

Provenance-aware judicial draft generation^{*}

Melissa Zorzanelli Costa^{1,*}, Vitor Lugon Cacciari Martinelli², Gabriel Negreiros Piffer², Hygor Oliveira Pancieri², Jean-Rémi Bourguet² and Victor Rodriguez Doncel³

¹Universidade Pitágoras UNOPAR Anhanguera, Rua Antônio Ataíde, 849, Centro - Vila Velha, ES, Brazil

²Universidade Vila Velha, Av. Comissário José Dantas de Melo, 21, Boa Vista, Vila Velha, ES, Brazil

³Universidad Politécnica de Madrid, Campus de Montegancedo, 28660 Boadilla del Monte, Madrid, Spain

Abstract

The growing demand for jurisprudential standardization combined with the increasing volume of decisions issued by the Appeals Panels of the Brazilian Federal Special Courts poses significant challenges to maintaining consistency and transparency in judicial decision-making. This paper presents a provenance-aware framework for the generation of judicial draft decisions to support decision-making based on the identification of representative themes established by the National Panel for Uniformization. We propose to combine automated data collection from real judicial sources, AI-assisted semantic classification of appellate decisions, rule-based determination of admissibility outcomes, and automatic generation of legal drafts in $\text{\LaTeX}$, subsequently compiled into PDF documents. Beyond document generation, the framework incorporates an ontological representation layer that integrates a legal ontology describing the admissibility procedure of Request for Standardizing the Interpretation of a Federal Law (RS) with PROV-O provenance modeling. This semantic layer enables the representation of legal assertions as well as the traceability of the agents, activities, entities, and artifacts involved throughout the generation process. As a result, the framework produces both human-readable legal documents and machine-interpretable RDF knowledge graphs describing their provenance. By promoting transparency, traceability, and interoperability, the proposed framework contributes to the adoption of trustworthy AI-based solutions within the Brazilian Judiciary.

Keywords

Judicial Draft Generation, Large Language Models, Legal Ontologies, Provenance Modeling, PROV-O, Decision-making support

1. Introduction

Following the resurgence of Artificial Intelligence (AI) after the so-called second AI Winter, ontologies became increasingly important as mechanisms for representing structured domain knowledge, enabling reusable conceptual schemas and interoperability across heterogeneous systems [1]. This ontological foundation also played a central role in the emergence of the Semantic Web by supporting shared vocabularies, machine-readable semantics, and standardized knowledge representation across domains [2]. The renewed progress of AI was also driven by advances in Machine Learning and Neural Networks, improvements in computational capacity, the growing availability of large-scale datasets, and the development of new sophisticated algorithms [3]. A major milestone in contemporary AI was the introduction of the Transformer architecture [4], which underpins modern Large Language Models (LLMs) such as GPT, BERT, and Claude. These models have demonstrated remarkable capabilities in natural language understanding and generation, significantly advancing the state of the art in multiple domains. Today, most AI applications operate within the paradigms of Artificial Narrow Intelligence (ANI), focusing on specialized tasks such as recommendation systems, language translation, summarization, and automatic document drafting [5].

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XII: The Tropical Spring of Ontology, co-located with the 16th International Conference on Formal Ontology in Information Systems (FOIS 2026), September 21-22, 2026, Vitória (ES), Brazil

*Corresponding author.

✉ melissa.inf.ufes@gmail.com (M. Z. Costa); vitor.cacciari@uvvnet.com.br (V. L. C. Martinelli); gabriel.piffer@uvvnet.com (G. N. Piffer); hygor.pancieri@uvvnet.com (H. O. Pancieri); jean-remi.bourguet@uvv.br (J. Bourguet); vrodriguez@fi.upm.es (V. R. Doncel)

ORCID 0000-0002-4695-5799 (M. Z. Costa); 0009-0004-8712-9874 (V. L. C. Martinelli); 0009-0008-5262-388X (G. N. Piffer); 0009-0000-9753-3185 (H. O. Pancieri); 0000-0003-3686-1104 (J. Bourguet); 0000-0003-1076-2511 (V. R. Doncel)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Within the legal domain, these advances have renewed interest in intelligent legal systems capable of supporting judicial and legal activities through automated reasoning, information retrieval, and document generation. A legal expert system is a specialized type of expert system that applies artificial intelligence techniques to emulate the reasoning and decision-making capabilities of legal professionals [6]. These systems typically rely on a knowledge base containing legal rules and an inference engine capable of processing and applying such knowledge to specific legal problems. Legal expert systems have long been proposed as a means of addressing the growing complexity and volume of legal information and judicial decisions, a challenge that has become increasingly significant with the expansion of digital legal infrastructures and AI technologies [7]. As part of the broader digital transformation of judicial systems worldwide, the Brazilian Judiciary increasingly employed AI technologies and AI-based systems began to be adopted to address the growing demand for efficiency in handling large volumes of electronic case files and unstructured data [8].

Particularly, the activities of the Appellate Panels (*Turmas Recursais* or TR) and Federal Regional Court (TRF) (*Tribunal Regional Federal* or TRFs) within the scope of the Federal Special Courts (*Juizados Especiais Federais* or JEFs) have consistently generated a significant body of case law, whose uniformity is essential to ensure legal certainty, equality, and the legitimate trust of litigants in the judicial system. However, the growing volume of decisions and the absence of efficient mechanisms for controlling, organizing, and systematizing jurisprudential divergences make it difficult to maintain consistency between regional interpretations and those established by the National Uniformization Panel (*Turma Nacional de Uniformização* or TNU), the court responsible for the uniform interpretation of the JEFs [9]. The lack of automated processes capable of comparing and correlating the issues decided by the TRFs with the precedents consolidated by the TNU reduces the effectiveness of the jurisprudential uniformization system. As a result, potentially conflicting decisions may persist across different regions, undermining the predictability of judicial decisions and leading to unequal treatment in legally similar cases. Despite the fact that the Brazilian Federal Regional Court of the 2nd Region (TRF2) has already officially adopted an AI-based decision-support tool called ApoIA¹, its functionality is exclusively limited to specific tasks such as case summarization, text revision, or headnote generation, and does not support jurisprudence retrieval or case-law uniformization activities. The system primarily assists judges and court staff in judicial case processing, while all legal analyses and judicial decision-making remain under human responsibility. Furthermore, in compliance with Resolution No. 615/2025² of the Brazilian National Council of Justice, the platform incorporates operational logging, internal monitoring, and governance mechanisms to ensure transparency, accountability, and auditability.

