International Journal of Future Engineering Innovations

National Citizen Intelligence Platform for Predictive Public Service Optimization and Grievance Resolution

Sankalp Pradip Thool [1*], Dipak Gahininath Mhaske [2], Atharv Jagannath Vethekar [3], Ruturaj Rajaram Konde [4], Dr. Khatal SS [5]
[1-5] Department of Computer Engineering, Sharadchandra Pawar College of Engineering, Savitribai Phule Pune University, Pune, Maharashtra 412409, India

* Corresponding Author: **Sankalp Pradip Thool**

Article Info

ISSN (online): 3049-1215
Volume: 02
Issue: 06
Received: 14-09-2026
Accepted: 15-10-2026
Published: 13-11-2026
Page No: 100-104

Abstract
Municipal offices deal with complaints about services people use every day, from roads and water supply to sanitation, transport, drainage, and electricity. When many reports arrive at once, reading and sorting each one by hand can delay the response. This paper reviews ways that Artificial Intelligence (AI), Machine Learning (ML), Natural Language Processing (NLP), sentiment analysis, and Large Language Models (LLMs) could assist staff with this work. It looks at TF-IDF, Logistic Regression, Random Forest, cosine similarity, topic modelling, and transformer-based models. With suitable data and careful testing, these methods could help categorize complaints, estimate urgency and resolution time, identify reports that may be related, and suggest the department to contact. The review also considers practical issues such as incomplete records, language variation, privacy, and prediction errors. The aim is to support public officials with useful information while keeping final decisions under human supervision.

Keywords: Artificial Intelligence, Machine Learning, Natural Language Processing, Citizen Grievance Management, Sentiment Analysis, Large Language Models, Predictive Analytics, Smart Governance

1. Introduction

As towns and cities grow, residents rely on public services for more of their everyday needs. They may report damaged roads, water shortages, poor sanitation, uncollected garbage, drainage problems, power cuts, transport issues, or faults in other public facilities. These reports offer authorities a direct view of local problems and residents' expectations. The difficulty begins when many complaints arrive close together. Staff have to understand each report, judge its urgency, and send it to the department that can act on it. If all of these steps are completed manually, some cases may wait longer than they should. AI and Machine Learning could help staff review a larger number of reports and notice patterns across them [1, 2].

Natural Language Processing (NLP) offers tools for working with the text residents submit. Once unnecessary elements are cleaned from a message, Term Frequency–Inverse Document Frequency (TF-IDF) can represent words that help distinguish one complaint from another. Cosine similarity can also compare a new report with earlier submissions and flag text that looks alike. This may save staff from checking the same material repeatedly, but similar wording does not prove that two reports describe one event. Possible matches should be checked in context before they are treated as duplicates [3, 4].

A category on its own doesn't tell officials how residents feel about a service. Sentiment analysis may identify if feedback is good, negative or neutral, while topic-modelling techniques can cluster comments that talk about similar things. Latent Dirichlet Allocation (LDA), Non-negative Matrix Factorization (NMF) and BERTopic find topics in distinct methods. "Their findings could point to services that are subject to repeated complaints or issues that require further investigation. Transformer models, including BERT and other LLMs, can use surrounding words to extract details from longer messages. Their output still needs review, particularly if it may influence a public decision [5-8].

For these methods to be useful to local authorities, they should be connected through a practical workflow. The workflow could classify complaints, estimate urgency, examine sentiment, check for duplicates, estimate resolution time, and recommend a department. Looking at these results together may give staff a fuller picture of a case. The quality of that picture will depend on the complaint records and on regular checks of model performance. The platform must also account for regional languages, privacy, explainability, and prediction errors. These safeguards are important because a wrong result can affect a resident or influence how public resources are distributed [9, 10].

This paper reviews selected AI- and NLP-based methods for managing citizen grievances and planning public services. It considers where the methods may help, where they have limitations, and which questions need further study. The review then presents a conceptual approach that connects complaint analysis with administrative decision support. It is intended as a starting point for a platform that helps officials assess reports consistently while retaining human judgement.

