Intelligent Agent for Legal Health Case Automation: A Case Study in the Brazilian Public Sector

L. Rocha[1], B. Amaro[1], V. Furtado[2], N. Duque[1], V. Menezes[1], J. Cláudio[3]

[1]Graphs and Computational Intelligence Laboratory – State University of Ceará (UECE)
Fortaleza – CE – Brazil

{leonardo.sampaio, bonfim.amaro, nathanael.duque, vinicius.menezes}@uece.br

[2]University of Fortaleza (UNIFOR)
Fortaleza – CE – Brazil

vasco@unifor.br

[3]Federal Institute of Ceará (IFCE)
Fortaleza – CE – Brazil

joao.claudio@ifce.edu.br

Abstract. *This work presents the design, implementation, and evaluation of an intelligent agent developed to support the State Health Attorney's Office of Ceará (PROSAÚDE) in managing over 3,000 monthly legal cases related to public health. The agent automates key decisions on legal contestation by interpreting court documents and applying internal legal norms. Built under the Judiciary vertical of Ceará's Chief Scientist Program, the system combines large language models (GPT-4o), regulatory rule sets, and robotic process automation to extract medical/legal data, retrieve external sources (e.g., ANVISA), and generate official dispatches. It ensures consistent application of PROSAÚDE's Internal Ordinance, which defines when appeals are required. Results show 98.4% precision in rule application and 93.8% recall in medication detection. In human-in-the-loop validation, 98.7% of dispatches required no revision. The solution highlights the potential of AI-driven legal automation to improve efficiency and accuracy in high-volume public workflows.*

1. Introduction

The progressive digitalization of public institutions has unfolded in waves, often unevenly across departments. As a result, digital systems have historically evolved in silos, with little or no coordination regarding data interoperability. This fragmented development poses persistent challenges to information exchange—particularly in the presence of Legacy Information Systems (LIS), which remain mission-critical but lack adherence to modern interoperability standards.

In such environments, public servants are frequently burdened with repetitive, manual tasks, including data entry and migration, report generation, backup creation, information retrieval, error correction, and security validation. These activities are not only time-consuming and error-prone but also lead to significant operational inefficiencies and, in some cases, occupational health issues such as musculoskeletal strain.

These challenges are further amplified in the context of digital transformation efforts in government, especially within e-government initiatives. A key enabler of this transformation is Robotic Process Automation (RPA), a technology designed to emulate human interactions with software systems to automate routine tasks. RPA offers a cost-effective and minimally invasive solution—particularly beneficial in scenarios where LIS replacement is impractical or economically prohibitive [Costa et al. 2022].

However, traditional RPA struggles with tasks involving unstructured data, complex decision-making, or semantically rich content. Recent advances in Artificial Intelligence (AI), especially with Large Language Models (LLMs), have led to the emergence of Intelligent Robotic Process Automation (IRPA) [Ng et al. 2021]. These systems extend traditional RPA by incorporating cognitive capabilities such as natural language understanding, semantic search, context-aware reasoning, and adaptive decision-making. IRPA enables more flexible and scalable automation, allowing organizations to adapt swiftly to evolving business rules or legacy constraints without costly reengineering.

This paper presents our practical experience in designing, deploying, and evaluating an IRPA system to support the Department of Health of the State Attorney's Office (PROSAÚDE) in managing legal demands related to public healthcare. The General State Attorney's Office of Ceará (Procuradoria Geral do Estado do Ceará - PGE) is the state body responsible for representing the government in legal matters across civil, labor, commercial, and criminal domains. The IRPA agent developed for PROSAÚDE was designed to automate workflows that previously depended on extensive manual effort.

Facing over 3,000 legal cases per month, the agent automates core decisions related to legal contestation by interpreting court documents and applying internal legal regulations. Powered by GPT-4o and guided by advanced prompt engineering, the agent combines IRPA functionalities to extract relevant medical and legal information, access external databases (e.g., ANVISA drug registries), and generate formal legal dispatches. Its primary objective is to ensure consistent and accurate application of PROSAÚDE's Internal Ordinance, which defines whether an appeal should be filed or waived.

