ANPC Panel: A Hybrid Framework for Legal Information Extraction, Transparency, and Human-on-the-Loop Validation

Vitória S. Santos[1], **Francisco Bortolanza**[1], **Carina F. Dorneles**[1]

[1]Departamento de Informática e Estatística
Universidade Federal de Santa Catarina (UFSC)
Caixa Postal 476 – 88.040-900 – Florianópolis – SC – Brazil

vit2santoss@gmail.com, bortolanza.francisco@gmail.com, dorneles@gmail.com

Abstract. *Legal document analysis plays an important role in supporting institutional oversight and public governance; however, manual extraction of information from large document collections is time-consuming and difficult to scale. This paper presents the ANPC Panel, a hybrid framework that combines Large Language Models (LLMs), retrieval-based extraction, and complementary rule-based strategies to transform Civil Non-Prosecution Agreements (ANPCs) into structured representations for institutional monitoring and large-scale analysis. Beyond automated extraction, the system incorporates algorithmic transparency and Human-on-the-Loop validation mechanisms, allowing users to inspect, contextualize, and audit AI-generated outputs. By combining automated information extraction, transparent analytics, and human supervision, the proposed approach contributes to strengthening efficiency, accountability, and evidence-based decision-making in public-sector environments.*

1. Introduction

In Brazil, a public civil action is "an action aimed at protecting diffuse or collective interests, holding accountable those who cause damage to the protected assets."[1] . Among the various types of civil actions, there is the civil action for administrative improbity, which refers to actions brought as a result of an "act of administrative improbity that violates the principles of public administration or willful action or omission that violates the duties of honesty, impartiality, and legality."[2]. In December 2019, Law 13,964, the Anti-Crime Law, was published, establishing that civil actions for administrative misconduct "allow the execution of a civil non-prosecution agreement" [Júnior 2021], the ANPC (*Acordo de Não Persecução Cível* in Portuguese, or Civil Non-Prosecution Agreement). The ANPC, therefore, is a negotiation instrument provided for in Brazilian law that allows for the consensual termination of proceedings related to acts of administrative misconduct.

An ANPC is a highly important instrument for both the Public Prosecutor's Office (MP - *Ministério Público*, in Portuguese) and society. For the Public Prosecutor's Office, the agreements validate its actions and allow for more strategic management in combating administrative misconduct. For society, the ANPC ensures a faster and more transparent recovery of public funds, avoiding the costs and delays of lengthy legal battles. By ensuring more visible and immediate justice, the tool also allows for a more intelligent use

[1]https://www.cnmp.mp.br/portal/glossario/8242-acao-civil-publica
[2]https://www.planalto.gov.br/ccivil_03/leis/l8429.htm

of public resources, as the Public Prosecutor's Office can focus on more complex and impactful cases.

For the Public Prosecutor's Office to strategically target its actions, it must conduct data-driven analyses regarding the effectiveness and impact of Civil Non-Prosecution Agreements. Such analyses require Business Intelligence (BI) tools capable of supporting institutional monitoring through financial, temporal, geographical, and operational indicators. However, the construction of these analytical environments depends on the availability of structured information extracted from legal documents. Although ANPCs are relatively short documents, relevant information does not follow a fixed structure and may appear in different textual contexts across agreements. Consequently, manual extraction becomes a repetitive and time-consuming process that is difficult to scale, limiting broader institutional analyses and continuous monitoring activities by prosecutors.

To address this challenge, this work presents the ANPC Panel, a hybrid framework designed to transform ANPC documents into structured and actionable information for institutional monitoring. The proposed approach combines Large Language Models (LLMs), retrieval-based strategies, and complementary rule-based procedures to support information extraction from legal documents. Beyond extraction itself, the framework integrates analytical dashboards, algorithmic transparency mechanisms, and Human-on-the-Loop validation interfaces that allow users to inspect and verify generated information against source evidence. By combining automated extraction with transparency and human supervision mechanisms, the ANPC Panel supports large-scale analysis while preserving traceability and reliability requirements associated with institutional use.

