

# From Legal Texts to Defeasible Deontic Logic via LLMs: A Study in Automated Semantic Analysis

Elias Horner<sup>1</sup>, Cristinel Mateis<sup>2</sup>, Guido Governatori<sup>3,4</sup> and Agata Ciabattoni<sup>1</sup>

<sup>1</sup>TU Wien, Vienna, Austria

<sup>2</sup>AIT Austrian Institute of Technology, Vienna, Austria

<sup>3</sup>School of Engineering and Technology, Central Queensland University, Rockhampton, QLD, Australia

<sup>4</sup>Artificial Intelligence and Cyber Futures Institute, Charles Sturt University, Bathurst, NSW, Australia

## Abstract

We present a novel approach to the automated semantic analysis of legal texts using large language models (LLMs), targeting their transformation into formal representations in Defeasible Deontic Logic (DDL). We propose a structured pipeline that segments complex normative language into atomic snippets, extracts deontic rules, and evaluates them for syntactic and semantic coherence. Our methodology is evaluated across various LLM configurations, including prompt engineering strategies, fine-tuned models, and multi-stage pipelines, focusing on legal norms from the Australian Telecommunications Consumer Protections Code. Empirical results demonstrate promising alignment between machine-generated and expert-crafted formalizations, showing that LLMs – particularly when prompted effectively – can significantly contribute to scalable legal informatics.

## Keywords

legal informatics, large language models, defeasible deontic logic, semantic formalization, prompt engineering, legal NLP

## 1. Introduction

The idea of automated legal reasoning has been one of the cornerstones of AI and Law for a long time, with many concrete attempts since the seminal paper by Sergot and Kowalski [1] on the formal representation of the British Nationality Act (see [2, 3] for an overview of some of the most influential approaches). A recent OECD report [4] outlines the benefits of the adoption of automated legal reasoning and encoding legal provisions in a format that is processable by machines. The major obstacle to this vision is the knowledge representation bottleneck. Anecdotal data from many large-scale encoding projects suggests that an experienced coder can only encode 4 to 5 pages per day, with serious burnout concerns (a recent empirical experiment on legal coding confirms the rate of encoding [5]). The issue is further exacerbated by the proliferation of legal information and the rising complexity of regulatory environments, which have intensified the need for automated tools that can interpret and formalize normative documents. Accordingly, there is a need to have tools that can help with the encoding of legal instruments.

The idea of using NLP techniques, specifically categorial grammar-based approaches, to encode norms was advanced by [6]. Then it has been extended in [7] where it was tested on small-scale examples. Successively, it was adopted in [8] that introduced a manually supervised pipeline for extracting a formal representation, relying on deterministic parsing rules. While this approach resulted in reasonable outcomes, a successful extraction required many iterations and was sensitive to the specific format of the input. Furthermore, it did not lead to a significant reduction in the time needed to create a complete and fully functional encoding. On the other hand, [9] explored the use of ML-based NLP techniques for the normative encoding process. The key findings were that these approaches required

---

*Proceedings of the Seventh International Workshop on Automated Semantic Analysis of Information in Legal Text (ASAIL 2025), 16 June, 2025, Chicago, USA.*

✉ elias.horner@tuwien.ac.at (E. Horner); cristinel.mateis@ait.ac.at (C. Mateis); g.governatori@cqu.edu.au (G. Governatori); agata@logic.at (A. Ciabattoni)

🆔 0009-0008-1080-2329 (E. Horner); 0000-0001-7502-0688 (C. Mateis); 0000-0002-9878-2762 (G. Governatori); 0000-0001-6947-8772 (A. Ciabattoni)



© 2025 Copyright for this paper by its authors. Published in the Proceedings of the Seventh International Workshop on Automated Semantic Analysis of Information in Legal Text (ASAIL 2025), 16 June 2025, Chicago, USA. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

very extensive training data (which was and still is not available), and that the performance was not comparable to the rule-based approach. More recent efforts have incorporated neural methods for legal information retrieval and summarization, but few have addressed the formalization task with the granularity and formal logic representation at the level we propose in this paper. For instance, the recent work [10] focuses on a single article from the Council Framework Decision 2002/584/JHA (European Arrest Warrant).

In recent years, Large Language Models (LLMs) have emerged as powerful tools for understanding and generating natural language. However, their application to legal texts remains underexplored, particularly in tasks requiring semantic precision, such as the conversion of legal norms into machine-interpretable representations.

This paper explores the feasibility and effectiveness of using LLMs to translate legal language into a formal representation. More specifically, we compare with the approach of [8], and we encode legal provisions in Defeasible Deontic Logic (DDL), a computational logic framework designed to reason about obligations, permissions, and prohibitions. The target corpus is the Australian Telecommunications Consumer Protections Code (TCP Code), characterized by complex, hierarchical rules. This dataset was also used in [8], and we compare our results with those reported there.

Our central hypothesis is that, with suitable prompting and architectural configurations, LLMs can assist in extracting semantically valid and logically coherent deontic rules from unstructured legal text. The novel contribution of this work lies in the integration of prompt engineering techniques, evaluation metrics grounded in logical correctness, and comparative studies of different LLM architectures and training strategies.

The remainder of this paper is structured as follows. Section 2 provides the necessary background on LLMs and the formal representation language DDL used in this study. Section 3 outlines the methodology, including the segmentation of legal texts into individual law snippets, their transformation into DDL, and the evaluation approach. Section 4 presents the experimental results, covering prompt engineering, multi-snippet processing strategies, fine-tuning, and two-stage pipelines. This section also includes a detailed comparison with the evaluation framework proposed by Dragoni et al. [8]. Section 5 discusses key limitations, such as challenges in legal implementation, handling of inter-snippet references, and atom reuse. Finally, Section 6 concludes the paper and outlines directions for future research.

## 2. Background

**Defeasible Deontic Logic** [11, 12] is a flexible and efficient rule-based non-monotonic formalism for the representation of legal norms and legal reasoning. The logic combines features of Defeasible Logic for the natural modeling of exceptions and defeasibility with concepts from Deontic Logic (i.e., obligations, permission, prohibition, compensatory obligations). A rule in DDL has the form

$$r: a_1, \dots, a_n \Rightarrow c$$

where  $r$  is the label (or name of the rule),  $a_1, \dots, a_n$  are the premises of the rule, and  $c$  is the conclusion of the rule.  $a_1, \dots, a_n, c$  are either literals or deontic literals, where a literal is either an atomic proposition or its negation, and a deontic literal is a literal in the scope of a deontic operator ([O] for obligation, [F] for forbidden or prohibition, and [P] for permission). Moreover, the logic is equipped with a superiority relation, a binary relation over the set of rules. The superiority relation is used when two conflicting rules are both applicable, and specifies which rule prevails over the other.

The DDL reasoning mechanism has an argumentation-like structure. To prove a conclusion, we need to have an applicable rule for it. Then we have to consider all possible counterarguments, namely the rules for the opposite. For each of such rules, we have to rebut them. Thus, we have to either discard it (show that the rule is not applicable) or defeat it, which means we have to show an applicable rule that defeats it (using the superiority relation).

**Large Language Models (LLMs)** are advanced machine learning systems designed to understand and generate human language. Trained on vast amounts of textual data, LLMs are capable of performing

a wide range of language-related tasks, including text generation, summarization, translation, question answering, and formal reasoning. They achieve this by learning complex statistical patterns and representations of language, enabling them to predict the most likely continuation of a given input.

