



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: <https://doi.org/10.22214/ijraset.2026.82085>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Towards Intelligent Government Grievance Redressal Systems: A Survey of Computational Methods and System Architectures

Navanidhi D J¹, Aadit Tomar², Varsha Virupaksha Gujjala³, Jayanth Gowda⁴, Basavaraj Jakkali⁵
Department of Computer Science & Engineering B.M.S College of Engineering, Bengaluru, Karnataka, India

Abstract: Citizen grievance redressal has long suffered from systemic inefficiencies rooted in manual triage, opaque routing, and delayed resolution cycles. Citizens routinely encounter portals where complaints go unacknowledged, routing decisions remain invisible, and resolution timelines stretch indefinitely, eroding public trust in governance itself. As governments increasingly deal with large volumes of structured and unstructured complaint data, the demand for intelligent automated processing has become hard to ignore.

This survey examines 20 research papers from 2017 to 2026, tracing the full arc of AI-driven grievance management, from early Naive Bayes and SVM classifiers to transformer-based architectures, zero-shot LLM pipelines, Graph Neural Networks (GNNs) and blockchain-integrated multimodal frameworks. The analysis reveals a clear technological progression: statistical models established baseline classification capability but ran into scalability and generalization limits; deep learning methods like CNNs and LSTMs improved accuracy at the cost of interpretability; transformer models such as BERT and RoBERTa brought complaint understanding far closer to human-level performance; and the latest graph-based and multimodal systems have further strengthened civic complaint analysis by capturing relational, spatial, and contextual signals that text-only approaches tend to miss.

Key findings include the persistent challenge of multilingual support, the absence of standardized public benchmarks, the ongoing tension between interpretability and performance, and the largely untapped potential of LLMs in zero-shot civic NLP. This survey identifies critical research gaps and outlines directions for building grievance intelligence systems that are scalable, transparent, and grounded in fairness and inclusivity.

Index Terms—Grievance Redressal, Natural Language Processing, Graph Neural Networks, Large Language Models, Complaint Classification, E-Governance

I. INTRODUCTION

Citizen grievance systems sit at the core of how the public interacts with governing institutions. Whether embedded in municipal portals, national e-governance platforms like CPGRAMS, or institutional complaint cells, these systems share a common bottleneck: the manual effort required to read, categorize, route, and track thousands of heterogeneous complaints. The fallout — delayed resolutions, misrouted complaints, missed escalation signals, and weakened public trust — has made the case for AI-augmented grievance redressal increasingly difficult to ignore.

Over the past decade, the AI methods applied to this domain have shifted considerably. What started with key-word matching and Naïve Bayes classifiers has grown to include BERT fine-tuning, zero-shot LLM prompting, Graph Attention Networks that model entire complaint ecosystems, and blockchain-integrated multimodal platforms capable of processing text, images, and video evidence together. Each technological wave has brought real gains alongside new trade-offs, yet the field still lacks a consolidated synthesis that maps this progression in any analytical depth.

This survey addresses that gap by examining 20 papers published between 2017 and 2026, spanning civic complaint portals, educational grievance systems, social media complaint detection, legal redressal chatbots, and threat-language analysis tools — domains where grievance intelligence directly bears on the effectiveness of government systems. Our contributions are:

- A structured analysis of methodological evolution in AI-driven complaint processing, ranging from classical machine learning to graph-based and multimodal systems.
- A taxonomy of approaches organized by architectural and functional distinctions.
- A comparative analysis of techniques, datasets, and performance metrics across the reviewed literature.

- Identification of concrete, literature-grounded research gaps and actionable directions for future work.

The remainder of this paper is organized as follows. Section II describes the review methodology; Section III presents the core literature review; Section IV provides the taxonomy; Section V offers comparative analysis; Sections VI and VII address research gaps and future directions; Section VIII concludes.