2. Related Work

Information extraction from legal documents has been widely investigated to support institutional analysis and decision-making. Traditional approaches commonly rely on supervised architectures and annotated datasets. In the Brazilian context, Fernandes et al. [Fernandes et al. 2022, Fernandes et al. 2020] applied deep learning techniques to extract information from judicial decisions. More recently, Large Language Models (LLMs) have emerged as an alternative due to their zero-shot capabilities. Studies such as Zambrano et al. [Zambrano 2024] and Pereira et al. [Pereira et al. 2024] explored LLM-based legal extraction workflows, while retrieval-enhanced approaches have shown improvements in extraction quality. Aquino et al. [Aquino et al. 2024] demonstrated the impact of retrieval configurations in Brazilian legal documents, and Li et al. [Li et al. 2024] showed that multi-query retrieval strategies can improve coverage and downstream generation performance.

Beyond extraction, analytical infrastructures and Business Intelligence (BI) platforms have increasingly supported transparency and judicial monitoring. Ruschel [Ruschel 2015] highlighted the use of BI dashboards for judicial modernization and decision support, while Lamarão Neto [Neto 2019] and Gonçalves [Gonçalves 2021] presented dashboard-based systems for process management and strategic monitoring in Brazilian courts. More recently, Aquino et al. [Justino et al. 2025] proposed a Business Intelligence platform to support legal data management at the Court of Justice of Maranhão (TJMA), reinforcing the potential of BI systems for legal analytics and institutional support. However, these solutions generally operate over previously structured repositories

rather than automatically extracting information from unstructured legal documents.

Another relevant direction concerns human supervision in AI-assisted systems. Human-in-the-loop approaches have been proposed to increase reliability and transparency in extraction workflows [Gebauer et al. 2023, Schroeder et al. 2025]. Nevertheless, existing studies typically investigate extraction, retrieval, dashboards, or validation mechanisms separately. Limited attention has been given to end-to-end frameworks that integrate automated extraction from unstructured legal documents with analytical dashboards and human validation interfaces. Therefore, this work investigates standard and multi-query RAG strategies combined with LLMs and rule-based extraction mechanisms to transform ANPC documents into structured information integrated into an interactive panel with human-on-the-loop validation support.

3. ANPC Panel Framework

The overall framework adopted in this work is illustrated in Figure 1. The process begins with the acquisition of Civil Non-Prosecution Agreements (ANPCs) from the SIG/MPSC system, where agreements are available in PDF format. These documents constitute a corpus of unstructured legal texts that must be transformed into structured information to support institutional monitoring and data-driven analysis. The proposed workflow converts legal documents into structured and auditable representations through a multi-stage pipeline comprising information extraction, ETL processing, database integration, and interactive dashboard visualization. Beyond enabling large-scale analytical exploration, the workflow was designed to preserve transparency and traceability by allowing users to inspect extraction outputs, access source documents, and validate AI-generated information through Human-on-the-Loop mechanisms integrated into the ANPC Panel.

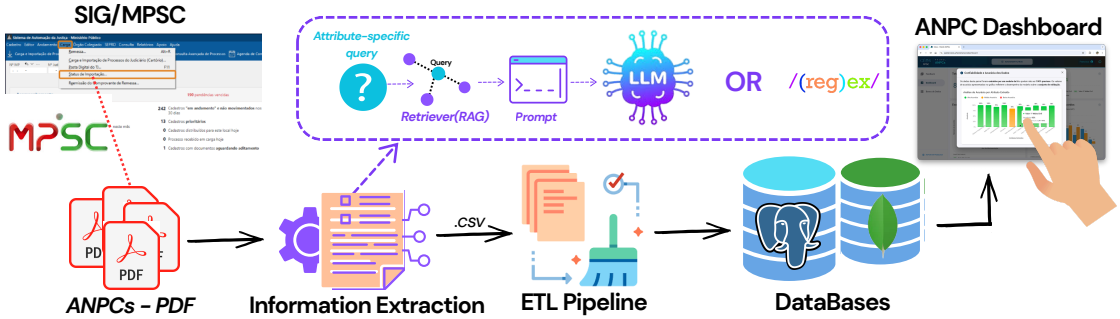

Figura 1. Overview of the proposed ANPC information extraction, ETL, and dashboard architecture.

