

AN INTELLIGENT AI-ENABLED SYSTEM FOR LARGE-SCALE PETITION ANALYSIS, REDUNDANT GRIEVANCE DETECTION, AND TRANSPARENT RESOLUTION TRACKING

R Sivaguru Dept. of CSE Knowledge Institute of Technology Salem-637504, Tamil Nadu rsgcse@kiot.ac.in

Y Indhuvadhani Dept. of CSE Knowledge Institute of Technology Salem-637504, Tamil Nadu 2k22cse053@kiot.ac.in

A Arutselvan Dept. of CSE Knowledge Institute of Technology Salem-637504, Tamil Nadu 2k22cse012@kiot.ac.in

M Gayathri Dept. of CSE Knowledge Institute of Technology Salem-637504, Tamil Nadu 2k22cse038@kiot.ac.in

K Balasakthi Dept. of CSE Knowledge Institute of Technology Salem-637504, Tamil Nadu 2k22cse015@kiot.ac.in

Abstract: Government portals and public grievance platforms receive thousands of citizen petitions every day, yet most of these systems are still handled through manual sorting, forwarding, and follow-up. This leads to slow processing, repeated complaints going unnoticed, and citizens left without any clear update on their case status. The present work addresses these gaps by developing a unified, AI-powered grievance management system that brings together several complementary technologies. A RASA-based conversational chatbot replaces traditional web forms and lets users file petitions through a natural dialogue. BERT-based transformer models handle the understanding and classification of petition text and also flag complaints that are semantically identical to earlier submissions. XGBoost is used to rank grievances by urgency, taking into account how often similar issues have appeared, how long past cases took to resolve, and current departmental load. When citizens attach photographs as supporting evidence, a ResNet-based deep learning model analyses those images automatically to check their relevance and weed out duplicates. Finally, a live dashboard tracks every petition from submission to resolution, giving both citizens and administrators clear visibility into the process. Tests on real and simulated grievance data show that this integrated approach outperforms conventional systems in detecting redundant complaints, prioritising urgent cases, and keeping stakeholders informed.

Keywords: Artificial Intelligence, RASA Chatbot, BERT, XGBoost, ResNet, Petition Analysis, Grievance Redressal System, e-Governance.

1. INTRODUCTION

Good governance depends on the ability of public institutions to hear and respond to citizen concerns in a timely and organised way. As government services have shifted online, the volume of petitions reaching administrative portals has grown sharply — many departments now deal with hundreds or even thousands of new complaints each working day. Despite this growth in volume, the underlying processes in most existing systems have not kept pace. Officials still manually read, categorise, and forward grievances, which creates bottlenecks, delays resolution, and makes it almost impossible to detect when the same complaint is submitted by multiple people. The result is an overloaded administration and frustrated citizens who see no progress on their cases [1], [2].

A core difficulty is the free-text nature of petitions. People describe the same problem using entirely different words, levels of detail, and sentence structures. Without a system that understands meaning rather than just matching keywords, duplicate or near-duplicate complaints pass through undetected, multiplying the workload without adding new information. Equally problematic is the absence of any smart ordering mechanism — whether a complaint is urgent or routine, it often joins the same queue and waits its turn [3], [4].

Researchers have explored several computational approaches to tackle these issues. Early systems relied on term-frequency methods such as TF-IDF combined with Naïve Bayes classifiers. While these worked reasonably well on small, clean datasets, they broke down when complaints were paraphrased or written informally [5]. Recurrent networks, particularly LSTM architectures, captured more context and improved classification accuracy, but they struggled with long passages and were slow to train at scale [6]. The arrival of transformer-based language models, especially BERT, changed things considerably. By encoding each word in relation to its full surrounding context, BERT can recognise that two differently worded complaints are in fact about the same issue — a crucial capability for any serious duplicate-detection system [7].

