# Learning Incident Prediction Models Over Large Geographical Areas for Emergency Response Systems

Sayyed Mohsen Vazirizade Vanderbilt University, Nashville, TN s.m.vazirizade@vanderbilt.edu

Said El Said Tennessee Department of Transportation said.elsaid@tn.gov Vanderbilt University, Nashville, TN ayan.mukhopadhyay@vanderbilt.edu

Ayan Mukhopadhyay

Hiba Baroud Vanderbilt University, Nashville, TN hiba.baroud@vanderbilt.edu Geoffrey Pettet Vanderbilt University, Nashville, TN geoffrey.a.pettet@vanderbilt.edu

Abhishek Dubey Vanderbilt University, Nashville, TN abhishek.dubey@vanderbilt.edu

Abstract—Principled decision making in emergency response management necessitates the use of statistical models that predict the spatial-temporal likelihood of incident occurrence. These statistical models are then used for proactive stationing which allocates first responders across the spatial area in order to reduce overall response time. Traditional methods that simply aggregate past incidents over space and time fail to make useful shortterm predictions when the spatial region is large and focused on fine-grained spatial entities like interstate highway networks. This is partially due to the sparsity of incidents with respect to the area in consideration. Further, accidents are affected by several covariates, and collecting, cleaning, and managing multiple streams of data from various sources is challenging for large spatial areas. In this paper, we highlight how this problem is being solved for the state of Tennessee, a state in the USA with a total area of over 100,000 sq. km. Our pipeline, based on a combination of synthetic resampling, non-spatial clustering, and learning from data can efficiently forecast the spatial and temporal dynamics of accident occurrence, even under sparse conditions. In the paper, we describe our pipeline that uses data related to roadway geometry, weather, historical accidents, and real-time traffic congestion to aid accident forecasting. To understand how our forecasting model can affect allocation and dispatch, we improve upon a classical resource allocation approach. Experimental results show that our approach can significantly reduce response times in the field in comparison with current approaches followed by first responders.

#### I. INTRODUCTION

A constant threat that plagues humans across the globe are incidents like traffic accidents, fires, and crimes. Such incidents result in loss of life, injuries, and damage to properties and are collectively labeled as *emergencies*, which are defined as incidents that threaten public safety, health, and welfare. Consider road accidents and calls for emergency medical services (EMS) as examples. Road accidents alone account for 1.25 million deaths globally and about 240 million EMS calls are made in the U.S. each year [1]. The large number of such incidents makes it imperative that principled methods be designed to ensure fast and effective response to incidents. At the same time, it is crucial to design infrastructure that mitigates and prevents the occurrence of such incidents. Indeed, it is well-documented that one of the most important responsibilities of federal, state, and local governments is mitigating and dealing with such events [2]. As a result, governments strive to make systematic plans, allocate resources, and take preventive measures in order to alleviate threats that such incidents pose.

Emergency response management (ERM) is defined as the set of procedures and tools that first responders use to deal with incidents such as road accidents. It includes specific mechanisms to forecast incidents, detect incidents, allocate resources like ambulances, dispatch resources, and finally mitigate the post-effects of incidents [1]. Arguably, the most important component of the pipeline is to understand the spatial and temporal dynamics of incident occurrence. Gaining such an understanding can aid resource allocation and dispatch, improve the understanding of factors that cause accidents, and improve the design of safety codes. While there are several ways to analyze spatial temporal incidents, learning data-driven forecasting models is particularly important since the fundamental goal of understanding the dynamics of accident occurrence is to aid response and dispatch. As a result, generative models conditional on relevant covariates are particularly relevant to the overall ERM pipeline. For example, consider a forecasting model for accident occurrence as a function of roadway speed limits. Understanding how the speed limit affects accidents helps in accurately capturing first-order effects that impact accidents, and forecasting future incidents helps shape better policy decisions pertaining to resource allocation (ambulances, for example) and response.

This paper discusses a framework for predicting extremely sparse spatial temporal incidents. Over the past year, our project has focused on developing principled approaches to address emergency response for Tennessee, a state in the United States with a population of approximately 6.9 million and a total area of over 100,000 sq. km. While we have extensively tackled emergency response in past collaborations with several government bodies restricted to cities [3], [4], [5], [6], [7], planning emergency response in extremely large geographical areas (an entire state, for example) is significantly more challenging. The problems are exacerbated when we limit the area of interest only to interstate highways across the state, which reduces the number of samples of positive incidents available across the road network, leading to extreme sparsity. This imbalance is particularly evident while creating forecasting models in high temporal resolution. However, our collaboration with first responders revealed that such forecasting models can be extremely beneficial for resource allocation and dispatch.

Our contributions in this paper are two-fold: 1) we develop a pipeline that can effectively forecast incidents that are sparsely scattered in space and time, which can be used by ERM pipelines to reduce the average response time to accidents. We show how a combination of synthetic resampling and non-spatial clustering can result in the creation of accurate spatial temporal models for short-term forecasting of road accidents. 2) Unlike the majority of forecasting models in literature, we evaluate how our forecasting pipeline affects response times to accidents. To this end, we modify the classical p-median problem for resource allocation. Our novel contribution balances the geographic spread of responders and their availability. Through extensive simulations, we show how our forecasting pipeline and allocation algorithm provide significant reduction in response times.

## **II. PRIOR WORK**

A variety of approaches have been used to understand the spatial and temporal dynamics of road accidents. One of the earliest methods, known as 'crash frequency analysis', uses the frequency of incidents in a specific discretized spatial area as a measure of the inherent risk the area possesses [8]. This approach also forms the basis of hotspot analysis [9], [10], which is widely used in practice as a tool to visualize historical accidents and make predictions. Statistical models have also been explored in this context. The most widely used approach, Poisson regression, models the expected value of the count of incidents in a given time period as a linear combination of the features. While it does not perform well on data with dispersion (where mean is not equal to the variance of the data) and sparse data, hierarchical Poisson models [11], [12], [13], [14], [15] and zero-inflated models can be used instead [16], [17], [18], [19]. In recent years, data mining models such as neural networks [20], [21], [22], [23], [24], [25] and support vector machines [26], [27], [28] have also been explored. We refer interested readers to our survey on emergency response for a detailed analysis of prior work in forecasting, allocation, and dispatch [1].

