International Journal On Engineering Technology and Sciences – IJETS
ISSN (P):2349-3968, ISSN (O): 2349-3976 Volume XIII - Issue VII, July - 2026

AI-Powered Intelligent Complaint Resolution and Management Platform

[1]Dr. Mary Bearly E, [2]Ahamed Bilal S, [3]Hari Jothishwar V, [4]Micheal Dony J, [5]Ponnudurai G
[1,2,3,4,5]*Department of Artificial Intelligence and Data Science, Loyola Institute of Technology and Science, Thovalai, Kanyakumari Dist, Tamil Nadu, India – 629302*

Abstract

Public complaint redressal in municipal and institutional settings continues to rely on phone calls, physical registers, and disconnected complaint books, leaving citizens without confirmation, tracking visibility, or assurance that a grievance will be acted upon. Administrative staff must manually read every complaint, judge urgency, identify the correct department, and remember to escalate stalled cases — a process that is slow, inconsistent, and dependent on individual attentiveness rather than systemic guarantees. This paper presents the design, architecture, and evaluation of an AI-Powered Intelligent Complaint Resolution and Management Platform that automatically classifies, routes, tracks, and escalates complaints through a web application built around two connected n8n automation workflows. The intake workflow forwards each submitted complaint to an OpenAI-based classifier equipped with short-term memory, so related complaints from the same citizen are interpreted in context; it identifies the department (water supply, electricity, or sanitation), persists the complaint with a system-generated ticket number, and returns a confirmation within seconds. A second, independently scheduled workflow compares every open complaint against a service-level threshold every five minutes and automatically escalates and reassigns any complaint that has exceeded its allotted time. Both workflows are implemented declaratively in n8n rather than hand-coded in a backend framework, so operational changes require editing workflow nodes rather than redeploying application code. This paper situates the proposed system against methodologies reported between 2021 and 2025 and identifies a consistent gap: most prior systems treat automation as an implementation detail bolted onto a hand-coded backend, leaving workflow logic opaque and costly to modify. By elevating the workflow engine to a first-class architectural component, the proposed platform offers a transparent, modular, low-maintenance alternative.

Keywords — Complaint Management System; n8n Automation; Artificial Intelligence; Workflow Automation; E-Governance; Natural Language Processing; SLA Escalation; PostgreSQL; REST API; Citizen Grievance Redressal

1. Introduction

Local government responsiveness is frequently judged not by the policies a municipality writes but by how it handles the small, everyday failures of civic infrastructure — a streetlight that stays dark, a leaking water pipe, an overflowing garbage bin. The mechanism through which citizens report such failures, the complaint or grievance-redressal process, carries outsized importance relative to its administrative simplicity.

In most municipal and institutional settings, this mechanism has not kept pace with digital transformation elsewhere in public services. Complaints are still commonly registered through phone calls, complaint books, or a static web form that deposits a message into an inbox no one is specifically responsible for monitoring. Three structural weaknesses recur: first, there is no acknowledgement loop,

so a citizen has no way of confirming a complaint was received; second, there is no visibility into progress, so the citizen's only recourse is a follow-up call; third, and most consequentially, there is no systemic guarantee against neglect — escalation depends on whether a staff member happens to notice and follow up, a dependency on personal attentiveness rather than institutional process.

Two adjacent technical developments make it newly practical to address these weaknesses without large IT budgets. First, transformer-based language models, including large language models (LLMs) accessed through commercial APIs, can now classify short, informally written complaint text into operationally meaningful categories with accuracy that rivals purpose-built classifiers. Second, workflow-automation platforms such as n8n now allow developers to compose webhook triggers, conditional branching, scheduled polling, database operations, and AI API calls as a visual, editable graph rather than as backend source code. This paper explores the intersection of these two developments — using an automation engine to orchestrate an AI classifier as one node among several.

