When Legal Articles Resist Formalisation

Ludi van Leeuwen[1][0000−0003−3165−4376], Tadeusz Zbiegień[2][0000−0001−9052−6978], and Cor Steging[1][0000−0001−6887−1687]

[1] Bernoulli Institute of Mathematics, Computer Science and Artificial Intelligence, University of Groningen
[2] Department of Legal Theory, Jagiellonian University

Abstract. Formal representations of legal text have long been central to research in AI and Law. In this study, we examine the process of formalising legal articles through a co-design approach conducted with a legal expert. Using the ANGELIC framework, we explore the formalisation of Japanese Civil Code articles and analyse where and why some articles 'resist' unambiguous formalisation. Through this analysis, we identify four lessons learnt about the limits of formalizations. These concern the inherent complexity and dynamic nature of the law, the concept of model granularity, and interpretive ambiguity in legal text. We hope these lessons can inform future efforts to design, apply, and evaluate formalisation methods in the legal domain.

Keywords: Legal formalisation · AI and Law · Knowledge representation · Legal reasoning · Limitations of formal models · Symbolic AI.

1 Introduction

There is a long-standing tradition of modelling legal knowledge in formal representations in the AI & Law domain. Such approaches include rule-based logics, case-based structures, ontologies, and argumentation frameworks [28,2,32,30]. Such representations can be used to reason with, to support domain understanding, and to enable automatic decision-making. Recently, while machine learning approaches are dominant, there has been a surge in neuro-symbolic approaches that combine data-driven AI techniques with formal representations [22]. This paper reflects on the practice of formalising legal articles. Drawing on concrete examples, we highlight situations in which the process of formalisation encounters resistance, revealing limitations, ambiguities, and structural challenges inherent to legal text.

In this paper, we report on the co-design of formal representations of legal articles, conducted with a legal expert, aiming to model a diverse set of provisions from the Japanese Civil Code. While several articles could be represented successfully, others proved difficult or nearly impossible to formalise unambiguously under our initial assumptions. This paper distills the lessons learned from these challenges and highlights design aspects that require explicit attention to produce usable formal models.

2 Background

Formalising law can be defined in terms of translation [29]: it is a process of translating legal provisions and norms into a chosen formal representation and has a long and extensive tradition in legal scholarship [26,24]. At the same time, legal provisions are known for the number of problems they generate. It seems that in many cases it is the specificity of legal language itself that makes it a rewarding subject of research. This includes in particular works on the issues of vagueness [12,23], defeasibility [16,8], context sensitivity [20] or reasoning based on values [6]. These issues have been reflected in a number of modern works that attempt to harness these properties in computational form. Previous work in this domain concerned, among others, the defeasibility of legal reasoning [7], application of fuzzy [10] and temporal logic [14].

2.1 ANGELIC

In our study, we represent the legal articles as ANGELIC Domain Models (ADMs). The ANGELIC methodology, later updated to ANGELIC II, is an approach for representing and reasoning about legal domains [1,3]. It has been used to model a variety of domains, such as US Trade Secrets [5] and the European Convention on Human Rights [9,4]. The ADMs that model legal articles are based on Abstract Dialectical Frameworks (ADFs), a generalisation of argumentation frameworks capable of modelling argumentation [11,19]. In these models, legal articles are organised hierarchically: the root node represents the article's verdict, which is determined by a set of issues, each in turn defined by a set of factors. Factors are legally relevant fact patterns that may or may not apply to a given case [2].

In ANGELIC, factors are divided into two types. Baselevel factors form the leaves of the hierarchy and are assigned a truth value (True', False', or 'Unknown') for each case, which is referred to as factor ascription. Abstract factors capture intermediary legal concepts and are defined by other abstract or baselevel factors. The values of abstract factors, as well as issues and the verdict, are calculated using accept and reject conditions along with a default value, applying logical operators (AND, OR, NOT) evaluated under three-valued logic [25].

2.2 Japanese Civil Code

In this study, we focus on modelling a selection of articles from the Japanese Civil Code, a comprehensive body of law governing private legal relations in Japan, including contracts, property, family, and obligations. The Civil Code is highly structured, with articles generally short and narrowly scoped, making it potentially suitable for formalisation. Our work is inspired in part by previous work in the COLIEE competition [15,31], a legal competition that provides benchmarks for legal information extraction and entailment tasks using the Japanese Civil Code.