Two major shortcomings of prior approaches are -1) the inability to deal with large sparsity in data, and 2) the inability to make accurate short-term predictions. Even zero-inflated models, the only class of statistical models that have been shown to work fairly well on sparse spatial temporal data, fail to perform well when trained on emergency incidents on



Fig. 1: Blue lines represent TN's roadway network. Yellow segments represent interstate highway segments under the jurisdiction of TDOT and are the area of study for this paper.

highways in large geographic areas (large areas often exhibit > 99% sparsity). We show later how zero-inflated models fail to work in such scenarios. Also, the majority of prior work in accident prediction either focuses on the spatial dynamics of incident occurrence or consider extremely coarse temporal resolutions. While long-term predictions can be useful to analyze policies (optimize road construction, for example), such predictions are not particularly useful for allocating and dispatching responders. We address these gaps in prior work by designing a pipeline for spatial temporal incident prediction that can adapt to extremely sparse data and aid resource allocation.

#### **III. PROBLEM FORMULATION**

Consider a spatial area of interest S, in which incidents (like accidents) occur in space and time. The decision-maker observes a set of samples (possibly noisy) drawn from an incident arrival distribution. These samples are denoted by  $\{(s_1, t_1, k_1, w_1), (s_2, t_2, k_2, w_2),$ 

 $\ldots, (s_n, t_n, k_n, w_n)$ , where  $s_i, t_i$  and  $k_i$  denote the location, time of occurrence, and reported severity of the *i*th incident, respectively, and  $w_i \in \mathbb{R}^m$  represents a vector of features associated with the environment defined by the location and time of the incident. We refer to this tuple of vectors as D, which denotes the input data that the decision-maker has access to. The vector w can contain spatial, temporal, or spatiotemporal features and it captures covariates that potentially affect incident occurrence. For example, w typically includes features such as weather, traffic volume, and time of day. The most general form of incident prediction can then be stated as learning the parameters  $\theta$  of a function over a random variable X conditioned on w. We denote this function by  $f(X \mid w, \theta)$ . The random variable X represents a measure of incident occurrence such as a *count* or *presence* of incidents during a specific time period. The goal of the incident prediction problem is to find the *optimal* parameters  $\theta^*$  that best describe D. This can be formulated as a maximum likelihood estimation (MLE) problem or an equivalent empirical risk minimization (ERM) problem.

In our problem setting (the Interstate Highway network of the state of Tennessee) (Fig. 1), it is intuitive to represent S in the form of a graph G = (V, E), where V is a set of vertices and E is a set of edges. Edges represent specific segments of highways. The prediction problem then reduces to learning fsuch that it captures the spatial temporal dynamics of incidents on E.

# A. Challenges

The problem described in previous section is hard due to the following challenges.

- 1) Irregular incident occurrence: It is well-established in literature that predicting road accidents is extremely difficult due to inherent randomness of accidents and spatially varying factors [29], [1]. While accidents are affected by various features, it is difficult to take all determinants into account while designing forecasting models. For example, consider the condition of a specific road. It is difficult to observe such features in real-time, thereby resulting in unobserved heterogeneity in the likelihood of incident occurrence across a large spatial region. Indeed, sophisticated models have under-performed in predicting accidents. For example, an approach particularly important to practitioners was developed by Bao et al. [25], who used a spatio-temporal convolution long short-term memory network (LSTM) to predict short-term crash risks. While the network structure was a combination of various complex sub-networks, the accuracy of hourly predictions was limited, highlighting the inherent difficulty of predicting crash frequency at low temporal and spatial resolutions.
- 2) Sparsity: It is also crucial to take into account the frequency of incident occurrence. While the frequency of road accidents is alarming, incidents are actually extremely sporadic when viewed from the perspective of total time and space in consideration. For example, there were a total of approximately 78,000 road accidents reported between 2017-2020 on interstate highways in Tennessee. Now consider the goal of learning the dynamics of incident occurrence. While historical data can be studied using hotspots to improve policy, short-term forecasting models are important for deploying ambulances, help-trucks, and other emergency responders. Based on our conversations with first-responders, short term deployment often occurs several times in a day, the most common frequency being once every four hours. Considering a total of about 5,000 road segments and time slots of four hours, our data shows > 99% sparsity. We represent this challenge schematically in Fig. 2 by randomly selecting 180 road segments for April 2019 and 180 four-hour time slots. Each pixel in the matrix denotes the presence (white) or absence (black) of an accident in a segment (denoted by rows) in a span of four hours (denoted by columns). We see that most of the matrix consists of black pixels (99.8%), making such problems extremely difficult from the perspective of datadriven modeling. In comparison, we have previously shown how standard statistical models can be used to predict incidents in small urban areas [3], [4], [5] (such situations typically exhibit < 90% sparsity).
- 3) Data Integration: Road accidents are affected by a large number of determinants which can be spatial, temporal, or spatial temporal in nature. For example, the geometry of a specific road segment does not change over time and is an example of a spatial feature. Time of day,



Fig. 2: Schematic overview of the sparsity of accident occurrence across space and time. The figure shows randomly selected 180 road segments for four hour time windows in April 2019.

on the other hand, is an example of a temporal feature. Some features can be affected by both space and time; for example traffic congestion in a specific area is determined by the spatial location of the area as well as time of day. For predicting accidents in large geographic areas, it is challenging to collect, clean, understand, and analyze data from different sources and integrate them into models for incident prediction.