The system gives citizens a single web-based entry point to submit and later track a complaint, while an n8n-orchestrated backend interprets the complaint using an OpenAI-based classifier, assigns it to the correct department, stores it with a traceable ticket number, and — through an independent scheduled workflow — automatically escalates any complaint that breaches its resolution window. Unlike systems where automation is an internal implementation detail of a hand-coded backend, this platform treats the workflow definition itself as the primary artifact of the system, keeping classification logic, routing rules, and escalation policy visible and editable — a property that matters in a government-facing context where citizens and auditors may reasonably ask why a complaint was routed or prioritized the way it was.

2. Objectives

The project is built around one central objective — replacing discretionary, staff-dependent complaint handling with a systemic, auditable pipeline — decomposed into the following specific objectives:

- To automatically classify free-form citizen complaint text into the correct civic department, removing the need for a human reader to interpret intent for every submission.
- To route each classified complaint consistently to the department empowered to act on it, eliminating the institutional-knowledge dependency that causes manual misrouting.
- To guarantee that no complaint can be silently neglected, by monitoring every open complaint against a bounded service-level threshold on a fixed schedule rather than relying on supervisor initiative.
- To give citizens a ticket-based tracking mechanism so they can verify receipt and check status without repeated follow-up phone calls.
- To express all classification, routing, and escalation logic as visible, editable workflow nodes rather than opaque backend code, so operational changes do not require redeployment.
- To preserve final resolution authority with human department officers, scoping the AI layer strictly to classification, priority judgement, and message drafting.

3. Literature Survey

Research relevant to this platform spans four overlapping strands: AI/NLP-based complaint classification, citizen grievance-redressal and e-governance platforms, image- and sensor-driven civic issue detection, and workflow-automation engines.

3.1 AI- and NLP-Based Complaint Classification

Shaukat [2] applied supervised classifiers to institutional complaint datasets, showing that even lightweight text classifiers reduce manual sorting effort, though the work was evaluated purely as a classification task without routing or escalation. A 2025 public-sector informatics study [1] used BERT-based embeddings with semantic clustering to detect latent issue patterns proactively, and noted that GPT-based LLMs fine-tuned on municipal data outperform classifiers such as Random Forest on multi-label triage, at higher computational cost — a trade-off directly relevant to this paper's choice of a hosted LLM API. A 2024 healthcare study [5] applied similar ML/NLP techniques to over 1,400 hospital complaint records, reinforcing that classification approaches transfer across institutional domains.

3.2 Grievance Redressal and E-Governance Platforms

Singh et al. [7] built an early NLP-based grievance system forwarding complaints to the appropriate authority. More recent systems add multimodal intake: an AI-Driven Citizen Grievance Redressal System [6] combines NLP, sentiment analysis, and voice-to-text so citizens can file complaints by writing, speaking, or chatbot. Kotwal et al. [8] report a 40% reduction in resolution time and 87% satisfaction from a three-tier web portal automating categorization and monitoring. The Autonomous Grievance Redressal platform [10] adds GPS-tagged evidence and Firebase real-time sync, while Nagarik Connect [12] uses a BERT-based model to determine the correct municipal department. A 2024 survey [13] observes that although NLP-based classification and routing are now common, most deployed platforms remain reactive, intervening only after a service-level violation has already occurred.

3.3 Image-Driven Civic Issue Detection

A parallel strand addresses complaints more naturally expressed visually than textually, such as potholes, garbage accumulation, or damaged streetlights. CitySolution [3] uses a deep-learning image classifier across paired citizen- and authority-facing Android applications to categorize photographed civic issues. CivicFix [11] similarly trains a convolutional neural network to classify photographs of civic issues, reporting over 92% classification accuracy, and demonstrates that integrating geolocation APIs substantially reduces location-reporting errors relative to citizen-entered addresses. An earlier IJARCCE study [4] proposed a comparable image-processing pipeline in 2021, an early example of the classification-plus-routing pattern this paper builds on.

