of the reasoning process in Figure 12. We propose an extremely complex but powerful reasoning process which is deployable within any software application equipped with application programing interfaces and access to ontologies, symbolic and non-symbolic computing [19,26,27].

However, this is not the only solution we may have. It is an effective and working solution which could find precedent cases based on questions we ask, without any help from GenAI tools. However, it involves multiple steps and inferences of individuals and object properties 'as we go'. It is computationally efficient but requires heavy involvement of legal professionals and computation scientists to work together on the way of showing this complex legal semantics within the SWT world. Therefore it is unlikely that Figure 12 will lose its complexity in future, regardless of which basic ontological design we may have. It may only change the reasoning steps and the order in which they are carried out.

At the time of writing, our ***Ontology A*** has already become a learning ontology by allowing new entries into our static vocabulary from the US Supreme Court webpages. We used GenAI for populating the ontology.

Our next research paper deals with the reasoning process from Figure 12 and its division into simple SWRL rules, which can be directly used for creating a semantic prompting mechanism in GenAI. This research will evaluate the proposal from Figure 12

https://doi.org/10.3390/cmsf2026013015

from a different perspective and evaluate the ways we use the semantic overlapping of ontological elements in the legal problem domain.

Unfortunately, this paper has no related work because of the lack of interest in addressing the existence of neurosymbolic ready legal SWRL ontologies. Currently KGs are taking the burden of addressing deficiencies of GenAI, and problems in industry, including the issues from the legal domain, are resolved by creating business specific KGs which do not necessarily have to involve full-scale ontological reasoning and results of reasoning inferences. Static vocabulary is very often sufficient for KGs and can play important role in quickly fixing deficiencies of using GenAI in the legal domain. However, if we really wish to utilize the power of ontologies in neurosymbolic computing, then the ideas from this research could be a safe step forward.

Author Contributions: This is interdisciplinary research with computer scientists and a paralegal specialist. Their work was interwoven. Conceptualization, methodology, formal analysis and investigation have been done by all authors. These activities have been based on previous work/publications. E.A. and R.J. have been working together in the field of symbolic computing with reasoning and FOL since 2014, which is evident from the reference list. A.M.C., from the legal domain, joined in 2025. Ontological reasoning was performed by E.A. using Protege tools. Data curation is not applicable in our case, software engineers/computer scientists were in charge of data extraction, formatting and modelling. However, the computational data have been debated by A.M.C. and the choice of data which populated computational model was dictated by the US Supreme Court webpages. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable because the authors are three independent researchers.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: Data used in this research is publicly available on the US SUPREEME COURT webpages.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Booch, G.; Fabiano, F.; Horesh, L.; Kate, K.; Lenchner, J.; Linck, N.; Loreggia, A.; Murgesan, K.; Mattei, N.; Rossi, F. Thinking Fast and Slow in AI. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtual, 19–21 May 2021; pp. 15042–15046.
2. Sheth, A.P.; Roy, K.; Gaur, M. Neurosymbolic Artificial Intelligence (Why, What, and How). *IEEE Intell. Syst.* **2023**, *38*, 56–62. [CrossRef]
3. Kelly, M.; Barron, A. The Best of Both Worlds: Dual Systems of Reasoning in Animals and AI. *Cognition* **2022**, *225*, 105118. [CrossRef] [PubMed]
4. Krämer, W.; Kahneman, D. (2011): Thinking, Fast and Slow. *Stat. Pap.* **2014**, *55*, 915. [CrossRef]
5. Hogan, A.; Blomqvist, E.; Cochez, M.; d'Amato, C.; Melo, G.D.; Gutierrez, C.; Kirrane, S.; Gayo, J.E.L.; Navigli, R.; Neumaier, S. Knowledge graphs. *ACM Comput. Surv. (CSUR)* **2021**, *54*, 1–37. [CrossRef]
6. Buehler, M.J. Accelerating Scientific Discovery with Generative Knowledge Extraction, Graph-Based Representation, and Multimodal Intelligent Graph Reasoning. *Mach. Learn. Sci. Technol.* **2024**, *5*, 035083. [CrossRef]
7. World Wide Web Consortium. Semantic Web Technology and Standards. Available online: https://www.w3.org/2001/sw/wiki/Main_Page (accessed on 1 May 2025).
8. World Wide Web Consortium. Semantic Web Rule Language (SWRL). Available online: https://www.w3.org/submissions/SWRL/ (accessed on 5 January 2015).
9. World Wide Web Consortium. OWL Web Ontology Language Overview. Available online: https://www.w3.org/TR/owl-features/ (accessed on 2 January 2025).
10. Herron, D.; Jimenez-Ruiz, E.; Weyde, T. On the Benefits of OWL-Based Knowledge Graphs for Neural-Symbolic Systems. In Proceedings of the 17th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy 2023), Siena, Italy, 3–5 July 2023.