For this study, we model the legal articles directly as a verdict determined by a set of factors, without explicitly representing intermediate issues. This simplification reflects the small scope of most articles. Despite this, modelling even these apparently straightforward provisions revealed challenges related to interpretive ambiguities, contextual dependencies and the interplay of multiple legal concepts, highlighting the limits of formalisation even in well-defined domains.

3 Constructing ADMs of Legal Articles

3.1 Overview of the approach

Our approach was based on a co-design process in collaboration with a legal expert. Together, we constructed ADMs for a selection of articles from the Japanese Civil Code. The co-design process involved iterative interpretation, where we discussed how each article could be represented in the ADM structure and how its legal concepts could be formalised as factors and logical relations. The goal was not necessarily to produce complete models but more so to examine how the process of formalisation unfolds when technical and legal perspectives interact. Through this collaboration, we explored both the strengths and the practical limits of formalising legal text.

3.2 Effective formalisations

Our approach proved to be effective for a number of articles. To illustrate a successful example, Figure 1 presents the ADM corresponding to Article 698 (see Figure 1a for the text of the article). The verdict, shown in yellow in Figure 1b, concerns whether the manager is liable to compensate for damages. This outcome is determined through the ascription of the base-level factors (blue), which in turn define the values of the abstract factors (green). Together, these factors, and the accept and reject conditions, yield the final value of the verdict.

We found that ADMs can be constructed most successfully for articles with a clear and logical structure that corresponds to the verdict. The methodology in these cases can be considered *effective* in the sense that it enabled the identification of the article's relevant characteristics and factors, yielding an ADM that can determine the verdict concerning it.

We have however encountered a number of informative edge cases. Those cases allowed us to draw several lessons concerning not so much the specific formalisation method used, but more broadly the general problem of formalising legal articles as such. We dive deeper into some of these edge cases in the next sections.

> If a manager engages in benevolent intervention in another's business in order to allow a principal to escape imminent danger to the principal's person, reputation, or property, the manager is not liable to compensate for damage resulting from this unless the manager has acted in bad faith or with gross negligence.

(a) Article 698 of the Japanese Civil Code.

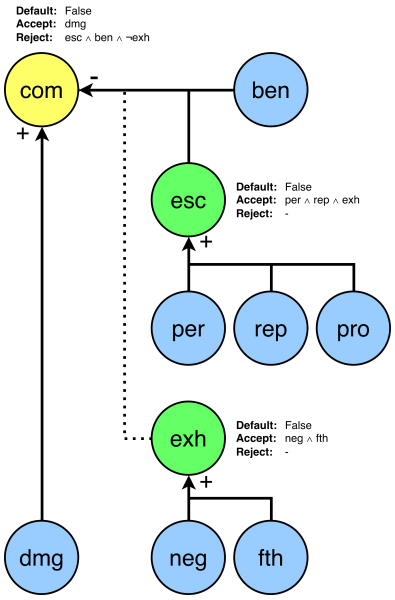

(b) ADM of Article 698

Verdict:

com The manager is liable to compensate for damages

Abstract factors:

esc The manager's actions aimed to allow the principal to escape imminent danger

exh The manager exhibited gross negligence or acted in bad faith

Baselevel factors:

ben The manager's intervention was benevolent

dmg There were damages resulting from the intervention

per The actions aimed to allow the escape of imminent danger to the person

rep The actions aimed to allow the escape of imminent danger to reputation

pro The actions aimed to allow the escape of imminent danger to property

neg The manager exhibited gross negligence

fth The manager acted in bad faith

(c) Verdict and factors of the ADM

Fig. 1: Article 698 of the Japanese Civil Code (a) alongside its ANGELIC Domain Model (b) and its verdict and factors (c).

3.3 Determining the verdict

In the case of Article 698 (Figure 1a), we determined that the *verdict* should be whether the manager is liable to compensate for damages. This answers the question, '*should the manager compensate for the damages?*'. Although this appears to be the most straightforward question to ask, a legal article can invite

several different questions. In such cases, it becomes more difficult to determine what the verdict of the ADM should represent.

We illustrate this concept using Article 593 of the Japanese Civil Code, shown in Figure 2a. Through a co-design session with the legal expert, we developed the ADM shown in Figure 2, where the verdict was defined as whether the loan for use is effective. This formulation allows reasoning about cases where the central question is: '*is the loan for use effective*?'.

