Traffic Detection and Forecasting from Social Media Data Using a Deep Learning-Based Model, Linguistic Knowledge, Large Language Models, and Knowledge Graphs

Wasen Melhem, Asad Abdi and Farid Meziane
Department of Computing and Mathematics, Faculty of Science and Engineering, University of Derby, U.K.

Keywords: Deep Learning, Large Language Models, Allen's Interval Algebra, Region Connection Calculus, Natural Language Processing, Knowledge Graphs, Retrieval Augmented Generation, Instruction Tuning, Fine Tuning.

Abstract: Traffic data analysis and forecasting is a multidimensional challenge that extracts details from sources such as social media and vehicle sensor data. This study proposes a three-stage framework using Deep Learning (DL) and natural language processing (NLP) techniques to enhance the end-to-end pipeline for traffic event identification and forecasting. The framework first identifies relevant traffic data from social media using NLP, context, and word-level embeddings. The second phase extracts events and locations to dynamically construct a knowledge graph using deep learning and slot filling. A domain-specific large language model (LLM), enriched with this graph, improves traffic information relevancy. The final phase integrates Allen's interval algebra and region connection calculus to forecast traffic events based on temporal and spatial logic. This framework's goal is to improve the accuracy and semantic quality of traffic event detection, bridging the gap between academic research and real-world systems, and enabling advancements in intelligent transport systems (ITS).

1 INTRODUCTION

Cities all over the world are experiencing severe traffic congestion and unexpected road conditions. These issues necessitate innovative traffic management strategies to improve road user's experience. Crowdsourcing traffic data from social networks like Twitter offers a cost-effective alternative to traditional sensor-based approaches.

The first step in utilizing social media for traffic management involves extracting relevant traffic data accurately classified to irrelevant or relevant traffic events (Suat-Rojas et al., 2022). Second, Events and locations are extracted using a variety of Information Retrieval and Machine Learning (ML) methods. However, privacy concerns limit the availability of latitude and longitude for tweets (Hodorog et al., 2022), Then events and locations are stored, often as raw data, but more structured methods like Knowledge Graphs (Wu et al., 2022)can be used for organizing information, improving the semantic representation of events. Moreover, Dynamic knowledge graphs using Deep Learning (DL) methods for event detection can enhance the quality of traffic data from social media (Bai et al., 2020) and can also support analysis and decision-making for traffic forecasting, enhancing the performance by representing high-dimensional spatial and temporal data (Zhang et al., 2019).

The development of ITS is an architecture that encompasses information and communication technology (ICT) between vehicles, users, and transportation networks. Recently, large language models (LLMs) such as ChatGPT (Zheng et al., 2023) have gained popularity for tasks like text completion and question answering, showing promise in various fields, this can significantly improve the efficiency and reliability of ITS.

To develop effective approaches for text relevancy, event detection, and forecasting, several methods are commonly employed: In text relevancy, accurately categorizing SM messages into relevant or irrelevant classes is challenged by data sparsity, imbalance, and ambiguity. Variations of Recurrent Neural Networks (RNN) have been used to tackle this issue. However, struggle with noisy, variable-length

Melhem, W., Abdi, A. and Meziane, F.
Traffic Detection and Forecasting from Social Media Data Using a Deep Learning-Based Model, Linguistic Knowledge, Large Language Models, and Knowledge Graphs.
DOI: 10.5220/0013066900003838

In *Proceedings of the 16th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2024) - Volume 2: KEOD*, pages 235-242
ISBN: 978-989-758-716-0; ISSN: 2184-3228

texts. In the phase of event detection methods relying on predefined classes lack adaptability for emerging events. Therefore, dynamically extracting and structuring relevant information from unstructured data is crucial for ITS rather than the static approach usually taken. Existing geoparsing algorithms struggle with accurately identifying and disambiguating location references in short texts, requiring robust approaches like Long Short-Term Memory (LSTM) and Transfer Learning (Das & Purves, 2020). It is important to Investigate models for capturing dynamic spatial dependencies to improve detection and prediction performance, highlighting the need to address the imbalance in traffic data and understand spatial and temporal dependencies necessitates advanced techniques such as Allen's Interval Algebra and Region Connection Calculus (Chuckravanen et al., 2017).

