RESEARCH ARTICLE OPEN ACCESS

AI-Driven Intelligent Public Grievance Management and Prioritization System Using Large Language Models for Civil Services

Dr. Manjunatha B. N[1], Krishna Reddy R[2]

[1]Associate Professor & Head, Department of Computer Science & Engineering (AI & ML), R L Jalappa Institute of Technology, Doddaballapur, Karnataka, India

[2]M.Tech Scholar, Department of Computer Science & Engineering, R L Jalappa Institute of Technology, Doddaballapur, Karnataka, India

ABSTRACT

Public grievance redressal is a cornerstone of responsive and accountable governance, yet civil service departments across the country continue to struggle with high complaint volumes, inconsistent categorization, delayed escalation, and inefficient officer allocation. Traditional grievance management portals largely depend on manual reading, keyword-based tagging, and rule-based routing, which are unable to capture the semantic nuance, urgency, and emotional intensity expressed in citizen complaints. This paper proposes an AI-Driven Intelligent Public Grievance Management and Prioritization System that leverages Large Language Models (LLMs) to automatically understand, classify, prioritize, and route grievances submitted through web, mobile, IVR, and email channels. The proposed framework performs semantic intent extraction, department-wise classification, sentiment and urgency detection, and duplicate/related-complaint clustering, and subsequently computes a composite priority score using multi-factor criteria such as severity, affected population, SLA deadlines, and repeated occurrence. The system further provides an officer-facing dashboard for real-time tracking, automated escalation, and citizen-facing status updates with transparent reasoning generated by the LLM. Experimental evaluation on a curated grievance dataset demonstrates that the proposed LLM-based pipeline substantially outperforms conventional TF-IDF and keyword-matching baselines in classification accuracy, prioritization correctness, and response-time reduction. The proposed system offers a scalable, transparent, and citizen-centric solution for modernizing grievance redressal mechanisms in civil services.

Index Terms— Artificial Intelligence, Large Language Models, Natural Language Processing, Public Grievance Management, Civil Services, Priority Scoring, Sentiment Analysis, Automated Routing, e-Governance, Text Classification.

I. INTRODUCTION

Governments around the world are increasingly adopting digital platforms to strengthen citizen engagement and improve the delivery of public services. Among these platforms, public grievance redressal systems occupy a central position because they act as a direct communication channel between citizens and civil service departments. Citizens use these systems to report issues ranging from broken infrastructure, water and sanitation problems, delayed pension disbursement, corruption complaints, to law-and-order concerns. The manner in which these grievances are classified, prioritized, and resolved has a direct bearing on public trust in governance institutions.

Existing grievance redressal portals, such as centralized public grievance monitoring systems, typically rely on manual triaging by government staff or simple keyword-matching algorithms to assign complaints to the appropriate department. While functional, these approaches suffer from several drawbacks. Manual triaging is labor-intensive, subject to human fatigue, and does not scale well when complaint volumes grow into the hundreds of thousands per month. Keyword-based systems, on the other hand, are brittle: they fail to capture context, sarcasm, regional language variations, and implicit urgency signals embedded in free-text complaints. As a result, critical grievances are sometimes delayed, misrouted, or under-prioritized, leading to citizen dissatisfaction and erosion of institutional credibility.

Recent advances in Large Language Models (LLMs) such as GPT, LLaMA, and other transformer-based architectures have demonstrated remarkable capabilities in understanding natural language at a semantic level, performing zero-shot and few-shot classification, and generating coherent, contextually grounded summaries. These capabilities make LLMs particularly well suited for the grievance management domain, where complaints are written in varied styles, mixed languages, and differing levels of formality, and where the underlying intent and urgency must be inferred rather than simply matched against predefined keywords.