However, other questions can be asked about this article, in which the constructed ADM cannot be applied. For example, two bar exam questions concerning Article 593 are shown in Figure 3, where examinees must determine whether a legal entailment exists between Article 593 and the statements in Figure 3. This task corresponds to Task 4 of the COLIEE competition [15].

When it comes to the first statement (H30-24-U), the implied question is: '*May B demand compensation for damages due to default based on the loan for use contract from A?*'. In the second statement statement (R1-23-3), the implied question is: '*Is a loan for use contract effective when not made in writing?*'. Strictly speaking, neither of these two questions can be answered directly ascribing factors to the ADM in Figure 2, as the questions do not match the question we had in mind for the verdict node when designing the ADM. To answer these questions, we would need a different ADM, tailored to those specific questions.

3.4 Complexity, dynamics and temporal effects

The ADM of Article 593 (Figure 2) also illustrates issues with regards to the symbolic structure and the affordances needed to represent legal articles. For example, we see that the modal verb *promise* is used. Additionally, the article contains reference to temporal aspects, such as a thing being returned *after* a contract is terminated, and the contract being terminated *after* the other party uses and makes profit of the thing. Furthermore, one needs to perform entity disambiguation to resolve references to *the borrowed thing*, *party* and *other party*. Hence, in this example alone, a formalism needs to account for logical connectives as well as modal, temporal and entity-level distinctions.

While ADMs are designed to correctly represent the structure of the logical connectors, they are in principle not designed to represent these other distinctions explicitly using deontic or temporal logic. Choosing a formalism hence means picking a (sub)set of relevant attributes of legal reasoning to represent, and leaving the rest implicit.

It is also the case that we often cannot reasonably answer legal questions only based on a single article. Instead, we need to read the article in the context of bigger network of interrelated articles, and consider *law as a complex and dynamic system* [21,17,18]. For example, to answer questions in Figure 3, the ADM in Figure 2 alone is insufficient and we would require additional articles to perform the legal reasoning necessary. Strictly speaking, in order for our formalisation to be able to answer legal questions adequately, it would have to take into account the fact that it is rare for a given legal question to be answered on the basis of a single provision, usually, a number of different regulations must be

A loan for use becomes effective if one of the parties promises to deliver a certain thing, and the other party promises to return the thing when the contract is terminated after the other party gratuitously uses and makes profit of the borrowed thing.

(a) Article 593 of the Japanese Civil Code.

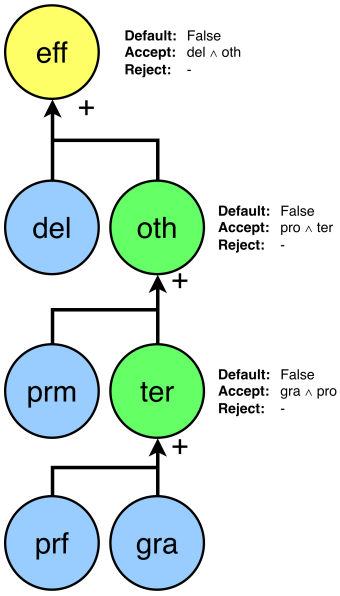

(b) ADM of Article 593

Verdict:

eff The loan for use is effective

Abstract factors:

oth The other party promises to return the thing when the contract is terminated after [the other party] gratuitously uses and makes profit of the borrowed thing.

ter The contract is terminated after the other party gratuitously uses and makes a profit from the borrowed thing.

Baselevel factors:

del One of the parties promised to deliver a certain thing

prm The other party promises to return the borrowed thing when the contract is terminated

prf The other party makes a profit from the borrowed thing

gra The other party gratuitously uses the borrowed thing

(c) Verdict and factors of the ADM

Fig. 2: Article 593 of the Japanese Civil Code (a) along with its ANGELIC Domain Model (b) and its verdict and factors (c).

taken into account, which combine into different structures depending on how the question is asked.

3.5 Model granularity

In the ADM of Article 593, the verdict is determined by two factors, one of which in turn are determined by sub-factors, and one of those by sub-sub-factors as shown in Figure 2. This raises a broader methodological question about the appropriate level of detail when decomposing factors in ADMs. For example, should factor *oth* be further broken down into its constituent sub-factors, or is it sufficient to treat it as a single baselevel factor? This leads to the open issue of

H30-24-U: A who is the owner of a building and B made a loan agreement where B borrows the building for free only while A is working abroad, and then, however, A made a loan contract with third party C, and transferred the building. In such case, B may not demand compensation for damages due to default based on the loan for use contract from A.