3.4 Prioritization, Escalation, and Workflow Orchestration

Fewer studies address prioritization and escalation with the same depth as classification. A 2026 reinforcement-learning framework for New York City Department of Buildings complaint intake [19] frames complaint triage as a sequential decision process, training agents to route complaints into escalate/batch/defer/inspect-now categories while explicitly optimizing for equitable classification coverage across neighborhoods, and reports that complaint recurrence is a stronger predictor of actionable violations than raw complaint volume. A zero-shot LLM framework for multimodal

grievance processing [20] treats urgency scoring as a distinct pipeline stage rather than folding it into category classification. On the automation-substrate side, a comparative study of open-source workflow orchestration engines presented at IEEE BigData 2022 [23] benchmarks tools such as Apache Airflow and Node-RED, establishing that visual, node-based orchestration can match hand-coded pipelines in reliability for well-scoped tasks. A 2025 study evaluating n8n specifically for enterprise integration and AI orchestration [21] reports linear scalability and greater than 98% reliability in production-style benchmarks, including case studies where n8n orchestrates calls to AI services such as OpenAI directly as workflow nodes — the same architectural pattern used in this paper's intake pipeline.

3.5 Research Gap

Across the reviewed literature, systems that classify complaints accurately rarely specify how classification output should drive routing, ticketing, and escalation as a unified pipeline. Systems that do implement a complete pipeline typically hard-code it into a conventional backend framework, making the routing and escalation policy opaque to non-developers and costly to modify. None of the surveyed complaint-management papers use a workflow-automation engine as the primary backend for a citizen-facing grievance system. Table 1 consolidates the comparison.

System	Method / Technology	Advantages	Limitations	Research Gap
Shaukat [2] (2022)	Supervised text classification (classical ML/NLP)	Reduces manual sorting effort	Classification only; no routing, ticketing, or escalation	No operational pipeline around classification output
Public-sector informatics study [1] (2025)	BERT embeddings + semantic clustering (CRISP-DM)	Enables proactive, pre-breach issue detection	High compute cost of GPT-based variants for small deployments	No integrated escalation / SLA mechanism
Singh et al. [7] (2022)	NLP-based grievance interpretation	Early, widely cited automated-routing baseline	Limited prioritization or escalation capability	No SLA-driven escalation
Kotwal et al. [8] (2025)	Three-tier web architecture with automated categorization	40% faster resolution; 87% satisfaction reported	Hand-coded backend; changes require redeployment	No visual / editable automation layer
CivicFix [11] (2025)	CNN-based image classification + geolocation routing	>92% classification accuracy; reduces location error	No text-based AI classification or escalation logic	Escalation and NLP intake both absent
Patient complaint study [5] (2024)	Supervised ML on dissatisfaction-term features; externally validated	Validated across institutions; domain-transferable method	Healthcare-specific; no citizen-facing ticketing or tracking	Lacks citizen-facing status tracking
AI-Driven Citizen Grievance Redressal [6] (2026)	NLP + sentiment analysis + voice-to-text + chatbot	Multi-channel intake (text, voice, chat); automated forwarding	Escalation logic not detailed; backend implementation opaque	Automation treated as implementation detail, not auditable policy
Autonomous Grievance	GPS-tagged multimedia intake +	Real-time sync; map-based clustering;	Rule-based prioritization, not	No time-bound automatic escalation

System	Method / Technology	Advantages	Limitations	Research Gap
Redressal (E-Gov) [10] (2025)	Firebase real-time sync	citizen feedback loop	learned; no SLA-timed escalation	mechanism
Nagarik Connect (ICGS) [12] (2025)	BERT-based department classification + GPS intake	Unified lifecycle from submission to resolution tracking	Escalation and SLA monitoring not central to design	Escalation treated as a secondary feature
CitySolution [3] (2024)	CNN image classification (Teachable Machine), Android	Simple citizen photo capture; authority-side dashboard	Text complaints unsupported; no explicit SLA logic	No unified text + image + escalation pipeline
n8n orchestration study [21] (2025)	n8n + OpenAI/Hugging Face integration nodes	>98% reliability; linear scalability; native AI-node support	Enterprise-focused case studies (SAP, Salesforce), not civic complaint systems	No application to citizen grievance redressal domain

As Table 1 shows, no reviewed system combines LLM-based text classification with conversational memory, automatic category-based routing, ticket-numbered persistence, and an independently scheduled SLA-driven escalation workflow, all within a single visual automation engine rather than a hand-coded backend.

