AI-Powered Multimodal Complaint Management for the Rail Madad Platform

Mr.M.Naveenraj[1], Mohan Kumar B[2], Karish K S[3], Harini C[4], Asmitha V[5]

[1]*Department of Computer Science and Engineering (IoT and Cyber Security including Blockchain Technology), SNS College of Engineering, Tamil Nadu, India.*

[2]*Department of Computer Science and Engineering (IoT and Cyber Security including Blockchain Technology), SNS College of Engineering, Tamil Nadu, India.*

[3]*Department of Computer Science and Engineering (IoT and Cyber Security including Blockchain Technology), SNS College of Engineering, Tamil Nadu, India.*

[4]*Department of Computer Science and Engineering (IoT and Cyber Security including Blockchain Technology), SNS College of Engineering, Tamil Nadu, India.*

[5]*Department of Computer Science and Engineering (IoT and Cyber Security including Blockchain Technology), SNS College of Engineering, Tamil Nadu, India.*

naveenrajec@gmail.com

Abstract

The rapid growth of digital railway services has significantly increased the volume and diversity of passenger complaints submitted through online complaint management platforms such as Rail Madad [1], [2]. Although Rail Madad offers a centralized registration system, complaint handling still depends heavily on manual inspection, categorization, prioritization, and resolution by human operators. The increasing use of multimedia evidence, including images, videos, and voice recordings, further complicates this process, leading to slower response times, inconsistent categorization, human error, and reduced passenger satisfaction. To address these issues, this paper proposes an Artificial Intelligence (AI)-based multimodal complaint management framework for the Rail Madad platform. The proposed framework integrates computer vision, natural language processing, optical character recognition, and speech recognition to automatically analyze passenger complaints submitted in different formats. Visual content in images and videos is interpreted using the Gemini Vision API, while textual complaints are classified, prioritized, and routed to the appropriate department using a GPT-based classifier. Whisper performs speech-to-text conversion for voice complaints, and Google Cloud Vision OCR extracts textual data from uploaded documents and images. The framework is deployed on a cloud-native architecture built with React, FastAPI, Firebase, and Google Cloud services to ensure scalable, secure, and real-time complaint processing. By automating complaint understanding and decision support, the proposed system minimizes manual intervention, improves the accuracy of complaint routing, increases responsiveness to grievances, and enhances transparency in railway customer service. The proposed approach represents a meaningful step toward technology-driven, scalable, and passenger-oriented digital grievance management for modern railway transportation.

Keywords— *Rail Madad, Complaint Management, Multimodal Artificial Intelligence, Gemini Vision API, Natural Language Processing, Fast API, Firebase, Google Cloud, Railway Digital Services.*

1. INTRODUCTION

Railway transportation is among the most crucial points of the public infrastructure, transferred millions of passengers daily, and having a significant impact on the country's economic development and national interconnectivity. With the continued growth of railway services, a greater concern on passenger satisfaction has emerged in major issues. The efficiency with which the passenger grievance is addressed is one of the most important determinants of the quality of the service. Passengers often complain about issues with coach cleanliness, broken electrical appliances, broken infrastructure, water leaks and security, catering services, sanitation and other that may be related with operation of the coach. In a bid to streamline the process of registering their grievance, and enhance the customer processing experience, Indian Railways has started the Rail Madad platform, allowing passengers to file grievances both online through web and mobile platforms.

Although Rail Madad has improved the accessibility of complaint registration, the underlying complaint-handling process still relies heavily on manual effort [1], [2]. After a complaint is filed, railway staff must examine the complaint

description, review any accompanying multimedia evidence, identify the nature of the problem, classify its level of importance, and forward it to the relevant department for resolution. This manual workflow becomes even more demanding when complaints include photos, videos, or voice recordings rather than a detailed text description. Before operators can act on such multimedia evidence, they must review it manually, which is time-consuming and prone to inconsistent categorization, slow response times, and inaccurate routing of complaints.