Current research often tackles text relevancy, event and location detection, and traffic forecasting in isolation. This study proposes a unified approach that evaluates these phases together to address issues identified in previous studies. The problem statement as follows:

1. Current social media event detection relies on rigid, predefined categories; this study proposes dynamic event discovery using online learning to capture novel events and store them in a dynamic knowledge graph for improved forecasting.
2. To address hallucinations and knowledge gaps in the traffic domain, this study explores fine-tuning and augmentation methods to enhance domain specificity, inject factual knowledge, and ensure response accuracy.
3. Previous research lacks robust temporal and spatial analysis; this study integrates Allen's Interval Algebra and Region Connection Calculus within a DL model for accurate spatiotemporal traffic forecasting.

The remainder of the paper is structured as follows. Section 2 gives an overview of related work. Section 3 describes the proposed method for Traffic Intelligence and Forecasting Methodology using NLP (TIFFNLP). Our system evaluation is explained in Section 4 Finally, section 5 concludes the work with a discussion as well as proposing some future research directions.

2 LITERATURE REVIEW

The literature for this research has been studied in a phased approach. Therefore, the review of the literature will first discuss text relevancy. Second, event detection and location detection, also using LLMs to query traffic related information. Third, traffic forecasting uses a novel approach.

2.1 Text Relevancy

Determining transport-related text relevance involves identifying whether a text includes traffic information. This process often relies on NLP techniques to analyze and interpret textual data. Traditional methods like Bag-of-Words, rule-based, and dictionary-based techniques often fall short due to their lack of semantic understanding and limited keyword coverage (Fontes et al., 2023). Modern approaches leverage supervised learning and word embeddings for better semantic representation. For instance, (Babbar & Bedi, 2023) discussed how word embedding methods have different purposes while FastText excels with rare words, Word2vec is preferable for short texts like tweets due to its lower memory trail.

Supervised machine learning (ML) methods, such as Support Vector Machine (SVM), Naïve Bayes commonly used to automate classification (Nirbhaya & Suadaa, 2023). However, deep learning models like CNN, RNN, and LSTM offer improved semantic enrichment and relationship identification. (Dabiri & Heaslip, 2019) Demonstrated that combining word embeddings with CNN, RNN, and LSTM models can effectively classify traffic-related tweets, achieving high precision. Transformer models like BERT have further advanced the field, as shown by (Fontes et al., 2023), who achieved significant results despite challenges with large dictionaries and term ambiguity. (Suat-Rojas et al., 2022)Combined doc2vec, TF-IDF, and BERT embeddings to classify tweets related to traffic, despite challenges with informal language and abbreviations. These studies show the evolving methodologies in improving the detection and classification of traffic-related events.

2.2 Event Detection

Classification to identify incidents like congestion, accidents, and weather issues. With the rise of social media, deep learning techniques have become crucial. (Hodorog et al., 2022) Utilized AWD-LSTM and ULMFiT, achieving 88.5% accuracy. (Chang et al., 2022) compared social media-detected events with official reports, using CNN and LSTM, achieving a

76% F1 score. (Yang et al., 2021)framed the problem as a slot-filling task, outperforming other models with a joint BERT-based approach. (Chang et al., 2022) used sentiment-enhanced KDE to prioritize accident-prone areas. (Sun et al., 2021)proposed ED-SWE, filtering tweets with word embedding and Relationship Assessment scoring. (Bok et al., 2023)introduced a graph-based scheme, improving accuracy by clustering semantically distinct event graphs and incorporating social activities. These methods demonstrate the potential of advanced NLP, slot filling, and dynamic knowledge graphs in traffic event detection.

2.3 Location Detection

Detecting location from social media has evolved from rule-based methods to advanced deep learning (DL) frameworks. Recent approaches, such as a study by (Tao et al., 2022), integrated ALBERT, BiLSTM, and CRF, achieving a 96.1% F1 score despite challenges with toponymic words and data imbalance. (Azhar et al., 2023)improved location detection accuracy (80%-94%) using reverse geocoding and Google API, addressing issues of accurate location naming and information reliability. (Zhou et al., 2022) proposed a three-stage model (classification, relation inference, entity pair recognition) to extract interrelated information from noisy social media data, enhancing semantic understanding with knowledge graphs. These studies highlight the progress and potential of integrating DL, semantic analysis, and geospatial techniques for accurate and efficient location detection from social media.