4. Existing System

Existing complaint-handling mechanisms in most municipal and institutional settings fall into two broad categories, both of which the proposed platform is designed to replace. The problem decomposes into four recurring sub-problems that any adequate replacement must solve concurrently: classification ambiguity, since complaints arrive as free-form text ("no water since yesterday") that a human reader must otherwise interpret and map to a department; routing inconsistency, since correctly directing a complaint to the department empowered to act on it requires institutional knowledge that varies between staff members; escalation blindness, since there is typically no systemic mechanism tracking how long a complaint has remained open relative to a reasonable resolution time; and lack of citizen-facing transparency, since without ticket-based tracking, citizens cannot verify receipt or check status.

4.1 Manual and Semi-Digital Complaint Handling

The dominant existing approach relies on phone calls logged in a diary, physical complaint books kept at a ward or department office, or a static web form that deposits messages into an unmonitored inbox. In these systems, there is no automated acknowledgement, no visibility into progress, and no systemic guarantee that a stalled complaint will be escalated — escalation depends entirely on whether a staff member notices, remembers, and has the authority to act.

4.2 Prior Automated and Research Systems

Where automation has been attempted, it takes two forms. Classification-focused systems [1], [2], [5] apply ML/NLP to sort complaints accurately but stop short of specifying the surrounding operational

pipeline — how a classified complaint is stored, tracked, escalated, or communicated back to the citizen. Full-pipeline systems [6], [8], [10], [12] implement intake, routing, and notification, but hard-code this logic into a conventional backend framework such as Express, Django, or a three-tier web architecture; every operational change (adding a department, adjusting an escalation window, adding a notification channel) therefore requires a developer to modify source code, test it, and redeploy the service. Image-driven systems [3], [11] extend intake to photographic evidence but generally lack text-based NLP classification and escalation logic entirely. Even where prioritization is addressed algorithmically [19], [20], it is evaluated as a research contribution without corresponding attention to how the logic would be deployed, maintained, or audited by a small technical team. In every case, escalation — where it exists at all — is treated as an implementation detail rather than a first-class, guaranteed system capability, leaving the same core weakness: a complaint can stall indefinitely without triggering any automatic response.

4.3 Reactive Orientation of Deployed Systems

A 2024 technological survey of citizen complaint management systems [13] synthesizes this body of work and observes that although NLP-based classification and routing have become common, most deployed platforms remain fundamentally reactive — they intervene only after a service-level violation or explicit citizen dissatisfaction has already occurred, with limited predictive monitoring or escalation logic built in as a first-class system capability rather than an afterthought. This reactive orientation is the single most consistent shortfall across the existing systems reviewed, and it is the specific gap the proposed platform is designed to close through an independently scheduled, systemic escalation workflow rather than a manually triggered or supervisor-dependent one.

5. Proposed System

The proposed AI-Powered Intelligent Complaint Resolution and Management Platform is designed around a single guiding principle: the automation engine, not a hand-coded backend, is the primary carrier of operational logic. The system consists of a lightweight citizen-facing web application, a workflow-automation layer implemented in n8n that performs all AI classification, routing, persistence, and notification logic, and a PostgreSQL database serving as the system of record.

At a conceptual level the system provides two citizen-facing capabilities and two administrative capabilities. Complaint submission lets a citizen describe an issue in natural language and immediately receive a ticket number, without selecting a department, category, or priority themselves, since these are inferred automatically. Complaint tracking lets a citizen look up a complaint by ticket number and see its current status (submitted, assigned, in progress, escalated, resolved). Complaint intake and routing shows department officers only the complaints relevant to their department, pre-classified and pre-prioritized. Escalation monitoring automatically alerts senior officers when a complaint has exceeded its expected resolution time, without requiring periodic manual audits.