2. Problem Statement

Local authorities receive a steady stream of complaints about roads, water supply, sanitation, transport, drainage, electricity, and other public services. Online systems often let residents submit a report and follow its status, but staff may still have to assess the details manually. As the number of submissions rises, urgent cases can be harder to spot, reports about the same issue may go unnoticed, and complaints may be sent to the wrong department. This paper considers how AI and NLP could assist with these tasks in one shared platform, while leaving important decisions open to staff review.

3. Objectives

- To review Artificial Intelligence and Natural Language Processing techniques for efficient citizen grievance management.
- To study Machine Learning approaches for automatic complaint classification and priority prediction.
- To examine sentiment analysis, topic modelling, and similarity-based techniques for identifying citizen concerns and duplicate complaints.
- To analyse predictive approaches for complaint resolution-time estimation and appropriate department recommendation.
- To identify existing research gaps and explore an integrated AI-based framework for predictive public-service optimization and grievance resolution.

4. Literature Survey

4.1. Sarzaeim *et al*. (2024) — a Framework for LLM-Assisted Smart Policing System

Sarzaeim *et al*. (2024) investigated if Large Language Models (LLMs) like BART, GPT-3 and GPT-4 could be exploited in text-based smart policing tasks. It examined zero-shot learning, few-shot learning and fine-tuning. The choices range in the amount and nature of task-specific examples used to inform a model. This work demonstrates that language models may be valuable for understanding complex content in public-sector contexts. The optimal way will rely on the task and on the data available and reliability required.

4.2. Kumi *et al*. (2024) — Uncovering Concerns of Citizens Through Machine Learning and Social Network Sentiment Analysis

Kumi *et al*. (2024) examined 114,390 Reddit comments about Saskatoon Canada to see what topics people mentioned NLP, topic modeling, clustering and sentiment analysis were used in their method. The algorithms used were Latent Dirichlet Allocation (LDA), Non-negative Matrix Factorization (NMF), BERTopic and K-means clustering. These techniques helped to arrange a big amount of comments, and to show reoccurring themes. One major drawback is that social media posters may not be representative of the community at large, therefore the results do not provide a thorough picture of popular opinion.

4.3. Zheng *et al*. (2023) — Generative Job Recommendations with Large Language Model

Zheng *et al*. (2023explored the capability of Large Language Models to provide job recommendations based on textual information and user information. Their work demonstrates how context can be exploited to give more relevant ideas for the individual. The study is about jobs, not public concerns but it is an example of suggestion backed by AI. A similar idea might be entertained in stating which department would be responsible for a complaint, or what action might be appropriate.

4.4. Yigitcanlar *et al*. — Artificial Intelligence and Urban Planning

Yigitcanlar *et al*.'s work explores the use of AI in urban planning and smart-city development . AI technologies can evaluate big data, find trends and provide knowledge that planners may use to make decisions. The feedback from citizens also helps to highlight concerns about infrastructure, transport, environmental conditions and other urban services. The value of such analysis depends on the reliability of the underlying information and its reflection of the individuals affected by the decisions.s.

4.5. Adikari and Alahakoon — Understanding Citizens' Perceptions of Smart City Development Through Social Media Analysis

Adikari and Alahakoon investigated social media information to understand the public perception on smart city development. Posts from residents may show expectations, pleasure, worries and reactions to urban projects. Authorities might analyze these messages methodically to discover common themes and assess the public's response to changes in local services. Many residents do not utilize social media therefore these posts should be viewed as one source of evidence, rather than a comprehensive reflection of public opinion.

4.6. Tran *et al*. — Natural Language Processing for Understanding Urban Transit Experiences

Tran *et al*. The analysis can bring together comments about delays, accessibility, service quality, and other passenger concerns. Grouping related reports may help transport authorities review large amounts of feedback more efficiently. This work also illustrates how text-analysis methods can be adapted to a particular public-service area.

