Uncertainty-aware framework for real-time traffic incident prediction

Thanh Tran[a], Dan He[a], Jiwon Kim[a,*] and Mark Hickman[a]

[a] School of Civil Engineering, The University of Queensland, St. Lucia, QLD 4072, Australia
thanh.tran@uq.edu.au
[*] Corresponding author

Extended abstract submitted for presentation at the Conference in Emerging Technologies in Transportation Systems (TRC-30)
September 02-03, 2024, Crete, Greece

April 28, 2024

Keywords: Real-time Traffic Incident Prediction, Explainable AI (XAI), Uncertainty Estimation

1 INTRODUCTION

The advent of advanced traffic data collection technologies, such as loop detectors and GPS devices, has revolutionized traffic management by enabling the generation of vast data amounts. This has facilitated the development of sophisticated machine learning (ML) models for real-time traffic incident prediction. These models, which are trained on historical traffic data, play a pivotal role in forecasting traffic incidents. However, the inherent variability and unpredictability of traffic dynamics and incident patterns necessitate that these models not only predict incidents with high accuracy but also gauge the certainty of their predictions. In critical safety applications, uncertainty estimation is vital, enabling traffic management authorities to measure the confidence levels of predictions and make informed decisions. This capability is essential for effective decision-making in road safety and for implementing these models in production environments. A particularly relevant use case for uncertainty estimation arises during adverse and rare event conditions, such as severe weather, which significantly alters road conditions and driver behavior. For example, if a model predicts a high risk of accidents on a specific road link during a storm but with notable uncertainty, traffic managers might choose to issue a general caution instead of a full closure. This informed strategy optimizes safety without unnecessary traffic disruption by considering both the predicted risks and the confidence levels of those predictions.

Research in traffic incident prediction increasingly utilizes Machine Learning (ML), especially Deep Learning techniques such as Graph Neural Networks (GNNs), due to their efficacy in handling complex traffic data patterns. GNNs are well-suited for modeling the graph structures of road networks, a strength demonstrated by Wang *et al.* (2021a)'s GSNet for regional risk assessment and further applied in city-wide risk evaluations (Wang *et al.*, 2021b). Our previous work advances these models through the Multi-structured Graph Neural Network (MSGNN) (Tran *et al.*, 2023), which integrates diverse data sources for 'network-wide' predictions. However, the dynamic nature of traffic conditions and behaviors poses a challenge to the reliability of ML models trained on static datasets, potentially leading to obsolete or overly confident predictions in real-world critical applications. Thus, ensuring the reliability and timely applicability of these models is critical, mirroring trends in safety-critical fields like healthcare (Dolezal, 2022) and autonomous driving (Dong *et al.*, 2023), where uncertainty quantification and explainable AI are key. Motivated by the necessity for reliable Traffic Incident Prediction (TIP), we diverge from existing works that emphasize model capability, focusing instead on reliability through our novel

Uncertainty Aware Traffic Incident Prediction (UATIP) framework. This enriches predictive models with an Uncertainty Estimation (UE) ability, enhancing real-time interpretability and reliability—a facet not thoroughly addressed in current research. UE aids in discerning highly confident predictions from those warranting caution, thus bolstering the practical utility of traffic prediction models. Our experiments on MSGNN with varied real-world data highlight the UE's role in augmenting prediction reliability. This advancement portends significant contributions to Explainable AI (Dong *et al.*, 2023) in TIP and supports the application of Trustworthy Transfer Learning (TTL) (Shen *et al.*, 2023), promising enhanced accuracy and robustness in a diverse array of traffic scenarios. Further details and results will be presented in the full paper.

2 METHODOLOGY

UATIP's Definition: UATIP involves extending the predictive capability of any ML incident prediction model to not only forecast traffic incidents but also estimate the uncertainty associated with each prediction. Formally, UATIP is defined as follows:

$$UATIP : (\mathcal{X}_l^t, \text{Features}) \rightarrow \{(y_l^t, u_l^t) | y_l^t \in \{0, 1\}\}, \tag{1}$$

