

CivicLens AI

Citizen Complaint to Policy Insight Engine

¹Dr. Vijayant Verma, ²Rana Padwar, ³Apurwa Chandrakar, ⁴Khushi Jaiswal, ⁵Palak Mishra

¹Assistant Professor, ^{2,3,4,5}Student

¹Computer Science & Engineering,

¹Bhilai Institute of Technology, Raipur, India

Abstract : Urbanization has led to a significant increase in the volume of citizen complaints directed toward municipal authorities. These complaints typically include infrastructural and service-related issues such as road damage, water leakage, electrical failures, waste management problems, and drainage inefficiencies. Existing complaint management systems primarily function as repositories that store and track complaints but fail to perform intelligent aggregation, pattern detection, or advanced analytics. As a result, repeated issues persist, and authorities lack actionable insights for effective governance and policy formulation.

To address these challenges, this paper proposes CivicLens AI – Citizen Complaint to Policy Insight Engine, an intelligent civic analytics platform designed to transform unstructured complaint data into structured governance insights. The proposed system integrates Natural Language Processing (NLP), sentiment analysis, clustering algorithms, image classification, and geospatial visualization to analyze multimodal complaint data. Furthermore, Generative AI is employed to generate policy recommendations and administrative strategies. The system is implemented using the MERN stack, with Mapbox for geospatial heatmap visualization. Experimental results demonstrate that the proposed system improves complaint resolution efficiency, identifies recurring urban issues, and supports data-driven policy-making, thereby enhancing overall governance effectiveness.

Index Terms—*Smart City, Natural Language Processing, Smart Governance, Complaint Analytics, Generative AI, Clustering Sentiment Analysis, Heatmap Visualization, MERN Stack.*

I. INTRODUCTION

Urbanization and rapid population growth have placed immense pressure on city infrastructure and governance systems. Municipal authorities are required to manage a vast number of citizen complaints daily, covering diverse issues such as road maintenance, waste management, water supply disruptions, drainage problems, and electrical failures.

Although many cities have adopted digital complaint management systems, these platforms are largely limited to recording and tracking complaints. They lack advanced analytical capabilities required to identify recurring issues, assess severity, and generate actionable insights. Consequently, authorities often respond reactively rather than proactively, leading to inefficiencies and repeated complaints.

Artificial Intelligence (AI) and data analytics provide an opportunity to transform traditional complaint management systems into intelligent decision-support platforms. By leveraging NLP, machine learning, and data visualization techniques, it is possible to extract meaningful information from large volumes of unstructured complaint data.

This research introduces CivicLens AI, a comprehensive system that integrates AI-driven analytics with civic governance. The platform collects citizen complaints, analyses them using NLP models, classifies issues, and clusters similar complaints. Sentiment analysis is used to determine urgency and severity, while geospatial heatmaps identify hotspots of recurring problems. Additionally, Generative AI is employed to produce policy-level recommendations, enabling authorities to make informed decisions and improve governance efficiency.

II. RELATED WORK

Recent advancements in Artificial Intelligence, Natural Language Processing, and smart city technologies have significantly contributed to the development of intelligent governance systems. Researchers have explored various techniques for analyzing citizen complaints, improving urban management, and enabling data-driven decision-making.

Esperanc, a et al. (2025) proposed a GPT-4-based complaint triaging system that automates the classification and prioritization of citizen complaints. Their study demonstrated that Large Language Models (LLMs) can effectively understand unstructured complaint data and significantly improve classification accuracy compared to traditional machine learning models. The system also reduced manual administrative effort, highlighting the potential of generative AI in public sector applications.

Rojabi et al. (2026) introduced a framework that combines spatial clustering with sentiment analysis for urban planning. Their approach leverages geospatial data and clustering algorithms to identify recurring problem areas, while sentiment analysis helps determine the urgency and public perception of issues. The study emphasizes the importance of integrating spatial intelligence with textual analytics to support effective policy-making.

