

Sociosphere: A Social Network Platform For Empowering Real-World Social Change

¹Purushotam Naidu k*, ²R.Srilatha, ³S.Gayathri, ⁴U.N. Harshitha, ⁵ P. Siri Chandana

¹Associate Professor, Dept. of CSE (AI&ML), GVP College of Engineering for Women, Visakhapatnam, AP, INDIA.

^{2,3,4,5}, Student, Dept. of CSE (AI&ML), GVP College of Engineering for Women, Visakhapatnam, AP, INDIA.

Abstract- Urban population growth creates new challenges related to civic infrastructure, and there is a need for efficient and smart complaint management systems. This paper describes the SocioSphere, which is an AI-based civic issue management platform that uses Natural Language Processing (NLP), machine learning, and high-performance web technologies to automatically process and route complaints. A report verification module (Fake/Real) built with Logistic Regression and engineered textual features can filter out spam and low-quality complaints. Valid complaints use a transformer model (RoBERTa) to identify the multi-class categories to which the complaint belongs. Furthermore, we have added a method of estimating the urgency of a complaint through the use of VADER-based sentiment analysis and heuristics for engagement, thus allowing for priority-based decisions. FastAPI is used to develop the backend API layer, offering high-speed (asynchronous/low latency) performance for model inferences and data processing. Complaints will be stored in the system's database and dynamically routed to appropriate authorities for final resolution. The experimental results demonstrate that the approach is effective for both classification and validation, as well as improving transparency and reducing manual work through the use of data-driven governance within smart city systems.

Keywords - Artificial Intelligence, Natural Language Processing (NLP), Civic Issue Management, RoBERTa, Logistic Regression, Sentiment Analysis, VADER, FastAPI, Smart Cities.

I. INTRODUCTION

Smart Cities are created in response to the fast-paced growth of technology (digital technologies) and the increasing number of people moving to cities (urbanisation). They are supposed to improve the quality of life within cities. For example, smart cities need to take on the tracking/management of community services, such as waste management, road maintenance, water supply, and public safety. As part of the smart city strategy, many different types of research and technology have been developed to engage citizens with their local governments to help find solutions to community-based issues. This includes creating and developing community complaint management systems and alternative ways of doing local government business through the use of data and evidence.

The literature survey focuses on analysing existing systems, research papers, and platforms related to civic issue reporting, participatory sensing, and intelligent complaint management. Traditional systems primarily rely on manual processes for registering and handling complaints, which often result in inefficiencies, delays, and a lack of

transparency. With the advent of web-based and mobile applications, platforms such as civic reporting systems have enabled citizens to directly communicate issues to authorities. However, many of these systems lack intelligent automation, requiring manual categorisation and prioritisation of complaints. Recent research has explored the integration of artificial intelligence and machine learning techniques to improve the efficiency of such systems. Natural Language Processing (NLP) models have been widely used for text classification tasks, enabling automatic categorisation of user-generated complaints. Transformer-based models, such as BERT and RoBERTa, have demonstrated superior performance in understanding contextual information and handling multi-class classification problems. Similarly,

sentiment analysis techniques have been employed to estimate the urgency or severity of reported issues. The literature review describes how we feel about existing civic issue reporting systems, academic literature related to civic issue reporting systems, and digital platforms connected to civic issues. The traditional way of dealing with complaints (doing everything manually) is very inefficient, time-consuming and leaves the public unsure as to what happens with their complaint. Web-based/mobile applications have made it easier for citizens to

communicate directly with the appropriate authorities to get their complaints addressed, but the majority of these systems don't have any automated capability to categorise or prioritise complaints. Recent studies have looked into using AI/ML/AI techniques in order to make civic issue reporting more efficient. NLP models have been utilised as a method of classifying text and, therefore, allow for automatic classification of user-submitted complaints. Transformer-based NLP (e.g., BERT, RoBERTa) have been shown to be more effective than other approaches with respect to understanding context/effectively addressing multi-class classification issues. Sentiment analysis has also been utilised as a means to determine the level of urgency/severity associated with reported issues.

The literature also identifies a number of important issues related to detecting fake spam and low-quality reports. There have been a variety of approaches suggested by researchers, including rule-based approaches, machine learning models, and deep learning methods, to assist with identifying irrelevant and/or misleading sources of information. By improving the reliability of the system, these methods will ensure only valid and actionable complaints are processed.

Although these advancements are positive steps toward helping establish a reliable system for processing complaints, there are several limitations to existing solutions. These limitations include: limited scalability; limited automation; insufficient validation mechanisms; and poor integration of multiple intelligent components within one system. Additionally, very few platforms provide any form of real-time feedback mechanism; therefore, the ability to efficiently handle vast amounts of complaint data is also greatly hindered because of a lack of effective methods for prioritising.

The SocioSphere project integrates the existing research work with a single unified platform, which incorporates AI-based issue categorising, prioritising, and validating. The system is based on a combination of current NLP methods, user-centred design processes and structured systems architectures, and is designed to overcome gaps identified in the literature to provide an efficient, scalable, and intelligent solution for civic issue management.

II. RELATED WORK

[1] The development of digital platforms is impacting how local governments and citizens interact by allowing people to report non-emergency urban problems (e.g., potholes, graffiti, etc.) directly to government agencies. This research will analyse participatory reporting data along with neighbourhood demographic variables, socioeconomic factors, pedestrian-friendly conditions, as well as transit and bike accessibility in different neighbourhoods within New York State's capital region. Our data-driven methodology provides a scalable, affordable, and functional alternative to traditional survey-based methodologies, providing an effective means of understanding citizen engagement, satisfaction levels, and producing public value through digital platforms. Findings from this research will provide local government service agencies with tools to develop collaboration with individual neighbourhoods while developing improved communication channels (online and offline) for the reporting of urban problems.