Information Extraction. To transform ANPC documents into structured information, the proposed workflow combines LLM-based extraction with retrieval mechanisms and complementary regular expression (regex) rules. While most attributes are extracted through the RAG+LLM pipeline, preliminary analyses revealed that some attributes exhibited lower extraction reliability due to formatting variability and document-specific patterns. For these cases, regex rules were incorporated as a complementary strategy for attributes with stable and predictable textual structures.

ETL Pipeline. The extracted outputs are converted into tabular `.csv` files and processed through an ETL pipeline responsible for normalization, standardization, and

validation procedures. This stage guarantees consistency before integration into the storage layer through operations including:

- normalization of monetary values;
- standardization of dates;
- formatting and cleaning of categorical attributes;
- replacement of invalid or missing entries;
- identifier validation and duplicate removal.

Databases. The storage layer comprises two complementary database systems designed to support different categories of information generated throughout the extraction workflow. Structured extracted attributes and their associated metadata are stored in PostgreSQL, enabling aggregation, filtering, and analytical queries. Original ANPC PDF documents are maintained in MongoDB, allowing efficient storage and retrieval of source files. This dual-storage architecture supports document-level traceability by preserving links between extracted values and source documents.

ANPC Dashboard (ANPC Panel). Finally, the processed information feeds the ANPC Dashboard, which provides interactive visualizations and analytical functionalities for institutional monitoring and decision support. The dashboard supports statistical exploration through financial indicators, temporal analyses, and geographic distributions of agreements. In addition, it incorporates transparency mechanisms that communicate extraction reliability and provenance information, as well as feedback functionalities that allow users to inspect source documents and validate extracted values. These features establish a human-on-the-loop workflow in which AI-generated outputs remain subject to institutional oversight and documentary verification.

4. ANPC Information Extraction

This section describes the information extraction workflow adopted to transform Civil Non-Prosecution Agreements (ANPCs) into structured data for integration into the ANPC Panel. Since dashboard indicators, transparency mechanisms, and human validation functionalities depend directly on extraction quality, different extraction strategies and two large language models were investigated before defining the final workflow. The proposed approach combines retrieval mechanisms, large language models (LLMs), and complementary regular expression (regex) rules to generate structured attribute-value representations suitable for ETL processing and dashboard integration.

4.1. Methodology

The dataset used in this study consists of Civil Non-Prosecution Agreements (ANPCs) in PDF format obtained from the institutional SIG system. At the time of collection, the repository contained 824 unique agreements. Since manually annotating the entire corpus would require substantial domain expertise and institutional effort, a probabilistic sampling procedure was adopted following the finite population methodology proposed by [Rea and Parker 2012]. Considering a confidence level of 95% and a target error margin of approximately 10%, a representative sample of 100 ANPCs was selected for experimentation. The selected documents were manually annotated by prosecutors from the Public Prosecutor's Office of Santa Catarina (MPSC), producing a gold-standard dataset used for extraction assessment.

The extraction workflow was implemented using two open-source models available through the Ollama framework: Gemma 3 (27B) and Llama 3.3 (70B) [Team et al. 2025, Grattafiori et al. 2024]. These models were selected because they support schema-constrained generation through `pydantic`, allowing outputs to follow predefined JSON structures and remain compatible with the extraction process. By adopting structured output generation, responses remain uniform across models while reducing extraneous text, hallucinated fields, and formatting inconsistencies, enabling direct conversion into `.csv` representations for database ingestion and dashboard visualization.

For information extraction, three strategies were evaluated:

- *Direct LLM extraction*: the complete ANPC document is processed without retrieval support, serving as a baseline for long-document understanding.
- *Standard RAG*: a single retrieval query identifies relevant document chunks that are appended to the prompt before model inference.
- *Multi-Query RAG*: independent retrieval queries are generated for each target attribute and the retrieved evidence is aggregated into a final prompt containing attribute-specific contextual information.