Table 1: Comparison Table

Author & Year	Method Used	Advantages	Limitations
Sarzaeim *et al.* (2024)	LLMs	Classifies text	Focuses on policing
Kumi *et al.* (2024)	NLP, sentiment analysis	Finds citizen concerns	Uses social media
Zheng *et al.* (2023)	LLMs	Gives recommendations	Focuses on jobs
Yigitcanlar *et al.*	AI	Supports urban planning	Limited grievance handling
Adikari and Alahakoon	Social-media analysis	Understands citizen opinions	Uses social-media data
Tran *et al.*	NLP	Identifies transport issues	Focuses on transportation

5. Working of System

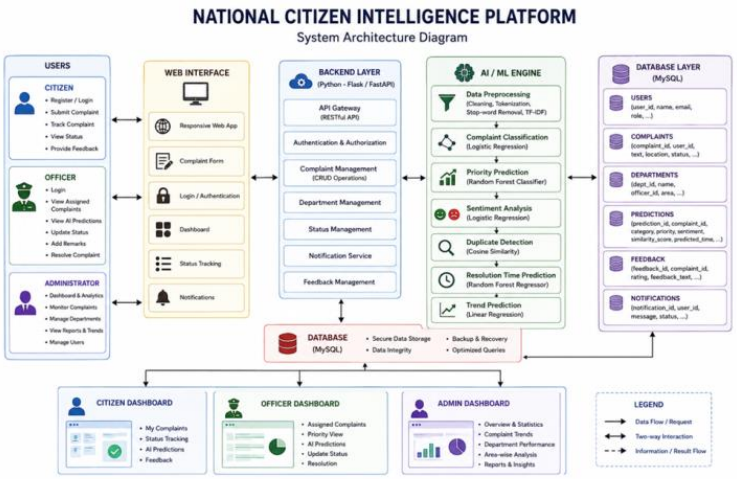


Fig 1: Design of the system

6. System Design

6.1. User Management Module

The platform offers three user roles: - Citizen - Officer - Admin Citizens may set up an account, file a complaint, track its progress and offer feedback. Officers review their assigned cases and change the status as work proceeds. Admins control user access and monitor activities on the platform. Separating these responsibilities allows each person to do the functions needed for their work.

6.2. Web Interface Module

The main features of the platform can be accessed through the online interface which includes complaint forms, login and authentication, dashboards, status tracking and notifications. The websites need to be easy to browse so residents may report an issue and staff can get to the information they need without extra steps.

6.3. Backend Layer

FastAPI links the user interface with the AI/ML components and the database. It receives complaint submissions, supports department assignment, records status changes, and manages notifications and feedback. In this way, the backend coordinates information exchanged between the platform's modules.

6.4. AI/ML Engine

In analyzing complaint data, the AI/ML engine applies a number of artificial intelligence and machine learning algorithms. Main tasks are: Data preparation: It cleans or normalizes the text components before analysis. Complaint Classification: Assigns each complaint to the right service type. Priority Prediction – Predicting the priority of a complaint; Sentiment Analysis - It identifies if the sentence carries a good, negative or neutral attitude. Duplicate Detection - Identifies reports that have the same information as current complaints. Resolution Time Prediction: Predicts the time to resolve a request. Trend Prediction Discovering recurring Service issues throughout time.

6.5. Database Layer

MySQL stores the platform's records, including user details, complaint information, department data, model outputs, feedback, and notification records. Proper record keeping helps in easy retrieval of information and report writing.

6. Dashboard Module

The platform includes a different dashboard for each user category • Citizen Dashboard: Shows complaints submitted, status, predictions and feedback choices. • Officer Dashboard – See assigned complaints, priority forecast, status changes and resolution data. • Admin Dashboard: Complaint volume, reoccurring issues, department activities, reports.