Empirical evaluation of the system demonstrated strong performance: 98.4% precision in applying legal rules, 93.8% recall in identifying prescribed medications, and 98.7% of generated dispatches validated by legal experts without revision. These results underscore the potential of AI-powered automation to enhance accuracy, efficiency, and procedural safety in large-scale public-sector legal workflows. The implementation of IRPA led to substantial productivity gains, with certain tasks executed *3*-times faster and yielding an estimated savings of *5376* annual work hours. These improvements have enabled staff to focus on higher-value, strategic activities.

The article is organized as follows. Section 3 introduces the architecture and core components of the intelligent agent, detailing how it processes legal documents, retrieves regulatory information, applies normative rules, and generates dispatches. Section 4 describes the datasets and strategies used for evaluation. Section 5 presents the experimental outcomes, including classification metrics, confusion matrices, and human-in-the-loop validation results. Finally, Section 2 reviews related works for comparison.

2. Related Works

Robotic Process Automation (RPA) has emerged as a foundational technology for automating routine and repetitive tasks traditionally performed by humans, particularly in contexts involving structured data and rule-based operations. As described by Lacity and Willcocks [Lacity and Willcocks 2016], RPA involves configuring software robots to replicate human interactions within digital environments, enabling the execution of predefined processes with high consistency and without the limitations of fatigue. These robots interface directly with existing systems, mimicking user actions without requiring modifications to legacy software.

In the study conducted by [Patrício et al. 2024], the authors highlight that the adoption of Robotic Process Automation (RPA) has accelerated across various sectors, primarily due to its relatively low implementation costs and its compatibility with Artificial Intelligence (AI) techniques. The integration of RPA with AI has significantly enhanced its capabilities, enabling the processing of semi-structured data and the execution of tasks involving cognitive functions. Modern RPA tools now support advanced functionalities, including web-based information retrieval, conversational interfaces, and the interpretation of unstructured textual inputs.

Khadka et al. [Khadka et al. 2013] conducted a systematic literature review on the evolution from legacy systems to service-oriented architectures (SOA), identifying key challenges in migrating monolithic applications to modular and interoperable platforms. Their findings highlight the importance of incremental transformation strategies and the necessity of automation tools in facilitating sustainable migration paths, a premise that aligns with the increasing application of RPA in modernization efforts.

Aguirre and Rodriguez [Aguirre and Rodriguez 2017] provided an early empirical account of RPA implementation in a business process. Their case study demonstrated how RPA can automate structured, repetitive tasks, yielding efficiency gains and error reductions. Complementing this practical insight, Radke et al. [Radke et al. 2020] showed that RPA can also address data consistency and quality challenges, particularly in complex domains such as item master data maintenance in logistics.

From a broader systems engineering perspective, Van der Aalst et al.[Van der Aalst et al. 2018] provided a theoretical lens on RPA, distinguishing it from traditional workflow automation by emphasizing its ability to operate across systems without requiring modifications to underlying software. Jovanović et al.[Jovanović et al. 2018] extended this by offering a comprehensive overview of RPA capabilities, deployment scenarios, and integration challenges, while Devarajan [Devarajan 2018] catalogued real-world use cases, highlighting RPA's relevance in finance, insurance, and retail.

Anagnoste [Anagnoste 2018] characterized RPA as the operating system of the digital enterprise, proposing that its integration within organizational structures signals a shift toward digitally native business operations. This perspective is echoed by Lamberton et al. [Lamberton et al. 2017], who analyzed the impact of RPA and AI in the insurance sector, revealing both efficiency opportunities and regulatory challenges.

Finally, Ivančić et al.[Ivančić et al. 2019] presented a systematic literature review of RPA research, providing an analytical framework for understanding its evolution, taxonomy, and maturity. This complements the conceptual framework proposed by Tsaih and Hsu[Tsaih and Hsu 2018] for AI in smart tourism, demonstrating the convergence of automation and intelligence in sector-specific transformations.

3. Agent Workflow and Component Design

The intelligent agent is designed with the goal of correctly analyzing and applying the legal criteria established in PROSAÚDE's Internal Ordinance, and of generating and routing the appropriate legal dispatch based on that analysis. To achieve this, the agent operates through a modular architecture that combines perception, reasoning, and action in an integrated decision-making loop.