R1-23-3: No contract of loan for use shall be effective unless it is made in writing.

Fig. 3: Two bar exam questions concerning Article 593.

model granularity. On the one hand, there is a natural inclination to subdivide the model into as many explicit conditions as possible. For example, in case law when determining whether a vague term such as 'gratuitous usage' applies. On the other hand, excessive granularity risks undermining clarity: the model may become unreadable, and each additional node, particularly when not directly supported by the text of the provision, and increases the risk of falling into the trap of explicitly encoding certain modelling choices.

3.6 Model uncertainty and scope ambiguity

In terms of ADMs, decision trees or graphs in general, *model uncertainty* in a legal context will refer to the uncertainty arising from the possibility of creating multiple (equally or unequally) justifiable edges with various logical operators assigned between the nodes representing the conditions. Simply put, it refers to a situation where it is possible to construct similarly justified models of the same article or legal norm resulting from it, which differ in their logical structure. Such model uncertainty can be caused by *scope ambiguity*, where logical operators like 'and' or 'or' are used without clear specification of their range, allowing multiple plausible ways to structure the logical relationships.

To illustrate this concept, we show an artificial article based on Article 178a of the Polish Criminal Code in Figure 4a [3]. In this article, a scope ambiguity is evident due to the use of or' and and' without clarity on the range of these connectors. This ambiguity permits at least two different ADMs, each representing one possible reading. Both readings are shown in Figures 4c and 4d. While seemingly superficial, this example illustrates model uncertainty due to scope ambiguity clearly, and reflects the kinds of interpretative disagreements that frequently arise in legal analysis. Importantly, this is not a classic problem of linguistic ambiguity and the scope of meaning of individual concepts. The uncertainty does not arise at the level of the meaning of individual concepts or

[3] This example comes from a webinar held by the Government Legislation Centre (RCL) on 21 August 2025[27]

```
Anyone who drives a vehicle
under the influence of alcohol
or intoxicating substances and
poses a threat is subject to
punishment.
```

(a) Artificial article based on Article 178a of the Polish Criminal Code.

Verdict:

pun Subject to punishment?

Baselevel factors:

veh Drives a vehile

alc Under influence of alcohol

tox Under influence of intoxicating substances

thr Poses a threat

(b) Verdict and factors of the ADMs

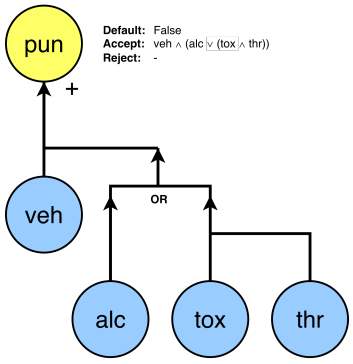

(c) First ADM interpretation

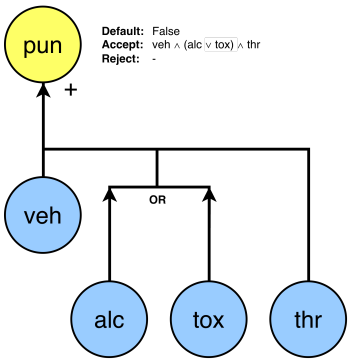

(d) Second ADM interpretation

Fig. 4: Artificial article based on Article 178a of the Polish Criminal Code (A) alongside two different ANGELIC Domain Models the (c & d) and its verdict and factors (b).

the scope of their meaning, but at the *level of the relationships between these concepts.* Because of this, formalising articles can also help expose such scope ambiguity, which can be useful for legal scholars and personnel.

3.7 Artificially generated formalisations

Since the advent of Large Language Models (LLMs), research has investigated the possibility of automatically extracting formal representations from legal texts [33]. In a previous study, we found that automatically generated ADMs were often ineffective [31]. Examining these artificial formalisations does further highlight the observations we have made during our co-design process with a human legal expert.

For example, when prompting ChatGPT to formalise an ADM of the article in Figure 4a, the LLM shows a persistent preference (10 out of 10 runs) for the second reading (see Figure 4d). Even with additional prompting, ChatGPT does not generate the first reading as shown in Figure 4c. We also see that the LLM takes a less granular approach, as it collapses the 'or' statement into one baselevel

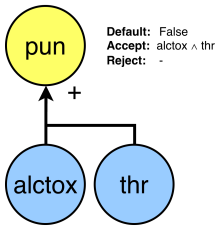

(a) GPT-generated ADM

Verdict:

pun Is the driver subject to punishment?