The extraction strategies were evaluated by comparing model outputs against the gold-standard dataset. For each attribute, extraction accuracy was computed by comparing predicted values with their corresponding annotated references, enabling assessment of retrieval strategies and model configurations. The resulting attribute-level accuracies were later incorporated into the ANPC Panel and exposed through transparency mechanisms to communicate the reliability associated with extracted information.

The attributes considered in this study were defined in collaboration with prosecutors from MPSC, who identified the information required for monitoring Civil Non-Prosecution Agreements. The resulting set of attributes supports financial oversight, temporal monitoring, geographical exploration, and the institutional indicators presented in the ANPC Dashboard:

- `Inquiry Number` (*Número do Inquérito Civil*): identifier associated with the civil inquiry process.
- `Reimbursement Value` (*Valor de Ressarcimento*): monetary value reimbursed to the public treasury.
- `Civil Fine 1`, `Civil Fine 2`, and `Civil Fine 3` (*Valores de Multa Civil 1, 2 e 3*): monetary penalties established in the agreement.
- `Forfeiture Value` (*Valor de Perdimento*): amount associated with asset forfeiture.
- `Movement Date` (*Data do Movimento*): agreement formalization date.
- `District` (*Comarca*): judicial district associated with the agreement.

Although retrieval-based approaches improved extraction quality for most attributes, some fields exhibited highly regular textual structures across documents. In particular, `Inquiry Number`, `Movement Date`, and `District` presented stable formatting patterns with low semantic variability. Consequently, regular expression (regex) rules were incorporated as complementary extraction mechanisms for these attributes, since deterministic pattern matching proved more reliable than semantic retrieval for recurrent textual structures.

4.2. Results

Table 1 summarizes attribute-level extraction accuracy across the evaluated models and extraction strategies. Each row corresponds to an extracted ANPC attribute, while grouped columns represent the evaluated extraction configurations for both large language models (LLMs). The highest value obtained for each attribute is highlighted in bold to facilitate identification of the best-performing extraction configuration.

Attribute	**LLM-only**		**Standard RAG**		**Multi-Query RAG**	
	L	G	L	G	L	G
Inquiry Number	0.38	0.39	0.52	0.48	**0.65**	0.68
Reimbursement Value	0.61	0.62	0.71	0.70	0.82	**0.84**
Civil Fine 1	0.57	0.68	0.72	0.70	0.89	**0.95**
Civil Fine 2	0.63	0.65	0.70	0.76	0.75	**0.90**
Civil Fine 3	0.80	0.76	0.86	0.84	**0.99**	0.98
Forfeiture Value	0.93	0.95	0.86	0.93	0.93	**0.95**
Movement Date	**0.87**	**0.87**	0.43	0.42	0.65	0.86
District	**0.96**	**0.96**	0.92	0.86	**0.96**	**0.96**

L = Llama 3.3 (70B) G = Gemma 3 (27B)

Tabela 1. Attribute-level extraction accuracy across models and extraction strategies.

Results reveal a progressive improvement in extraction quality as retrieval mechanisms become more specialized. Direct LLM extraction presented the lowest performance for most attributes, while Standard RAG reduced irrelevant contextual information and Multi-Query RAG further improved performance through attribute-specific evidence retrieval. The largest gains were observed for financial attributes, suggesting that retrieval mechanisms improve localization of dispersed information within lengthy legal documents.

Based on these observations, the final extraction workflow adopted in the ANPC Panel combines Multi-Query RAG and regex-based extraction. Attributes exhibiting highly regular textual structures, namely `Inquiry Number`, `Movement Date`, and `District`, were extracted through regex rules, resulting in final accuracies of 100%, 98%, and 100%, respectively. The remaining attributes were extracted through Multi-Query RAG, producing the reliability estimates later incorporated into the dashboard transparency mechanisms, including 84% for `Reimbursement Value`, 95% for `Civil Fine 1`, 90% for `Civil Fine 2`, 98% for `Civil Fine 3`, and 95% for `Forfeiture Value`.