II. REVIEW METHODOLOGY

This survey employs a systematic literature review framework. Primary sources were aggregated from IEEE Xplore, Google Scholar, Scopus, the ACM Digital Library, arXiv, and Nature Scientific Reports using targeted search strings such as grievance redressal, complaint classification, civic NLP, e-governance AI, and escalation detection. Studies were included if they (a) proposed automated systems for complaint routing or classification in civic and social contexts; (b) utilized machine learning methodologies; and (c) were published between 2017 and 2026. We excluded papers focusing on general NLP without a specific grievance application or those lacking technical depth. The final corpus of 20 papers was curated based on relevance, methodological rigor, and domain diversity.

III. LITERATURE REVIEW

A. Early Machine Learning and Rule-Based Grievance Systems

The early AI-driven grievance systems were built using classical machine learning models applied to preprocessed complaint text. Models like Naive Bayes, Logistic Regression, and Support Vector Machines showed the automatic routing of complaints. However, they did struggle with two major issues such as they depended heavily on small, carefully prepared datasets and they didn't generalize well to real-world scenarios.

For instance, one system trained on around 170 records with the accuracy of 92%. While that sounds impressive, such a small dataset makes the result unreliable and highly sensitive to how the data was split [11]. During this phase, most systems relied on manually labeled and domain-specific data, which worked fine in controlled environments but often failed in practical use.

At the same time, early mobile grievance applications such as the first versions of All India grievance apps focused more on digitizing workflows rather than adding intelligence. Complaints were stored using Firebase real-time databases and submitted through simple category-based forms without any automated classification. Although these systems lacked AI capabilities to understand the urgency of the complaint, they played an important role in building an initial digital infrastructure for citizen services. [16].

Some systems have also tackled the problem of multi-department routing, where a single complaint needs to reach the right authority without the citizen having to figure out which department handles what [3].

B. Deep Learning Era: CNNs, LSTMs, and the Race for Accuracy

As deep learning approaches have become more prevalent, CNN and LSTM hybrid architectures and GloVe/Word2Vec embeddings are gradually replacing traditional bag-of-words representations. Word2Vec-based approaches paired with recurrent models also picked up during this period, showing that oral and text complaints could be sorted into categories without much manual feature engineering [5]. When applied to the Pune Municipal Corporation complaint dataset, a CNN-LSTM model achieved an accuracy of 81.61%, which is relatively modest for a four-class problem. This can be partly explained by the fact that only 2,983 out of the available 100,000 records were used for training [15]. This example persists as a problem in grievance AI research: the mismatch between dataset scale and model capacity. In contrast, experiments on government railway complaint data used a much larger and more balanced dataset, with 186,940 complaints across 14 categories. Here, LSTM and CNN models achieved over 95% accuracy [12]. Although these results appear strong, the dataset does not reflect real-world class imbalance, where common issues such as main tenance dominate while sensitive complaints like bribery are rare. As a result, actual deployment performance may be lower than reported. Overall, this phase demonstrated that deep learning can outperform classical methods when sufficient data is available, but also showed that data quality and representativeness are often more important than model complexity. Social media complaint detection introduced an additional challenge: identifying complaints within informal and noisy text. A nearly Twitter-based classifier trained on just 100 tweets [8] reported high accuracy (95.8%) within the sample, but did not establish whether the model could generalize beyond this limited dataset. This highlights a broader concern in the literature—the gap between controlled experimental results and real-world applicability—which later approaches attempt to address more systematically.

C. *TheTransformerInflectionPoint*

The introduction of the Transformer architecture [18] marked a turning point in natural language processing, and complaint intelligence systems were no exception. Pre-trained contextual representations from models like BERT and its variants proved to transfer well to downstream complaint classification tasks. A systematic comparison of BERT, RoBERTa, ALBERT, and XLNet on multi-domain social media complaints [13] reported an 8% improvement in macro F1 over the previous LR-BOW baseline, with BERT reaching 88.0% accuracy and a macro F1 of 87.0. Notably, the study also demonstrated cross-domain transfer capability — something earlier statistical and LSTM-based systems consistently struggled with.

