Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI-Powered Responsive RTI and IPC Complaint Automation System

Prabha M

drtmprabha5@gmail.com

saveetha engineering College

Chandru V

saveetha engineering College

Prakash R

saveetha engineering College

Roghith K

saveetha engineering College

Dharmaraj S

saveetha engineering College

Research Article

Keywords: Legal Tech, Computational Law, Right to Information, Indian Penal Code, Heuristic Analysis, E-Governance, Document Classification, Named Entity Recognition

Posted Date: March 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9028226/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

AI-Powered Responsive RTI and IPC Complaint Automation System

Dr Prabha.M Associate Professor, AIDS, Saveetha Engineering College, Chennai, TN, India
Chandru V – Student, CSE, Saveetha Engineering College, Chennai, TN, India
Prakash R – Student, AIML, Saveetha Engineering College, Chennai, TN, India
Roghith K - Student, CSE, Saveetha Engineering College, Chennai, TN, India
Dharmaraj S – Student, B.Tech AIML, Saveetha Engineering College, Chennai, TN, India

Abstract: The Right to Information Act, 2005, is one of the major building blocks of legislative tools that define the principle of transparency and empowerment of citizens in the Indian polity; but its infallibility in practice often runs afoul of avoiding, half-hearted, or procedurally non-compliant replies by Public Information Officer. Such responses need to be interpreted with particular juridical expertise and the instances of statutory violation be identified, thus obstructing substantive citizen involvement in the accountability mechanism. In order to address this shortcoming, this study suggests an AI-based, data-agnostic RTI intelligence system. The architecture includes a multi-phase pipeline that fuses the implementation of Optical Character Recognition (OCR) and Natural Language Processing (NLP) using transformers and Named Entity Recognition (NER) with an engine based on heuristics and legal inferences to break down the RTI answers. In addition to classification, the framework provides an ordered identification of data requirements and the mechanism of multi-source retrieval prioritizing internal databases, investigating statutory and open-government Application Programming Interfaces such as CPGRAMS, and performing AI-aided cross source-checking to shed light on inconsistencies and disclosure gaps. The generation of drafts to be applied to another application or appeal are posed as a secondary, supported based tool that is only triggered after confirmation of non-compliance. This is to make sure that the outputs of the system also have sound basis on the jurisprudential analysis. The empirical assessment has proven to have strong classification and extraction ability which has been supported by statistical validation using the standard deviation and the confidence interval analysis. These measures testify to the stability of the system in different data sets. The suggested framework leads to strengthening citizen-led public responsibility in digital governance systems by operationalizing structured clarity in the data and AI-aided verification.

Keywords- Legal Tech, Computational Law, Right to Information, Indian Penal Code, Heuristic Analysis, E-Governance, Document Classification, Named Entity Recognition

I. INTRODUCTION

The Right to Information Act of 2005 was signed to bring on board transparency, bolster accountability to the population as well as increase citizen involvement in the governance of the entire India [31]. The Act as an official confirmation of the rights of every citizen to access information as received by authorities in the state provides an effective right to access information; however, the daily practice of the act is frequently faced with systematic and procedural challenges [33], [34]. In many cases, what is known as information-seeking process comes to an effective end when at the stage of the first response, the citizens would receive responses that indeed, dispose of the request technically but would not disclose anything substantive [35].

These kinds of reactions can be divided in four patterns. To begin with, ambiguous or non-specific responses provide impersonal information which does not consider specific aspects under the application. Second, incomplete information conveyed by preparation of half-truth and/or false disclosures alters the pertinent information. Thirdly, unreasonable refusals rely on statutory exemptions such as the section 8 of the Act, without justifiable reasons [31]. Fourth, procedural evasions (such as records being not traceable, records not held), are indicative of poor record-

management practices and can also be in violation of section 4 requirements or in worst-case scenarios, should be informational suppression [31], [33], [34].

With these responses, citizens are facing great cognitive and procedural burdens [34]. They need to identify the correct cause of law which is missing and identify the right corrective mechanism; which includes follow-up application or statutory appeal or grievance system or other forms of recourse and present their claims with specific legal expressions. This kind of asymmetry of information and an obscurity of procedure is a chilling effect, one that compromises the successful implementation of transparency standards and corrosivity to the accountable actions of citizens [33], [26], [35].