Two properties distinguish this design from prior systems. First, classification is context-aware: the AI node retains short-term memory of a citizen's recent submissions, so a follow-up message such as "still not fixed" is interpreted in relation to the prior complaint rather than as an isolated fragment. Second, escalation is systemic rather than discretionary: because the escalation check runs on a fixed five-minute

schedule against every open complaint, no complaint can be silently forgotten for longer than one scheduler cycle beyond its SLA threshold.

Architecturally, the system comprises five layers. The presentation layer (HTML5/CSS3/vanilla JavaScript) exposes the complaint submission form and tracking views and contains no business logic. The API/webhook layer consists of n8n Webhook Trigger nodes that accept requests and hand them to the appropriate workflow. The automation layer is the architectural core, consisting of two independently triggered workflows: an Intake & Classification Workflow, triggered per submission, which uses a memory node, an OpenAI chat model, a Switch/Router node, and a Postgres insert to classify, route, persist, and confirm each complaint within seconds; and an SLA Escalation Workflow, triggered every five minutes, which queries all open complaints, compares elapsed time against department- and priority-specific thresholds, and automatically escalates, reassigns, and notifies for any complaint that has breached its threshold. Both workflows read from and write to the same PostgreSQL persistence layer — comprising complaints, citizens, departments, status_history, and escalations tables — allowing them to operate independently while sharing a single source of truth. The notification layer currently delivers confirmations and alerts through the web application itself, a deliberate initial scope kept simple before expanding to email or SMS.

The system currently supports three departments — water supply, electricity, and sanitation — chosen as the highest-volume, most time-sensitive categories of civic complaint, with the architecture designed so that additional departments, notification channels, and escalation policies can be added by extending workflow nodes rather than by re-architecting the system. Final resolution authority remains with human department officers; the AI layer is deliberately scoped to classification, priority judgement, and message drafting only.

5.1 Workflow Summary

The Intake & Classification Workflow executes on every submission: a Webhook Trigger receives the payload; a memory node retrieves recent context for that citizen; an OpenAI chat model returns a structured department, priority, and drafted acknowledgement; a Switch node routes by department; a Postgres node inserts the complaint and generates a human-readable ticket number (e.g., WS-000241); and a Respond-to-Webhook node returns the confirmation to the citizen, typically within a few seconds.

The SLA Escalation Workflow runs independently on a fixed five-minute schedule: it queries all complaints not yet resolved or closed, computes elapsed time since each complaint's last status change, and compares it against the department- and priority-specific SLA threshold. Complaints within their allotted time are skipped; complaints that have breached the threshold are marked escalated, reassigned to the department's senior officer, logged, and the relevant parties notified. Because the two workflows share no direct trigger dependency but read and write the same complaints table, any complaint inserted by the intake workflow becomes visible to the very next escalation cycle, bounding detection delay to at most five minutes past any SLA breach — a concrete, quantifiable guarantee that a manually monitored process cannot offer.

REFERENCES

[1] "Proactive Complaint Management in Public Sector Informatics Using AI: A Semantic PatternRecognitionFramework,"*AppliedSciences*,vol.15,no.12,art.no.6673,2025.doi: 10.3390/app15126673.

[2] A. Shaukat, "NLP Based Model for Classification of Complaints: Autonomous and Intelligent System,"in *Proc.IEEEInt.Conf.onAdvancesinComputing*,2022,pp.234–239.

[3] F. Shama, A. Aziz, and L. B. M. Deya, "City Solution: A complaining task distributive mobile application for smartcity corporation using deep learning," *Array*(Science Direct),2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352711024002000

[4] "Public Complaint Sorting for Smart City using Image Processing," *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE)*, vol. 10, no. 5, May 2021.

[5] "An Intelligent System for Classifying Patient Complaints Using Machine Learning and NaturalLanguageProcessing:DevelopmentandValidationStudy,"2024.[Online].Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11754990/

[6] "AI-DrivenCitizenGrievanceRedressalSystem,"*InternationalJournalofEngineering Research & Technology (IJERT)*, vol. 15, no. 7, 2026. [Online]. Available: https://www.ijert.org/ai-driven-citizen-grievance-redressal-system-ijertv15is070169

[7] A.P.Singh,A.Goel,A.Goel,andD.Arya, "NLP based Grievance Redressal System,"

International Journal of Computer Applications (IJCA),vol.184,no.12,pp.44–48, 2022.