Perception involves continuously monitoring for new judicial intimations, retrieving the full legal process from external judiciary platforms, selecting the relevant document for analysis, and querying external systems, such as ANVISA and official price databases, to collect updated regulatory information.

Reasoning consists of interpreting legal texts using large language models (LLMs), extracting structured information, and evaluating whether the legal and medical conditions match the waiver criteria defined in the ordinance.

Action is performed by generating formal legal dispatches that justify the decision and routing each case to the appropriate internal unit, whether administrative staff (for uncontested cases) or legal advisors and attorneys (for further legal response).

3.1. Intimation Intake and Case Retrieval

The agent continuously monitors the internal systems of the State Attorney General's Office (PGE) to detect new judicial intimations that require action. For each identified intimation, it autonomously accesses external judiciary platforms (e.g., the State Court of Ceará or the Federal Court of the 5th Region), using automated web access to retrieve the full case history. This ensures that the most recent and relevant legal documents are available for analysis.

3.2. Document Selection and Summarization

From the set of retrieved case documents, the agent analyzes each item to identify the document that contains the core legal or medical content related to the current intimation. This involves distinguishing between different types of procedural movements—such as petitions, judicial decisions, or administrative communications—and selecting the most relevant document for downstream processing. It then applies natural language processing (NLP) techniques to generate a structured summary with key information such as the cause value, the type of document (e.g., sentence or petition), and the list of requested health items and services. See Figure 1 for the complete prompt.

3.3. Prompt-Based Information Extraction

With the summary in place, specialized modules use prompt-based interactions with GPT-4o to extract fine-grained legal and medical information. These modules identify the names of medications, nutritional supplements, and medical devices; detect the presence of moral damages or attorney's fees; and capture legal deadlines or urgent compliance requirements. The specialized extraction modules are: (1) medications, (2) nutritional compounds and special diets, (3) medical consultations, exams, and procedures, (4) general medical supplies, and (5) attorney’s fee sentencing.

You are a legal assistant and need to summarize the content of a document requesting or ordering the provision of items and/or services in the medical field.

Consider as medications only substances or pharmaceutical compounds used exclusively to treat, prevent, or cure diseases. Other medical or support items, such as diapers, nutritional compounds, syringes, gloves, oximeters, hospital beds, or thermometers are not medications. Under no circumstances should you include in the list medications that were not mentioned in the decision. If there are no medications, simply state that there are none.

Consider as hospitalization or transfer to a hospital bed any type of admission or transfer to specialized health units. This should include ICUs (Intensive Care Units), UCEs (Special Care Units), and psychiatric or substance treatment facilities, among others.

Consider consultations as medical evaluations performed by healthcare professionals for diagnosis, follow-up, or treatment of health conditions.

Consider exams as medical investigation procedures, such as laboratory or imaging exams, to assist in diagnosing diseases or monitoring the patient.

Consider clinical or surgical procedures as therapeutic or diagnostic interventions, which may be minor (such as dressing or suture removal) or invasive (such as surgeries), aimed at treating or improving the patient’s condition.

Consider nutritional compounds and/or special diets as supplements or specific dietary regimes recommended to meet the nutritional needs of patients with particular health conditions, such as allergies, intolerances, nutritional deficiencies, or chronic illnesses, aimed at promoting or restoring health. They must not be confused with medications.

The summary must be written in Brazilian Portuguese and must contain the following structure and information:
*Document Summary**

- **Document Type:** Indicate whether it is a Sentence, Interlocutory Decision, or Initial Petition
- **Cause Value:** Inform the cause value
- **Items to be Provided:**
- **Medications:**
- List the medications, each with dosage, quantity, and treatment duration
- **Nutritional Compounds:**
- List the nutritional compounds with dosage, quantity, and treatment duration
- **Hospitalization or Transfer to Hospital Bed:**
- List hospital admissions or transfers
- **Consultations, Exams, or Procedures:**
- List the consultations, exams, or procedures
- **Other Items or Services:**
- List other necessary items or services not included in the categories above

