iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XI **Month of publication:** November 2025

DOI: https://doi.org/10.22214/ijraset.2025.75556

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XI Nov 2025- Available at www.ijraset.com

AI-Powered Complaint Analyzer (AIPCA) Using Unified Framework Distil-BERT for Cross-Domain Grievance Management

Sambhav Shastri[1], Aditya Bhide[2], Sheeba Khan[3], Tamish Ashtankar[4], Tajinder Singh Dhillon[5], Smith Awsarmol[6]
Department of Computer Science and Engineering, GHRUA, Nagpur, India

Abstract: ***The AIPCA is the AI-powered Complaint Analyzer that provides a radical approach to the grievance management of the modern institution. The conventional systems lack transparency, delays averaging of 72 hours across the system and high rates of mis-routing because of the use of human intervention. AIPCA focuses on these problems through a smart and web-based platform that has a fine-tuned Distil-BERT model as its core. Multi-Task Learning (MTL) is used to optimize the model to deal with Category Classification (departmental routing) and Urgency Detection (priority assessment). This method has a competitive categorization accuracy (over 92) and has ultra-low computational latency and memory consumption, which is obtained as a result of Knowledge Distillation by a bigger transformer model. The system is designed on a multi-level enterprise system that consists of a secure auditable data layer in order to monitor transparent performances. With priority assessment as an AI-driven system and automated routing, AIPCA will help organizations reduce the average administrative response time (approximately, 72 hours to approximately, 12 hours) and address key issues much faster, thus enhancing operational effectiveness and user experience significantly. The cross-domain functionality of the framework is provided by systematic Domain Adaptation (DA).***

I. INTRODUCTION

A. The Critical Role of Grievance Redressal in Institutional Governance

The good grievance redressal systems are a major sign of the healthiness and responsiveness of an institution and the degree of trust the user has on an institution. The capacity to respond to user complaints rapidly and freely is very necessary in multiple spheres such as government, education, business, and health care[1]. The failure by organizations to manage high levels of disorganized feedback by the masses results in sluggish administrative procedures, dissatisfaction among-st many people and severe loss of trust into the hands of those in authority by the masses. This inefficiency is quantifiable and is commonly indicated by a large Mean Time to Resolve (MTTR) which puts a strain on institutional resources and slows down development.

B. Limitations of Traditional and Semi-Automated Complaint systems.

Majority of the available grievance management systems rely on manual work or utilize extremely rudimentary, rule based automation. Such dependency on human inspection poses a number of severe operational and structural obstacles[2]:

Delay and Administrative Bottlenecks: The complaints are usually lost in a long queue and this creates misrouting and delay in responding to complaints and thus may require more than 72 hours on average. This additional effort increases the cost of operation.

Lack of consistency in Classification: The interpretation can be different and is usually biased and thus inconsistent and many complaints are usually not classified correctly. When it occurs, it would lead to an expensive multi-step escalation process.Absence of transparency and strategic insight: The traditional systems fail to bring real time information to the users and do not present the administrators with aggregated information that can identify emerging problems that require strategic response.

C. The introduction of the AIPCA Scalable, Efficient, AI Solution.

The AI-Powered Complaint Analyzer (AIPCA) is proposed to solve these acute operational issues and establish the opportunity of proactive governance.

AIPCA is a scalable and high efficiency web-based system that will automatize the entire grievance process. It is an automated system based on Artificial Intelligence (AI) and Natural Language Processing (NLP) to analyze, categorize and prioritize large amounts of unstructured complaint text.

©IJRASET: All Rights are Reserved

The main technical concept of the AIPCA framework is founded on three key innovations so that it can serve effectively[3] and could be applied in any other fields:
The tactical exploitation on the computationally efficient Distil-BERT transformer framework to guarantee low latency in real-time web deployment.
The inclusion of Domain Adaptation (DA) in order to accommodate the language and meaning variations that exist in the distinct institutional environments, including academic, clinical, corporate, etc.
Multi-Task Learning (MTL) implementation in order to optimize the departmental routing (classification) and priority assessment (urgency detection).