5. ANPC Panel Architecture, Transparency, and Human-on-the-Loop Validation

This section presents the system architecture and transparency mechanisms adopted in the ANPC Panel. Since the dashboard supports institutional monitoring of Civil Non-Prosecution Agreements (ANPCs), its design extends beyond conventional visualization requirements and incorporates accountability, traceability, and human supervision mechanisms. The panel was designed not only to provide analytical capabilities, but also to ensure that AI-generated information remains transparent, auditable, and suitable for institutional use.

5.1. System Architecture

The ANPC Panel follows a distributed client-server architecture implemented as a full-stack web application. The frontend was developed using React within a Vite environment, providing a dynamic interface for analytical exploration and validation functionalities. The backend consists of a REST API implemented with FastAPI, responsible for mediating communication between the interface and the storage layer. FastAPI was selected due to its compatibility with the Python ecosystem adopted by the extraction pipeline and its support for asynchronous request handling.

The storage layer comprises two complementary database systems designed to support distinct data access patterns. Structured outputs and extraction metadata are stored in PostgreSQL, enabling filtering, aggregation, and statistical operations required by dashboard indicators. Original PDF documents are maintained in MongoDB, whose document-oriented structure supports efficient storage and retrieval of legal files without schema restrictions.

This dual-database architecture establishes the foundation for traceability within the system. By maintaining structured attributes and source documents independently while preserving links between them, the panel enables extracted values to remain associated with their corresponding source evidence, supporting document-level inspection and auditability. Figure1 illustrates the overall system architecture and its integration with the extraction workflow.

5.2. Transparency and Statistical Analysis

The analytical component of the ANPC Panel consists of a dashboard that aggregates structured information extracted from Civil Non-Prosecution Agreements into statistical indicators designed to support institutional monitoring and strategic decision-making. By transforming unstructured legal documents into queryable information, the system enables large-scale analysis of agreement data, allowing the Public Prosecutor's Office to assess financial, temporal, and geographical aspects of administrative misconduct settlements. From a social perspective, this analytical capability contributes to public governance by facilitating evidence-based oversight and improving access to institutional information.

Since dashboard information originates from AI-based extraction procedures, transparency mechanisms were incorporated to support trustworthy use of generated information in institutional contexts. Figure 2(a) presents the dashboard disclosure interface, which informs users that displayed information was generated through automated extraction procedures and exposes metadata associated with the extraction process. Rather than presenting analytical outputs as isolated values, the system explicitly communicates their provenance, allowing users to contextualize information according to its generation process.

Figure 2(b) illustrates the main dashboard environment and its statistical analysis functionalities. The indicators were defined in collaboration with professionals from the Public Prosecutor's Office to reflect institutional demands associated with monitoring Civil Non-Prosecution Agreements. The dashboard supports multiple analytical dimensions, including agreement statistics, judicialization rates, total financial recovery values, annual agreement evolution, comparisons between judicial and extrajudicial agreements,

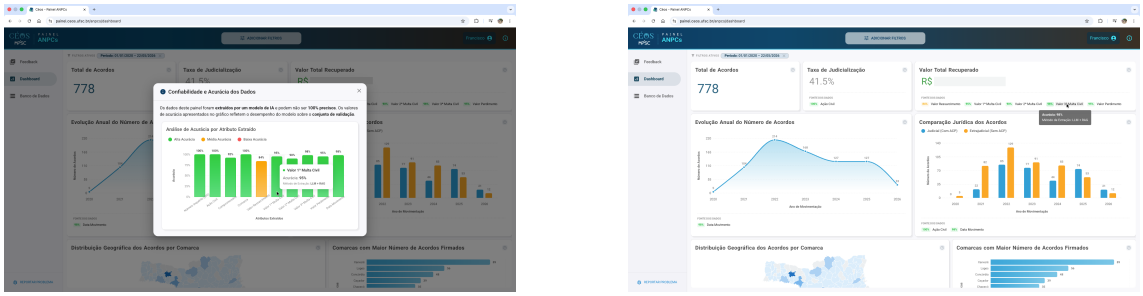

(a) Transparency interface presenting extraction metadata and disclosure information

(b) Dashboard view displaying institutional indicators and provenance panels

Figura 2. Algorithmic transparency mechanisms adopted in the ANPC Panel

and geographical distributions by district and municipality. Users may further explore the data through filters combining agreement identifiers, districts, municipalities, and temporal ranges, enabling flexible multi-dimensional analysis of the ANPC corpus.