- **Is there Hospitalization or Transfer to a Hospital Bed:** ’Yes’ or ’No’
- **Are there Consultations, Exams and/or Procedures:** ’Yes’ or ’No’
- **Are there Moral Damages:** ’Yes’ or ’No’
- **Are there Attorney’s Fees:** Indicate if there is a fee order, specifying who was sentenced and the amount
- **Is there a Request for Preliminary Injunction or Anticipatory Relief:** ’Yes’ or ’No’
- **Was the Request for Injunction or Anticipatory Relief Granted or Denied:** Indicate whether it was granted, denied, or not applicable

- **Was the Case Dismissed Without Judgment on the Merits:** ’Yes’ or ’No’
- **Is this a Decision to Enforce a Sentence:** ’Yes’ or ’No’
- **Is this a Decision to Block Assets or Accounts:** ’Yes’ or ’No’. If blocking is only mentioned as a possibility, but not determined, the answer is ’No’
- **Is this a Document Entitled Monocratic Decision:** ’Yes’ or ’No’
- **Is this a Decision or Sentence in which the Judge Requested Further Action (Diligence):** ’Yes’ or ’No’

*Additional Details:**
- **Claimant:** Provide the claimant’s name and, if applicable, whom they represent
- **Defendants:** List the defendants
- **Legal Grounds:** Briefly summarize the subject of the document

Figure 1. Prompt used for structured summarization of legal documents (translated from the original version in Brazilian Portuguese).

One notable case involves the detection of attorney's fees imposed on the State of Ceará. These cases can be difficult, as fees may be shared between different entities (e.g., State and Municipality), or defined as a percentage of the total cause value rather than a fixed amount. In earlier versions of the system, these scenarios often led to incorrect or incomplete extractions. To address this, we introduced Chain-of-Thought prompting [Wei et al. 2022], guiding the language model to reason step-by-step through the necessary legal interpretation:

1. Check whether the document includes a sentence imposing attorney's fees;
2. Determine whether the State of Ceará is explicitly among the parties sentenced;
3. Identify whether the fee is expressed as a fixed amount or as a percentage of the cause value;
4. If a percentage is used, calculate the corresponding monetary amount;
5. Output the structured result with the conclusion and monetary value.

3.4. Use of External Regulatory Sources

In addition to document analysis and internal rule application, the intelligent agent relies on external regulatory data to support its decision-making. These sources are essential for determining whether the requested health items fall within the scope of public health policies and for validating registration or pricing information that may affect the application of the internal ordinance. A particularly notable case is the analysis of medication requests, which often requires the integration of multiple external databases:

- **ANVISA (National Health Surveillance Agency)**: Brazil's federal agency responsible for regulating the approval and registration of medications and medical products.[1] The agent checks whether the requested medications are registered with ANVISA to determine their legal availability in the national healthcare system.
- **CMED (Drug Market Regulation Chamber)**: A division of ANVISA that publishes the official price ceiling table for medications marketed in Brazil.[2] The agent uses this table to assess whether the price of a requested medication is within the legal limits for public acquisition.
- **RENAME (National List of Essential Medicines)**: A federal list that defines the essential medications provided by the Brazilian Unified Health System (SUS).[3] Medications included in RENAME are generally guaranteed by the public system and may not require litigation, while those absent are subject to further legal evaluation.
- **RESME (State List of Medicines for Ceará)**: A list maintained by the State Health Department of Ceará specifying the medications available through state-level SUS programs.[4] This list is used to verify whether the item is already available through local supply channels.

These resources are accessed through automated web navigation using the Selenium framework. Selenium enables the agent to access and extract data from dynamic web portals through automated browser control. This approach ensures that data retrieved is current, complete, and aligned with official sources.

[1] `https://consultas.anvisa.gov.br`
[2] `https://www.gov.br/anvisa/pt-br/assuntos/medicamentos/cmed/precos`
[3] `https://www.gov.br/saude/pt-br/composicao/sectics/rename`
[4] `https://www.saude.ce.gov.br/download/publicacoes-resme/`

3.5. Normative Rule Evaluation

The agent applies a rule-based decision module that encodes PROSAÚDE's Internal Ordinance (PIO). These rules determine whether a legal contestation is warranted or whether the case should be closed administratively and forwarded to the health department for execution.