On the prioritisation side, gradient boosting methods have shown strong results in decision-support applications. XGBoost in particular handles large, imbalanced datasets well and can model the kind of multi-factor interactions needed to score grievances by urgency [8]. Meanwhile, the RASA framework has made it practical to replace static complaint forms with chatbots that guide users conversationally, collect structured data, and handle follow-up questions automatically [9]. Where citizens also attach photographs — pictures of a broken road, a flooded area, or a faulty installation — convolutional networks based on the ResNet architecture have proven effective at validating and filtering such image evidence [10].

What is missing from existing deployments is a system that combines all of these components under one roof. Most current solutions treat text classification, chatbot interaction, prioritisation, and image analysis as separate concerns. This paper presents a unified framework that integrates all four, together with a real-time resolution-tracking dashboard, into a single coherent platform. The sections that follow review related work, describe the shortcomings of current systems, set out the proposed architecture, detail the implementation, and report experimental results.

2. LITERATURE SURVEY

The management of citizen grievances through digital platforms has attracted growing research interest as e-Governance adoption has widened. Early work on automated complaint handling drew heavily on classical text mining. TF-IDF-weighted features fed into Naïve Bayes or support vector machine classifiers could sort complaints into broad categories, but these models had little tolerance for the spelling variations, abbreviations, and informal phrasing common in real petitions [5]. Performance degraded quickly on large or multilingual datasets.

The transition to neural methods brought meaningful gains. LSTM-based models learned sequential dependencies in text and were applied to both classification and sentiment scoring of complaints [6]. Their weakness was handling very long grievance descriptions, where the relevant information might appear far apart in the text, causing the model to lose track of the context. Attention mechanisms helped to some extent, but it was the full transformer architecture that proved most effective.

BERT-based models represented a step change. Pre-trained on massive corpora and fine-tuned on domain-specific data, BERT produces rich contextual embeddings that capture synonymy and paraphrase. Studies applying BERT to grievance datasets found it substantially more accurate at identifying duplicate submissions than any previous method [3], [7]. This is especially valuable in e-Governance settings where the same road pothole or waterlogging problem gets reported independently by dozens of residents within days of each other.

For deciding how to rank and route complaints, machine learning classifiers have been explored. Decision trees and Random Forest models offer interpretability but tend to underperform on imbalanced data where urgent grievances are outnumbered by routine ones. XGBoost addresses these weaknesses through gradient-boosted ensembles with built-in regularisation, and has been shown to handle complex, multi-feature prioritisation tasks reliably [8].

Conversational interfaces have been studied separately as a way of improving the user experience. RASA, an open-source dialogue framework, allows developers to build domain-specific chatbots that classify user intent, extract key entities, and maintain context across multiple turns [9]. Deployed chatbots have been shown to reduce form-filling errors and improve completion rates compared with traditional web forms. Image evidence validation using convolutional networks has also been explored in adjacent domains such as infrastructure monitoring and public safety, with ResNet architectures consistently achieving strong accuracy on visual classification tasks [10].

What the literature lacks is a single system that pulls these strands together. Individual papers have addressed chatbots, or BERT-based classification, or XGBoost-based prioritisation in isolation. A fully integrated platform that covers all stages — submission, deduplication, prioritisation, image validation, and resolution tracking — has not yet been demonstrated at scale, which is the gap this work aims to fill [2].

3. EXISTING SYSTEM

Most grievance redressal portals in operation today follow a broadly similar workflow. Citizens fill in a web or mobile form describing their complaint, optionally attach photographs, and submit. The submission lands in an inbox where a human officer reads it, decides which department it belongs to, and forwards it manually. The officer may also assign a rough priority based on personal judgement. Resolution updates, if any, are posted manually and are often sparse.

This model was adequate when petition volumes were low and departments were small. At modern scale it creates several recognisable problems. Manual categorisation is slow and inconsistent — different officers apply different criteria, and the sheer number of incoming submissions means some petitions wait days before anyone even reads them. There is no systematic check for duplicates, so the same complaint might be forwarded to the same department multiple times and resolved multiple times, wasting resources [2].