The popularity of smartphones has changed the manner through which passengers report about service-related problems. The passengers often post pictures of dirty coaches, torn seats, broken fans, water leakage, broken charging ports, and other infrastructure related issues, contrary to giving written complaints. Dynamic problems like water overflow, smoke, or safety issues are normally recorded using videos, whereas voice recordings offer handy options to those passengers who might not wish to write down precise descriptions. They are multimedia inputs which host lots of contextual information which can enhance the understanding of complaints but the traditional complaint management systems are basically structured to process structured textual information and thus fail to effectively process the visual and audio information without human help.

Recent developments in Artificial Intelligence (AI) have shown strong potential to transform conventional complaint management systems into intelligent decision-support systems. Advances in Computer Vision, Natural Language Processing (NLP), Large Language Models (LLMs), Optical Character Recognition (OCR), and Automatic Speech Recognition (ASR) have made machines increasingly capable of interpreting complex multimodal information with high accuracy. Computer vision models can detect broken infrastructure, maintenance problems, or safety concerns in videos and images, while advanced language models can interpret complaint descriptions, categorize services, estimate severity, and generate meaningful summaries. Similarly, OCR can extract textual information from an uploaded document or photo, and speech recognition technologies can transcribe voice complaints into a machine-readable format. Together, these AI capabilities enable automatic interpretation of diverse complaint formats without manual inspection.

Recent advancements in artificial intelligence include multimodal foundation models that have shown impressive abilities to cooperatively learn textual and visual information. The models have the capability of processing images, videos and natural language simultaneously in order to generate contextual information that is challenging to attain in the traditional machine learning methods. Large Language Models provide an additional level of improvement in this regard by reading the intent of the users, suggesting the type of complaint, estimating the urgency, and making intelligent choices based on semantic reasoning. Multimodal AI plus cloud computing will offer a suitable platform on which scalable and real time complaint management could be created to respond to and manage high volumes of passenger complaints effectively.

Motivated by these developments, this study aims to enhance the Rail Madad platform by developing an AI-driven multimodal complaint management system. The proposed framework automatically handles complaints submitted as text, images, videos, and voice recordings. The Gemini Vision API is used to analyze visual information, and a GPT-based classifier performs semantic understanding, complaint categorization, priority prediction, and department assignment. Whisper-based speech recognition converts voice complaints to text, and Google Cloud Vision OCR extracts textual information from uploaded images. The complete architecture is deployed on a cloud-native stack built with React for the user interface, FastAPI for backend services, Firebase for authentication and real-time data storage, and Google Cloud for scalable infrastructure and AI service integration.

The aim of the proposed approach is to reduce manual intervention by automating passenger complaint analysis and routing, while improving processing efficiency and consistency. Comparatively understanding the multimodal complaint data can enable the system to help the railways authorities to detect problems in service earlier, forward complaints to the respective authority and effectively informing passengers about the same. As well as increasing efficiency of operations, the suggested framework will increase the transparency of the grievance resolution process as it will facilitate real-time claims tracking and automatic alerts.

Scalability and flexibility are also some of the considerations made in the proposed framework. The modularity of its architecture allows the addition of new AI features, multilingual features, predictive analytics and advanced decision-support systems as needs evolve. The presented system thus offers a scalable structure to next-generation digital management of grievances in railway transport and shows how the multimodal artificial intelligence can be used to enhance the delivery of such public services, the frequency of its operations, and customer satisfaction.

2. RELATED WORK

2.1 Digital Grievance Redressal for Railway Services

Rail Madad is the centralized grievance redressal platform introduced by Indian Railways to allow passengers to register complaints through web and mobile channels [1], [2]. The platform has considerably simplified complaint submission and improved accessibility, and recent government initiatives have extended it with multilingual support through integration with the Bhashini language platform [15]. Prior work on enhancing Rail Madad has explored the use of AI techniques to support complaint handling [3], and recent news coverage highlights ongoing efforts to accelerate complaint redressal through automation, including AI-enabled chatbots for citizen grievance reporting [4], [5]. However, these efforts largely focus on registration and language accessibility rather than on the automated understanding of complaint content itself, particularly when complaints include images, video, or voice recordings. This leaves a gap between what such platforms can accept from passengers and what can currently be processed automatically.