Rules consider:

- The nature of the decision (preliminary, final, interlocutory);
- Whether the requested item is part of official public health protocols;
- The existence of jurisprudence or previously settled positions.

3.6. Dispatch Generation and Routing

Based on the applied rules and extracted data, the agent automatically generates a formal legal dispatch, including standardized justification and institutional language. It then routes the case to the appropriate unit:

- **Administrative support** when no legal action is required;
- **Legal advisors or attorneys** when further legal response is needed.

3.7. Technical Foundations

The system is implemented in Python and integrates several technologies: Selenium for web automation, spaCy and OpenAI APIs for natural language processing and prompt-based extraction, and Pandas for structured data handling. The architecture is modular and extensible, allowing new legal rules, document types, and external integrations to be added as regulations evolve.

4. Methodology

To evaluate the effectiveness of the proposed intelligent agent in supporting the legal workflow at PROSAÚDE, we conducted a series of validation experiments using a curated set of real judicial documents. The objective was to assess the system's ability to extract relevant information and apply internal legal rules with high accuracy and consistency.

4.1. Use of Large Language Models

At the core of the system lies the GPT-4o model from OpenAI, which is responsible for interpreting legal documents and extracting structured information from unstructured text. Prompt engineering techniques were applied to generate summaries, extract key facts (e.g., medication names, document type, presence of attorney fees), and support legal reasoning aligned with PROSAÚDE's internal ordinance. All prompts were iteratively refined based on expert feedback, and the model's outputs were post-processed and validated against internal rules and ontologies.

4.2. Gold Standard Dataset

To support system evaluation, we constructed a **gold standard dataset** from previously analyzed legal cases. Each case had been manually reviewed and annotated by legal experts from PROSAÚDE, including legal advisors and state attorneys. As such, the dataset serves as a reliable ground truth for benchmarking the performance of the intelligent agent.

The dataset comprises three distinct document types commonly received by PROSAÚDE:

- **Set 1: Final Judgments** – 150 documents;
- **Set 2: Interlocutory Decisions** – 150 documents;
- **Set 3: Initial Petitions** – 140 documents.

Each document was annotated with three key labels:

1. Whether the internal ordinance should be applied;
2. Whether the case includes a request for medication supply;
3. Whether there is a condemnation for attorney's fees above R$1,500.

4.3. Tasks and Metrics

Three core classification tasks were defined:

- **Rule Application Detection:** Determine whether the internal ordinance applies based on the document content;
- **Medication Request Detection:** Identify the presence of requested medications;
- **Attorney Fee Detection:** Detect whether there is a financial condemnation for attorney's fees above the R$1,500 threshold.

Each task was evaluated using standard classification metrics: **accuracy**, **precision**, and **recall**.

5. Results

The results for each of the three classification tasks are presented below, including updated performance metrics (accuracy, precision, recall, and F1-score) and their corresponding confusion matrices.

5.1. Application of the internal ordinance

The detection of whether internal ordinance applies to a judicial case is the central element guiding PROSAÚDE's response strategy. This internal ordinance establishes clear criteria for when the state should abstain from filing defenses or appeals. When the ordinance is applicable, the case no longer requires legal action and can be handled by assistants instead of attorneys. In such situations, the assistants proceed with archiving the case and issuing a formal notice to the *Secretaria de Saúde do Ceará* (State Health Department), which becomes responsible for executing the judicial order and ensuring the supply of the requested items.

Table 1. Performance Metrics – Application of Ordinance 01/2017

Metric	Value
Accuracy	85.89%
Precision	98.41%
Recall	64.92%
F1-Score	78.23%

Table 2. Confusion Matrix – Application of Ordinance 01/2017

Predicted \ Actual	TRUE	FALSE
TRUE	124	2
FALSE	67	296

The model achieves notable precision (98.41%) in identifying cases where the ordinance should be applied. This is the most important metric for this task, as it ensures that cases predicted as exempt from legal action are indeed correct, minimizing the risk of mistakenly omitting a necessary defense. A high-precision model significantly reduces legal exposure by preventing the misclassification of borderline cases as administratively resolvable.