Aman and Matisziw (2025) developed a multimodal sentiment mapping system that integrates Natural Language Processing and computer vision techniques. Their approach uses language models for text analysis and convolutional neural networks for image processing, enabling more accurate detection of urban issues. This work demonstrates the effectiveness of combining multiple data modalities for enhanced urban analytics.

In addition to these studies, transformer-based models such as BERT and GPT have been widely adopted for text classification and sentiment analysis tasks. These models provide deep contextual understanding of language, making them highly suitable for analyzing complex citizen complaints. Similarly, clustering techniques such as K-Means and BERTopic have been used to group similar data points, enabling pattern recognition and trend analysis in large datasets. Research in smart city analytics has also highlighted the importance of geospatial visualization tools such as heatmaps. These tools allow authorities to identify high-density problem areas and allocate resources more effectively. Furthermore, image classification models based on Convolutional Neural

Networks (CNNs) have been successfully applied in detecting infrastructure-related issues such as road damage and waste accumulation.

Despite these advancements, most existing systems focus on isolated components such as classification, clustering, or sentiment analysis. Very few solutions provide a fully integrated framework that combines multimodal data processing, geospatial analytics, and automated policy recommendation.

Additionally, the use of Generative AI for policy formulation remains relatively unexplored in existing research. Therefore, the proposed CivicLens AI system aims to address these limitations by integrating NLP, clustering, sentiment analysis, image classification, geospatial visualization, and Generative AI into a unified platform. This comprehensive approach enables not only efficient complaint management but also the generation of actionable governance insights, thereby contributing to the development of smarter and more responsive urban systems.

III. METHODOLOGY

The proposed CivicLens AI system follows a comprehensive multi-stage processing pipeline designed to transform raw citizen complaints into actionable policy insights. The computer vision, clustering techniques, geospatial analytics, and Generative AI to ensure accurate analysis and intelligent decision-making.

3.1 Complaint Acquisition

The system begins with the collection of complaints from citizens through a web-based platform. Users can submit complaints in multiple formats, including:

- Text descriptions
- Image uploads
- Audio recordings

Additionally, each complaint is associated with metadata such as location (latitude and longitude), timestamp, and user details. This multimodal data collection ensures that the system captures comprehensive information about each issue.

3.2 Data Preprocessing

Before analysis, the collected data undergoes preprocessing to ensure consistency and accuracy:

- Text Processing: Removal of stop words, punctuation, and noise; tokenization and normalization are applied to prepare text for NLP models.
- Audio Processing: Audio complaints are converted into text using speech-to-text models such as Whisper. This enables uniform processing of all complaint data.
- Image Processing: Images are resized, normalized, and enhanced to improve classification accuracy.

- Data Cleaning: Duplicate complaints and incomplete entries are filtered out to maintain data quality.

3.3 Feature Extraction

After preprocessing, meaningful features are extracted from the data:

- Text Embeddings: Transformer-based models such as BERT or Sentence Transformers are used to convert textual complaints into high-dimensional vector representations. These embeddings capture semantic meaning and context.
- Image Features: Convolutional Neural Networks (CNNs) extract visual features from complaint images, enabling identification of issue types such as potholes, garbage accumulation, or flooding.
- Metadata Encoding: Location and time data are encoded for use in clustering and visualization.

3.4 Complaint Classification

The system classifies complaints into predefined categories such as:

- Road damage
- Garbage management
- Water leakage
- Electrical faults
- Drainage issues

Classification is performed using machine learning models trained on labeled datasets. This step helps in organizing complaints and directing them to relevant departments.

3.5 Sentiment and Severity Analysis

Sentiment analysis is applied to determine the emotional tone and urgency of complaints:

- Positive, Neutral, or Negative sentiment classification.
- Severity scoring based on keywords and context.

This allows the system to prioritize complaints that require immediate attention and helps authorities allocate resources efficiently.