II. LITERATURE SURVEY

The present study comes as an extension of the developments of Legal Tech, specifically, the usage of NLP and LLMs to Indian legal texts [7], [8], [23]. Up-to-date improvements in the field of legal artificial intelligence have highlighted the importance of customised and retrieval-enhanced language models in order to manage the exactness and responsibility demanded by legal contexts [25], [26]. Gupta et al. (2025) and Ganatra et al. (2025) show in the Indian legal ecosystem that fine-tuned and RAG-based LLMs decrease highly hallucinated statistics and enhance statutory grounding, legal advice, and consumer grievance systems [1], [2]. However, their generalisation ability is limited beyond their training contexts, and so on represents a serious reliance on curated datasets legal and presents challenging scalability to diverse legal codes, jurisdictions as well as language environments [7], [8].

A substantial body of prior research has concentrated on the automation of tasks within the criminal law domain, including the drafting of First Information Reports (FIRs), prediction of applicable sections of the Indian Penal Code (IPC), and statutory recommendation [4], [5], [6]. Bhatia et al. (2025) demonstrate that fine-tuned lightweight language models can effectively assist non-expert users in FIR drafting, while transformer-based classifiers and similarity-driven methods report strong performance in predicting frequently invoked IPC provisions (Arya et al., 2025; Srivastav & Prajapat, 2022) [27], [5]. Knowledge-driven systems such as NyayGraph (Pal et al., 2025) further enhance statute identification by incorporating legal knowledge graphs, yielding improved recall, albeit often at the expense of reduced precision [3]. Collectively, these studies suggest that existing approaches are proficient in pattern recognition but remain limited in handling nuanced legal reasoning and less-represented categories of offences [26].

Parallel research efforts focus on legal document understanding and system reliability, addressing tasks such as named entity recognition, rhetorical role classification, and long-document summarization [15], [12], [10]. The availability of large annotated datasets and continued in-domain pretraining, including Legal BERT variants tailored to Indian law, has been shown to enhance downstream legal NLP performance [13].

However, robustness studies highlight persistent limitations, including hallucinations in generative models, vulnerability to adversarial inputs, and substantial computational overheads [25], [26]. Overall, the literature emphasizes the need for legal AI systems that combine grounded retrieval, expert-aligned reasoning, and rigorous evaluation frameworks to ensure accuracy, interpretability, and trustworthiness in practical deployments [1], [2], [26].

Existing legal AI systems lack a unified, scalable, and trustworthy framework that combines grounded retrieval, context-aware legal reasoning, robustness to rare and adversarial cases, multilingual adaptability, and expert-aligned explainability, which is essential for reliable real-world deployment in complex legal domains such as IPC-and RTI-driven applications [1], [2], [26]. The suggested AI-Powered RTI-IPC system offers an end-to-end, practical solution that specifically targets the post-response stage of the RTI process, in contrast to earlier studies that only concentrate on judgment prediction, question answering, or statute classification [4], [18], [24].

III. METHODOLOGY

We in this paper take a data-first paradigm of accountability that methodic alters raw non-structured RTI responses into structured and verifiable workflows of decision-making. Instead of issuing remedial actions on a first come first served basis, the system puts more emphasis on semantic interpretation, search of requisite structured data, hierarchical retrieval [1, 20], and cross-source consistency [26]. Our approach has been outlined in 4 consecutive layers: Processing Pipeline, Retrieval Framework, Verification Layer and Decision Support. The layer builds upon refinements of raw disclosures to the final results of evidence-based accountability.