In this paper, we present a provenance-aware semantic framework for judicial draft generation in the context of Request for Standardizing the Interpretation of a Federal Law (RS) (*Pedido de Uniformização Nacional* or PU) within the JEFs. The proposed approach combines LLMs, legal knowledge representation, and semantic provenance technologies to support the identification of representative themes established by the TNU, the determination of the corresponding admissibility outcome, and the automatic generation of judicial draft decisions based on templates provided by domain experts. Considering the previously discussed growth in the volume of judicial decisions and the amount of time required to manually verify whether the central issue of a given TR decision (i.e., its procedural matter) has already been addressed within one of the TNU's established themes, as well as the position of the National Council of Justice (CNJ) regarding the adoption of AI-based solutions within the Judiciary³, our proposal seeks to support jurisprudential uniformity, improve procedural efficiency, and assist legal practitioners in identifying precedents and potentially related themes in a more consistent and scalable manner. Beyond the use of LLMs for legal classification and document generation, the framework incorporates an ontological representation layer that combines a legal ontology representing the admissibility procedure of RSs [10, 11] with PROV-O provenance modeling [12, 13]. While the legal ontology provides a structured representation of the legal concepts and decision outcomes involved in

¹<https://www.trf2.jus.br/trf2/consultas-e-servicos/apoia-inteligencia-artificial>

²<https://atos.cnj.jus.br/files/original1555302025031467d4517244566.pdf>

³<https://atos.cnj.jus.br/atos/detalhar/3613>

the appellate process, PROV-O captures the agents, activities, entities, and temporal information associated with the generation workflow. As a result, the framework produces both human-readable judicial drafts, represented in $\text{\LaTeX}$ and PDF formats, and machine-interpretable RDF knowledge graphs capable of supporting traceability, auditability, and explainability. This combination of LLM-based decision support, legal ontologies, and semantic provenance aligns with the broader vision of neuro-symbolic AI, which seeks to combine the learning capabilities of data-driven models with the transparency, structure, and interpretability provided by symbolic representations [14].

The remainder of this paper is organized as follows. Section 2 presents the proposed framework, describing the overall workflow for judicial decision analysis. Section 3 details the implementation of the framework, including data collection, LLM-assisted classification, draft generation, and ontology-based provenance modeling. Section 4 discusses related work on legal AI, judicial document generation, legal ontologies, and provenance-aware systems. Finally, Section 5 concludes the paper and outlines directions for future research.

2. Proposal

The proposed framework aims to support the generation of judicial draft decisions to support decision-making for Request for Standardizing the Interpretation of a Federal Law (RS) (PUs) by combining LLMs, representative legal themes established by the National Uniformization Panel (TNU), legal ontologies, and provenance-aware semantic technologies. The framework operates as a multi-stage pipeline composed of: (i) the extraction of real judicial decisions and representative themes from judicial public brazilian data sources; (ii) the semantic classification of judicial decisions according to TNU representative themes using an LLM; (iii) the determination of the appropriate admissibility outcome of the RSs based on the identified theme, its procedural status, and its consonance or dissonance with the analyzed decision; (iv) the automatic generation of judicial draft decisions in $\text{\LaTeX}$ and PDF formats, based on templates provided by domain experts and currently used in real judicial practice; and (v) the production of RDF knowledge graphs that combine legal-domain representations with PROV-O provenance information, enabling traceability and auditability throughout the document-generation process.

In this article, we decided to consider three possible legal outcomes that, according to the norm, may result from the admissibility assessment of a Request for Standardizing the Interpretation of a Federal Law (RS) (The admissibility analysis process of an RS appeal was modeled using BPMN and subsequently validated by domain specialists [15]). Such assessment must be conducted sequentially and requires the evaluation of representative themes and binding precedents established by the TNU.

More specifically, the approach assists judges and judicial clerks in identifying whether an RS is associated with a representative TNU theme and whether that theme has already reached a final and unappealable status. If the representative theme is still pending judgment, the pipeline supports the generation of a draft decision suspending the proceeding until the TNU issues a final ruling. Conversely, if a final ruling exists, the pipeline assists in assessing whether the challenged decision is in consonance or dissonance with the representative theme.

According to the result of this assessment, the pipeline automatically generates a draft decision corresponding to one of the three legal outcomes considered in this work: suspension of the appeal, not entertain the appeal, or granting the RS to revoke the challenged decision and align it with the binding TNU precedent (TNU theme).

2.1. About the data extraction

Our benchmark was extracted from real decisions issued by the Appeals Panels (TRs) of the Federal Regional Court of the 2nd Region (TRF2) from Espirito Santo (1ªTRES and 2ªTRES), in Brazil, through the court’s jurisprudence consultation portal⁴ by using the *Selenium* library. A total of 2,000 of the most

⁴https://eproc.trf2.jus.br/eproc/externo_controlador.php?acao=jurisprudencia@jurisprudencia/listar_resultados


recent judicial decisions (*acórdãos*) available at the time of extraction (April 2, 2025) were collected. The following information regarding each decision was considered relevant for subsequent analysis: Case Class; Type of Judgment; Subject(s); Jurisdiction; Original Rapporteur; Filing Date; Judgment Date; and Case Judicial Process Number. In addition, the full-text version (*Inteiro Teor*) of each decision was extracted and stored in PDF format.

Subsequently, information regarding representative themes (*temas representativos*) was collected from the TNU website⁵. The following information concerning the representative themes was extracted for subsequent analysis: Theme Number; Theme Status; Branch of Law; Legal Question Submitted for Judgment; Established Thesis; Case Number; Date of the Referral Decision; Rapporteur; Judgment Date; Date of Publication of the Appellate Decision; and Final and Unappealable Decision Date. The full-text document of the referral decision (*Inteiro Teor da Decisão de Afetação*) was also extracted and stored. In total, 378 representative themes were collected.

2.2. About the comparison and categorization

Our framework compares the aforementioned decisions issued by the TRs of the TRF2 and categorizes them according to the foregoing representative themes of the TNU by using an LLM. The comparison process followed the criterion that the subject matter discussed in the full text of the TRF2 decision should first be compared against the LEGAL QUESTION SUBMITTED FOR JUDGMENT associated with each extracted TNU representative theme. As a result of this first stage, the model was instructed to return either the corresponding theme number, when a representative theme matching the subject of the decision was identified, or the label NO THEME when no suitable match was not found. Subsequently, when a theme was identified, the model compared the full text of the TRF2 decision with the ESTABLISHED THESIS associated with the selected theme, whenever such a thesis existed. In this case, the model determined whether the judicial decision and the established thesis were in consonance (alignment) or dissonance (disagreement). All comparison results were stored together with a unique identifier for each judicial decision, the TNU representative theme that best matched the subject of the decision, and an indication of whether the corresponding established thesis was in consonance or dissonance with the analyzed decision. Subsequently, the validity of each analysis must be assessed. LLM models may exhibit failures, such as not responding to the requested task, producing outputs that do not comply with the required formatting, or even hallucinating and providing incorrect information [16].