2.4 Large Language Models

Large Language Models (LLMs) can generate "hallucinations," or incorrect responses, categorized into intrinsic (contradictions within training data) and extrinsic (unverifiable information) types (Mihindukulasooriya et al., 2023). This reduces trust, especially in safety-critical areas (Zheng et al., 2023). Efforts to improve domain-specific accuracy include prompt engineering, Reinforcement Learning from Task Feedback (RLTF), fine-tuning (Balaguer et al., 2024), and Retrieval Augmented Generation (RAG) (Fan et al., 2024). Fine-tuning uses labelled data to adapt models for specific tasks but is costly. Techniques like Parameter-Efficient Fine-Tuning (PEFT) reduce computational demand (Houlsby et al., n.d.)while Localized Fine-Tuning (LOFIT) uses sparse attention subsets to improve truthfulness and reasoning (Yin et al., 2024).

For specialized domains, methods like FinGPT and Fin-LLaMA enhance LLMs in finance, excelling in predictive analysis and financial tasks (Yang et al., 2023). TrustLLM improves smart contract auditing with iterative cause selection, achieving over 91% in F1 score and accuracy (Ma et al., 2024). RAG enhances LLMs by retrieving external information based on input prompts, improving tasks like drug discovery and financial analysis (Wang et al., 2024). For example, RAG-guided molecule generation shows promise for SARS-CoV-2 compound design, while it also enhances financial sentiment analysis by incorporating external sources like news and social media (S. Zhang et al., 2023)Comparisons between RAG and fine-tuning in agriculture-specific contexts reveal that fine-tuning significantly improves knowledge and task-specific accuracy (Balaguer et al., 2024).

2.5 Traffic Forecasting

Social media data, like geo-tagged Twitter, impacts traffic prediction, and deep learning (DL) extracts relevant features due to social networks' graph structure (Yuan et al., 2021). Spatiotemporal forecasting has leveraged Graph Neural Networks (GNN), CNN, and Graph Convolutional Networks (GCN) with attention mechanisms. LSTM effectively captures high-dimensional temporal features and is widely used for traffic flow prediction (Lu et al., 2020; Chen & Chen, 2022). Lu et al. (2020) addressed RNN limitations in modeling spatial aspects, introducing multi-diffusion convolution (MDC) to overcome them. Chen & Chen (2022) proposed using GCN with absolute value matrices to capture dynamic spatial patterns. Recent advances highlight ST-GAT's superiority in real-world traffic forecasting, demonstrating scalability and robustness, with future efforts aimed at incorporating additional factors and addressing missing data (H. Zhang et al., 2019). Allen's Interval Algebra (AIA) defines 13 temporal relations between events (Allen, 1983), and Region Connection Calculus represents regions via 8 possible relations in topological space (Randell et al., 1992). AIA has been used in smart homes, planning, and scheduling (Chuckravanen et al., 2017). Advancements in text relevancy, event detection, and traffic forecasting are largely driven by deep learning, though challenges like semantic understanding and data imbalance continue to persist. While large language models (LLMs) show potential for traffic-related queries, they also face issues such as hallucinations, necessitating approaches like fine-tuning and retrieval-augmented generation to improve accuracy and context-awareness.

This study addresses these challenges by tackling each phase separately. The first phase focuses on improving text relevancy, resolving the limitations of prior research that predominantly relied on word-level embeddings for classification. By incorporating character, word, sentence, and concept embeddings, this approach aims to enhance classification accuracy through better contextual understanding. Additionally, advanced deep learning models will be employed to overcome the shortcomings of traditional CNNs and RNNs, further boosting classification performance.

The second phase involves extracting events from classified text and constructing a dynamic KG, which will serve as domain-specific knowledge for LLMs. This KG will improve the efficacy and reliability of LLMs in generating accurate responses, particularly for traffic forecasting. By leveraging multiple approaches to knowledge induction, this phase aims to significantly enhance the performance of LLMs in this domain.

Finally, the third phase optimizes traffic forecasting using novel algebraic methods, addressing the temporal and spatial dimensions of traffic data. These methods, which have not been previously applied in traffic forecasting, aim to offer innovative solutions for improving prediction accuracy and handling the complexities of traffic data.

3 TIFFNLP METHODOLOGY

This methodology aims to enhance traffic forecasting from English social media text. Ensuring accurate identification of traffic-related information, addressing challenges such as language irregularities and contextual nuances in social media texts by employing advanced word embeddings and deep learning models for accurate classification and event detection. The first phase focuses on text relevancy, event detection, and location detection, In the Second phase, a robust solution is provided for predicting. The architecture is represented in Figure 1.