[8] R.Kotwal,J. Naik,K. Mane, A.Shinde, and K.Palande," Grievance Redressal System:A Web-Based Public Grievance Portal," *ITM Web of Conferences*, vol. 81, 2025.

[9] S. Iyer, D. Gandhi, D. Mishra, R. Trivedi, and T. Ansari, "Solve My Problem – Grievance RedressalSystem,"in*Proc.SpringerConf.Volume*,2024.doi:10.1007/978-3-031-50993-3_11.

[10] "AutonomousGrievanceRedressalSystemforE-GovernanceUsingAIandAutomation," 2025. [Online]. Available: https://www.researchgate.net/publication/397517090

[11] "CivicFix:SmartComplaintRoutingforUrbanSolutions,"2025.[Online].Available: https://www.researchgate.net/publication/398400172

[12] "NagarikConnect–IntegratedCitizenGrievanceSystem,"*InternationalJournalfor Multidisciplinary Research (IJFMR)*, 2025. [Online]. Available: https://www.ijfmr.com/papers/2025/6/64107.pdf

[13] "A Technological Survey on Citizen Complaint Management Systems and Future Advances,"2024.[Online].Available:https://www.researchgate.net/publication/381716725

[14] "AIBasedSmartCivicComplaintManagementSystem," *InternationalJournalof Engineering & Extended Technologies Research (IJEETR)*, 2024.

[15] "SmartComplaintManagementSystemwithAI-EnabledPrioritization," *International Journal of Creative Research Thoughts (IJCRT)*, 2025. [Online]. Available: https://ijcrt.org/download.php?file=IJCRT2510624.pdf

[16] M. Musthafa and A. Muthukumaran, "Task Management System with AI-based Prioritization,"*InternationalJournalofEngineeringResearch&Technology(IJERT)*,vol.13, no. 4, art. 40324, 2024.

[17] B.Gelman,S.Taoufiq,T.Vörös,andK.Berlin,"ThatEscalatedQuickly:AnML Framework for Alert Prioritization," *arXiv preprint* arXiv:2302.06648, 2023.

[18] K. T. Devi, R. Aditya, P. Yeswanth Raj, Y. Thirupathi Rao, T. HemaVenkata Sai, and A. Lakshman,"SmartComplaintPortal,"*InternationalJournalofInnovativeScienceandResearch Technology (IJISRT)*, 2023.

[19] "ScalingtheQueue:ReinforcementLearningforEquitableCallClassificationCapacityin NYC Municipal Complaint Systems," *arXiv preprint* arXiv:2605.06482, 2026.

[20] "Azero-shotLLMframeworkformultimodalgrievanceclassification,urgencyscoring,and abuse detection in civic feedback systems," *Scientific Reports*, 2025.

[21] "n8n: An Open-Source Workflow Automation Platform for Enterprise Integration and AI-DrivenOrchestration,"*InternationalJournalofComputerApplications(IJCA)*,vol.187,no.63, 2025.

[22] J.Chenetal.,"LangChain:Buildingcontext-awarereasoningpipelines,"*arXivpreprint*

arXiv:2307.01952, 2023.

[23] M.Arendtetal.,"Comparativestudyofopen-sourceworkfloworchestrationengines,"in *Proc. IEEE Int. Conf. on Big Data (BigData)*, Osaka, Japan, 2022, pp. 4745–4754. doi: 10.1109/BigData55660.2022.10021136.

[24] OWASPFoundation,"OWASPAPISecurityTop10,"2023.[Online].Available: https://owasp.org/API-Security/

[25] M.Jones,J.Bradley,andN.Sakimura,"JSONWebToken(JWT),"IETFRFC7519,May 2015.

[26] R.FieldingandJ.Reschke,Eds.,"HypertextTransferProtocol(HTTP/1.1):Semanticsand Content," IETF RFC 7231, June 2014.