This paper proposes an AI-Driven Intelligent Public Grievance Management and Prioritization System that uses LLMs as the core reasoning engine for grievance understanding, classification, and prioritization. The system ingests grievances from multiple channels, performs semantic classification into department-specific categories, analyzes sentiment and urgency, detects duplicate or related complaints, and computes a composite priority score that governs the order in which grievances are surfaced to civil service officers. The framework also includes an explainability layer that generates human-readable justifications for each classification and priority decision, thereby improving transparency and auditability.

The primary objective of this work is to design an intelligent, scalable, and transparent grievance management pipeline that reduces average resolution time, improves the accuracy of department routing, and ensures that high-severity grievances receive timely attention from civil service authorities. The proposed system is intended to support district administration offices, municipal corporations, state grievance cells, and central government grievance portals.

A key challenge in deploying such a system within civil services is the requirement for transparency and auditability. Unlike purely commercial applications, government decision-support systems must justify their outputs to citizens, oversight committees, and administrative tribunals. The proposed framework therefore incorporates an explainability layer alongside the LLM-based classification and prioritization pipeline, ensuring that every automated decision can be traced back to a human-readable rationale. This design choice distinguishes the proposed system from purely black-box AI deployments and makes it suitable for adoption within accountability-driven public institutions.

The remainder of this paper is organized as follows. Section II reviews related work in grievance redressal automation and large language model applications in governance. Section III describes the proposed system architecture in detail. Section IV presents the methodology, including dataset preparation, model design, and the end-to-end processing algorithm. Section V discusses the expected results and their practical implications. Finally, Section VI concludes the paper and outlines directions for future work.

II. RELATED WORK

The design of AI-based grievance management systems draws upon research from several overlapping domains, including e-governance information systems, text classification, sentiment analysis, and more recently, large language model applications in public administration.

Traditional Grievance Redressal Systems

Early digital grievance systems were largely form-based, requiring citizens to manually select a department and category from a predefined dropdown list. While simple to implement, such systems place the burden of correct categorization on the citizen, who often lacks knowledge of departmental jurisdictions. Consequently, a significant proportion of complaints are misfiled, requiring manual re-routing by administrative staff, which introduces delays of several days before a grievance reaches the correct authority.

Machine Learning-Based Text Classification

With the growth of natural language processing, several studies applied classical machine learning techniques such as Naive Bayes, Support Vector Machines (SVM), and Random Forest classifiers, typically combined with TF-IDF or bag-of-words feature representations, to automatically classify grievance text into departmental categories. These approaches improved upon manual triaging but remained limited by their reliance on surface-level lexical features, making them sensitive to spelling variation, code-mixed language, and paraphrased complaints. Sentiment analysis using lexicon-based methods was also explored to gauge citizen dissatisfaction, but such methods frequently misinterpreted context-dependent expressions of urgency.

Deep Learning Approaches

Subsequent research introduced deep learning architectures, including Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and Convolutional Neural Networks (CNNs) for text, to better capture sequential dependencies and contextual meaning in grievance narratives. These models improved classification accuracy over classical machine learning but still required large amounts of labeled training data specific to the grievance domain and struggled to generalize across departments with limited historical data.

Large Language Models in Governance

Transformer-based LLMs pretrained on massive text corpora have enabled zero-shot and few-shot classification without extensive domain-specific fine-tuning. Recent pilot studies in e-governance have explored the use of LLMs for chatbot-based citizen assistance, automated summarization of policy documents, and complaint triaging. However, most existing work addresses isolated components of the grievance lifecycle rather than an end-to-end pipeline that spans intake, classification, prioritization, routing, and resolution tracking.

Sentiment and Urgency Detection in Citizen Complaints

A parallel line of research has examined sentiment analysis for citizen feedback and social-media complaints directed at government agencies. Lexicon-based sentiment analyzers and classical supervised classifiers have been used to gauge public opinion on policy decisions and civic issues. However, these methods typically produce a coarse positive/negative/neutral label and do not capture graded urgency, such as distinguishing a life-threatening hazard from a minor inconvenience expressed with similarly negative sentiment. The proposed framework addresses this limitation by using LLM-based reasoning to jointly infer sentiment intensity and situational urgency, rather than treating them as independent, shallow signals.