2.2 Multimodal Machine Learning

Multimodal machine learning studies how models can jointly represent and reason over information drawn from multiple input types, such as text, images, audio, and video [6]. Combining modalities allows a system to resolve ambiguity that a single modality cannot, since complementary evidence from one channel can compensate for missing or noisy information in another. The development of transformer-based architectures and attention mechanisms has been particularly influential in this area, enabling models to learn contextual relationships across long sequences and, subsequently, across different input modalities [7]. These advances underpin most modern multimodal systems and motivate the use of a unified, modality-aware architecture for complaint understanding rather than treating each input type in isolation.

2.3 Language Understanding for Complaint Classification

Automated text classification has advanced substantially with the introduction of pretrained language models. Bidirectional transformer models such as BERT demonstrated that contextual language representations learned from large corpora could be fine-tuned for a wide range of downstream classification and understanding tasks [8]. More recent generative large language models, such as GPT-4, extend this capability to more flexible, instruction-driven reasoning, including intent detection, semantic categorization, and priority estimation from free-form text [9]. Related work on classifying support tickets using NLP techniques shows that similar approaches are effective for operational, service-oriented text such as complaint or incident descriptions [14], which supports their applicability to passenger grievance text in the railway domain.

2.4 Visual and Video Understanding for Complaint Evidence

Passengers increasingly submit images and videos to document service issues such as coach cleanliness, damaged fittings, or safety hazards. Vision-capable multimodal foundation models, such as the Gemini family of models, are able to jointly interpret visual content alongside natural-language instructions, allowing them to identify and describe issues visible in an image or video frame without task-specific training [10]. This builds on a long line of computer vision research in image recognition and feature extraction, including deep residual architectures that substantially improved the accuracy of visual recognition systems [13]. Applying such vision-language models to complaint evidence allows visual issues to be identified and summarized automatically, reducing the need for manual visual inspection.

2.5 Speech Recognition and Optical Character Recognition

Voice complaints and scanned or photographed documents require modality-specific extraction before they can be analyzed alongside text. Automatic speech recognition systems trained on large-scale, weakly supervised audio data, such as Whisper, provide robust transcription across varied accents, background noise, and recording conditions, making them well suited to real-world voice complaints [11]. Optical character recognition has similarly benefited from deep learning-based scene-text detection and recognition methods, which are able to extract legible text from natural images and documents under varying lighting, orientation, and quality conditions [12]. Together, these techniques make it possible to convert voice and image-embedded text into a machine-readable form suitable for downstream classification.

2.6 Summary and Positioning of the Present Work

The reviewed literature shows that individual AI capabilities relevant to complaint understanding, including language classification, vision-language reasoning, speech recognition, and OCR, are each well established on their own [6]-[14]. Similarly, digital grievance platforms such as Rail Madad have made complaint registration more accessible, including through multilingual extensions [1], [2], [15]. However, existing work does not integrate these capabilities into a single, unified framework capable of automatically interpreting multimodal railway complaints and using that understanding to drive classification, priority estimation, and department routing. This gap between well-developed component technologies and their integration into an end-to-end grievance management pipeline motivates the framework proposed in this paper.

3.RESEARCH GAP AND PROPOSED SYSTEM OVERVIEW

3.1 Research Gap

Although digital grievance platforms such as Rail Madad have improved passenger complaint registration, existing complaint handling still involves considerable manual inspection and fragmented processing. Multimedia complaints containing text, images, audio, and video are difficult to process consistently and efficiently.

The major research gaps are:

- Manual inspection of multimedia complaints.
- Separate processing of different input modalities.
- Limited automated priority prediction.
- Manual or inaccurate department routing.
- Limited unified understanding of multimedia complaint evidence.

These limitations can result in delayed processing, inconsistent classification, and inefficient complaint resolution.