Although the recall (64.92%) is lower, its impact is less critical in this context. A false negative, i.e., a case where the Ordinance could have been applied but the system failed to detect it — results only in the case being forwarded for human review by legal advisors or state attorneys. Since these professionals will still verify whether the ordinance applies, no harm is done; the process simply follows a more conservative, attorney-involved path. Thus, the current balance between high precision and moderate recall favors legal safety and operational efficiency.

5.2. Detection of Medication Requests

Proper detection of medication requests is a critical component of the system, as it directly impacts the subsequent legal and financial evaluation of each case. Many medications requested via judicial processes involve exceptionally high costs for the public health system. Additionally, Brazilian law does not obligate the state to supply medications that are not officially registered with ANVISA. As a result, accurately identifying the presence of medications in judicial decisions is essential to trigger external validations and determine whether the request falls within the scope of enforceable public obligations.

Table 3. Performance Metrics – Detection of Medication Requests

Metric	**Value**
Accuracy	90.80%
Precision	87.19%
Recall	93.78%
F1-Score	90.36%

Table 4. Confusion Matrix – Detection of Medication Requests

Predicted \ Actual	TRUE	FALSE
TRUE	211	31
FALSE	14	233

The model achieves strong overall performance in identifying medication-related requests, with both high recall (93.78%) and precision (87.19%). In this context, both metrics are important: precision ensures that non-medication entities are not misclassified, avoiding unnecessary or incorrect validations; recall ensures that the system detects all true instances of medications, enabling the application of internal rules and external checks such as ANVISA registration lookup and cost verification.

Among the two, recall is the most critical: if the system fails to detect a true medication, it will omit the ANVISA verification step, potentially leading to the wrongful assumption that the case is not subject to legal objection or further analysis. In contrast, a false positive would only result in an unnecessary check, which is operationally tolerable. Therefore, the model's high recall suggests that it is well suited to the task, reducing legal risk and ensuring coverage of all medication-related demands.

5.3. Detection of Attorney's Fees Above R$1,500

Identifying cases in which the judicial decision imposes attorney's fees above R$1,500 is relevant for the strategic legal response by PROSAÚDE. When such financial penalties are involved, especially in high volumes, they can significantly impact the budget and accountability of the State Attorney's Office. In these situations, legal teams are often more motivated to challenge the decision through appeals or other procedural defenses. Therefore, accurate detection of high attorney fee condemnations is essential to avoid mistakenly waiving the right to contest or appeal when the internal ordinance would otherwise recommend dismissal of legal action.

Table 5. Performance Metrics – Attorney's Fees Above R$1,500

Metric	Value
Accuracy	92.43%
Precision	56.00%
Recall	91.30%
F1-Score	69.42%

Table 6. Confusion Matrix – Attorney's Fees Above R$1,500

Predicted \ Actual	TRUE	FALSE
TRUE	42	33
FALSE	4	410

The model demonstrates strong recall (91.30%), meaning it is highly capable of detecting nearly all cases in which fees exceed the defined threshold. This is crucial to avoid situations in which attorney fee condemnations go unnoticed and the case is inappropriately closed under the terms of the ordinance.

Although precision is low (56.00%), this is an acceptable trade-off in this context. False positives—cases predicted to include high fees when they actually do not—will still be routed to legal advisors or attorneys for further review. Importantly, these cases will not be automatically exempted from legal response by the system. In practice, this means that the ordinance will not be applied prematurely, and human validation will safeguard against improper dismissals.

Overall, the model effectively prioritizes legal security by maximizing coverage of risky cases while leaving the final validation of borderline situations to human experts.

5.4. Validation of Dispatch Generation

The quality of the legal dispatches generated by the intelligent agent was evaluated through a human-in-the-loop validation process. PROSAÚDE's team of professionals manually reviewed each dispatch to assess whether it was complete, correct, and aligned with PROSAÚDE's institutional standards and language.