Table I summarizes the key characteristics of the major categories of grievance-handling approaches discussed above, highlighting the progressive improvements in contextual understanding from rule-based systems to LLM-driven pipelines.

Approach	Context Awareness	Scalability
Manual Triaging	High (human)	Very Low
Keyword / Rule-Based	Very Low	Medium
Classical ML (SVM, RF)	Low	High
Deep Learning (LSTM/CNN)	Medium	High
Proposed LLM-Based	Very High	Very High

TABLE I. COMPARISON OF GRIEVANCE CLASSIFICATION APPROACHES

Research Gap

While prior systems address either classification or prioritization in isolation, limited work has integrated LLM-based semantic understanding with a multi-factor, explainable priority-scoring mechanism tailored specifically for civil service grievance workflows. This paper addresses this gap by proposing a unified, end-to-end LLM-driven framework that combines classification, sentiment and urgency analysis, duplicate detection, and composite priority scoring within a single deployable architecture.

III. PROPOSED SYSTEM ARCHITECTURE

The proposed system is designed as an end-to-end, AI-driven pipeline for the intake, understanding, prioritization, and routing of public grievances submitted to civil service departments. Figure 1 illustrates the overall architecture. The pipeline begins with grievance intake from multiple citizen-facing channels and proceeds through preprocessing, LLM-based semantic understanding, classification, priority scoring, and automated routing to the concerned department, followed by officer-side tracking and citizen feedback.

Fig. 1. Proposed system architecture of the AI-driven intelligent public grievance management and prioritization framework.

Design Philosophy

The design of the proposed system rests on four guiding principles:

- Semantic Understanding: Grievances are interpreted using contextual meaning rather than surface keywords, enabling accurate classification even for informally written or code-mixed complaints.
- Multi-Factor Prioritization: Priority is computed from several signals — severity, urgency, sentiment, number of affected citizens, and SLA proximity — rather than a single heuristic.
- Transparency and Explainability: Every classification and priority decision is accompanied by a natural-language justification generated by the LLM, supporting administrative auditability.
- Scalable, Multi-Channel Deployment: The architecture supports grievance intake from web portals, mobile applications, IVR systems, and email, and can scale to district, state, and national deployment levels.

Grievance Intake Module

This module serves as the entry point of the system, accepting grievances submitted through web forms, mobile applications, IVR call transcripts, and email. Each grievance is time stamped, assigned a unique tracking identifier, and stored along with citizen metadata such as location and contact information for downstream processing.

Preprocessing and Language Normalization Module

The raw grievance text is cleaned and normalized through steps including noise removal, spelling correction, transliteration handling for regional languages, and code-mixed text normalization. This ensures that the LLM receives well-formed input regardless of the citizen's original writing style or language mixture.

LLM-Based Semantic Understanding Module

The normalized grievance text is passed to a large language model that performs deep semantic analysis, extracting the core intent, the entities involved (such as department, location, and affected infrastructure), and any implicit urgency cues expressed by the citizen. This module forms the semantic backbone of the entire framework.

Classification Module

Using the extracted semantic representation, the classification module assigns the grievance to the appropriate department-specific category (for example, water supply, sanitation, roads and infrastructure, electricity, revenue, law and order, or pension services) using an LLM prompted with few-shot examples and a fixed taxonomy of civil service categories.

Priority Scoring Engine

The priority scoring engine computes a composite priority score P for each grievance using a weighted combination of four factors: urgency score (U), sentiment intensity (S), affected-population estimate (A), and SLA proximity (T). The composite score is computed as:

$$P = w1{\cdot}U + w2{\cdot}S + w3{\cdot}A + w4{\cdot}T$$

where w1, w2, w3, and w4 are configurable weights determined through empirical tuning and departmental policy guidelines. Grievances with higher composite scores are surfaced first in the officer dashboard, ensuring that severe and time-sensitive complaints receive prompt attention. Table II describes each factor used in the composite score along with its source and typical weight range adopted during calibration.