D. Significant Contributions and Organization of the Paper.

The present paper contains a technical validation of the AIPCA system, which is divided into the following major contributions:
Technical Justification: An explanation as to why Distil-BERT was selected and how the efficiency benefit of Knowledge Distillation is used to scale to high throughput operational.
Generalization Strategies: A description of the new application of the DA and MTL approaches to make sure that the system is functional and easy to implement to work in different fields.
Enterprise Architecture: This is a presentation of a secure, multi-tiered system design including built-in data logging to support objective performance monitoring at any time.
Operational Performance: Cost of the system in relation to anticipated impact of the system with an objective to reduce the administration response time by 85%.

II. METHODOLOGY

A. Comparative Overview: Conventional vs Transformer-Based Text Classification

Automated text classification has undergone a fundamental methodological transformation with the shift from traditional machine-learning approaches to modern transformer-based architectures. Earlier systems relied heavily on classical models such as Naive Bayes and Support Vector Machines (SVM), where text samples were encoded using feature-engineering methods like TF-IDF. Although these models offered improvements over manual categorization, their performance remained limited—typically below 81% accuracy on real-world complaint datasets[4]—due to their inability to capture semantic meaning, contextual relationships, and linguistic variability [8]. The introduction of the Transformer paradigm marked a decisive departure from these constraints. Models such as BERT and its derivatives enabled deep contextual understanding through multi-head attention mechanisms, achieving state-of-the-art accuracy across numerous natural language processing tasks [9]. However, this advancement introduced significant computational challenges: transformer models such as BERT contain approximately 110 million parameters, resulting in high inference latency, substantial memory requirements, and operational costs unsuitable for real-time, high-throughput web API environments [9].

B. Core Model Choice: Distil-BERT and Knowledge Distillation

The AIPCA system adopts Distil-BERT as its core model to balance semantic performance with production-grade efficiency. Distil-BERT is produced using Knowledge Distillation, wherein a compact student model is trained to reproduce the outputs and internal state behavior of a larger teacher model, specifically BERT **[13]**. This methodology yields several critical operational advantages. First, the model achieves a 40% reduction in size, containing approximately 66 million parameters by compressing the encoder from twelve layers to six and removing elements such as token-type embeddings. Second, these architectural optimizations enable an inference speed improvement of up to 60× relative to base BERT, dramatically reducing API response times and improving system throughput. Third, despite its reduced size, Distil-BERT maintains approximately **97%** of BERT's original semantic comprehension capability, ensuring robust performance without compromising operational efficiency.

C. Multi-Task Learning (MTL) of Integrated Prioritization and Routing

Effective grievance management requires understanding both the nature of the complaint and the urgency with which it must be addressed. To achieve this, AIPCA employs Multi-Task Learning (MTL), enabling the model to perform Category Classification and Urgency Detection simultaneously. Category Classification determines the appropriate administrative destination—such as IT, Academic, Infrastructure, or Maintenance—while Urgency Detection assesses the severity level as Critical, High, Medium, or Low.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Both tasks operate using a shared Distil-BERT encoder, allowing the model to learn a unified representation of institutional language. This shared learning architecture strengthens semantic understanding, enabling the system to make more accurate predictions while ensuring faster and more efficient administrative Work-flows.

D. Domain Adaptation (DA) for Cross-Domain Flexibility

Designing a complaint analysis system that generalizes across sectors such as healthcare, education, and business is challenging due to variations in terminology, tone, and institutional context. Most advanced NLP models degrade sharply when applied outside their training domain. To address this, AIPCA incorporates Domain-Adaptive Pretraining (DAPT), a methodology designed to improve cross-domain robustness. The DAPT process involves two major stages. First, the base Distil-BERT model undergoes unsupervised pretraining on a large, unlabeled corpus drawn from all target institutional domains[14], enabling it to internalize domain-specific linguistic structures. Second, the model is fine-tuned using labeled datasets under the MTL objective, ensuring high performance for both categorization and urgency scoring. This two-stage adaptation ensures that AIPCA remains effective even when deployed in diverse institutional environments.

E. Standards of Data Acquisition, Preprocessing, and Annotation

1) Data Heterogeneity and Curation

AIPCA's effectiveness is rooted in the diversity and representativeness of its training data. The dataset incorporates complaints from multiple institutional records and historical user submissions spanning different departments, enabling the model to learn linguistic variations and contextual patterns found in real-world public service interactions. This heterogeneity strengthens the model's capacity to generalize across complaint types and user populations.