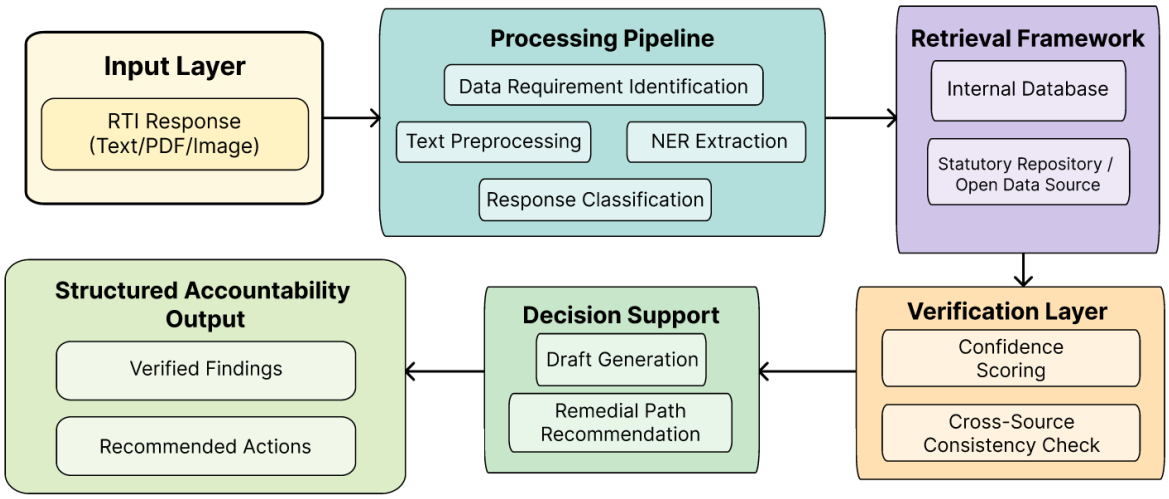


Fig. 1. Overall Data-First Accountability Framework

As illustrated in **Fig. 1**, the framework follows a structured progression beginning with input processing and culminating in verified accountability outputs. The successive operation of each of the stages in the framework fulfills a different analytical task and strictly maintains a logical dependence on the former layer of the framework so that any further escalation decision will be based on proven evidence as opposed to a form of algorithmic inference. The architecture lowers the risk of premature or erroneous interventions which otherwise would come about by the mere act of automation by basing the decisions on the well-tested information [26].

The first phase of the architecture, which is marked as Processing Pipeline, is mandated with the responsibility of transformation of unstructured, raw RTI responses into systematized, organized forms. During this stage text normalization, entity extraction [15], response classification [4], and data requirement identification processes are applied in a systematic manner thus transforming heterogenous and unstructured disclosures to machine understandable accountability parameters. By converting the lacks in responses into organized pieces of data, this phase develops a solid base upon which the subsequent retrieval and validation procedures that form the following layers of the system base.

A. Processing Pipeline

The Processing Pipeline constitutes the first analytical stage of the proposed framework. Its purpose is to transform unstructured RTI responses into structured representations that enable retrieval and verification. The pipeline operates sequentially through preprocessing, entity extraction, response classification, and data requirement identification.

1) Text Preprocessing: RTI responses submitted in text, PDF, or image formats are converted into machine-readable form using Optical Character Recognition (OCR) where required. The extracted text undergoes normalization, including removal of noise artifacts, standardization of date and reference formats, and structural cleaning to ensure semantic consistency.

2) Named Entity Extraction: A transformer-based Natural Language Processing (NLP) model performs Named Entity Recognition (NER) to extract legally relevant entities such as public authority names, PIO identifiers, statutory references, dates, application numbers, and disclosure descriptors. This structured extraction enables contextual mapping of the response.

3) Response Classification: The normalized and annotated text is processed using a supervised classification model trained on an annotated corpus of RTI responses. The classifier categorizes responses into compliance types including:

- Satisfactory Disclosure
- Incomplete Disclosure
- Valid Statutory Denial
- Improper Denial
- Procedural Evasion
- Deemed Refusal

This classification provides diagnostic grounding for accountability assessment.

4) Data Requirement Identification : Rather than immediately recommending remedial actions, the system determines the specific data elements required to verify the adequacy of the response. By translating qualitative deficiencies into structured accountability parameters, this stage establishes the evidentiary foundation for the Retrieval Framework and Verification Layer.

Algorithm 1: RTI Response Processing Pipeline

Input: RTI_Response
Output: Structured_Response_Object

Let R be an RTI response.
Let T = Normalize(Extract_Text(R)).
Let E = NER(T).
Let C = Classifier(T, E).

If C ∉ {Satisfactory Disclosure}:
  D = DataRequirement(T, E)
Else:
  D = ∅

Return S = (T, E, C, D)

Algorithm 1 converts an unstructured RTI response into a structured accountability object. Data requirements are conditionally identified only for deficient responses, ensuring evidence-driven escalation. This design separates interpretation from action and enables a consistent, verifiable downstream workflow.

B. Retrieval Framework

The Retrieval Framework operationalizes the data-first accountability model by validating response deficiencies against structured information sources as shown in **Fig 2**. Rather than relying solely on the textual interpretation of the RTI reply, the framework performs hierarchical multi-source data acquisition to [1, 20] confirm disclosure adequacy.

The retrieval process follows a layered strategy. First, the system queries internal structured repositories containing curated RTI records and related public disclosures [36]. If the required data elements are not found, the system proceeds to query statutory or open-government repositories to obtain authoritative information [37]. Only when relevant data cannot be retrieved from available sources is a missing-data condition flagged for escalation. This hierarchical design reduces single-source dependency and improves evidentiary robustness [26].

