2024 International Conference on Intelligent and Innovative Computing Applications (ICONIC)
7th-9th November 2024, Pearle Beach Resort & Spa, Flic en Flac, Mauritius
https://doi.org/10.59200/ICONIC.2024.016

Spatio-Temporal Modelling of Road Traffic Fatalities in the Western Cape

H. Gool.
Tanzanite Management Services
Cape Town, South Africa
haroong@tanzanite-ms.com

B. Lindner
Spatial Edge
Stellenbosch, South Africa
brian@spatialedge.co.za

L.E. Burger
Stellenbosch University
Stellenbosch, South Africa
eldonburger@sun.ac.za

Abstract
A machine learning model capable of predicting the probability of road fatal events in both time and space is developed. By aggregating relevant features of the Western Cape into an H3 grid, the model learns patterns of fatal events. Traditional machine learning and deep learning techniques are employed to understand the relationship between aggregated features and road fatal events. The models are compared against each other and against historical average models currently used in the industry. This study represents the first attempt to use machine learning techniques to model road traffic fatalities in South Africa, specifically in the Western Cape.

Keywords: Road Traffic Fatalities, Spatio-Temporal Modeling, Machine Learning, South Africa

1. Introduction

Road traffic accidents claim approximately 1.35 million lives annually, cause injuries to an estimated 50 million people, and are the leading cause of death for people aged 5 to 29 years old (*Global Status Report on Road Safety 2018*, 2018). To reduce road traffic fatalities, the World Health Organization, the United Nations Regional Commissions, and other stakeholders compiled a global plan which aims to reduce road fatalities by at least 50% by 2030 (Global Plan Decade of Action for Road Safety 2021-2030, 2021).

To achieve a drastic reduction in road traffic fatalities, low- and middle-income countries must be prioritized. Ninety percent of road traffic fatalities currently occur in low- and middle-income countries, even though less than 60% of the world's motor vehicles are found in these countries (Global Plan Decade of Action for Road Safety 2021-2030, 2021).

South Africa has managed to reduce the road fatality rate per 100,000 inhabitants from 24.72 in 2017 to 19.39 in 2023 (State of Road Safety in South Africa: January 2023 to December 2023, 2024). Despite this decrease, the current road fatality rate in South Africa is still higher than the global average of 15 road fatalities per 100,000 inhabitants and higher than the African continent's road fatality rate of 19 (State of Road Safety in South Africa: January 2023 to December 2023, 2024).

As a participant in the global road safety action plan, South Africa has undertaken to reduce road fatalities by at least 50% by 2030. This means that the annual number of road fatalities in South Africa needs to decrease to 6,984 by 2030 from the current 11,883 annual road fatalities (State of Road Safety in South Africa: January 2023 to December 2023, 2024). Similar to other low- and middle-income countries, the main challenge in reducing road fatalities in South Africa is funding (State of Road Safety in South Africa: January 2023 to December 2023, 2024).

Given limited funding, innovative solutions are required that optimally use the limited resources available. The emergence of big data and the development of artificial intelligent algorithms offer promising solutions to predict road traffic accidents across time and space (Bao et al., 2019). By anticipating areas and timeframes with high road fatality risks, local authorities can implement targeted safety measures and improve post-accident response.

Traditional approaches used to predict road traffic fatalities primarily rely on historical accident data to forecast future incidents over the long term. Apart from poor generalization, long-term predictions considerably limit the potential for intervention, as it is unclear when events are likely to occur. Machine learning-based approaches can predict road traffic fatalities across space in the short term by integrating multiple data sources and learning complex patterns. Machine learning-based road fatality prediction models have historically been applied on small road segments, or, when evaluated on extensive road networks, it is commonly applied in the context of high-income countries.

In this study, a machine learning-based road fatality prediction model is developed for the Western Cape province, one of the nine provinces of South Africa. Road fatality event data is integrated with road infrastructure data and weather data across time and space. The aggregated data is used to develop and compare the appropriateness of multiple vector-based and raster-based models to predict road fatalities across time and space. The developed approach outperforms historical average models currently used in the industry. Unlike previous approaches, data availability is explicitly considered in the design of the approach to enable countries to deploy the approach using limited data.

The paper proceeds as follows. Related approaches used to predict road traffic accidents are discussed in Section 2. Section 3 outlines the steps followed to develop a road traffic fatality model for the Western Cape. The performance of the approach outlined is discussed in Section 4. The practical considerations of using the model to reduce road fatalities are discussed in Section 5, followed by the conclusion in Section 6.

2. Related Work

Road traffic accidents have been defined, using various terms and definitions (Al-Hasani, 2021). Peden et al. (2004) differentiate between a road traffic crash, a road traffic accident and a road traffic fatality. Road traffic crashes encompass all events which involve at least one motor vehicle, road traffic accidents include all crashes where at least one injury or fatality occurred, and road traffic fatalities refer to events resulting in death within 30 days of the crash (Peden et al., 2004). In South Africa, road traffic accidents are only considered road traffic fatalities when a victim(s) dies within six days of the event (Arrive Alive, n.d.).

When considering a predictive modeling approach, the target feature must be clearly defined with explicit consideration of data availability. While predicting all road traffic accidents offers additional benefits, it requires the collection and confirmation of all traffic accident data which might be infeasible in a resource-constrained environment. Even in unconstrained environments, collecting all traffic accident events can be challenging since less severe traffic accidents are often not reported (Abdulhafedh, 2017). In Africa, road traffic accidents tend to be underreported (Road Transport Accident Deaths in South Africa, 2007-2019, 2024). Data collection in South Africa is further complicated by the involvement of multiple stakeholders (State of Road Safety in South Africa: January 2023 to December 2023, 2024). Given these challenges, focusing on road traffic fatalities is more appropriate in countries with limited resources.