LLMs can be broadly divided into two categories: *traditional LLMs* and *reasoning LLMs*. Traditional LLMs, such as the GPT series by OpenAI, DeepSeek-V3 or similar models, are primarily optimized for fluent language generation and general-purpose tasks. Their strength lies in producing coherent and contextually appropriate text. However, their capabilities in logical reasoning and structured problem-solving are limited. Reasoning LLMs represent a newer generation of models that are explicitly designed to perform structured reasoning tasks more effectively. These models consider additional training objectives, architectural innovations, or fine-tuning procedures that enhance their ability to perform logical inference, complex decision-making, and consistent multi-step problem-solving. Reasoning LLMs aim not only at linguistic fluency but also at improved logical accuracy and reliability in formal contexts.

Recent advancements in LLMs offer opportunities to automatically extract formal semantics like DDL from legal documents. However, achieving logical coherence and semantic validity remains non-trivial, motivating the need for careful experimental design.

In this study, we consider models from both categories: (i) Traditional LLMs: GPT-4o, GPT-4o mini, DeepSeek-V3, and (ii) Reasoning LLMs: OpenAI o3, OpenAI o1, OpenAI o4-mini, OpenAI o3-mini, DeepSeek-R1. These models will be evaluated and compared based on their performance in formalizing legal texts into DDL representations.

### 3. Methodology

To reduce hallucinations and promote deterministic behavior, all LLMs were assessed under conservative decoding settings:

- `temperature = 0`  
Controls randomness; higher values yield more diverse outputs [13]. Set to 0 to prioritize consistency over creativity [14].
- `top-p = 1`  
Governs nucleus sampling [15]; left at the default to avoid compounding effects with temperature, per API guidelines [16, 17].
- `frequency penalty = 0`  
Penalizes repeated tokens by frequency [18]. Disabled to ensure consistent terminology use.
- `presence penalty = 0`  
Penalizes repeated tokens after first occurrence [18]. Also disabled for terminological consistency.

These settings were applied both during legal text segmentation and formalization into DDL. For OpenAI's reasoning models, which do not accept these parameters, the `reasoning_effort` option was set to high.

We evaluated three main strategies:

- **Chain-of-Instructions (CoI)** prompting with varying configurations and shot counts.
- **Fine-tuning of GPT-4o** using a limited dataset of annotated examples to enhance task-specific performance.
- **Two-stage pipelines**, using separate LLMs for atom extraction and rule generation to enhance consistency and limit error propagation

The rest of this section presents the core aspects that form the basis of our approach.

### 3.1. Segmentation into Law Snippets

Legal texts are initially segmented into manageable "law snippets" using DeepSeek R1. Enumerations in legal provisions are split into individual rules where appropriate, aiming at a balance between contextual completeness and token constraints.

A key challenge in instructing the LLMs was determining the optimal length for law snippets: overly long snippets risked losing critical information during formalization, whereas overly short ones hindered atom reuse. To address this, we instructed the model to split enumerations containing more than two elements into separate law snippets while preserving shorter ones intact.

Note that pre-processing might not be required for all legal texts. In some normative acts, paragraphs are sufficiently short and need no further subdivision. However, in documents like the Australian Telecommunications Consumer Protections Code, where individual articles can span 4-5 pages, splitting the text into smaller segments helps the LLM systematically analyze each component without overlooking critical details.

### 3.2. Transformation into DDL

Each law snippet is transformed into DDL rules via various prompting strategies. These include Chain-Of-Instructions (CoI) prompting and few-shot learning using prompt variants with progressively enhanced instructions. We also evaluate a pipeline approach where atom extraction and DDL rule generation are handled by different LLMs in sequence.

Note that despite following OpenAI's guidelines for achieving reproducible outputs [19], e.g., fixing the seed and temperature parameters, we observed non-deterministic behavior. This phenomenon can be attributed to inherent LLM stochasticity [20].

### 3.3. Evaluation

We evaluate the generated rules across six dimensions, each operationalized as a concrete question: completeness (Q1), syntactic (Q2) and semantic correctness (Q3), deontic modality accuracy (Q4), precondition appropriateness (Q5), and meaningfulness/reuse of atom names (Q6).

It is important to note that a single law snippet may lead to the generation of multiple rules. Furthermore, Q1 is assessed based on the law snippet as a whole, taking into account all rules derived from it. In contrast, Q2 through Q6 are evaluated individually for each generated rule. The questions are ordered such that earlier ones address more general and fundamental aspects of correctness, while later ones examine increasingly fine-grained details. Importantly, the evaluation follows a short-circuiting scheme: if for some rule  $r$  a question  $Q_i$  with  $i \geq 2$  is evaluated as `false`, then all subsequent questions  $Q_j$  with  $j > i$  are not considered and are implicitly assigned the value `false`.

#### [Q1: Completeness.] Are all aspects of the law text formalized?

Consider, for instance, the following formalization of law snippet 8.2.1.a.xiv:

```
complaint(X), consentConsumer(X) ⇒ [P] closeComplaint(X)
complaint(X), complied8.2.1.c(X) ⇒ [P] closeComplaint(X)
complaint(X), complied8.2.1.d(X) ⇒ [P] closeComplaint(X)
complaint(X), complied8.2.1.e(X) ⇒ [P] closeComplaint(X)
```

This is not a complete formalization of the facts, as the following rule is missing:

```
complaint(X) ⇒ [O] -closeComplaint(X)
```

This initial check is crucial to prevent the LLM from achieving a high score merely by formalizing only the simplest aspects of a problem.

#### [Q2: Syntactic Validity.] Is the rule syntactically valid and non-redundant?

An example of a rule that fails the syntactic validity check is the following:

$\text{closeComplaint}(X), \text{-consent}(X), \text{-clausesCDEComplied}(X) \Rightarrow [\text{O}] \text{closeComplaint}(X)$

Note that the consequence of this rule also appears as its antecedent. However, this issue was later resolved by adding a corresponding instruction to the prompt.

**[Q3: Semantic Correctness.] Is the rule semantically valid and non-redundant?**

This question serves as a “catch-all” check that applies when no other question describes the problem better, for example, when hallucinations of the LLMs occur. The following rules fail this check, as the atoms  $\text{informResolution}(X)$  and  $\text{informNoResolution}(X)$  are unrelated to the facts described in the legal text.

$\text{informResolution}(X) \Rightarrow [\text{P}] \text{closeComplaint}(X)$

$\text{informNoResolution}(X) \Rightarrow [\text{P}] \text{closeComplaint}(X)$

However, there are also more subtle issues filtered by this question, for example, when a LLM combines several aspects with a logical “and”, even though they should be connected with a logical “or” according to the legal text. This question also verifies whether formalizations that are not syntactically identical to another rule convey the same meaning and are therefore redundant.

**[Q4: Deontic Modality Accuracy.] Are the Deontic modalities and negations correctly placed?**

In this example, a permission is incorrectly formalized as an obligation:

$\text{complaint}(X), \text{consentConsumer}(X) \Rightarrow [\text{O}] \text{closeComplaint}(X)$

Hence, the question would be answered with `false`, and no further checks performed. Note that this output stems from an early variation of the prompt. Such an error did not occur in later iterations.

**[Q5: Precondition Appropriateness.] Is the precondition appropriate?**

A common problem was that the precondition of the rules contained either too many, too few or wrong atoms. This question should cover precisely these cases.