Baselevel factors:

alctox Is the driver under the influence of alcohol or intoxicating substances?

thr Does the driver pose a threat to others or public safety?

(b) Verdict and factors of the ADM

Fig. 5: The ANGELIC Domain Models of the article from Figure 4a as generated by ChatGPT (a) and its generated verdict and factors (b).

factor The structure of the generated ADMs is consistent across runs, with only minor paraphrasing of the text within factors, as would be expected from a language model. The consistent preference of ChatGPT for the second reading implies that it does not expose the scope ambiguity of this legal article, which can be harmful if the formalisation is used to reason with by legal professionals.

4 Lessons and Design Implications

In our analysis, we used the process of formalisation through ADMs as a prototypical example to illustrate broader challenges in representing legal articles formally. In many cases, the ADM structure enables the construction of models that can be used to reason about legal questions in a structured and interpretable manner. However, several fundamental challenges remain. From our observations, we distil several key lessons and corresponding design implications for future work

4.1 Lesson 1: No single formalisation fits all legal questions

As discussed in Section 3.3, no single ADM or formalisation can capture all possible questions that may arise from a given article. Although the content of legal articles often appears to imply a specific question they are meant to address, the reality is more complex. In practice, the same article may give rise to multiple legitimate questions, each requiring a distinct formalisation. The same article can also often also be a building block to answer other questions relating to other articles.

4.2 Lesson 2: Law is complex, temporal and dynamic

In Section 3.4, we showed that although legal provisions are written in texts that at first appear clearly structured and stable, in reality, the law, the interpretation of articles, the reconstruction of norms, and the arguments surrounding them are highly susceptible to change and instability. Following a long-standing tradition

in legal theory, we distinguish between an *article* (or *legal provision*), understood as an editorial unit of legal text, and the corresponding *legal norm*, which describes the structure of an obligation or duty and is reconstructed through interpretation from one or more provisions. Law is therefore a complex and dynamic system. Despite a long tradition of formalisation in law, the actual, effective modelling of these two facts remains a crucial challenge, and future work is needed that focuses on these temporal and referential aspects of legal norms.

4.3 Lesson 3: Granularity is an explicit design choice

As shown in Section 3.5, the level of detail, or granularity, of a model is an explicit design choice with significant implications for its structure and effectiveness. For a formalisation to be effective, however understood, the granularity of the model must therefore be taken into account. There is always a trade-off between descriptive accuracy and interpretability. In striving for a more precise representation, we may lose readability and, consequently, practical applicability.

The appropriate level of granularity depends on the purpose of the formalisation. Is the goal to construct an algorithm capable of automating a process defined by legal norms? Is it to study the nature of legal norms or to model legal reasoning itself? Or is the aim to optimise performance on a computational benchmark? In some cases, the purpose may even be exploratory or conceptual. These questions should precede any formalisation effort, as the intended function of the model directly determines the depth and structure of its design.

4.4 Lesson 4: A single article can yield multiple valid models

Model uncertainty arises when a single legal article can be represented in multiple equally justified ways. One source of this uncertainty, as shown in Section 3.6, is scope ambiguity, where connectors such as and' or or' can be interpreted in different ways, producing distinct ADMs.

Recognising that legal norms are derived from many different sources, that their meaning may evolve over time, and that the linguistic boundaries of legal terms are rarely precise, is only a starting point. It should be noted that all these building blocks can be combined in various ways, none of which should be considered "the one and only correct" (especially before a decision is made to settle this type of legal dispute in a given case, and even then we only obtain certainty *ex post* in that one case – the decision-making pattern does not necessarily have to be repeated in subsequent similar or even identical cases, even if there are a number of legal and procedural safeguards in place to ensure this). It is precisely this inherent indeterminacy that we refer to as *model uncertainty*, and it is one of the features that makes law an especially challenging and intellectually rich domain. In this context, research addressing uncertainty directly seems to be warranted. Methods that avoid creating a false (and, in the context of law, dangerous) sense of stability and instead capture the uncertainty intrinsic to legal reasoning are essential. In particular, research on uncertainty quantification in legal models appears to be a promising direction.