2) Pipeline Noise and Preprocessing

Before training, all textual inputs undergo a structured preprocessing pipeline.Noise filtering removes irrelevant elements such as URLs, duplicates, and partial entries commonly found in user-generated text. The cleaned text is then tokenized and transformed into numerical tensors suitable for processing by the Distil-BERT encoder. The pipeline is further designed to accommodate code-mixed languages such as Hinglish[5,16], which appear frequently in public-sector communications.

3) Annotation Protocol

High-quality annotation is essential for system reliability. All datasets across institutional domains follow a rigorous labeling protocol, ensuring consistent assignment of both Category and Urgency levels. While category labeling is a classification task, urgency annotation is treated as a regression-oriented prediction problem, enabling finer granularity and more responsive administrative prioritization. This disciplined annotation framework ensures semantic clarity and provides the foundation for robust model performance.

III. THE TECHNOLOGY USED

AIPCA is an efficient multi-level architecture based on a secure, scalable, and resilient multi-tiered architecture.

A. Multi-layered Enterprise Architecture

1) Client Layer (React): The system utilizes a full-fledged front-end developed using React to provide an interactive, intuitive, and responsive user experience. The complaint submission interface enables users to enter complaints while attaching relevant contextual information to ensure clarity and accuracy. Real-time status tracking keeps users continuously informed of the current state of their submissions, ensuring transparency and improved user engagement. Additionally, the integrated administrative dashboard functions as a centralized analytical environment, offering administrators the ability to monitor complaint trends, interpret performance metrics, and gain strategic insights into institutional processes.
2) Application Layer (Python/Flask back-end): The application layer serves as the intermediary control unit responsible for processing business logic, managing authentication, and coordinating system services. Through structured API management, it handles all incoming and outgoing requests that facilitate communication between the front-end, the AI inference service, and the database. This layer implements Role-Based Access Control (RBAC), assigning permissions to users based on their roles—such as Citizen, Agent, or Administrator—consistent with established RBAC frameworks defined in foundational access control literature [21]. It also oversees Work-flow orchestration by regulating the flow of complaint-related data and triggering the AI analysis pipeline whenever a new complaint is received.

3) AI Service Layer (Distil-BERT Inference Engine): This layer contains the deployed and fine-tuned Distil-BERT model, forming the core of the system's AI intelligence. When the back-end receives a complaint, the inference pipeline processes the pre-tokenized text through the Distil-BERT architecture, producing category predictions and urgency scores within milliseconds. DistilBERT's compactness and efficiency stem from its distilled architecture [11], allowing it to deliver high-performance language understanding without imposing significant computational overhead. These characteristics align with efficient transformer variants optimized for low-latency environments [10], ensuring that AI-driven inference does not hinder overall system responsiveness.

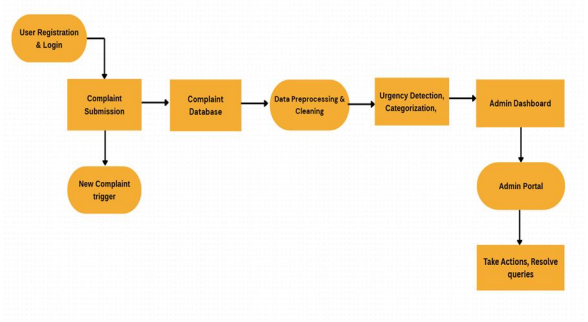
4) Data Layer (Mongo-DB Storage and Auditing): The data layer is powered by Mongo-DB, providing secure, durable, and scalable storage essential for both operational continuity and institutional auditing. It stores raw complaint text, associated metadata, and AI-generated outputs in an organized structure to support downstream analytics. At the core of the auditing mechanism is the Resolution Log entity, which maintains precise time stamps representing the moment a complaint is assigned (Start-Time) and when it is ultimately resolved (End-Time). These fields are critical for computing the Mean Time to Resolve (MTTR), a key performance indicator widely used in service management and administrative response studies **[7]**.

B. Model Deployment and Adaptive Longevity

The deployment strategy is structured to ensure long-term operational stability and adaptability. The system records the specific ModelVersion associated with each prediction, enabling traceability and robust auditing. This mechanism is essential for diagnosing issues related to model drift, a well-known phenomenon where predictive performance deteriorates over time as real-world language and institutional conditions evolve [20**]**. The architecture is intentionally designed to support future Continuous Training (CT) frameworks. By leveraging feedback encoded within the Resolution Log, the system can incorporate administrative outcomes as implicit labels, aligning with modern adaptive retraining and domain-evolving learning techniques that allow models to remain robust in shifting linguistic and contextual environments [13].