# IV. DATA

Before describing our pipeline for learning the spatial and temporal dynamics of highway accidents, we describe the covariates used and their sources (Table I). We highlight the importance of this stage in real-world machine learning pipelines; in fact, the availability of multiple streams of data has been noted as being particularly important for predicting accidents [1]. Our exercise of collecting data for features that affect accidents was guided by the invaluable domain expertise of first responders and our collaborators at the Tennessee Department of Transportation.

# A. Features

This section describes the features we extract from the base data (Table I) and use as covariates in our pipeline (Table V).

1) *Roadway Information*: To learn a predictive model for accidents over a graph of roadways, it is imperative to first define the edges and the vertices of the graph. We collect roadway information from INRIX [30], a private entity that provides location-based data and analytics, such as traffic and parking, to automakers, cities, and road authorities worldwide. We retrieved information for about 80,000 roadway segments in the state of Tennessee, out of which about 5,000 are interstate highway segments. We also retrieved *static* features associated with each segment that are immutable (relatively) over time. For example, for each road segment, we collected the number of lanes, length, and coordinates.

In order to evaluate how roadway shape affects accidents, we introduce a feature called the *inverse stretch factor* 

Dataset	Range	Size	Rows	Features	Source	Frequency	Туре	Description
-	-	-	-	Time of day	derived	-	Temporal	We divide each day into six 4-hour time windows.
-		-		Weekend	derived	-	temporal	A binary feature that denotes weekdays.
Incident	02/01/2017 to 05/01/2020	21MB	80,000	Past Incidents in the last window	derived	-	Spatio-temporal	Number of incidents on the segment in the last time window of 4 hours
				Past Incidents in a day	derived	-	Spatio-temporal	Number of incidents on the segment in the last day
				Past Incidents in a week	derived	-	Spatio-temporal	Number of incidents on the segment in the last week
				Past Incidents in a month	derived	-	Spatio-temporal	Number of incidents on the segment in the last month
Weather	02/01/2017 to 06/01/2020	300MB	1,400,000	Visibility	Weatherbit	1 hour	Spatio-temporal	A measure of the distance at which an object or light can be clearly discerned.
				Wind Speed	Weatherbit	1 hour	Spatio-temporal	Speed of wind.
				Precipitation	Weatherbit	1 hour	Spatio-temporal	Amount of precipitation.
				Temperature	Weatherbit	1 hour	Spatio-temporal	It is the reported temperature.
Traffic	04/01/2017	2017 1.2TB 2020	30,000,000,000	Congestion	derived	5 minutes	Spatio-temporal	Congestion is the ratio of the difference between free flow speed and the current speed to free flow speed
	to			Free Flow Speed	INRIX	5 minutes	spatial	The speed at which drivers feel comfortable if there is no traffic and adverse weather condition.
	12/01/2020			Traffic Confidence	INRIX	5 minutes	Spatio-temporal	A confidence score regarding the accuracy of the traffic data (we collect this directly from INRIX).
Roadways	Static	81MB	80,000	Lanes	INRIX	static	Spatial	Number of lanes for a roadway segment.
				Miles	derived	static	Spatial	Length of a roadway segment.
				iSF	derived	static	Spatial	Inverse scale factor which represents the the curvature of a roadway segment.



Fig. 3: A combination of the length of a curve  $(\overline{AB})$  and the shortest path between the two ends of the curve  $(\overline{AB})$  can be used to denote its curvature

(iSF), that represents the curvature of road segments. We show an example for calculating iSF in Fig. 3. For the segment in consideration (between points A and B), iSF can be calculated as the length the straight line  $\overline{AB}$  divided by that of the curve  $\widehat{AB}$ ).

- 2) Traffic: The correlation of traffic and road accidents is wellestablished [1]. We collected traffic data for each of the road segments through INRIX at a temporal resolution of 5-minute intervals for about three years. Specifically, we retrieved the free flow speed of traffic, the estimated current speed of the vehicles, and the confidence scores of the estimates. Effective congestion can be calculated from our data as the ratio of the difference between the free flow speed and the current speed to the free flow speed.
- 3) *Weather*: Weather is inherently spatial temporal, and can play an important role in accident rates [1]. We collected hourly weather data (temperature, precipitation, visibility, and wind) from 40 different weather stations in and around the state of Tennessee. To use weather data to forecast accidents on a given road segment, we use the weather station that is the closest to that particular segment.
- 4) Incidents: We look at every accident reported in Tennessee from January 2017 to May 2020. Incident data for this project is provided by the Tennessee Department of Transportation (TDOT). Our data consists of approximately 78,000 accidents. The accuracy of the incident data was verified with the Enhanced Tennessee Roadway Information Management System (E-TRIMS).

# V. Approach

We now describe how we design a pipeline to predict roadway accidents in space and time. To begin with, we filter out road segments that exhibit no accidents or extremely small number of accidents over the temporal period in consideration



Fig. 4: Overview of our approach. We extract spatial temporal information from a variety of data sources, focus on heterogeneity not explicitly modeled in the feature space by identifying clusters, perform synthetic sampling to address sparsity, and learn models on incident occurrence.

(about three years). Our analysis is done on 77% of the observed accidents with a total sparsity of 98%. Recall our goal is to learn a function f (section III) that outputs the likelihood of incident occurrence on a road segment conditional on a set of features. A straightforward way to do so is to learn a separate model over each segment. However, such an approach results in overfitting; each segment contributes a relatively small amount of data which ignores structural similarities between patterns of incident occurrence across the entire spatial region in consideration. The other approach is to learn one model for the entire area. However, a universal model fails to capture any heterogeneity that is not explicitly modeled in the feature space. In order to balance these considerations, we try to identify segments that observe similar patterns for incident occurrence [3], [4]. While it is possible to identify distinct spatial regions (hotspots) and learn a separate model for each area, it is possible that there exists generalizable information in the entire area that is spatially invariant. In order to do so, we seek to identify common areas irrespective of spatial contiguity by clustering all the available segments based on their frequency of incident occurrence. In this study, we used the well known k-means algorithm [31] to group the segments into distinct clusters.