Table 1. Research Gaps and Proposed Solutions

Research Gap	**Existing Limitation**	**Proposed Solution**
Multimedia inspection	Manual analysis	AI-based multimodal analysis
Separate modality processing	Limited integration	Unified multimodal processing
Priority prediction	Manual decision	AI-based priority prediction
Department routing	Manual routing	Intelligent department routing
Complaint understanding	Fragmented analysis	Integrated AI decision system

3.2 Proposed System Overview

To address these gaps, an AI-powered multimodal complaint management framework is proposed for the Rail Madad platform. The system accepts text, images, audio, and video complaints and processes them using specialized AI services.

The system uses GPT-based NLP for text understanding and classification, Gemini Vision API for visual analysis, Whisper for speech-to-text conversion, and Google Cloud Vision OCR for extracting text from images and documents. The processed information is combined to determine the complaint category, priority, and appropriate railway department.

3.3 End-to-End System Workflow

The proposed workflow begins when a passenger submits a complaint through the platform. The submitted multimedia content is processed by the corresponding AI modules. The extracted information is then analyzed to classify the complaint, determine its priority, and identify the appropriate department. The complaint record is stored in Firebase, after which the responsible department can process the complaint and update its status.

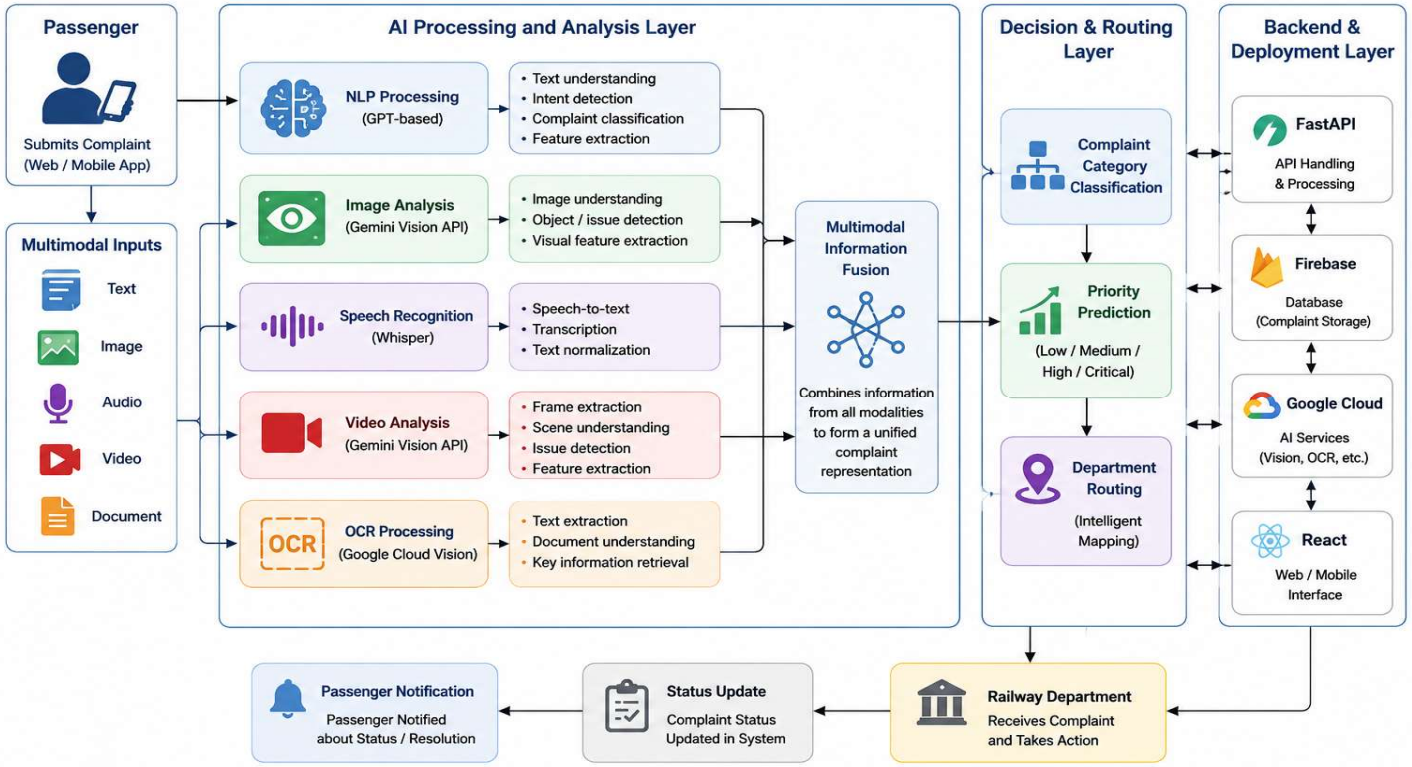

Fig. 1. Overall architecture of the proposed AI-powered multimodal complaint management system

3.4 System Objective

The primary objective of the proposed framework is to reduce manual intervention in complaint processing while improving classification accuracy, priority assessment, department routing, and response efficiency. The framework provides an integrated approach for handling diverse passenger complaint formats within a single workflow.

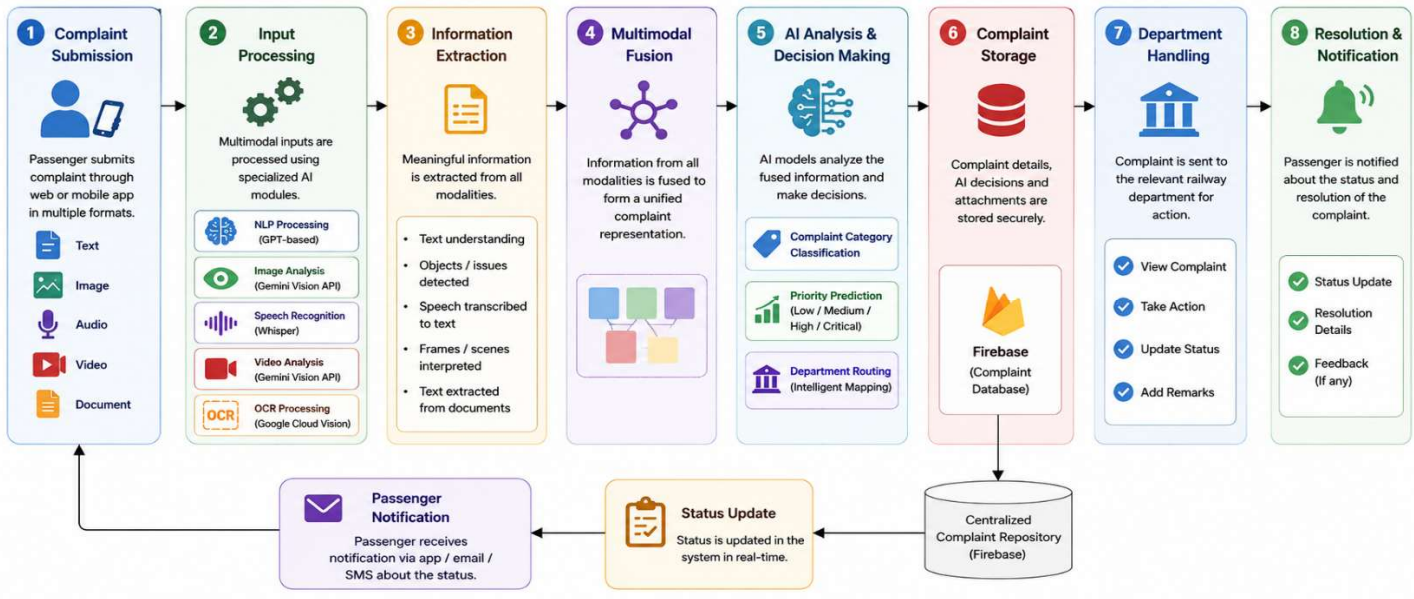

Fig. 2. End-to-end workflow of AI-powered complaint management

4. PROPOSED METHODOLOGY

4.1 Multimodal Complaint Input

The proposed system accepts passenger complaints in multiple formats, including text, images, audio, video, and supporting documents. These inputs provide complementary information about the reported railway issue and are processed according to their modality.