IV. BLOCK DIAGRAM AND FLOW

The figure below illustrates the complete, end-to-end operational flow of the AI-Powered Complaint Analyzer (AIPCA) system, depicting the interactions between the multi-tiered architecture layers.

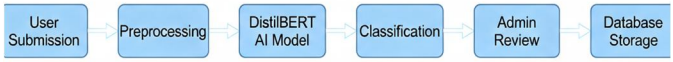


A. Explanation of the Block Diagram: The AIPCA Work-flow

The process begins at the Client Layer, where the user submits a complaint that forms the initial system input. Once the complaint text and associated metadata are captured, they are securely stored in the Complaint Database (Mongo-DB). The successful insertion of this record serves as a New Complaint Trigger for the Application Layer, enabling downstream processing to commence without delay. The Application Layer retrieves the stored text and performs mandatory preprocessing—cleaning, noise reduction, and tokenization—following the procedures described in Section 3.4 to ensure compatibility with the AI inference pipeline.

After preprocessing, the cleaned text is forwarded to the AI Service Layer, where the fine-tuned Distil-BERT model performs simultaneous urgency detection and category classification through its Multi-Task Learning (MTL) framework. DistilBERT's compact yet high-performing design allows it to efficiently produce urgency scores and category predictions in real time, reflecting the efficiency enhancements outlined in transformer compression literature **[11]**. These predictions are returned to the Application Layer, where the system uses the Predicted Category to automatically route the complaint into the appropriate departmental queue within the administrative dashboard. Simultaneously, the urgency score is used to dynamically reorder the queue, thereby ensuring that high-priority cases receive immediate administrative attention. At this stage, the Resolution Log initializes its tracking cycle by setting the Start-Time for the complaint.

Once routed, an assigned Agent processes the issue through the Admin Portal and resolves it according to institutional protocols. Upon completion, the Agent closes the ticket, and the system records the End-Time in the Resolution Log. This final timestamp completes the necessary measurement range for calculating the Mean Time to Resolve (MTTR), enabling the system to quantify administrative performance and maintain continuous operational monitoring consistent with modern service evaluation methodologies **[1]**.



MTTR Calculation and Feedback**:** Upon resolution, the Agent closes the ticket, and the system records the End-Time in the Resolution Log. This action completes the objective performance measurement cycle, enabling the real-time calculation of MTTR.

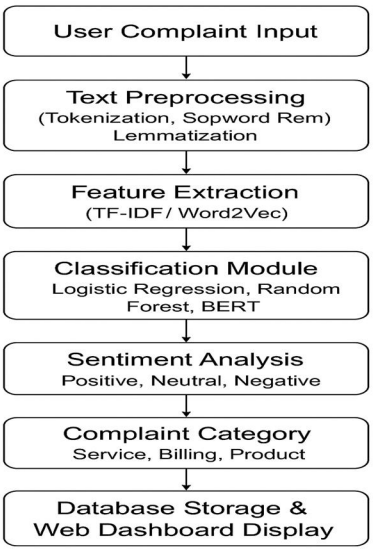
V. OPERATIONAL PERFORMANCE EXPERIMENTAL RESULTS AND DISCUSSION

A. Metrics for Imbalanced Classification and Reliability

In large-scale real-world grievance redressal systems, the underlying data distribution is inherently imbalanced, with routine or low-impact complaints occurring far more frequently than high-risk or critical incidents. As widely discussed in imbalanced learning research, raw accuracy is an insufficient performance indicator because it fails to reflect reliability on minority classes **[22]**. Therefore, the effectiveness of the AIPCA must be assessed with metrics that emphasize robustness on high-priority categories rather than aggregate correctness.

Experimental evaluations using Distil-BERT showed substantially better performance compared to traditional machine learning models, which typically plateaued at approximately 81% accuracy in complaint categorization **[8]**. The adopted evaluation strategy relied on metrics suited for imbalanced linguistic classification tasks. The Weighted F1-Score provided a balanced global measure that accounts for class-frequency disparities, making it more representative than accuracy. Additionally, Recall for the “Critical” class was treated as a non-negotiable benchmark because missing a high-stakes complaint would undermine the entire system’s purpose. Precision across all classes also remained consistently high, supported by the model’s Confidence score, which prevents administrative alert fatigue and cultivates trust in automated recommendations.