[2] The purpose of this document is to provide a framework for constructing and assessing a value network for mobile services in the public sector. This document uses an already established and validated framework and adds additional criteria that are characteristic of public authorities, such as city administrators, when they are involved in a Value Network. The rapid pace of change in the mobile telecommunications industry makes this framework a valuable resource for both theoretical and practical purposes, and will allow for comparison and analysis of complex mobile city service business models using public sector stakeholders. The framework is validated through a case study using FixMyStreet.

[3] Platforms for social media use personalised experiences to increase user engagement, with machine learning (ML) being one of the key technologies utilised to create this personalisation. This paper discusses how different machine learning techniques (e.g., collaborative filtering, content-based filtering, hybrid frameworks and natural language processing) are used to analyse user behaviours, preferences and sentiment to provide personalised content and improve how users interact with the platform. Also discussed are advanced ML technologies, including deep learning and reinforcement learning and their effectiveness at maximising the accuracy and adaptability of

recommendations. In addition, this paper identifies the challenges hindering successful incorporation of these technologies — including problems regarding privacy of data, algorithmic bias, and ethical considerations — thus emphasising the necessity for responsible use of data. Ultimately, this study demonstrates how ML technology-driven personalisation leads to increased user satisfaction, engagement and satisfaction with an overall social media experience.

[4] This paper sets out a decentralised incident reporting system (IRS) design utilising blockchain technology and IPFS to detect and manage incidents more efficiently in terms of urban areas. Citizens will be able to report incidents to the IRS in real time, which will allow for greater transparency, security, and data integrity, as well as reduced storage on the blockchain as a result of using a decentralised file storage system. This paper also presents a proposed incentive model that will encourage users to participate and further engage, and promote usability through mobile integration. Through the use of the Quorum-Raft blockchain, our experimental results have demonstrated a high level of performance, an average latency of under 15 seconds, scalability up to 200 incidents per second, and an overall low cost to implement and maintain. The IRS we have developed is a highly capable, scalable solution to meet the needs of modern urban emergency management.

[5] The author's work examines how AI assists in delivering more personalised social media marketing and how this impacts customer experience among people within the MENA region. Data collected from 893 participants revealed that AI helps to improve the effectiveness of marketing with tailored offerings delivered to customers, the ability to efficiently select an influencer, and allows for real-time interactions with customers. By utilising these capabilities, companies will increase customer awareness and increase their likelihood of purchasing (the purchase intention) while also providing a more enjoyable experience for the user through the delivery of their marketing message or offer. Tools utilised to engage consumers, like influencer marketing, optimising the content of their offers, and customising the offer, will be effective; however, future studies should be conducted to research even further into the ways AI can help companies effectively enhance the way they engage consumers. The researchers also discovered that AI is capable of assisting companies by analysing massive databases and supporting the development of

targeted marketing strategies, and providing useful information to help companies improve their marketing methods through the use of AI.

[6] A smart city has been created due to the integration of ICT within traditional infrastructure to improve urban management and operation digitally. This article lists five essential objectives to help create smart cities: to be able to identify and analyse the many challenges that urban environments face; to coordinate all technologies used; to utilise data generated from urban regions more efficiently; to improve governance in cities; and lastly, to address the various risk and sustainability issues that exist in cities. The article concludes by noting four large research challenges that still exist to aid in creating smart cities: Improving the efficiency of infrastructure, allowing for new and innovative ways to provide services, achieving equity, and allowing more citizens to participate in creating their communities. It also discusses several different smart city scenarios and project types: Urban modelling; Mobility; Data systems, and participatory governance, as well as looking towards future developments and paradigm shifts for the continued progress of smart city development.

[7] This article sees four major issues remaining to be solved to help develop smart cities: Improved infrastructure performance, new, innovative means of delivering services, equity, and increased citizens' involvement in developing their communities. Other topics explored include numerous examples of smart city scenarios and types of projects, including urban modelling, mobility, data systems, and participatory governance, as well as the future direction and change in paradigm to continue the advancement of smart city development.

[8] The introduction of BERT (Bidirectional Encoder Representations from Transformers) was presented in the paper "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". With the introduction of BERT, the authors of the paper created a language model capable of learning high-quality representations through the use of both left and right context. This means that the model can be fine-tuned with very little effort on multiple tasks, such as language inference (LI) and question answering (QA). BERT is a very simple (yet powerful) model that has achieved state-of-the-art results on multiple natural language processing (NLP)

benchmarks, including GLUE, MultiNLI, and SQuAD.

[9] Although pretraining language models have made substantial advancements, it is still challenging to compare different methodologies due to high computation requirements, multiple dataset sizes, and considerable variability in hyperparameter selection. This re-creation of the BERT pretrained model (Devlin et al., 2019) serves as a platform to systematically evaluate how fundamental hyperparameters and data scale affect the final performance of the models based upon a stable dataset. BERT was originally undertrained and has subsequently demonstrated the ability to be trained to rates of performance equal to or exceeding those of later developed pretrained models. The models developed by the authors were shown to be state-of-the-art on multiple benchmarks (e.g., GLUE, RACE, SQuAD). The authors highlight how many design decisions that have generally been overlooked are very important, and suggest the basis for recent improvements and develop additional databases of models with codes available for download.

[10] Traditional sequence transduction models use nearly all of the complex recurrent and/or convolutional (CNN) neural networks to create end-to-end models encoded via an encoder-decoder architecture and typically have an attention mechanism to improve their performance. In this paper, we propose a new architecture called Transformer, which is a simpler architecture that uses only attention and does not use recurrent or convolutional layers. In empirical studies of machine translation tasks, the Transformer model outperformed traditional sequence transduction models and allowed for greater parallelism and faster training. The improved performance of the Transformer model, it set new benchmarks on English-to-German and English-to-French translation tasks based on BLEU scores, as well as having a significant reduction in training time. The Transformer model was also able to generalise to tasks outside of machine translation, such as English constituency parsing, even with limited training data.