11. Juric, R. Could Semantic Web Technologies Create New Computational Models Outside Semantic Web. In Proceedings of the 21th International Conference on Society for Design and Process Science, (SDPS '16), Orlando, FL, USA, 4–6 December 2016.
12. Lewis, S. Precedent and the Rule of Law. *Oxf. J. Leg. Stud.* **2021**, *41*, 873–898. [CrossRef] [PubMed]
13. Somodevilla, M.; Ayala, D.; Pineda, I. An Overview on Ontology Learning Tasks. *Comput. Sist.* **2018**, *22*, 137–146. [CrossRef]
14. Palagin, O.; Kaverinsky, V.; Litvin, A.; Malakhov, K. OntoChatGPT Information System: Ontology-Driven Structured Prompts for ChatGPT Meta-Learning. *Int. J. Comput.* **2023**, *22*, 170–183.
15. Garcez, A.d.A.; Lamb, L.C. Neurosymbolic AI: The 3rd wave. *Artif. Intell. Rev.* **2023**, *56*, 12387–12406. [CrossRef]
16. Juric, R.; Almami, E.; Almami, I. Evaluation of GenAI in Biomedical Knowledge Discovery. In Proceedings of the to Appear in HICSS 2025, Waikoloa Village, HI, USA, 7–10 January 2025.
17. Sowa, J. Without Ontology, LLMs are clueless. In Proceedings of the Ontology Summit, Virtual, 8 March 2024.
18. Bennett, M.; Westerinen, A. Ontologies and Large Language Models: Related but Different. In Proceedings of the Ontology Summit 2024 Fall Series, Virtual, 4 October–8 November 2023. Available online: https://www.researchgate.net/publication/374445501_Kickoff_of_Ontology_Summit_2024_Fall_Series_on_Ontologies_and_Large_Language_Models_Related_but_Different (accessed on 1 September 2025).
19. Almami, E. A Reference Computational Model to Manage the Semantics of Learning Environments Using SWRL-Enabled OWL Ontologies. Ph.D. Thesis, University of Plymouth, Plymouth, UK, 2017.
20. Almami, E.; Juric, R.; Ahmed, M.Z. Software Applications Built Upon SWRL Enabled OWL Ontologies. In Proceedings of the 21th International Conference on Society for Design and Process Science, (SDPS '16), Orlando, FL, USA, 4–6 December 2016.
21. Washington University in St. Louis. The Supreme Court Database. Available online: http://scdb.wustl.edu/ (accessed on 1 September 2025).
22. Supreme Court of the United States. Supreme Court Case Documents. Available online: https://www.supremecourt.gov/case_documents.aspx (accessed on 1 September 2025).
23. Almami, E.; Juric, R.; Ibtesam, A. Semantic Overlapping of OWL Models in Biomedicine. In Proceedings of the SDPS 2018 Workshop on Accountability in AI, Bologna, Italy, 6–7 December 2018; pp. 128–133.
24. Idelfonso, R.; Vsevolodovna, M.; Monti, M. Enhancing Large Language Models through Neuro-Symbolic Integration and Ontological Reasoning. *arXiv* **2023**, arXiv:2504.07640.
25. Juric, R.; Almami, E.; Almami, I.; Ahmed, Z. Semantic Computational Models for Polypharmacology: Applications in Drug Repurposing. In Proceedings of the 56th Hawaii International Conference on System Sciences (HICSS), Maui, HI, USA, 3–6 January 2023.
26. Patadia, R.; Kataria, P.; Juric, R.; Kim, I.; Sayal, P. Conceptual Design of Semantic Software Applications. In Proceedings of the SDPS 2011 Conference, Jeju Island, Republic of Korea, 9–10 June 2011.
27. Tarabi, M.; Juric, R. Software Architectures for Smart Applications in the Management of Chronic Diseases: A Study of Reversibility of Diabetes Type 2. In Proceedings of the 51st Hawaii International Conference on System Sciences (HICSS 2018), Waikoloa Village, HI, USA, 3–8 January 2018.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.