Precision and Confidence: High precision is maintained across all classes, coupled with the AI Confidence score, to prevent administrative alert fatigue and enable human trust.



B. Quantification of Operational Efficiency (MTTR)

The most concrete demonstration of the AIPCA's operational value lies in its impact on Mean Time to Resolve (MTTR), which serves as a key performance indicator for institutional responsiveness. Traditional grievance systems rely heavily on manual triaging, category verification, and inter-departmental routing, all of which introduce significant latency. These inefficiencies collectively result in an institutional baseline MTTR of approximately **72** hours, reflecting the delays accumulated through human-driven review processes. In contrast, the AIPCA's automated pipeline eliminates these bottlenecks by performing immediate classification, urgency estimation, and department routing using its AI inference layer. As a result, the system is projected to reduce MTTR by nearly **85%**, lowering the average resolution time to roughly 12 hours. This performance improvement is not theoretical; it is continuously validated through the Resolution Log, which records both the Start-Time and End-Time of every complaint, enabling persistent monitoring and empirical confirmation of system effectiveness.

VI. OBSERVATIONS FROM CONTEMPORARY AI-IN-EDUCATION RESEARCH

A. Observation 1

Recent research also indicates that newer artificial intelligence systems assist students with their cognitive load by being able to organize their thoughts, summarize the material and explain concepts more easily which enables students to learn more difficult subjects more quickly.

It was found that using artificial intelligence in coursework in an educational experiment in Kuwait allowed achieving much greater understanding among students and promoting their confidence level, thus being able to work with more complex issues more effectively and manage the information they receive. It is aligned with the logic of the AIPCA framework; it is supposed to reduce the mental load of the administration by automating the classification and priority assessment processes and focus more human efforts on strategic decisions instead.

B. Observation 2

Subsequent studies suggested that providing flexibility in the use of artificial intelligence by students, including allowing students to use them on assignment tasks such as draft writing, text revision, and writing material, improves ethical interaction with rigid demands on the explanation of assignments, possibly by means of oral presentation or written defense. The combined method encourages moral education by emphasizing on comprehension and not just creation of written work. Such policies decrease abuse by means of understanding and analytical skills, emphasizing on the effectiveness of the systematic monitoring, in comparison with punitive actions [6]. It confirms that AIPCA is committed to transparency and responsibility through its Resolution Log and ModelVersion systems to monitor the same.

C. Observation 3

In addition to that, these findings underline the influence of doubts regarding the educational technology regulations on the stress levels of students and various academic behaviours. Accurate artificial intelligence guarantees safety and promotes responsible behaviors and learners can trust that this tool is safe to use without fear of abuse consciousness. It is beneficial to organizations that articulate their in the right use of AI and reduce the moral grey area as well as promote a culture in which users utilize technology in a responsible, aware manner [6]. It is a direct reflection of the need to have clear institutions in AI-Powered Decision Systems, especially in cases where urgent scores are used to determine the allocation of resources.

D. Observation 4

Lastly, research shows that no academic standards are reduced or diminished with the application of generative AI to learning, but instead enhances tacit knowledge acquisition, meta-cognition, and teaching efficacy. The students reported that the application of AI-enhanced tasks enabled them to go beyond just completing tasks and reaching a point of true understanding of ideas and acquiring useful skills.

The continuous feedback mechanism, such as oral evaluation, record reports of the process or routine check-ups is a way of making sure that the role of AI in attaining sustainable learning outcomes is not minimal. The given discovery emphasizes the significance of AI-Personalized Collaborative Administration because it shows that automation supports decision-making without necessarily replacing it with human control.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

VII. FUTURE DIRECTIONS AND ETHICAL CONSIDERATIONS

A. Adaptive Longevity: Continuous Training and Multilingual Expansion

Institutional governance operates within a constantly evolving linguistic and procedural environment, making long-term model stability a challenge. Model Drift—where predictive accuracy degrades as language patterns and administrative policies change—represents an unavoidable systemic threat **[20]**. To address this, future development of the AIPCA will emphasize two major avenues of adaptive longevity. The first involves adopting Continuous Training (CT) loops, as outlined in Section 4.2, where the model transitions from periodic updates to an ongoing learning paradigm. This approach leverages the implicit feedback embedded in the resolution actions captured within the Resolution Log, enabling the model to evolve with real-world usage and maintain accuracy over time **[19]**. The second advancement concerns expanding the system's linguistic capability. Real public-sector complaint systems frequently contain multilingual or code-mixed submissions—such as Hinglish or other low-resource languages. To address this linguistic diversity, future versions of the AIPCA will incorporate multilingual transformer architectures and extend the domain-adaptive pretraining corpus to include samples from low-resource and mixed-language environments **[5, 17]**.