Text mining approaches that were introduced in earlier automation attempts helped to some degree. Keyword matching and TF-IDF-based classifiers could auto-route complaints with moderate accuracy, but they failed on anything paraphrased or informally written. A complaint describing a "pothole on Main Street" and another about a "big hole in the road near the

market" would be treated as entirely different issues [5]. LSTM-based classifiers improved on this by capturing more context, but still struggled with long or complex complaints [6]. More recently some portals have adopted BERT-based classifiers for routing, which perform better, but these deployments typically focus only on categorisation and do not tackle deduplication [7].

Prioritisation in current systems is generally handled by fixed rules — complaints tagged with certain keywords like "emergency" or "urgent" jump the queue, while everything else follows arrival order. This ignores factors like how many people have raised the same issue or how long resolution typically takes for a given complaint type. Machine learning prioritisation models such as XGBoost have been tested in research settings but have not been widely deployed in production grievance platforms [4], [8].

Citizen-facing interfaces remain static. Some portals have experimented with simple chatbots for initial triage, but these are typically shallow FAQ bots rather than full dialogue systems integrated with the backend analytics [9]. Image evidence, where accepted, is reviewed manually. Even when convolutional networks have been demonstrated to automate this step effectively, most deployed platforms have not adopted them [10]. The overall picture is a fragmented landscape where useful technologies exist but have not been combined into a single, end-to-end solution.

4. PROPOSED SYSTEM

The system proposed here is built around the idea that the five key steps in grievance handling — submission, text analysis, deduplication, prioritisation, image validation, and tracking — should be treated as a single pipeline rather than as separate modules bolted together. Each component feeds data to the next, and the whole system shares a common database so that nothing is lost or duplicated across the handoffs.

At the front end, citizens interact with a RASA-based chatbot deployed on both web and mobile interfaces. Rather than presenting a blank form, the chatbot opens a conversation, asks targeted questions to collect the petition type, location, urgency indicators, and supporting details, and clarifies ambiguous answers before finalising the submission. This approach yields more complete and consistently structured data than self-completed forms, which in turn improves the performance of the downstream analytics [9].

Once a petition is submitted, its text is processed by a fine-tuned BERT model. The model produces a sentence embedding — a numerical representation of the semantic content — and uses this for two tasks simultaneously: classifying the petition into the appropriate department category, and checking whether any semantically similar petition is already on record. Similarity is measured using cosine distance between embeddings. If the new petition falls within a configurable threshold of an existing one, it is flagged as a probable duplicate and linked to the original rather than creating a fresh case. This single step can substantially cut the administrative workload in scenarios where a widespread local problem generates dozens of independent reports [7].

Grievances that pass the deduplication check, or that are confirmed as unique, are then scored by the XGBoost priority model. The model takes as input a feature vector that includes urgency keywords from the text, the frequency with which similar complaints have appeared over the past weeks, the average resolution time historically associated with that complaint category, current departmental workload, and any prior citizen feedback scores. The model outputs one of three priority levels — high, medium, or low — which determines where the petition is placed in the administrative queue [8].

Where a citizen has attached photographs, the ResNet-based image analyser processes these in parallel with the text pipeline. Images are resized, normalised, and passed through the network, which extracts high-level visual features and checks whether the image is relevant to the stated complaint type, whether it duplicates an image already associated with an existing case, and whether it meets a minimum quality threshold for administrative use. Photographs that fail these checks are flagged for manual review rather than being silently discarded, so the citizen is informed if resubmission is needed [10].

Every petition, whether new or linked to an existing case, is tracked through a real-time dashboard accessible to both citizens and administrators. Citizens can see the current status of their submission, the expected resolution timeline based on historical data, and any requests for additional information from the handling officer. Administrators see an aggregated view of departmental performance, backlog levels, and resolution rates. This transparency is the piece most consistently missing from current systems, and it is what turns the platform from a submission tool into a genuine accountability mechanism [1].