[11] Modern mobile devices produce much useful confidential data that has varying privacy policies. This presents numerous obstacles for centralised training methods. An alternative solution that addresses these issues is called Federated Learning,

where models are trained locally on user devices with only the aggregate updates being sent. In this study, we propose an implementation of the Federated Learning concept using an iterative version of model averaging, and we demonstrate that our method is effective across a variety of model architectures and datasets, including datasets that are heavily imbalanced and/or have non-IID distributions of data. Additionally, the proposed method results in a reduction in communication costs of between 10 and 100 times relative to conventional approaches used for training machine learning models.

[12] This paper aims to analyse how well different types of linguistic features work together when determining the sentiment within Tweets. The research will focus on two main areas: Type A - existing lexical resources; and Type B - new linguistic features which represent the many types of informal and creative language used on microblogs. The research will utilise a supervised learning process and develop training data from the Twitter posts using hashtags contained in the tweets.

[13] Social media content makes sentiment analysis challenging due to its informal nature. This paper introduces VADER, a simple rule-based model designed for sentiment analysis in microblog contexts. By combining specialised lexical features with rules for grammar and sentiment intensity, VADER outperforms several traditional methods and even human raters, demonstrating strong accuracy and better generalisation across different contexts.

[14] We present two novel model architectures capable of generating continuous representations (vector embeddings) of words based on extensive data sets. The produced vector embeddings are evaluated for quality by performing word similarity tasks and compared to the previous state-of-the-art neural networks. The two proposed models produced 'significantly more accurately' than all previously published models and show clear evidence of performance improvements with an accompanying 'increase in computational' ease of use, completing training (on 1.6 billion words) in less than one day. Furthermore, the vector embeddings that these two models learn represent the best known values from literature, outperforming all other state-of-the-art methods in capturing word-related (syntactic and semantic) attributes.

[15] This paper presents two novel model architectures capable of generating continuous representations (vector embeddings) of words based on extensive data sets. The produced vector embeddings are evaluated for quality by performing word similarity tasks and compared to the previous state-of-the-art neural networks. The two proposed models produced 'significantly more accurately' than all previously published models and show clear evidence of performance improvements with an accompanying 'increase in computational' ease of use, completing training (on 1.6 billion words) in less than one day. Furthermore, the vector embeddings that these two models learn represent the best-known values from literature, outperforming all other state-of-the-art methods in capturing word-related (syntactic and semantic) attributes.

[16] The concept of a Smart City (SC) is designed to enhance the quality of life for its residents, yet there is no single clear definition of what this term encompasses. In order to achieve an improved understanding of the initiatives associated with smart cities, this paper proposes several domains—e.g., energy, transport, buildings, governance, economy and people—upon which a taxonomy is constructed. As this paper demonstrates through experience and research, the smart development of a city is based on local characteristics, but not limited to, the economic climate, urban form or layout, geography and population density. The results of this research can be utilised by public sector leaders in the planning of smart city strategies that are effective and contextualised to the environment in which they will be implemented.

[17] Administrative challenges in urban settings include report processing and managing civic concerns. This paper will provide an AI-Multimodal brief to explore the benefits of using MLLM's to evaluate geotagged photos, provide descriptive text for images, categorise the problem with an automatic prioritised solution and provide an AI-based before-and-after validity method to verify whether the problem has been solved via semantic change. Additionally, this will result in increasing transparency, consistency and efficiency of the

management of civic issues with continued oversight from humans.

[18] The Smart Civic Complaint Analyser uses AI to help citizens submit complaints and also helps officers resolve those complaints sooner. The analyser uses Natural Language Processing (NLP) and machine learning (TF-IDF and Naive Bayes) to classify citizen complaints by category, determine how urgent the complaint is, and find out where the complaint is coming from by reading what was originally submitted via text. Citizens can submit and track their complaints through the dashboard, while officers will have access to visual representations of those complaints on the dashboard so they can see trends in the overall number of complaints submitted. The system was built using Python; therefore, its utilisation of Streamlit provides better efficiency, increased transparency, and more citizen involvement in the urban governance process.

[19] AI can reduce the need for humans to perform tasks manually by using powerful computer technology to quickly analyse large amounts of data, including any type of written complaint. This research paper proposes an automatic method for classifying complaints using machine learning and natural language processing techniques. The proposed automatic method for classifying complaints uses several techniques, such as tokenisation, stemming, and multiple classifying methods, to classify complaints (more than 70% accuracy on a unique dataset created for the purpose and 86% accuracy on an already created complaint dataset). This would significantly reduce manual input and improve the speed/efficiency of the process compared to manually classifying complaints.

[20] Developing smart cities requires an analysis of the complaints and an understanding of how to use the information gathered through that analysis to make proactive decisions. The use of urban computing to cluster complaints by location and category allows us to begin to identify the underlying issues causing these civic issues. By using a two-phase clustering technique applied to data collected from New York and Bangalore, it has allowed us to develop solutions that are specific to a region that will assist urban planners and create a more vibrant, livable place to call home.

TABLE 1. RELATED WORKS ON SOCIOSPHERE PROJECT

Category	Researcher	Method	Description	Limitation
Civic Engagement Systems	Yong et al. (2019) [1]	Social Network Analysis	Analysed civic issue reporting using SeeClickFix data to understand participation patterns	Lacks automation and AI-based categorisation
	Walravens (2013) [2]	Platform-based Model (FixMyStreet)	Studied digital civic platforms for citizen-government interaction	Limited scalability and intelligent processing
NLP-Based Classification	Devlin et al. (2019) [8]	BERT	Transformer-based model for contextual text classification	High computational cost
	Liu et al. (2019) [9]	RoBERTa	Optimised transformer model improving NLP classification accuracy	Requires large datasets and resources