In addition to statistical exploration, each analytical visualization is associated with provenance information identifying which extracted attributes contributed to the corresponding indicator and their extraction procedures. By combining statistical analysis with explicit extraction metadata, the ANPC Panel transforms AI-generated outputs into a mechanism for transparent and accountable public governance, ensuring that the algorithmic nature of the underlying information remains visible and interpretable to institutional users.

5.3. Human-on-the-Loop Validation

Although transparency mechanisms provide information regarding the origin and reliability of extracted data, institutional use of AI-generated information also requires mechanisms that allow users to inspect and verify individual extraction results. In public governance contexts, analytical outputs may support decisions with institutional consequences, making document-level verification an important requirement. To address this challenge, the ANPC Panel incorporates a Human-on-the-Loop validation workflow designed to support direct inspection and auditing of extraction outputs.

The Human-on-the-Loop validation process is organized into two complementary interfaces. Figure 3(a) presents the extraction management environment, where processed agreements are organized according to validation status categories, including pending, under validation, and validated records. Each agreement is represented through an extraction card displaying extracted attributes and explicitly highlighting non-identified fields. Additional counters and filtering mechanisms provide users with an overview of validation progress and support structured navigation throughout the corpus.

Figure 3(b) illustrates the validation environment, which presents the original PDF document alongside a structured form containing extracted attributes. Users can inspect source evidence, confirm extracted values, or modify incorrect information before validation. This side-by-side interaction establishes a direct connection between extracted information and documentary evidence, enabling traceable auditing of AI-generated outputs.

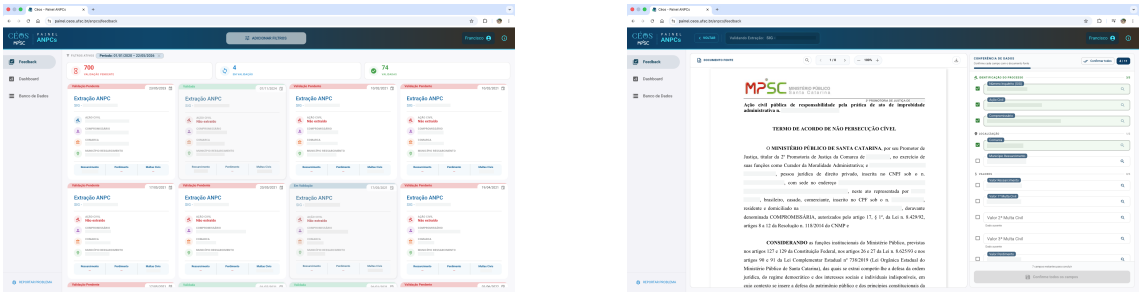

(a) Extraction management interface showing validation status and extracted attribute cards

(b) Validation environment displaying source documents and attribute confirmation forms

Figura 3. Human-on-the-Loop validation workflow adopted in the ANPC Panel

Beyond verification, the Human-on-the-Loop process progressively creates a subset of human-audited records that can support future analyses and increase confidence in institutional use of extracted information. From a broader social perspective, this mechanism reinforces transparency and accountability principles by ensuring that automated extraction and institutional expertise operate as complementary components rather than competing processes.

6. Conclusion and Future Work

This work presented the ANPC Panel, a hybrid framework designed to transform Civil Non-Prosecution Agreements into structured information for institutional monitoring within the Public Prosecutor's Office of Santa Catarina. By combining automated extraction procedures with transparency and Human-on-the-Loop validation mechanisms, the proposed system enables large-scale analysis of legal agreements while preserving traceability and auditability requirements associated with institutional use. For the MPSC, the resulting analytical environment provides a practical mechanism for supporting prosecutors in monitoring ANPCs through centralized access to institutional indicators and structured information that would otherwise remain dispersed across legal documents. By facilitating data-driven analyses and broader visibility over agreement characteristics and outcomes, the system contributes to strengthening transparency, accountability, and institutional oversight.