Consider for instance the following formalization generated in an experiment:

$\text{consentConsumer}(X) \Rightarrow [\text{P}] \text{closeComplaint}(X)$

$\text{compliedWithClauseC}(X) \Rightarrow [\text{P}] \text{closeComplaint}(X)$

$\text{compliedWithClauseD}(X) \Rightarrow [\text{P}] \text{closeComplaint}(X)$

$\text{compliedWithClauseE}(X) \Rightarrow [\text{P}] \text{closeComplaint}(X)$

$\text{-consentConsumer}(X), \text{-compliedWithClauseC}(X), \text{-compliedWithClauseD}(X),$   
 $\text{-compliedWithClauseE}(X) \Rightarrow [\text{F}] \text{closeComplaint}(X)$

In the last rule, it is not necessary that all these atoms are included in the precondition. A simple  $\text{complaint}(X)$  would have been enough – that the prohibition to close the complaint does not hold when there is consent from the consumer already follows from the first rule.

**[Q6: Meaningfulness/Reuse of Atom Names.] Are the atom names meaningful and, if appropriate, reused?**

Consider again the above formalization, for example, the atom  $\text{compliedWithClauseC}(X)$ . Unfortunately, it is not fully clear from the atom name to which clause the name is referring – a better name would be  $\text{clause8.2.1cComplied}(X)$ .

In the success score calculation, we represent the outcomes of questions Q1 through Q6 using binary values: 1 for `true` (satisfied) and 0 for `false` (not satisfied).

We define  $Q_i(r) \in \{0, 1\}$  as the evaluation of question  $Q_i$  on rule  $r$ . For a given law snippet  $l$ , let  $\mathcal{R}(l)$  denote the set of rules generated from  $l$ . We first introduce a **modifier function**  $m(l)$ , which reduces the success score by half if Q1 is not satisfied over the entire snippet:

$$m(l) = \begin{cases} 0.5 & \text{if } Q1(l) = 0, \\ 1 & \text{if } Q1(l) = 1, \end{cases}$$

where  $Q_1(l)$  evaluates the overall satisfaction of Q1 across the law snippet  $l$  (i.e., by considering all generated rules together). The **success score**  $S(l)$  for an individual law snippet  $l$  is then defined as:

$$S(l) = m(l) \times \frac{1}{|\mathcal{R}(l)|} \sum_{r \in \mathcal{R}(l)} \frac{1}{5} \sum_{i=2}^6 Q_i(r).$$

The **overall success score**  $S(\mathcal{L})$  for a set of law snippets  $\mathcal{L}$  is the average success score of the individual snippets:

$$S(\mathcal{L}) = \frac{1}{|\mathcal{L}|} \sum_{l \in \mathcal{L}} S(l).$$

In addition, we define a stricter evaluation  $S^*$  where only perfect formalizations contribute to the success score:

$$S^*(l) = \begin{cases} 1 & \text{if } S(l) = 1, \\ 0 & \text{otherwise} \end{cases}, \quad S^*(\mathcal{L}) = \frac{1}{|\mathcal{L}|} \sum_{l \in \mathcal{L}} S^*(l).$$

## 4. Results

A series of experiments has been conducted to identify the LLM configuration that is the most promising for the formalization task. These experiments are conducted on real legal content from Sections 8.2.1(a)–(c) of the TCP Code.

First, we start with an initial experiment, where the LLMs are given a prompt with detailed instructions how to solve the problem, and the respective law snippet to formalize. We also perform different variations of the experiment, including different output formats and varied prompts. We then evaluate fine-tuned models. Finally, we implement a pipeline where two LLMs work together to solve the task. Specifically, one LLM is responsible for extracting the atoms and another for the actual formalization of the DDL rules.

In each of the following experiments, we provided LLMs with a prompt containing step-by-step instructions to guide the model through the extraction process. This approach is called Chain-of-Instructions (CoI) prompting [21]. Hence, the model is encouraged to solve each subtask step by step until the final answer is reached [22]. This method contrasts with Chain-of-Thoughts (CoT), which usually depends more on implicit reasoning [21] – especially for Zero-Shot-CoT, where just a sentence like “*Let’s think step by step*” is appended to the prompt [23].

Moreover, we use few-shot learning [24], where we provide the LLM with a few examples (input-output pairs) in the prompt to demonstrate how to solve the task.

In all experiments conducted, the prompt was passed to the LLMs via a system message.

### 4.1. Prompt Development

The prompt employed in our experiments was derived through a series of iterative refinements. Beginning with an initial two-shot learning prompt, successive modifications were made to enhance the clarity of the instructions and the quality of the generated outputs. These iterations involved the inclusion of additional guidance and examples to better align the model’s behavior with the desired output format. The final version utilizes a three-shot learning approach. Listing 1 provides the complete and final prompt used in our evaluation.

Transform legal text in natural language to expressions in Defeasible Deontic Logic (DDL) in XML format. Each atom should end with "(X)". If you want to represent a conjunction, separate the atoms by a comma. If you want to represent a disjunction, please use multiple rules and do not write it as a single rule. Output only a single <Paragraph> element with multiple <Rule> elements if necessary. Make sure to output valid XML. Represent obligations with [O], permissions with [P], and prohibitions with [F]. If you want to negate an atom, use the negation symbol "-" before the atom or the deontic operator. Each rule should have only one consequence. If you want to represent multiple consequences, please use multiple rules. Since the law snippet are talking

about complaint handling, in most preconditions, there will be an atom like complaint(X). Make sure to keep the atoms in the precondition as simple as possible. If it is possible to break down the atoms into smaller parts, please do so. For example, instead of urgentComplaint(X), write complaint(X), urgent(X). Moreover, NEVER put an atom in the antecedent if it also appears in the consequence, because this would be syntactically invalid.

Work in the following steps:

1. Define the atoms that will be used in the rules.
2. Define the if-then structure of the rules.
3. Identify deontic modalities.
4. Formalize the rules in the given format using Defeasible Deontic Logic (DDL).

# Example 1

## Input

8.1.1.1 A Supplier must take the following actions to enable this outcome:

(c) Ensure awareness and visibility: ensure their staff who have direct contact with Consumers or former Customers, including personnel working for contractors, understand the Supplier's Complaint handling process, their responsibilities under it and are able to identify and record a Complaint.

## Output

```
<Paragraph paragraphLabel="8.1.1.c">
  <Rules>
    <Rule ruleLabel="tcpc.8.1.1.c.1">
      complaintHandlingProcess(X) => [O] relevantStaffAwareComplaintHandlingProcess(X)
    </Rule>
    <Rule ruleLabel="tcpc.8.1.1.c.2">
      complaintHandlingProcess(X) => [O] relevantStaffAbleToHandleComplaint(X)
    </Rule>
  </Rules>
</Paragraph>
```

# Example 2

## Input

8.1.1.1 A Supplier must take the following actions to enable this outcome:

(a) Implement a process: implement, operate and comply with a Complaint handling process that:  
(x) is transparent, including:  
D. requiring Consumers or former Customers to be advised of the Resolution of their Complaint; and

## Output

```
<Paragraph paragraphLabel="8.1.1.a.x.D">
  <Rules>
    <Rule ruleLabel="tcpc.8.1.1.a.x.D">
      complaint(X), resolution(X) => [O] informResolution(X)
    </Rule>
  </Rules>
</Paragraph>
```

# Example 3

## Input

8.5.1.1 A Supplier must take the following actions to enable this outcome:

(e) Maintain confidentiality: Suppliers not subject to the requirements of the Privacy Act must ensure personal information concerning a Complaint is not disclosed except as required to manage a Complaint with the TIO or with the express consent of the Consumer.