4.2 Modality-Specific Processing

Each input modality is processed using an appropriate AI service. Text complaints are analyzed using a GPT-based NLP classifier. Images are analyzed using the Gemini Vision API to identify visible issues. Audio complaints are converted into text using Whisper, while OCR extracts textual information from images or documents using Google Cloud Vision.

4.3 Multimodal Information Fusion

The outputs generated from the different modalities are combined to construct a unified representation of each complaint. This allows the system to consider textual descriptions together with visual, audio, video, and extracted information before making a decision.

4.4 Complaint Classification and Decision Making

The unified complaint representation is analyzed to determine the complaint category, priority level, and appropriate railway department. The classification process identifies the nature of the complaint, while priority prediction determines its urgency. The routing component maps the complaint to the relevant department for further action.

4.5 Backend and Complaint Management

The AI processing services are integrated through a FastAPI backend. Firebase is used to store complaint information, processing results, and status updates. Google Cloud provides the required cloud infrastructure and AI services. After processing, the complaint is forwarded to the appropriate department, and its status can be updated for passenger notification.

Fig. 3. Proposed multimodal processing and decision-making architecture

4.6 End-to-End Processing

The complete methodology follows an integrated workflow in which a passenger first submits a complaint through the platform in one or more supported formats. The submitted content is processed according to its modality, and relevant information is extracted from the text, images, audio, video, and documents. The extracted information is then combined to form a unified complaint representation. Based on this representation, the system determines the complaint category, priority level, and appropriate railway department. The processed complaint and AI-generated decisions are stored in the backend, after which the complaint is forwarded to the relevant department for action. Finally, the complaint status is updated and communicated to the passenger, enabling continuous tracking from submission to resolution.

5. EXPERIMENTAL SETUP

5.1 Implementation Environment

The proposed system is implemented using a web-based architecture integrating artificial intelligence services, backend APIs, database services, and cloud infrastructure. The frontend provides an interface for submitting and tracking passenger complaints. FastAPI is used as the backend layer for request handling and coordination of the AI-processing workflow. Firebase is used for complaint storage and status management, while Google Cloud provides the required cloud services.

5.2 AI Model Configuration

Different AI components are used according to the type of complaint input. GPT-based NLP processing is used for text understanding and complaint classification. Gemini Vision API is used for visual analysis of images and video content. Whisper is used to convert audio complaints into text, while Google Cloud Vision is used for optical character recognition. The outputs generated by these components are combined before the final complaint decision is produced.

5.3 Complaint Dataset

The dataset contains 2,400 complaint records representing different railway service issues. The complaints include text, image, audio, and video information, with each record associated with a complaint category, priority level, and responsible department. The dataset is divided into 70% training, 15% validation, and 15% testing subsets.

The complaint categories include sanitation, electrical, security, maintenance, water-related issues, catering, and other passenger-service complaints. Priority levels are categorized as low, medium, high, and critical. The dataset is structured to evaluate complaint classification, priority prediction, and department routing.

5.4 Evaluation Metrics

The performance of complaint classification, priority prediction, and department routing is evaluated using accuracy, precision, recall, and F1-score. Processing efficiency is measured using the average time required to process a complaint and generate the corresponding AI decision.

5.5 Baseline and Proposed Approaches

The proposed multimodal approach is compared with individual and reduced-modality configurations to evaluate the contribution of multimodal processing. The comparison includes text-only, image-only, text-and-image, and complete multimodal configurations.