The rise of transformer models coincided with a broader shift toward more integrated complaint routing architectures. Systems began combining NLP with rule-based routing engines, SLA-driven escalation mechanisms, and sentiment analysis to support end-to-end complaint lifecycle management. A university grievance system [20] that brings together NLP-based categorization, Random Forest-based resolution time prediction, and SLA-aware escalation reported 89% categorization accuracy alongside a 40% reduction in average resolution time — a clear indication that pairing classification with workflow automation can produce tangible operational gains. Educational institutions have seen similar efforts, with automated categorization and department-level routing being applied to student grievance management to cut down on manual handling [4]. Similar hybrid approaches combining NLP with standard ML classifiers have shown that complaint routing across departments is achievable without deep learning, provided the pipeline is well designed [7].

The Retrieval-Augmented Generation (RAG) paradigm pushed transformer capabilities further into domain-specific legal applications. A system built for Indian consumer rights query handling [9] assembled a curated legal knowledge base and introduced the HAB (Helpfulness, Accuracy, Brevity) evaluation framework as a domain-specific metric set for legal conversational AI. Despite strong helpfulness scores, the system managed only 3.61/5 on accuracy — a telling limitation. In high-stakes legal contexts, where imprecise responses can carry real consequences, large language models still fall noticeably short.

D. *TowardStructuralIntelligence:Graph-BasedandMulti-Model system*

The latest developments point to a clear shift in how grievance systems are being designed. Rather than treating complaints as isolated text, they are increasingly understood as part of larger relational, spatial, and temporal networks. Two GNN-based systems — one focused on civic escalation detection [1] and another on grievance network analysis — reflect this shift well. They show that escalation risk [1] is better predicted by looking at a complaint's position within a network — geographic proximity to other complaints, departmental overlap, user behavior patterns — than by relying on text alone. A Graph Convolutional Network with attention mechanisms achieved 91.2–92.2% accuracy for escalation prediction, roughly 9 percentage points ahead of LSTM-based models.

The zero-shot LLM framework [2] represents a parallel line of development. Instead of training task-specific classifiers, it prompts a pre-trained multimodal LLM to handle grievance classification, urgency scoring, and abuse detection without any domain-specific fine-tuning. Since it processes both text and image input through a single prompt interface, it removes the dependency on labeled datasets — a meaningful advantage in real-world deployments where annotated data is hard to come by.

Blockchain integration entered the picture with Citizen Connect [19], which brings together NLP classification, computer vision-based evidence analysis, cryptographic complaint logging, and geospatial heatmap visualization within a Flutter mobile platform. Its evaluation is qualitative only, which limits direct performance comparison, but the architectural contribution stands on its own — blockchain-backed complaint immutability with optional citizen anonymity tackles accountability gaps that purely ML-based systems are not built to handle. Efforts toward fully autonomous redressal for government services have also emerged, attempting to handle the entire complaint lifecycle from intake to resolution with minimal human involvement [6].

A more recent direction involves unsupervised topic modeling on large-scale government portal data [10]. Rather than slotting complaints into predefined categories, BERTopic and NMF applied to 91,866 SewaSetu grievance records surface latent thematic structures that administrators may not have anticipated. This is particularly useful for policy discovery and proactive governance, where the goal is understanding complaint landscapes broadly rather than routing individual tickets.

IV. TAXONOMY OF APPROACHES

Based on the reviewed literature, we derive a five-tier taxonomy that captures the architectural and functional evolution of grievance AI systems, as summarized in Table I.