where $\mathcal{X}_l^t$ represents any form of data input (e.g., traffic graphs, vectors, images) corresponding to a location l at time t, and u_l^t is the uncertainty measure associated with the prediction y_l^t. This measure quantifies the model's confidence in its prediction, providing a nuanced interpretation of the predictive output that aids traffic management authorities in making informed decisions based on both the prediction and its reliability. The challenge lies in the precise estimation of this uncertainty, considering the dynamic and complex nature of traffic patterns and incident determinants at different locations. In our experiments, we apply UATIP to **Sub-area Traffic Incident Prediction** (STIP) (Tran *et al.*, 2023) where l is defined as a sub-area. Given a large traffic network, let S_c denote a set of all sub-areas with radius c within the network. For each sub-area $s_j \in S_c$, and at any given time t, consider m different sources of data that create m subgraphs $\mathcal{G}_{s_j}^t = \{G_{s_j}^1, G_{s_j}^2, \ldots, G_{s_j}^m\}^t$, each representing one of the m data sources pertinent to sub-area s_j and timestamp t. STIP aims to process this set of subgraphs $\mathcal{G}_{s_j}^t$, along with time-variant traffic features (e.g., flow, occupancy, speed) from the period $[t - T_b, t]$ and time-invariant features (e.g., link characteristics, day of the week, public holidays), to generate a binary output indicator $y_{s_j}^t$. This indicator is 1 if at least one traffic incident occurs within s_j during the time period $[t + T_p, t + T_p + T_d]$ and 0 otherwise, where T_p is the prediction horizon and T_d is the width of the prediction window. With this application context, our proposed model is briefly illustrated in Figure 1. Particularly, given any trained incident predictor $\mathbf{f}()$, which was trained on X_{train} including i different training incident and non-incident cases in total, UATIP takes input as conventional TIP models. For simplicity, we assume the use of only loop detector data sources; hence, the input would be $G_{s_j}^t$ with respect to UATIP's definition. UATIP extracts feature (encoded) vector $\mathbf{f}(G_{s_j}^t)$ from the traffic graph input $G_{s_j}^t$ using $\mathbf{f}()$'s last hidden layer. Subsequently, $\mathbf{f}(G_{s_j}^t)$ is used as input to a K-Nearest Neighbors (K-NN) algorithm, which assesses the proximity of the input's feature vector $\mathbf{f}(G_{s_j}^t)$ to the training set X_{train}, given the feature vectors $\mathbf{f}(X_{\text{train}})$. Although different distance measures can be employed, K-NN computes K minimum Euclidean distances $d(i)$ between feature vector $\mathbf{f}(G_{s_j}^t)$ and feature vectors $\mathbf{f}(X_{train})$. AGGREGATION then combines (SUM or MEAN) these K distances into an uncertainty score, informing the certainty level $(u_{s_j}^t)$ of the prediction for the given input $G_{s_j}^t$. Finally, this score can be used to determine whether the current prediction is confident, given a decision threshold.

3 PRELIMINARY RESULTS

How UATIP performs in estimating uncertainty? We applied the MSGNN incident predictor, as delineated in Tran *et al.* (2023), using traffic data from loop detectors in Brisbane