5 Discussion and conclusion

While this study focused primarily on hard cases, we do not want to suggest that law consists only of such difficult cases subject to varying levels and types of uncertainty. A legal system in which every interpretation required complex logical or interpretative operations would be difficult to imagine. In practice, many articles and cases are relatively straightforward [13], which enables formalisation and supports the development of e-government services and the automation of administrative procedures in many countries. However, it is the edge cases that provide the most insight into both the formalisation process and the law itself. While we used the ANGELIC methodology in this study, our aim is not to evaluate the specific limitations or capabilities of this framework, but rather to examine the properties of law itself as an object of study. Different formalisation methods might highlight other issues, but our focus is on the structure of norms and the relationships between them, rather than on the fulfilment of specific conditions, which has already received substantial attention.

In this study, we analysed the process of formalising legal provisions through a co-design approach with a legal expert. Our experience demonstrated that, while many articles can be effectively formalised, certain provisions resist formalisation due to inherent characteristics of law. From these cases, we derived four key lessons: there is not a single formalisation suitable for every question; law is complex, temporal, and dynamic; model granularity is an explicit design choice; and multiple representations of the same article can exist, giving rise to model uncertainty. We argue that these lessons are broadly applicable and should inform future research on legal formalisation, guiding both the development of formal methods and the interpretation of law as a structured domain.

References

1. Al-Abdulkarim, L., Atkinson, K., Bench-Capon, T.: A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. Artif. Intell. Law **24**(1), 1–49 (Mar 2016). https://doi.org/10.1007/s10506-016-9178-1
2. Aleven, V.: Teaching Case-based Argumentation through a Model and Examples. Ph.D. thesis, University of Pittsburgh (1997)
3. Atkinson, K., Bench-Capon, T.: Angelic ii: An improved methodology for representing legal domain knowledge. In: Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law. pp. 12–21. ICAIL '23, Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3594536.3595137
4. Atkinson, K., Collenette, J., Bench-Capon, T., Dzehtsiarou, K.: Practical tools from formal models: the echr as a case study. In: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law. pp. 170–174. ICAIL '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3462757.3466095
5. Bench-Capon, T., Gordon, T.F.: Implementing a theory of a legal domain. In: Legal Knowledge and Information Systems, pp. 13–22. IOS Press (2022)

6. Bench-Capon, T., Modgil, S.: Norms and value based reasoning: Justifying compliance and violation. Artificial Intelligence and Law **25**, 29–64 (2017). https://doi.org/10.1007/s10506-017-9194-9
7. Billi, M., Calegari, R., Contissa, G., Lagioia, F., Pisano, G., Sartor, G., Sartor, G.: Argumentation and defeasible reasoning in the law. J **4**(4), 897–914 (2021). https://doi.org/10.3390/j4040061
8. Brożek, B.: Law and defeasibility. Revus (23) (2014). https://doi.org/10.4000/revus.3110
9. Collenette, J., Atkinson, K., Bench-Capon, T.: Explainable AI tools for legal reasoning about cases: A study on the European Court of Human Rights. Artificial Intelligence **317**, 103861 (2023)
10. da Costa Pereira, C., Tettamanzi, A.G.B., Liao, B., Malerba, A., Rotolo, A., van der Torre, L.: Combining fuzzy logic and formal argumentation for legal interpretation. In: Proceedings of the 16th Edition of the International Conference on Articial Intelligence and Law. p. 49–58. ICAIL '17, Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/10.1145/3086512.3086532
11. Dung, P.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2), 321–357 (1995)
12. Endicott, T.A.O.: Vagueness and legal theory. Legal Theory **3**(1), 37–63 (1997). https://doi.org/10.1017/S135232520000063X
13. Fischman, J.B.: How many cases are easy? Journal of Legal Analysis **13**(1), 595–656 (2021). https://doi.org/10.1093/jla/laaa010
14. Ghari, M.: A formalization of the protagoras court paradox in a temporal logic of epistemic and normative reasons. Artificial Intelligence and Law **32**(2), 325–367 (2024). https://doi.org/10.1007/s10506-023-09351-0
15. Goebel, R., Kano, Y., Kim, M., Rabelo, J., Satoh, K., Yoshioka, M.: Overview of Benchmark Datasets and Methods for the Legal Information Extraction/Entailment Competition (COLIEE) 2024. In: Suzumura, T., Bono, M. (eds.) New Frontiers in Artificial Intelligence. pp. 109–124. Springer Nature Singapore, Singapore (2024)
16. Governatori, G.: An ASP implementation of defeasible deontic logic. Künstliche Intelligenz **38**, 79–88 (2024). https://doi.org/10.1007/s13218-024-00854-9
17. Hage, J., Verheij, B.: The law as a dynamic interconnected system of states of affairs: a legal top ontology. International Journal of Human-Computer Studies **51**(6), 1043–1077 (1999). https://doi.org/10.1006/ijhc.1999.0297
18. Jones, G.T.: Dynamical jurisprudence: Law as a complex system. Georgia State University Law Review **24**(4) (2012)
19. Keshavarzi Zafarghandi, A., Verbrugge, R., Verheij, B.: Discussion games for preferred semantics of abstract dialectical frameworks. In: Kern-Isberner, G., Ognjanović, Z. (eds.) Symbolic and Quantitative Approaches to Reasoning with Uncertainty. vol. 11726, pp. 62–73. Springer (2019)
20. Kompa, N.: The role of vagueness and context sensitivity in legal interpretation. In: Keil, G., Poscher, R. (eds.) Vagueness and Law: Philosophical and Legal Perspectives. Oxford University Press (2016). https://doi.org/10.1093/acprof:oso/9780198782889.003.0010
21. Moreso, J.J.: Legal dynamics. In: Legal Indeterminacy and Constitutional Interpretation, Law and Philosophy Library, vol. 37, pp. 101–116. Springer, Dordrecht (1998). https://doi.org/10.1007/978-94-015-9123-2_4