Figure 1: LLM Legal analyses classification workflow




⁵<https://www.cjf.jus.br/cjf/corregedoria-da-justica-federal/turma-nacional-de-uniformizacao/temas-representativos>

For this reason, the analyses returned by the model were classified according to the following workflow depicted in Figure 1: analyses were considered *Invalid* when they failed to comply with the requested output format; *Incomplete* when they failed to indicate whether the judicial decision and the established thesis were in consonance or dissonance in cases where a thesis had already been established; and *Valid* when the response adhered to the required formatting and provided complete information.

2.3. About the documents generation

When an analysis is classified as valid and a corresponding TNU representative theme is identified for a judicial decision, the decision draft on the admissibility of the appeal (*minuta de decisão de admissibilidade recursal*) can be generated to determine the procedural path of a case involving a RS. Depending on the outcome of the analysis, three procedural actions may be taken regarding this kind of request: i- to suspend the appeal (*sobrestamento*); ii- to not entertain the appeal (*negação de seguimento*); or iii- appeal RS Granted To Revoke Decision of the Appeal Panel (TRES) judgment (*adequação do julgado*).

Figure 2: LLM Legal decisions workflow




As depicted in Figure 2, for each of these actions, a $\text{\LaTeX}$ template file was created to allow the insertion of data related to the identified TNU representative theme and subsequent compilation into a PDF document containing a draft of the legal decision. A request should be suspended when i- the associated TNU representative theme has not yet been adjudicated; or ii- when it has been adjudicated but has not yet reached final and unappealable status *trânsito em julgado*. A request (appeal) should be not entertained further processing when the representative theme has reached final status and the appellate decision is in consonance with the thesis of this theme. A request should be granted to revoke decision when the representative theme has reached final status and the appellate decision is in dissonance with the thesis of the TNU Theme.

2.4. About the production of knowledge graphs

In addition to the draft generation, our framework incorporates an ontological layer aimed at improving the transparency, traceability, and interoperability of the generated artifacts. This layer combines a domain ontology [10, 11] developed in OntoUML [17] and serializable in OWL [18] describing the admissibility procedure with a provenance ontology [12, 13] capable of capturing the generation process of LLM-assisted legal artifacts.

The domain ontology proposed in [10] was slightly extended to support the proposed framework. Specifically, three concepts that were not required in the original admissibility model were introduced: *THEME*, representing the representative themes established by the TNU; *THESIS*, representing the binding

legal thesis associated with a representative theme; and UNAPPEALABLE DECISION, representing the final decision concluding the adjudication of a representative theme. These concepts are used to represent the legal assertions produced by the classification stage and the resulting judicial decisions inferred from the identified TNU representative themes, the assessment of consonance or dissonance between the request and the corresponding established thesis, or the procedural status of the representative theme. The resulting ontological model with its core legal concepts and their relationships throughout the admissibility procedure is illustrated in Figure 3.



To complement the legal representation, our framework employs the PROV-O ontology [12, 13] to capture the provenance of the generated artifacts. This provenance layer explicitly records the agents, activities, entities, and timestamps involved in the generation process of the decision draft from the analyzed request.

As a result, the system produces both a consolidated knowledge graph describing each pipeline execution and individual provenance graphs associated with each generated draft. The ontology contributes to transparency by making the legal concepts underlying each generated draft explicit instead of leaving them embedded only in textual documents. It contributes to traceability because each generated artifact is semantically connected to the representative theme, legal thesis, judicial decision, and provenance records that led to its production. Finally, it improves interoperability by representing

legal knowledge using OWL and RDF vocabularies that can be queried, exchanged, and integrated with other Semantic Web applications. Thus, by explicitly representing how each draft was produced, our framework supports explainability, auditability, and trustworthiness in AI-assisted legal decision support.

3. Implementation

The implemented system⁶ consists of a command-line Python pipeline organized into modules with well-defined responsibilities and each component is responsible for a specific stage of the workflow. The command-line execution interface is concentrated in `app/cli`, workflow configuration and pipeline orchestration are handled in `app/core`, domain abstractions are defined in `app/domain`, and legal reasoning is implemented in `app/services`. Specifically, data acquisition is handled by `app/services/collectors.py`, decision classification is performed by `app/services/analysis.py` and `app/services/gemini.py`, appellate decision rules are implemented in `app/services/actions.py`, judicial draft generation is managed by `app/services/documents.py`, and semantic graph construction and provenance representation are provided by `app/services/semantic.py`. This organization enabled the project to be structured as a data pipeline composed of inspectable and verifiable intermediate artifacts, thereby facilitating validation, reproducibility, and process transparency and provides intermediate outputs to ensure traceability across processing stages.

3.1. Data Collection

The data collection process was formalized to support multiple operational modes. The *sample* mode was designed for local validation without network dependencies, the *import* mode supported the ingestion of locally available datasets, including compressed ZIP packages containing internal files and, the *live* mode enabled direct retrieval from TNU and TRF2 data sources. The *live* mode supports effective pagination and structured parsing of HTML tables with an accurate parser for metadata extraction, including case class, jurisdiction, original rapporteur, and case number. Additionally, an optional fallback mechanism based on browser automation using *Playwright* was incorporated to support scenarios in which data sources relied heavily on JavaScript-based rendering.

3.2. Analysis Layer

The system was designed to support two complementary processing approaches. The first, operating locally, relies on contextual similarity and text normalization techniques to mitigate encoding-related noise and improve the consistency of the matching process between the subjects of judicial decisions and TNU representative themes. The second approach leverages the Gemini API to perform semantic classification and produce structured outputs. The API call received as input parameters: (i) the reference to the file containing information about the TNU representative themes, (ii) the reference to the PDF file containing the full text of the TRF2 decision under analysis, and (iii) a set of instructions specifying how the categorization should be performed. The system explicitly distinguishes among the alternatives SUSPENDED RS, NOT ENTERTAINED RS, RS GRANTED TO REVOKE DECISION, and NO ACTION, while restricting document generation to cases in which a corresponding procedural decision is applicable.