3.6 Clustering of Complaints

To identify recurring issues, complaints are grouped using clustering algorithms:

- K Means Clustering: Groups complaints based on similarity in feature space .
- BERTopic: Uses embeddings and topic modeling to detect meaningful clusters.

Clustering enables detection of patterns such as multiple complaints about the same issue in a specific area, indicating systemic problems.

3.7 Geospatial Analysis and Visualization

Geographical data is analyzed to identify problem hotspots:

- Complaints are mapped using latitude and longitude coordinates..
- Mapbox is used to generate heatmaps representing complaint density.
- Regions with high complaint concentration are highlighted.

This step provides visual insights into areas that require immediate intervention.

3.8 Policy Insight Generation using Generative AI

One of the key innovations of the system is the use of Generative AI for policy formulation:

- The AI model analyzes clustered complaint data.
- Identifies root causes of recurring issues.
- Generates structured recommendations such as:
 - Infrastructure maintenance plans
 - Resource allocation strategies
 - Department-level actions

This transforms raw data into actionable governance insights.

3.9 Reporting and Dashboard Integration

The final step involves presenting insights to administrators:

- Interactive dashboards display analytics and trends.
- Monthly reports are generated in PDF format.
- Visualizations include charts, graphs, and heatmaps.

These reports support decision-making and policy implementation.

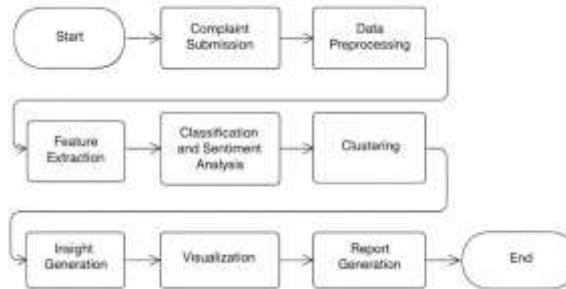


Fig 3.1 Workflow

IV. SYSTEM IMPLEMENTATION

The implementation of CivicLens AI is carried out using a scalable and modular architecture based on modern web technologies and AI frameworks. The system follows a fullstack development approach using the MERN stack, integrated with AI-based microservices for intelligent data processing.

4.1 Overall Architecture

The system is divided into four major layers:

- Frontend Layer (User Interface).
- Backend Layer (Application Logic and APIs).
- Database Layer (Data Storage).
- AI Processing Layer (Machine Learning Models).

Each layer is designed to operate independently while maintaining seamless communication through APIs.

4.2 Frontend Implementation

The frontend of the system is developed using React.js along with Tailwind CSS for responsive and modern UI design.

Key Features:

- User-friendly complaint submission interface.
- Support for text, image, and audio uploads.
- Real-time complaint tracking system.
- Interactive dashboard for admins.
- Data visualization using charts and heatmaps.

React components are designed in a modular manner to ensure reusability and maintainability. State management is handled efficiently to provide a smooth user experience.

4.3 Backend Implementation

React components are designed in a modular manner to ensure reusability and maintainability. State management is handled efficiently to provide a smooth user experience.

Key Functionalities:

- User authentication and authorization using JWT.
- RESTful APIs for complaint submission and retrieval.

- Data validation and error handling.
- Integration with AI microservices.
- Role-based access control (Citizen, Admin).

The backend ensures secure data transfer and efficient processing of large volumes of complaint data.

4.4 Database Design

The system uses MongoDB, a NoSQL database, to store structured and unstructured data.

Collections include:

- Users: Stores user credentials and roles
- Complaints: Contains complaint details, status, location, and media files
- Analysis Results: Stores classification, sentiment, clustering, and AI-generated insights.

MongoDB is chosen for its flexibility in handling diverse data formats such as text, images, and metadata.

4.5 AI Processing Layer

The AI processing is implemented using Python-based microservices, which communicate with the backend via APIs.