TABLE I
TAXONOMY OF AI APPROACHES IN GRIEVANCE REDRESSAL

Tier	Category	Key Characteristics
T1	Rule-Based & Classical ML	Keyword routing, Naïve Bayes, SVM, Logistic Regression; low scalability; high interpretability
T2	Deep Neural Networks	CNN, LSTM, GloVe/Word2Vec; improved accuracy; limited cross-domain transfer
T3	Transformer-Based	BERT, RoBERTa, RAG; contextual embeddings; strong transfer learning; moderate interpretability
T4	Structural & Graph-Based	GCN, GAT, models relational complaint networks; captures escalation patterns; computationally intensive
T5	Multimodal & LLM-Integrated	Zero-shot LLMs, Vision-Language Models, Blockchain; multimodal input; minimal labeled data; increased architectural complexity

Each tier reflects not just a performance jump but a conceptual expansion in what aspects of a complaint are actually being modeled. T1 systems work with lexical features of complaint text; T2 brings in sequential context; T3 taps into transfer learning from pre-trained corpora; T4 factors in inter-complaint relational structure; and T5 adds modality richness alongside institutional accountability mechanisms. This progression is not strictly linear — recent systems often blend multiple tiers, such as GNNs built on BERT-based node embeddings — but the taxonomy helps clarify what each generation of approaches is actually contributing.

V. COMPARATIVE ANALYSIS

Table II summarizes representative systems across the re-viewed literature along key comparative dimensions. For space efficiency, we present one to two systems per methodological tier.

A. Discussion of Trade-offs

Several cross-cutting trade-offs emerge from this analysis. First, there is a persistent tension between *accuracy and interpretability*: GNN and transformer-based models achieve the highest reported accuracies but generate predictions that are

TABLE II
COMPARATIVE ANALYSIS OF REPRESENTATIVE GRIEVANCE AI SYSTEMS

System/Reference	Method Type	Core Technique	Dataset Type	Best Metric	Strengths	Limitations
NB Classifier [11]	Classical ML	Naïve Bayes	Organic (170 records)	Acc: 92%	Simple; interpretable	Tiny dataset; no generalization
CNN-LSTM [15]	Deep Learning	CNN+LSTM, GloVe	Municipal (2,983 of 100K)	Acc: 81.61%	Real govt. data	Underutilized dataset; modest accuracy
Railway LSTM [12]	Deep Learning	LSTM, CNN	Govt. (186,940 balanced)	Acc: 95%+	Large real dataset	Artificial class balance; low interpretability
BERT/Transformer	Transformer	BERT,	Twitter (1,971)	F1: 87.0	Cross-domain	Small dataset; social

[13]		RoBERTa, M-BERT	tweets)		transfer	mediabias
RAGChatbot[9]	Transformer +RAG	RAG,LLM, BERTScore	LegalQA(4 novelsets)	HABAcc: 3.61/5	Novevaluation framework	Legalprecisiongaps; English-only
GNNEscalation[1]	Graph-Based	GCN+Attention, BERT	Civic (10,000+)	Acc:91.3%; F1:0.88	Capturesrelational context	Computationallyin- tensive;black-box
Zero-ShotLLM[2]	LLM (Zero-Shot)	MultimodalLLM, prompting	Civic(text+ images)	Qualitative	Nolabeleddata required	Limitedevaluation; hallucinationrisk
CampusCare[20]	Transformer +Rules	NLP,Random Forest,SLA	Institutional (2,500+)	Acc:89%; -40%ART	SLAescalation; sentiment	Slangsensitivity;no multimedia
CitizenConnect[19]	Multimodal+ Blockchain	NLP,CV, Blockchain	Prototype(9 complaints)	Qualitative only	Immutability; anonymity	Noquantitativeeval- uation
BERTopic/ [10]	NMF Unsupervised	LDA,NMF, BERTopic	Govt.(91,866 records)	Cv coherence: BERTopic best	Largestreal dataset; privacy-aware	Nocomplaintrout- ingoutput
Grievance Dictionary[14]	Lexicon- Based	Psycholinguistic lexicon,NLP	Multi-corpus (novels,Reddit, manifestos)	Precision: high	Threat-language analysis	Narrowscope;no routing