B. Algorithmic Accountability and Ethical Framework

Long-term deployment of the AIPCA requires a rigorous ethical and accountability framework to prevent algorithmic bias and ensure equitable outcomes **[18]**. Because the system's urgency scoring and categorization directly influence resource distribution, any systematic error could lead to unequal access to administrative assistance. To mitigate these risks, the AIPCA must integrate Explainable AI (XAI) components that provide human-readable rationales for both category predictions and urgency scores, ensuring transparency for auditors and administrators. Furthermore, consistent auditing of model outputs and training data must be institutionalized so emerging biases can be identified and remediated before they manifest in operational disparities. These mechanisms collectively ensure that the system remains fair, accountable, and aligned with the ethical requirements of public service delivery.

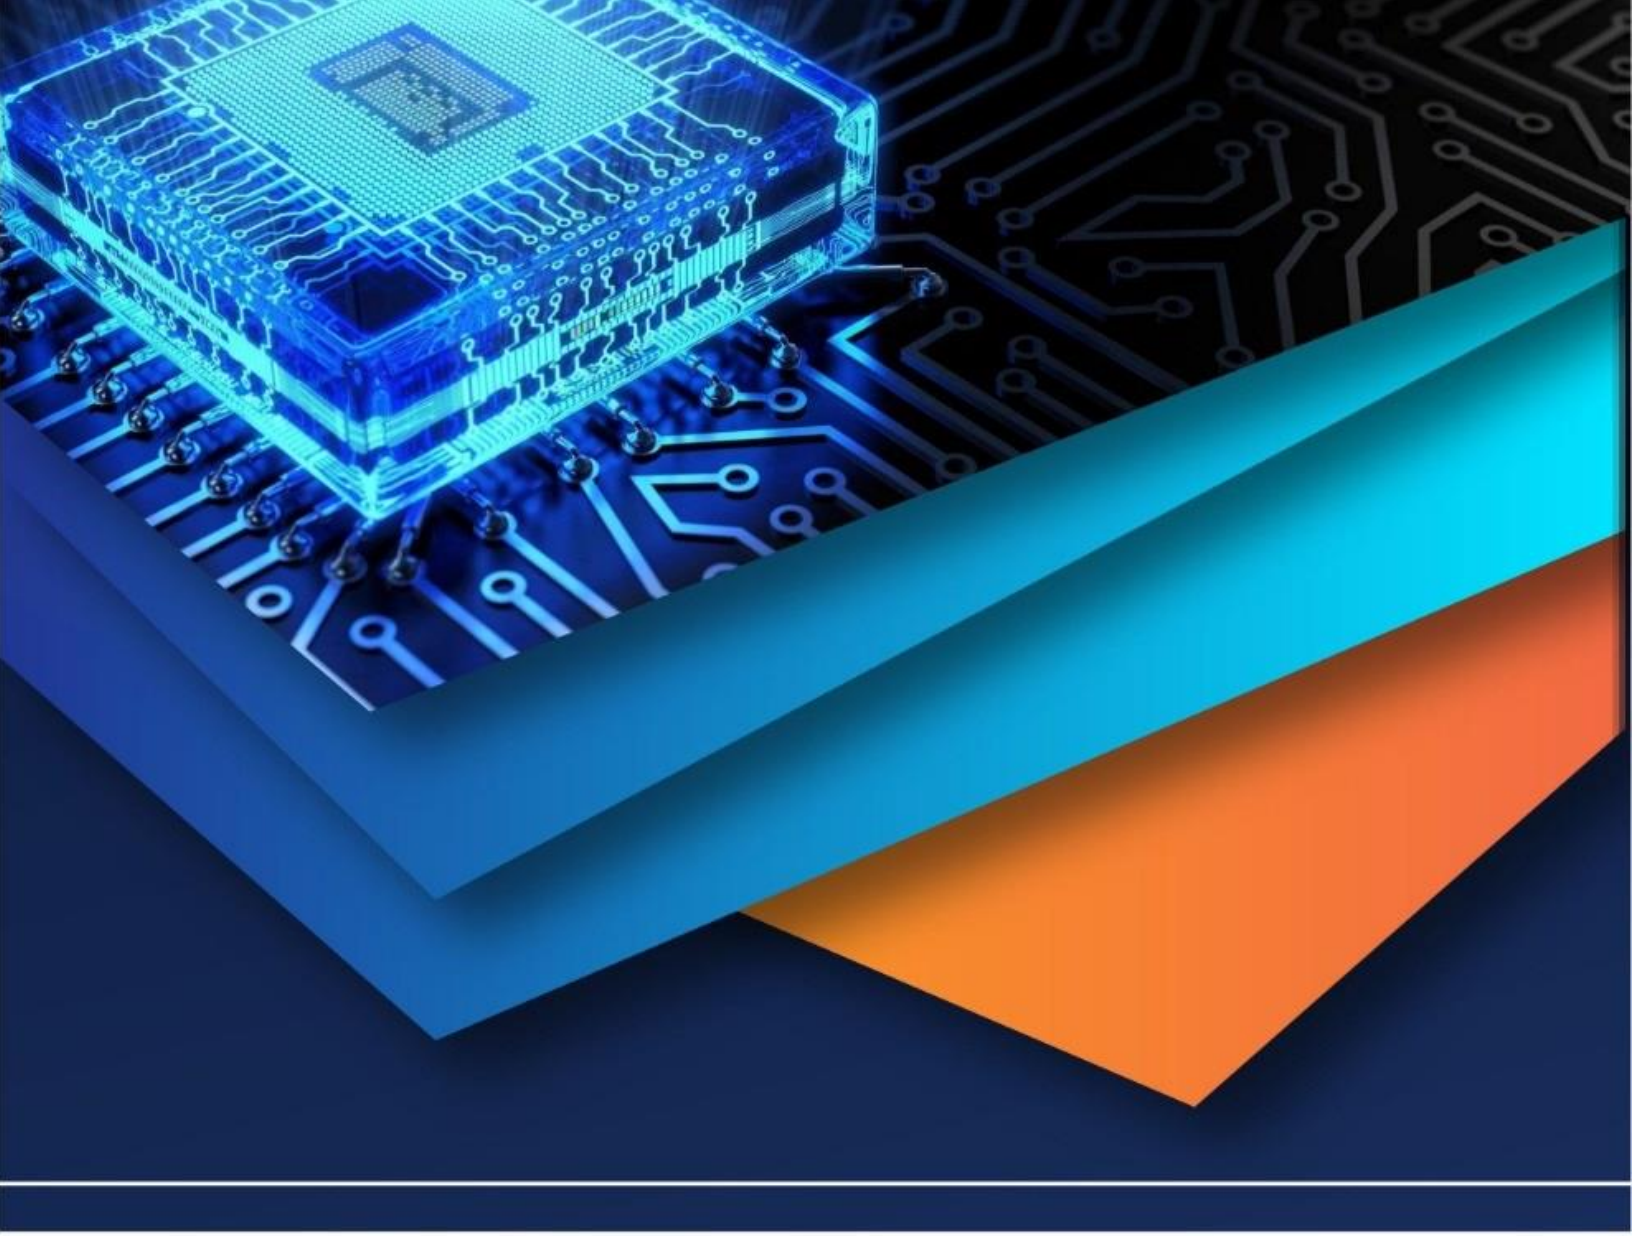
REFERENCES

[1] Wang, "An intelligent system for classifying patient complaints using ML-based NLP," Journal of Medical Internet Research, Jan. 2025. DOI: 10.2196/55721. Accessible at: https://doi.org/10.2196/55721
[2] Filgueiras and L. Barbosa, "Complaint analysis and classification for economic and food safety," ACL Workshop on Economics & NLP, Nov. 2019. DOI: 10.18653/v1/D19-5107. Accessible at: https://doi.org/10.18653/v1/D19-5107
[3] Sara et al., "Proactive complaint management in public sector informatics using AI," Applied Sciences (MDPI), vol. 15, no. 12, 2024. DOI: 10.3390/app15126673. Accessible at: https://doi.org/10.3390/app15126673
[4] Airani and N. Pipada, "Classification of consumer complaints text data by ensembling large language models (LLM)," SCIT-EPress Proceedings, May 2025. DOI: 10.5220/0013173900003890. Accessible at: https://doi.org/10.5220/0013173900003890