The timing of road traffic fatality data collection should also be considered when determining an appropriate modeling approach. Models based on sequential data are often recommended. For instance, Ren et al. (2018) proposed using a long short-term memory (LSTM) artificial neural network. This method however requires a sequence of recent confirmed traffic accidents which is impractical considering that the definition of a road traffic fatality often includes deaths that occur up to 30 days after the event (Peden et al., 2004).

While road traffic fatality data might be more readily available than road traffic accident data, the number of road traffic fatality events is significantly lower compared to road accidents or non-events. Events that occur at a significantly lower rate compared to common events are referred to as rare events (Lazarević et al., 2004) and introduce several modeling challenges that must be explicitly considered (Weiss, 2004). For instance, measuring the performance of a road traffic prediction model using accuracy is inappropriate, as it is straightforward to maximize accuracy by simply predicting non-events. Maalouf & Trafalis (2011) suggested using sampling techniques to artificially create a more balanced training dataset, thereby reducing the likelihood of the model simply predicting the majority event to optimize performance. Similarly, the loss function used to optimize the parameters of error-based models can also be adjusted by assigning larger weights to the minority class.

Although road traffic accident data is primarily used and essential for predicting road traffic accidents, including exposure data can lead to an improved understanding of these accidents and consequently improved models (Mikkonen & Peltola, 1997). Mikkonen & Peltola (1997) suggested that information such as population, the number of registered vehicles, and road hours can enhance road safety.

With the emergence of big data and the Internet of Things, various other sources of exposure data have been considered, such as weather data (Bao et al., 2019; Moosavi et al., 2019), human mobility data (Bao et al., 2019; Chen et al., 2016; Moosavi et al., 2019) and road network data (Bao et al., 2019; Moosavi et al., 2019). Like road accident data, however, the timely availability of these features must be considered. For instance, Chen et al. (2016) suggested using global positioning system (GPS) measurements to improve the accuracy of road traffic accident predictions. While the approach recommended by Chen et al. (2016), showed promising results, the approach required the GPS data of 1.6 million users. Collecting such extensive data would require an ongoing effort, as the predictions of the road accident model are based on real-time GPS data.

Given the potential range of exposure data, various types of road traffic accident predictive modeling approaches have been developed. These modeling approaches aim to learn the underlying spatial and temporal patterns to generate accurate predictions.

Historically, predicting the frequency of road traffic crashes has been attempted using statistical methods (Lord & Mannering, 2010). To improve the accuracy of these statistical models, several types of regression models have been proposed that specifically aim to address under-dispersion and/or over-dispersion (Abdulhafedh, 2017; Lord & Mannering, 2010). These approaches often perform poorly, which Yuan et al. (2018) attribute to the use of limited data and too simple modeling approaches that do cannot account for the properties of the underlying data, such as spatial heterogeneity.

Machine learning models can potentially address the limitations of statistical techniques in modeling road traffic accidents, such as assumptions about data distribution that may not hold and consequently lead to erroneous predictions (Chang & Chen, 2005). To establish the appropriateness of machine learning for predicting road traffic accidents, Silva et al. (2020) conducted a systematic literature review, covering machine learning approaches published from 2003 to 2020. Several of the identified studies found that machine learning based approaches are superior to traditional statistical models in predicting the number of road accidents, their severity, or a combination of both (Silva et al., 2020). The papers reviewed did not use the same dataset or consistent evaluation metrices, which makes it difficult to establish what the most appropriate machine learning technique is (Silva et al., 2020).

A significant shortcoming of the studies reviewed by Silva et al. (2020) is that most were confined to a single road segment or a small area. This limitation prevents the modeling or assessment of spatial heterogeneity, which is crucial in understanding the broader factors influencing road traffic fatalities (Yuan et al., 2018).

Limited research has employed traditional machine learning models over large areas. Parra et al. (2020) utilized the H3 hexagon technique (Brodsky, 2018) to spatially aggregate road traffic accident data. To incorporate, temporal data, the month, day of the month and the time of the day were included. Traditional machine learning models were then used to predict road traffic accidents across time and space. Road traffic accidents have also been modeled using Geographic Information System (GIS) techniques, which specifically consider spatial properties. For instance, Al-Dogom et al. (2019) analyzed traffic accidents that occurred in the United Kingdom to identify potential hotspots. Identifying these road traffic accident hotspots can assist in implementing targeted risk reduction strategies (Zahran et al., 2021). Historically, GIS techniques have had limited success when applied to small areas (Zahran et al., 2021), but have been more effective for larger areas (Al-Dogom et al., 2019; Nteziyaremye, 2018).

Very few studies have considered predicting road traffic fatalities over a short period, mainly due to the limitations of traditional models in explaining complex, non-linear and hierarchical data (Bao et al., 2019). Deep learning models have the potential to overcome these limitations by effectively capturing intricate patterns in tabular and non-tabular data. Chen et al. (2016) were among the first to use deep learning to predict road traffic accidents. However, as previously discussed their approach relies on human mobility data which may not be readily available. To model non-linear temporal patterns Ren et al. (2018), Yuan et al. (2018) and Bao et al. (2019) suggested using sequential deep learning models. While effective, using sequential models assumes that data is readily available.