3.3. Draft Generation

Judicial draft documents were produced in $\text{\LaTeX}$ to ensure typographic control and alignment with the desired legal document format. The pipeline was subsequently extended with automatic PDF compilation through either *Tectonic* or *pdf $\text{\LaTeX}$* . Mechanisms supporting reproducible document generation were also incorporated through metadata control using `SOURCE_DATE_EPOCH` and `FORCE_SOURCE_DATE`, with

⁶<https://github.com/LabDeborah/judicial-draft-decisions-project>

the objective of reducing variations among equivalent artifacts generated across different executions. Figure 5 presents an example of a judicial draft generated by the proposed framework recommending the suspension of a RS. In this case, the system identified a similarity between the *Appellate Decision on SFCA* extracted and TNU Theme 172 as shown in Figure 4, which had already been selected for uniformization but had not yet reached a final and unappealable decision (*trânsito em julgado*). According to the admissibility rules⁷, the appropriate procedural action is therefore to suspend the appeal. Based on this classification, a draft decision to suspend RS was automatically generated (*sobrestamento*) representing an artifact for future decision-making.

Figure 4: Representation of Theme 172 on the official TNU jurisprudence portal.

Tema	172	Situação do tema	Sobrestado - Tema 999/STJ (RE 1276977 e RE 1276976 admitidos, no STJ, como representativos da controvérsia)	Ramo do direito	DIREITO PREVIDENCIÁRIO
Questão submetida a julgamento	Saber se é possível ou não aplicação da regra prevista no art. 29, I e II, da Lei 8.213/91, quando mais favorável que a regra de transição prevista no art. 3º da Lei 9.876/99.				
Tese firmada	Tese firmada no Tema 999/STJ: Aplica-se a regra definitiva prevista no art. 29, I e II da Lei 8.213/1991, na apuração do salário de benefício, quando mais favorável do que a regra de transição contida no art. 3o. da Lei 9.876/1999, aos Segurado que ingressaram no Regime Geral da Previdência Social até o dia anterior à publicação da Lei 9.876/1999.				
Processo	Decisão de afetação	Relator (a)	Julgado em	Acórdão publicado em	Trânsito em julgado
PEDILEF 0514224-28.2017.4.05.8013/AL	29/05/2018 26/10/2018 - manutenção da afetação	Juiz Federal Jairo Gilberto Schafer			

Figure 5: Example of a generated suspension draft based on a real template of a legal decision.

Processo n.: 5056703-85.2023.4.02.5101

Verifica-se que a questão objeto da controvérsia jurídica suscitada pelo recurso manejado pela parte recorrente encontra-se afetada ao Tema n. 172 da jurisprudência da TNU – PU no processo n. PEDILEF 0514224-28.2017.4.05.8013/AL, pendente de julgamento (afetado em 29/05/2018 26/10/2018 - manutenção da afetação) –, por meio do qual se elucidará a seguinte questão:

”Saber se é possível ou não aplicação da regra prevista no art. 29, I e II, da Lei 8.213/91, quando mais favorável que a regra de transição prevista no art. 3º da Lei 9.876/99.”

Nesse contexto, determino o **SOBRESTAMENTO** do PU interposto nestes autos – o qual versa sobre a mesma matéria –, na forma do **art. 14, II, b**, da Resolução CJF n. 586/2019 c/c 1.030, III, do CPC, até o julgamento final do representativo da controvérsia, o PU no processo n. PEDILEF 0514224-28.2017.4.05.8013/AL (**Tema n. 172 TNU**).

Intimem-se. Cumpra-se.


⁷https://www.cjf.jus.br/cjf/corregedoria-da-justica-federal/turma-nacional-de-uniformizacao/regimento_interno

3.4. Graph Generation

Our framework implements provenance information as RDF/Turtle graphs generated automatically at the end of each pipeline execution. The provenance model complements the legal ontology by representing legal documents and intermediate results as `PROV:ENTITY` instances, computational processing steps as `PROV:ACTIVITY` instances, and both software components and human participants as `PROV:AGENT` instances. For each analyzed decision, the framework creates a provenance chain linking the original judicial decision, the representative TNU theme identified during classification, the analysis result, and the corresponding judicial draft generated by the system.

Figure 6 presents an excerpt of the RDF provenance graph corresponding to the NOT ENTERTAIN RS draft (*negar seguimento ao Pedido Nacional de Uniformização*) shown in Figure 5. The graph was rendered using the Ontology Visualization tool⁸, which generates Graphviz⁹ visualizations from RDF/OWL data. We have considered the JUDICIAL PROCESS NUMBER as a decision number.

Figure 6: RDF provenance graph of the generated suspension draft (*sobrestamento*).



To validate the practical value of the generated knowledge graphs, competency questions were elicited from judicial law clerks regarding common auditing and explanation tasks. These competency questions were translated into SPARQL queries¹⁰ executed over the generated knowledge graphs. Representative queries include identifying the TNU representative theme used during classification, retrieving the legal thesis supporting a generated draft, determining which software agent participated in the generation process, and retrieving all appellate decisions associated with a specific TNU representative theme.

⁸<https://github.com/fatestigma/ontology-visualization>

⁹<https://graphviz.org>

¹⁰<https://github.com/LabDeborah/judicial-draft-decisions-project/tree/main/cq>

4. Related Work

As judicial and governmental institutions increasingly rely on digital systems, the growing volume of legal and administrative documents has intensified the demand for automated methods capable of supporting classification, organization, and routing tasks. For example, the use of machine learning techniques to automate the classification and routing of electronic petitions has been largely investigated [19, 20]. These works demonstrate how Natural Language Processing (NLP) and Machine Learning techniques can improve the efficiency of legal and administrative workflows by automating repetitive tasks associated with large volumes of textual data.

More recently, the emergence of LLMs has significantly expanded the scope of AI applications in the legal domain. LLM-based systems are capable of generating legal text, answering legal questions, summarizing documents, and supporting legal reasoning activities. The transition from traditional NLP pipelines toward generative AI architectures represents an important shift in Computational Law research, enabling systems not only to classify legal information but also to interact with users and generate complex legal content [21, 22, 23].

Beyond document processing and legal assistance, recent research has increasingly explored the use of AI for legal reasoning, judgment prediction, and judicial decision support. These systems attempt to emulate aspects of legal reasoning by combining subsymbolic techniques with structured legal knowledge. Several studies investigate the prediction of criminal charges, prison terms, and judicial outcomes using deep learning and transformer-based architectures [24, 25, 26, 27, 28, 29, 30]. Together, these studies illustrate the growing interest in AI systems capable of supporting judicial reasoning processes through the integration of legal knowledge, explainability mechanisms, and predictive models.

As AI systems become increasingly integrated into legal environments, concerns regarding reliability, explainability, fairness, and regulatory compliance have also become central research topics. The lack of human involvement may compromise the courts' ability to consider the specific nuances of each case, making judicial decision-making more mechanized and less sensitive to the individual contexts of litigants. AI can support tasks such as information screening, precedent identification, and the organization of complex data, allowing judges to dedicate more time to legal analysis and interpretation. For example, in [31], the author discusses the impacts of AI on fair trial guarantees distinguishing between fully automated judicial systems and judge-support models in which AI acts only as an assisting tool.