	Qurat-ul-Ain et al. (2022) [19]	TF-IDF + ML	Automated complaint classification using traditional ML techniques	Limited contextual understanding
Sentiment & Prioritization	Hutto & Gilbert (2014) [13]	VADER Sentiment Analysis	Rule-based sentiment scoring for social media text	Less effective for complex linguistic contexts
	Kouloumpis et al. (2011) [12]	Sentiment Classification	Analysed sentiment in social media for opinion mining	Limited domain-specific accuracy
Fake/Spam Detection	Essa et al. (2023)	BERT + ML Hybrid	Detects fake content using deep learning and ML models	High computational complexity
	Naveen et al. (2026) [17]	AI-based Validation	Multimodal AI framework for validating civic complaints	Still conceptual, limited real deployment
Smart City & Data Analytics	Bansal & Toshniwal (2016) [20]	Data Mining & Clustering	Analysed civic complaints for proactive maintenance	No real-time automation
	Neirotti et al.	Trend Analysis	Studied smart city initiatives and data-	Lacks implementation-

	. (2014) [16]		driven governance	level details
--	------------------	--	-------------------	---------------

III. METHODOLOGY

A. Roberta (Robustly Optimised Bert Pretraining Approach)

The RoBERTa (Robustly Optimised BERT Pre-training Approach) model is a state-of-the-art transformer-based architecture that was developed by Facebook AI for Natural Language Processing (NLP) tasks. It represents an optimised version of the original BERT (Bidirectional Encoder Representations from Transformers) model and improves on the performance of BERT by optimising the pre-training process in several ways, including the removal of the Next Sentence Prediction (NSP) task, employing dynamic masking, and training on larger datasets using larger batch sizes. As a result of these changes,

RoBERTa has an improved capability to represent the contextual relationships between

words in a given sentence.

The RoBERTa model is built upon the Transformer encoder-based architecture, which processes all of the input sequence at once (in parallel) instead of one input token at a time (in sequence). This process allows the RoBERTa model to learn and represent long-range dependencies between words in the same sentence, as well as the contextual meaning associated with those words, more effectively than if only one of the words in the input sequence were processed before its dependent word.

The overall architecture of RoBERTa consists of multiple (i.e., stacked) layers of encoder structures, where each encoder structure contains a multi-headed self-attention mechanism and a feed-forward neural network (FFNN) layer.

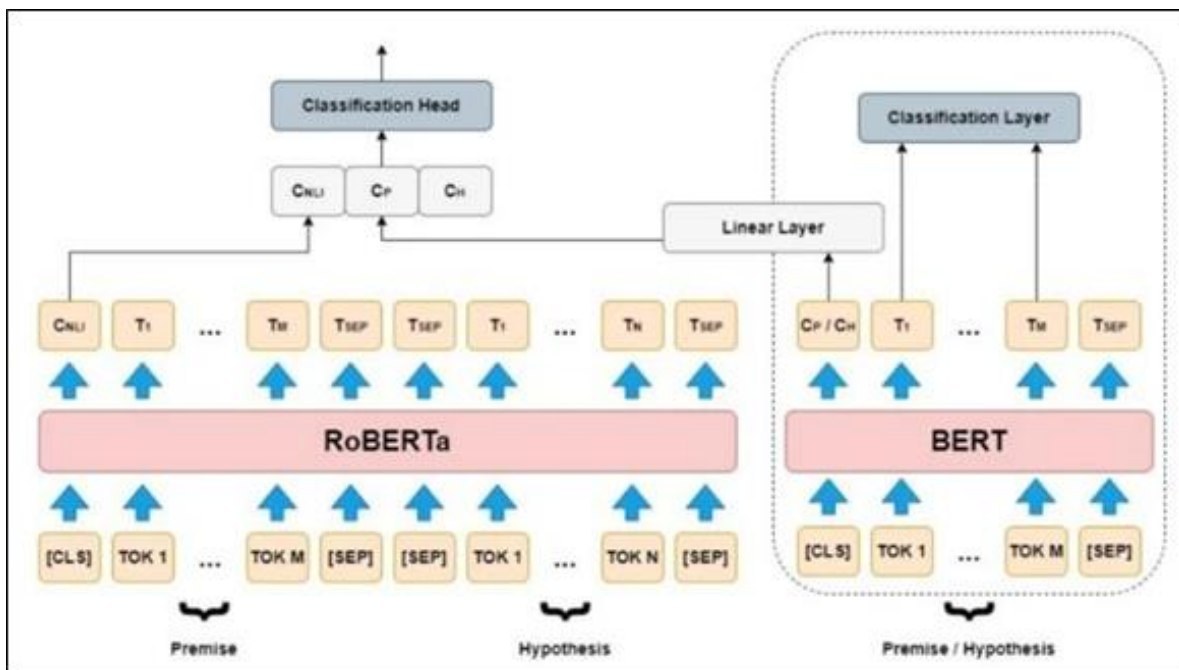


Fig 1. Structure of RoBERTa Architecture

Figure 1 shows how the RoBERTa architecture uses the Transformer's encoder model. Initially, input text has been tokenised into tokens (tokens) before they are converted into feature vectors (embeddings) that contain both Token Embedding and Positional Embedding.

These embeddings' features then undergo successive transformations within a series of encoder layers (shown in the figure above). During each layer, their features are changed through two processes: (1) Self-attention is applied to identify how the two types of features relate to each other, and (2) Feed-Forward Networks (FFN) refine those initial values into the resultant, final feature values that will be used to represent the original input textual content.

Once this process is complete, the final layer of your RoBERTa model creates a representation of your original input textual content, which will be used to perform various tasks associated with classification and prediction.

Let the input sequence be represented as $X = (x_1, x_2, \dots, x_n)$. The embeddings are passed into the self-attention mechanism, where Query (Q), Key (K), and Value (V) matrices are computed. The attention mechanism is defined as:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V$$

(1)

The final contextual representation is obtained after passing through multiple encoder layers.

The RoBERTa architecture comprises numerous Transformer encoder layers, each of which is responsible for four key tasks:

Multi-head self-attention enables encoders to focus on various elements of the same input sentence, helping the model to represent different contextual relationships between the inputs.