In South Africa, limited research has been conducted on predicting road traffic fatalities, most studies tend to focus on causal factors (Du Toit, 2022). Causal studies suggest that both temporal and spatial relationships exist in South Africa (Nteziyaremye, 2018). While in terms of predictive modelling, Twala (2014), Mokoatle et al. (2019) and Du Toit (2022) have investigated road traffic accident severity classification. Rather than predicting where and when road fatalities will occur, these approaches aim to establish the severity of a road traffic accident. In this section, multiple approaches that can be used to predict road traffic fatalities were reviewed within the context of South Africa. Establishing clear definitions of the target feature and considering data availability is crucial for effective predictive modeling of road traffic fatalities. Several historical studies overlooked data availability aspects, making the approaches impractical in the context of this study.

To accurately predict road traffic fatalities, spatial and temporal patterns must be considered. In general, the related work suggests that these patterns are complex, requiring advanced modeling techniques such as deep learning. In the context of South Africa, limited research has been conducted on predicting road traffic fatalities. Considering these observations, this study aims to valuate whether road traffic fatalities can be predicted across time and space for the Western Cape province. Since it is unclear which modeling approach is best suited, both traditional machine learning models and deep learning models will be considered.

3. Methodology

This section outlines the methodology used to predict road traffic fatalities in the Western Cape. The methdology involves sourcing road fatality data, road network data, and weather data, which are then mapped to a hexagonal representation of the Western Cape using the H3 segmentation approach (Brodsky, 2018). The data is prepared for modeling by dividing the H3 segmented data into subsets and converting the data into both vector and raster representations. Vector data is utilized to develop traditional machine learning models, while raster-based data is used to develop deep learning models. Finally, the measures used to evaluate the performance of the road fatality prediction models are discussed.

3.1 Data Collection

Road traffic fatality data for the Western Cape was sourced from the Forensic Pathology Service. The dataset contained a row for each fatality. Fatalities related to the same traffic accident were grouped and considered a single binary event. Weather data was sourced from 62 weather stations and included average daily precipitation, temperature, visibility, and wind speed measurements. Missing and erroneous weather values were imputed using the closest reading from the same weather station. Road network data was assumed to be static and sourced from OpenStreetMap using the OSMnx (Boeing, 2017) Python package. The road network data retrieved is structured as a graph, where an edge represents a road segment, and each node represents either a dead-end or an intersection of roads. The maximum speed limit was missing for approximately 61.84% of the road segments. For segments with missing speed limits, the maximum speed limit was inferred using OSMnx, which assumes that similar roads have similar maximum speeds. Figure 1 shows an example of the road fatality event data, the location of weather measurements and the extent of the road traffic network of the Western Cape.

3.2 Data Aggregation

Before aggregating the road fatality data, the weather data and the road network data, the Western Cape Province was segmented using the H3 segmentation approach (Brodsky, 2018). The area of each hexagon obtained is approximately 36 km^2, with an average edge length of 8.54 km. The road fatality data was mapped to hexagons based on the latitude and longitude coordinates of the road fatality event. In urban areas, when the coordinates of a road fatality were unavailable, the fatality can be geo-coded because the margin for location error is relatively low. A similar approach cannot be used for rural areas. Fatalities in rural areas will be geo-coded to the nearest town which is imprecise, as the actual event may have occurred several dozen kilometres from the geocoded result. As a result, the research team conducted an exercise to ensure the validity and precision of the location information by cross-correlating the fatal event with government vehicle telemetry position data. This exercise does not fall within the scope of this paper.

Weather data was only collected from 62 weather stations, which is significantly fewer than the number of hexagons in the H3 grid. To assign weather data to the unsampled hexagons, spatial interpolation using the kriging method (Oliver & Webster, 1990) was used. The kriging method interpolates values based on a weighted combination of observed values.

Since weather tends to decrease smoothly and continuously as distance increases a Gaussian variogram model was selected. After spatial interpolation, outliers were identified based on threshold values and corrected using backward fill.

(a) Example of road fatality events
(b) Weather measurement locations
(c) Road network of the Western Cape

Figure 1: Raw data collected

Road network data was assigned to hexagons in a similar manner as road fatality data. After the road network data was assigned to hexagons, aggregated road statistics, such as the average maximum speed were calculated and assigned to each hexagon. Hexagons without any public road network data were not further considered. The spatial aggregation of data is similar to the approach used by Parra et al. (2020), but considered significantly smaller areas and also incorporates road network data.

3.4 Temporal Analysis

Various temporal patterns were observed in road traffic accidents covering the period from January 2019 to December 2022. In general, a sharp decrease in the number of fatal events was evident between January 2020 and July 2020. This decrease in road fatalities is mainly attributed to the travel restrictions imposed by the government during this period to curb the spread of coronavirus disease 2019 (COVID-19).

The total number of fatal events remained relatively consistent between Monday and Thursday, increased on Fridays, and peaked during weekends. The mean number of fatal events per day type was 1.62 for festive days, 1.37 for Friday to Sunday, and 1.26 for weekdays from Monday to Thursday. To account for the typical increase in traffic flow around holidays, the days before and after a holiday were also considered festive days. Road fatalities peaked around 7 am and 8 pm. The evening peak in fatalities was specifically high during weekends (Figure 2).
The observed temporal patterns suggested the need for careful temporal data aggregation. Since it was not explicitly clear how granular temporal aggregation should be performed, four different levels of temporal aggregation were considered

namely: (i) daily temporal aggregation over 4-hour timeslots, (ii) daily temporal aggregation over 8-hour timeslots, (iii) type of day aggregation over 4-hour timeslots and (iv) type of day aggregation over 8-hour timeslots. Here type of day aggregation refers to aggregation by festive days, Monday to Thursday and Friday to Sunday.