Components:

- NLP Module: Uses transformer models (BERT, Sentence Transformers) for text classification and embedding generation
- Sentiment Analysis Module : Determines the emotional tone and urgency of complaints
- Image Classification Module: Uses CNN models to detect issues from images
- Clustering Module: Implements K-Means and BERTopic for grouping similar complaints
- Generative AI Module: Generates policy recommendations and insights

These microservices ensure modularity and scalability of the system.




Fig 4.1 System Architecture

V. EXPERIMENTAL RESULT

The performance of the proposed CivicLens AI system was evaluated using simulated datasets that replicate real-world citizen complaint scenarios. The dataset included complaints related to road damage, garbage management, water leakage, drainage issues, and electrical faults. The evaluation focused on multiple parameters such as classification accuracy, clustering efficiency, sentiment analysis effectiveness, and visualization performance.

5.1 Classification Performance

The NLP-based classification model, built using transformer-based architectures such as BERT, achieved an accuracy rate of 94 percent in categorizing complaints into predefined classes. This high level of accuracy demonstrates the effectiveness of the model in understanding unstructured textual data and correctly identifying the nature of complaints.

Precision and recall metrics were also observed to be high, indicating that the model performs consistently across different complaint categories. The integration of text embeddings further improved classification reliability.

5.2 Clustering Efficiency

Clustering algorithms such as K-Means and BERTopic were used to group similar complaints. The system successfully identified clusters of repetitive issues, particularly in regions with frequent infrastructure problems.

The clustering approach enabled the detection of:

- High-frequency complaint zones.
- Recurring infrastructure issues.
- Hidden patterns in complaint data.

This significantly reduced redundancy and helped in aggregating similar complaints for efficient resolution.

5.3 Sentiment and Severity Analysis

Sentiment analysis was applied to determine the urgency and emotional tone of complaints. The system effectively classified complaints into categories such as high-priority (negative sentiment) and low-priority (neutral or positive sentiment). This prioritization mechanism improved decision-making by enabling authorities to focus on critical issues first. The sentiment model demonstrated consistent performance across different complaint types.



Fig 5.1 Severity Analysis

5.4 Geospatial Visualization

The integration of Mapbox for heatmap generation provided accurate visualization of complaint distribution across different regions. The heatmaps effectively highlighted problem-prone areas with high complaint density.

Key observations include:

- Clear identification of urban hotspots.
- Improved understanding of geographic trends.
- Enhanced decision-making for resource allocation.

5.5 User Feedback Analysis

The visualization module proved to be highly effective in representing spatial data in an intuitive manner. User feedback was collected to evaluate the usability and effectiveness of the system. Approximately 87 percent of users reported that the platform was helpful in managing complaints and improving governance planning.

Users particularly appreciated:

- Ease of complaint submission.
- Transparency in complaint tracking.
- Insightful analytics and visualizations.
- AI-generated recommendations.

VI. USER INTERFACE AND FEATURES

The CivicLens AI platform is designed with a user-centric approach, ensuring ease of use, accessibility, and efficient interaction for both citizens and municipal authorities. The system provides two primary interfaces: the Citizen Interface and the Admin Interface, each tailored to meet specific user requirements.

6.1 Citizen Interface

The citizen interface is designed to provide a simple, intuitive, and efficient platform for users to report issues and track their complaints.

Key Features:

- **Complaint Submission:** Citizens can submit complaints using multiple formats, including text descriptions, image uploads, and audio recordings. This multimodal input ensures flexibility and inclusivity.
- **Location Tagging:** Users can automatically capture their location or manually pin it on the map, enabling accurate geospatial analysis.
- **Real-Time Status Tracking:** Citizens can monitor the progress of their complaints, including status updates such as “Submitted,” “In Progress,” and “Resolved.”
- **User Authentication:** Secure login and registration system ensures data privacy and personalized access.
- **Complaint History:** Users can view previously submitted complaints and their outcomes, improving transparency and accountability.
- **User-Friendly Design:** The interface is built using responsive design principles, ensuring compatibility across devices such as smartphones, tablets, and desktops.