In high-stakes domains such as law, inaccurate or biased outputs may directly affect judicial decisions and access to justice, making evaluation and validation essential components of legal AI research. In [32, 22], the authors discuss the importance of domain-specific benchmarks and evaluation methodologies for legal AI systems. In particular, [32] presents a comprehensive evaluation framework for LLMs in legal tasks such as similar-case retrieval, legal argument extraction, and legal element recognition, with benchmark questions reviewed by legal experts to ensure reliability and consistency. Similarly, [22] highlights several open challenges involving hallucinations, explainability, bias, privacy, and regulatory compliance in legal-domain LLMs. The survey emphasizes that, despite recent advances, AI systems in the legal domain still face important limitations regarding robustness, transparency, and trustworthiness.

The growing adoption of AI in legal systems has also strongly influenced the Brazilian Judiciary, particularly within broader initiatives related to judicial modernization, digital transformation, and procedural efficiency. In this context, AI technologies have increasingly been explored in the literature as mechanisms for supporting judicial and administrative activities, especially in response to challenges involving excessive caseloads, procedural delays, and operational costs [33, 34, 35, 36]. The studies demonstrate that the adoption of AI in the Brazilian Judiciary is closely associated not only with technological modernization and efficiency gains, but also with ongoing concerns regarding transparency, accountability, human oversight, and the preservation of fundamental legal principles.

Legal ontologies have long played an important role in the intersection between AI and Law, particularly as mechanisms for structuring legal knowledge, enabling semantic interoperability, and supporting reasoning over legal concepts. Early on, ontologies have been established as essential components

for representing legal domains in a machine-readable and logically consistent manner [37]. Ontology became important semantic structures supporting legal knowledge management, information retrieval, and decision support systems, improving interoperability among heterogeneous legal systems and governmental applications [38]. LKIF Core ontology introduced a reference model designed to represent fundamental legal concepts such as norms, actions, roles, legal agents, and procedural relations [39]. In parallel, Akoma Ntoso emerged as a semantic XML standard for the structured representation and interchange of legislative, parliamentary, and judicial documents, enabling interoperable legal document management and semantic annotation across institutions [40]. More recently, efforts have focused on representing legal norms and rules in a way that combines symbolic reasoning with web-based standards like LegalRuleML [41]. This approach is particularly relevant for contemporary AI applications because it enables the explicit representation of legal logic, traceability of normative sources, and interoperability across legal knowledge systems.

As AI systems become increasingly integrated into data-intensive environments, concerns regarding explainability, reproducibility, traceability, and semantic interoperability have motivated the adoption of Semantic Web technologies and provenance models. In this context, ontologies have also emerged as important tools for structuring, documenting, and explaining complex computational workflows and decision-making processes. One of the most influential standards in this area is the PROV Data Model (PROV-DM), which defines a domain-independent provenance representation framework capable of describing how entities, activities, and agents participate in the generation and transformation of data [42]. PROV-O [12] is an OWL2 ontology that allows a mapping of PROV to RDF. Building upon provenance representations, several studies have explored mechanisms for making provenance information more accessible and interpretable to users. In [43], the author investigates methods for automatically generating domain-independent natural language explanations from PROV provenance graphs, demonstrating how structured provenance information can support explainability and human understanding of computational processes. The SPAR (Semantic Publishing and Referencing) Ontologies [44] provide a modular suite of OWL ontologies designed to represent scholarly publishing information in machine-readable RDF form, enabling semantic interoperability and structured representation of bibliographic and publishing data. Similarly, MetaScriptura [45] proposes a general framework for recording, storing, and analyzing provenance information by integrating technologies such as PROV, RDF graphs, ontologies, SPARQL, and blockchain mechanisms. The framework supports semantic querying, traceability, explainability, and trust across heterogeneous application domains. Other works focus specifically on provenance-aware computational workflows. D-PROV [46] extends the W3C PROV model to represent not only retrospective provenance generated during workflow execution, but also the structure and specification of the workflows themselves. This distinction enables more complete representations of computational processes by combining execution traces with workflow descriptions. OntExtract [47] combines NLP pipelines, LLM orchestration, and PROV-O provenance tracking to support reproducible semantic analysis workflows. The framework automatically records provenance information describing which tools, parameters, workflows, and decisions contributed to the generation of analytical results, improving transparency and reproducibility in AI-driven document analysis.

In the context of modern Legal AI, these ontological approaches have become increasingly important due to the rise of machine learning and LLMs. While LLMs provide strong capabilities for legal text generation and classification, they often lack explicit semantic grounding and explainability. Legal ontologies offer a complementary symbolic layer capable of structuring legal concepts, representing provenance, and supporting interpretable reasoning processes. Such integration is particularly relevant for applications involving legal document generation, decision classification, and auditable AI-assisted legal workflows. Recent research has also explored neuro-symbolic approaches as alternatives to purely subsymbolic legal reasoning systems by combining such capabilities with structured symbolic reasoning, improving transparency, explainability, and auditability in legal tasks (see for example [48, 49]).

While the previous studies establish the theoretical foundations of Legal AI, semantic technologies, and provenance modeling, several recent initiatives have also explored the practical application of artificial intelligence within the Brazilian Judiciary. These systems generally address specific stages

of the judicial workflow, such as precedent retrieval, legal document generation, or decision classification. However, they remain focused on isolated functionalities rather than providing an integrated, provenance-aware decision-support framework.

Within the Federal Regional Court of the 2nd Region (TRF2), the *Intelligentia* project [50] assists judges by comparing procedural themes with previously established judicial precedents, supporting decisions regarding whether proceedings should be suspended or continue. This functionality is closely related to the theme-matching component of our proposed pipeline. Another important initiative is the already mentioned in section 1, *ApoIA*, the generative AI platform made available through the Brazilian Judiciary Digital Platform [51]. *ApoIA* supports the generation of judicial drafts and other legal documents using large language models, improving productivity in judicial activities. Similarly, the recently introduced STJ Logos platform employs generative AI to assist the production of judicial decisions within the Superior Court of Justice (STJ) [52]. Like *ApoIA*, its primary objective is to support legal text generation and judicial productivity. Finally, Gregório [53] proposes the classification of judicial decisions according to procedural themes using TRF2 decisions as reference data.

Overall, the reviewed systems demonstrate that the Brazilian Judiciary has made significant advances in applying artificial intelligence to support judicial activities. Nevertheless, existing solutions generally address individual tasks, such as precedent retrieval, judicial text generation, or semantic knowledge representation, without integrating these capabilities into a single auditable workflow.