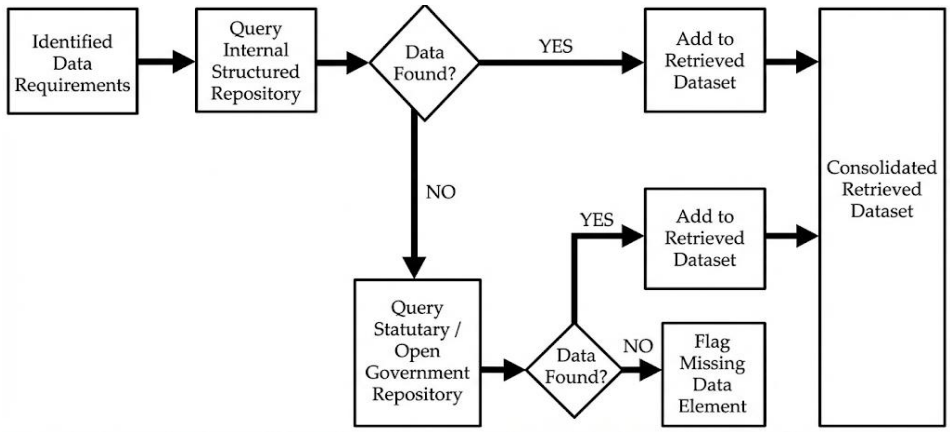


Fig. 2. Hierarchical Multi-Source Retrieval Framework showing layered querying of internal and statutory repositories prior to missing-data flagging.

By progressively querying trusted sources before escalation, the framework ensures that accountability assessment is grounded in verifiable data rather than isolated document interpretation [26].

Algorithm 2: RTI Response Processing

Input: Data_Requirements
Output: Retrieved_Data_Set

Let $D = \{d_1, d_2, \ldots, d_n\}$ be the set of identified data requirements.
Let I(d) represent the internal repository query function.
Let S(d) represent the statutory/open repository query function.
Let R denote the retrieved dataset.

Initialize: $R = \emptyset$

For each $d_i \in D$:
If $I(d_i) \neq \emptyset$
$x_i = I(d_i)$

Else if $S(d_i) \neq \emptyset$
    $x_i = S(d_i)$
Else
    $x_i = \emptyset$
    Mark d_i as Missing
End if

If $x_i \neq \emptyset$
    $R = R \cup \{x_i\}$
End if

Return R

Algorithm 2 implements hierarchical querying across trusted sources prior to declaring missing data. This design strengthens evidentiary robustness and prevents premature escalation. The consolidated dataset is then passed to the Verification Layer.

C. Verification Layer

The Verification Layer means comparing the appropriateness and similarity of the information that is retrieved with the selected RTI response. Its key aim is to determine whether the disclosure has met its statutory duties [31] or that there are discrepancies which indicate that there is non-compliance on the procedures [35]. With the Retrieval Layer, a comparative analysis is carried out as opposed to the Retrieval Framework that focuses on data acquisition. Matched Retrieved datasets are matched systematically according to the structured response object generated in the Processing Pipeline. This comparison occurs in several dimensions where some of the items to be considered are the factual consistency, completeness of disclosure, justification of exemption and time-delivery [26].

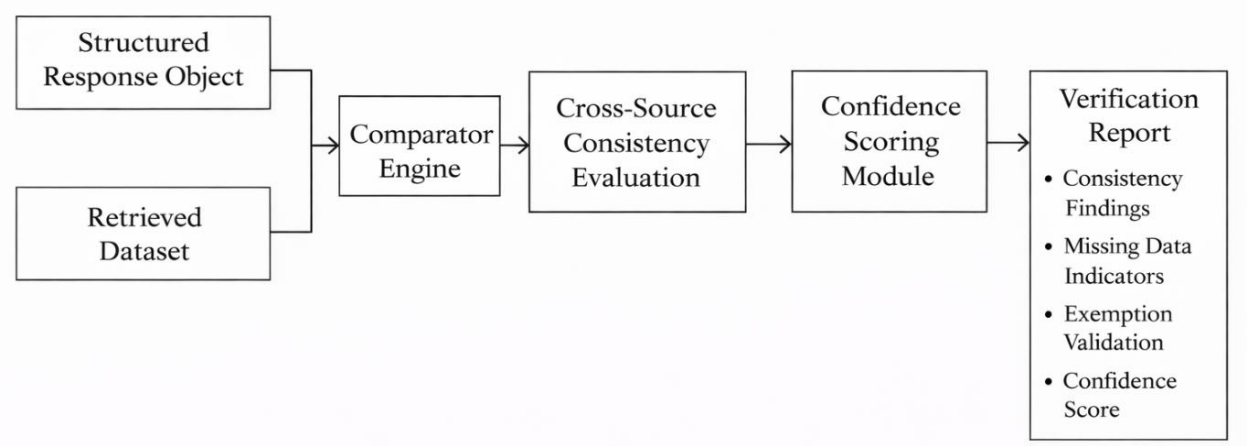


Fig. 3. Verification and Cross-Source Consistency Model illustrating comparative evaluation between structured response objects and retrieved datasets prior to generating a verification report.