Feedforward Neural Networks will apply non-linear transformations to all tokens to produce deeper representations of the input than would be possible without such transformations. The mathematical definition of the feedforward function is as follows:

$$FFN(x) = \max(0, xW_1 + b_1)W_2 + b_2$$

2

Each of the four layers will perform Layer Normalisation to ensure the stability and speed of the training process.

The Residual Connection preserves past learning and avoids the vanishing gradient problem.

Using Masked Language Model (MLM), during the training of RoBERTa, certain input tokens are randomly selected and then masked out of the input stream. The model attempts to predict masked tokens mostly from their context, which would be the remaining non-masked data.

The model loss function used for RoBERTa training is:

$$\mathcal{L} = -\sum \log P(w_i | context)$$

(3)

RoBERTa builds on BERT in several important ways, including the removal of the Next Sentence Prediction task; the use of dynamic masking (where masked tokens are different each time they are seen during training); training on larger data sets; and larger batch sizes.

Because of its strong capacity for capturing deeper contextual relationships, RoBERTa performs extremely well when used within applications like detecting fake reports, classifying issues, performing sentiment analyses, or classifying text generally. RoBERTa's strong training strategy also makes it extremely accurate and dependable when completing natural language processing tasks in the real world.

B. Logistic Regression

The Logistic Regression Model is a supervised machine learning algorithm that helps solve classification problems, typically binary classification. The term 'Logistic' in Logistic Regression can be confusing because it is actually a classification algorithm rather than a regression algorithm. Through the use of the sigmoid or logistic function, Logistic Regression can compute the probability that a particular input will belong to a certain classification. Logistic Regression has many uses, including spam detection, disease diagnosis, and systems that detect fraudulent reports.

Logistic Regression establishes a connection between the input variables/attributes and the probability that a given class output occurs. Logistic Regression does not estimate a continuous number; it produces a decimal between 0 and 1, which is then converted into its corresponding class according to its threshold value (usually set at 0.5).

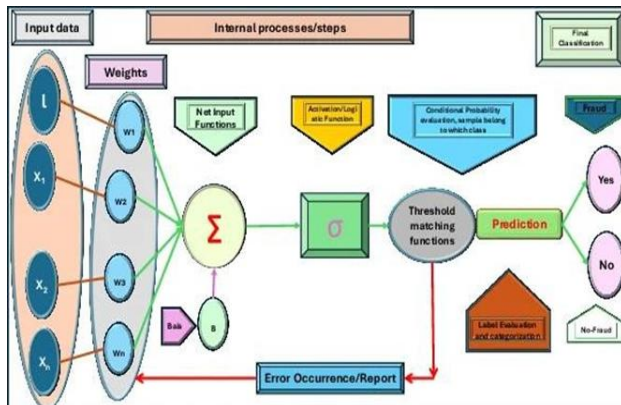


Fig 2. Structure of Logistic Regression Model

This diagram depicts how the Logistic Regression function works. The input features are combined linearly using weights and bias. The result is then passed through the sigmoid function, which converts the linear output into a probability value. Based on this probability, the model assigns the input to a particular class.

Let the input feature vector be represented as $X = (x_1, x_2, \dots, x_n)$. The linear combination of inputs is given by:

$$z = w_1x_1 + w_2x_2 + \dots + w_nx_n + b$$

where w_1, w_2, w_n are the weights and b is the bias. The output of the Logistic Regression model is obtained using the sigmoid function:

$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

$$\hat{y} = \begin{cases} 1 & \text{if } \sigma(z) \geq 0.5 \\ 0 & \text{if } \sigma(z) < 0.5 \end{cases} \quad (6)$$

Logistic Regression is trained using a loss function known as Binary Cross-Entropy (Log Loss), which measures the difference between predicted and

actual values:

$$\mathcal{L} = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})]$$

Components of a logistic regression model include:
(7)

Input Features: Independent variables used to predict values.

Weights and Bias: Parameters that are learned during training and indicate how important each feature is to determining the output.

Sigmoid Function: A function that converts the linear output of the model into a probability.

Decision Boundary: A threshold that differentiates between classes.

Loss Function: A metric that measures the performance of the model and provides guidance for optimising it.

Logistic regression can also be extended to multiclass classification using techniques such as softmax.

Regression (Multinomial Logistic Regression).

Logistic regression is a very popular method used in many practical applications, such as fraud detection, medical diagnosis, and the classification of social issue reports (i.e., identifying whether the report is fake or real), partially due to its simplicity, interpretability, and efficiency.

C. TF-IDF Vectorizer

A TF-IDF (Term Frequency-Inverse Document Frequency) vectorisation is an approach to extracting features that transforms textual data into numerical representations for use in machine learning models, which allows these models to interact with text. The model assigns a score to each type of word according to the frequency at which that word occurs within a document, relative to how many different types of documents

have at least one occurrence of that word. TF-IDF is useful for recognising word importance while minimising the weighting of so-called stop words (e.g., “the”, “is”, and “and”).

In text classification, information retrieval, false report detection, and issue classification systems, TF-IDF is commonly used to reflect a meaningful numerical representation for textual data input.

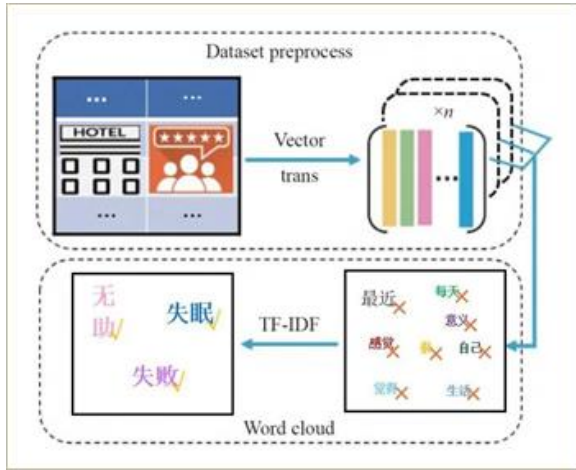


Fig. 3 Structure of TF-IDF Vectorisation Process

In the above diagram, you can see how the TF-IDF Vectorizer works. To begin, the input text documents need to go through pre-processing steps, such as creating tokens for words, removing "stop" words, and normalising all words. After that, the TF and IDF of each word are computed, which gives the correlation between the two vectors. Finally, these two vectors are multiplied together to form a final score called the TF-IDF score for each word in relation to the entire corpus. Let a term appear in a document from a collection of documents D .