5. Conclusion

The rapid advancement of Large Language Models has introduced a new paradigm for the development of intelligent legal systems, enabling the automation of tasks that traditionally required substantial human effort in legal analysis and document preparation. In this context, the literature increasingly recognizes that the advancement of AI in legal systems depends not only on improving predictive performance, but also on ensuring accountability, interpretability, and alignment with legal and ethical principles.

This paper presents a framework for the generation of judicial draft decisions based on the identification of representative legal themes (biding precedents) and the application of admissibility rules within the context of the Appellate Panels and the Regional Uniformization Panels of the Brazilian Federal Special Courts. A central contribution of the work lies in the integration of ontological and provenance-based representations into the document-generation workflow providing traceability and auditability mechanisms for the generated outputs, in accordance with the transparency, governance in Brazil, and oversight requirements established by CNJ Resolution No. 615/2025¹¹. By combining a legal ontology describing the admissibility procedure with PROV-O provenance modeling, the framework is capable of representing not only the legal assertions underlying each generated draft but also the agents, activities, entities, and temporal information involved in its production. As a result, the generated artifacts become both human-readable and machine-interpretable, supporting greater transparency and traceability of AI-assisted legal workflows. The incorporation of this semantic layer transformed the project from a document-generation pipeline into a provenance-aware system capable of providing an auditable semantic trail for all generated artifacts. Beyond supporting transparency and traceability, this approach contributes to reproducibility, facilitates future integration with knowledge graphs and semantic-web technologies. This modular organization facilitates maintainability, reproducibility, and future extensions of the framework, while preserving a clear separation between data acquisition, legal analysis, document generation, and semantic provenance management.

While the combination of the legal ontology and PROV-O provides sufficient support for representing legal decisions and their provenance, future work may investigate the incorporation of additional legal vocabularies, richer ontological representations, and semantic reasoning mechanisms capable of further enhancing the expressiveness of the generated knowledge graphs. Moreover, the generated drafts remain subject to human review and validation before any practical use within judicial environments.

¹¹<https://atos.cnj.jus.br/files/original1555302025031467d4517244566.pdf>

Rather than replacing judicial decision-making, the proposed framework is intended to support legal practitioners by providing AI-assisted document generation together with transparent and traceable representations of the underlying decision process.

A further direction concerns the governance layer that sits above provenance. While PROV-O records how each draft was produced, it does not express under which conditions the generated artifacts, and the judicial data consumed to produce them, may be reused or redistributed. Expressing such conditions as machine readable policies, for instance in ODRL [54], would attach permissions, prohibitions and duties to the very entities already identified in the provenance graph, turning an auditable trail into an enforceable one and giving operational form to the governance and accountability requirements of CNJ Resolution No. 615/2025. Portability is a second direction: encoding the admissibility rules in LegalRuleML over documents represented in Akoma Ntoso would allow the pipeline to be retargeted to other jurisdictions with analogous uniformization mechanisms.

Overall, the integration of Large Language Models, legal ontologies, and semantic provenance demonstrates the potential of neuro-symbolic approaches for the development of intelligent legal systems. By combining automated judicial draft generation with a provenance-aware semantic layer, the framework provides an auditable semantic trail linking legal decisions, classification activities, AI agents, and generated artifacts, thereby contributing to more transparent, explainable, and trustworthy AI-assisted judicial workflows. As future work, we intend to conduct a comprehensive evaluation involving legal experts to assess both the deterministic legal workflows and the AI-assisted classification process. This evaluation will include both quantitative and qualitative analyses, enabling the measurement of classification performance as well as an investigation of the factors that led to correct and incorrect predictions, thereby providing insights into the strengths and limitations of our approach.

Declaration on Generative AI

During the preparation of this work, the authors used GPT-5.5 for grammar, spelling check. After using these services, they reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] N. Guarino, Formal ontology in information systems: Proceedings of the first international conference (FOIS'98), June 6-8, Trento, Italy, volume 46, IOS press, 1998.
- [2] O. Lassila, J. Hendler, T. Berners-Lee, The semantic web, *Scientific american* 284 (2001) 34–43.
- [3] M. Mohri, A. Rostamizadeh, A. Talwalkar, *Foundations of machine learning*, MIT press, 2018.
- [4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, 2017*, pp. 5998–6008.
- [5] J. Guerrero, *Artificial narrow intelligence*, 2023, pp. 163–185. doi:10.1016/B978-0-12-820119-0.00010-8.
- [6] R. E. Susskind, Expert systems in law: A jurisprudential approach to artificial intelligence and legal reasoning, *The modern law review* 49 (1986) 168–194.
- [7] D. H. Berman, C. D. Hafner, The potential of artificial intelligence to help solve the crisis in our legal system, *Commun. ACM* 32 (1989) 928–938. URL: <https://doi.org/10.1145/65971.65972>. doi:10.1145/65971.65972.
- [8] F. Bragança, L. F. da FPG Bragança, *Revolução 4.0 no poder judiciário: levantamento do uso de inteligência artificial nos tribunais brasileiros*, *Revista da Seção Judiciária do Rio de Janeiro* 23 (2019) 65–76.