## Output

```
<Paragraph paragraphLabel="8.5.1.e">
  <Rules>
    <Rule ruleLabel="tcpc.8.5.1.e.1">
      complaintHandlingProcess(X), personalInformation(X), -subjectPrivacyAct(X) =>
        [O] -discloseInformation(X)
    </Rule>
    <Rule ruleLabel="tcpc.8.5.1.e.2">
      personalInformation(X), requestFromTIO(X) => [O] discloseInformation(X)
    </Rule>
  </Rules>
</Paragraph>
```

```

<Rule ruleLabel="tcpc.8.5.1.e.3">
  consentDisclosurePersonalInformation(X) => [P] discloseInformation(X)
</Rule>
</Rules>
</Paragraph>

```

Listing 1: Best prompt

The final prompt was evaluated across multiple LLMs. Two diagrams summarize the results: one displaying the standard success scores  $S$  (s. Figure 1a) and another illustrating the success scores  $S^*$  under the stricter criterion of perfect formalizations (s. Figure 1b).

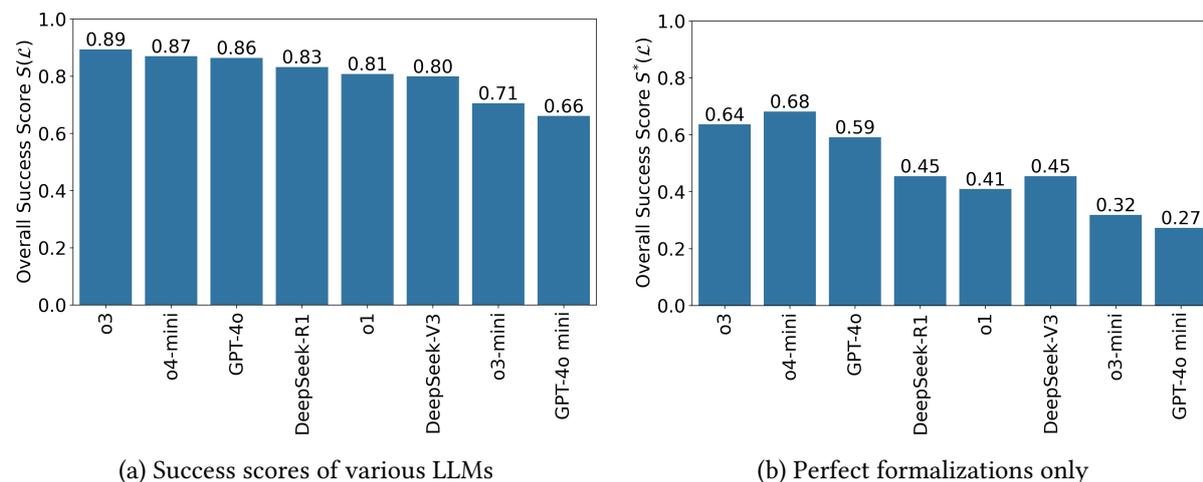


Figure 1: Success scores across models.

## 4.2. Consideration of Multiple Law Snippets Simultaneously

In the experiments described in Section 4.1, the prompt was sent together with an individual law snippet to the LLM. This approach minimized token consumption but limited the reuse of atom names across snippets. Here, we investigate whether providing multiple law snippets simultaneously enhances formalization performance.

### 4.2.1. Incorporating the Formalization History

In one variant, the complete formalization history was included by alternating user and assistant messages for all prior snippets, aiming to encourage more consistent reuse of atom names across different law texts.

However, no improvement was observed compared to the single-snippet baseline; in fact, the success scores were marginally lower. A plausible explanation is that the additional context overwhelmed the models, hindering their ability to focus effectively on the current snippet.

### 4.2.2. Providing Only Previously Formalized Atoms

In a second variant, previously extracted atom names were provided collectively in a single user message, rather than replicating the entire prior dialogue history. For each new law snippet, three messages were sent to the LLM: (1) a system prompt, (2) a user message listing previously formalized atom names (cf. Listing 2), and (3) a user message containing the new law snippet to be formalized.

Although an increased reuse of atom names was observed, the atoms were often applied in inappropriate or irrelevant contexts. As a result, this approach led to a greater number of hallucinations rather than an improvement in the formalizations. Consequently, no further evaluation of this strategy was pursued.

```

Try to reuse the following atoms you have used for the formalization of previous paragraphs:
* complaint(X)
* madeInPerson(X)
* acknowledgeImmediately(X)
...

```

Listing 2: user message including previous atom names

### 4.2.3. Formalizing All Law Snippets in a Single Interaction

In a final approach, all law snippets were provided together within a single user prompt to the LLM. While the input text contained multiple snippets, the division into distinct law snippets was preserved to encourage the model to treat each snippet individually.

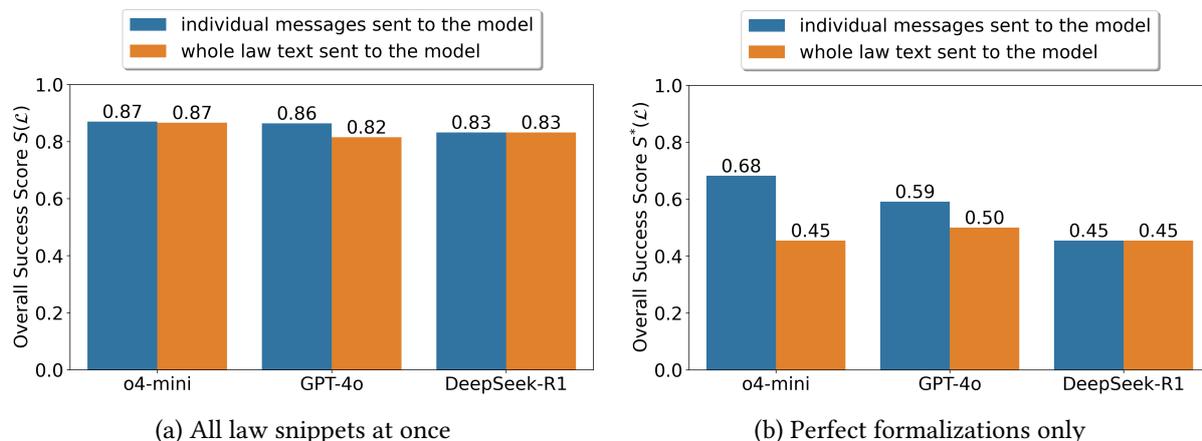


Figure 2: Success scores when formalizing all law snippets in a single interaction.

Note that it was not possible to evaluate OpenAI’s o3 model in this experiment, as this model did not adhere to the predefined output structure and issued rules without further structuring in law snippets.

Figures 2a and 2b present the results for this setting, with the latter considering only perfect formalizations.

Although this approach led to a slight increase in atom reuse across snippets, it exhibited a major drawback: the generated formalizations were often less detailed compared to the baseline obtained in Section 4.1. In particular, important facts were frequently merged into a single rule, even in cases where separate rules would have been necessary for a proper and precise formalization.

## 4.3. Fine-Tuning

Fine-tuning is a transfer learning technique where pretrained model weights are adapted to a new task through further training. By leveraging knowledge acquired during pretraining, fine-tuning can substantially enhance model performance, particularly in scenarios characterized by limited training data. Prior work has demonstrated the effectiveness of fine-tuning LLMs in improving task-specific outcomes [25].