As future work, we intend to integrate an automated data ingestion workflow directly with the institutional systems of the Public Prosecutor's Office, enabling continuous acquisition and processing of newly generated ANPC documents. Such integration would allow the extraction pipeline and dashboard infrastructure to operate as a continuously updated environment, automatically incorporating new agreements and refreshing analytical indicators as additional documents become available. Future developments also include leveraging newly validated records produced through the Human-on-the-Loop process to improve extraction quality for attributes that still present lower performance. This evolution would transform the proposed workflow from a batch-oriented analytical process into a dynamic monitoring platform, further strengthening transparency, institutional oversight, and data-driven decision support in public governance contexts.

Referências

Aquino, I., dos Santos, M. M., Dorneles, C., and Carvalho, J. T. (2024). Extracting information from brazilian legal documents with retrieval augmented generation. In *Anais Estendidos do XXXIX Simpósio Brasileiro de Bancos de Dados*, pages 280–287, Porto Alegre, RS, Brasil. SBC.

Fernandes, W. P. D., Frajhof, I. Z., da Franca Couto Fernandes de Almeida, G., Rodrigues, A. M. B., Barbosa, S. D. J., Konder, C. N., Nasser, R. B., de Carvalho, G. R., and Lopes, H. C. V. (2022). Extracting value from brazilian court decisions. *Information Systems*, 106:101965.

Fernandes, W. P. D., Silva, L. J. S., Frajhof, I. Z., da Franca Couto Fernandes de Almeida, G., Konder, C. N., Nasser, R. B., de Carvalho, G. R., Barbosa, S. D. J., and Lopes, H. C. V. (2020). Appellate court modifications extraction for portuguese. *Artificial Intelligence and Law*, 28(3):327–360.

Gebauer, M. et al. (2023). A human-in-the-loop approach for information extraction from privacy policies under data scarcity.

Gonçalves, J. P. (2021). Desenvolvimento de painéis bi/eproc de apoio à gestão no tribunal de justiça do tocantins. Technical report, Universidade Federal do Tocantins.

Grattafiori, A., Dubey, A., Jauhri, A., and et al. (2024). The llama 3 herd of models.

Justino, A. F. et al. (2025). A business intelligence platform to support legal data management at the tjma.

Júnior, J. B. T. (2021). Considerações sobre o acordo de não persecução cível – anpc. *Revista Procuradoria Geral do Estado do Mato Grosso - PGE/MS*, 1(17):24.

Li, Z. et al. (2024). Dmqr-rag: Diverse multi-query rewriting for retrieval-augmented generation. *arXiv preprint arXiv:2411.13154*.

Neto, H. L. (2019). O mapeamento das ações de saúde pelo tribunal de justiça do estado do pará. *Revista de Política Judiciária, Gestão e Administração da Justiça*, 5(2):1–17.

Pereira, J., Assumpcao, A., Trecenti, J., Airosa, L., Lente, C., Cléto, J., Dobins, G., Nogueira, R., Mitchell, L., and Lotufo, R. (2024). Inacia: Integrating large language models in brazilian audit courts: Opportunities and challenges.

Rea, L. and Parker, R. (2012). *Designing and Conducting Survey Research: A Comprehensive Guide*. Wiley.

Ruschel, A. J. (2015). Governo eletrônico: Business intelligence para a modernização do judiciário.

Schroeder, N., Jaldi, C., and Zhang, S. (2025). Large language models with human-in-the-loop validation for data extraction.

Team, G., Kamath, A., Ferret, J., Pathak, S., and et al. (2025). Gemma 3 technical report.

Zambrano, G. (2024). Case law as data : Prompt engineering strategies for case outcome extraction with large language models in a zero-shot setting. *Law, Technology and Humans*, 6(3):80–101.