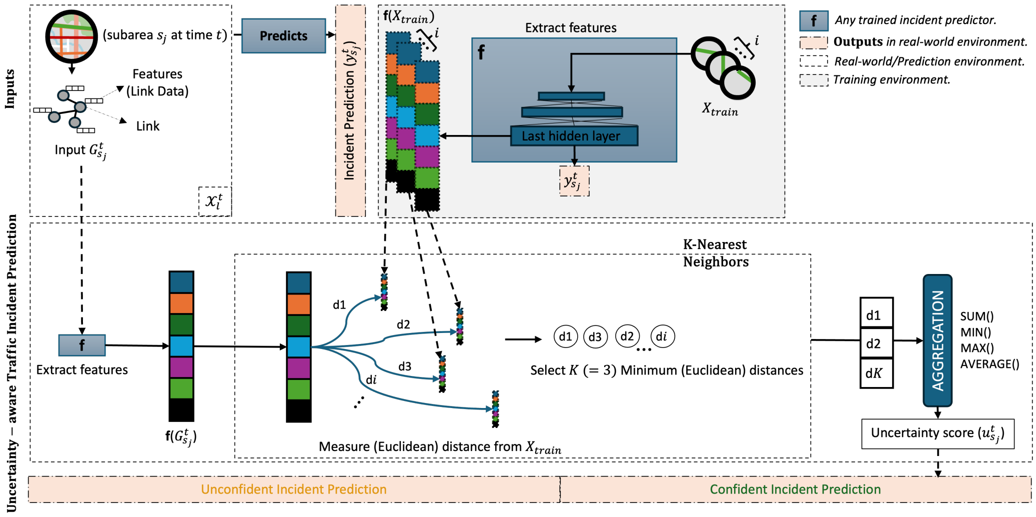

Figure 1 – *Overall framework of Uncertainty-aware Traffic Incident Prediction*

(BRIS) and Gold Coast (GC), Australia, in 2017. Our model training utilized a 15-minute prediction horizon (T_p) and a 5-minute prediction window (T_d), with X_{train} comprising 758 cases and X_{test} with 324 cases, split in a 7:3 ratio as instructed by Yu *et al.* (2021), Tran *et al.* (2023). A 'missing rate' **F** simulated sensor outages by introducing '-1' to denote missing data for link data (i.e., node features of inputs G), reflecting real-world data incompleteness. UATIP's discrimination between prediction certainties was tested under varied data availabilities and regional differences between Brisbane and Gold Coast networks. To test UATIP, we needed two types of datasets: uncertain and certain datasets. We used inputs from X_{test} and labeled them as 'certain predictions' $X_{certain}$ since X_{test} is in-distribution data with respect to X_{train}. To construct uncertainty testing datasets, we labeled 324 instances as one uncertainty dataset $X_{uncertain}$, including inputs from X_{test} but at missing rates of $F = 10\%$, and similarly, two more uncertainty sets $X_{certain}$ with $F = 50\%; 80\%$. Another uncertainty dataset was constructed by obtaining 324 cases from Gold Coast, used to represent inter-regional uncertainty, indicative of disparate network structures and traffic behaviors. The model's capability to accurately detect uncertain from certain predictions for inputs from $X_{certain}$ and $X_{uncertain}$ is demonstrated in Figure 2. Particularly, across low to high uncertainty scenarios, the model consistently identified actual certain predictions from $X_{certain}$ with lower uncertainty scores and uncertain predictions with higher scores (closer to 1.0). The inter-regional uncertainty analysis revealed comparable distribution patterns, indicating that the model trained on BRIS data may be generalizable to GC data, suggesting its broad utility across different regions, which can be explained by the cross-region applicability of MSGNN.

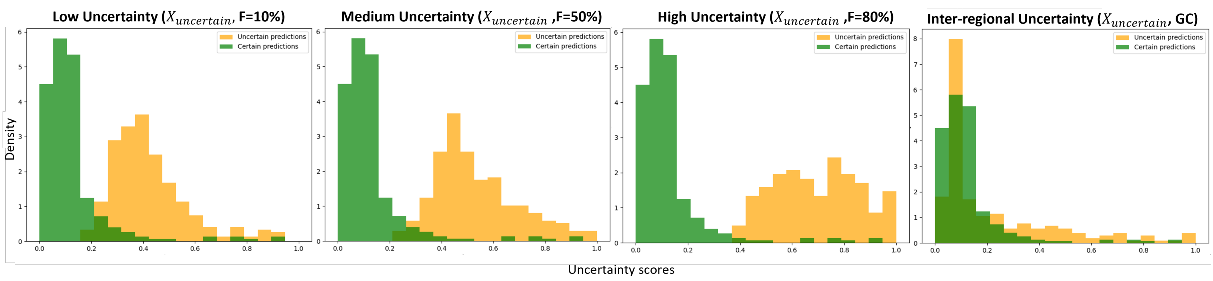

Figure 2 – *UATIP's uncertainty scores distribution by* ***true labels*** *under uncertainty scenarios.*

UATIP was employed to evaluate the uncertainty of various predictions. The predictor $f()$ extracted high-dimensional feature vectors from X_{train}, $X_{certain}$, and $X_{uncertain}$. Utilizing t-