3.1 Text Relevancy

The objective of the first step is to use multi-level embeddings to classify social media messages focusing on the character, word, and concept of the textual data. Focusing on the challenges previously faced in SM traffic relevancy such as 1) data sparsity 2) data Imbalance and 3) Ambiguity due to SAB terms. This section will use a transformer-based model, BERT, to capture contextual semantics effectively through its multi-headed attention mechanism and compared tit o embeddings like word2Vec, FastText, and GloVe to evaluate their efficacy. Additionally, we will explore the application of CNN, and RNN, including Temporal Convolutional Networks (TCN) and LSTM. The combination of character, word, and concept embeddings in the model boosts classification accuracy by capturing both detailed linguistic Features. Character embeddings handle variations in spelling and abbreviations, while word embeddings provide contextual meaning. Concept embeddings disambiguate words with multiple meanings based on context. Together, they improve the model's ability to classify informal, nuanced social media text. To address the challenges of text irregularities and data imbalance, adversarial networks and transfer learning

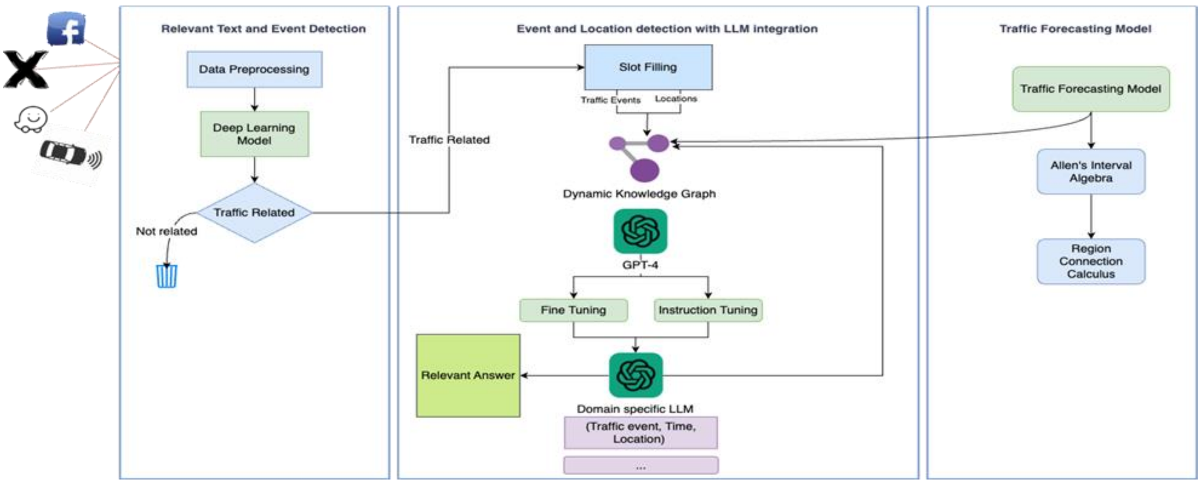

Figure 1: TIFFNLP Methodology.

techniques will be employed, leveraging pre-trained models to improve classification performance. Furthermore, the integration of domains from external sources such as Probase will enrich semantic representation at both word and concept levels.

3.2 Event Detection

This phase addresses the challenges of relying on predefined classes that may overlook emerging events and the difficulty of structuring relevant information from unstructured data. To overcome these challenges, Dynamic Knowledge Graph Embedding (DKGE) (Wu et al., 2022) and slot filling (Yang et al., 2021) to dynamic embedding and clustering techniques will be used. The approach begins with the dynamic population of a KG using an online learning approach to capture Spatio-temporal events. This KG is tailored to the traffic domain through ontology integration, enhancing its ability to support Question Answering (QA) via an LLM.

Addressing challenges such as geo-ambiguities and unseen place names due to limited context, as well as the informal features of tweets. Geocoding enhancements involve clustering methods to group tweets of the same topic, expanding context and improving geocoding accuracy. Integration with LLMs and global gazetteers further enhances geocoding by considering entity co-occurrence within Twitter networks, and overall traffic forecasting capabilities.

3.3 Traffic-Domain LLM

The last phase is to enhance the reliability and accuracy of LLMs within the traffic domain by addressing their limitations through techniques such as fine-tuning, Retrieval Augmented Generation, and instruction tuning. This methodology aims to develop a traffic-specific model capable of generating accurate and contextually relevant responses.