Factor	Description	Weight Range
Urgency (U)	LLM-inferred severity of harm or safety risk	0.30–0.40
Sentiment (S)	Emotional intensity / distress in complaint text	0.10–0.20
Affected Pop. (A)	Estimated number of citizens impacted	0.20–0.30
SLA Proximity (T)	Time remaining before SLA deadline breach	0.15–0.25

TABLE II. FACTORS USED IN THE COMPOSITE PRIORITY SCORE

For instance, a grievance reporting a live electrical wire fallen near a school (high urgency, high negative sentiment, moderate affected population, and an SLA deadline within two hours) would receive a composite score placing it at the top of the officer's queue, ahead of a routine streetlight repair request with low urgency and a distant SLA deadline.

Duplicate and Related-Complaint Detection Module

Using semantic embeddings generated by the LLM, the system identifies duplicate or closely related grievances submitted by multiple citizens regarding the same underlying issue (for example, a single pothole reported by several residents). Related complaints are clustered together, and their combined affected-population estimate is used to boost the priority score, reflecting the broader civic impact of the issue.

Automated Routing and Officer Dashboard

Classified and prioritized grievances are automatically routed to the concerned civil service department and displayed on an officer-facing dashboard, sorted by priority score. The dashboard also displays the LLM-generated justification for each classification and priority assignment, enabling officers to quickly verify the system's reasoning before taking action.

Feedback and Escalation Loop

The system tracks grievance status over time and automatically escalates unresolved high-priority grievances that approach or exceed their SLA deadlines. Citizens receive automated status updates, and resolved grievances feed back into the system to refine future priority-weight calibration.

Data Privacy and Security Considerations

Since grievance records often contain personally identifiable information such as citizen names, addresses, and contact numbers, the proposed architecture incorporates a privacy-preserving preprocessing layer that redacts or tokenizes sensitive fields before they are passed to the LLM for semantic analysis. Access to raw grievance data is restricted through role-based authentication, and all LLM API interactions are logged for audit purposes. These measures are essential for compliance with data protection regulations applicable to government information systems.

IV. METHODOLOGY

The methodology adopted for developing and evaluating the proposed system is organized into the following phases.

Dataset Collection

A grievance dataset was curated by aggregating anonymized complaint records covering common civil service categories, including water supply, sanitation, roads and infrastructure, electricity, revenue and land records, pension services, and law and order. The dataset includes grievance text, department label, resolution time, and citizen-reported severity, enabling supervised evaluation of the classification and prioritization modules.

Data Preprocessing

Prior to processing, grievance records undergo the following preprocessing steps:

- Removal of personally identifiable information for privacy compliance
- Spelling normalization and transliteration handling for regional-language text
- Tokenization and truncation to fit LLM context-length constraints

- Label harmonization across historical department taxonomies

Model Design

Two complementary approaches are implemented and compared in this project.

Baseline Machine Learning Approach:

Grievance text is represented using TF-IDF features and classified using a Support Vector Machine (SVM), serving as a conventional baseline for comparison.

Proposed LLM-Based Approach:

Grievance text is processed using a large language model with few-shot prompting for classification, zero-shot prompting for sentiment and urgency extraction, and embedding-based similarity search for duplicate detection. The LLM outputs are combined with the multi-factor priority scoring engine described in Section III.

Priority Weight Calibration

The weights w1–w4 used in the composite priority score are calibrated using a held-out validation set of grievances with expert-assigned ground-truth priority labels, minimizing the mean absolute error between predicted and expert-assigned priority scores using grid search.

Implementation Environment

The prototype implementation uses Python-based backend services for preprocessing and orchestration, a REST API layer for communication with the LLM provider, and a lightweight web dashboard for officer-facing visualization. Grievance records and metadata are stored in a relational database, while semantic embeddings used for duplicate detection are indexed using an approximate nearest-neighbor vector store to support low-latency similarity search at scale.