Given clusters of roadway segments that share similar patterns of spatial-temporal incident occurrence, learning the patterns is still challenging due to the sparsity of the data. To address this concern we perform synthetic under-sampling and over-sampling to balance our data. However, naive synthetic sampling performs poorly in our case since the relative frequencies of incident occurrence are markedly different among the clusters. Therefore it is impractical to 'balance' data in each cluster in the same manner. To alleviate this, we start with the cluster with the highest frequency of incident occurrence (cluster A, say) and perform synthetic sampling such that the number of positive data points (spatial segments in temporal windows that *have* accidents) is the same as the number of negative data points (spatial segments in temporal windows that *do not* have accidents). Then, we perform synthetic sampling in the other clusters such that the ratio of accidents occurring for any given cluster (cluster B, say) to the frequency in A is the same as in the original dataset.<sup>1</sup>

Clustering and synthetic sampling provide the foundation for learning spatial temporal forecasting models over accident occurrence. We use the following well-known models to this end.

- Logistic Regression: There are two classes of approaches that can be used to forecast the chances of accidents on road segments. First, one can try to model the count of accidents as a binary variable and use well-known count-based regression models like Poisson regression, zero-inflated Poisson regression, and negative binomial regression [1]. The other approach is to treat the occurrence of accidents as a dichotomous output and model the likelihood that *any* accidents occur. We start with the latter, and use logistic regression, which models the log-odds of the probability of incident occurrence as a linear combination of the features *w*.
- 2) Zero-Inflated Poisson Regression: We also use countbased models to model accident occurrence conditional on spatial temporal features. While Poisson regression has been widely used to model accident occurrence [1], hierarchical Poisson models and zero-inflated models have demonstrated significantly improved predictive power [17]. Zero-inflated models can be described as having dual states, one of which is the *normal* state, and the other the *zero* state [32].
- 3) Random Forests: Random forest classifiers are a decision tree ensemble method where each tree is constructed from independently bootstrapped samples [33]. They reduce model variance and are less likely to overfit compared to standard decision trees due to bootstrap aggregation and the use of a random selection of features to split nodes when constructing each tree (called "feature bagging"). In addition to synthetic sampling, random forests can address sparsity using the Balanced Random Forest method [34]. This works by assigning weights to each class inverse-proportionally to their frequency in the dataset, giving a heavier penalty to misclassifying the minority class.
- 4) Artificial Neural Networks: Finally, we also use simple artificial neural networks to learn a model over incident occurrence. Neural networks consist of a set of *layers*, each of which further consists of *neurons* or computing units. The output of each layer is fed as input to the next layer [35]. Each neuron uses a non-linear function

(called the activation function) of the sum of its inputs, and produces an output. The network can be trained by stochastic gradient descent. We use fully connected layers in this study. An important note to practitioners is the noninterpretability of neural networks can be a barrier when deploying systems in the real-world that affect government policies.

It is natural to compare forecasting approaches through metrics like likelihood values on test data, error rates, precision, and recall. However, our conversations with first responders revealed that it is particularly beneficial for them to understand if forecasting models can *rank* roadway segments based on risk. This is intuitive since accurately forecasting the risk at each segment relative to other segments is important for allocating resources. Therefore, besides standard statistical metrics (accuracy, precision, recall, F1-score), we also report the correlation of each model's marginal accident likelihood distribution over space with the real accident distribution. We report both Pearson and Spearman correlation values.

## VI. ALLOCATION AND DISPATCH

The primary purpose of incident prediction models is to make informed resource allocation decisions. However, prior literature rarely evaluates ERM pipelines in their entirety. Our goal in this project is to ensure that helper vehicles and ambulances controlled by the state of Tennessee save vital response time when dealing with accidents. Therefore, we seek to evaluate the impact incident models have on response time outcomes. Our evaluation process is guided by the following steps:

- Understanding existing policies: Through our collaboration with first responders, we first understand how emergency resources are allocated and deployed in practice. We describe this approach below and use it as our baseline.
- 2) Resource Allocation: In practice, discrete location models like the well-known p-median formulation [36], [37], [38], [39] are widely used to allocated emergency resources [1]. A shortcoming of such approaches is that the service time of resources (for example, the time that ambulances are busy responding to accidents) is not taken directly into account in the allocation process. We introduce a novel modification to the p-median problem to design a heuristic approach that bridges this crucial gap.
- 3) **Evaluation**: Using the proposed allocation model, we compare the performance of existing prediction models and our pipeline by creating a black-box simulator that imitates emergency response.

## A. Resource Allocation

In practice, our interaction with first responders revealed that resource allocation is based on identifying *hotspots* of incident occurrence. First, a map based on historical accidents is created. Then a group of experienced engineers determine the location of the responders; typically, responders are placed in areas with the highest historical accident rates.

<sup>&</sup>lt;sup>1</sup>we also show results without synthetic sampling and clustering.

The allocation formulation we use is based on the p-median problem, which is commonly applied to ambulance allocation. The objective of the standard p-median problem is to locate p facilities (i.e. responders) such that the average demandweighted distance between edges and their nearest facility is minimized. One shortcoming of the p-median formulation is that it does not account for responders becoming unavailable when attending to incidents. To address this, we modify the standard p-median formulation by adding a balancing term to the objective function. Intuitively, this balancing term penalizes responders that cover disproportionately large demand compared to other facilities, encouraging multiple responders to congregate near high demand areas. This effect is schematically demonstrated in Fig. 5. In the figure, the values in the cells correspond to the chance of accident occurrence for the location and the green points show the allocated locations of responders (p=2 in this case). By considering  $\alpha = 0$ (alternative a), the problem is equivalent to the simple pmedian formulation, which seeks to minimize the weighted distance between allocations and points of demand. However, by increasing  $\alpha$  (alternative b), the optimizer seeks to avoid assigning high risk cells to a single responder. Formally, we solve the following optimization problem:

$$\min \sum_{i=1}^{|E|} \sum_{j=1}^{|L|} a_i d_{ij} Y_{ij} b_j$$
(1a)