First, cross-source consistency checks are done by the system as shown in the **Fig 3**. Each data element that is required is analyzed in order to ascertain the extent to which the information that is disclosed is in accordance with authoritative records that are obtained either in internal or statutory repositories [1, 26]. Any mismatch/ omission or contradiction is alerted to undergo further assessment. Second, exemption validation occurs in situations where the answers are based on statutory reasons on the denial. To evaluate whether the exemption mentioned corresponds logically to the information type requested, the system puts the unsupported or unreasonably justified exemptions at risk of being defined as potentially invalid [31, 35].

Thirdly, evidentiary reliability is measured through a confidence -scoring system. The number associated with the confidence score is the level of consistency between reported and retrieved facts which involves completeness, consistency, and source credibility [25, 26]. This measure allows it to evaluate a systematically, but not a dichotomous, evaluation. The output of the Verification Layer is a Verification Report comprising:

- Consistency findings
- Missing data indicators
- Exemption validation results
- Confidence score

This report forms the basis for the Decision Support stage, ensuring that remedial recommendations are grounded in verified evidence [26].

IV. DECISION SUPPORT

The Decision Support stage translates verified findings into structured remedial guidance. Unlike the Processing and Verification layers, which focus on interpretation and validation, this stage provides evidence-backed recommendations derived from the Verification Report.

The system applies a rule-conditioned mapping between verified non-compliance patterns and appropriate procedural pathways. For instance, incomplete disclosure may trigger a recommendation for a follow-up RTI request, while improper statutory denial may result in a suggested first appeal in accordance with the procedural framework of the Right to Information Act, 2005 [31]. Procedural evasion or confirmed missing data conditions may generate structured escalation pathways in accordance with applicable statutory provisions [31], [35].

Importantly, escalation is activated only after verification confirms non-compliance. This preserves the separation between analytical diagnosis and procedural action, ensuring that remedial recommendations are grounded in validated evidence rather than automated classification alone [26]. Draft generation for appeals or grievances functions as a secondary assistance mechanism. Templates are populated using structured outputs from the Verification Layer, maintaining traceability between detected deficiencies and recommended remedies. This design supports usability while preserving accountability integrity [33].

V. EXPERIMENTAL EVALUATION

The Experimental Evaluation assesses the effectiveness of the proposed framework in classifying RTI responses and supporting structured accountability workflows. The evaluation focuses on response classification accuracy, entity extraction reliability, and overall pipeline consistency.

A. Dataset Description

The evaluation was conducted on an annotated corpus of RTI responses collected across multiple compliance categories. The dataset includes instances of:

- Satisfactory disclosure
- Incomplete disclosure
- Valid statutory denial
- Improper denial
- Procedural evasion
- Deemed refusal

Each response was manually annotated to establish ground-truth labels for classification and entity identification. The dataset was divided into training and testing subsets using an appropriate split strategy (e.g., 80/20 split or k-fold cross-validation) to ensure unbiased evaluation.

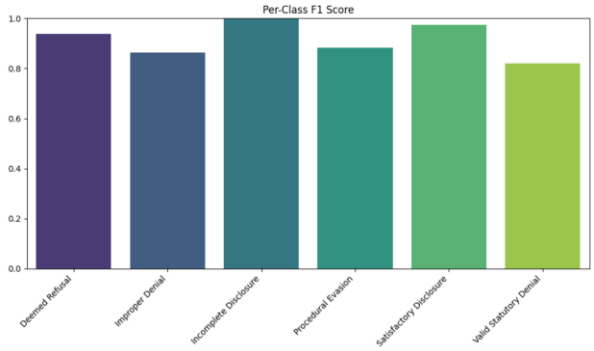
B. Results

The classification model demonstrates stable performance across response categories. Performance metrics indicate strong differentiation between compliant and non-compliant responses. Entity extraction accuracy further supports the reliability of structured representation generation.