3.5 Spatial Analysis

Road fatalities in the Western Cape were dispersed across both rural and urban areas, with approximately 826 road fatality events in the City of Cape Town, the only metropolitan municipality in the region. The number of road fatalities in the five district municipalities in the Western Cape ranged between 49 and 155 events. A spatial correlation analysis revealed both global and local spatial autocorrelation of road fatalities.

The global spatial autocorrelation was assessed using Moran's I value (Moran, 1950). Moran's I value ranges from negative one to positive one and quantifies how similar or dissimilar values of a variable are across a geographic space. A Moran's I value of 0.82 with a p-value of 0.001 was obtained which indicated that hexagons with a high number of road fatality events are located close to hexagons with a high number of road fatality events. The same holds for hexagons with a low number of road fatality events. To determine whether specific hexagons belong to a cluster and the nature of the cluster, the local indicators of spatial association (LISA) algorithm (Anselin, 1995) was used. Several groups of high-fatality hotspots and low-fatality hotspots were observed when a p-value of 0.05 was used to determine significance (Figure 3). Hexagons assigned to the high fatality cluster typically occurred in population-dense areas or within cities.

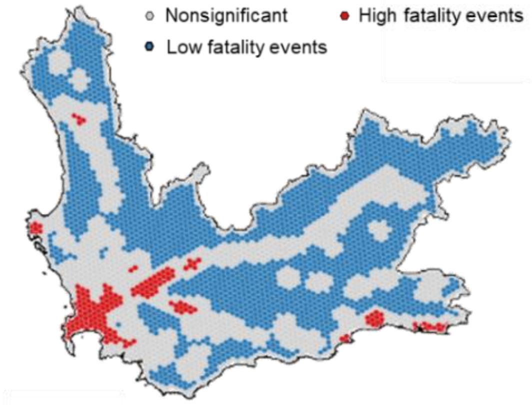

Figure 3: Spatial dispersion of road fatality events

3.6 Data Preparation

Given a hexagon H and a time frame T, the modeling task is to predict whether a road fatality occurred in H and within T. The descriptive features consist of (i) global and local time features, (ii) weather features and (iii) road network features.

For vector-based modeling, each hexagon and time slice is considered an instance. For raster-based modeling, the vector-based data is rasterized at a resolution of (-0.03, 0.03), creating a separate image encoding for each non-global feature, including the target feature. To consistently compare vector-based modeling with raster-based modeling, raster-based predictions are re-projected to the original hexagons. The aggregated dataset is divided into three non-overlapping subsets using out-of-time sampling to prevent data leakage. The training dataset covers

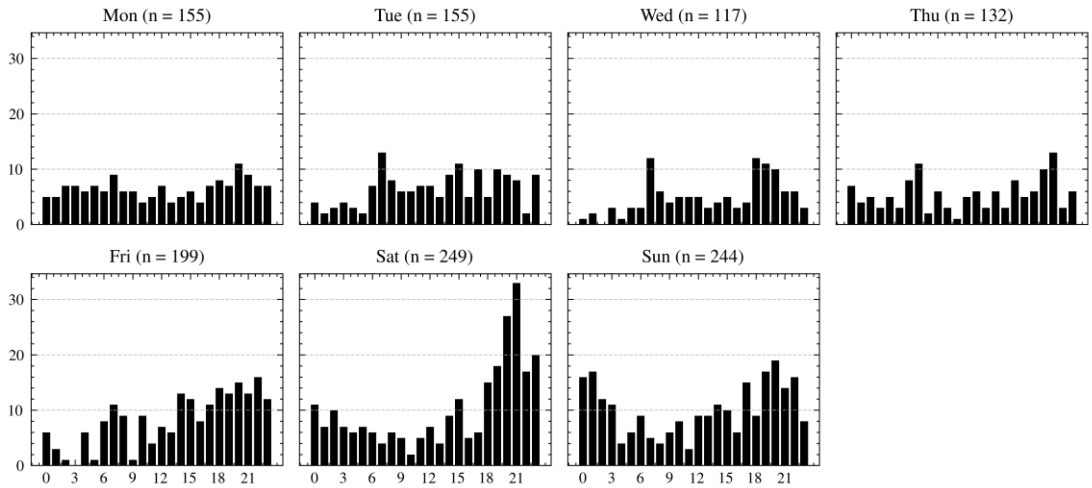

Figure 2: Temporal dispersion of road fatality events

the period from 18 November 2019 to 31 December 2021, the validation dataset spans from 1 January 2022 to 31 August 2022, and the test dataset includes data from 1 September 2022 to 28 February 2023.

3.7 Machine learning models

The following vector-based models were considered: (i) a logistic regression model with L2 regularization, (ii) a random forest model with a maximum depth of five, a minimum sample split set to 13, and a maximum of 204 trees, and (iii) a Naive Bayes model. Since road fatality events are rare, artificially increasing the rare events using the synthetic minority over-sampling technique (SMOTE) (Chawla et al., 2002) were considered. SMOTE was specifically considered since Du Toit (2022) obtained the best classification of road accident severity in South Africa when SMOTE was used.

Two raster-based models were considered: (i) an adapted U-Net (Ronneberger et al., 2015) model and a custom road fatality fully convolutional neural network (RFFCNN). Both raster-based models accept two inputs: spatial information represented as a raster where each feature represents a channel, and spatially independent information represented as a vector. The raster-based models were developed and implemented using Pytorch. Experiments were conducted in Python 3.9.16 in a Windows 10 operating system with 32 GB RAM and a Nvidia GeForce RTX 3070 GPU to accelerate the deep learning procedure. Both the adapted U-Net and the RFFCNN averaged 15-20 seconds per epoch.