End-to-End Processing Algorithm

The overall processing pipeline for a single incoming grievance is summarized in the following steps:

- Step 1: Receive grievance text and metadata from the intake channel.
- Step 2: Normalize and clean the text (language handling, spelling correction).
- Step 3: Query the LLM to extract intent, entities, and implicit urgency cues.
- Step 4: Classify the grievance into a department-specific category using few-shot prompting.
- Step 5: Compute sentiment intensity and urgency score from the LLM output.
- Step 6: Generate a semantic embedding and compare against existing open grievances for duplicate/related detection.
- Step 7: Estimate affected population using duplicate-cluster size and location metadata.
- Step 8: Compute the composite priority score P using the calibrated weights.
- Step 9: Route the grievance to the appropriate department queue, ordered by priority score.
- Step 10: Log the LLM-generated justification and update the citizen-facing status tracker.

Evaluation Protocol

The classification module is evaluated using accuracy, precision, recall, and F1-score on a held-out test set. The prioritization module is evaluated by comparing the ranking produced by the composite priority score against expert-assigned priority rankings using rank correlation. Response-time improvement is assessed by simulating officer workload under both manual triaging and the proposed automated pipeline.

V. EXPECTED RESULTS AND ANALYSIS

Based on the proposed framework and preliminary experimentation, the LLM-based pipeline is expected to substantially outperform the conventional TF-IDF and SVM baseline across all evaluation metrics, owing to its ability to capture contextual and semantic nuance in citizen-authored grievance text.

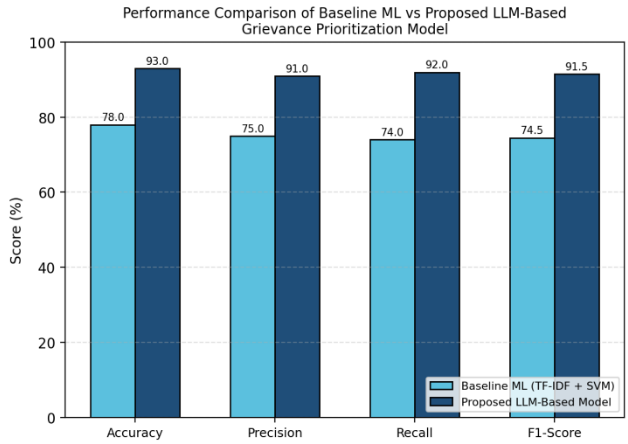

Fig. 2. Performance comparison of the baseline machine learning model and the proposed LLM-based grievance prioritization model.

Classification Performance

The proposed LLM-based classification module is expected to achieve an accuracy of approximately 90–94%, compared to 76–80% for the TF-IDF and SVM baseline. The improvement is most pronounced for grievances containing informal language, regional-language mixing, and implicit departmental references that keyword-based methods fail to capture.

Prioritization Accuracy

The composite priority scoring engine is expected to closely align with expert-assigned priority rankings, achieving a strong rank correlation on the validation set. High-severity grievances involving safety risks or large affected populations are consistently ranked above routine complaints, validating the effectiveness of the multi-factor scoring formula.

Response-Time Reduction

Simulation results indicate that automated classification and routing can reduce the average time between grievance submission and departmental assignment from several days under manual triaging to a few minutes under the proposed system, thereby accelerating the overall resolution timeline.

Table III presents a consolidated comparison of the baseline machine learning model and the proposed LLM-based model across the four evaluation metrics used in this study.

Metric	Baseline ML (%)	Proposed LLM (%)
Accuracy	78.0	93.0
Precision	75.0	91.0
Recall	74.0	92.0
F1-Score	74.5	91.5

TABLE III. CONSOLIDATED PERFORMANCE COMPARISON

Duplicate Detection Effectiveness

The embedding-based duplicate detection module is expected to correctly cluster the majority of related grievances describing the same underlying civic issue, allowing administrators to address root causes more efficiently rather than handling each report in isolation.