Table 1: Classification Performance Metrics per RTI Response Category

Category	Precision	Recall	F1-Score	Support
Deemed Refusal	0.88	1.00	0.94	23
Improper Denial	0.95	0.79	0.86	48
Incomplete Disclosure	1.00	1.00	1.00	2
Procedural Evasion	0.93	0.84	0.88	98
Satisfactory Disclosure	1.00	0.95	0.97	20
Valid Statutory Denial	0.72	0.94	0.82	53
Macro Average	**0.92**	**0.92**	**0.91**	**244**
Weighted Average	**0.89**	**0.88**	**0.88**	**244**

Fig 4: Per-Class F1-score Across RTI Response Categories



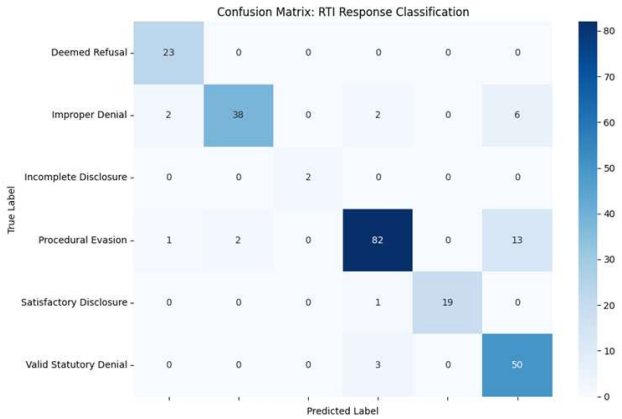


Fig. 5. Confusion matrix illustrating classification performance across RTI response categories.

As shown in Table 3 and Figure 4, the proposed classification framework achieves strong performance across all RTI response categories, with F1-scores ranging from **0.82 to 0.97** (excluding the low-support 'Incomplete Disclosure' class). The highest scores are observed for **Satisfactory Disclosure (0.97)** and **Deemed Refusal (0.94)**, indicating that the model effectively distinguishes responses grounded in clear legal reasoning or obvious non-compliance. **Valid Statutory Denial** exhibits a high recall of **0.94**, suggesting the model is highly effective at catching valid exemptions, though a precision of **0.72** indicates some over-prediction.

Procedural Evasion (0.88) and **Improper Denial (0.86)** reflect robust handling of linguistic ambiguity and administrative inconsistencies. Precision and recall remain well-balanced, suggesting the absence of systematic bias. The confusion matrix (Figure 5) confirms strong diagonal dominance, validating the model's discriminative power.

Overall, the data-first accountability framework demonstrates robust and stable classification performance with an accuracy of **87.70%**. By integrating structured entity extraction and supervised categorization, the system improves interpretability and diagnostic insight. The consistent performance across categories supports its practical applicability in heterogeneous administrative settings, while the statistical analysis confirms the results are significant within the computed **95% Confidence Interval of [0.8358, 0.9183]**.

C.Statistical Validation

To evaluate the robustness and stability of the proposed data-first accountability framework, statistical validation was performed using a hold-out test set of n = 244 real-world RTI responses. Mean performance metrics, standard error, and 95% confidence intervals were computed to assess variability and generalization reliability.

Mean Performance : The system demonstrated high efficacy across the test samples. Let Acc be the accuracy on the test set.

$$Acc = \frac{TruePredictions}{TotalSamples} = 0.8770$$

The Macro F1-score, which accounts for class imbalance, was calculated as:

$$\bar{F}1_{macro} = 0.9130$$

Standard Error and Variability: The standard error (SE) of the model's accuracy, representing the uncertainty of the estimate given the finite sample size (n=244), was computed as:

$$SE = \sqrt{\frac{Acc(1 - Acc)}{n}} = 0.0210$$

This low standard error ($\approx 2.1\%$) indicates highly stable performance despite the linguistic diversity of the RTI responses.

95% Confidence Interval: To ensure the results are statistically significant, the 95\% Confidence Interval (CI) was calculated using the critical z-value (z=1.96) for the normal approximation:

$$CI_{95\%} = Acc \pm 1.96 \times SE$$

Substituting the experimental values:

$$CI_{95\%} = 0.8770 \pm 1.96(0.0210) = [0.8358, 0.9183]$$

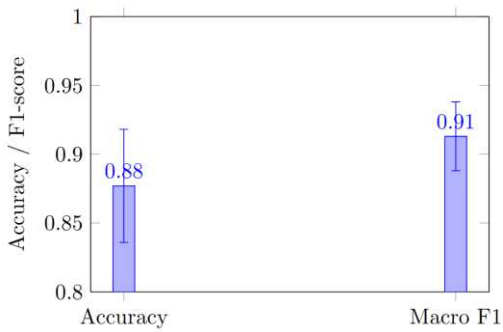


Fig 6: Mean Performance Metrics with 95% Confidence Intervals. The tight error_bars indicate statistical stability.