[5] Rani et al., "Automated classification of cybercrime complaints using transformer-based language models for Hinglish texts," arXiv preprint, Dec. 2024. No DOI. Accessible at: https://arxiv.org/abs/ (add arXiv ID when known)
[6] J. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed., Prentice Hall, 2020. No DOI. Accessible at: https://aima.cs.berkeley.edu/
[7] Hajihashemi, "Understanding and Managing Customer Complaints," Ph.D. dissertation, Univ. of Washington, 2023. No DOI. Accessible at: https://digital.lib.washington.edu/researchworks/handle/1773/ (add specific handle if known)
[8] Adapted from analysis of conventional ML methods in enterprise complaint data. No DOI or publicly accessible link.
[9] Devlin, J. et al., "BERT: Pre-training of deep bidirectional transformers for language understanding," NAACL, 2019. No DOI. Accessible at: https://aclanthology.org/N19-1423/
[10] Sun, T. et al., "MobileBERT: A compact task-agnostic BERT for resource-limited devices," ICASSP, 2020. DOI: 10.1109/ICASSP40776.2020.9054696. Accessible at: https://doi.org/10.1109/ICASSP40776.2020.9054696
[11] Sanh, V. et al., "DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter," arXiv preprint, 2019. No DOI. Accessible at: https://arxiv.org/abs/1910.01108
[12] Ziser, T. and Reichart, R., "Neural models for domain adaptation in text classification: A comprehensive survey," ACM Computing Surveys, 2021. DOI:
[13] 10.1145/3446775. Accessible at: https://doi.org/10.1145/3446775
[14] Gururangan, S. et al., "Don't stop pretraining: Adapt language models to tasks and domains," ACL, 2020. DOI: 10.18653/v1/2020.acl-main.740. Accessible at: https://doi.org/10.18653/v1/2020.acl-main.740