22. Mumford, J., Atkinson, K., Bench-Capon, T.: Reasoning with Legal Cases: A Hybrid ADF-ML Approach (12 2022). https://doi.org/10.3233/FAIA220452
23. Poscher, R.: Ambiguity and vagueness in legal interpretation. In: Solan, L.M., Tiersma, P.M. (eds.) The Oxford Handbook of Language and Law. Oxford University Press (2012). https://doi.org/10.1093/oxfordhb/9780199572120.013.0010
24. Prakken, H., Sartor, G.: Law and logic: A review from an argumentation perspective. Artificial Intelligence **227**, 214–245 (2015). https://doi.org/10.1016/j.artint.2015.06.005
25. Priest, G.: Many-valued Logics, p. 456–475. Cambridge Introductions to Philosophy, Cambridge University Press (2008)
26. Royakkers, L.M.M.: Extending Deontic Logic for the Formalisation of Legal Rules, Law and Philosophy Library, vol. 36. Springer Science+Business Media Dordrecht, Dordrecht, 1 edn. (1998). https://doi.org/10.1007/978-94-015-9099-0
27. Rządowe Centrum Legislacji: Rola funktorów prawdziwościowych w prawidłowym konstruowaniu przepisów [the role of truth-functional operators in the proper drafting of legal provisions]. Webinarium [Webinar] (Aug 2025), accessed: 7 November 2025
28. Satoh, K., Asai, K., Kogawa, T., Kubota, M., Nakamura, M., Nishigai, Y., Shirakawa, K., Takano, C.: Proleg: An implementation of the presupposed ultimate fact theory of japanese civil code by prolog technology. In: Onada, T., Bekki, D., McCready, E. (eds.) New Frontiers in Artificial Intelligence. pp. 153–164. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
29. Schafer, B.: Formalising law, or the return of the golem. In: Research Handbook on Law and Technology, chap. 5. Edward Elgar Publishing, Cheltenham, UK (2023). https://doi.org/10.4337/9781803921327.00012, retrieved November 4, 2025
30. Sergot, M.J., Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P., Cory, H.T.: The British Nationality Act as a logic program. Communications of the ACM **29**(5), 370–386 (May 1986). https://doi.org/10.1145/5689.5920
31. Steging, C., van Leeuwen, L.: A hybrid approach to legal textual entailment. In: Eighteenth International Workshop on Juris-Informatics (JURISIN 2024). Hamamatsu, Japan (5 2024)
32. Verheij, B.: Formalizing arguments, rules and cases. In: ICAIL '17: Proceedings of the Sixteenth International Conference for Artificial Intelligence and Law. pp. 199–208. ACM, New York, London, United Kingdom (6 2017)
33. Zin, M.M., Satoh, K., Borges, G.: Leveraging llm for identification and extraction of normative statements. In: Legal Knowledge and Information Systems, pp. 215–225. IOS Press (2024)