s.t. 
$$\sum_{j=1}^{|L|} Y_{ij} = 1, \quad \forall i \in \{1, \dots, |E|\}$$
 (1b)

$$\sum_{j=1}^{|L|} X_j = p \tag{1c}$$

$$\substack{j=1\\Y_{ij} \leq X_j, \quad \forall i \in \{1, \dots, |E|\}, \forall j \in \{1, \dots, |L|\}$$
 (1d)

$$X_j, Y_{ij} \in \{0, 1\}, \quad \forall i \in \{1, \dots, |E|\}, \forall j \in \{1, \dots, |L|\}$$
 (1e)

where E is the set of demand edges from graph G, L is the set of possible responder locations, p is the number of responders to be located,  $a_i$  is the likelihood of accident occurrence on edge  $e_i \in E$ , and  $d_{ij}$  is the distance between edge  $e_i \in E$ and location  $l_j \in L$ .  $Y_{ij}$  and  $X_j$  are two sets of decision variables;  $X_j = 1$  if a responder is located at  $l_j \in L$  and 0 otherwise, and  $Y_{ij} = 1$  if edge  $e_i \in E$  is covered by a responder located at  $l_j \in L$  (i.e. the responder at j is the nearest placed responder to e) and 0 otherwise. The balancing term we add is denoted by  $b_j = (\frac{\sum_{e \in E} a_e Y_{ej}}{\sum_{e \in E} a_e})^{\alpha}$ , and represents the proportion of total demand covered by a responder located at j. The influence of the balancing term is controlled by the hyper-parameter  $\alpha$ ; intuitively, as  $\alpha$  increases, responders are more 'tightly packed' around high demand areas, and if  $\alpha = 0$ our formulation reduces to the standard p-median formulation. Constraint (1b) expresses that the demand of each edge must be met, (1c) ensures that p responders are located, and (1d) shows that edges must be covered only by locations where responders have been located.

The p-median problem is known to be NP-hard on general networks [40], therefore heuristic methods are employed to



Fig. 5: Illustrating the impact of  $\alpha$ : a) standard p-median ( $\alpha = 0$ ). b) modified p-median with  $\alpha > 0$ . Notice as  $\alpha$  increases responders (green dots) are tightly packed around high demand areas.

Algorithm 1: Greedy-Add Algorithm input : Demand Edges E, Potential Responder Locations L, Segment Incident Likelihoods  $a_i \quad \forall e_i \in E$ , Segment to Location Distances  $d(i, j) \quad \forall e_i \in E, \forall l_j \in L$ , Number of Responders p, Balance Factor  $\alpha$ output: Responder Locations X 1 Initialize  $k := 0, X_k := \emptyset$ ; 2 while k < p do k := k + 1;3 for location  $l_{j'} \in L$ , where  $j' \notin X_{k-1}$  do 4 
$$\begin{split} X_k' &:= X_{k-1} \cup l_{j'}; \\ \text{Find nearest facilities } y_i \; \forall e_i \in E, \text{ where } \end{split}$$
5 6  $y_{e_i} \in X'_k;$  $\begin{array}{l} g_{e_i} \in \mathcal{I}_k^{k_i},\\ \text{Compute balance terms}\\ b_j := (\frac{\sum_{e_i \in E} a_i \psi}{\sum_{e_i \in E} a_i})^{\alpha} \quad \forall l_j \in L \text{ where } \psi := 1 \text{ if }\\ y_i = l_j, \ \psi := 0 \text{ otherwise;}\\ \text{Compute } Z_{j'}^k := \sum_{e_i \in E} a_e d(e_i, y_i) b_{y_{e_i}}; \end{array}$ 7 8 end 9 Best location  $l_j^* := \operatorname{argmin}_j Z_j^k;$ 10  $X_k := X_{k-1} \cup j^*;$ 11 12 end 13 Return  $X_k$ 

find approximate solutions in practice. We use the Greedy-Add algorithm [41] to optimize the locations of responders. We show the algorithm in Algorithm 1. First we initialize the iteration counter k and the set of allocated responder locations  $X_k$  to the empty set (step 1). Then, as long as there are responders awaiting allocation, we iterate through the following loop: (1) update counter k current iteration (step 3), (2) for each potential location not already in the allocation, compute the modified p-median score (equation 1a) of the allocation which includes the potential location (steps 5 - 8), and (3) find the location that minimizes the modified p-median score (step 10) and add it to the set of allocated responder locations (step 11). While myopic, this algorithm is scalable to large allocation problems.

Rather than restricting responders to the roadway segments E, we allow them to be located anywhere across the state. To accomplish this, we define the set of possible responder locations L as a grid of spatial cells over Tennessee. We define each grid cell as being 0.1 degrees latitude by 0.1

degrees longitude, which is approximately 9 km x 11 km in Tennessee. This results in 1445 possible locations across the state. The center of each cell is used when calculating the distance between it and each edge in E.

Given an allocation of responders, we simulate response to real incidents to evaluate the efficacy of our model. Response to emergency incidents is typically greedy [1]; the closest available responder to the scene of the incident is dispatched to attend to it. This a direct consequence of the critical nature of the incidents that emergency responders address. We use a simulator that imitates greedy dispatch and evaluate the performance of different predictive models.

# VII. EXPERIMENTAL EVALUATION

To evaluate our models, we use actual historical incident data, roadway geometry, traffic data, and weather data. We train each model based on a rolling temporal window. We start with a train set of 10 months and use the next month as the evaluation set. Then, we use a rolling temporal window and add one month to the training set and use the subsequent month as the test set. Our code, a synthesized dataset, and high-resolution figures and tables are all available online (see the Appendix).