- [9] N. M. C. d. Araujo, Turma nacional de uniformização: organização, estrutura e funcionamento, Dissertação (mestrado em direito), Universidade Católica de Pernambuco, Recife, Brasil, 2012. URL: <http://tede2.unicap.br:8080/handle/tede/496>, acesso aberto.
- [10] C. Melissa Zorzaneli, G. Guizzardi, J. P. A. Almeida, On capturing legal knowledge in ontology and process models combined, in: E. Francesconi, G. Borges, C. Sorge (Eds.), Legal Knowledge and Information Systems - JURIX 2022: The Thirty-fifth Annual Conference, Saarbrücken, Germany, 14-16 December 2022, Frontiers in Artificial Intelligence and Applications, IOS Press, 2022, pp. 267–272. doi:10.3233/FAIA220478.
- [11] M. Z. Costa, T. B. P. Vieira, J. R. Bourguet, G. Guizzardi, J. P. A. Almeida, Enhancing access to legal data through ontology-based representation: A case study with brazilian judicial appeals, in: CEUR workshop proceedings, volume 3564, CEUR, 2023, pp. 79–92.
- [12] L. Moreau, P. Groth, The prov Ontology, Springer International Publishing, Cham, 2013, pp. 21–38. doi:10.1007/978-3-031-79450-6_3.
- [13] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, Prov-o: The prov ontology (2013).
- [14] B. P. Bhuyan, A. Ramdane-Cherif, R. Tomar, T. P. Singh, Neuro-symbolic artificial intelligence: a survey, Neural Comput. Appl. 36 (2024) 12809–12844. doi:10.1007/S00521-024-09960-Z.
- [15] M. Z. Costa, D. D. de Souza, R. S. S. Guizzardi, J. P. A. Almeida, Can business process modeling support legal decision-making in practice?, in: 2025 27th International Conference on Business Informatics (CBI), IEEE, 2025, pp. 21–31.
- [16] H. Alkaissi, S. I. McFarlane, Artificial hallucinations in chatgpt: implications in scientific writing, Cureus 15 (2023).
- [17] G. Guizzardi, Ontological foundations for structural conceptual models, Phd thesis - research ut, graduation ut, University of Twente, 2005.
- [18] J. P. A. Almeida, G. Guizzardi, T. P. Sales, C. M. Fonseca, gufo: A gentle foundational ontology for semantic web knowledge graphs, CoRR abs/2603.20948 (2026). doi:10.48550/ARXIV.2603.20948. arXiv:2603.20948.
- [19] F. Sun, Y. Zuo, Autonomous classification and decision-making support of citizen e-petitions based on bi-lstm-cnn, Mathematical Problems in Engineering 2022 (2022) 9451108. doi:<https://doi.org/10.1155/2022/9451108>.
- [20] D. Buryakov, M. Kovács, U. Serdült, V. V. Kryssanov, A multi-label classifier for online petition systems, in: H. Liao, D. Duenas-Cid, M. A. Macadar, F. Bernardini (Eds.), Proceedings of the 25th Annual International Conference on Digital Government Research, DGO 2024, Taipei, Taiwan, June 11-14, 2024, ACM, 2024, pp. 156–164. doi:10.1145/3657054.3657250.
- [21] C. V. Chien, M. Kim, A. Raj, R. Rathish, How generative ai can help address the access to justice gap through the courts, Loyola of Los Angeles Law Review (2024) 1–63. URL: <https://ssrn.com/abstract=4719217>, forthcoming.
- [22] C. He, H. Hu, Y. Li, H. Zhang, Q. Zhang, A survey of large language models for legal tasks: Progress, prospects and challenges, Comput. Sci. Rev. 60 (2026) 100906. doi:10.1016/J.COSREV.2026.100906.
- [23] Z. Zhou, K. Yu, S. Tian, X. Yang, J. Shi, P. Song, Y. Jin, L. Guo, Y. Li, Lawgpt: Knowledge-guided data generation and its application to legal LLM, CoRR abs/2502.06572 (2025). doi:10.48550/ARXIV.2502.06572.
- [24] B. Luo, Y. Feng, J. Xu, X. Zhang, D. Zhao, Learning to predict charges for criminal cases with legal basis, in: M. Palmer, R. Hwa, S. Riedel (Eds.), Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Copenhagen, Denmark, 2017, pp. 2727–2736. doi:10.18653/v1/D17-1289.
- [25] Z. Hu, X. Li, C. Tu, Z. Liu, M. Sun, Few-shot charge prediction with discriminative legal attributes, in: E. M. Bender, L. Derczynski, P. Isabelle (Eds.), Proceedings of the 27th International Conference on Computational Linguistics, Association for Computational Linguistics, Santa Fe, New Mexico, USA, 2018, pp. 487–498.
- [26] H. Chen, D. Cai, W. Dai, Z. Dai, Y. Ding, Charge-based prison term prediction with deep gating

- network, in: K. Inui, J. Jiang, V. Ng, X. Wan (Eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Association for Computational Linguistics, Hong Kong, China, 2019, pp. 6362–6367. doi:10.18653/v1/D19-1667.
- [27] L. Kang, J. Liu, L. Liu, D. Ye, Label definitions augmented interaction model for legal charge prediction, in: *Advances in Information Retrieval: 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 – April 1, 2021, Proceedings, Part I*, Springer-Verlag, Berlin, Heidelberg, 2021, p. 270–283. doi:10.1007/978-3-030-72113-8_18.
- [28] S. K. Nigam, A. Deroy, S. Maity, A. Bhattacharya, Rethinking legal judgement prediction in a realistic scenario in the era of large language models, in: N. Aletras, I. Chalkidis, L. Barrett, C. Goanță, D. Preoțiuc-Pietro, G. Spanakis (Eds.), *Proceedings of the Natural Legal Language Processing Workshop 2024*, Association for Computational Linguistics, Miami, FL, USA, 2024, pp. 61–80. doi:10.18653/v1/2024.nllp-1.6.
- [29] W. Yuan, K. Song, Z. Jiang, J. Cao, Y. Zhang, C. Liu, J. Lin, J. Zhang, K. Kuang, X. Liu, A multi-agent framework with legal event logic graph for multi-defendant legal judgment prediction, *Inf. Process. Manag.* 63 (2026) 104319. doi:10.1016/J.IPM.2025.104319.
- [30] W. Shi, Q. Wang, Y. Pan, Q. Liu, K. Zhao, Legalchainreasoner: A legal chain-guided framework for criminal judicial opinion generation, *CoRR abs/2509.00783* (2025). doi:10.48550/ARXIV.2509.00783.
- [31] P. M. Nowotko, AI in judicial application of law and the right to a court, in: J. Watróbski, W. Salabun, C. Toro, C. Zanni-Merk, R. J. Howlett, L. C. Jain (Eds.), *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 25th International Conference KES-2021, Virtual Event / Szczecin, Poland, 8-10 September 2021*, Procedia Computer Science, Elsevier, 2021, pp. 2220–2228. doi:10.1016/J.PROCS.2021.08.235.
- [32] H. Li, Y. Chen, Q. Ai, Y. Wu, R. Zhang, Y. Liu, Lexeval: A comprehensive chinese legal benchmark for evaluating large language models, in: A. Globersons, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. M. Tomczak, C. Zhang (Eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024.
- [33] M. Da Costa Abreu, B. Silva, A critical analysis of 'law 4.0': The use of automation and artificial intelligence and their impact on the judicial landscape of brazil, *Revista de Direitos Fundamentais e Tributação* 1 (2020) 1–16. doi:10.47319/rdft.v1i3.30.
- [34] C. Rocha, J. Carvalho, Artificial intelligence in the judiciary: Uses and threats (short paper), in: M. Janssen, C. Csáki, M. R. Johannessen, R. Krimmer, T. J. Lampoltshammer, I. Lindgren, E. N. Loukis, U. Melin, P. Parycek, G. V. Pereira, M. P. R. Bolívar, G. Schwabe, E. Tambouris, J. Ubacht (Eds.), *Proceedings of Ongoing Research, Practitioners, Posters, Workshops, and Projects. International Conference EGOV-CeDEM-ePart 2022, Hybrid Event, Linköping University, Linköping, Sweden, September 6-8, 2022*, CEUR Workshop Proceedings, CEUR-WS.org, 2022.
- [35] A. E. N. P. d. Mattos, L. V. Curto, M. S. Mussallam, Inteligência artificial e o direito digital, *Revista Políticas Públicas & Cidades* 13 (2024) e1201. doi:10.23900/2359-1552v13n2-186-2024.
- [36] P. Büttner, Precedentes judiciais, inteligência artificial e a (in)dispensável atividade interpretativa do juiz, *Revista da ESMESC* 31 (2024) 199–224. doi:10.14295/revistadaesmesec.v31i37.p199.
- [37] T. J. M. Bench-Capon, P. R. S. Visser, Ontologies in legal information systems: The need for explicit specifications of domain conceptualisations, in: J. Zeleznikow, D. Hunter, L. K. Branting (Eds.), *Proceedings of the Sixth International Conference on Artificial Intelligence and Law, ICAIL '97, Melbourne, Victoria, Australia, June 30 - July 3, 1997*, ACM, 1997, pp. 132–141. doi:10.1145/261618.261646.
- [38] T. M. van Engers, A. Boer, J. Breuker, A. Valente, R. Winkels, Ontologies in the legal domain, in: H. Chen, L. Brandt, V. Gregg, R. Traunmüller, S. S. Dawes, E. H. Hovy, A. Macintosh, C. A. Larson (Eds.), *Digital Government: E-Government Research, Case Studies, and Implementation, Integrated Series In Information Systems*, Springer, 2008, pp. 233–261. doi:10.1007/978-0-387-71611-4_13.