Fine-tuning adapts pre-trained LLMs to specific tasks using a labelled dataset. The process involves selecting a task, preprocessing the dataset, experimenting with models, fine-tuning the best one, and evaluating its performance.

Instruction tuning further refines LLMs to follow specific human instructions, enhancing model controllability and predictability by extracting instruction-output pairs from annotated datasets and to generate outputs for specific instructions. The base model is then refined using the constructed instruction dataset.

RAG enhances LLMs by integrating external knowledge from a dynamically constructed knowledge graph, improving response accuracy, real-time relevance, and explainability. The process involves: 1) curating external sources, 2) retrieving context-relevant data, and 3) integrating it with the LLM for response generation. Fine-tuning prepares the model, while instruction tuning refines it for predictable outputs, enabling RAG to deliver more precise, context-aware results.

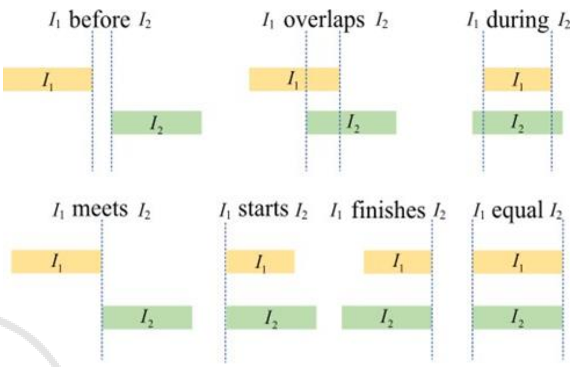

Figure 2: Allen's Intervals to represent temporal logic.

3.4 Traffic Forecasting

By integrating these techniques, the methodology significantly improves the accuracy, reliability, and contextual relevance of LLMs for domain-specific applications. Ensuring that LLMs generate accurate responses but also adjust dynamically to user instructions and external Knowledge graphs. The final phase aims to forecast traffic events using a KG) enriched with historical traffic data and external factors. This phase focuses on representing and analysing temporal relationships to enhance traffic prediction accuracy. The process begins by collecting historical traffic data, including location coordinates, timestamps, traffic flow, and road structure information from previous phases, as well as external data such as weather conditions and driver behaviours. These data points are represented as time intervals [start, end] to capture the temporal duration of events.

The methodology includes the following steps: First begins by setting up the interval representation and relationships, which includes representing traffic events as time intervals, then using Allen's Interval Algebra in Figure 2. Allen's 13 interval relationships understand how different time intervals interact. This helps analyse the temporal relations between events (Allen, 1983).

Once the events have been represented with Allens intervals, a feature matrix to capture interval relationships will be constructed, ensuring that reasoning paths are temporally consistent. This will allow us to model the relationships between traffic events using LSTM and utilize GNN (Graph Neural Network) with Attention Mechanism: Integrate semantic correlations between potentially distant roads to improve prediction accuracy. This DL model will be trained to forecast missing entities or relationships within specified time intervals on Temporal Knowledge Graphs (TKGs). Therefore, by starting with initial events such as a traffic incident at a specific location and time interval, queries to predict potential concerns, such as increased traffic congestion following a known event can be calculated using the reasoning algebra through the analysis of interval relationships and identify a reasoning path connecting the initial event to the predicted outcome within the specified time interval, then through a trained model to predict new potential concerns based on the given query. For example, predict traffic concerns at Location A based on a prior incident.

By combining spatiotemporal data, knowledge graphs, interval algebra, and advanced machine learning techniques, this methodology can develop a robust framework for predicting traffic concerns. This approach addresses the complexity and the evaluation of TIFFNLP will be conducted through the performance evaluation by comparing the output of the TIFFNLP with the model generated manually.

4 SYSTEM EVALUATION

The evaluation of TIFFNLP will be conducted through the performance evaluation by comparing the output of the TIFFNLP with the model generated manually. For this purpose, different case studies from different domains have been used. The purpose of the system evaluation is to assess the text relevancy, detected events, locations, forecasted events, and LLM generated text concerning its semantic quality measured by semantic conformance with accuracy and completeness. Furthermore, this section aims to answer the following questions: How can the use of words enhance the accuracy and relevance of data embeddings and deep learning models, addressing the dynamic nature of social media data, be developed for detection and classification? How can models be developed to accurately interpret and classify informal language, including slang, and abbreviations to improve the precision of traffic detection mechanisms in social media? How can Allen's interval algebra and region connection calculus be used to improve traffic sourced from social media and finally, how accurate the results of TIFFNLP Framework compared to previous state of the art studies?