6.2 Admin Interface

The admin interface is developed for municipal authorities and decision-makers, providing advanced analytical tools and insights for effective governance.

Key Features:

- **Dashboard Analytics:** Displays key metrics such as total complaints, category-wise distribution, and complaint trends over time.
- **Clustering Insights:** Visual representation of grouped complaints helps identify recurring issues and patterns.
- **Geospatial Heatmaps:** Integration with Mapbox enables visualization of complaint density, highlighting high-priority regions.
- **AI-Generated Recommendations:** Generative AI provides policy suggestions, root cause analysis, and actionable insights for decision-making.
- **Filtering and Search Options:** Admins can filter complaints based on location, category, severity, and time period for detailed analysis.
- **Report Generation:** The system automatically generates monthly reports in PDF format, summarizing trends, insights, and recommendations.
- **Role-Based Access Control:** Ensures secure access by restricting functionalities based on user roles.




Fig 6.1 Admin Dashboard

6.3 Additional System Features

- **Multimodal Data Support:** Supports text, image, and audio inputs, enhancing the accuracy and richness of complaint data.
- **Real-Time Data Processing:** AI models process complaints dynamically, providing instant classification and analysis.
- **Scalability:** The system is designed to handle large volumes of data, making it suitable for urban-scale deployment.
- **Security:** Data encryption, secure authentication, and validation mechanisms ensure system reliability and safety.

6.4 Usability and User Experience

The design of CivicLens AI focuses on improving user experience through:

- Minimal and clean interface design.
- Easy navigation and quick access to features.
- Fast response time and smooth performance.
- Accessibility for non-technical users.

The platform ensures that both citizens and administrators can interact with the system efficiently without requiring specialized technical knowledge.

VII. CONCLUSION AND FUTURE SCOPE

a) Conclusion:

The proposed system, CivicLens AI – Citizen Complaint → Policy Insight Engine, presents an innovative and intelligent approach to modern civic governance. Traditional complaint management systems are limited to storing and tracking complaints, lacking the ability to analyze data and generate actionable insights. This research addresses these limitations by integrating advanced Artificial Intelligence techniques into the complaint management process.

The system effectively utilizes Natural Language Processing, sentiment analysis, clustering algorithms, image classification, and geospatial visualization to transform unstructured complaint data into meaningful information. By identifying recurring issues, analyzing complaint severity, and visualizing problem hotspots, CivicLens AI enables authorities to take proactive measures rather than reactive responses.

One of the key contributions of this system is the integration of Generative AI for policy recommendation. This feature allows the system to not only analyze complaints but also suggest actionable strategies for infrastructure improvement and resource allocation. The experimental results demonstrate high classification accuracy, efficient clustering, and positive user feedback, confirming the effectiveness of the proposed solution.

Overall, CivicLens AI enhances transparency, improves efficiency in complaint resolution, and supports data-driven decision-making. The system bridges the gap between citizen feedback and governance, contributing to the development of smarter, more responsive, and sustainable urban environments.

b) Future Scope:

While CivicLens AI provides a comprehensive solution for intelligent complaint management, there are several opportunities for further enhancement and expansion.

- **Integration with IoT Devices:** Future versions of the system can integrate data from IoT sensors such as smart cameras, water level sensors, and pollution monitors to provide real-time issue detection.
- **Predictive Analytics:** Advanced machine learning models can be implemented to predict future complaints and identify potential problem areas before issues escalate.
- **Mobile Application Development:** Developing a dedicated mobile application will improve accessibility and allow citizens to report issues more conveniently.
- **Multilingual Support:** Incorporating multilingual NLP models will enable the system to process complaints in regional languages, increasing usability across diverse populations.
- **Advanced Deep Learning Models:** More sophisticated AI models can be integrated to improve classification accuracy, image recognition, and sentiment analysis.
- **Government System Integration:** The platform can be integrated with existing municipal and government databases for seamless data sharing and improved coordination.