difficult for administrators and policy-makers to meaningfully interpret. Classical ML systems, while less accurate, offer considerably more transparency—a property that still matters in governance contexts where algorithmic decisions need to be auditable. Second, *dataset scale and realism* heavily shape reported performance, often overstating confidence in otherwise limited systems. Several studies claiming 92–96% accuracy are built on datasets with fewer than 200 samples or artificially balanced class distributions — conditions far removed from real-world deployment. Meanwhile, the largest dataset in this survey (91,866 grievances from the SewaSetu portal) is used for unsupervised topic discovery, where supervised accuracy cannot be reported at all. This points to a fundamental gap: there is no large, publicly available, labeled benchmark for civic complaint classification. Third, *modality richness versus computational overhead* marks the current frontier. Zero-shot LLMs and multimodal architectures reduce or eliminate the need for labeled data, but they bring inference latency, hallucination risk, and infrastructure demands that may be impractical for low-resource municipal environments. Finally, *evaluation methodology* remains inconsistent across the literature. Some systems are assessed using standard metrics like accuracy and F1, while others lean on operational indicators such as resolution time reduction or user satisfaction scores. Recent LLM-based systems introduce their own evaluation frameworks—HAB metrics being one example making direct comparison harder still. This inconsistency weakens cross-study comparability and complicates evidence-based system selection for practitioners.

VI. RESEARCH GAPS AND CHALLENGES

Lack of Transparency and Trust Mechanisms. Despite advances in automated grievance processing, very few systems incorporate meaningful mechanisms for verifying complaint authenticity or ensuring transparency in resolution workflows. Most reviewed platforms rely on internal administrative processes, offering citizens no verifiable proof of action, no audit trails, and no tamper-resistant records. This erodes public trust — especially when complaints are delayed, misrouted, or closed without clear justification. Blockchain-integrated systems [19] make a start on immutability and optional anonymity, but the absence of standardized verification and accountability frameworks remains a real barrier to trustworthy large-scale deployment.

Absence of Standardized Benchmarks. The most structurally damaging gap in the reviewed literature is the lack of a publicly available, large-scale, labeled benchmark for civic grievance classification. Studies building organic datasets of 100–170 records [11], or training on just 2,983 of 100,000 available records [15], produce accuracy figures that cannot be meaningfully compared or reproduced. Without a shared benchmark, the field has no reliable way to measure actual progress. Duplicate complaint detection and smart escalation remain largely separate from classification pipelines, yet both are necessary in practice—without them, systems get flooded with repeats and miss genuinely urgent cases [17].

Multilingual Limitations. Every system in this survey is English-only or English-dominant—despite target deployment contexts like India, civic municipalities, and multi-district portals being inherently multilingual. Systems described as having “rural applicability” consistently push multilingual support to future work. This is not a minor gap: code-switching in informal complaint text—mixing Hindi, Kannada, or Telugu with English—degrades even transformer models that individually perform well on monolingual inputs. **Interpretability Deficit.** GNN-based escalation models deliver the strongest performance on relational complaint tasks but offer no human-interpretable explanation for their risk scores. In governance contexts where AI-driven decisions must be defensible to citizens and oversight bodies, black-box outputs are institutionally unworkable. XAI techniques such as GNN Explainer or attention visualization have not been applied to grievance escalation in any reviewed paper.

Real-Time Deployment Constraints. Transformers and GNNs require inference infrastructure that sits well beyond the resource constraints of district-level governments or small institutions. Edge deployment optimization, model quantization, and distillation for grievance classifiers remain entirely unexplored in the reviewed literature—despite being practical prerequisites for deployment at any meaningful scale.

Dataset Bias and Class Imbalance. Several systems trained on artificially balanced datasets [12] or small curated corpora [11] are prone to performance degradation under real complaint distributions, where common categories vastly outnumber rare ones—maintenance complaints versus bribery or emergency reports, for instance. No paper in this corpus treats class imbalance as an explicit design constraint.