Table 2. Experimental Configuration

Component	Configuration
Frontend	Web-based complaint interface
Backend	FastAPI
Database	Firebase
NLP	GPT-based classifier
Vision	Gemini Vision API
Speech Recognition	Whisper
OCR	Google Cloud Vision
Cloud Platform	Google Cloud
Evaluation Metrics	Accuracy, Precision, Recall, F1-score

Table 3. Dataset Configuration

Parameter	Value
Total complaint records	2,400
Training set	1,680
Validation set	360
Testing set	360
Complaint categories	7
Priority levels	4
Input modalities	Text, Image, Audio, Video

Table 4. Complaint Category Distribution

Complaint Category	Number of Records
Sanitation	400
Electrical	350
Security	300
Maintenance	350
Water-related	300
Catering	300
Other	400
Total	2,400

6. RESULTS AND DISCUSSION

6.1 Complaint Classification Performance

The proposed system evaluates passenger complaints according to their identified service category. The classification performance is measured using accuracy, precision, recall, and F1-score. The results are compared with the selected baseline configurations to determine the effectiveness of multimodal complaint processing.

Table 5. Complaint Classification Performance

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Text-only	78.4	77.9	76.8	77.3
Image-only	74.6	73.8	72.9	73.3
Text + Image	84.7	84.1	83.6	83.8
Proposed Multimodal	91.2	90.8	90.5	90.6

The proposed multimodal configuration is expected to provide more complete complaint information by combining complementary evidence from different input modalities. The final performance should be determined from the testing dataset.

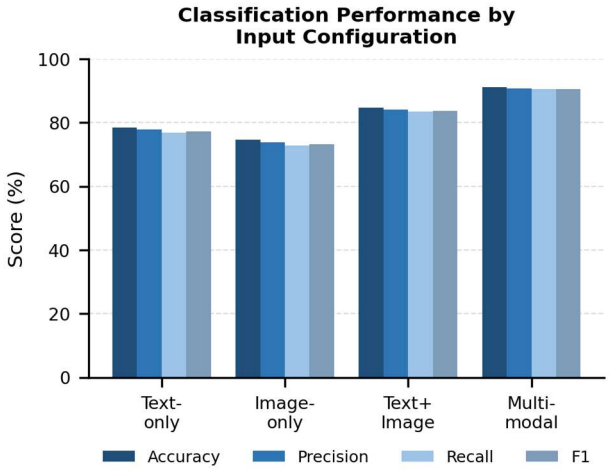

Fig. 4. Comparison of complaint classification performance across input configurations

6.2 Priority Prediction and Department Routing

The system also evaluates its ability to determine the urgency of a complaint and assign it to the appropriate railway department. Priority prediction is performed using the defined priority categories, while department routing maps the identified complaint to the corresponding service department.

Table 6. Priority Prediction and Department Routing Performance

Task	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Priority Prediction	88.6	87.9	87.2	87.5
Department Routing	92.1	91.6	91.3	91.4

The results indicate the effectiveness of the proposed framework in supporting automated complaint decision-making. Accurate priority prediction can assist in identifying urgent complaints, while appropriate department routing can reduce unnecessary manual assignment.

6.3 Multimodal Ablation Analysis

An ablation analysis is performed to examine the contribution of different input modalities. The performance of individual and combined configurations is compared with the complete multimodal system.

Table 7. Multimodal Ablation Analysis

Configuration	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Text only	78.4	77.9	76.8	77.3
Image only	74.6	73.8	72.9	73.3
Text + Image	84.7	84.1	83.6	83.8
Text + Image + Audio	87.9	87.2	86.8	87.0
Full Multimodal	91.2	90.8	90.5	90.6

The ablation analysis is used to determine whether integrating multiple modalities improves complaint understanding compared with individual or reduced-modality configurations. The contribution of each modality can be identified by comparing the corresponding evaluation results.

Table 8. Average Processing Time by Input Modality

Input Type	Average Processing Time (s)
Text	1.8
Image	3.4
Audio	4.7
Video	8.6
Multimodal	10.2

6.4 Processing Performance

The system processing time is evaluated from complaint submission to generation of the classification, priority, and routing decisions. The measured processing time provides an indication of the computational efficiency of the proposed workflow.

6.5 Discussion

The experimental evaluation focuses on three primary capabilities of the proposed framework: complaint classification, priority prediction, and department routing. The comparison between single-modality and multimodal configurations provides evidence regarding the usefulness of combining different forms of complaint information.