The adapted U-Net model consists of an encoder, a bottleneck and a decoder, each composed of convolution blocks which consist of a two-dimensional (2D) convolutional layer, batch normalisation, rectified linear unit (ReLU) activation, and skip connections. The encoder consists of four convolution blocks, each followed by 2D max pooling. Parameters for the network are in line with the architecture outline described in (Ronneberger et al., 2015). The tabular data is passed through a feed-forward fully connected network which consists of three dense layers with ReLU activations and a dropout layer with a probability of 50%. The first two dense layers output 32 and 64 features, respectively. The final dense layer produces a vector with dimensions that match the output of the encoder.

The output of the feed-forward fully connected network is reshaped and joined with the output of the encoder. The decoder performs up-sampling using transposed 2D convolution layers with a stride of two. The kernel sizes and padding of the 2D convolution layers are selected such that the size of the feature maps passed from the skip-connected is the same as the up-sampled feature maps. Transpose convolution layers as opposed to traditional upsampling are employed in the decoder, as transpose convolution layers have parameters that can be learned. The adapted U-Net, produces a raster with predictions, matching the input raster.

The RFFCNN is a straightforward fully convolutional neural network, inspired by similar deep learning architectures in literature but without an encoder-decoder approach. Using a simpler model helps prevent memorization and encourages learning, as it has significantly fewer parameters. The RFFCNN processes spatial features through four convolutional layers, gradually increasing channel depth, i.e., 32, 64, 128, 256, with batch normalization. Increasing depth is a common practice in deep learning and is more efficient in terms of parameter usage.

Tabular features are processed by a feed-forward fully connected neural network of five dense layers, gradually increasing in output features, i.e., 32, 64, 128, 256, with ReLU and dropout at 50%. The final dense layer produces a vector with dimensions equal to the output of the fourth convolution layer. The outputs of the fully connected network are concatenated with the output of the fourth convolutional layer. The concatenated values are then passed through a final convolutional layer with a single channel, a kernel size of one, and zero padding, producing a single raster image with predictions.

Both the adapted U-Net and RFFCNN are trained using binary cross entropy with logit loss. To handle class imbalance the weight of positive event is set to 1750:1. The models were trained using the adaptive moment estimation (Adam) optimizer (Kingma & Ba, 2014), for a maximum of 40 epochs with a batch size of 3. To prevent overfitting, training was stopped when the validation loss did not decrease for three consecutive epochs. After training a model was reset to the parameter values associated with the smallest validation loss.

Both the vector-based models and raster-based models are compared against the baseline model to determine if they provide significant improvements over the intuitive historical average approach. The predictions of the baseline model are based on historical fatality events at the same hexagon, with a lag of one year. This approach is similar to the MEDIC model developed by Zhou & Matteson (2015), which averages historical counts over 4 to 5 years to establish ambulance demand.

3.8 Evaluation measures

To evaluate the effectiveness of the models considered, multiple evaluation metrics are employed. Precision is calculated by determining the percentage of correctly predicted fatal events out of all predicted fatal events. Recall is calculated by determining the percentage of correctly predicted fatal events out of all actual fatal events. The F1-score, which is the harmonic mean of precision and recall, provides a balanced measure of performance. Additionally, the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) is used to establish the model's ability to distinguish between fatal and non-fatal events. All four performance metrics range from zero to one, with larger values indicating better performance. In the case of AUC-ROC, a value of 0.5 corresponds to a model that cannot differentiate between fatal and non-fatal events.

In the context of road traffic fatalities, maximizing recall is crucial since it measures the ability of the model to identify actual road traffic fatal events. High recall ensures that most fatal events are detected, which is vital for effective intervention and prevention. High recall paired with low precision can however lead to the generation of multiple false positives, which means that fatal events are predicted where there are none. Consequently, the limited resources available will be used ineffectively.

Although all four of the metrices considered are reported, the different algorithms are ranked using the AUC-ROC performance since the AUC-ROC provides a more holistic view of the model’s effectiveness in distinguishing between fatal and non-fatal events.

4. Results and discussion

The performance of the models on the validation dataset is provided in Table 1. Table 1 excludes the Random Forest model and the Logistic Regression model trained without SMOTE, as they never predicted any road fatality events.

The baseline model outperformed all other models in terms of precision for all four-time aggregations considered. However, the recall of the baseline model was considerably worse compared to the other models. The low number of predictions made by the baseline model suggests that no learning is performed - historical average is simply used. Consequently, it cannot extrapolate learnings from one area and apply these learnings to other areas. The Naïve Bayes model without SMOTE consistently achieved the second-highest precision across the various scenarios. However, this model also had the second lowest recall, meaning it missed many actual fatal events. Since road traffic fatalities are rare, the probability of predicting road traffic fatalities is inherently low, which leads to fewer rare event predictions. Additionally, using SMOTE did not lead to significant improvements in the performance of the models. The Logistic Regression model using SMOTE achieved the highest recall scores for all experiments and was typically followed by the Random Forest model using SMOTE.