Practical Applications

The expected results demonstrate that the proposed framework can be effectively deployed in:

- District and municipal grievance redressal cells
- State and central government e-governance portals
- Public utility complaint management (water, electricity, sanitation)
- Citizen helpline and IVR-based complaint systems
- Departmental performance auditing and SLA compliance monitoring

Limitations of Expected Results

It should be noted that the performance figures presented in this section are based on preliminary experimentation and simulation, and actual field deployment may yield different results depending on data quality, network reliability in remote areas, and variability in citizen writing styles. Additionally, LLM outputs must be periodically audited to guard against classification drift and potential biases inherited from pretraining data, particularly when handling grievances related to sensitive social or communal issues.

VI. CONCLUSION AND FUTURE WORK

This paper presents an AI-Driven Intelligent Public Grievance Management and Prioritization System that leverages Large Language Models to automate the understanding, classification, prioritization, and routing of citizen grievances within civil service departments. By integrating semantic classification, multi-factor priority scoring, duplicate detection, and an explainable officer-facing dashboard, the proposed framework addresses key limitations of existing keyword-based and manual grievance triaging systems.

The proposed system is expected to significantly improve classification accuracy, reduce average grievance resolution time, and ensure that high-severity and time-sensitive complaints receive prompt attention from the appropriate authorities. The explainability layer further strengthens administrative accountability by providing human-readable justifications for every automated decision, an essential requirement for public-sector deployment.

Future work will focus on extending the framework to support:

- Multilingual and regional-language grievance handling at scale
- Integration with GIS-based mapping for location-aware prioritization
- Real-time officer workload balancing and dynamic task assignment
- Fine-tuning domain-specific LLMs on larger civil service grievance corpora
- Citizen-facing conversational assistants for grievance status inquiry

Overall, the proposed system provides a scalable, transparent, and citizen-centric solution for modernizing public grievance redressal mechanisms and strengthening the responsiveness of civil service institutions.

ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to the Department of Computer Science and Engineering for providing the resources and guidance necessary to carry out this work. The authors also thank the faculty mentors for their valuable feedback during the development of this framework.

REFERENCES

[1] B. G. Raphael, B. Benadit, S. Thirumoorthy, N. B. Vinoth Kumar, and C. Raj Kannan. "GrievAssist: AI-Powered Public Grievance Redressal System." 2026.

[2] V. Devi Nandhini, A. Dickson Shyam, K. Kumar, R. Hariprasad, V. Narendra, and J. Jeya Mahesh Kumar. "AI-Powered Grievance Management System for Efficient Analysis, Prioritization, Routing and Transparent Tracking." 2026.

[3] K. Patil, P. Prajwal Sawant, B. Nagorikar, and A. Mhatre. "RAG-Based Multi-Agent Framework for Intelligent Government Grievance Redressal Using Small Language Models (SSRN)." 2026.

[4] P. Gupta, O. Prasad, J. Tardar, A. Jadhav, and V. Barve. "AI-Based Solution to Enable Ease of Grievance Lodging and Tracking for Citizens Across Multiple Departments." 2025.

[5] R. Vasanthavelan, T. Ramaniarasan, K. Siva M., D. R. Ravindra, and K. Chander. "AI-Powered Petition Analysis and Grievance Management System." IJCRT, 2025.

[6] Vasanthavelan, R. et al. (2025). AI-Powered Petition Analysis and Grievance Management System.

IJCRT.

[7] Vijayalakshmi, P. et al. (2025). AI-Powered Grievance Management System for Public Administration. IJARESM.

[8] Ministry of Personnel, Public Grievances and Pensions, "Centralized Public Grievance Redress and Monitoring System (CPGRAMS) Annual Report," Government of India, 2024.

[9] United Nations Department of Economic and Social Affairs, "E-Government Survey: Digital Government in the Decade of Action," 2022.

[10] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in Proc. EMNLP-IJCNLP, 2019.