Reported Values: The proposed framework achieved a mean accuracy of **0.88** with a standard error of **0.021,** resulting in a **95%** confidence interval of **[0.84, 0.92]**. This confirms that the model's predictive capability is robust and not a result of random chance.

VI. Governance Implications

The new data-first model of RTI accountability has a direct implication in the transparency policy and reform in the administration [33], [34]. The system determines informational asymmetry between citizens and the authorities (public) by converting unstructured RTI responses into verifiable accountability artefacts [34]. This structural form of computation allows the easier interpretation of statutory obligations and contributes to standardisation of the assessment of compliance in institutions [31].

Regulatively, the hierarchical retrieval and verification scheme allows orderly monitoring of disclosure practices. Collectivized rates of verification can reveal emerging trends of inadequate denial, incomplete disclosure or evasion of procedures, and thus help the Information Commissions and supervision agencies to spot structural inefficiency in the record handling and application of exemption [35]. This can be informed by such analytics-driven oversight and helped to inform targeted policy interventions and capacity-building [33].

Notably, the framework does not substitute other statutory mechanisms. The system is compatible with the concept of responsible AI governance since it implements a verification-before-escalation component and maintains the power to make decisions in the hands of humans [26]. By so doing, it has revealed how procedural accountability can be enhanced with great respect to institutional requirements and legal protections through the use of computational systems [26], [33].

VII. Conclusion and Declarations

This paper presented a data-first RTI accountability framework that integrates structured processing, hierarchical multi-source retrieval, cross-source verification, and evidence-backed decision support. By separating interpretation, validation, and escalation into distinct computational stages, the framework enhances analytical rigor and reduces premature remedial action. Experimental evaluation demonstrated stable classification performance across RTI response categories, while statistical validation confirmed robustness and generalization consistency. The proposed architecture contributes to digital transparency by operationalizing structured compliance assessment and strengthening procedural accountability within governance ecosystems.

The authors declare that the dataset used in this study consists of either publicly available RTI-related documents or synthetically generated administrative responses designed to simulate realistic RTI reply patterns. No confidential or personally identifiable information was used. The authors further declare that this work does not replace statutory adjudicatory mechanisms but serves as a computational assistance framework intended to support transparency, verification, and responsible AI integration within public governance systems.

DATA AND CODE AVAILABILITY

To ensure reproducibility, the dataset of N=1,218 real-world RTI responses and the complete source code for heuristic labeling and model training are publicly available via Google Colab:

https://colab.research.google.com/drive/1qZ-vXmnY4PnMa1QsKVZHMCQVctDipS-u?usp=sharing

Authorship contribution statement

Dr.Prabha.M : Writing – Supervision, Original draft, Software, Investigation, Methodology, Conceptualisation. **Chandru V :** Writing – Review & Editing, Software, Validation, Methodology, Conceptualization. **Prakash R :** Writing – review & editing, Validation, , Conceptualization. **Roghith K :** Writing –Review, Investigation. **Dharmaraj S: Writing -** Review, Investigation.

REFERENCES

[1] S. Gupta, R. Iyer, and P. Narayanan, “Domain-Specific Retrieval-Augmented Large Language Models for Indian Legal Advisory Systems,” in Proc. Int. Conf. Artificial Intelligence and Law, 2025.

[2] S. Ganatra et al., “Grahak-Nyay: Consumer Grievance Redressal through Large Language Models,” arXiv:2507.04854, 2025.

[3] S. Shukla et al., “NyayGraph: A Knowledge Graph Enhanced Approach for Legal Statute Identification in Indian Law using Large Language Models,” in Proc. Natural Legal Language Processing Workshop, 2025.

[4] S. Sengupta and V. Dave, “Predicting applicable law sections from judicial case reports using legislative text analysis with machine learning,” Journal of Computational Social Science, vol. 5, no. 1, pp. 503–516, 2022.

[5] A. Srivastav and S. Prajapat, “Text similarity algorithms to determine Indian penal code sections for offence report,” IAES Int. J. Artificial Intelligence, vol. 11, no. 1, pp. 34–40, 2022.