The Term Frequency (TF) is defined as:

$$TF(t, d) = \frac{\text{Number of times term } t \text{ appears in document } d}{\text{Total number of terms in document } d} \quad (8)$$

The **Inverse Document Frequency (IDF)** is defined as:

$$IDF(t, D) = \log \left(\frac{N}{1 + df(t)} \right)$$

where:

N is the total number of documents

$df(t)$ is the number of documents containing the term t

The TF-IDF score is calculated as:

$$TF-IDF(t, d, D) = TF(t, d) \times IDF(t, D)$$

A TF-IDF Vectorizer consists of the following components:

- Tokenisation: Splitting text into individual words or terms.
- Vocabulary Creation: Building a set of unique words from all documents.
- Term Frequency Calculation: Measuring how often a term appears in a document.
- Inverse Document Frequency Calculation: Measuring how important a term is across the corpus.
- Vector Representation: Converting each document into a numerical vector based on TF-IDF scores.

TF-IDF reduces the impact of common, yet not uniquely informative, terms, and increases the impact of seldom, yet informative, terms; thus, increasing the effectiveness of Text Classification, Clustering & Information Retrieval Machine Learning models in Text Data driven use cases.

D. Description Of The Dataset

The data referred to in connection with the SocioSphere project is derived from a set of complaints made by citizens to local government offices. This dataset was developed from an open-source project called Open City, which collects citizen complaints pertaining to a variety of urban-related problems (e.g., litter, potholes, insufficient water supply, burned-out streetlights). This data primarily consists of written reports of urban problems.

Dataset Characteristics Issue Categorisation:

- Total Records: ~16,000 (initial dataset)
- After Cleaning & Balancing: ~16,000 records
- Data Type: Structured + Unstructured (Text-based)
- Domain: Urban civic issues / Smart city applications

Fake Report Detection

- Initial Shape: (14791, 3)
- After Cleaning: (6731, 4)
- Removing rows with labels that have only one instance: ['label']
- Shape after removing single instance labels: (6730, 4)]

Features (Attributes)

- Title: Summary of the issue
- Description: Detailed explanation of the complaint
- Category: Target label (e.g., garbage, roads, water supply)
- Combined Text: Title + Description (used for model input)
- Comment Count: Used for urgency/priority calculation

Preprocessing Steps

- Handling missing values
- Text cleaning and normalisation
- Label standardisation (merging similar categories)
- Removal of noisy and irrelevant labels
- Filtering rare classes (minimum sample threshold)
- Dataset balancing (sampling per class)

Final Output

- Number of Categories: 17–18 classes
- Balanced Dataset: Ensures uniform distribution
- Model-Ready Input: Preprocessed text for AI models

IV. EXPERIMENTAL ANALYSIS

A. Evaluation Metrics

To evaluate the performance of the proposed models, standard classification metrics are used. These metrics provide a comprehensive understanding of accuracy, correctness, and reliability.

- **Accuracy**
Measures the overall correctness of the model.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN}$$

- **Precision**

Measures how many predicted positive cases are actually

$$Precision = \frac{TP}{TP+FP} \quad (12)$$

- **Recall**

Measures how many actual positive cases are correctly identified.

$$Recall = \frac{TP}{TP+FN} \quad (13)$$

- **F1-Score**

Harmonic mean of Precision and Recall. (14)

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

Where:

- **TP (True Positive):** Correct positive predictions
- **TN (True Negative):** Correct negative predictions
- **FP (False Positive):** Incorrect positive predictions
- **FN (False Negative):** Missed positive cases

Additional Analysis

- **Confusion Matrix:** Visualises classification performance
- **Macro Average:** Equal weight to all classes

- Weighted Average: Accounts for class imbalance

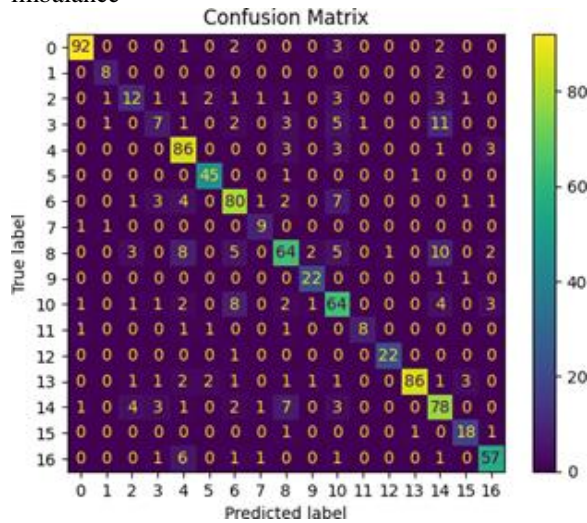


Fig 4. Confusion Matrix of Issue Categoriser and Prioritizer Model

Overall Model Performance:

Table 2. Issue categorizer

Metric	Value
Accuracy	0.79
Precision	0.79
Recall	0.79
F1-Score	0.79

Table 3. Fake/ Real Detection Model

Metric	Value
Accuracy	0.9837
Precision (Avg)	0.985

Recall (Avg)	0.985
F1-score (Avg)	0.98

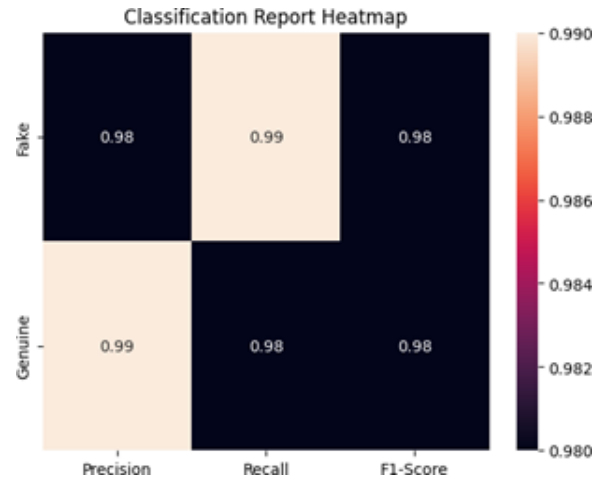


Fig 5. Classification Report Heatmap Analysis

B. Comparison With Other Models

The accuracy of the Issue Categoriser is 79%, with equal precision, recall and F1 values indicating stable and consistent performance across multiple issue categories. Therefore, the model can classify the vast majority of civic issues accurately, although there may occasionally be slight misclassifications caused by overlapping categories.

The Fake/Real Detection model has a very high performance at 98.37% accuracy, along with strong (~0.98) precision, recall and F1 values; therefore, the model can correctly identify real and fake reports most of the time with very few misclassifications.

In conclusion, the complete system allows for accurate classification of issues along with very high confidence in the validation of the classifications, thus being an effective tool to help manage civic issues in the real world.

Table 4. Comparison with other models

Model	Type	Approach	Accuracy	Advantages	Limitations
Logistic Regression	ML (Linear)	Statistical classification using weighted features	~70–75%	Simple, fast, interpretable	Poor performance on complex text
Naïve Bayes	ML (Probabilistic)	Based on the probability distribution of words	~65–72%	Works well on small data, fast	Assumes feature independence
Support Vector Machine (SVM)	ML (Kernel-based)	Finds optimal decision boundary	~72–74%	Effective for high-dimensional data	Computationally expensive
Random Forest	Ensemble ML	Multiple decision trees	~73–74%	Handles non-linearity, robust	Slower, less interpretable
LSTM	Deep Learning	Sequential text processing	~73–75%	Captures sequence information	Training complexity, slower
BERT	Transformer	Context-aware bidirectional model	~75–76%	Strong contextual understanding	High computational cost
RoBERTa (Proposed)	Transformer	Optimised BERT with better training	~76–80%	High accuracy, better generalisation	Requires large amounts of data and resources

C Validation

Comparison between logistic regression, naive Bayes, and support vector classifier all indicate these models will produce moderate accuracy when performing machine-learning tasks without sophisticated reasoning on contextually relevant issues with text datasets. Because of the limitations of traditional ML models at understanding the sequential structure of longer documents, implementations of deep learning techniques demonstrate more ability for learning sequences but require more computational resources. The relative power of transformer devices like BERT and RoBERTa makes them functionally superior to other classes of machine-learning models that are not transformers. Between both the BERT & RoBERTa transformer models available for use, the RoBERTa transformer model should be chosen over BERT in order to achieve optimal accuracy and generalisation

performance during the civic issue classification activity in this project.

Table 5. Categoriser & Prioritizer Validation

Metric	Value
Accuracy	0.79
Macro F1-Score	0.77
Weighted F1-Score	0.79

Observations:

- Structured categories (e.g., roads, garbage, streetlights) have performed the strongest.
- Overlapping categories (e.g., mobility vs. roads) have performed moderately.

- The confusion matrix indicates most of the predictions are along the diagonal (i.e. correct classification of categorical data)
- Misclassification occurs within semantically similar categories.
 - **Prioritisation Validation:**
- Negative sentiment means there is a higher priority for urgent needs.
- High user engagement means that an urgent need will have a greater priority.
- Ensures there is a true way to prioritise the most critical issues.

Table 6. Fake/Real Detection Validation

Metric	Value
Accuracy	0.9837
Precision (Avg)	0.985
Recall (Avg)	0.985
F1-Score (Avg)	0.98

Table 7. Confusion Matrix Summary:

Actual \ Predicted	Fake	Genuine
Fake	676	7
Genuine	13	650

Observations: Very low misclassification rate (<2%) and balanced precision and recall for both classes, showing strong capability for detecting spam, random and out-of-scope input, with high robustness in performance on noisy and informal texts.

Generalisation: Perform well on unseen data and can handle a wide variety of linguistic variations and edge cases effectively Both classifications and validations confirm that the SocioSphere system is capable of performing its intended tasks reliably. Additionally, the Issue Categoriser and Prioritiser model is able to perform multi-class classification accurately while estimating urgency realistically. In contrast, the Fake/Real Detection model has a very high degree of accuracy and is also highly robust to spam and invalid inputs.

The combination of these two models provides an effective and reliable system to process, categorise, prioritise, and validate

civic complaints, making it stable, scalable and deployable to live environments, with the ability for future advancement in more complicated and ambiguous cases.

V. CONCLUSION

The system demonstrates both successful development and realisation of an intelligent, Artificial Intelligence (AI)-driven platform for reporting and managing civic issues. The platform uses a combination of front-end interfaces, back-end services and sophisticated deep learning models to enable users to report issues in an efficient manner and automatically categorises and prioritises all issues reported.

The application of transformer-based Natural Language Processing (NLP) techniques allows the system to better understand the user inputs received in the form of unstructured text and classify this input into appropriate categories of issues (i.e. garbage, roadways, water supply, street lighting). This reduces the manual effort required to process issues, and results in the processing of complaints faster and more accurately.

The modular architecture of the system promotes the interoperability between components of the platform – including user interfaces, back-end Application Programming Interfaces (APIs), AI models, and data storage. The application's implementation has validated the ability of the system to ensure that all complaints reported are routed to the appropriate authorities for action based on their category of classification, resulting in a more structured and transparent workflow.