Evaluation Framework Fragmentation. The absence of a shared evaluation framework—one that spans classification metrics, operational impact, and user trust—makes evidence-based comparison between systems practically impossible. The HAB framework introduced by [9] is a step forward for conversational systems, but nothing equivalent exists for routing, escalation, or end-to-end redressal platforms.

VII. FUTURE WORK

The gaps identified above point to several concrete and practical research directions. In particular, the lack of transparency and verifiable trust mechanisms in existing systems highlights the need for grievance platforms that not only make accurate predictions but also provide auditable workflows, clear resolution trails, and mechanisms for complaint verification. Addressing this alongside technical improvements will be critical for real-world adoption.

One of the most impactful near-term contributions would be the creation and public release of a multilingual, multi-domain, large-scale civic grievance benchmark—analogue to GLUE for general NLP—that enables reproducible cross-system comparison. Such a dataset should span multiple Indian languages, include metadata (e.g., location, timestamp, department), and capture both resolved and escalated complaint labels.

On the modeling side, combining the relational capabilities of GNNs with the contextual understanding of transformer encoders in an end-to-end architecture represents a promising but underexplored direction. Hybrid Graph-BERT approaches, which have shown success in domains such as citation networks and fraud detection, could be effectively adapted for complaint escalation prediction. Improving explainability in graph-based grievance systems remains both a technical and institutional priority. Techniques such as GNN Explainer or Integrated Gradients can be extended to generate complaint-level justifications, ideally expressed in natural language, enabling non-technical administrators to better understand and trust model outputs.

For deployment in low-resource settings, lightweight transformer models such as DistilBERT or MobileBERT, fine-tuned on civic complaint data and optimized through quantization, offer a practical path toward achieving high accuracy without reliance on centralized infrastructure.

Finally, federated learning across multiple municipal bodies presents an opportunity to improve model performance collaboratively while preserving data privacy. Such an approach allows systems to learn from diverse complaint distributions without requiring centralized storage of sensitive citizen data, addressing both privacy concerns and data scarcity challenges.

VIII. CONCLUSION

This survey traces the arc of AI-driven grievance redressal from early rule-based and statistical classifiers through deep learning and transformer models, to the current frontier of graph-based, zero-shot LLM, and blockchain-integrated multimodal architectures. Each stage has pushed both the performance ceiling and the conceptual scope of what automation can do: early systems focused on categorizing complaints, while modern ones can predict escalation risk from relational complaint networks, process image evidence, and handle tasks without any labeled data through zero-shot prompting.

Three themes run consistently through this progression. First, data quality and scale tend to matter more than model complexity — systems that pair strong models with large, realistic datasets hold up reliably, while those trained on small curated corpora often overfit. Second, the tension between performance and interpretability does not go away: the most accurate models, GNNs and transformers in particular, offer limited transparency for institutional decision-making. Third, multilingual support stays largely unaddressed across the literature, despite being a basic requirement for equitable and inclusive civic AI. The field is at something of an inflection point. Methodological capabilities have matured considerably, but the supporting empirical infrastructure — standardized benchmarks, unified evaluation frameworks, real-world deployment studies — has not kept pace. Closing that gap is both the central challenge and the most consequential direction for future work in grievance intelligence.