In evaluating the TIFFNLP model, a diverse dataset, rich in real-world traffic events, will be utilized. The dataset comprises traffic incident reports sourced from social media, particularly Twitter, encompassing various details such as date, time, city, location, latitude, longitude, accuracy, direction, event type, lanes blocked, vehicles involved, tweet content, and source. For instance, the dataset includes incidents like a vehicular accident in Pasig City at Ortigas Emerald, involving a taxi and motorcycle, with the tweets" MMDA ALERT: Vehicular accident at Ortigas Emerald EB involving taxi and MC as of 7:55 AM. 1 lane occupied." This sample illustrates the dataset's capacity to provide detailed and varied traffic scenarios. Our proposed method is under development and Figure 3 displays the output of the traffic forecasting step.

To assess the performance of the TIFFNLP model, we employ three standard evaluation metrics: Precision (P), Recall (R), and F- measure. These metrics provide a comprehensive evaluation of the model's ability to accurately identify and extract relevant traffic events from social media data.

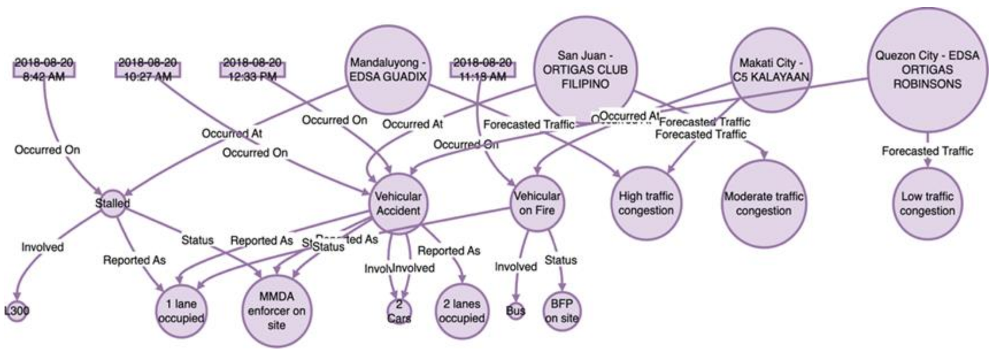

Figure 3: Sample intermediate diagram produced by TIFFNLP (Forecast event).

Precision measures the proportion of positive identifications made by the model out of all positive cations it made. In this context, it is the fraction of traffic events correctly identified by TIFFNLP out of all events labeled as relevant. It is calculated using the formula:

$$\text{Precision} = \frac{TP}{TP+FP} \qquad (1)$$

Recall assesses the fraction of actual relevant traffic events that the model successfully identifies. It is the proportion of true positive identified by the model out of all actual positive cases present in the dataset. Recall is calculated using the formula:

$$Recall = \frac{TP}{TP+FN} \qquad (2)$$

F-measure provides a mean of Precision and Recall, offering a single metric that balances both concerns. The F-measure is calculated using the formula:

$$\text{F}-\text{measure} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision}+\text{Recall}} \qquad (3)$$

When an incident, like a taxi and motorcycle accident at Ortigas Emerald, Pasig City, is reported, it is logged with details such as location, time, and vehicles involved, the KG is adjusted dynamically to ease congestion to predict and manage future traffic disruptions. In evaluating TIFFNLP's performance, we focus on its ability to accurately extract, and model key elements related to traffic events using deep learning, Online learning, and NLP techniques.

Table 1 outlines the number of elements identified by experts versus TIFFNLP, describing any discrepancies. (P) Precision, Recall (R), and F-measure metrics are then calculated using formulas tailored to count M (true positives), N (false positives), and K (false negatives). These metrics provide a quantitative assessment of TIFFNLP's accuracy. Additionally, similar evaluation tables can be structured to assess the detection of attributes, methods, and relationships to evaluate across all phases.

Table 1: Precision, Recall, and F-measure results.