While there are ongoing efforts to enhance the platform by adding additional advanced features, such as detecting fake/spam reports and full-scale implementation of the database, the existing platform currently provides a robust functional infrastructure to support the implementation of scalable solutions to facilitate civic engagement.

In summary, SocioSphere has shown how Artificial Intelligence in combination with Smart City applications has the potential to enhance governance,

increase citizen engagement, and improve the efficiency of resolving issues. The functionality of SocioSphere could be increased through the addition of features like image-based analysis, real-time notifications, and large-scale implementation for greater benefits.

VI. REFERENCES

- [1] C. Yong, C. Chelmiss, W. Lee, and D. -S. Zois, "Understanding Online Civic Engagement: A Multi-Neighbourhood Study of SeeClickFix," 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), Vancouver, BC, Canada, 2019, pp. 1048-1055, doi: 10.1145/3341161.3345330. URL: <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9073580&isnumber=9072951>
- [2] N. Walravens, "Validating a Business Model Framework for Smart City Services: The Case of FixMyStreet," 2013 27th International Conference on Advanced Information Networking and Applications Workshops, Barcelona, Spain, 2013, pp. 1355-1360, doi: 10.1109/WAINA.2013.11. URL: <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6550584&isnumber=6550285>
- [3] Sreeprasad Govindankutty, Kratika Jain "Machine Learning Algorithms for Personalised User Engagement in Social Media" Iconic Research And Engineering Journals,8(5). URL: <https://www.irejournals.com/paper-details/1706713>
- [4] Diallo, E.-h.; Abdallah, R.; Dib, M.; Dib, O. Decentralised Incident Reporting: Mobilising Urban Communities with Blockchain. Smart Cities 2024, 7, 2283-2317. URL: <https://doi.org/10.3390/smartcities7040090>
- [5] Beyari, H.; Hashem, T. The Role of Artificial Intelligence in Personalising Social Media Marketing Strategies for Enhanced Customer Experience. Behav. Sci. 2025, 15, 700. URL: <https://doi.org/10.3390/bs15050700>
- [6] A. Nam and T. A. Pardo, "Conceptualizing Smart City with Dimensions of Technology, People, and Institutions," Proceedings of the 12th Annual International Digital Government Research Conference, College Park, MD, USA, 2011, pp. 282–291, doi: 10.1145/2037556.2037602. URL: <https://dl.acm.org/doi/10.1145/2037556.2037602>
- [7] M. Batty et al., "Smart Cities of the Future," The European Physical Journal Special Topics, vol. 214, no. 1, pp. 481–518, 2012, doi: 10.1140/epjst/e2012-01703-3. URL: <https://link.springer.com/article/10.1140/epjst/e2012-01703-3>
- [8] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," NAACL-HLT, 2019, pp. 4171–4186. URL: <https://arxiv.org/abs/1810.04805>
- [9] Y. Liu et al., "RoBERTa: A Robustly Optimised BERT Pretraining Approach," arXiv preprint, 2019. URL: <https://arxiv.org/abs/1907.11692>
- [10] A. Vaswani et al., "Attention Is All You Need," Advances in Neural Information Processing Systems (NeurIPS), 2017. URL: <https://arxiv.org/abs/1706.03762>
- [11] H. B. McMahan et al., "Communication-Efficient Learning of Deep Networks from Decentralised Data," AISTATS, 2017. URL: <https://arxiv.org/abs/1602.05629>
- [12] E. Kouloumpis, T. Wilson, and J. Moore, "Twitter Sentiment Analysis: The Good, the Bad and the OMG!," ICWSM, 2011. URL: <https://ojs.aaai.org/index.php/ICWSM/article/view/14185>
- [13] C. Hutto and E. Gilbert, "VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text," ICWSM, 2014. URL: <https://ojs.aaai.org/index.php/ICWSM/article/view/14550>
- [14] T. Mikolov et al., "Efficient Estimation of Word Representations in Vector Space," ICLR, 2013. URL: <https://arxiv.org/abs/1301.3781>
- [15] R. Kitchin, "The Real-time City? Big Data and Smart Urbanism," GeoJournal, vol. 79, pp. 1–14, 2014. URL: <https://link.springer.com/article/10.1007/s10708-013-9516-8>
- [16] P. Neirotti et al., "Current Trends in Smart City Initiatives: Some Stylised Facts," Cities, vol. 38, pp. 25–36, 2014. URL: <https://www.sciencedirect.com/science/article/pii/S0264275113001935>
- [17] Naveen. V, Umesh. I, Raghavender. M, Vaishnavi. B, Devi. G, Dr. Venkataramana. B, 2026, A Theoretical Framework for AI-Assisted Civic Issue Reporting and Validation in Smart Cities, INTERNATIONAL JOURNAL OF ENGINEERING

RESEARCH & TECHNOLOGY (IJERT) Volume 15,
Issue 01, January – 2026,
10.17577/IJERTV15IS010731 DOI :

<https://doi.org/10.5281/zenodo.18482374>

[18] K.Vani and C.Ashwitha, "Smart Civic Complaint Analyzer Using Natural Language Processing," International Journal of Innovative Research in Technology (IJIRT), vol. 12, no. 6, pp. 5408–5415, 2025.

URL:

<https://ijirt.org/article?manuscript=187517>

[19] Qurat-Ul-Ain, A. Shaukat and U. Saif, "NLP based Model for Classification of Complaints: Autonomous and Intelligent System," 2022 2nd International Conference on Digital Futures and Transformative Technologies (ICoDT2), Rawalpindi, Pakistan, 2022, pp. 1-6, doi: 10.1109/ICoDT255437.2022.9787456.

[20] P. Bansal and D. Toshniwal, "Analyzing civic complaints for proactive maintenance in smart city," 2016 IEEE/ACIS 15th International Conference on Computer and Information Science (ICIS), Okayama, Japan, 2016, pp. 1-6, doi: 10.1109/ICIS.2016.7550747