[15] Wallach, H. M. and Mimno, D., "Modeling the historical dynamics of scientific fields," ICML, 2013. No DOI. Accessible at: https://proceedings.mlr.press/v28/wallach13.html
[16] Zhang, Y. and Wallace, B. C., "A sensitivity analysis of (and alternatives to) classifier's output on noisy features," AAAI, 2020. No DOI. Accessible at: https://ojs.aaai.org/index.php/AAAI/article/view/5954
[17] Alam, F. et al., "A survey on code-mixed language processing," ACM Transactions on Asian and Low-Resource Language Information Processing, 2020. DOI: 10.1145/3417987. Accessible at: https://doi.org/10.1145/3417987
[18] Hovy, D. and Loper, E., "Experiments with a unified model of text normalization and part-of-speech tagging for English tweets," ACL Workshop on Noisy User-generated Text, 2014. No DOI. Accessible at: https://www.aclweb.org/anthology/W14-5917/

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

[19] Mehrabi, N. et al., "A survey on bias and fairness in machine learning," ACM Computing Surveys, 2021. DOI: 10.1145/3457607. Accessible at: https://doi.org/10.1145/3457607.
[20] Fowler, M., Patterns of Enterprise Application Architecture, Addison-Wesley, 2002. No DOI. Accessible at: https://martinfowler.com/books/eaa.html.
[21] Sculley, D. et al., "Hidden technical debt in machine learning systems," NIPS, 2015. No DOI. Accessible at: https://papers.nips.cc/paper/2015/hash/86df7dcfd896fcaf2674f757a2463eba-Abstract.html.
[22] Sandhu, R. S. et al., "Role-based access control models," IEEE Computer, 1996. DOI: 10.1109/2.485845. Accessible at: https://doi.org/10.1109/2.485845.
[23] He, H. and Garcia, E. A., "Learning from imbalanced data," IEEE Transactions on Knowledge and Data Engineering, 2009. DOI: 10.1109/TKDE.2008.239. Accessible at: https://doi.org/10.1109/TKDE.2008.239.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



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)