#### A. Model hyper-parameters

We tune hyper-parameters for each model by crossvalidation. For models based on random forests and neural networks, we keep the architecture fixed based on the largest training sample we have; classification thresholds are tuned for every training window based on a validation set. We describe our model parameters below:

- 1) Random forests: Each random forest consists of 250 decision trees. We use Gini impurity to measure the quality of a node split, and consider  $\sqrt{|w|}$  random features for each split, where w is the total number of features. The following hyper-parameters are tuned for each model: the maximum depth of each tree, the minimum number of observations in a node required to split it, and the minimum number of samples required to be at a leaf node to split it's parent.
- 2) Neural networks: We use a sequential architecture with fully connected layers. We use a total of three hidden layers. The size of the first layer equals twice the number of input features w (the number of neurons in the input layer). The second and third layers consist of neurons equal to the size of the input layer. The output layer conists of a single neuron. We use the 'ReLU' activation function [42] for all hidden layers and the sigmoid activation function for the output layer. We minimize the cross-entropy loss between true labels and predicted labels and use the *adam* algorithm [43] for training the network.
- 3) Clustering: We use the *k-means* algorithm [31] to group the segments into clusters. We use k = 2. A higher value for k renders an extremely small number of segments in some of the clusters, thereby hampering overall performance. Naturally, we recommend practitioners to tune all hyperparameters based on the specific dataset in consideration.

TABLE II: Summary of performance evaluation metrics for each model in percentage (the performance in each column is color coded; green is the best and red is the worst)



# B. Forecasting

We first evaluate the performance of the forecasting pipeline. We refer to the following abbreviations for brevity: LR (logistic regression), NN (neural networks), RF (random forests), ZIP (zero-inflated Poisson), RUS (random under sampling), ROS (random over sampling), NoC1 (No clustering) and KM2 (k-means clustering). Our baseline is based on the actual forecasting model that aids first responders in Tennessee. We refer to the baseline as the naive model. The naive model essentially creates an empirical distribution based on historical incident data. Then, given a segment, a specific point in time, and the set of covariates induced by them, a realization of incident occurrence is sampled from the empirical distribution conditional on the covariates. We present results for each of the approaches in Table II. To understand the role and efficacy of each component of the pipeline, we present results with and without synthetic resampling and clustering.

We present the major observations first — neural networks and random forests outperform the naive model, logistic regression, and the zero-inflated Poisson regression model. Also, based on Table II, we see that while the naive model is fairly accurate, its accuracy is based on under-predicting accidents, as shown by its poor F-1 score. Also, clustering (even in isolation) generally improves the F1-score and accuracy of the forecasting models (for each method, compare the two rows that denote no resampling). We further observe a similar trend with synthetic sampling, which even in isolation usually results in an improvement in accuracy as well as F1-score (for each method, see the set of rows that denote no clustering and compare the rows that show resampling). The efficacy of the combination of clustering and oversampling is somewhat unclear though. We observe that typically, the combination slightly under-performs in comparison to using one of the two approaches. We present three major takeaways from this observation — first, synthetic sampling and clustering enable forecasting in sparse datasets significantly more than approaches that do not use them. However, we recommend practitioners to carefully evaluate each component of the



Fig. 6: Total travel distance of responders per accident (a) 10 responders (b) 15 responders, and (c) 20 responders; for example, for the naive model, when p = 10 and  $\alpha = 0$ , the min., median, and max. of the average travel distance per accident is 2.65 km, 29.57 km, 147.26 km.

TABLE III: Average and maximum number accidents in each 4-h window that responders are not able to immediately respond because all responders are busy (the performance in each column is color coded; green is the best and reed is the worst.) Also, since all values for p = 0 are zero, it is not included in this table.



proposed incident prediction pipeline on unseen data (test set) before deployment. Second, count-based models (zeroinflated Poisson regression) do not perform as well as binary classification models on sparse data. Third, it is important to note that while the resulting F1 scores might seem low in comparison to approaches on other data-driven problems, we claim that the improvement is significant in the context of extremely sparse and inherently random incidents like road accidents. We show the validity of this claim by simulating allocation and dispatch to accidents.

#### C. Allocation and Dispatch

Our final goal is to enable our community partners save crucial response time to accidents. We now discuss how our forecasting models aid allocation and dispatch. We evaluate the entire combination of forecasting and dispatch on 186 temporal windows of 4 hours each, having a total of 1,865 incidents. We also vary the hyper-parameter  $\alpha$  and the number of available responders p (Tennessee currently has 10 responders, but we perform experiments on other realizations of p nonetheless). We use two metrics to evaluate performance the average distance traveled by the responders and the number of incidents that cannot be attended due to unavailability of responders. We present the results in Fig. 6 and Table III.

An important observation is that our allocation approach, which adds a balancing term to the classical p-median problem, improves resource allocation in general; indeed, we observe this improvement in general across the spectrum of forecasting model used and the number of available responders. The maximum improvement we observe is a reduction of 3 km traveled by responders per incident (on average). This observation is particularly important for practitioners and first responders — while it is important to allocate resources in areas with (relatively) high historical rates of occurrence, assigning a small number of responders to cover large areas can be detrimental to the overall goal of reducing response times. Intuitively, our approach penalizes additional burden on responders. However, we also note that a large penalty (value of  $\alpha$ ) can result in increased response times. This is expected — as  $\alpha$  grows, it discourages the geographic spread of responders. Fig.7 shows the influence of  $\alpha$  on the performance of different models with a varying number of responders. Our empirical results show that  $0.5 < \alpha < 1$ results in the optimal allocation of responders.



Fig. 7: Effect the of the hyper parameter  $\alpha$  in allocation model on the performance of different models (the grey lines denote each learned model).

We also observe that our forecasting pipeline results in **significant** savings in response times (upto a maximum long-term reduction in response times by 10%) (see Figure 6)

and a reduction in the number of accidents that cannot be attended to (see Table III). It is important to understand the importance of this reduction. Prior work reports that a saving of only ten minutes of response time can reduce deaths due to road accidents by 33% [44]. Allocating responders based on the random forest model provides the maximum reduction in response times (with neural networks being a close second). We provide three major takeaways based on our experiments on allocation and dispatch. First, forecasting models that provide the highest accuracy might not be the best candidates for allocation. This observation shows the importance of using a metric that focuses on false negatives and false positives (like the F-1 score) for sparse emergency incidents. Second, we point out that while traditional allocation models based on long-term (temporal) hotspots are widely used in the field, accurate short-term forecasting models can result in significant reduction in response times to accidents. Finally, leveraging the structure of the problem to improve classical resource allocation formulations can aid emergency response in the field.