5. IMPLEMENTATION AND METHODS

The system is built as a modular architecture where each component — chatbot, NLP engine, prioritisation model, image analyser, and dashboard — operates as a distinct service communicating through a shared database and a lightweight API layer. This design makes it straightforward to update or replace individual components without rebuilding the whole system.

A. System Architecture and Workflow

The architecture is organised into four layers. The user interaction layer hosts the RASA chatbot and the citizen-facing dashboard. Below it, the processing layer handles incoming data: text is normalised (lowercased, punctuation-stripped, stop-words removed) and images are resized to a standard resolution before being passed downstream. The analytics layer runs the BERT embedding service, the XGBoost scorer, and the ResNet image model. Finally, the presentation layer serves the administrator dashboard and generates status notifications for citizens. Data flows unidirectionally from submission through analysis to resolution, with the database acting as the shared state store across all layers.

B. Chatbot-Based Petition Submission

The RASA framework was chosen for the chatbot because it supports fully local deployment, meaning sensitive citizen data never has to leave the institution's infrastructure. The chatbot is trained on a domain-specific NLU dataset covering intent categories such as `complaint_submission`, `status_enquiry`, and `clarification_response`. Entity extraction pulls out key slot values — location, complaint type, urgency — from free-form user input. RASA Core manages multi-turn dialogue, allowing the chatbot to ask follow-up questions when required slots are missing and to handle unexpected turns in the conversation gracefully.

C. Textual Analysis and Duplicate Grievance Detection

A BERT-base model pre-trained on English text was fine-tuned for 10 epochs on a labelled grievance dataset covering eight common complaint categories. The fine-tuned model is used both for classification and for generating 768-dimensional sentence embeddings via mean-pooling over the final hidden layer. For deduplication, incoming embeddings are compared against a stored index of embeddings from the past 90 days using cosine similarity. A threshold of 0.88 was selected through cross-validation; petitions above this threshold are marked as linked rather than standalone, and the citizen is informed that a similar case is already being handled.

D. Grievance Prioritisation Using XGBoost

The XGBoost model was trained on 12 months of historical grievance records labelled by experienced officers. Features include TF-IDF urgency scores derived from the complaint text, a 7-day rolling count of similar complaints, the median resolution time for the predicted category, the current open-case count for the target department, and the submitting citizen's historical feedback score on past resolutions. The model was tuned using 5-fold cross-validation and optimised for macro-F1 across the three priority classes. In testing, it achieved around 84% macro-F1, substantially higher than the rule-based baseline at 61%.

E. Image-Based Evidence Validation

A ResNet-50 backbone pre-trained on ImageNet was fine-tuned on a labelled dataset of grievance images covering categories such as road damage, drainage blockage, building defects, and irrelevant or inappropriate content. Input images are resized to 224×224 pixels, normalised by channel mean and standard deviation, and augmented during training with random flips and colour jitter. The network outputs both a relevance label and a confidence score; images with a relevance confidence below 0.7 are routed to manual review. A perceptual hashing step run before the network catches near-identical duplicate images without consuming GPU resources.

F. Resolution Tracking and Dashboard Module

The dashboard is built as a web application that queries the grievance database in real time. Citizens who log in with their submission reference can see the current stage of their petition, the name and department of the assigned officer, and an estimated completion date calculated from historical resolution times for similar cases. The administrator view shows heat maps of complaint density by area, bar charts of open versus resolved cases by department, and alert flags for cases that have exceeded expected resolution time. All views update automatically every five minutes without requiring a page refresh.