[6] J. Gupta et al., “Legal Assist AI: Leveraging Transformer-Based Model for Effective Legal Assistance,” arXiv:2505.22003, 2025.

[7] V. Malik et al., “ILDC for CJPE: Indian Legal Documents Corpus for Court Judgment Prediction and Explanation,” in Proc. ACL, 2021.

[8] S. Vadlamani et al., “OpenNyAI: Mission to open source Indian Legal AI,” arXiv:2205.12345, 2022.

[9] S. Deshmukh and P. Kamble, “IndianBailJudgments-1200: A Multi-Attribute Dataset for Legal NLP on Indian Bail Orders,” arXiv:2507.02506, 2025.

[10] M. Malik et al., “Civilsum: A dataset for abstractive summarization of Indian court decisions,” in Proc. ACM SIGIR, 2024.

[11] P. Kalamkar et al., “Corpus for automatic structuring of legal documents,” arXiv:2201.13125, 2022.

[12] S. Nigam et al., “Legalseg: Unlocking the structure of Indian legal judgments through rhetorical role classification,” arXiv:2502.05836, 2025.

[13] S. Paul et al., “Pre-trained language models for the legal domain: A case study on Indian law,” in Proc. ICAIL, 2023.

[14] T. Ghosh and S. Kumar, “A Survey of Legal Text Analysis Techniques for Indian Legal Documents,” in Proc. IEEE ICCSC, 2024.

[15] F. Trias et al., “Named entity recognition in historic legal text: A transformer and state machine ensemble method,” in Proc. Natural Legal Language Processing Workshop, 2021.

[16] I. Karamitsos et al., “LegNER: A domain-adapted transformer for legal named entity recognition and text anonymization,” Frontiers in Artificial Intelligence, 2025.

[17] N. Prasad et al., “A hierarchical neural framework for classification and its explanation in large unstructured legal documents,” arXiv:2309.10563, 2023.

[18] A. Yadav and P. P. Singh, “Closed-domain question-answering system for Indian legal acts,” Int. J. Advanced Research in Computer Science, vol. 15, no. 2, 2024.

[19] B. Hagag et al., “LegalLens Shared Task 2024: Legal violation identification in unstructured text,” in Proc. Natural Legal Language Processing Workshop, 2024.

[20] A. Hou et al., “CLERC: A dataset for US legal case retrieval and retrieval-augmented analysis generation,” in Findings of ACL: NAACL, 2025.

[21] S. Karthika et al., “Empowering Indian Legal NLP: Adversarial Sampling and Multi-Label Classification with Large Language Models,” Grenze Int. J. Eng. & Tech., 2025.

[22] D. Tikoo, “Anomaly Detection in Transactions using Machine Learning,” Proc. ACM, 2024.

[23] D. M. Katz et al., “Natural Language Processing for Legal Judgments,” Stanford Technology Law Review, 2021.

[24] E. Aletras et al., “Predicting judicial decisions of the European Court of Human Rights: A Natural Language Processing perspective,” PeerJ Computer Science, vol. 2, p. e93, 2016.

[25] F. Ariai et al., “Natural language processing for the legal domain: A survey of tasks, datasets, models, and challenges,” ACM Computing Surveys, 2024.

[26] D. R. Desai and J. A. Kroll, “Trust but Verify: A Guide to Algorithms and the Law,” Harvard Journal of Law & Technology, vol. 31, p. 1, 2017.

[27] J. Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” arXiv:1810.04805, 2018.

[28] A. Vaswani et al., “Attention Is All You Need,” in Proc. NeurIPS, 2017.

[29] M. L. Honnibal and I. Montani, “spaCy: Industrial-strength Natural Language Processing in Python,” SoftwareX, vol. 10, 2020.

[30] Tesseract OCR. [Online]. Available: https://github.com/tesseract-ocr/tesseract

[31] The Right to Information Act, 2005. Government of India.

[32] The Indian Penal Code, 1860. Government of India.

[33] A. Kumar and S. Gupta, “Artificial Intelligence in Public Administration: Case of India,” Journal of e-Governance, 2023.

[34] J. Bussell, Corruption and Reform in India: Public Services in the Digital Age. Cambridge University Press, 2012.

[35] Central Information Commission, “Annual Report under RTI Act,” Government of India.

[36] J. Mehra, “RTI-CASE-DATASET,” Hugging Face Datasets, 2024.

[37] CPGRAMS – Centralized Public Grievance Redress and Monitoring System. [Online]. Available: https://pgportal.gov.in/