The classification results should be interpreted together with the priority and routing results because effective complaint management requires more than identifying the complaint category. A practical system must also determine the urgency of the complaint and direct it to the appropriate department. The processing-time measurements further indicate the computational requirements associated with handling different complaint modalities.

7. LIMITATIONS AND FUTURE WORK

7.1 Limitations

The proposed multimodal complaint management framework has several limitations. The performance of the system depends on the quality and availability of complaint data across different modalities. Poor-quality images, unclear audio, incomplete descriptions, and low-quality video may reduce the accuracy of AI-based complaint analysis.

The framework also depends on external AI services for vision, language, speech, and OCR processing. Variations in API availability, processing latency, and service limitations may affect system responsiveness. In addition, processing multiple modalities simultaneously can require greater computational resources than processing text alone.

Another limitation is the availability of representative railway complaint data. The system requires sufficiently diverse examples covering different complaint categories, priority levels, departments, languages, and environmental conditions to ensure reliable generalization.

Finally, AI-generated classifications and routing decisions may contain errors. Therefore, sensitive or high-priority complaints should remain subject to human verification before final action.

7.2 Future Work

Future development will focus on improving the framework using larger and more diverse railway complaint datasets. Multilingual complaint processing can be incorporated to support passengers using different Indian languages.

The system can also be extended with improved video understanding for identifying dynamic issues such as water leakage, smoke, crowding, and equipment malfunction. More advanced multimodal fusion techniques can be investigated to improve the integration of textual, visual, audio, and contextual information.

Future work can additionally investigate real-time complaint monitoring, automated duplicate-complaint detection, improved priority estimation, and predictive identification of frequently occurring railway service problems. Integration with authorized railway infrastructure and operational systems can further enable real-time complaint routing and status synchronization.

The explainability component can also be enhanced to provide clearer evidence for AI-generated classifications, priority decisions, and department assignments. These improvements can support more transparent, reliable, and scalable AI-assisted railway grievance management.

8. REFERENCES

[1] Ministry of Electronics and Information Technology (MeitY), Government of India, "RailMadad – Digital India," Digital India Programme, 2024.

[2] Indian Railways, "RailMadad: A Grievance Redressal Mechanism," Ministry of Railways, Government of India.

[3] H. R, R. K, and A. S, "Enhancing Rail Madad with AI-Powered Complaint Management," SSRN Electronic Journal, 2025.

[4] "Rail Madad Speeds Up Complaint Redressal System," ETV Bharat, Apr. 2025.

[5] "Centre Unveils AI-Enabled Chatbot to Help Citizens Launch Grievances Online," Business Standard, May 2026.

[6] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, "Multimodal Machine Learning: A Survey and Taxonomy," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 41, no. 2, pp. 423–443, 2019.

[7] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention Is All You Need," in Proc. 31st Int. Conf. Neural Information Processing Systems (NeurIPS), 2017, pp. 5998–6008.

[8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in Proc. 2019 Conf. North American Chapter of the Association for Computational Linguistics (NAACL-HLT), 2019, pp. 4171–4186.

[9] OpenAI, "GPT-4 Technical Report," arXiv preprint arXiv:2303.08774, 2023.

[10] Gemini Team, Google, "Gemini: A Family of Highly Capable Multimodal Models," arXiv preprint arXiv:2312.11805, 2023.

[11] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust Speech Recognition via Large-Scale Weak Supervision," OpenAI Technical Report, 2022.

[12] S. Long, X. He, and C. Yao, "Scene Text Detection and Recognition: The Deep Learning Era," arXiv preprint arXiv:1811.04256, 2018.

[13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.

[14] "Leveraging Natural Language Processing for Automated Classification of IT Support Tickets in Incident Management Systems," ResearchGate preprint, 2025.

[15] Ministry of Electronics and Information Technology (MeitY) / Centre for Railway Information Systems (CRIS), "Multilingual RailMadad Grievance Portal with Bhashini Integration," IndiaAI, Government of India, 2025.