#### VIII. CONCLUSION

Emergency response to incidents like road accidents is a major concern for first responders. Standard approaches to predicting road accident rarely scale to large geographic areas due to extremely high sparsity in data and difficulties in gathering data. In collaboration with the Tennessee Department of Transportation, we present a framework for forecasting extremely sparse spatial and temporal incidents like road accidents. We show how our approach for forecasting, based on a combination of non-spatial clustering, synthetic resampling, and learning from multiple data sources, outperforms forecasting methods used in the field. We also present a novel modification to a classical formulation for resource allocation. Through extensive simulations, we show how our pipeline results in significant reduction in response times to emergency incidents. Our implementation is opensource and can be used by other organizations that seek to optimize emergency response.

#### IX. ACKNOWLEDGEMENT

We acknowledge the support from the National Science Foundation (IIS-1814958), the Tennessee Department of Transportation (TDOT) for funding the research, and Google for providing computational resources.

#### REFERENCES

- A. Mukhopadhyay, G. Pettet, S. Vazirizade, D. Lu, S. E. Said, A. Jaimes, H. Baroud, Y. Vorobeychik, M. Kochenderfer, and A. Dubey, "A review of incident prediction, resource allocation, and dispatch models for emergency management," 2021.
- [2] "Disaster sequence of events," https://training.fema.gov/emiweb/ downloads/is208sdmunit3.pdf, accessed: 2018-04-01.
- [3] A. Mukhopadhyay, C. Zhang, Y. Vorobeychik, M. Tambe, K. Pence, and P. Speer, "Optimal allocation of police patrol resources using a continuous-time crime model," in *International Conference on Decision* and Game Theory for Security, 2016, pp. 139–158.

- [4] A. Mukhopadhyay, Y. Vorobeychik, A. Dubey, and G. Biswas, "Prioritized allocation of emergency responders based on a continuous-time incident prediction model," in *International Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 168–177.
- [5] A. Mukhopadhyay, G. Pettet, C. Samal, A. Dubey, and Y. Vorobeychik, "An online decision-theoretic pipeline for responder dispatch," in *International Conference on Cyber-Physical Systems*, 2019, pp. 185–196.
- [6] G. Pettet, A. Mukhopadhyay, M. Kochenderfer, Y. Vorobeychik, and A. Dubey, "On algorithmic decision procedures in emergency response systems in smart and connected communities," in *International Conference on Autonomous Agents and Multiagent Systems*, 2020, pp. 1046– 1054.
- [7] G. Pettet, A. Mukhopadhyay, M. Kochenderfer, and A. Dubey, "Hierarchical planning for resource allocation in emergency response systems," *arXiv preprint arXiv:2012.13300*, 2020.
- [8] J. A. Deacon, C. V. Zegeer, and R. C. Deen, "Identification of hazardous rural highway locations," *Transportation Research Record*, vol. 543, 1974.
- [9] W. Cheng and S. P. Washington, "Experimental evaluation of hotspot identification methods," *Accident Analysis & Prevention*, vol. 37, no. 5, pp. 870–881, 2005.
- [10] J. Eck, S. Chainey, J. Cameron, M. Leitner, and R. Wilson, "Mapping crime: Understanding hot spots," National Institute of Justice, Tech. Rep., 2005.
- [11] M. Deublein, M. Schubert, B. T. Adey, J. Köhler, and M. H. Faber, "Prediction of road accidents: A Bayesian hierarchical approach," Accident Analysis & Prevention, vol. 51, pp. 274–291, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/ pii/S0001457512004101
- [12] M. A. Quddus, "Modelling Area-Wide Count Outcomes with Spatial Correlation and Heterogeneity: An Analysis of London Crash Data," *Accident Analysis and Prevention*, vol. 40, pp. 1486–1497, 2008.
- [13] D. Akin, "Analysis of highway crash data by negative binomial and poisson regression models." International Symposium on Computing in Science and Engineering (ISCSE), 2011.
- [14] F. Ladron de Guevara, S. P. Washington, and J. Oh, "Forecasting crashes at the planning level: Simultaneous negative binomial crash model applied in tucson, arizona," *Transportation Research Record*, vol. 1897, no. 1, pp. 191–199, 2004.
- [15] W. Ackaah and M. Salifu, "Crash prediction model for two-lane rural highways in the ashanti region of ghana," *IATSS Research*, vol. 35, no. 1, pp. 34–40, 2011.
- [16] X. Qin, J. N. Ivan, and N. Ravishanker, "Selecting exposure measures in crash rate prediction for two-lane highway segments," *Accident Analysis & Prevention*, vol. 36, no. 2, pp. 183–191, 2004.
- [17] D. Lord, S. Washington, and J. N. Ivan, "Further notes on the application of zero-inflated models in highway safety," *Accident Analysis & Prevention*, vol. 39, no. 1, pp. 53–57, 2007.
- [18] H. Huang and H. C. Chin, "Modeling road traffic crashes with zeroinflation and site-specific random effects," *Statistical Methods & Applications*, vol. 19, no. 3, pp. 445–462, 2010.
- [19] D. Lord, S. P. Washington, and J. N. Ivan, "Poisson, poisson-gamma and zero-inflated regression models of motor vehicle crashes: Balancing statistical fit and theory," *Accident Analysis & Prevention*, vol. 37, no. 1, pp. 35–46, 2005.
- [20] A. Pande and M. Abdel-Aty, "Assessment of freeway traffic parameters leading to lane-change related collisions," *Accident Analysis & Prevention*, vol. 38, no. 5, pp. 936–948, 2006.
- [21] H. T. Abdelwahab and M. A. Abdel-Aty, "Artificial neural networks and logit models for traffic safety analysis of toll plazas," *Transportation Research Record*, vol. 1784, no. 1, pp. 115–125, 2002.
- [22] L. Y. Chang, "Analysis of freeway accident frequencies: negative binomial regression versus artificial neural network," *Safety Science*, vol. 43, no. 8, pp. 541–557, 2005.
- [23] C. Riviere, P. Lauret, J.-F. Ramsamy Manicom, and Y. Page, "A Bayesian neural network approach to estimating the energy equivalent speed," *Accident Analysis and Prevention*, vol. 38, no. 2, pp. 248–259, 2006.
- [24] L. Zhu, F. Guo, R. Krishnan, and J. W. Polak, "The use of convolutional neural networks for traffic incident detection at a network level," in *Transportation Research Board Annual Meeting*, 2018.
- [25] J. Bao, P. Liu, and S. V. Ukkusuri, "A spatiotemporal deep learning approach for citywide short-term crash risk prediction with multi-source data," *Accident Analysis & Prevention*, vol. 122, pp. 239–254, 2019.