	Model		The value of			Evaluation metrics		
Case study	Human	TIFFNLP	M	N	K	P	R	F-Measure
1	13	16	13	16	13	0.45	0.50	0.47
...	…	…	…	…	…	…	…	…
n	12	9	10	9	12	0.53	0.45	0.49
					Average:	0.47	0.48	0.47

5 CONCLUSION

The TIFFNLP framework advances traffic forecasting by integrating text relevancy, event detection, location detection, and predictive modelling. It leverages NLP and deep learning to classify and predict traffic information from social media. The three-phase framework addresses text relevancy with transformer models, improves event and location detection using slot filling and knowledge graphs, and enhances traffic forecasting with interval algebra and spatial reasoning. Evaluated using Precision, Recall, and F-measure, TIFFNLP offers valuable insights for urban planning, authorities, and the public, providing a comprehensive approach to traffic management.

REFERENCES

Allen, J. F. (1983). Maintaining Knowledge about Temporal Intervals.

Azhar, A., Rubab, S., Khan, M. M., Bangash, Y. A., Alshehri, M. D., Illahi, F., & Bashir, A. K. (2023). Detection and prediction of traffic accidents using deep learning techniques. Cluster Computing, 26(1), 477–493. https://doi.org/10.1007/s10586-021-03502-1

Babbar, S., & Bedi, J. (2023). Real-time traffic, accident, and potholes detection by deep learning techniques: a modern approach for traffic management. Neural Computing and Applications, 35(26), 19465–19479.

Bai, L., Yao, L., Wang, X., & Wang, C. (2020). *Adaptive Graph Convolutional Recurrent Network for Traffic Forecasting*.

Balaguer, A., Benara, V., Cunha, R. L. de F., Filho, R. de M. E., Hendry, T., Holstein, D., Marsman, J., Mecklenburg, N., Malvar, S., Nunes, L. O., Padilha, R., Sharp, M., Silva, B., Sharma, S., Aski, V., & Chandra, R. (2024). RAG vs Fine-tuning: Pipelines, Tradeoffs, and a Case Study on Agriculture.

Bok, K., Kim, I., Lim, J., & Yoo, J. (2023). Efficient graph-based event detection scheme on social media. Information Sciences, 646.

Chang, H., Li, L., Huang, J., Zhang, Q., & Chin, K. S. (2022). Tracking traffic congestion and accidents using social media data: A case study of Shanghai. Accident Analysis and Prevention, 169.

Chen, Y., & Chen, X. (Michael). (2022). A novel reinforced dynamic graph convolutional network model with data imputation for network-wide traffic flow prediction. Transportation Research Part C: Emerging Technologies, 143.

Chuckravanen, D., Daykin, J. W., Hunsdale, K., Seeam, A., & Business School, W. (2017). Allen's Interval Algebra and Smart-type Environments. www.iaria.org

Dabiri, S., & Heaslip, K. (2019). Developing a Twitter-based traffic event detection model using deep learning architectures. Expert Systems with Applications, 118, 425–439. https://doi.org/10.1016/j.eswa.2018.10.017

Das, R. D., & Purves, R. S. (2020). Exploring the Potential of Twitter to Understand Traffic Events and Their Locations in Greater Mumbai, India. IEEE Transactions on Intelligent Transportation Systems, 21(12), 5213–5222.

Fan, W., Ding, Y., Ning, L., Wang, S., Li, H., Yin, D., Chua, T.-S., & Li, Q. (2024). A Survey on RAG Meeting LLMs: Towards Retrieval-Augmented Large Language Models. http://arxiv.org/abs/2405.06211

Fontes, T., Murcos, F., Carneiro, E., Ribeiro, J., & Rossetti, R. J. F. (2023). Leveraging Social Media as a Source of Mobility Intelligence: An NLP-Based Approach. IEEE Open Journal of Intelligent Transportation Systems, 4, 663–681. https://doi.org/10.1109/OJITS.2023.3308210

Hodorog, A., Petri, I., & Rezgui, Y. (2022). Machine learning and Natural Language Processing of social media data for event detection in smart cities. Sustainable Cities and Society, 85. https://doi.org/10.1016/j.scs.2022.104026

Houlsby, N., Giurgiu, A., Jastrze¸bski, S. J., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., & Gelly, S. (n.d.). Parameter-Efficient Transfer Learning for NLP. https://github.com/google-research/

Lu, H., Huang, D., Song, Y., Jiang, D., Zhou, T., & Qin, J. (2020). St-trafficnet: A spatial-temporal deep learning network for traffic forecasting. Electronics (Switzerland), 9(9), 1–17. https://doi.org/10.3390/electronics9091474

Ma, W., Wu, D., Sun, Y., Wang, T., Liu, S., Zhang, J., Xue, Y., & Liu, Y. (2024). Combining Fine-tuning and LLM-based Agents for Intuitive Smart Contract Auditing with Justifications.