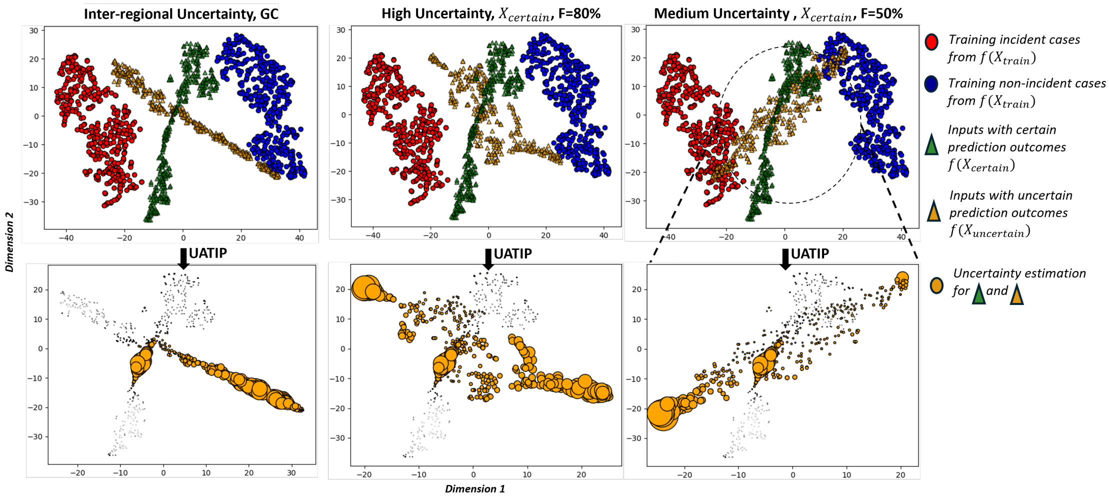

Figure 3 – *2-Dimension visualisation of uncertainty scores for multiple inputs.*

distributed Stochastic Neighbor Embedding (t-SNE) for dimensionality reduction, these feature vectors were visualized in a 2D space as shown in Figure 3. In this figure, the top plots show the training data with red and blue dots for incidents and non-incidents, respectively, while predictions are marked with green and orange triangles for certain and uncertain outcomes. Following UATIP's processing, the bottom plots display predicted uncertainties with larger orange circles indicating higher uncertainty. Notably, these larger circles correlate with the orange triangles, identifying inputs that lead to uncertain predictions. This visual indicator is especially valuable when a deployed model, denoted as f, records a significant number of uncertain predictions, marked by many large orange circles, within a specific timeframe such as a day, surpassing a pre-established threshold. This could then serve as a prompt for selective retraining.

4 DISCUSSION & FUTURE RESEARCH

Results with MSGNN demonstrate UATIP's ability to effectively estimate uncertainty scores across various settings, as shown in Figure 2. This highlights its potential as an interpretive tool in large-scale incident prediction, distinguishing high-certainty predictions from those requiring cautious interpretation, as depicted in Figure 3. In the future, it could enhance the model's efficiency by advocating selective retraining of highly uncertain cases rather than all inputs.

References

Dolezal, James M. et al. 2022. Uncertainty-informed deep learning models enable high-confidence predictions for digital histopathology. *Nature Communications*.

Dong, Jiqian, Chen, Sikai, Miralinaghi, Mohammad, Chen, Tiantian, Li, Pei, & Labi, Samuel. 2023. Why did the AI make that decision? Towards an explainable artificial intelligence (XAI) for autonomous driving systems. *Transportation Research Part C: Emerging Technologies*.

Shen, Maohao, Bu, Yuheng, Sattigeri, Prasanna, Ghosh, Soumya, Das, Subhro, & Wornell, Gregory. 2023. Post-hoc uncertainty learning using a dirichlet meta-model. *In: AAAI*.

Tran, Thanh, He, Dan, Kim, Jiwon, & Hickman, Mark. 2023. MSGNN: A Multi-structured Graph Neural Network model for real-time incident prediction in large traffic networks. *Transportation Research Part C: Emerging Technologies*.

Wang, Beibei, Lin, Youfang, Guo, Shengnan, & Wan, Huaiyu. 2021a. GSNet: Learning Spatial-Temporal Correlations from Geographical and Semantic Aspects for Traffic Accident Risk Forecasting. *In: AAAI*.

Wang, Senzhang, Zhang, Jiaqiang, Li, Jiyue, Miao, Hao, & Cao, Jiannong. 2021b. Traffic Accident Risk Prediction via Multi-View Multi-Task Spatio-Temporal Networks. *IEEE Transactions on KDE*.

Yu, Le, Du, Bowen, Hu, Xiao, Sun, Leilei, Han, Liangzhe, & Lv, Weifeng. 2021. Deep spatio-temporal graph convolutional network for traffic accident prediction. *Neurocomputing*.