In the present study, fine-tuning experiments were conducted with GPT-4o. Although the proprietary nature of GPT-4o does not allow direct access to the model weights, OpenAI provides fine-tuning capabilities for non-reasoning models via its platform. Note that OpenAI’s reasoning models are not amendable to fine-tuning.

Given that only 22 law snippets from the dataset presented in [8] correspond to Sections 8.2.1(a)–8.2.1(c), the remaining 44 snippets from unrelated sections were utilized as training data.

Three distinct fine-tuning configurations were evaluated, as summarized in Table 1.



```
(a) Implement a process: implement, operate and comply with a Complaint handling process that:
(v) clearly states that Consumers or former Customers have a right to make a Complaint
and that a proposed Resolution must be accepted by a Consumer or former Customer before a
Supplier is required to implement it;
```

```
## Output
```

```
informRightToMakeComplaint(X): Supplier informs customer of right to make a complaint.
informComplaintHandlingProcess(X): Supplier informs Customer of Complaint handling process.
complaintHandlingProcess(X): Supplier has a complaint handling process as per TCPC section 8.
```

Listing 3: Prompt for atom extraction (2 examples omitted)

Four separate experiments on atom extraction were conducted, involving the models DeepSeek-R1, a fine-tuned variant of GPT-4o, OpenAI o3, and OpenAI o4-mini. As in Section 4.3, the 44 law snippets from [8] not associated with Sections 8.2.1(a)–8.2.1(c) were used as training data for the fine-tuning variant.

In the subsequent stage, the generation of DDL rules was performed based on the legal text and the previously extracted atom definitions.

The results of these experiments are presented in Figures 4a and 4b. Figure 4a compares the success scores of the two-stage pipeline against those achieved with the best prompt from Section 4.1 (blue bars). Figure 4b shows the corresponding comparison when only perfect formalizations are considered.

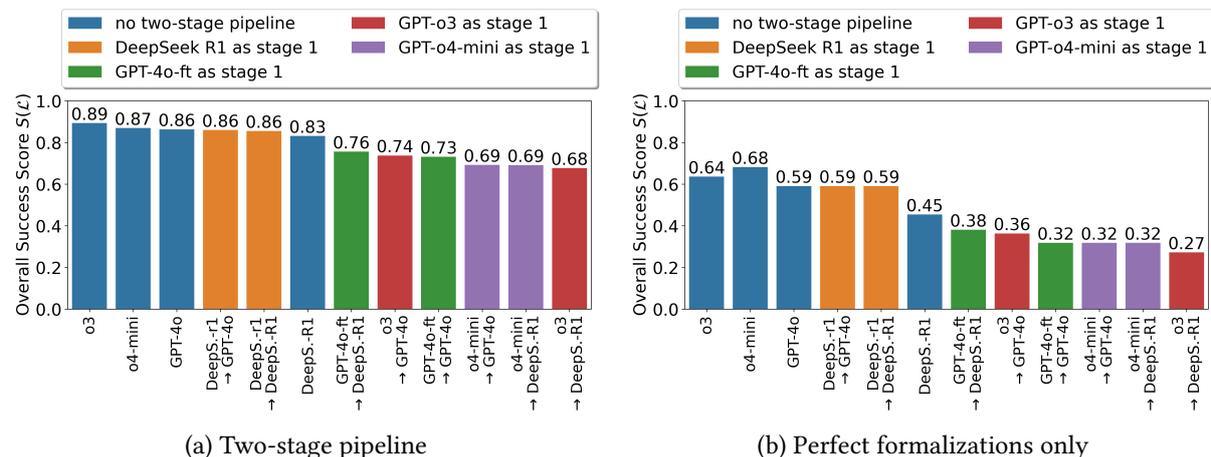


Figure 4: Success scores of the two-stage pipeline.

The results indicate that employing DeepSeek-R1 for atom extraction yields the most favorable outcomes, although they are inferior to the best results obtained without using a two-stage pipeline.

Furthermore, since overall performance remained relatively stable regardless of the model used in the second stage, it can be inferred that the primary source of error originates from the atom extraction step. Consequently, further optimization efforts should prioritize improving atom extraction rather than refining the second stage of the pipeline.

#### 4.5. Comparison with Dragoni et al. [8]

In this section, we compare our experimental findings with the results presented by Dragoni et al. [8]. To ensure methodological consistency, we restrict the comparison to the *lower branch evaluation* reported in their study. For this purpose, we employ the standard metrics of *Precision*, *Recall*, and their harmonic

mean, the *F1-score*, defined as follows:

$$\begin{aligned} \text{Precision} &= \frac{\text{TP}}{\text{TP} + \text{FP}} = \frac{\text{Matched Items}}{\text{Generated Items}} \\ \text{Recall} &= \frac{\text{TP}}{\text{TP} + \text{FN}} = \frac{\text{Matched Items}}{\text{Gold Standard Items}} \\ \text{F1} &= 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \end{aligned}$$

In this context, *Items* refers to either atoms or rules, depending on the evaluation. True positives (TP) are items generated by the LLM that are also found in the gold standard. False positives (FP) are generated items not present in the gold standard. False negatives (FN) are not explicitly observed. The denominator in the classical recall formula, TP + FN, serves as a proxy for the total number of positives in the labeled set. However, in our case, the gold standard itself is the labeled set and contains only positives. Therefore, the total number of positives is directly known and equals the number of items in the gold standard.

Our analysis reveals an immediate discrepancy: Dragoni et al. report 65 terms and 36 rules in the gold standard, whereas we identified 69 terms and 52 rules across Sections 8.2.1(a)–8.2.1(c) in the gold standard.

Importantly, in the following analysis, we evaluate the precision and recall of the formalizations produced by the LLMs based on our counts. As a result, the reported metrics are not directly comparable to those presented in the work of Dragoni et al.

#### 4.5.1. Evaluation of Term Identification

The first level of comparison involves the correct identification of legal terms (referred to as *atoms*). Dragoni et al. report a precision of 83.05% and a recall of 90.78%. Table 2 summarizes the precision and recall achieved across various configurations and models in our study. The best precision, 86.21%, is achieved with DeepSeek-R1 when all law snippets were provided together within a single user prompt. The best recall, 84.06%, is achieved in the baseline setting with GPT-o3.

**Table 2**  
Precision and Recall for Term Identification

Experiment	Model	Generated Atoms	Matched Atoms	Precision	Recall	F1
Results from [8]				83.05%	<b>90.78%</b>	<b>0.87</b>
Simple Col (Section 4.1)	o3	93	58	62.37%	84.06%	0.72
	o4-mini	96	57	59.38%	82.61%	0.69
	DeepSeek-R1	84	56	66.67%	81.16%	0.73
	GPT-4o	77	56	72.73%	81.16%	0.77
Single Interaction (Section 4.2.3)	o4-mini	64	52	81.25%	75.36%	0.78
	DeepSeek-R1	58	50	<b>86.21%</b>	72.46%	0.79
	GPT-4o	64	54	84.38%	78.26%	0.81
Fine-tuned (Section 4.3)	GPT-4o (Config 1, Epoch 1)	57	42	73.68%	60.87%	0.67
	GPT-4o (Config 2, Epoch 1)	71	57	80.28%	82.61%	0.81
	GPT-4o (Config 3, Epoch 1)	72	57	79.17%	82.61%	0.81
Two-Stage (Section 4.4)	DeepSeek-R1 → DeepSeek-R1	73	49	67.12%	71.01%	0.69
	DeepSeek-R1 → GPT-4o	72	49	68.06%	71.01%	0.70

### 4.5.2. Evaluation of Deontic Annotation Accuracy

The second dimension of analysis concerns the accurate assignment of deontic modality (i.e., obligation, permission, or prohibition). In the benchmark study, 47 of 49 correctly identified atoms were accurately annotated, yielding a deontic annotation precision of 95.92%. In contrast, across all experiments conducted in this work, 100% of atoms – both correctly identified atoms and such without a counterpart in the gold standard – were annotated with the correct deontic label. Thus, the deontic annotation precision in our experiments is consistently 100%.