Mihindukulasooriya, N., Tiwari, S., Enguix, C. F., & Lata, K. (2023). Text2KGBench: A Benchmark for Ontology-Driven Knowledge Graph Generation from Text. http://arxiv.org/abs/2308.02357

Nirbhaya, M. A. W., & Suadaa, L. H. (2023). Traffic Incident Detection in Jakarta on Twitter Texts Using a Multi-Label Classification Approach. Proceedings - 2023 10th International Conference on Computer, Control, Informatics and Its Applications: Exploring the Power of Data: Leveraging Information to Drive Digital Innovation, IC3INA 2023, 290–295. https://doi.org/10.1109/IC3INA60834.2023.10285731

Randell, David. A, Zhan Cui, & Anthony G. Cohn. (1992). A spatial logic based on regions and connection. KR, 92.

Suat-Rojas, N., Gutierrez-Osorio, C., & Pedraza, C. (2022). Extraction and Analysis of Social Networks Data to Detect Traffic Accidents. Information (Switzerland), 13(1). https://doi.org/10.3390/info13010026

Sun, X., Liu, L., Ayorinde, A., & Panneerselvam, J. (2021). ED-SWE: Event detection based on scoring and word embedding in online social networks for the internet of people. Digital Communications and Networks, 7(4), 559–569. https://doi.org/10.1016/j.dcan.2021.03.006

Tao, L., Xie, Z., Xu, D., Ma, K., Qiu, Q., Pan, S., & Huang, B. (2022). Geographic Named Entity Recognition by Employing Natural Language Processing and an Improved BERT Model. ISPRS International Journal of Geo-Information, 11(12). https://doi.org/10.3390/ijgi11120598

Wang, M., Pang, A., Kan, Y., Pun, M.-O., Chen, C. S., & Huang, B. (2024). LLM-Assisted Light: Leveraging Large Language Model Capabilities for Human-Mimetic Traffic Signal Control in Complex Urban Environments. http://arxiv.org/abs/2403.08337

Wu, T., Khan, A., Yong, M., Qi, G., & Wang, M. (2022). Efficiently embedding dynamic knowledge graphs. Knowledge-Based Systems, 250. https://doi.org/10.1016/j.knosys.2022.109124

Yang, X., Bekoulis, G., & Deligiannis, N. (2021). Traffic Event Detection as a Slot Filling Problem. http://arxiv.org/abs/2109.06035

Yang, X., Yan, J., Cheng, Y., & Zhang, Y. (2023). Learning Deep Generative Clustering via Mutual Information Maximization. IEEE Transactions on Neural Networks and Learning Systems, 34(9), 6263–6275. https://doi.org/10.1109/TNNLS.2021.3135375

Yin, F., Ye, X., & Durrett, G. (2024). LoFiT: Localized Fine-tuning on LLM Representations. http://arxiv.org/abs/2406.01563

Yuan, Z., Liu, H., Liu, J., Liu, Y., Yang, Y., Hu, R., & Xiong, H. (2021). Incremental spatio-temporal graph learning for online query-poi matching. The Web Conference 2021 - Proceedings of the World Wide Web Conference, WWW 2021, 1586–1597.

Zhang, H., Liu, Z., Xiong, C., & Liu, Z. (2019). Grounded Conversation Generation as Guided Traverses in Commonsense Knowledge Graphs.

Zhang, S., Zhu, K., & Zhang, W. (2023). Multivariate Correlation Matrix-Based Deep Learning Model With Enhanced Heuristic Optimization for Short-Term Traffic Forecasting. IEEE Transactions on Knowledge and Data Engineering, 35(3), 2847–2858.

Zheng, G., Chai, W. K., Duanmu, J. L., & Katos, V. (2023). Hybrid deep learning models for traffic prediction in large-scale road networks. Information Fusion, 92, 93–114. https://doi.org/10.1016/j.inffus.2022.11.019

Zhou, S., Thomas Ng, S., Huang, G., Dao, J., & Li, D. (2022). Extracting interrelated information from road-related social media data. Advanced Engineering Informatics, 54.