- [26] Y. Zhang and Y. Xie, "Forecasting of short-term freeway volume with v-support vector machines," *Transportation Research Record*, vol. 2024, no. 1, pp. 92–99, 2007.
- [27] X. Li, D. Lord, Y. Zhang, and Y. Xie, "Predicting motor vehicle crashes using support vector machine models," *Accident Analysis and Prevention*, vol. 40, no. 4, pp. 1611–1618, 2008.
- [28] R. Yu and M. Abdel-Aty, "Utilizing support vector machine in realtime crash risk evaluation," *Accident Analysis & Prevention*, vol. 51, pp. 252–259, 2013.
- [29] Y. Qi, B. L. Smith, and J. Guo, "Freeway accident likelihood prediction using a panel data analysis approach," *Journal of transportation engineering*, vol. 133, no. 3, pp. 149–156, 2007.
- [30] "INRIX." [Online]. Available: https://inrix.com/
- [31] J. MacQueen *et al.*, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, no. 14, 1967, pp. 281–297.
- [32] C.-S. Li, J.-C. Lu, J. Park, K. Kim, P. A. Brinkley, and J. P. Peterson, "Multivariate zero-inflated poisson models and their applications," *Technometrics*, vol. 41, no. 1, pp. 29–38, 1999.
- [33] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [34] C. Chen, A. Liaw, and L. Breiman, "Using random forest to learn imbalanced data," University of California, Berkeley, 2004.
- [35] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.
- [36] A. B. Calvo and D. H. Marks, "Location of health care facilities: an analytical approach," *Socio-Economic Planning Sciences*, vol. 7, no. 5, pp. 407–422, 1973.
- [37] D. Serra and V. Marianov, "The p-median problem in a changing network: the case of barcelona," *Location Science*, vol. 6, no. 1-4, pp. 383–394, 1998.
- [38] M. Dzator and J. Dzator, "An effective heuristic for the p-median problem with application to ambulance location," *Opsearch*, vol. 50, no. 1, pp. 60–74, 2013.
- [39] L. Caccetta, M. Dzator *et al.*, "Heuristic methods for locating emergency facilities," in *Proceedings Modsim*. Citeseer, 2005.
- [40] O. Kariv and S. L. Hakimi, "An algorithmic approach to network location problems. i: The p-centers," *SIAM Journal on Applied Mathematics*, vol. 37, no. 3, pp. 513–538, 1979.
- [41] M. S. Daskin, *Network and discrete location: models, algorithms, and applications.* John Wiley & Sons, 1995.
- [42] A. F. Agarap, "Deep learning using rectified linear units (relu)," arXiv preprint arXiv:1803.08375, 2018.
- [43] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.
- [44] R. Sánchez-Mangas, A. García-Ferrrer, A. De Juan, and A. M. Arroyo, "The probability of death in road traffic accidents. how important is a quick medical response?" *Accident Analysis & Prevention*, vol. 42, no. 4, pp. 1048–1056, 2010.

#### APPENDIX

We have made the code used in this study publicly available at the following  $link^2$ . While detailed instructions on how to use the code can be found on the link, there are a few key files worth mentioning.

- **config/params.conf**: This file contains metadata and parameters that must be customized to a user's specific deployment. This includes configuration information such as file paths and dataset information. It also configures which models to run, model hyperparameters, the features to use in regression, clustering parameters, and the synthetic sampling to apply. More detail on the specific metadata can be found in the readme file.
- run sample.py: This is the main script for fitting and evaluating the forecasting models. Based on the metadata in the *params.conf* file, this script loads and formats the data, calls sub-routines for clustering, synthetic sampling, tuning model specific hyper-parameters, and finally fit the desired models for each sliding test window. It then evaluates each model on a test set and outputs the following result files to the output/ directory: DF\_Results.pkl pandas dataframe which contains the overall evaluation metrics (accuracy, etc.) for each of the models averaged over space and time, DF\_Test\_spacetime.pkl - dataframe containing the models' predictions for each 4-hour time window (used for resource allocation), report.html html file which visualizes the evaluation results from the DF\_results file. There is an example of the html results visualization provided in the repository.
- allocation/run\_allocation.py: This script evaluates the models' impact when integrated with an allocation model. It uses the output from *run\_sample.py* as prediction inputs. Given a set of test incidents, the script performs allocation based on each model's prediction output (*DF\_Test\_spacetime.pkl*), and then simulates dispatch to calculate the distance between incidents and their nearest

The data used in this study is proprietary, but we release a synthesized example dataset (in the *sample\_data* folder) to demonstrate the expected data format. The data is formatted as a csv document; each row represents the features and incident counts for a 4 hour time window at a particular road segment. The specific feature names are detailed on the link provided.

responders.

<sup>&</sup>lt;sup>2</sup>https://github.com/StatResp/smart-comp\_IncidentPrediction