G. Experimental Setup and Evaluation Metrics

The system was evaluated on two datasets: a real-world set of 8,400 petitions drawn from a municipal grievance portal, and a synthetic set of 5,000 complaints generated by paraphrasing real submissions to create controlled duplicate pairs. The BERT classifier achieved 91.3% accuracy on the real dataset and detected duplicates with a precision of 89.7% and a recall of 86.4%, compared to a TF-IDF cosine baseline of 74.2% precision and 68.9% recall. The XGBoost prioritiser reached 84.1% macro-F1 versus the rule-based baseline's 61.3%. User engagement metrics from a pilot deployment showed that petition completion rates through the chatbot interface were 23 percentage points higher than through the previous web form. These results collectively indicate meaningful improvements across all key dimensions of grievance management.

6. CONCLUSION AND FUTURE WORK

A. Conclusion

This paper set out to address a genuine operational problem: the gap between the large volumes of citizen grievances that modern e-Governance platforms receive and the limited capacity of manual or semi-automated systems to handle them well. The proposed system closes this gap by combining four AI components — a RASA conversational chatbot, a BERT-based text analyser and deduplicator, an XGBoost prioritisation model, and a ResNet image validator — with a real-time tracking dashboard, all within a single integrated pipeline.

The experimental results confirm that this combination delivers real benefits. Duplicate detection accuracy is substantially higher than keyword-based baselines, the prioritisation model reliably identifies urgent cases that fixed rules would miss, and the chatbot interface measurably improves petition completion rates. The tracking dashboard addresses the transparency deficit that citizens consistently cite as a major source of dissatisfaction with existing grievance portals. Together, these improvements point toward a governance model where citizens can trust that their complaints will be heard, fairly prioritised, and visibly acted upon.

B. Future Work

Several directions remain open for further development. Extending the system to handle petitions in multiple languages — particularly Indian regional languages — would broaden its applicability significantly, and multilingual BERT variants offer a clear path toward this goal. Privacy and data security are increasingly important given tightening data protection regulations; federated learning approaches, where models train on data that never leaves the local device or server, deserve investigation in this context.

On the analytics side, integrating sentiment scoring into the prioritisation model could help distinguish between complaints that are factually urgent and those where the citizen's tone signals particular distress, even if the underlying issue appears routine. Predictive models that anticipate seasonal or event-driven spikes in certain complaint types — waterlogging after heavy rain, power outages during festivals — could allow administrators to prepare resources in advance. Blockchain-based audit trails would provide tamper-evident records of every action taken on a grievance, strengthening accountability further. Finally, a reinforcement-learning chatbot that adapts its dialogue strategy based on past interaction outcomes could progressively improve the quality and completeness of petition data collected at submission. A large-scale longitudinal deployment is planned to assess how these enhancements affect system robustness and governance outcomes over time.

REFERENCES

- [1] World Bank Group, "Using Artificial Intelligence to Improve Public Service Delivery," Technical Report, 2021.
- [2] P. Gupta, M. Kumar, and R. Agrawal, "AI-Based Smart Grievance Redressal System for E-Governance Applications," *International Journal of e-Government Research*, vol. 16, no. 3, pp. 1-18, 2020.
- [3] R. Mihalcea, C. Corley, and C. Strapparava, "Corpus-Based and Knowledge-Based Measures of Text Semantic Similarity," in *Proc. AAAI*, 2006.
- [4] N. Kumar and R. Ranjan, "Sentiment Analysis and Complaint Classification for Public Grievance Systems," in *Proc. IEEE Int. Conf. on Smart Computing*, 2019.
- [5] A. McCallum and K. Nigam, "A Comparison of Event Models for Naive Bayes Text Classification," in *Proc. AAAI Workshop*, 1998.
- [6] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, 1997.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proc. NAACL-HLT*, 2019.
- [8] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proc. ACM SIGKDD*, 2016.
- [9] A. Bocklisch et al., "Rasa: Open Source Language Understanding and Dialogue Management," arXiv:1712.05181, 2017.
- [10] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE CVPR*, 2016.