In terms of deep-learning models, the RFFCNN achieved higher AUC-ROC values than the U-Net model in all scenarios except for the most granular aggregation. For the most granular time aggregation, the U-Net achieved the best AUC-ROC. Each of the four different time aggregations had a different model which achieved the highest AUC-ROC. The U-Net model performed the best on the daily 4-hour aggregation, the Logistic Regression model with SMOTE performed the best on the daily 8-hour aggregation, the Random Forest with SMOTE performed the best on the weekday 4- hour aggregation, and the RFFCNN achieved the highest AUC-ROC on the weekday 8-hour aggregation. The models performed similarly on the test dataset, indicating that they do not simply memorize the historical values but instead generalize to new data.

Table 1: Comparison of model performance on the test subset

Model	Day	Time	Precision	Recall	F1-score	AUC-ROC
U-Net	Daily	4-hour	0.60	70.21	1.18	88.16
Random Forest*	Daily	4-hour	0.47	76.43	0.93	84.35
Logistic regression*	Daily	4-hour	0.25	82.51	0.51	83.63
RFCNN	Daily	4-hour	0.64	70.55	1.28	82.73
Naïve Bayes*	Daily	4-hour	1.32	59.32	2.58	78.61
Naïve Bayes	Daily	4-hour	1.47	58.17	2.86	78.16
Baseline	Daily	4-hour	2.45	2.38	2.42	-
Logistic regression*	Daily	8-hour	0.29	79.85	0.58	83.85
Random Forest*	Daily	8-hour	0.51	72.62	1.01	83.21
RFCNN	Daily	8-hour	0.58	66.09	1.15	80.60
Naïve Bayes*	Daily	8-hour	1.21	59.31	2.38	78.60
Naïve Bayes	Daily	8-hour	1.37	58.17	2.67	78.17
U-Net	Daily	8-hour	0.72	59.93	1.41	78.17
Baseline	Daily	8-hour	2.65	2.61	2.63	-
Random Forest*	Weekly	4-hour	0.60	73.00	1.20	83.44
Logistic regression*	Weekly	4-hour	0.37	75.66	0.74	82.70
RFCNN	Weekly	4-hour	0.69	68.51	1.37	81.81
U-Net	Weekly	4-hour	0.77	64.64	1.52	80.25
Naïve Bayes*	Weekly	4-hour	1.44	59.70	2.81	78.81
Naïve Bayes	Weekly	4-hour	1.58	57.80	3.06	77.98
Baseline	Weekly	4-hour	2.92	2.83	2.89	-
RFCNN	Weekly	8-hour	0.72	73.27	1.44	83.42
Random Forest*	Weekly	8-hour	0.73	72.24	1.45	82.80
Logistic regression*	Weekly	8-hour	0.41	76.43	0.82	81.95
U-Net	Weekly	8-hour	0.60	70.79	1.18	81.62
Naïve Bayes*	Weekly	8-hour	1.93	59.70	3.73	78.82
Naïve Bayes	Weekly	8-hour	2.07	57.80	4.00	77.97
Baseline	Weekly	8-hour	3.93	3.52	3.72	-

* Models trained on a dataset sampled using the Synthetic Minority Over-sampling Technique (SMOTE)

While the metrics considered this far capture the global performance of each model, they fail to reveal biases in predictions. For instance, a model may perform well overall but still exhibit biases towards certain locations (Yuan et al., 2018). To examine the potential biases of each model, the predictions of the different models for similar timeframes were compared (Figure 4).

The Naive Bayes model typically repeats predictions, often focusing on areas within the City of Cape Town metropolitan area across different day types and shifts. This behaviour explains its relatively high precision, as it makes fewer predictions outside this area. The Logistic Regression with SMOTE offered slight variability, with predictions often targeting national and provincial routes during festive days. Similarly, the Random Forest model also demonstrated slight variability, highlighting areas of high risk. In contrast, the deep learning models exhibited shifts in predictions based on temporal elements, with noticeable variations across different shifts and day types. An example of the predictions of the different models is illustrated in Figure 4 for a festive day during the time 20:00 to 24:00 with predictions indicated in yellow. During this timeslot, only one road fatality event occurred. Apart from manually examining predictions, the performance of models must also be compared at a district level. In the case of the Western Province, the majority of road fatalities occurred in the City of Cape Town which can bias the predictions towards this district. Performing a district-level analysis supported the insights derived from examining individual predictions. It was particularly evident that the Naïve Bayes model did not generate predictions for more than one district.

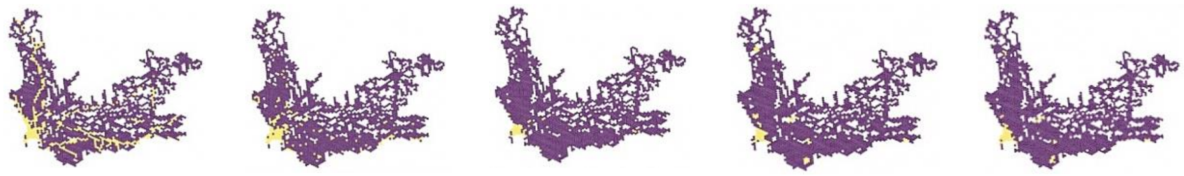

Logistic Regression with SMOTE | Random Forest with SMOTE | Naïve Bayes with SMOTE | U-Net | RFFCNN

Figure 4: Comparison of model predictions for a festive day during the timeframe 20:00 to 24:00

5. Practical Implementation for Authorities

The road fatality predictions model is intended to be used as a decision support tool for intervention planning by traffic law enforcement. When used, operational bias must be specifically considered. For instance, when traffic law enforcement identifies a specific hotspot area and applies constant high amounts of surveillance in that area, the model may start to predict fatalities in a different area and not the original hotspot area. This does however not imply that law surveillance must be stopped in the area. To avoid the operational risk of reallocating resources, traffic operational literature must be incorporated in conjunction with road traffic accident literature.