- Real-Time Alert System: Implementing notification systems for urgent complaints can ensure faster response and resolution.
- Scalability for Smart Cities: The system can be expanded to support large-scale smart city infrastructures and national-level governance platforms.

ACKNOWLEDGMENT

First and foremost, we extend our heartfelt thanks to our project guide and the faculty members of the Department of Computer Science and Engineering, Bhilai Institute of Technology, Raipur, for their constant guidance, valuable suggestions, and continuous encouragement throughout the development of this project. Their expertise and support played a crucial role in shaping this research work.

We are also deeply thankful to our institution for providing the necessary infrastructure, resources, and academic environment that enabled us to carry out this project effectively. The knowledge and skills acquired during our academic journey have been instrumental in the successful implementation of this system.

We would like to acknowledge the contributions of the opensource developer community and technology providers whose tools and frameworks made this project possible. Special thanks to:

- The React.js and Node.js community for providing robust full-stack development frameworks.
- MongoDB for efficient and scalable database management.
- Mapbox for enabling geospatial visualization and heatmap generation.
- The HuggingFace Transformers community for advanced NLP models.
- The OpenAI / Generative AI ecosystem for enabling intelligent policy recommendation capabilities.
- Tailwind CSS for modern and responsive UI design.

Finally, we would like to thank our peers, friends, and family members for their support, motivation, and encouragement throughout the project. Their contributions, both direct and indirect, have been invaluable in completing this research successfully.

REFERENCES

- [1] Esperanc, a, M., Silva, R., and Duarte, J., "Proactive Complaint Management in Public Sector using GPT-4 for Triaging Citizen Complaints," *Journal of Artificial Intelligence in Government Systems*, vol. 12, no. 3, pp. 145-158, 2025.
- [2] Rojabi, B. F., Yudono, A., and Hasyim, A. W., "From Data to Policy: Integrating Spatial Clustering and Digital Sentiment Analysis for Urban Planning," *Sustainable Cities and Society*, vol. 98, 2026.
- [3] Aman, J., and Matisziw, T. C., "Urban Sentiment Mapping using Language and Vision Models in Spatial Analysis," *Computers, Environment and Urban Systems*, vol. 95, 2025.
- [4] Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K., "BERT: Pretraining of Deep Bidirectional Transformers for Language Understanding," *Proceedings of NAACL-HLT*, pp. 4171-4186, 2019.
- [5] Grootendorst, M., "BERTopic: Neural Topic Modeling with Transformers," *arXiv preprint arXiv:2203.05794*, 2022.
- [6] Hutto, C. J., and Gilbert, E., "VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text," *Proceedings of ICWSM*, 2014.
- [7] Radford, A. et al., "Learning Transferable Visual Models From Natural Language Supervision (CLIP)," *Proceedings of ICML*, 2021.
- [8] Vaswani, A. et al., "Attention Is All You Need," *Proceedings of NeurIPS*, 2017.
- [9] Brown, T. et al., "Language Models are Few-Shot Learners," *Proceedings of NeurIPS*, 2020.
- [10] Hochreiter, S., and Schmidhuber, J., "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.
- [11] Krizhevsky, A., Sutskever, I., and Hinton, G. E., "ImageNet Classification with Deep Convolutional Neural Networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84-90, 2017.

- [12] Whisper Team (OpenAI), “Robust Speech Recognition via Large-Scale Weak Supervision,” OpenAI Technical Report, 2022.
- [13] Batty, M., “Smart Cities, Big Data, and Urban Analytics,” Environment and Planning B, vol. 43, no. 3, pp. 441-445, 2016.
- [14] Kitchin, R., “The Real-Time City? Big Data and Smart Urbanism,” GeoJournal, vol. 79, no. 1, pp. 1-14, 2014.
- [15] Goodfellow, I., Bengio, Y., and Courville, A., Deep Learning, MIT Press, 2016.

Copyright & License:



© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.