### 4.5.3. Identification of Rule Counterparts

The third criterion evaluates the number of generated rules that have a semantically corresponding rule in the gold standard. Following the method defined in [8], a rule is considered a *counterpart* if there is a semantic match in its consequent with a rule in the manually curated set of the gold standard. In their evaluation, 33 out of 36 gold standard rules had counterparts, resulting in a precision of 80.49% and a recall of 91.67%. Our results are presented in Table 3. The best precision, 93.33%, is achieved with a fine-tuned version of GPT-4o, whereas the best recall, 84.62%, with GPT-o3 in the baseline setting.

**Table 3**  
Precision and Recall for Counterpart Identification

Experiment	Model	Generated Rules	Matched Rules	Precision	Recall	F1
Results from [8]				80.49%	<b>91.67%</b>	0.86
Simple Col (Section 4.1)	o3	57	44	77.19%	84.62%	0.81
	o4-mini	51	39	76.47%	75.00%	0.76
	DeepSeek-R1	47	42	89.36%	80.77%	0.85
	GPT-4o	47	42	89.36%	80.77%	0.85
Single Interaction (Section 4.2.3)	o4-mini	47	41	87.23%	78.85%	0.83
	DeepSeek-R1	31	27	87.10%	51.92%	0.65
	GPT-4o	36	30	83.33%	57.69%	0.68
Fine-tuned (Section 4.3)	GPT-4o (Config 1, Epoch 1)	33	30	90.91%	57.69%	0.71
	GPT-4o (Config 2, Epoch 1)	45	41	91.11%	78.85%	0.85
	GPT-4o (Config 3, Epoch 1)	45	42	<b>93.33%</b>	80.77%	<b>0.87</b>
Two-Stage (Section 4.4)	DeepSeek-R1 → DeepSeek-R1	44	36	81.82%	69.23%	0.75
	DeepSeek-R1 → GPT-4o	44	36	81.82%	69.23%	0.75

### 4.5.4. Evaluation of Full Rule Correspondence

Finally, we assess the degree of *full correspondence*, where a rule in the generated set semantically matches the gold standard in both antecedents and consequent. Dragoni et al. report that 24 rules in their generated set fully matched semantically their counterparts in the gold standard, with a resulting precision of 58.54% and recall of 66.67%. The corresponding results from our evaluation are displayed in Table 4. The best precision (80.56%) is achieved with a fine-tuned version of GPT-4o, while the best recall value (69.23%) is reached by multiple approaches simultaneously.

### 4.5.5. Limitations of the Compared Evaluation Metrics

While the evaluation framework used by Dragoni et al. has certain benefits – such as penalizing insufficient reuse of atoms across legal clauses – it also presents some limitations.

A key issue lies in the one-to-one rule mapping constraint: each rule in the gold standard can be matched to at most one rule in the generated set and vice versa. This restriction becomes problematic

**Table 4**  
Precision and Recall for Full Rule Correspondence

Experiment	Model	Generated Rules	Matched Rules	Precision	Recall	F1
Results from [8]				58.54%	66.67%	0.62
Simple Col (Section 4.1)	o3	57	36	63.16%	<b>69.23%</b>	0.66
	o4-mini	51	30	58.82%	57.69%	0.58
	DeepSeek-R1	47	31	65.96%	59.62%	0.63
	GPT-4o	47	32	68.09%	61.54%	0.65
Single Interaction (Section 4.2.3)	o4-mini	47	36	76.60%	<b>69.23%</b>	0.73
	DeepSeek-R1	31	23	74.19%	44.23%	0.55
	GPT-4o	36	29	<b>80.56%</b>	55.57%	0.66
Fine-tuned (Section 4.3)	GPT-4o (Config 1, Epoch 1)	33	24	72.73%	46.15%	0.56
	GPT-4o (Config 2, Epoch 1)	45	36	80.00%	<b>69.23%</b>	<b>0.74</b>
	GPT-4o (Config 3, Epoch 1)	45	35	77.78%	67.31%	0.72
Two-Stage (Section 4.4)	DeepSeek-R1	44	27	61.36%	51.92%	0.56
	GPT-4o	44	29	65.91%	55.77%	0.60

in cases where an LLM produces multiple valid rules which all together are equivalent to a single rule in the gold standard. In such scenarios, semantically correct rules are penalized due to a lack of corresponding entries in the gold standard.

For instance, in the formalization of law snippet 8.2.1.a.i, the gold standard specifies only three rules, while the LLM-generated formalizations typically consist of six rules, offering a more detailed representation. Nonetheless, these additional rules are considered false positives in the evaluation, thus lowering precision.

The same limitation applies to term identification: valid atoms that are not present in the gold standard reduce the measured precision, even if their extraction is semantically justified.

An additional shortcoming is that LLMs are penalized for formalizing additional information from the law text that is not represented in the gold standard. Consider, for example, law snippet 8.2.1.b, which contains the following sentence:

”If a Consumer tells the Supplier that they are dissatisfied with the timeframes that apply to the management of a Complaint or seek to have a Complaint treated as an Urgent Complaint, the Supplier must tell the Consumer about the Supplier’s internal prioritisation and internal escalation processes.“

In the gold standard, this clause is formalized as:

$$\begin{aligned} \text{customerDissatisfiedTimeframe}(X) &\Rightarrow [O] \text{informInternalPrioritisation}(X) \\ \text{customerDissatisfiedTimeframe}(X) &\Rightarrow [O] \text{informInternalEscalationProcess}(X) \\ \text{escalation}(X), \text{internalPrioritisation}(X) &\Rightarrow [O] \text{informExternalDisputeResolution}(X) \end{aligned}$$

In contrast, some LLMs produced the following additional formalizations:

$$\begin{aligned} \text{complaint}(X), \text{consumerRequestsUrgent}(X) &\Rightarrow [O] \text{informInternalPrioritisation}(X) \\ \text{complaint}(X), \text{consumerRequestsUrgent}(X) &\Rightarrow [O] \text{informInternalEscalation}(X) \end{aligned}$$

Although these rules are arguably justified by the legal text, they are penalized under the evaluation framework due to their absence from the gold standard. Thus, the metric fails to distinguish between semantically valid additions and hallucinated additions, undermining its reliability in assessing true model performance.

## 5. Limitations

We describe both the intrinsic limitations and the technical challenges we encountered.