Table 7. Human Validation of Generated Dispatches

Metric	Value
Dispatches reviewed	157
Correct dispatches	155
Accuracy	98.7%

Out of 157 dispatches reviewed, 155 were considered correct by the legal reviewers, resulting in an accuracy of 98.7%. This high level of precision confirms that the system is capable of generating legally sound texts with minimal need for human intervention. The remaining two dispatches were flagged for minor inconsistencies or adjustments, but none required full rework. These results reinforce the system's readiness for operational deployment and its potential to significantly reduce manual workload in the legal handling of health-related cases.

6. Conclusion

This work presented an intelligent agent that combines LLM-based information extraction, rule evaluation, and robotic process automation to support high-volume legal workflows in the Brazilian public sector. In controlled evaluation, the system achieved **98.41% precision** for detecting when the internal ordinance applies, **93.78% recall** for medication request detection, and **98.7% correctness** in human validation of generated dispatches.

Beyond accuracy, the system delivers organizational value by converting unstructured case documents into structured data, enabling demand monitoring, early pattern detection, and improved coordination between the Attorney General's Office and the State Health Secretariat. It functions as both an automation layer and a decision-support pipeline.

Operationally, IRPA adoption produces notable productivity gains, with some tasks performed three times faster and an estimated 5,376 annual work hours saved, demonstrating that IRPA combined with LLMs can reliably enhance legal workflows.

References

Aguirre, S. and Rodriguez, A. (2017). Automation of a business process using robotic process automation (rpa): A case study. In *Applied Computer Sciences in Engineering: 4th Workshop on Engineering Applications (WEA 2017)*, pages 65–71. Springer.

Anagnoste, S. (2018). Robotic automation process – the operating system for the digital enterprise. In *Proceedings of the International Conference on Business Excellence*, volume 12, pages 54–69.

Costa, D., Mamede, H., and Mira da Silva, M. (2022). Robotic process automation (rpa) adoption: A systematic literature review. *Engineering Management in Production and Services*, 14(2):1–12.

Devarajan, Y. (2018). A study of robotic process automation use cases today for tomorrow's business. *International Journal of Computer Techniques*, 5(6):12–18.

Ivančić, L., Suša Vugec, D., and Bosilj Vukšić, V. (2019). Robotic process automation: Systematic literature review. In *Business Process Management: Blockchain and Central and Eastern Europe Forum*, pages 280–295. Springer.

Jovanović, S. Z., Đurić, J. S., and Šibalija, T. V. (2018). Robotic process automation: Overview and opportunities. *International Journal Advanced Quality*, 46(3–4):34–39.

Khadka, R., Saeidi, A., Idu, A., Hage, J., and Jansen, S. (2013). Legacy to soa evolution: a systematic literature review. In *Migrating Legacy Applications: Challenges in Service Oriented Architecture and Cloud Computing Environments*, pages 40–70.

Lacity, M. and Willcocks, L. (2016). Robotic process automation at telefonica o2. *MIS Quarterly Executive*, 15(1). Article 4.

Lamberton, C., Brigo, D., and Hoy, D. (2017). Impact of robotics, rpa and ai on the insurance industry: Challenges and opportunities. *Journal of Financial Perspectives*, 4(1).

Ng, K., Chen, C.-H., Lee, C., Jiao, J., and Yang, Z.-X. (2021). A systematic literature review on intelligent automation: Aligning concepts from theory, practice, and future perspectives. *Advanced Engineering Informatics*, 47:101246.

Patrício, L., Varela, L., and Silveira, Z. (2024). Integration of artificial intelligence and robotic process automation: Literature review and proposal for a sustainable model. *Applied Sciences*, 14:9648.

Radke, A. M., Dang, M. T., and Tan, A. (2020). Using robotic process automation (rpa) to enhance item master data maintenance process. *LogForum*, 16(1).

Tsaih, R. H. and Hsu, C. C. (2018). Artificial intelligence in smart tourism: A conceptual framework.

Van der Aalst, W. M., Bichler, M., and Heinzl, A. (2018). Robotic process automation. *Business & Information Systems Engineering*, 60:269–272.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., and Zhou, D. (2022). Chain of thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24877–24891.