REFERENCES

- [1] A. Saranya, V. Pushpalatha, and C. E. Thinega, "Graph neural network based escalation detection in grievance network," in Proc. 2nd Int. Conf. Artificial Intelligence and Knowledge Discovery in Concurrent Engineering (ICECONF), 2025, doi: 10.1109/ICE-CONF65644.2025.11379675.
- [2] N. R. Sivakumar, "A zero-shot LLM framework for multimodal grievance classification, urgency scoring, and abuse detection in civic feedback systems," Sci. Rep., vol. 16, no. 1, p. 2299, Dec. 2025, doi:10.1038/s41598-025-32079-7.
- [3] A. Tiwari, P. Sharma, A. Gupta, Anmol, and A. Sharma, "AI-based solution to enable ease of grievance lodging and tracking for citizens across multiple departments," in Proc. Int. Conf. Innovations and Emerging Technologies in AI & Communication Systems (IETACS), 2025, doi:10.1109/IETACS68750.2025.11385682.
- [4] A. S. Manek et al., "Grievance redressal system for educational institution," Int. J. Elect., Electron. Comput., Special Issue, Jan.–Feb. 2026.
- [5] M. Rathore, D. Gupta, and D. Bhandari, "Complaint classification using Word2Vec model," Int. J. Eng. Technol., vol. 7, no. 4.5, pp. 402–404, 2018, doi: 10.14419/ijet.v7i4.5.20192.
- [6] S. Suganya et al., "Autonomous grievance redressal system for government services," in Proc. Int. Conf. Data Science and Business Systems (ICDSBS), 2025, doi: 10.1109/ICDSBS63635.2025.11031909.
- [7] A. Saha et al., "Automated complaint classification and routing using NLP and machine learning," Int. J. Innovative Res. Technol., vol. 11, no. 6, Nov. 2024.
- [8] A. P. Singh et al., "NLP-based grievance redressal system," Int. J. Comput. Appl., vol. 184, no. 12, May 2022.
- [9] S. Ganatra et al., "Grahak-Nyay: Consumer grievance redressal through large language models," arXiv preprint arXiv:2507.04854, 2025.
- [10] R. Deka and R. K. Srivastava, "Improving public service grievance analysis: A comparative study of topic modelling techniques," J. Inf. Technol., Cybersecurity, Artif. Intell., vol. 2, no. 3, pp. 41–53, 2025, doi: 10.70715/jitcai.2025.v2.i3.031.
- [11] R. S. Rao et al., "Implementing NLP to categorize grievances received via a voice input mechanism," Int. J. Eng. Res. Technol., vol. 12, no. 04, Apr. 2023.
- [12] M. Gupta et al., "Multi-class railway complaints categorization using neural networks: RailNeural," J. Rail Transp. Planning Manage., vol. 20, 2021, doi: 10.1016/j.jrtpm.2021.100265.
- [13] M. Jin and N. Aletras, "Complaint identification in social media with transformer networks," arXiv preprint arXiv:2010.10910, 2020.
- [14] I. van der Vegt et al., "The grievance dictionary: Understanding threatening language use," Preprint, 2020.
- [15] S. Sawant et al., "Design and implementation of grievance filing application using automatic text classification," Int. J. Innovative Technol. Exploring Eng., vol. 9, no. 7, May 2020, doi: 10.35940/iji-tee.G5852.059720.
- [16] V. Patel, D. Kapadia, D. Ghevariya, and S. Pappu, "All India grievance redressal app," J. Inf. Technol. Digit. World, vol. 2, no. 2, pp. 91–99, 2020, doi: 10.36548/jitdw.2020.2.002.
- [17] S. K. Reddy et al., "Complaint and grievance redressal system with duplicate detection and smart escalation," Int. J. Sci., Eng. Technol., vol. 14, no. 2, 2026.
- [18] A. Vaswani et al., "Attention is all you need," in Proc. 31st Conf. Neural Information Processing Systems (NeurIPS), 2017.
- [19] K. G. Sankar et al., "CitizenConnect: An AI and blockchain-powered mobile grievance redressal system," Int. Res. J. Adv. Eng. Hub, vol. 4, no. 2, pp. 783–791, Feb. 2026, doi: 10.47392/IRJAEH.2026.0112.
- [20] S. N. Sivanandam et al., "CampusCare: An intelligent grievance redressal system using AI and web technologies," Int. J. Eng. Res. Technol., Mar. 2026.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)