For any given period, the model predicts many possible fatality locations. It is not feasible for traffic law enforcement to intervene at all locations, so we suggest future work on an optimization layer to determine the most effective allocation of limited traffic law enforcement resources to the predicted locations. Rosenfeld et al. (2017) formulated the traffic enforcement as $\sum_t \sum_h p_{t,h} \times (1 - F(o_{t,h}))$, where p is the probability of a fatal event in a hexagon and F calculates the effectiveness of an intervention o in a hexagon h in a timeslot t. This optimization approach, visualized in Figure 5, can take into account the type of intervention, given constraints such as the number of officers available, distance of location, and mandatory operations. The arcs in Figure 5 are placed from each traffic centre to an accident hotspot predicted by the fatality model, suggesting specific interventions with each colour.

Finally, road traffic fatalities models need to be deployed and researched with a focus on ethical artificial intelligence. RTFs in South Africa disproportionately affect lower-income groups. This is partly due to systemic issues, such as the negative historical Apartheid spatial planning policies. The disproportionate representation of some groups in the historical data may cause bias in a model, and decision-making based on these models may cause adverse impact. Model reproducibility, explainability, and traceability are therefore critical to mitigating bias.

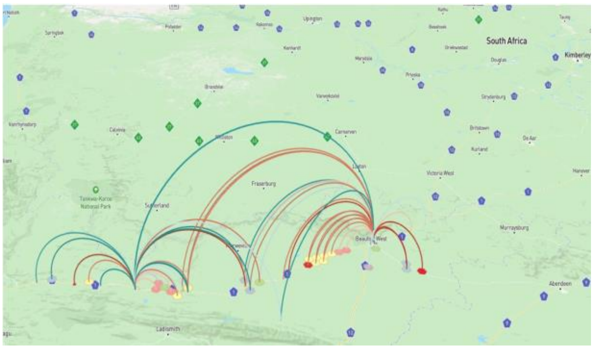

Figure 5: Optimizing operational interventions with predictions

6. Conclusion

The primary limitation of this study was data, in quantity and quality of data. Training a spatio-temporal model solely on road traffic fatalities, a subset of road traffic accidents resulted in a highly unbalanced dataset. We believe that increasing the dataset and expanding it to include all crash data would improve results, particularly precision. Likewise, using more granular exposure data could improve the results, especially the weather data which was aggregated daily. Despite these data limitations, the results are promising and indicate the potential for improved predictive accuracy with enhanced data collection and integration efforts. When selecting a model for implementation, it is suggested to carefully examine the predictions since models with similar global performance can behave very differently due to biases and the data available. For instance, if a government entity has access to a pool of resources, data streams, and clean data, the deep learning approach would be suggested at high levels of temporal granularity. On the other hand, when resources are limited and predictions are required on a less frequent basis, a Random Forest model could be more useful than relying on historical averages.

7. References

Abdulhafedh, A. (2017). Road Traffic Crash Data: An Overview on Sources, Problems, and Collection Methods. Journal of Transportation Technologies, 07(02),206–219. https://doi.org/10.4236/jtts.2017.72015

Al-Dogom, D., Aburaed, N., Al-Saad, M., & Almansoori, S. (2019). Spatio-temporal Analysis and Machine Learning for Traffic Accidents Prediction. 2019 2nd International Conference on Signal Processing and Information Security (ICSPIS), 1–4. https://doi.org/10.1109/ICSPIS48135.2019.9045892

Al-Hasani, G. (2021). Spatio-Temporal Modelling, Analysis and Forecasting of Road Traffic Accidents in Oman [Doctor of Philosophy]. Staffordshire University.

Anselin, L. (1995). Local Indicators of Spatial Association—LISA. Geographical Analysis, 27(2), 93–115. https://doi.org/10.1111/j.1538-4632.1995.tb00338.x

Arrive Alive. (n.d.). Definitions And Explanatory Notes. Retrieved August 4, 2023, from https://www.arrivealive.co.za/definitions-and-explanatory-notes

Bao, J., Liu, P., & Ukkusuri, S. V. (2019). A spatiotemporal deep learning approach for citywide short-term crash risk prediction with multi-source data. Accident Analysis & Prevention, 122, 239–254. https://doi.org/ 10.1016/j.aap.2018.10.015

Boeing, G. (2017). OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. Computers, Environment and Urban Systems, 65, 126–139. https://doi.org/10.1016/j.compenvurbsys.2017.05.004

Brodsky, I. (2018, June 27). H3: Uber's Hexagonal Hierarchical Spatial Index. https://www. uber.com/en-ZA/blog/h3/

Chang, L.-Y., & Chen, W.-C. (2005). Data mining of tree-based models to analyze freeway accident frequency. Journal of Safety Research, 36(4), 365–375. https://doi.org/10.1016/j.jsr.2005.06.013

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. Journal of Artificial Intelligence Research, 16, 321–357. https://doi.org/10.1613/jair.953

Chen, Q., Song, X., Yamada, H., & Shibasaki, R. (2016). Learning Deep Representation from Big and Heterogeneous Data for Traffic Accident Inference. Proceedings of the Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence, 30(1), 338–334. https://doi.org/10.1609/aaai.v30i1.10011

Du Toit, C. (2022). Cape Town road traffic accident analysis: Utilising supervised learning techniques and discussing their effectiveness [Master Thesis]. University of Cape Town.