7. Complaint Processing and Decision Support

When a resident submits a complaint, the backend saves the report and passes its text to the AI/ML engine. The engine estimates the category and priority, checks for similar reports, and predicts a possible resolution time. Staff can use these results when choosing a department and deciding which cases may need attention first. Officers should review the predictions before making a final decision.

8. Notification and Feedback Module

When a complaint is made, when a complaint's status changes, and when a resolution is logged, citizens are notified. They can give feedback on the service when the matter is closed.

7. Expected Outcome

The planned National Citizen Intelligence Platform would help authorities categorize complaints by topic and probable urgency. This could help identify urgent reports and direct them to the department responsible for the issue. How useful it will be in practice will rely on the accuracy of the model and the extent to which the platform can be integrated into existing administrative processes.

The technology will also be able to detect similar reports, analyze the tone of the citizen input and forecast how long a case could take to resolve. Such elements could help officials see repeated service problems and give them a sense of what's on residents' minds. The forecasts are estimations not guaranties and should be taken in conjunction with staff judgment.

Separate screens for people, officers and administrators would make complaint status, agency activities and reoccurring issues easy to examine. Sharing this information could improve collaboration and assist residents know the status of their complaints. The key anticipated results are summarized here.

Table 2: The projected key effects of the proposed platform are given below.

Sr. No.	Expected Outcome
1	Automatic classification of citizen complaints.
2	Prediction of complaint priority levels.
3	Detection of duplicate and similar complaints.
4	Sentiment analysis of citizen feedback.
5	Prediction of complaint resolution time.
6	Identification of recurring public-service issues.
7	Efficient department assignment and complaint tracking.
8	Interactive dashboards for monitoring and decision-making.

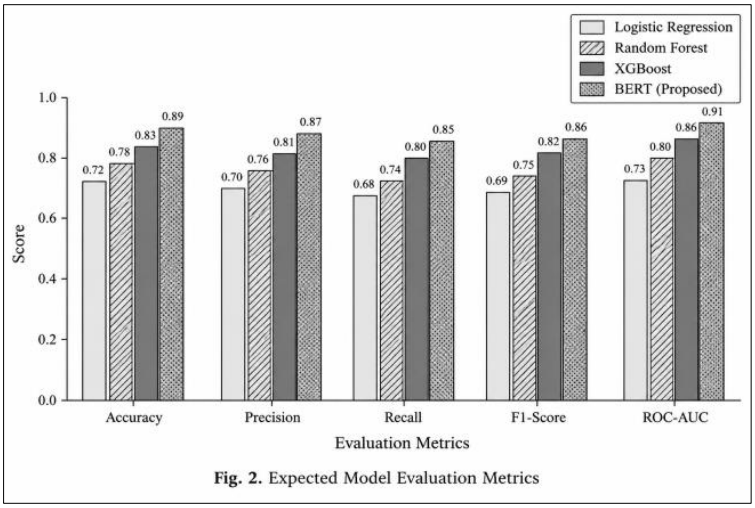


Fig. 2. Expected Model Evaluation Metrics

Fig 2: Expected Model Evaluation Metrics

Table 3: Expected Model Evaluation Metrics

Sr. No.	Evaluation Metric	Logistic Regression	Random Forest	XGBoost	BERT (Proposed)
1	Accuracy	0.72	0.78	0.83	0.89
2	Precision	0.70	0.76	0.81	0.87
3	Recall	0.68	0.74	0.80	0.85
4	F1-Score	0.69	0.75	0.82	0.86
5	ROC-AUC	0.73	0.80	0.86	0.91

8. Conclusion

The National Citizen Intelligence Platform leverages AI, Machine Learning and NLP for public complaints processing. It is offered for functions such as complaint classification, priority prediction, duplication detection, sentiment analysis and resolution-time estimation. Dashboards for residents, officers and administrators would make it easy to track case progress and broader patterns. These services can assist departments in structuring their work and utilizing complaint information for decision processes. But the design itself doesn't ensure the platform will perform reliably. It should be tested on actual complaint records and evaluated for its accuracy, fairness, privacy and use in day-to-day management.