### 5.1. Legal Interpretation

A major challenge for the formal encoding of legal documents is that each encoding is an interpretation, and the gold standard should correspond to the authentic interpretation. However, in some jurisdictions, it will not be possible to have a true gold standard. The gold standard would correspond to the authentic interpretation of the legal provision, and the only authority able to provide an authentic interpretation is the judiciary. Moreover, this is possible only for cases disputed in court, and it would be limited to the provisions effectively used in the legal proceeding. The second issue is that a legal interpretation depends on the understanding of the legal intent, legal context and the encoding style of the coders. [28] reports on an empirical experiment where three (experienced) coders were asked to model in DDL the same set of legal provisions (from the Australian Copyright Act). The experiment had two phases; in the first phase, the coders did the encoding fully independently. In the second phase, the coders agreed on a common set of terms and then encoded independently the provisions as rules. In the first phase, the degree of agreement varied from 0% to 10% for terms, and 0% of rules using a perfect match, and around 50% for terms and 3% on rules with a semantics correspondence. In phase two, the term agreement was between 30% and 56% for the full correspondence and 85% for semantic correspondence; the rule similarity ranged from 10% to 30% for full correspondence and 26% to 53% with semantic correspondence.

### 5.2. Inter-Paragraph References

A common strategy for managing the complexity of legal documents is the use of references, which may be internal (linking to other sections within the same document) or external (referring to other documents) [29]. Accordingly, the dataset used in this study included occasional references between different legal paragraphs. For example, law snippet 8.2.1.a.xiv mandates that a complaint must only be closed “*with the consent of the Consumer or former Customer or if clauses 8.2.1(c),(d) or (e) below have been complied with*”.

The LLMs generated suitable atoms such as `clause8.2.1.c.complied(X)`, but these were not reused in the formalizations of the referenced paragraphs (i.e., 8.2.1(c), (d), or (e)). As a consequence, the preconditions set in the formalization of law snippet 8.2.1(a)(xiv) could not be met, rendering this rule ineffective within the formal system.

This problem persisted even when using the methodology presented in Section 4.2. This limitation suggests that prompt engineering alone is insufficient to fully address the challenge of reference resolution. Instead, it requires the incorporation of additional procedural components into the methodology. One possible solution is a refinement phase following the initial generation process, designed to ensure semantic coherence across references.

### 5.3. Atom Reuse Across Legal Snippets

Effective formalization of legal texts requires the consistent reuse of atoms across rules; otherwise, reasoning with the resulting rule set containing many redundant atoms would not be useful. However, the experimental results revealed that the LLMs rarely reused atoms across different law snippets in the current setup, except for `complaint(X)` which was explicitly required in the prompt. For instance, while the gold standard consistently employed the atom `resolvable15Days(X)`, most LLMs generated varying alternatives such as `resolveIn15Days(X)`, `resolvedBy15Days(X)` and `cannotResolveIn15Days(X)` across different law snippets. Although OpenAI’s newer reasoning models (e.g., o3 and o4-mini) have shown success in various evaluations, this issue is particularly significant in their outputs, as they produced up to 96 atoms compared to the 69 present in the gold standard.

Possible strategies to address this problem include:

(i) As noted in Section 4.2.2, simply providing LLMs with all previously generated atoms did not improve reuse; instead, it increased hallucinations. A more effective approach may involve identifying the most relevant atoms in advance and selectively providing only those. Alternatively, previously generated atoms may be supplied exclusively to the initial atom extraction phase within a multi-phase pipeline, while subsequent phases focus on generating logically coherent DDL rules, potentially correcting earlier hallucinations.

(ii) We could introduce an intermediate step between atom generation and DDL rule formalization, in which similar atoms could be clustered and evaluated by an LLM to determine whether they should be merged, thereby promoting consistency and reuse.

## 6. Conclusion and Future Work

This paper showed that LLMs can be leveraged to transform legal norms into formal DDL rules with substantial fidelity, and with performance similar to the approach proposed in [8]. Our evaluations confirm that prompt engineering – especially few-shot learning with Chain-of-Instructions – can significantly improve the semantic precision of extracted rules. Fine-tuning offers benefits but risks overfitting, while multi-stage pipelines are promising but sensitive to the quality of initial atom extraction.

Future work includes integrating active learning and expert-in-the-loop feedback to continuously refine LLM outputs. Expanding the domain beyond TCP Code and adapting the pipeline to multilingual legal corpora could further validate the generalizability of our approach. Moreover, the formalization of the superiority relationship – currently omitted due to limited occurrences in the dataset – deserves further investigation, potentially via prompt engineering or a dedicated pipeline stage. Finally, embedding these methods in end-user tools for compliance and regulatory auditing represents a practical next step.

## Acknowledgments

This work has been partially funded by the Vienna Science and Technology Fund (WWTF) [Grant ID: 10.47379/ICT23030].

## Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

## References

- [1] M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, H. T. Cory, The British Nationality Act as a logic program, *Communications of the ACM* 29 (1986) 370–386. doi:10.1145/5689.5920.
- [2] T. J. M. Bench-Capon, M. Araszkiwicz, K. D. Ashley, K. Atkinson, F. Bex, F. Borges, D. Bourcier, P. Bourguine, J. G. Conrad, E. Francesconi, T. F. Gordon, G. Governatori, J. L. Leidner, D. D. Lewis, R. P. Loui, L. T. McCarty, H. Prakken, F. Schilder, E. Schweighofer, P. Thompson, A. Tyrrell, B. Verheij, D. N. Walton, A. Z. Wyner, A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law, *Artificial Intelligence and Law* 20 (2012) 215–319.
- [3] G. Governatori, T. Bench-Capon, B. Verheij, M. Araszkiwicz, E. Francesconi, M. Grabmair, Thirty years of Artificial Intelligence and Law: The first decade, *Artificial Intelligence and Law* 30 (2022) 481–519. doi:10.1007/s10506-022-09329-4.
- [4] J. Mohun, A. Roberts, *Cracking the code: Rulemaking for humans and machines*, OECD Working Papers on Public Governance, OECD, Paris, France, 2020. doi:10.1787/3afe6ba5-en.

- [5] M. Cristani, G. Governatori, F. Olivieri, M. Palmirani, G. Buriola, Explainability by design: an experimental analysis of the legal coding process, 2025. URL: <https://arxiv.org/abs/2505.01944>. arXiv: 2505.01944.
- [6] A. Z. Wyner, W. Peters, On rule extraction from regulations, in: K. Atkinson (Ed.), *Legal Knowledge and Information Systems - JURIX 2011: The Twenty-Fourth Annual Conference*, University of Vienna, Austria, 14th-16th December 2011, volume 235 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2011, pp. 113–122. URL: <https://doi.org/10.3233/978-1-60750-981-3-113>. doi:10.3233/978-1-60750-981-3-113.
- [7] A. Wyner, G. Governatori, A study on translating regulatory rules from natural language to Defeasible Logics, in: P. Fodor, D. Roman, D. Anicic, A. Wyner, M. Palmirani, D. Sottara, F. Lévy (Eds.), *RuleML (2)*, volume 1004 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2013.
- [8] M. Dragoni, S. Villata, W. Rizzi, G. Governatori, Combining natural language processing approaches for rule extraction from legal documents, in: U. Pagallo, M. Palmirani, P. Casanovas, G. Sartor, S. Villata (Eds.), *AI Approaches to the Complexity of Legal Systems - AICOL International Workshops 2015-2017: AICOL-VI@JURIX 2015, AICOL-VII@EKAW 2016, AICOL-VIII@JURIX 2016, AICOL-IX@ICAIL 2017, and AICOL-X@JURIX 2017*, Revised Selected Papers, volume 10791 of *Lecture Notes in Computer Science*, Springer, 2017, pp. 287–300. URL: [https://doi.org/10.1007/978-3-030-00178-0\\_19](https://doi.org/10.1007/978-3-030-00178-0_19). doi:10.1007/978-3-030-00178-0\_19.
- [9] G. Ferraro, H.-P. Lam, S. Colombo Tosatto, F. Olivieri, M. B. Islam, N. van Beest, G. Governatori, Automatic extraction of legal norms: Evaluation of natural language processing tools, in: M. Sakamoto, N. Okazaki, K. Mineshima, K. Satoh (Eds.), *New Frontiers in Artificial Intelligence. JSAI-isAI 2019*, volume 12331 of *LNCS*, Springer, Cham, 2019, pp. 64–81. doi:10.1007/978-3-030-58790-1\_5.
- [10] M. Billi, G. Pisano, M. Sanchi, Fighting the knowledge representation bottleneck with Large Language Models, in: J. Savelka, J. Harasta, T. Novotná, J. Mísek (Eds.), *Legal Knowledge and Information Systems - JURIX 2024*, volume 395 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2024, pp. 14–24. URL: <https://doi.org/10.3233/FAIA241230>. doi:10.3233/FAIA241230.
- [11] G. Governatori, F. Olivieri, A. Rotolo, S. Scannapieco, Computing strong and weak permissions in Defeasible Logic, *J. Philos. Log.* 42 (2013) 799–829. URL: <https://doi.org/10.1007/s10992-013-9295-1>. doi:10.1007/s10992-013-9295-1.
- [12] G. Governatori, A. Rotolo, G. Sartor, Logic and the Law: Philosophical foundations, Deontics, and Defeasible Reasoning, in: D. M. Gabbay, J. Horty, X. Parent, R. van der Meyden, L. van der Torre (Eds.), *Handbook of Deontic Logic and Normative Reasoning*, volume 2, College Publications, London, 2021, pp. 655–760.
- [13] M. Peepkorn, T. Kouwenhoven, D. Brown, A. Jordanous, Is temperature the creativity parameter of Large Language Models?, in: K. Grace, M. T. Llano, P. Martins, M. M. Hedblom (Eds.), *Proceedings of the 15th International Conference on Computational Creativity, ICCC 2024*, Jönköping, Sweden, June 17-21, 2024, Association for Computational Creativity (ACC), 2024, pp. 226–235. URL: [https://computationalcreativity.net/iccc24/papers/ICCC24\\_paper\\_70.pdf](https://computationalcreativity.net/iccc24/papers/ICCC24_paper_70.pdf).
- [14] E. Manjavacas, F. Karsdorp, B. Burtenshaw, M. Kestemont, Synthetic literature: Writing science fiction in a co-creative process, in: *Proceedings of the Workshop on Computational Creativity in Natural Language Generation, CC-NLG@INLG 2017*, Santiago de Compostela, Spain, September 4, 2017, Association for Computational Linguistics, 2017, pp. 29–37. URL: <https://doi.org/10.18653/v1/w17-3904>. doi:10.18653/v1/w17-3904.
- [15] A. Holtzman, J. Buys, L. Du, M. Forbes, Y. Choi, The curious case of neural text degeneration, in: *8th International Conference on Learning Representations, ICLR 2020*, Addis Ababa, Ethiopia, April 26-30, 2020, OpenReview.net, 2020. URL: <https://openreview.net/forum?id=rygGQyrFvH>.
- [16] OpenAI, OpenAI API Reference, 2025. URL: <https://platform.openai.com/docs/api-reference/chat/create>.
- [17] DeepSeek-AI, DeepSeek API Reference, 2025. URL: <https://api-docs.deepseek.com/api/create-chat-completion>.
- [18] C. Arora, A. I. Sayeed, S. A. Licorish, F. Wang, C. Treude, Optimizing Large Language Model hyperparameters for code generation, *CoRR abs/2408.10577* (2024). URL: <https://doi.org/10.48550/>

arXiv:2408.10577. doi:10.48550/ARXIV.2408.10577. arXiv:2408.10577.

- [19] OpenAI, OpenAI Reproducible Outputs Reference, 2025. URL: <https://platform.openai.com/docs/advanced-usage/reproducible-outputs>.
- [20] R. E. Blackwell, J. Barry, A. G. Cohn, Towards reproducible LLM evaluation: Quantifying uncertainty in LLM benchmark scores, 2024. URL: <https://arxiv.org/abs/2410.03492>. arXiv:2410.03492.
- [21] M. M. Zin, K. Satoh, G. Borges, Leveraging LLM for identification and extraction of normative statements, in: J. Savelka, J. Harasta, T. Novotná, J. Mísek (Eds.), Legal Knowledge and Information Systems - JURIX 2024: The Thirty-seventh Annual Conference, Brno, Czech Republic, 11-13 December 2024, volume 395 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2024, pp. 215–225. URL: <https://doi.org/10.3233/FAIA241247>. doi:10.3233/FAIA241247.
- [22] S. A. Hayati, T. Jung, T. Bodding-Long, S. Kar, A. Sethy, J. Kim, D. Kang, Chain-of-instructions: Compositional instruction tuning on Large Language Models, CoRR abs/2402.11532 (2024). URL: <https://doi.org/10.48550/arXiv.2402.11532>. doi:10.48550/ARXIV.2402.11532. arXiv:2402.11532.
- [23] M. Besta, F. Memedi, Z. Zhang, R. Gerstenberger, N. Blach, P. Nyczyk, M. Copik, G. Kwasniewski, J. Müller, L. Gianinazzi, A. Kubicek, H. Niewiadomski, O. Mutlu, T. Hoefler, Topologies of reasoning: Demystifying chains, trees, and graphs of thoughts, CoRR abs/2401.14295 (2024). URL: <https://doi.org/10.48550/arXiv.2401.14295>. doi:10.48550/ARXIV.2401.14295. arXiv:2401.14295.
- [24] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL: <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>.
- [25] J. Wei, M. Bosma, V. Y. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, Q. V. Le, Finetuned language models are zero-shot learners, in: *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*, OpenReview.net, 2022. URL: <https://openreview.net/forum?id=gEZrGCozdqR>.
- [26] OpenAI, OpenAI Fine-tuning Reference, 2025. URL: <https://platform.openai.com/docs/guides/fine-tuning>.
- [27] W. Fungwacharakorn, H. Nguyen, M. M. Zin, K. Satoh, Layer-of-Thoughts Prompting (LoT): Leveraging LLM-based retrieval with constraint hierarchies, CoRR abs/2410.12153 (2024). URL: <https://doi.org/10.48550/arXiv.2410.12153>. doi:10.48550/ARXIV.2410.12153. arXiv:2410.12153.
- [28] A. Witt, A. Huggings, G. Governatori, J. Buckley, Encoding legislation: A methodology for enhancing technical validation, legal alignment and interdisciplinarity, *Artificial Intelligence and Law* 32 (2024) 293–324. URL: <https://rdcu.be/dI0KN>. doi:10.1007/s10506-023-09350-1.
- [29] G. Governatori, F. Olivieri, Unravel legal references in defeasible deontic logic, in: J. Maranhão, A. Z. Wyner (Eds.), *ICAAIL '21: Eighteenth International Conference for Artificial Intelligence and Law, São Paulo Brazil, June 21 - 25, 2021*, ACM, 2021, pp. 69–78. URL: <https://doi.org/10.1145/3462757.3466080>. doi:10.1145/3462757.3466080.