Global plan decade of action for road safety 2021-2030.(2021). https://cdn.who.int/media/docs/ default-source/documents/health-topics/road-traffic-injuries/global-plan-for-road-safety.pdf?sfvrsn=65cf34c8_35&download=true

Global status report on road safety 2018. (2018). World Health Organization.

Kingma, D. P., & Ba, J. (2014). Adam: A Method for Stochastic Optimization. ArXiv Preprint ArXiv:1412.6980.

Lazarević, A., Srivastava, J., & Kumar, V. (2004, May). Data Mining for Analysis of Rare Events: A Case Study in Security, Financial and Medical Applications. Pacific-Asia Conference on Knowledge Discovery and Data Mining.

Lord, D., & Mannering, F. (2010). The statistical analysis of crash-frequency data: A review and assessment of methodological alternatives. Transportation Research Part A: Policy and Practice, 44(5), 291–305. https://doi.org/ 10.1016/j.tra.2010.02.001

Maalouf, M., & Trafalis, T. B. (2011). Rare events and imbalanced datasets: an overview. International Journal of Data Mining, Modelling and Management, 3(4), 375. https://doi.org/10.1504/IJDMMM.2011.042935

Mikkonen, V., & Peltola, H. (1997). Road safety principles and models: Review of descriptive, predictive, risk and accident consequence models.

Mokoatle, M., Vukosi Marivate, Dr., & Michael Esiefarienrhe Bukohwo, Professor. (2019). Predicting Road Traffic Accident Severity using Accident Report Data in South Africa. Proceedings of the 20th Annual International Conference on Digital Government Research, 11–17. https://doi.org/10.1145/3325112.3325211

Moosavi, S., Samavatian, M. H., Parthasarathy, S., Teodorescu, R., & Ramnath, R. (2019). Accident Risk Prediction based on Heterogeneous Sparse Data. Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems,33–42. https://doi.org/10.1145/3347146.3359078

Moran, P. A. P. (1950). Notes on Continuous Stochastic Phenomena. Biometrika, 37(1/2),17. https://doi.org/10.2307/2332142

Nteziyaremye, P. (2018). Investigating the link between the built environment and the incidence of pedestrian crashes in Cape Town, South Africa [Doctor of Philosophy]. Stellenbosch University.

Oliver, M. A., & Webster, R. (1990). Kriging: a method of interpolation for geographical information systems. International Journal of Geographical Information Systems, 4(3), 313–332. https://doi.org/10.1080/02693799008941549

Parra, C., Ponce, C., & Rodrigo, S. F. (2020). Evaluating the Performance of Explainable Machine Learning Models in Traffic Accidents Prediction in California. 2020 39th International Conference of the Chilean Computer Science Society, 1–8. https://doi.org/10.1109/SCCC51225.2020.9281196

Peden, M., Scurfield, R., Sleet, D., Mohan, D., Hyder, A. A., Jarawan, E., & Mathers, C. (Eds.). (2004). World report on road traffic injury prevention. WHO.

Ren, H., Song, Y., Wang, J., Hu, Y., & Lei, J. (2018). A Deep Learning Approach to the Citywide Traffic Accident Risk Prediction. 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 3346–3351. https://doi.org/10.1109/ITSC.2018.8569437

Road Transport Accident Deaths in South Africa, 2007-2019.(2024). https://www.statssa.gov.za/publications/Report-03-09-07/Report-03-09-072019.pdf

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, 234–241.https://doi.org/10.1007/978-3-319-24574-4_28

Rosenfeld, A., Maksimov, O., & Kraus, S. (2017). Optimizing Traffic Enforcement: From the Lab to the Roads (pp. 3–20). https://doi.org/10.1007/978-3-319-68711-7_1

Silva, P. B., Andrade, M., & Ferreira, S. (2020). Machine learning applied to road safety modeling: A systematic literature review. Journal of Traffic and Transportation Engineering, 7(6), 775–790. https://doi.org/10.1016/ j.jtte.2020.07.004

State of road safety in South Africa: January 2023 to December 2023. (2024). https://www.rtmc.co.za/images/rtmc/docs/CalendarReports/STATE-OF-ROAD-SAFETY-REPORT-2023.pdf

Twala, B. (2014). Extracting grey relational systems from incomplete road traffic accidents data: the case of Gauteng Province in South Africa. Expert Systems, 31(3), 220–231. https://doi.org/10.1111/exsy.12035

Weiss, G. M. (2004). Mining with rarity. The Association for Computing Machinery's Special Interest Group on Knowledge Discovery in Data Explorations Newsletter, 6(1), 7–19. https://doi.org/10.1145/1007730.1007734

Yuan, Z., Zhou, X., & Yang, T. (2018). Hetero-ConvLSTM. Proceedings of the 24th Association for Computing Machinery's Special Interest Group on Knowledge Discovery and Data Mining International Conference on Knowledge Discovery and Data Mining, 984–992. https://doi.org/10.1145/3219819.3219922

Zahran, E.-S. M. M., Tan, S. J., Tan, E. H. A., Mohamad 'Asri Putra, N. A. 'Atiqah B., Yap, Y. H., & Abdul Rahman, E. K. (2021). Spatial analysis of road traffic accident hotspots: evaluation and validation of recent approaches using road safety audit. Journal of Transportation Safety & Security, 13(6), 575–604. https://doi.org/10.1080/19439962.2019.1658673

Zhou, Z., & Matteson, D. S. (2015). Predicting Ambulance Demand. Proceedings of the 21th Association for Computing Machinery's Special Interest Group on Knowledge Discovery and Data Mining International Conference on Knowledge Discovery and Data Mining, 2297–2303. https://doi.org/10.1145/2783258.2788570