- [39] R. Hoekstra, J. Breuker, M. D. Bello, A. Boer, The LKIF core ontology of basic legal concepts, in: P. Casanovas, M. A. Biasiotti, E. Francesconi, M. Sagri (Eds.), *Proceedings of the 2nd Workshop on Legal Ontologies and Artificial Intelligence Techniques* June 4th, 2007, Stanford University, Stanford, CA, USA, CEUR Workshop Proceedings, CEUR-WS.org, 2007, pp. 43–63.
- [40] G. Barabucci, L. Cervone, M. Palmirani, S. Peroni, F. Vitali, Multi-layer markup and ontological structures in akoma ntoso, in: P. Casanovas, U. Pagallo, G. Sartor, G. Ajani (Eds.), *AI Approaches to the Complexity of Legal Systems. Complex Systems, the Semantic Web, Ontologies, Argumentation, and Dialogue - International Workshops AICOL-I/IVR-XXIV Beijing, China, September 19, 2009 and AICOL-II/JURIX 2009, Rotterdam, The Netherlands, December 16, 2009 Revised Selected Papers, Lecture Notes in Computer Science*, Springer, 2009, pp. 133–149. doi:10.1007/978-3-642-16524-5_9.
- [41] M. Palmirani, G. Governatori, A. Rotolo, S. Tabet, H. Boley, A. Paschke, Legalruleml: Xml-based rules and norms, in: F. Olken, M. Palmirani, D. Sottara (Eds.), *Rule-Based Modeling and Computing on the Semantic Web, 5th International Symposium, RuleML 2011- America, Ft. Lauderdale, FL, Florida, USA, November 3-5, 2011. Proceedings, Lecture Notes in Computer Science*, Springer, 2011, pp. 298–312. doi:10.1007/978-3-642-24908-2_30.
- [42] K. Belhajjame, R. B'Far, J. Cheney, S. Coppens, S. Cresswell, Y. Gil, P. Groth, G. Klyne, T. Lebo, J. McCusker, et al., *Prov-dm: The prov data model*, W3C Recommendation 14 (2013) 15–16.
- [43] D. P. Richardson, *The domain agnostic generation of natural language explanations from provenance graphs*, Ph.D. thesis, University of Southampton, UK, 2018. URL: <https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.759231>.
- [44] S. Peroni, *The semantic publishing and referencing ontologies*, in: *Semantic web technologies and legal scholarly publishing*, Springer, 2014, pp. 121–193.
- [45] M. J. Israel, *MetaScriptura: A General Data Provenance Framework*, Ph.D. thesis, Santa Clara University, 2022.
- [46] P. Missier, S. C. Dey, K. Belhajjame, V. Cuevas-Vicentín, B. Ludäscher, *D-PROV: extending the PROV provenance model with workflow structure*, in: A. Meliou, V. Tannen (Eds.), *5th Workshop on the Theory and Practice of Provenance, TaPP'13, Lombard, IL, USA, April 2-3, 2013*, USENIX Association, 2013.
- [47] C. B. Rauch, H. W. Choi, M. Kelly, Ontextract: PROV-O provenance tracking for document analysis workflows, in: H. Alhoori, J. S. Downie, M. Kelly, S. R. Choudhury, I. Frommholz, J. Chen (Eds.), *ACM/IEEE Joint Conference on Digital Libraries, JCDL 2025, Dekalb, IL, USA, December 15-19, 2025*, IEEE, 2025, pp. 249–252. doi:10.1109/JCDL67857.2025.00038.
- [48] C. Melissa Zorzanelli, D. F. Robson, T. B. P. Vieira, J. Bourguet, G. Guizzardi, J. P. A. Almeida, Automated semantic annotation pipeline for brazilian judicial decisions, 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*, *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2024, pp. 226–238. doi:10.3233/FAIA241248.
- [49] S. Sójka, W. Kowalczyk, Accurate legal reasoning at scale: Neuro-symbolic offloading and structural auditability for robust legal adjudication, in: *The 64th Annual Meeting of the Association for Computational Linguistics – Industry Track*, 2026. URL: <https://openreview.net/forum?id=i15v9zgp6i>.
- [50] M. Paumgartten, J. S. C. Cavalcante, R. de Freitas Monteiro, C. M. de Bacco Ramalho, C. M. P. de Araújo, G. C. Branco, *Panorama do uso de inteligência artificial nos tribunais do estado do rio de janeiro*, *Revista Ibmec de Direito-ISSN 3085-704X* 1 (2025).
- [51] G. R. S. Bertoni, *O impacto da inteligência artificial no âmbito jurídico: obsolescência da advocacia e adaptação para uma nova era* (2025).
- [52] L. P. F. Vasconcelos, *Inteligência artificial no judiciário: analisando risco, impactos e regulamentação* (2026).
- [53] F. da Silva Gregorio, *Análise de processos judiciais via processamento de linguagem natural* (2024).
- [54] B. Esteves, V. Rodríguez-Doncel, Analysis of ontologies and policy languages to represent information flows in GDPR, *Semantic Web* 15 (2024) 709–743. doi:10.3233/SW-223009.