9. Future Scope

The platform may be scaled to take complaints in regional languages with proper multi-lingual NLP models. Further work with deep learning techniques and LLMs could be used to improve categorization, priority predictions and draft replies. By linking to government websites, it may be able to coordinate the registration of complaints and the status of complaints across departments. Other possible features include explainable AI, location-based analysis and tools to identify emerging service difficulties. The system should be tested on a range of real world datasets in an effort to evaluate performance, uncover vulnerabilities and inform changes prior to the wider deployment of the system.

References

1. Sarzaeim P, Mahmoud QH, Azim A. A Framework for LLM-Assisted Smart Policing System. IEEE Access. 2024;12:74915-74929.
2. https://doi.org/10.1109/ACCESS.2024.3404862
3. Kumi S, Snow C, Lomotey RK, Deters R. Uncovering Concerns of Citizens Through Machine Learning and Social Network Sentiment Analysis. IEEE Access. 2024;12:94885-94913.
4. https://doi.org/10.1109/ACCESS.2024.3426329
5. Kim N, Hong S. Automatic classification of citizen requests for transportation using deep learning: Case study from Boston city. Information Processing & Management. 2021;58(1):102410. doi:10.1016/j.ipm.2020.102410
6. Chen S, Zhang Y, Song B, Du X, Guizani M. An Intelligent Government Complaint Prediction Approach. Big Data Research. 2022;30:100336.
7. https://doi.org/10.1016/j.bdr.2022.100336
8. Verma S. Sentiment analysis of public services for smart society: Literature review and future research directions. Government Information Quarterly. 2022;39(3):101708. doi:10.1016/j.giq.2022.101708
9. Lidiawaty BR, Suryani E, Vinarti RA. Traffic Urgency Model: A Novel Approach to Prioritize Complaint Texts Using Enhanced Named Entity Recognition for Bahasa Indonesia Cases. IEEE Access. 2025;13. doi:10.1109/ACCESS.2025.3556423
10. Blei DM, Ng AY, Jordan MI. Latent Dirichlet Allocation. Journal of Machine Learning Research. 2003;3:993-1022.
11. Grootendorst M. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. arXiv preprint. 2022. doi:10.48550/arXiv.2203.05794
12. Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of NAACL-HLT. 2019:4171-4186.
13. Liu B. Sentiment Analysis and Opinion Mining. Morgan & Claypool Publishers; 2012.
14. Manning CD, Raghavan P, Schütze H. Introduction to Information Retrieval. Cambridge University Press; 2008.
15. Pedregosa F, *et al*. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research. 2011;12:2825-2830.
16. Breiman L. Random Forests. Machine Learning. 2001;45:5-32.
17. Cortes C, Vapnik V. Support-vector networks. Machine Learning. 1995;20:273-297.
18. Vaswani A, *et al*. Attention Is All You Need. Advances in Neural Information Processing Systems. 2017;30.
19. Joachims T. Text categorization with Support Vector Machines: Learning with many relevant features. Proceedings of ECML. 1998:137-142.
20. Salton G, Buckley C. Term-weighting approaches in automatic text retrieval. Information Processing & Management. 1988;24(5):513-523.
21. Hutto CJ, Gilbert E. VADER: A parsimonious rule-based model for sentiment analysis of social media text. Proceedings of ICWSM. 2014;8(1):216-225.
22. Sokolova M, Lapalme G. A systematic analysis of performance measures for classification tasks. Information Processing & Management. 2009;45(4):427-437.
23. Fawcett T. An introduction to ROC analysis. Pattern Recognition Letters. 2006;27(8):861-874.

How to Cite This Article

Thool SP, Mhaske DG, Vethekar AJ, Konde RR, Khatal SS. National citizen intelligence platform for predictive public service optimization and grievance resolution. International Journal of Future Engineering Innovations. 2026;2(6):100–104.

Creative Commons (CC) License

This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms