A Review of Emergency Incident Prediction, Resource Allocation, and Dispatch Models

Ayan Mukhopadhyay*, Geoffrey Pettet†, Sayyed Vazirizade†,
Yevgeniy Vorobeychik‡, Mykel Kochenderfer*, Abhishek Dubey†
*Stanford University {ayanmukh, mykel}@stanford.edu
†Vanderbilt University {geoffrey.a.pettet,s.m.vazirizade,abhishek.dubey}@vanderbilt.edu
‡Washington University in St. Louis {yvorobeychik}@wustl.edu

***Abstract*—Emergency response is one of the most pressing problems faced by communities across the globe. In the last fifty years, researchers have developed statistical, analytical, and algorithmic approaches for designing emergency response management (ERM) systems. In this survey, we present models for incident prediction, resource allocation, and dispatch concerning urban emergency incidents like accidents and crimes. We highlight the strengths and weaknesses of prior work in this domain and explore the similarities and differences between different incident types. Finally, we present future research directions. To the best of our knowledge, our work is the first comprehensive survey that explores the entirety of ERM systems.**

I. Introduction

Effective emergency response management (ERM) is a challenge faced by communities across the globe. First responders need to respond to a variety of incidents, such as fires, traffic accidents, and crimes. They must respond quickly to incidents to minimize the risk to human life [1, 2]. Consequently, considerable attention in the last several decades has been devoted to studying emergency incidents and responses. Data-driven models can help reduce both human and financial loss, as well as improve design codes, traffic regulations, and safety measures. Such models are increasingly being adopted by government agencies. Nevertheless, emergency incidents still cause thousands of deaths and injuries, as well as result in losses worth more than billions of dollars directly or indirectly each year [3]. This is in part due to the fact that emergency incidents (like accidents, for example) are perhaps inevitable in the modern world, and also because of the mismatch between the number of incidents and the number of responders.

The overall pipeline for ERM can be divided into four major components: 1) mitigation, 2) preparedness, 3) response, and 4) recovery [4, 5]. Mitigation involves sustained and continuous efforts to ensure safety and reduce long-term risks to people and property. It also involves understanding *where* and *when* incidents occur and designing predictive models of both risk and spatial-temporal incident occurrence. Preparedness involves creating infrastructures that enables emergency response management. This stage involves selecting stations for housing responders, ambulances, and police vehicles as well as designing plans for response. The third phase, arguably the most crucial, involves dispatching responders when incidents happen or are about to occur. Finally, the recovery phase ensures that the broader community or impacted individuals can cope with the effects of incidents. While much prior work in ERM has studied these problems independently, these stages are inter-linked. Frequently, the output of one stage serves as the input for another. For example, predictive models learned in the *preparedness* stage are used in planning *response* strategies. Therefore, it is crucial that ERM pipelines are designed keeping the intricate inter-dependencies in mind. In this survey, we cover prior work on some of the most widely explored approaches that fall into the categories of mitigation, preparedness, and response, and we explain how the overall ERM pipeline functions.

One way to categorize emergency incidents is by their frequency of occurrence. The first kind involves the more frequent incidents and addressing them is a part of day to day operations of first-responders. Examples of such incidents include crimes, accidents, calls for medical services, and urban fires. The second category consists of the comparatively less frequent incidents, which include natural calamities like floods and cyclones. While response management to disasters is an active area of research and extremely important for communities, we focus on principled approaches to address frequent urban incidents.

Our primary reason to focus on urban emergency incidents is simply the alarming frequency of their occurrence. Globally, about 3,200 people die every day from road accidents alone, leading to a total of 1.25 million deaths annually [6]. In fact, it is noted that without appropriate measures, road accidents are set to be the fifth largest cause of death worldwide by 2030 [7]. Calls for emergency medical services (EMS) are also a major engagement for first-responders, and there are more than 240 million[1] such calls made annually in the United States alone [8]. Therefore, it is imperative that we design principled approaches to understand the spatial and temporal characteristics of such incidents and investigate algorithmic methods that can mitigate their effects.

Another important and frequent type of incident that plagues urban areas are crimes. While crimes share common characteristics with other urban incidents in some ways, they are a fundamentally different problem in others. For example, similar to accidents, once a crime incident is reported, response must be dispatched as soon as possible. Further, responding to

[1]This includes all calls made to the emergency number 911.

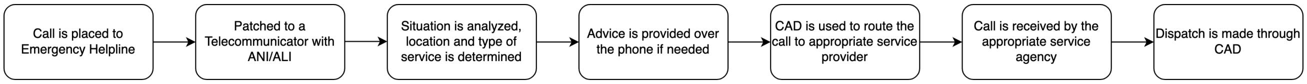

Figure 1: Typical Emergency Dispatch Helpline Model

crimes also suffers from a shortage of resources. For example, the United States experiences over a million violent crimes and over eight million property crimes a year, but it has only nine-hundred thousand law enforcement personnel in total [9]. The major difference between crimes and other emergency incidents is that the former is caused by deliberate and planned actions by individuals or groups, unlike the latter. The United Nations office on drugs and crime presents a detailed report about the current crime situation in the world and highlights that for many countries, crimes related to homicides, drugs, burglary, and robbery are on the rise [10]. Also, emergency responders are often shared in urban areas; for example, it is common for police vehicles to attend to accidents in conjunction with ambulances. Consequently, it is crucial to understand the requirements and dynamics of ERM systems pertaining to crimes, along with accidents and EMS calls.

Our goal is to review existing work on urban emergency incidents and understand commonalities and differences among them in order to provide a unified perspective on ERM systems. There are comprehensive reviews on crash prediction models [11, 12, 13], emergency facility location approaches [14] as well as dispatch strategies [15]. In particular, the doctoral thesis of Kiattikomol [13] and the work by Lord and Mannering [16] provide particularly insightful summaries of crash prediction models. There are detailed reviews on crime prediction approaches as well [17, 18, 19, 4, 20, 21]. However, to the best of our knowledge, there is no comprehensive study that links prediction models from different incident types like accidents and crimes, investigates covariates of relevance, and discusses planning approaches comprehensively. We treat the ERM system in its entirety and relate predictive models with algorithmic approaches in mitigation and planning. This survey provides a framework for future research on integrated emergency incident pipelines for smart and connected communities.

II. System Model

We study the problem of optimally responding to emergency incidents in urban areas. Incidents are reported to central emergency response agencies, which have streamlined mechanisms for processing the request. For example, in the United States, the emergency helpline calls are placed by dialing 911. We show the steps that follow such a call in Fig.1 [22]. The call is appended with automatic name and location information (ANI/ALI), and patched to a trained telecommunicator. The telecommunicator analyzes the situation and the type of response needed (police, EMS, or fire, for example). In some cases, such as those requiring cardiopulmonary resuscitation (CPR), guidance might be provided through the phone before first-responders reach the scene. The call is then transferred to the concerned agency (such as the police or fire department) by a computerized mechanism. The agency then uses its computer-aided dispatch (CAD) system to dispatch a responder to the scene. This set of events defines an ERM system, and it governs the pipeline of incident response, including detecting and reporting incidents, monitoring and controlling a fleet of response vehicles, and finally dispatching responders when incidents occur. In many cases, there could be more than one central body governing this pipeline for an urban area; for example, ambulances and police cars might be dispatched from different authorities.

The agents, who respond to incidents like crimes and accidents include ambulances, police vehicles, and fire trucks, among others. We refer to such agents as responders. Responders are typically equipped with devices that facilitate communication to and from central control stations. In many cases, especially in the US, responders like ambulances are equipped with computational devices like laptops as well. Once an incident is reported, responders are dispatched by a human agent to the scene of the incident (guided by some algorithmic approach like a CAD system). This process typically takes a few seconds.[2] but can be longer if dispatchers are busy.

Each responder is allocated to a specific *depot*, which are stations located at various points in the spatial area under consideration. Once a responder has finished servicing an incident, it is directed back to its depot and becomes available to be re-dispatched while en route. An exception to this paradigm is patrol vehicles, which are deployed on specific routes to deter crimes. A key aspect, that plays an important role in dispatch algorithms is that if there are any free responders available when an incident is reported, then one must be dispatched to attend to the incident. This constraint is a direct consequence of the bounds within which emergency responders operate, as well as of the critical nature of the incidents. If an incident takes place and there are no free responders available, then the incident typically enters a waiting queue and is attended to when a responder becomes free.

The components of ERM that we focus on are shown in Fig. 2. ERM pipelines typically use data from historical incidents and environment, including weather, geometry of roads, traffic patterns, and socio-economic data. We divide an ERM system into four major types: 1) predictive models about incident occurrence, 2) models for environmental features like traffic and weather, 3) allocation models to optimize the spatial locations of responders and stations, and 4) dispatch models to create algorithmic approaches to respond to incidents when they occur. The components are intricately linked, and the performance of each component plays a crucial role in the performance of the overall ERM pipeline.

Incident prediction models form the basis of an ERM system. In order to mitigate the effects of incidents, it is important to understand *where* and *when* such incidents occur. Incident models are typically designed using historical incident data, but such models often use historical environmental data

[2]This is based on our communication with fire departments in the United States [23]; time taken to dispatch responders presumably varies across the globe.

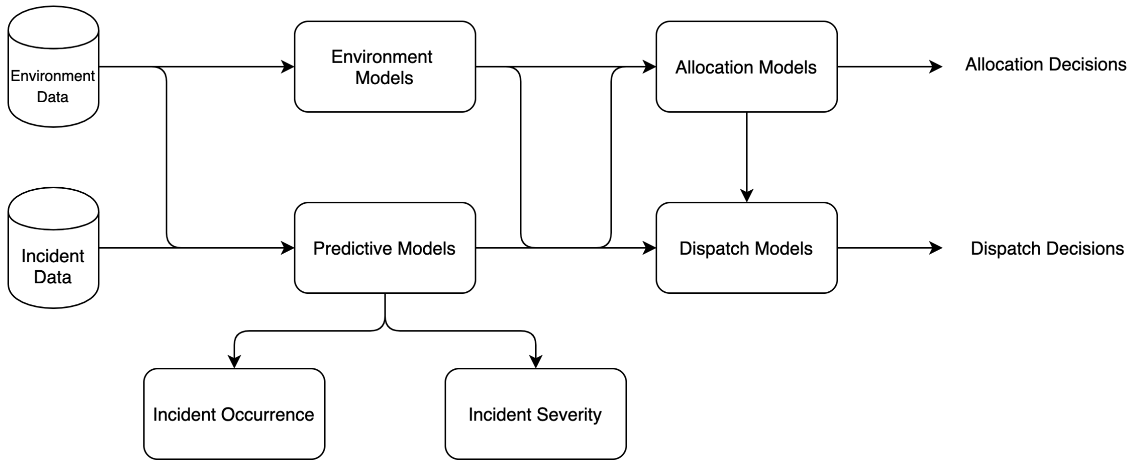

Figure 2: ERM System Pipeline

as well; for example, it is common for crime prediction models to use socio-economic data and for accident prediction models to use traffic data. Allocation models are then used to allocate responders in time and space in anticipation of future incidents. Finally, allocation and prediction models are used to create dispatch models, which can be thought of as a policy that guides real-time response.

Significant prior work has focused on understanding and designing algorithmic approaches for each of the modular components. This article studies models for incident prediction, allocation, and dispatch. While we do not discuss models of relevant environmental factors, they are important to the development of the overall pipeline.

We focus this survey on two major types of incidents: accidents and crimes. The reason for such a categorization is two-fold. First, most prior work in incident analysis has focused on accidents and crimes. Second, these two types of incidents exhibit the characteristics of emergency response in general. Accidents represent the category of incidents where EMS services are essential and efficient response is mandatory. There is also no strategic interaction between the the person involved in the accident and first-responders. Crimes represent incidents where the strategic interaction is plausible and secondary objectives of prevention and deterrence are important. Much of our discussion on accidents can be broadened to EMS calls in general, but focusing on one particular type allows us to discuss various technical approaches involved in greater detail.

III. Formulation

We start by defining the formulation of incident prediction and planning problems that we use throughout this survey. Given a spatial area of interest S 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 ith incident, respectively, and $w_i \in \mathbb{R}^m$ represents a vector of features associated with the incident. These features can be spatial, temporal, or spatio-temporal as discussed in section IV-B . The features capture covariates that potentially affect incident occurrence. For example, w typically includes features such as weather, population density, and time of the day. The most general form of incident prediction can then be stated as learning the parameters θ of a function over 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 *count* of incidents (the number of incidents in S during a specific time period) or *time* between successive incidents. The decision-maker seeks to find the *optimal* parameters θ^* that best describe D. This can be formulated as a maximum likelihood estimation (MLE) problem or an equivalent empirical risk minimization (ERM) problem.

We review prior work focused on modeling the function $f(X \mid w, \theta)$. There have been many different approaches for modeling f. It can be modeled as an explicit probability density or mass (e.g., Poisson distribution), or a function that does not strictly conform to such definitions (e.g., a linear regression approach with X being the dependent variable. Nonetheless, such functions typically have probabilistic interpretations, and we present different approaches to modeling f in section IV. We highlight different modeling choices for both accidents and crimes and highlight similarities and differences. Then, we focus on the vector w. Arguably, the most crucial part in learning a model over incident occurrence involves choosing w, and we review various covariates in section IV-B.

The next step in an emergency response pipeline is to plan in anticipation of incidents. This involves stationing responders spatially and dispatching them as incidents occur. This process can be broadly represented by the optimization problem $\max_y G(y \mid f)$, where y represents the decision variable (which typically denotes the location of emergency responders in space), G is a reward function chosen by the decision-maker, and f is the model of incident occurrence. For example, G might measure the total coverage (spatial spread) of the responders, or the expected response time to incidents. Therefore, given f, the decision-maker seeks to maximize the function G. While this formulation accurately represents planning models for accidents, response against crimes is tricky because police allocation affects the distribution of crime. In other words, the decision variable y affects f, but f determines the choice of y. This circular dependency makes it challenging to deploy police units. We focus on this challenge and show how it has (or at times has not) been tackled in prior work.

There are two major paradigms for modeling the response problem. First, the planning problem can be represented as a stochastic control process. For example, the planning problem can be formulated as a Markov decision process (MDP) [24]. This formulation is particularly relevant for problems seeking to find policies for dispatch. The aim is to find an optimal policy (control choices for every possible state of the system) that maximizes the expected sum of rewards. The other approach is to directly model the planning problem as an optimization problem according to a specific measure of interest (for example, a lot of prior work has focused on maximizing the coverage of emergency responders).

IV. Incident Forecasting

We divide the discussion on incident forecasting into three major parts – a) approaches to incident prediction, b) predicting incident severity, and c) features used in incident prediction.

A. Approaches to Incident Prediction

Prior work has involved learning spatial-temporal models of incident occurrence. From our definition of incident prediction models described in section III, forecasting models correspond to the function f. We review literature from accident prediction and crime prediction separately, and then identify similarities and differences.

1) Accident Prediction: An important method in incident prediction is known as 'crash frequency analysis', which uses the frequency of incidents in a specific discretized spatial area as a measure of the inherent risk the area possesses [25]. Deacon, Zegeer, and Deen [25] identified key questions that practitioners should answer while designing predictive models for incident occurrence, and their work is still relevant to decision-makers and policy designers. This approach also forms the basis of *hotspot* analysis [26, 27], which is widely used today as a relatively simple and fast method to visualize incident data. A shortcoming of frequency analysis is that it neglects fluctuations in incident occurrence, and requires a large volume of incident data to infer accurate characteristics of occurrence [28, 29]. Nonetheless, the core idea behind frequency analysis continues to be in use today; although it is common to use it in conjunction with other covariates of relevance and frame the overall problem as a regression model.

One of the earliest regression models used to model incident occurrence involved multiple linear regression models with Gaussian errors [30, 31]. However, modeling accident count by linear regression can be inaccurate, as the response variable is discrete and strictly positive. In addition, it has also been shown that linear regression models fail to model the sporadic nature of emergency incidents [32, 33]. Linear regression models with multiplicative effects have also been investigated but have shown to be inaccurate compared to other models [32]. The inaccuracies of linear regression methods in the context of accident prediction is investigated and summarized by Miaou and Lum [32]. Rakha et al. [34] revisited this problem recently, and used data aggregation techniques to satisfy assumptions made by linear regression. While such an approach has shown performance on par with other regression models (Poisson regression, for example), the authors admit that it needs further validation before it is widely adopted.

The inaccuracies of linear regression and the suitability of Poisson models for count data led to the widespread use of Poisson regression for modeling incident data [31]. Each incident is considered a result of an independent Bernoulli trial. Given that all the trials are generated by the same stochastic process, the series of trials can be modeled by a binomial distribution. As the number of trials becomes large and the probability of success is very small, the probability distribution over the count of incidents takes the form of a Poisson distribution [35]. To accommodate the feature vector w, Poisson regression assumes that the logarithm of the expected value of the distribution is a linear combination of w. This methodology has been used extensively for emergency incident analysis [36, 37, 38, 33, 32].

An issue with using Poisson regression is that the expected value of the response variable (count of incidents) equals its variance. This is typically not the case with crash data, which is over-dispersed, meaning that the variance of the data is greater than its mean [35]. Also, there are examples of incident data being under-dispersed as well [39]. Therefore, the broader argument against the use of Poisson regression is that it can accommodate neither under-dispersed nor over-dispersed data. An approach to accommodate over-dispersion is to use Poisson-hierarchical models [40]. Poisson-hierarchical models (as well as Poisson models) fall under the broader category of generalized-linear models (GLM), which is a family of distributions used widely in statistics and machine learning. From this family, the Poisson-gamma (also called negative binomial) and Poisson-lognormal models are particularly relevant. The Poisson-gamma is a Poisson distribution whose mean parameter follows a gamma distribution. It has been shown that the Poisson-gamma model fits crash data better than Poisson models, and it has been extensively used for crash prediction [41, 42, 43, 44, 45, 46]. While the Poisson-gamma model solves the problem of over-dispersion, it performs poorly on under-dispersed data and is particularly problematic to use with small sample sizes and with data with low sample mean [47, 48]. The Poisson-lognormal model is the same as Poisson-gamma model, but it uses the lognormal distribution for the mean parameter rather than the gamma distribution [49, 50, 51, 52]. The lognormal distribution is a heavy tail distribution and provides more flexibility for over-dispersion. Recently, the Poisson-inverse-gamma model has been used in crash modeling [53]. However, such models do not have closed-form MLE solutions unlike Poisson-gamma models [16].

Despite the success of Poisson and Poisson-hierarchical models, a common shortcoming is that both models fail to adequately handle the prevalence of zero counts in crash data [35]. A remedy to this problem is to use zero-inflated models, and both zero-inflated Poisson and zero-inflated Poisson-gamma models have been used to model accident data [54, 55, 56]. Zero-inflated models can be described as having dual states, one of which is the *normal* state, and the other the *zero* state. The excess zeros that cannot be explained by standard count-based models can then be considered to have arisen due to the presence of a separate state. Zero-inflated models result in improved statistical fit to accident data. Lord, Washington, and Ivan [35] note that most prior works justify the use of zero-inflated models by improved likelihood, and therefore automatically assume that crash data is generated by a dual-state process (except work by Miaou and Lum [32], which uses a zero-inflated model to justify misreporting of incidents). Through empirical data and simulations, they show that excess zeros could arise due to various other factors like low traffic exposure and the choice of spatial and temporal scales by the model designer. As a result, it is not clear if the statistical backing to using dual-state models is accurate or not. In our opinion, the work by Lord, Washington, and Ivan [35] is particularly profound, and the argument that statistical fit should not be the only consideration for fitting models to crash data (and other data in general) is extremely cogent.

A somewhat different approach in predicting emergency incidents is to directly model inter-incident time as a function of relevant covariates. In this case, the variable X corresponds to the time between consecutive incidents. Mukhopadhyay et al. [57] describe an example of such models by using uncensored (parametric) survival models to estimate time between accidents. It has been since used to model different incident types [58, 59, 60]. A key advantage of such methods is that planning problems are often modeled as continuous-time processes, and as a result, the incident prediction models can be easily used by planning models.

While time-based models are not the most commonly used approaches to model the occurrence of crashes, continuous-time models are often used to predict the duration of crashes and the delay that crashes cause in traffic and congestion. While estimating traffic delay is crucial to the overall planning problem [61, 62, 63, 64, 65], it is outside the scope of this paper.

Bayesian methods [66, 67] are often used for parameter estimation. Such models result in a distribution over parameters rather than point estimates, which can result in greater robustness to outliers and small sample sizes [68]. The empirical Bayes method (also known as maximum marginal likelihood) has been used in traffic engineering [69, 70, 71, 72] (the method as applied to crash prediction is explained particularly well by Hauer et al. [73]). Bayesian modelling techniques have also been used to assess potential risk factors of spatial regions [74, 58] and to estimate expected crash frequencies [75].

Hierarchical Bayesian estimation (also known as full Bayesian models) of safety performance models have also been explored over the last two decades [50, 51, 47, 76, 77, 78]. Recently, the Poisson-gamma and Poisson-lognormal models have also been estimated using Bayesian methods [41, 42, 79, 43, 44, 45, 46, 49, 52, 53]. A caveat regarding Bayesian models is the role that the choice of priors play in the predictive models. The underlying information for designing priors might be available from previous models, engineering judgement, etc., and prior distributions can also be chosen to be non-informative or weakly informative. An important investigation in this context, specifically regarding crash prediction, was done by Song et al. [80], who study the performance of various Bayesian multivariate spatial models with different prior distributions. It was shown that using the non-informative prior may result in a high bias for the dispersion parameter for small numbers of observations [81].

With improved sensor technology and easier storage, data-mining methods have successfully been used for crash prediction. Random forests [82, 83], support vector machines [84, 85, 86], and neural networks [87, 88, 89, 90] have recently been used to model crashes. Bayesian neural networks have also been explored, specifically to account for over-fitting of neural-networks in crash modeling [91]. Deep learning techniques have also been used in various studies [92, 93]. One specific model that is of interest to practitioners was developed by Bao, Liu, and Ukkusuri [93], who used a spatio-temporal convolution long short-term memory network (LSTM) to predict short-term crash risks, including propagation of traffic congestion [94]. While the network structure was a combination of various complex networks, the accuracy of hourly predictions was limited, which highlights the inherent difficulty of predicting crash frequency at low temporal and spatial resolutions. It also makes a case against the use of complex models in this domain because are harder to generalize.

Dynamics of urban environments change frequently. As a result, it is important that such changes are taken into account by forecasting methodologies. This consideration applies to all models created for emergency incidents. Recently, the development of online models for predicting accidents has been explored, which can work with an incoming stream of data and update the model continuously based on new information [59].

2) Crime Prediction: Crime prediction has generated substantial interest over the last decades. There are two major categories of models in this context. The first seeks to measure the likelihood of crime occurrence given a set of environmental (spatio-temporal) features. The second seeks to predict the likelihood that a specific individual is likely to be an offender. We choose to focus only on the former category; the latter is outside the scope of this survey. Predictive policing has faced numerous ethical issues in the recent years. In 2016, an investigative journalism agency ProPublica reported about the inherent bias that can arise in predictive policing tools. It spoke about one such tool that had been used to identify the likelihood of individuals committing future crimes [95]. While this was disputed by the designers of the algorithm, it created a channel of widespread discussion and analysis on how such algorithmic approaches need to be evaluated. The ethical issues regarding predictive policing are well-summarized in prior work [96, 97, 98]. Mukhopadhyay et al. [99] recently showed that algorithmic approaches to predictive policing (even the ones that model the likelihood of a set of environmental conditions to be susceptible to crimes) can increase the likelihood of police interaction with citizens. It is important that such effects be carefully considered before implementing policies. In this survey, our focus is solely on the technical aspects of predictive methods.

Approaches to model the likelihood of crime occurrence can further be sub-divided in four groups: 1) purely spatial models, which identify spatial features of previously observed crime, such as hot spots (or crime clusters), 2) spatial-temporal models, which attempt to capture dynamics of attractiveness of a discrete set of locations on a map, 3) risk-terrain models, which identify key environmental determinants (risk factors) of crime and create an associated time-independent risk map, and 4) game-theoretic models, which seek to identify the strategic interaction between people who intend to commit crimes and law enforcements authorities [4, 100, 101]. In each of the these categories, extensive research has been made to understand the occurrence of crimes.

Spatial techniques have been extensively used to identify hot-spots of criminal activities. Levine et al. [102] present an accurate sub-categorization and an extremely detailed review of spatial methods used in crime prediction. We use the same ideas but present a more coarse grouping. A straightforward way to find hot-spots is to discretize the spatial area of interest S, and then identify locations that have the most number of reported incidents. Such an approach is called point-based clustering. It is also possible to find hot-spots dynamically by maintaining an exogenous search parameter that lets the decision-maker decide the granularity of the clusters [103, 104]. Another way to find hot-spots involves clustering the spatial area under consideration by a partitioning algorithm, like the well-known k-means algorithm [102, 105, 106, 107, 100, 108]. While the k-means algorithm has been used extensively in crime mapping, it suffers from two well-known issues. First, the number of clusters must be specified beforehand. Second, the algorithm performs poorly at identifying non-convex shaped clusters [109]. As a result, density based methods such as the density based spatial clustering algorithm (DBSCAN) and kernel density estimation (KDE) have been used for crime mapping [110, 111, 112]. Hierarchical clustering has also been used to learn hot-spots, in which smaller clusters are aggregated into larger clusters iteratively based on an appropriate similarity measure [57]. Crime counts in discretized regions can also be modeled using count data, which seek to learn a distribution for the number (or frequency) of crime incidents, similar to accidents. Both Poisson and negative-binomial models have been used in this regard [113, 114, 115].

An alternative approach, risk-terrain modeling, seeks to identify and study quantifiable environmental factors as determinants of spatial crime incidence, rather than looking at crime correlation [116]. Risk Terrain Modeling started as a tool to identify behavioral settings for crimes in the city of New York but has been adopted by many law enforcement agencies to combat crime. Caplan and Kennedy [117] analyse specific risk factors for fourteen different types of crimes, along with specific case-studies from the field. It is a particularly useful resource for practitioners, not only for implementing risk-terrain models but also to identify useful covariates for other types of models.

A limitation of approaches that focus only on spatial mapping is that they ignore the temporal dynamics of crimes. Prevailing theories of crime suggest strong temporal correlation between crimes. For example, the well-known *repeat victimization* describes elevated risks of crime incidents following an initial incident [118, 119], and the theory of *broken-windows* suggests that tangible signs of past crime occurrence result in increased risk of future occurrences [120]. As a result, it is crucial to take into account the temporal dynamics of crimes. To this end, there has been significant work to create approaches that consider the spatio-temporal dynamics of crimes in its entirety.

An important branch of such models was presented by Short et al. [121], who proposed using a spatio-temporal differential equation model to capture spatial and temporal crime correlation. Later, self-exciting point-processes were also used to capture spatio-temporal clusters in crimes [122]. Leading indicator models have also been used in this regard, which identified temporal and spatial correlation with historical data to predict future crimes [123]. A disadvantage of such models is that they do not naturally capture crime covariates. One way to model the spatial-temporal patterns in crimes is to identify spatial and temporal separately and then use the combined model. This general paradigm was used to create the dynamic spatial disaggregation approach (DSDA) [124]. This approach combines an autoregressive model to capture temporal crime patterns, and spatial clustering techniques to model spatial correlations. A model recently proposed by Mukhopadhyay et al. [100] combines hierarchical clustering and parametric survival analysis to learn a continuous-time model over crime occurrence. It outperformed DSDA and game-theoretic approaches and is fairly intuitive to understand.

An orthogonal approach to crime prediction involves studying the strategic interaction between law enforcement authorities and people who commit crimes, formulating the crime prediction problem as a game. The paradigm of Stackelberg games [125] has been used extensively in crime prediction. Stackelberg games incorporate a *leader-follower* model, which makes it particularly suitable for modeling crimes. In such a model, it is assumed that the *defender* allocates resources (typically police patrols) first, and the *follower* (people with malicious intention of committing crimes) observes the defender's strategy and plans accordingly. Stackelberg games have been used to deploy air marshals in flights [126], protect biodiversity in conservation areas [127], and screen passengers in airports [128]. An extension to this paradigm, known as green security games, models the repeated interaction between criminals and law enforcement agencies by extending the leader-follower paradigm to multiple rounds [129, 130]. In such games, the attacker behavior in previous rounds can be used to make better policing decisions in subsequent rounds. This notion of strategic interaction can also be used to capture how criminals potentially respond to arbitrary predictive models [4]. Recently, the field of robustness in crime prediction methods has also received attention. A potential issue with forecasting crimes is that people with malicious intentions of committing crimes can potentially change their preference over spatial locations in response to deployed patrols. Mukhopadhyay et al. [99] present a principled framework for ensuring that arbitrary predictive models (with a convex likelihood function) are robust against such shifts. However, such models can potentially increase the interaction of police with citizens, and it is important to evaluate the effects of such approaches before deploying them in practice.

B. Feature Selection

1) Features in Accident Prediction: An important part of developing predictive models is feature engineering. The accuracy of models depends highly on the selected features, and as a result, they should be chosen strategically. Features for accident prediction can be categorized into temporal, spatial, or combination of both. For example, one can choose to use time of day as a feature in order to understand how it affects accident rates. This is an example of a temporal feature. The geometry of a specific road segment, on the other hand, is a spatial feature, as it is a characteristic property of a particular spatial unit. Spatio-temporal features measure spatial properties that change with time. For example, traffic congestion in a specific part of the city falls under this category since it is characterized by both space and time. Generally, the features available for crash analysis are restricted to the information on the crash report, weather and environmental conditions, roadway geometry, and traffic information. It is also possible to categorize features into static or dynamic [131], but we choose to follow the categorization with respect to spatio-temporal characteristics of the features.

i) *Temporal Features:* Weather [132, 57, 131] and visibility range [133] have been proven to be useful in predicting accident rates, especially features like fog, rain, and snow. Weather data can also include seasonality features, temperature, light, etc. Time of day and day of week are also important predictors of accident rates [57, 134, 131].

ii) *Spatial Features:* Roadway geometry is also known to be an effective predictor of crash frequency [135, 53, 136, 137]. The most commonly used features in this regard are the number of lanes, width of the lanes, features regarding shoulders, horizontal turns and slopes [51, 138], the presence of uncontrolled left-turn lane, the presence of bus stops, median widths, speed limit [53, 134], and features specific to intersections [134, 137]. Road infrastructure [41] and socio-economic features [139] are also studied to be important such as density of the bars in the region.

iii) *Spatio-Temporal Features:* Crashes exhibit strong spatial-temporal incident correlation. Past incidents are an important predictor of future incidents. For example, areas that have typically experienced a relatively high concentration of incidents in the recent past are more likely to have incidents in the future [57, 60, 59]. Traffic congestion also plays a crucial role since its combination with other features may cause different effects. For example, traffic congestion naturally increases the likelihood of one specific type of accident (rear-end crash) [140], while there have been studies showing that congestion has no or negative effect on crash frequency [141, 142]. The other features, which fall into this category and are different representations of congestion, include peak hour [131], traffic volume [91], and average speed of vehicles [140].

2) Features in Crime Prediction: Similar to accidents, feature engineering in crime prediction is also crucial. We use a similar categorization for features used in crime prediction. As with crime prediction models, covariates used in such models can be divided into features pertaining to individuals and features pertaining to environmental conditions. We only focus on the latter since the former is outside the scope of this paper.

i) *Temporal Features:* The correlation between weather and crime occurrence was studied as early as the beginning of the twentieth century [143], and it has since been studied extensively [144, 145, 146, 147, 148, 100]. The effect of weather has been shown to differ on different types of crimes. Lauritsen and White [148] study this systematically, and their work is particularly relevant for policy makers and model designers. Time of day is also commonly used as a covariate to predict crime [149, 150, 100]. There are counter-examples of this effect as well; for example, Bernasco, Ruiter, and Block [151] found no effect of time of day on street crimes in the city of Chicago. We recommend that model designers evaluate the effects of specific covariates on their region of interest.

ii) *Spatial Features:* The effect of socio-demographic variables on crime occurrences is well-explored. This includes the effect of establishments like liquor availability [152, 153, 154, 100], presence of pawn shops [155, 156, 157], and homeless shelters [158, 100, 159], as well as socio-economic features like population density, income levels, and unemployment rate [160, 161, 162, 100]. The well-known approach of *Risk Terrain Modeling* (RTM) seeks to use geographic mapping with specific spatial risk covariates like the presence of bars, foreclosures, etc. [163]. RTM has been successfully used in multiple cities in the United States.

iii) *Spatial-Temporal Features:* The most important spatial-temporal feature used in crime is locations and times of previous incidents. This feature is ubiquitous, and almost all prediction approaches model future crime occurrence using correlation with historical crime data. Another feature of relevance is police presence, but its use in prediction models is somewhat rare. This is intuitive, since one of the goals of prediction algorithms is to deploy police patrols. As a consequence, using the decision variable of the overall ERM pipeline as part of the underlying prediction algorithm creates a circular dependency. Nonetheless, it has been used in crime prediction models; for example, Mukhopadhyay et al. [100] use police patrol data in different spatial-temporal resolutions to predict future crimes.

An important consideration in using features in prediction models is that one must design appropriate prediction models for the features themselves. For example, consider the role of weather in predicting accidents. In order to design effective policies, the decision-maker must be able to forecast accidents, which makes it important to forecast weather. Accurate models over features are immensely important in practice, but this form of forecasting is beyond the scope of this survey.

C. Incident Severity

Prediction of severity of incidents is usually defined in the context of accidents and crashes. Severity of accidents plays a crucial role in planning approaches for allocation as well as for dispatching resources when incidents occur. Naturally, decision-makers plan to prioritize incidents with higher severity over the ones with relatively lower severity. Since it is difficult to gauge the severity of an incident based on a call for assistance, it is common in practice to dispatch the responder closest to the scene of the incident. However, understanding spatial and temporal patterns in severity and its relationship with incident occurrence models is crucial in planning. Understanding covariates that affect severity, and creating models for predicting severity of crashes have attracted a lot of attention. While there are different definitions of severity, it can usually be categorized into five levels: 1) no-injury or just property damage, 2) possible injury, 3) non-incapacitating injury, 4) incapacitating injury, 5) and fatal [164]. Most of the prior work in severity prediction has focused on using similar ordinal categorization of severity. Savolainen et al. [164] present a detailed review regarding severity of accidents, which is self-contained, complete, and comprehensive. Much of this section is informed by their work; we identify crucial insights from it and also focus on models that have been introduced since then.

Let incident severity be represented by the random variable K. From the perspective of the formulation in section III, designing models for incident severity can be represented in two ways. First, there is significant work on creating marginal models over severity. These models have the form $h(K \mid w, \theta)$, where h is a distribution over K, w is a set of covariates that impact incident severity, and θ denotes the model parameters. Note that w could include information about the crash itself, such as information from post-crash reports. The other approach is to model a joint distribution that governs incident occurrence and the resulting in severity. In this scenario, given incident data, the decision-maker seeks to learn a joint distribution over incident occurrence and severity, which can be represented by $h(X, K \mid w, \theta)$.

The relationship between traffic flow and accident severity is well-explored [165, 166, 167, 168, 169]. Crash severity has been explored using multinomial logit and probit models [170, 171, 172, 173], decision trees [174], random forests [175, 176, 177], and neural networks [178, 179].

One natural way to account for correlation between crash frequency and severity is to learn an independent regression model for each category of severity. Multiple regression models [180, 181, 80] as well as neural networks [182] have been used to this end. Although such a paradigm captures inherent correlation (to some extent) between incident arrival and severity, it does not model an explicit joint distribution. Mukhopadhyay et al. [57] present an approach that forms a bridge between marginal and joint models. They assume that the joint distribution can be decomposed into a marginal distribution over incident arrival, followed by a conditional distribution over severity given incident arrival.

In the last two decades, there has also been significant interest in jointly modeling incident arrival (frequency) and severity [181, 51, 50, 75, 183]. This includes multivariate Poisson regression [181] and multivariate Poisson log-normal regression models [50]. Pei, Wong, and Sze [183] model the joint distribution explicitly and use a fully Bayesian approach to learn the model. While such models are promising, a crucial (potential) limitation is identified out by Savolainen et al. [164]. Jointly modeling crash arrival and severity limits the use of data related to the specific crash while learning the model. On the other hand, marginal models can use detailed post-crash data to infer insights about severity [164].

Finally, there are two orthogonal directions of work in severity prediction that can be combined with both marginal models or joint models. The first approach is rather recent and focuses to identify spatial relationships between different levels of severity [139]. The other approach seeks to tackle inherent heterogeneity in crash data by identifying clusters of incidents (not necessarily spatial) to better understand the relationship between crash data and covariates [184, 185, 186].

D. Key Takeaways

Having discussed prediction models in context of both accidents and crimes, we now summarize key takeaways. There are two major differences between predictive models for the two categories of incidents. First, in practice, it appears that there is a significantly greater focus on spatial models that create risk maps for crime prediction than in accident prediction. We hypothesize that this is primarily due to the simplicity and use of such models in designing patrol routes. Another major difference is the lack of strategic modeling in accident prediction. This is fairly straight-forward, since accidents are not caused by deliberate planning, there is no need for strategic models in the context of crash prediction.

Despite the differences, there are key similarities too. There are modeling paradigms that work well for both crimes and crashes. First, arrival models (distributions over frequency, count or inter-arrival time) over a discretized spatial area, such as Poisson regression, negative-binomial models and survival analysis have been used widely in both contexts. Online models are needed in all forms of emergency response, since urban environments change frequently. The use of data-mining is increasingly gaining more popularity in both fields. Also, hierarchical clustering has shown to balance spatial heterogeneity and model variance while predicting both crimes and accidents. Finally, the choice of covariates has been shown to be of utmost importance.

We recommend practitioners, model designers, and planners to:

1) Be aware of advances made in predictive modeling in the context of different types of incidents.
2) Seek the help of domain experts (researchers, fire-fighters, policemen, etc.) to design the feature space w, which is a crucial factor in the performance of predictive models.
3) Start by using well-defined paradigms that have been shown to work on multiple datasets, and are backed by assumptions that are statistically sound.
4) Be aware of flaws and shortcomings of models, and carefully evaluate the possible costs of inaccurate predictive models.
5) In case of crime prediction, be aware of inherent biases in historical crime data and evaluate the consequences of

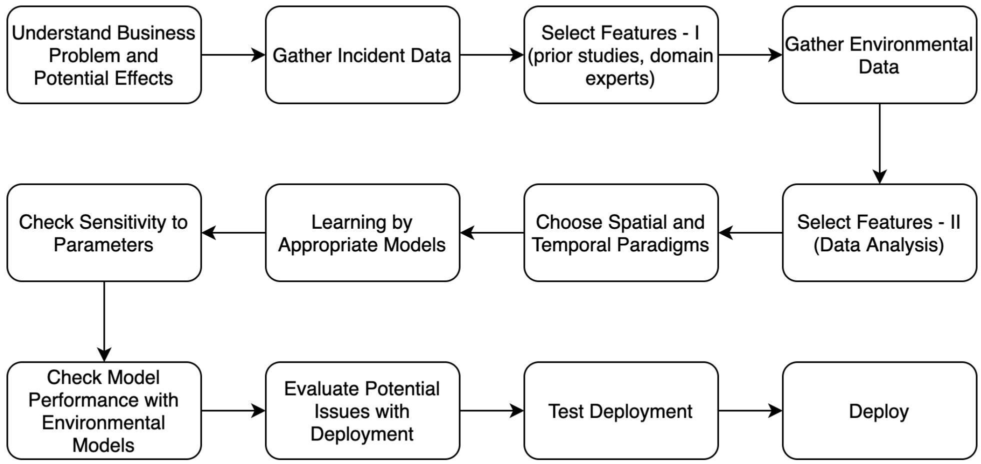

Figure 3: Incident Prediction Model Design Pipeline

using such models that increase the interaction between citizens and law-enforcement authorities.

We combine our experience in designing ERM pipelines with prior work in this field and summarize steps that practitioners and model designers should take in Fig 3.

V. Responder Allocation and Dispatch

There are two crucial steps in an ERM system that come into effect *after* the decision-maker gains an understanding of when and where incidents happen. These involve allocating resources (also referred to as the stationing problem [187]) in expectation of incidents and dispatching resources when calls for service are received. While prediction problems are primarily formulated as *learning* problems, allocation and response are commonly modeled as optimization problems. As discussed in section III, an allocation or response problem can be represented as $\max_y G(y \mid f)$, where y represents the decision variable, G is a reward function chosen by the decision-maker, and f is the model of incident occurrence. For allocation problems, y typically refers to the location of emergency responders in space. For response problems, the decision variable is a mapping between responders and specific calls for service.

The distinction between allocation and response problems can be hazy since the solution to the allocation problem implicitly creates a policy for response. For example, consider an algorithm that ambulances have been allocated to stations across the city in a manner that minimizes expected response times to incidents according to an incident arrival model f. Now, when an incident occurs in the jurisdiction of a specific station, naturally, a responder (if available) is dispatched from the station, without the need for an explicit dispatch model. While this is generally true for allocation models, there are finer subtleties involved. As noted by Mukhopadyay et.al. [57, 60, 59], implicit response strategies are not always optimal. For example, consider a situation where an incident occurs close to a station that has no available responders. Should the incident enter a waiting queue? How does the potential severity of the concerned incident affect this decision? If a nearby station has a free responder, should it be dispatched? How do response time guarantees from the allocation model change in such scenarios? Answering such questions is critical for an efficient ERM system. This section discusses algorithmic approaches to both allocation and response.

A. Allocation and Response – EMS

We first introduce the metrics used to allocate emergency response stations and responders. The three most common metrics are coverage [188, 189, 190], distance between facilities and demand locations [59], and patient survival [191, 192, 193]. *Coverage* measures the proportion of spatial locations that are within some predefined distance of the responders (or depots). It is measured with respect to demand nodes, which are discretized spatial units that can potentially generate calls for service. Of the three metrics, it is the most straightforward to examine as it is generally binary. The demand node is considered covered by some facility if it is within the predefined distance, and otherwise considered to be uncovered. It also lines up well with the broader objective of many EMS providers, which is to limit the number of calls that are responded to *late*, i.e. that have a response time higher than some threshold (the distance often serves as a proxy for the response time, for example see Mukhopadhyay et al. [57]). These factors contributed to coverage being a prevalent metric in early EMS allocation research.

The distance between potential demand nodes and their nearest facilities is another metric that can be used for optimization of the spatial distribution of stations and responders. These metrics are more difficult to use since they are not binary, but recent advances in computational capability have made them more accessible. Both coverage and distance to potential demand locations actually approximate the true objective of EMS policies, which is increasing patient survival. Erkut, Ingolfsson, and Erdoğan [191] argued that it is more appropriate to use expected patient survival directly by incorporating a survival function that captures the relationship between response times and survival rates.

Most early ERM allocation approaches modeled the allocation problem as an integer or linear optimization problem [188, 189, 190]. These models are relatively straightforward and can be solved by a large body of optimization techniques. Exact methods such as branch-and-bound have been applied to small instances of the problem [194, 195] but do not easily scale to realistic environments. As a result, most prior work relies on heuristic approaches, such as genetic algorithms [196, 197] and tabu search [190, 197, 198, 199]. Recently, decision theoretic models such as Markov decision processes (MDPs) have gained traction as efficient solution methods have evolved [200, 59].

Early allocation approaches also generally tackled *static* allocation. Facilities are assumed to be immobile, so the model determines the optimal locations for the facilities without allowing for temporal redistribution. In such models, responders are often used synonymously with facilities. The two seminal static facility allocation models are the Location Set Covering Problem (LSCP) [188] and the Maximal Covering Location Problem (MCLP) [189]. Both models have similar assumptions, including that stations act independently, response is deterministic, that at most one ambulance is at each facility, and that there is one type of ambulance. The primary difference between the two is in the optimization objective. LSCP finds the least number of facilities that *cover* all demand nodes, while MCLP maximizes the demand covered by a given number of facilities. LSCP can be useful for planning a lower bound on the number of facilities needed for a given coverage standard, while MCLP better captures the constraints of real world use cases where the number of facilities is heavily constrained by cost. It is also common to introduce constraints on secondary objectives like waiting times in optimization problems that seek to maximize coverage. For example, Silva and Serra [201] and Mukhopadhyay et al. [57] define optimization frameworks for maximizing coverage with upper bounds on waiting times, and can accommodate different levels of incident severity.

There are a number of extensions to LSCP and MCLP, many of which relax some of their strong assumptions. Aly and White [202] consider a spatially continuous demand model, rather than the discrete demand nodes. Jia, Ordóñez, and Dessouky [196] introduce different quality levels for facilities (which can represent each facility's available services or equipment), with demand points having different coverage constraints for each level. Erkut, Ingolfsson, and Erdoğan [191] incorporated a survival function into the optimization function of MCLP which maps response times to survival rates.

LSCP, MCLP, and many of their extensions all have a common shortcoming in that they assume deterministic system behavior in regards to response. Resources at a facility are considered to be always available, and the models assume that a station is able to service all demand nodes that it covers. In the real world, there are finite resources at each station, and calls from a specific demand node might need to be answered by a station other than the closest one. For example, it is common for other stations to respond to a call if the closest one is busy. One way to address this is by increasing the number of stations that cover each demand point, i.e. using a *multiple coverage* metric.

A key example is the Double Standard Model (DSM) [190], which incorporates two distance standards r_1 and r_2, where $r_1 < r_2$. The model adds the constraint that all demand must be covered within r_2, similarly to LSCP, ensuring that each point has *some* coverage. It also specifies that some proportion α of the demand is covered within r_1. Given those constraints, the objective is to maximize the demand covered *by at least two stations* within r_1. Essentially, this maximizes the demand nodes that have nearby facilities while ensuring that all demand nodes have adequate coverage. While this approach helps mitigate the issue of station unavailability, there can still be situations where both facilities covering some demand point are busy. Accounting for such situations requires modeling facility availability explicitly.

There is a large body of research on probabilistic coverage models, which model the stochastic nature of station availability. Two foundational probabilistic models are the Maximum Expected Covering Location Model (MEXCLP) and Maximum Availability Location Problem (MALP). MEXCLP was introduced by Daskin [203] and extends MCLP, modifying the optimization function to account for station availability. It assumes that each facility has the same probability of being busy, which simplifies computation but does not accurately represent the real world where facilities near incident hot spots are unavailable for a greater proportion of the time. Also, it inherits many of the assumptions of MCLP, and assumes that facilities act independently. MALP, proposed by Revelle and Hogan [205], maximizes the demand covered by facilities with some exogenously specified probability. The first version, MALP-I [205] is similar to MEXCLP in that it assumes equal probabilities for being busy for facilities. MALP-II [205], however, removes this assumption. The proportion of time that facilities are busy is computed as a ratio between the total demand generated by demand points and the availability of facilities covering them.

There have been several extensions to the above probabilistic models to relax some of their simplifying assumptions and make them better match the real world. TIMEXCLP, developed by Repede and Bernardo [206], introduces temporal variations in travel times between points to MEXCLP. Adjusted MEXCLP (AMEXCLP) [207] relaxes MEXCLP's assumption that facilities are independent by treating them as servers in a hypercube queuing system [208] with equal busy factions. The Queuing Probabilistic Location Set Covering Problem (QPLSCP) [195] makes a similar extension to MALP by computing each individual facility's busy fraction using a queuing model and feeding them into MALP-II.

An alternate approach to modeling allocation and response problem is to model the problem as a stochastic control problem, and then optimizing over the set of control choices to maximize expected reward. The most commonly used model in this regard is the Markov decision process (MDP). A variety of models and approaches have been explored in this space. Keneally, Robbins, and Lunday [209] modeled the optimal dispatch problem as a continuous-time MDP, and used canonical policy iteration to solve the problem. A shortcoming of such a model is that it assumes memoryless transitions, which reduces the computation of state transitions to closed-form expressions. Real-world transitions are not necessarily memoryless, and this was addressed by Mukhopadhyay et.al. [60], who formulate the problem as a semi-Markovian decision problem (SMDP) instead, and use a simulator to estimate the transition probabilities. However, it does not scale to real-world problems. An approach to alleviate this problem is to focus on finding an action for the current state of the world instead of aiming to find a policy for the entire state-space [59].

Recently, a potential shortcoming of algorithmic dispatch approaches has been pointed out which is important to ponder over. Based on conversations with first responders, Pettet et al. [187] point out that the moral constraints in emergency response dictate that the nearest responder be dispatched to the scene of an incident. This observation explains why algorithmic approaches to response often do not get implemented in the field. Pettet et al. [187] create an approach to optimize over the spatial distribution of responders *between* incidents, while always dispatching the closest available responder to attend to incidents. This process alleviates two major issues. First, it does not waste crucial time *after* an incident has occurred to optimize over which responder to dispatch. Second, the moral constraint of always sending the closest responder to an incident is not violated. It remains to be seen if such an approach gets accepted by first-responders and is tested in the field.

B. Allocation and Response – Crimes

Response pertaining to crime can be broadly split into two categories: response to specific service calls and response to create deterrence [210, 211]. This combination of proactive policing and reactive response makes police patrolling particularly difficult. The consequences of proactive policing are debatable though; experimental results from two famous studies done in Kansas City [212] and Newark [213] showed that varying police presence had no effects on crime rates. Nonetheless, proactive patrolling is a major undertaking of most police departments [214, 211, 215] because of two reasons. First, proactive policing helps with the maintenance of police presence and enhances the role of police and helps in secondary objectives like recover stolen automobiles and maintain traffic regulations [216]. Secondly, non-experimental studies have demonstrated the use of proactive policing in reducing crime rates [217, 218].

We discuss responder placement with respect to policing in both ways. We start by discussing algorithmic approaches to proactive policing, with the goal of deterrence. Then, we discuss allocation algorithms to better prepare responders to service calls efficiently.

Early work in using algorithmic approaches to create deterrence focused on maximizing the probability of police patrols intercepting a crime in progress. Chelst [219] used heterogeneous weights and duration for different types of crimes and solved the resulting optimization problem by iteratively assigning patrols to regions ranked by the objective function. Olson and Wright [210] formulate patrol dynamics in a specific region (or street segment) as a Markov chain, and use the formulation to maximize interceptions.

An alternate approach to proactive policing is called *hot-spot policing*. It specifically seeks to direct proactive policing towards clusters of crime that show clearly elevated frequency of incidents than others [220, 221]. Hotspot-policing is arguably the most widely technique used for of proactive policing today. Indeed, a survey conducted in 2008 in the USA revealed that about 90% of the police departments used some form of hotspot-policing (a total of 176 departments were surveyed) [220, 222]. Braga et al. [220] present a detailed summary of hotspot-police policing and highlight that there is strong evidence that supports its efficacy in dealing with clusters of crimes. In general, the use of GIS techniques and visual analytics has increased in proactive policing [223, 224, 225].

Similar to predictive models for crime, an orthogonal approach to crime prediction involves modeling the strategic interaction between patrols and criminals. As mentioned in section IV, the most commonly used game-theoretic formulation in this context is the Stackelberg game model, which has been widely used to create patrol policies [226, 227, 228]. Several variations exist for such models as well. For example, patrolling security games [229] take into account the possibility that resources could be mobile; green security games are played over multiple rounds and take into account information availability for crimes like poaching and illegal fishing [230, 231, 232]. Models for opportunistic crimes have also been explored in this regard, which use dynamic Bayesian networks to model the interaction between criminals and patrols [233, 234].

There are some important caveats that must be alluded while discussing proactive policing. First, such an approach to policing might result in diffusion of crimes to nearby areas [220], which calls for creating robust models of incident prediction [99]. Secondly, the possible adverse effects of increased interactions of police with citizens have also been widely discussed and studied [235, 236, 237, 238, 239, 240, 241]. We refrain from discussing this in detail, since this review is specifically focused on the algorithmic aspects of ERM systems. Nonetheless, we urge practitioners to carefully consider citizens' expectations and possible effects of proactive policing before implementing it in practice.

The second type of police response is reactive. Patrols need to respond to specific calls for service. From an algorithmic perspective, a response of this kind is very similar to ambulances responding to accidents but has some important differences. First, ambulances often need to transfer people affected by accidents to hospitals, which might not be the case for police patrols. This is a constraint that must be taken into account while planning response strategies for ambulances; police patrolling algorithms can relax the constraint. Second, ambulances responding to accidents do not affect future distribution of accidents, while police patrols might impact the future distribution of crimes. While this constraint is usually not taken into account while planning police patrols, there are exceptions.

A report prepared by Chaiken and Dormont [216] for government bodies pertaining to urban development in the United States is one of the earliest works on police response. It provides a nuanced treatment for patrol allocation, and has served as a building block for future work (for example, [57]). It considered the effects of service calls being missed due to unavailability and the use of queuing models, down-times experienced by responders, and the effect of different levels of severity on police patrols. The use of multi-server queuing models have also been explored to model police dispatches, since in practice, it is common for more than one police car to attend to an incident [242]. The most widely used approach in designing reactive police patrols has primarily focused on expected response time to incidents, by explicitly minimizing it, placing acceptable upper bounds on time to service [243], or evaluating risk of violating response time guarantees [244]. The explicit effect of reactive patrols on future crime distribution has also been considered. Mukhopadhyay et al. [100] model the response problem as a two-stage optimization problem, which is solved by iterative stochastic programming.

C. Key Takeaways

Allocation and response models are a crucial component of ERM pipelines, and a variety of algorithmic approaches have been used for allocating responders in anticipation of crimes and accidents. The most apparent difference between stationing responders for crimes and accidents is the consideration of secondary objectives like deterrence. Ambulances, by the sole virtue of their presence, cannot deter or prevent accidents. On the other hand, a major engagement of police departments is to perform proactive patrolling to deter crime. Patient survival is also a vital consideration that ambulances need to take into account while designing response models, since ambulances need to transport patients to medical facilities, which in turn increases the overall service time. This effect is naturally manifested in the choice of objectives and variables for allocation models. Despite this difference, there are high-level similarities in response modeling that apply to all emergency incidents (especially in reactive response). Models focusing on increasing coverage and reducing wait-times are common objectives that have been widely used in practice. We recommend model designers and practitioners to:

1) Be well-versed with the different objectives that have been used in response and allocation models, and carefully choose the one that suits the specific needs of the concerned area.
2) Seek the help of domain experts (researchers, fire-fighters, policemen, etc.) to understand problems that responders face in the field. For example, the nearest ambulance might be heading in the opposite direction from the demand node on a highway, without the scope of making a turn. This makes it important to consider features that might not be intuitive to researchers.
3) Seek to bridge the gap between theoretical models and realistic environmental constraints. For example, there is a rich body of work that makes the assumption that the environment in which ERM systems operate is static. While such assumptions simplify computational challenges, they might not truly capture the dynamics of actual ERM pipelines.
4) Static models fail to take into account the changing dynamics of urban areas. As a consequence, there is a need to create online models for emergency response.
5) Be aware of flaws and shortcomings of models, and carefully evaluate the possible costs of inaccurate predictive models.
6) In the case of crime patrolling, be aware that proactive policing increases the interaction between police and citizens. It is crucial that the effects of such models be considered and evaluated.

VI. Challenges and Opportunities

The field of designing emergency response pipelines has seen tremendous growth in the last few decades. Several factors have contributed to this growth. Wider availability of data, the development of data-driven methodologies, increased cognizance, dependence and trust over algorithmic approaches by governments, and increase in computational power are a few reasons for this growth. However, there are still challenges in this field that need to be addressed. As we have pointed out, an EMS pipeline consists of an intricate combination of several components for its smooth functioning. There is a need for more research groups to: i) study EMS pipelines in their entirety, and consider the broader impact of their modular work on ERM systems, ii) consider and acknowledge the challenges and constraints that first responders face in the field, and iii) iteratively develop ERM tools by having first responder organizations in the loop. There are nuances that describe such needs throughout this paper. For example, an improved statistical fit for the prediction models does not necessarily mean an overall improvement for the ERM pipeline if the underlying model does not capture the true dynamics of incident occurrence. There is also a need for researchers to make their data and tools available to both the research community and ERM organizations. In a comprehensive review of statistical methods of crash prediction, Lord and Mannering [16] pointed out that the wider availability of data is extremely promising for the field of crash prediction. This is particularly true now. Vast volumes of real-time data are now available from electric scooters, automobiles, ambulances, and police patrols. There is also wider coverage of sensors like video-cameras throughout urban areas. This promise of increased availability of richer data holds true not only for incident data but also for data regarding covariates that potentially affect incident occurrence, like traffic congestion. The net result of an increased stream of data promises a finer understanding of the effect of covariates on incident occurrence. This benefit can be utilized by sharing data and algorithmic approaches between research groups and first responders.

Urban dynamics of crashes and crimes are continuously changing, and both fields hold opportunities. The increase in the number of automobiles and the arrival of autonomous vehicles in the markets across the globe presents the scope of re-evaluating existing models of crash occurrence and designing newer models that accommodate the changing landscape. Litman [245] lists the various additional planning constraints that need to be taken into account while developing transit systems that can accommodate autonomous vehicles, as well as additional causes for crashes, like software failure and increased overall travel volume. The potential risk factors caused by the interaction between autonomous and non-autonomous vehicles also pose challenges [246] and the need to design newer models of incident prediction.

Crimes also evolve continuously. However, we think that the biggest challenge in crime prediction comes from understanding and evaluating the effects of predictive policing, and possible bias that such methodologies present. The AI Now Institute, an organization dedicated to evaluate and understand the social implications of artificial intelligence pointed out the no public safety organization should use black-box AI models due to ethical concerns [247]. This calls for work in interpretable models of prediction, rather than black-box models [248]. Another potential area of future work involves understanding the robustness of crime prediction models, and more modeling as well as empirical trials are required to better understand the benefits as well as shortcomings of such approaches.

Incident response also poses fresh challenges and opportunities. First, there is a need to combine the different metrics used in designing dispatch and allocation models. There are several interesting threads of research (cooperative coverage, survival metrics, gradual coverage decay, incorporating multiple resource types with different functionalities, etc.) that, to the best of our knowledge, have not been combined and evaluated together. Also, there has not been much focus on explicitly incorporating measures of patient survival directly in response models. We think that it is crucial that patient survival be studied in more detail and included as a part of objective functions for optimization approaches used in designing allocation and dispatch systems.

A recent development in emergency response systems has been the computational ability of agents. Most modern ambulances and police vehicles are now equipped with laptops, which presents the scope of fast and decentralized decision-making, a particularly exciting area for multi-agent systems. Decentralized decision-making has been explored in the context of urban ERM systems [58], but such approaches are probably more relevant for disaster scenarios like floods and hurricanes, where agents might lose connectivity to the central decision-making authority. Algorithmic approaches to aid the strategic redistribution of responders between incidents is extremely promising. While post-incident planning presents many technical challenges, such approaches rarely get implemented in the field. Inter-incident planning, on the other hand, respects the inherent challenges that emergency response faces. As urban areas grow and witness a rise in population density, the need to design principled approaches to aid emergency response grows as well. This survey identifies how the field has evolved over the last few decades, with the view to aid researchers, policy-makers and first-responders in designing better ERM pipelines.

VII. Acknowledgements

We would like to thank the Center of Automotive Research at Stanford (CARS), the National Science Foundation (Grant: CNS-1640624, IIS-1814958) and the Tennessee Department of Transportation (TDOT) for funding the research. We would also like to express our gratitude towards the Nashville Fire Department (NFD), the Metropolitan Nashville Police Department (MNPD) and the Tennessee Department of Transportation (TDOT) for invaluable feedback and domain expertise that made us understand the subtleties and intricate challenges of emergency response. Finally, we would like to thank El Said, Hiba Baroud, Saideep Nannapaneni and Hemant Purohit for giving feedback about the paper.

References

[1] Henrik Jaldell. “How important is the time factor? Saving lives using fire and rescue services”. In: *Fire Technology* 53.2 (2017), pp. 695–708.

[2] Henrik Jaldell, Prachaksvich Lebnak, and Anurak Amornpetchsathaporn. “Time is money, but how much? The monetary value of response time for Thai ambulance emergency services”. In: *Value in Health* 17.5 (2014), pp. 555–560.

[3] Shana Hertz Hattis. *Crime in the United States*. Lanham Bernan Press, 2015.

[4] Ayan Mukhopadhyay. “Robust Incident Prediction, Resource Allocation and Dynamic Dispatch”. PhD thesis. Vanderbilt University, 2019, p. 141.

[5] USA Deparment of Homeland Security. *Phases of Emergency Management*. https://www.hsdl.org/?view&did=488295. [Online; accessed 14-Mar-2020]. 2019.

[6] Center of Disease Control and Prevention. *Road Traffic Injuries and Deaths — A Global Problem*. https://www.cdc.gov/injury/features/global-road-safety/index.html. 2019.

[7] Association for Safe International Road Travel. *Road Safety Facts*. https://www.asirt.org/safe-travel/road-safety-facts/. [Online; accessed 14-Mar-2020]. 2019.

[8] National Emergency Number Association. *9-1-1 Statistics*. https://www.nena.org/page/911Statistics. 2019.

[9] U.S. Department of Justice. *2014 Crime in the United States*. https://ucr.fbi.gov/crime-in-the-u.s/2014/crime-in-the-u.s.-2014. [Online; accessed 14-Mar-2020]. 2014.

[10] Stefan Harrendorf, Markku Heiskanen, and Steven Malby. “International Statistics on Crime and Justice”. In: *Publication Series (Helsinki Institute for Crime Prevention and Control)* (2010).

[11] B Nambuusi, Tom Brijs, and Elke Hermans. *A Review of Accident Prediction Models for Road Intersections*. Tech. rep. Policy Research Centre for Traffic Safety, 2008.

[12] George Yannis et al. “Road Traffic Accident Prediction Modelling: A Literature Review”. In: *Proceedings of the Institution of Civil Engineers-Transport*. Vol. 170. 5. Thomas Telford Ltd. 2017, pp. 245–254.

[13] Vasin Kiattikomol. “Freeway crash prediction models for long-range urban transportation planning”. PhD thesis. University of Tennessee, Knoxville, 2005, p. 144.

[14] Xueping Li et al. “Covering models and optimization techniques for emergency response facility location and planning: a review”. In: *Mathematical Methods of Operations Research* 74.3 (2011), pp. 281–310.

[15] K Bohm and Lisa Kurland. “The Accuracy of Medical Dispatch - A Systematic Review”. In: *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine* 26.1 (2018), p. 94.

[16] Dominique Lord and Fred Mannering. “The Statistical Analysis of Crash-Frequency Data: A Review and Assessment of Methodological Alternatives”. In: *Transportation Research Part A: Policy and Practice* 44.5 (2010), pp. 291–305.

[17] Ubon Thongsatapornwatana. “A Survey of Data Mining Techniques for Analyzing Crime Patterns”. In: *2016 Second Asian Conference on Defence Technology (ACDT)*. IEEE. 2016, pp. 123–128.

[18] Vikas Grover, Richard Adderley, and Max Bramer. “Review of Current Crime Prediction Techniques”. In: *International Conference on Innovative Techniques and Applications of Artificial Intelligence*. Springer. 2006, pp. 233–237.

[19] Peter Lynn and Dave Elliot. *The British Crime Survey: A Review of Methodology*. National Centre for Social Research London, 2000.

[20] Richard Berk. “Forecasting Methods in Crime and Justice”. In: *Annual Review of Law and Social Science* 4 (2008), pp. 219–238.

[21] Hossein Hassani et al. “A Review of Data Mining Applications in Crime”. In: *Statistical Analysis and Data Mining: The ASA Data Science Journal* 9.3 (2016), pp. 139–154.

[22] Ciy of Rochester. *How 911 Works - What Happens When You Dial 911*. https://www.cityofrochester.gov/article.aspx?id=8589935579. 2019.

[23] Private Communication. Nashville Fire Department. USA, 2018.

[24] Mykel J Kochenderfer. *Decision Making Under Uncertainty: Theory and Application*. MIT press, 2015.

[25] John A Deacon, Charles V Zegeer, and Robert C Deen. “Identification of Hazardous Rural Highway Locations”. In: *Transportation Research Record* 543 (1974).

[26] Wen Cheng and Simon P Washington. “Experimental evaluation of hotspot identification methods”. In: *Accident Analysis & Prevention* 37.5 (2005), pp. 870–881.

[27] John Eck et al. *Mapping Crime: Understanding Hot Spots*. Tech. rep. National Institute of Justice, 2005.

[28] Hao Yu et al. “Comparative analysis of the Spatial Analysis Methods for Hotspot Identification”. In: *Accident Analysis & Prevention* 66 (2014), pp. 80–88.

[29] Benjamin Ryder et al. “Spatial Prediction of Traffic Accidents with Critical Driving Events – Insights from a Nationwide Field Study”. In: *Transportation Research Part A: Policy and Practice* 124 (2019), pp. 611–626.

[30] J Frantzeskakis, V Assimakopoulos, and G Kindinis. “Interurban Accident Prediction by Administrative Area Application in Greece”. In: *Institution of Transportation Engineers journal* 64.1 (1994), pp. 35–42.

[31] Paul P Jovanis and Hsin-Li Chang. “Modeling the Relationship of Accidents to Miles Traveled”. In: *Transportation Research Record* 1068 (1986), pp. 42–51.

[32] Shaw-Pin Miaou and Harry Lum. “Modeling vehicle accidents and highway geometric design relationships”. In: *Accident Analysis & Prevention* 25.6 (1993), pp. 689–709.

[33] Sarath C Joshua and Nicholas J Garber. “Estimating Truck Accident Rate and Involvements using Linear and Poisson Regression Models”. In: *Transportation Planning and Technology* 15.1 (1990), pp. 41–58.

[34] Hesham Rakha et al. “Linear Regression Crash Prediction Models: Issues and Proposed Solutions”. In: *Efficient Transportation and Pavement Systems: Characterization, Mechanisms, Simulation and Modeling* (2010), pp. 241–256.

[35] Dominique Lord, Simon P Washington, and John N Ivan. “Poisson, Poisson-Gamma and Zero-Inflated Regression Models of Motor Vehicle Crashes: Balancing Statistical Fit and Theory”. In: *Accident Analysis & Prevention* 37.1 (2005), pp. 35–46.

[36] James A Bonneson and Patrick T McCoy. “Estimation of Safety at Two-Way Stop-Controlled Intersections on Rural Highways”. In: *Transportation Research Record* 1401 (1993), pp. 83–89.

[37] Michael J Maher and Ian Summersgill. “A Comprehensive Methodology for the Fitting of Predictive Accident Models”. In: *Accident Analysis & Prevention* 28.3 (1996), pp. 281–296.

[38] Tarek Sayed and Felipe Rodriguez. “Accident Prediction Models for Urban Unsignalized Intersections in British Columbia”. In: *Transportation Research Record: Journal of the Transportation Research Board* 1665 (1999), pp. 93–99.

[39] Zhirui Ye, Yueru Xu, and Dominique Lord. “Crash Data Modeling with a Generalized Estimator”. In: *Accident Analysis & Prevention* 117 (2018), pp. 340–345.

[40] Markus Deublein et al. “Prediction of road accidents: A Bayesian hierarchical approach”. In: *Accident Analysis & Prevention* 51 (2013), pp. 274–291.

[41] Mohammed A. Quddus. “Modelling Area-Wide Count Outcomes with Spatial Correlation and Heterogeneity: An Analysis of London Crash Data”. In: *Accident Analysis and Prevention* (2008).

[42] Darcin Akin. “Analysis of Highway Crash Data by Negative Binomial and Poisson Regression Models”. In: 2nd International Symposium on Computing in Science and Engineering (ISCSE). 2011.

[43] Felipe Ladron de Guevara, Simon P Washington, and Jutaek Oh. “Forecasting Crashes at the Planning Level: Simultaneous Negative Binomial Crash Model Applied in Tucson, Arizona”. In: *Transportation Research Record* 1897.1 (2004), pp. 191–199.

[44] Ciro Caliendo, Maurizio Guida, and Alessandra Parisi. “A Crash-Prediction Model for Multilane Roads”. In:

Accident Analysis & Prevention 39.4 (2007), pp. 657–670.

[45] Williams Ackaah and Mohammed Salifu. "Crash Prediction Model for Two-Lane Rural Highways in the Ashanti Region of Ghana". In: *IATSS research* 35.1 (2011), pp. 34–40.

[46] Sunanda Dissanayake and I Ratnayake. "Statistical Modelling of Crash Frequency on Rural Freeways and Two-Lane Highways using Negative Binomial Distribution". In: *Advances in Transportation Studies* 9 (2006).

[47] Dominique Lord and Luis F Miranda-Moreno. "Effects of Low Sample Mean Values and Small Sample Size on the Estimation of the Fixed Dispersion Parameter of Poisson-Gamma Models for Modeling Motor Vehicle Crashes: A Bayesian Perspective". In: *Safety Science* 46.5 (2008), pp. 751–770.

[48] Jonathan Aguero-Valverde and Paul P Jovanis. "Analysis of road crash frequency with spatial models". In: *Transportation Research Record* 2061.1 (2008), pp. 55–63.

[49] Mohammadali Shirazi and Dominique Lord. "Characteristics-based heuristics to select a logical distribution between the Poisson-gamma and the Poisson-lognormal for crash data modelling". In: *Transportmetrica A: Transport Science* 15.2 (Nov. 2019), pp. 1791–1803.

[50] Eun Sug Park and Dominique Lord. "Multivariate Poisson-Lognormal Models for Jointly Modeling Crash Frequency by Severity". In: *Transportation Research Record: Journal of the Transportation Research Board* 2019.1 (Jan. 2007), pp. 1–6.

[51] Jianming Ma, Kara M. Kockelman, and Paul Damien. "A multivariate Poisson-lognormal regression model for prediction of crash counts by severity, using Bayesian methods". In: *Accident Analysis & Prevention* 40.3 (May 2008), pp. 964–975.

[52] Jonathan Aguero-Valverde. "Full Bayes Poisson gamma, Poisson lognormal, and zero inflated random effects models: Comparing the precision of crash frequency estimates". In: *Accident Analysis and Prevention* 50 (2013), pp. 289–297.

[53] S. Hadi Khazraee, Valen Johnson, and Dominique Lord. "Bayesian Poisson hierarchical models for crash data analysis: Investigating the impact of model choice on site-specific predictions". In: *Accident Analysis & Prevention* 117 (Aug. 2018), pp. 181–195.

[54] Xiao Qin, John N Ivan, and Nalini Ravishanker. "Selecting exposure measures in crash rate prediction for two-lane highway segments". In: *Accident Analysis & Prevention* 36.2 (2004), pp. 183–191.

[55] Dominique Lord, Simon Washington, and John N Ivan. "Further notes on the application of zero-inflated models in highway safety". In: *Accident Analysis & Prevention* 39.1 (2007), pp. 53–57.

[56] Helai Huang and Hong Chor Chin. "Modeling road traffic crashes with zero-inflation and site-specific random effects". In: *Statistical Methods & Applications* 19.3 (2010), pp. 445–462.

[57] Ayan Mukhopadhyay et al. "Prioritized Allocation of Emergency Responders based on a Continuous-Time Incident Prediction Model". In: *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. International Foundation for Autonomous Agents and Multiagent Systems. 2017, pp. 168–177.

[58] Geoffrey Pettet et al. "Incident analysis and prediction using clustering and Bayesian network". In: *2017 IEEE SmartWorld*. IEEE. IEEE, 2017, pp. 1–8.

[59] Ayan Mukhopadhyay et al. "An Online Decision-theoretic Pipeline for Responder Dispatch". In: *Proceedings of the 10th ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*. Montreal, Quebec, Canada: ACM, 2019, pp. 185–196.

[60] Ayan Mukhopadhyay, Zilin Wang, and Yevgeniy Vorobeychik. "A Decision Theoretic Framework for Emergency Responder Dispatch". In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, (AAMAS)*. 2018, pp. 588–596.

[61] Rui Jiang, Ming Qu, Edward Chung, et al. "Traffic incident clearance time and arrival time prediction based on hazard models". In: *Mathematical Problems in Engineering* (2014).

[62] Hasti Tajtehranifard et al. "Motorway crash duration and its determinants: do durations vary across motorways?" In: *Journal of Advanced Transportation* 50.5 (2016), pp. 717–735.

[63] Ruimin Li. "Traffic incident duration analysis and prediction models based on the survival analysis approach". In: *IET Intelligent Transport Systems* 9.4 (2014), pp. 351–358.

[64] Younshik Chung and Wilfred W Recker. "A methodological approach for estimating temporal and spatial extent of delays caused by freeway accidents". In: *IEEE Transactions on Intelligent Transportation Systems* 13.3 (2012), pp. 1454–1461.

[65] Sanchita Basak et al. "Data-Driven Detection of Anomalies and Cascading Failures in Traffic Networks". In: *Proceedings of the Annual Conference of the PHM Society*. Vol. 11. 1. 2019.

[66] Walter R Gilks, Sylvia Richardson, and David J Spiegelhalter. "Introducing markov chain monte carlo". In: *Markov Chain Monte Carlo in Practice* 1 (1996), p. 19.

[67] Harvey Goldstein. "Multilevel Statistical Models". In: *Institute of Education, Multilevel Models Project* (1995).

[68] Shaw-Pin Miaou and Dominique Lord. "Modeling traffic crash-flow relationships for intersections: dispersion parameter, functional form, and Bayes versus empirical Bayes methods". In: *Transportation Research Record* 1840.1 (2003), pp. 31–40.

[69] Ezra Hauer. "On the estimation of the expected number of accidents". In: *Accident Analysis & Prevention* 18.1 (1986), pp. 1–12.

[70] Ezra Hauer. "Empirical Bayes approach to the estimation of "unsafety": the multivariate regression method". In: *Accident Analysis & Prevention* 24.5 (1992), pp. 457–477.

[71] Ezra Hauer and Bhagwant Persaud. "Common bias in before-and-after accident comparisons and its elimination". In: *Transportation Research Record* 905 (1983), pp. 164–174.

[72] B G Heydecker and J Wu. "Identification of sites for road accident remedial work by Bayesian statistical methods: an example of uncertain inference". In: *Advances in Engineering Software* 32.10 (2001), pp. 859–869.

[73] Ezra Hauer et al. "Estimating safety by the empirical Bayes method: a tutorial". In: *Transportation Research Record* 1784.1 (2002), pp. 126–131.

[74] Ying C MacNab. "Bayesian spatial and ecological models for small-area accident and injury analysis". In: *Accident Analysis & Prevention* 36.6 (2004), pp. 1019–1028.

[75] Jonathan Aguero-Valverde and Paul P Jovanis. "Bayesian multivariate Poisson lognormal models for crash severity modeling and site ranking". In: *Transportation Research Record* 2136.1 (2009), pp. 82–91.

[76] Shaw-Pin Miaou and Joon Jin Song. "Bayesian ranking of sites for engineering safety improvements: decision parameter, treatability concept, statistical criterion, and spatial dependence". In: *Accident Analysis & Prevention* 37.4 (2005), pp. 699–720.

[77] P J Schlüter, J J Deely, and A J Nicholson. "Ranking and selecting motor vehicle accident sites by using a hierarchical Bayesian model". In: *Journal of the Royal Statistical Society: Series D (The Statistician)* 46.3 (1997), pp. 293–316.

[78] Gary A Davis and Shimin Yang. "Bayesian identification of high-risk intersections for older drivers via Gibbs sampling". In: *Transportation Research Record* 1746.1 (2001), pp. 84–89.

[79] Sujata Basu and Pritam Saha. "Regression models of highway traffic crashes: a review of recent research and future research needs". In: *Procedia Engineering* 187 (2017), pp. 59–66.

[80] Joon Jin Song et al. "Bayesian multivariate spatial models for roadway traffic crash mapping". In: *Journal of Multivariate Analysis* 97.1 (2006), pp. 246–273.

[81] Byung-Jung Park, Dominique Lord, and Jeffrey D Hart. "Bias properties of Bayesian statistics in finite mixture of negative binomial regression models in crash data analysis". In: *Accident Analysis & Prevention* 42.2 (2010), pp. 741–749.

[82] Mohamed Abdel-Aty et al. "Assessing Safety on Dutch Freeways with Data from Infrastructure-Based Intelligent Transportation Systems". In: *Transportation Research Record* 2083.1 (Jan. 2008), pp. 153–161.

[83] Rongjie Yu and Mohamed Abdel-Aty. "Analyzing crash injury severity for a mountainous freeway incorporating real-time traffic and weather data". In: *Safety Science* 63 (2014), pp. 50–56.

[84] Yunlong Zhang and Yuanchang Xie. "Forecasting of Short-Term Freeway Volume with v-Support Vector Machines". In: *Transportation Research Record* 2024.1 (2007), pp. 92–99.

[85] Xiugang Li et al. "Predicting motor vehicle crashes using Support Vector Machine models". In: *Accident Analysis and Prevention* 40.4 (2008), pp. 1611–1618.

[86] Rongjie Yu and Mohamed Abdel-Aty. "Utilizing support vector machine in real-time crash risk evaluation". In: *Accident Analysis & Prevention* 51 (Mar. 2013), pp. 252–259.

[87] Anurag Pande and Mohamed Abdel-Aty. "Assessment of freeway traffic parameters leading to lane-change related collisions". In: *Accident Analysis & Prevention* 38.5 (Sept. 2006), pp. 936–948.

[88] Hassan T Abdelwahab and Mohamed A Abdel-Aty. "Artificial Neural Networks and Logit Models for Traffic Safety Analysis of Toll Plazas". In: *Transportation Research Record* 1784.1 (Jan. 2002), pp. 115–125.

[89] Li-Yen Chang. "Analysis of freeway accident frequencies: Negative binomial regression versus artificial neural network". In: *Safety Science* 43.8 (2005), pp. 541–557.

[90] C Riviere et al. "A Bayesian Neural Network approach to estimating the Energy Equivalent Speed". In: *Accident Analysis and Prevention* 38.2 (2006), pp. 248–259.

[91] Yuanchang Xie, Dominique Lord, and Yunlong Zhang. "Predicting motor vehicle collisions using Bayesian neural network models: An empirical analysis". In: *Accident Analysis & Prevention* 39.5 (Sept. 2007), pp. 922–933.

[92] Lin Zhu et al. *The use of convolutional neural networks for traffic incident detection at a network level*. Tech. rep. 2018.

[93] Jie Bao, Pan Liu, and Satish V. Ukkusuri. "A spatiotemporal deep learning approach for citywide short-term crash risk prediction with multi-source data". In: *Accident Analysis & Prevention* 122 (Jan. 2019), pp. 239–254.

[94] Sanchita Basak, Abhishek Dubey, and Leao Bruno. "Analyzing the Cascading Effect of Traffic Congestion Using LSTM Networks". In: *2019 IEEE International Conference on Big Data (Big Data)*. IEEE. 2019, pp. 2144–2153.

[95] Julia Angwin et al. *Machine Bias*. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing. 2016.

[96] Oskar Josef Gstrein, Anno Bunnik, and Andrej Zwitter. "Ethical, Legal and Social Challenges of Predictive Policing". In: *Católica Law Review (2019)* (2019).

[97] Peter M Asaro. "AI ethics in predictive policing: From models of threat to an ethics of care". In: *IEEE Technology and Society Magazine* 38.2 (2019), pp. 40–53.

[98] Tero Karppi. ""The Computer Said So": On the Ethics, Effectiveness, and Cultural Techniques of Predictive Policing". In: *Social Media+ Society* 4.2 (2018), p. 2056305118768296.

[99] Ayan Mukhopadhyay et al. "Robust Spatial-Temporal Incident Prediction". In: *36th Conference on Uncertainty in Artificial Intelligence, Toronto (UAI 2020)*. CONF. AUAI. 2019 (to appear).

[100] Ayan Mukhopadhyay et al. "Optimal Allocation of Police Patrol Resources Using a Continuous-Time Crime Model". In: *Conference on Decision and Game Theory for Security*. 2016.

[101] Hongjian Wang et al. "Crime rate inference with big data". In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 2016, pp. 635–644.

[102] Ned Levine et al. "CrimeStat III: a spatial statistics program for the analysis of crime incident locations (version 3.0)". In: *Houston (TX): Ned Levine & Associates/Washington, DC: National Institute of Justice* (2004).

[103] Martin Kulldorff. "A spatial scan statistic". In: *Communications in Statistics-Theory and methods* 26.6 (1997), pp. 1481–1496.

[104] Martin Kulldorff and Neville Nagarwalla. "Spatial disease clusters: detection and inference". In: *Statistics in medicine* 14.8 (1995), pp. 799–810.

[105] Shyam Varan Nath. "Crime pattern detection using data mining". In: *2006 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology Workshops*. IEEE. 2006, pp. 41–44.

[106] Jyoti Agarwal, Renuka Nagpal, and Rajni Sehgal. "Crime analysis using k-means clustering". In: *International Journal of Computer Applications* 83.4 (2013).

[107] Fateha Khanam Bappee, Amilcar Soares Junior, and Stan Matwin. "Predicting crime using spatial features". In: *Canadian Conference on Artificial Intelligence*. Springer. 2018, pp. 367–373.

[108] Bashar Aubaidan, Masnizah Mohd, and Mohammed Albared. "Comparative study of k-means and k-means++ clustering algorithms on crime domain". In: *Journal of Computer Science* 10.7 (2014), pp. 1197–1206.

[109] Charu C Aggarwal. *Data classification: algorithms and applications*. CRC press, 2014.

[110] SP Chainey. "Examining the influence of cell size and bandwidth size on kernel density estimation crime hotspot maps for predicting spatial patterns of crime". In: *Bulletin of the Geographical Society of Liege* 60 (2013), pp. 7–19.

[111] Timothy Hart and Paul Zandbergen. "Kernel density estimation and hotspot mapping". In: *Policing: An International Journal of Police Strategies & Management* (2014).

[112] S Sivaranjani, S Sivakumari, and M Aasha. "Crime prediction and forecasting in Tamilnadu using clustering approaches". In: *2016 International Conference on Emerging Technological Trends (ICETT)*. IEEE. 2016, pp. 1–6.

[113] Yongwan Chun. "Analyzing space–time crime incidents using eigenvector spatial filtering: an application to vehicle burglary". In: *Geographical Analysis* 46.2 (2014), pp. 165–184.

[114] D Wayne Osgood. "Poisson-based regression analysis of aggregate crime rates". In: *Journal of quantitative criminology* 16.1 (2000), pp. 21–43.

[115] Eric L Piza. "Using Poisson and negative binomial regression models to measure the influence of risk on crime incident counts". In: *Rutgers Center on Public Security* (2012).

[116] Leslie W Kennedy, Joel M Caplan, and Eric Piza. "Risk clusters, hotspots, and spatial intelligence: risk terrain modeling as an algorithm for police resource allocation strategies". In: *Journal of Quantitative Criminology* 27.3 (2011), pp. 339–362.

[117] Joel M Caplan and Leslie W Kennedy. "Risk terrain modeling compendium". In: *Rutgers Center on Public Security, Newark* (2011).

[118] Shane D Johnson, Kate Bowers, and Alex Hirschfield. "New insights into the spatial and temporal distribution of repeat victimization". In: *The British Journal of Criminology* 37.2 (1997), pp. 224–241.

[119] Graham Farrell and Kenneth Pease. *Repeat victimization*. Vol. 12. Criminal Justice Press, 2001.

[120] Allison T Chappell, Elizabeth Monk-Turner, and Brian K Payne. "Broken windows or window breakers: The influence of physical and social disorder on quality of life". In: *Justice Quarterly* 28.3 (2011), pp. 522–540.

[121] M. B. Short et al. "A statistical model of criminal behavior". In: *Mathematical Models and Methods in Applied Sciences* (2008), pp. 1249–1267.

[122] G. O. Mohler et al. "Self-exciting point process modeling of crime". In: *Journal of the American Statistical Association* 106.493 (2011), pp. 100–108.

[123] Jacqueline Cohen, Wilpen L Gorr, and Andreas M Olligschlaeger. "Leading indicators and spatial interactions: A crime-forecasting model for proactive police deployment". In: *Geographical Analysis* 39.1 (2007), pp. 105–127.

[124] Christian Ivaha et al. "The Dynamic Spatial Disaggregation Approach: A Spatio-Temporal Modelling of Crime." In: *World Congress on Engineering*. 2007, pp. 961–966.

[125] Bernhard Von Stengel and Shmuel Zamir. *Leadership with commitment to mixed strategies*. Tech. rep. Citeseer, 2004.

[126] Jason Tsai et al. "IRIS-a tool for strategic security allocation in transportation networks". In: (2009).

[127] Fei Fang et al. "Deploying PAWS: Field optimization of the protection assistant for wildlife security". In: *Twenty-Eighth IAAI Conference*. 2016.

[128] Matthew Brown et al. "One size does not fit all: A game-theoretic approach for dynamically and effectively screening for threats". In: *Thirtieth AAAI Conference on Artificial Intelligence*. 2016.

[129] William Haskell et al. "Robust protection of fisheries with compass". In: *Twenty-Sixth IAAI Conference*. 2014.

[130] Matthew Paul Johnson, Fei Fang, and Milind Tambe. "Patrol strategies to maximize pristine forest area". In: *Twenty-Sixth AAAI Conference on Artificial Intelligence*. 2012.

[131] Yi (Grace) Qi, Brian L. Smith, and Jianhua Guo. "Freeway Accident Likelihood Prediction Using a Panel Data Analysis Approach". In: *Journal of Transportation Engineering* 133.3 (Mar. 2007), pp. 149–156.

[132] Praprut Songchitruksa and Kevin Balke. "Assessing weather, environment, and loop data for real-time freeway incident prediction". In: *Transportation Research Record: Journal of the Transportation Research Board* 1959 (2006), pp. 105–113.

[133] Mohamed A. Abdel-Aty et al. "Real-time prediction of visibility related crashes". In: *Transportation Research Part C: Emerging Technologies* 24 (Oct. 2012), pp. 288–298.

[134] Helai Huang, Hoong Chor Chin, and Md. Mazharul Haque. "Severity of driver injury and vehicle damage in traffic crashes at intersections: A Bayesian hierarchical analysis". In: *Accident Analysis & Prevention* 40.1 (Jan. 2008), pp. 45–54.

[135] Mark Poch and Fred Mannering. "Negative binomial analysis of intersection-accident frequencies". In: *Journal of Transportation Engineering* 122.2 (1996), pp. 105–113.

[136] Venkataraman Shankar, Fred Mannering, and Woodrow Barfield. "Effect of roadway geometrics and environmental factors on rural freeway accident frequencies". In: *Accident Analysis & Prevention* 27.3 (1995), pp. 371–389.

[137] Hoong Chor Chin and Mohammed Abdul Quddus. "Applying the random effect negative binomial model to examine traffic accident occurrence at signalized intersections". In: *Accident Analysis & Prevention* 35.2 (2003), pp. 253–259.

[138] Qiang Zeng, Jiaren Sun, and Huiying Wen. "Bayesian Hierarchical Modeling Monthly Crash Counts on Freeway Segments with Temporal Correlation". In: *Journal of Advanced Transportation* 2017 (2017). Ed. by Francesco Bella, p. 5391054.

[139] Pei-Fen Kuo and Dominique Lord. "Applying the colocation quotient index to crash severity analyses". In: *Accident Analysis & Prevention* 135 (Feb. 2020), p. 105368.

[140] Qi Shi and Mohamed Abdel-Aty. "Big Data applications in real-time traffic operation and safety monitoring and improvement on urban expressways". In: *Transportation Research Part C: Emerging Technologies* 58 (Sept. 2015), pp. 380–394.

[141] Chao Wang, Mohammed A. Quddus, and Stephen G. Ison. "Impact of traffic congestion on road accidents: A spatial analysis of the M25 motorway in England". In: *Accident Analysis & Prevention* 41.4 (July 2009), pp. 798–808.

[142] A Baruya. "Speed-accident relationships on European roads". In: *9th International Conference on Road Safety in Europe*. 1998, pp. 1–19.

[143] Edwin Grant Dexter. *Weather influences*. Macmillan, 1904.

[144] Cesare Lombroso. *Crime, its causes and remedies*. 3. Little, Brown, 1918.

[145] Gerhard J Falk. "The influence of the seasons on the crime rate". In: *J. Crim. L. Criminology & Police Sci.* 43 (1952), p. 199.

[146] Hyeon-Woo Kang and Hang-Bong Kang. "Prediction of crime occurrence from multi-modal data using deep learning". In: *PloS one* 12.4 (2017).

[147] Alexander Stec and Diego Klabjan. "Forecasting Crime with Deep Learning". In: *arXiv preprint arXiv:1806.01486* (2018).

[148] Janet L Lauritsen and Nicole White. *Seasonal patterns in criminal victimization trends*. US Department of Justice, Office of Justice Programs, Bureau of Justice . . ., 2014.

[149] Sherry Towers et al. "Factors influencing temporal patterns in crime in a large American city: A predictive analytics perspective". In: *PLoS one* 13.10 (2018).

[150] Marcus Felson and Erika Poulsen. "Simple indicators of crime by time of day". In: *International Journal of Forecasting* 19.4 (2003), pp. 595–601.

[151] Wim Bernasco, Stijn Ruiter, and Richard Block. "Do street robbery location choices vary over time of day or day of week? A test in Chicago". In: *Journal of Research in Crime and Delinquency* 54.2 (2017), pp. 244–275.

[152] Richard A Scribner, David P MacKinnon, and James H Dwyer. "The risk of assaultive violence and alcohol availability in Los Angeles County." In: *American journal of public health* 85.3 (1995), pp. 335–340.

[153] Paul W Speer et al. "Violent crime and alcohol availability: relationships in an urban community". In: *Journal of public health policy* 19.3 (1998), pp. 303–318.

[154] Dennis M Gorman et al. "Spatial dynamics of alcohol availability, neighborhood structure and violent crime." In: *Journal of studies on alcohol* 62.5 (2001), pp. 628–636.

[155] Rocco d'Este. "The effects of stolen goods markets on crime: pawnshops, property theft, and the gold rush of the 2000s". In: *The Journal of Law & Economics* (2019).

[156] Simon M Fass and Janice Francis. "Where have all the hot goods gone? The role of pawnshops". In: *Journal of Research in Crime and Delinquency* 41.2 (2004), pp. 156–179.

[157] Rocco D'Este. *The Effect of Stolen Goods Markets on Crime: Evidence from a Quasi Natural Experiment*. Tech. rep. 2014.

[158] Sara-Laure Faraji, Greg Ridgeway, and Yuhao Wu. "Effect of emergency winter homeless shelters on property crime". In: *Journal of Experimental Criminology* 14.2 (2018), pp. 129–140.

[159] Sean N Fischer et al. "Homelessness, mental illness, and criminal activity: Examining patterns over time". In: *American Journal of Community Psychology* 42.3-4 (2008), p. 251.

[160] Paolo Buonanno. "The socioeconomic determinants of crime. A review of the literature". In: (2003).

[161] Oana-Ramona Lobonţ et al. "The effect of socioeconomic factors on crime rates in Romania: a macro-level analysis". In: *Economic research-Ekonomska istraživanja* 30.1 (2017), pp. 91–111.

[162] Rajindar K Koshal and Manjulika Koshal. "Crimes and socio-economic environments". In: *Social Indicators Research* 2.2 (1975), pp. 223–227.

[163] Joel M Caplan et al. "Risk terrain modeling for spatial risk assessment". In: *Cityscape* 17.1 (2015), pp. 7–16.

[164] Peter T. Savolainen et al. "The statistical analysis of highway crash-injury severities: A review and assessment of methodological alternatives". In: *Accident Analysis & Prevention* 43.5 (Sept. 2011), pp. 1666–1676.

[165] KS Jadaan and AJ Nicholson. "Relationships between road accidents and traffic flows in an urban network". In: *Traffic engineering and control* 33.9 (1992), pp. 507–511.

[166] Jean-Louis Martin. "Relationship between crash rate and hourly traffic flow on interurban motorways". In: *Accident Analysis & Prevention* 34.5 (2002), pp. 619–629.

[167] DJ Turner and R Thomas. "Motorway accidents: an examination of accident totals, rates and severity and their relationship with traffic flow". In: *Traffic Engineering and Control* 27.7-8 (1986), pp. 377–383.

[168] David Ronald Dickson McGuigan. "An examination of relationships between road accidents and traffic flow". PhD thesis. Newcastle University, 1987.

[169] JW Hall and OJ Pendleton. "Rural accident rate variations with traffic volume". In: *Transportation Research Record* 1281 (1990).

[170] Fan Ye and Dominique Lord. "Comparing three commonly used crash severity models on sample size requirements: Multinomial logit, ordered probit and mixed logit models". In: *Analytic Methods in Accident Research* 1 (Jan. 2014), pp. 72–85.

[171] Asad J Khattak, Paula Kantor, and Forrest M Council. "Role of adverse weather in key crash types on limited-access: roadways implications for advanced weather systems". In: *Transportation research record* 1621.1 (1998), pp. 10–19.

[172] Akira Shibata and Katsuhiro Fukuda. "Risk factors of fatality in motor vehicle traffic accidents". In: *Accident Analysis & Prevention* 26.3 (1994), pp. 391–397.

[173] Wei David Fan et al. *Identifying and quantifying factors affecting vehicle crash severity at highway-rail grade crossings: models and their comparison*. Tech. rep. 2016.

[174] Li-Yen Chang and Hsiu-Wen Wang. "Analysis of traffic injury severity: An application of non-parametric classification tree techniques". In: *Accident Analysis & Prevention* 38.5 (2006), pp. 1019–1027.

[175] Zijian Zheng, Pan Lu, and Brenda Lantz. "Commercial truck crash injury severity analysis using gradient boosting data mining model". In: *Journal of Safety Research* 65 (June 2018), pp. 115–124.

[176] Amirfarrokh Iranitalab and Aemal Khattak. "Comparison of four statistical and machine learning methods for crash severity prediction". In: *Accident Analysis & Prevention* 108 (2017), pp. 27–36.

[177] Somayeh Mafi, Yassir Abdelrazig, and Ryan Doczy. "Machine learning methods to analyze injury severity of drivers from different age and gender groups". In: *Transportation research record* 2672.38 (2018), pp. 171–183.

[178] F Rezaie Moghaddam, Sh Afandizadeh, and M Ziyadi. "Prediction of accident severity using artificial neural networks". In: *International Journal of Civil Engineering* 9.1 (2011), pp. 41–48.

[179] Deo Chimba and Thobias Sando. "The prediction of highway traffic accident injury severity with neuromorphic techniques". In: *Advances in transportation studies* 19 (2009), pp. 17–26.

[180] FD Bijleveld. "The covariance between the number of accidents and the number of victims in multivariate analysis of accident related outcomes". In: *Accident Analysis & Prevention* 37.4 (2005), pp. 591–600.

[181] Jianming Ma and Kara M Kockelman. "Bayesian multivariate Poisson regression for models of injury count, by severity". In: *Transportation Research Record* 1950.1 (2006), pp. 24–34.

[182] Qiang Zeng et al. "Modeling nonlinear relationship between crash frequency by severity and contributing factors by neural networks". In: *Analytic methods in accident research* 10 (2016), pp. 12–25.

[183] Xin Pei, SC Wong, and Nang-Ngai Sze. "A joint-probability approach to crash prediction models". In: *Accident Analysis & Prevention* 43.3 (2011), pp. 1160–1166.

[184] Ming Sun, Xiaoduan Sun, and Donghui Shan. "Pedestrian crash analysis with latent class clustering method". In: *Accident Analysis & Prevention* 124 (2019), pp. 50–57.

[185] Lekshmi Sasidharan, Kun-Feng Wu, and Monica Menendez. "Exploring the application of latent class cluster analysis for investigating pedestrian crash injury severities in Switzerland". In: *Accident Analysis & Prevention* 85 (2015), pp. 219–228.

[186] Sathish Kumar Sivasankaran and Venkatesh Balasubramanian. "Exploring the severity of bicycle–Vehicle crashes using latent class clustering approach in India". In: *Journal of Safety Research* 72 (2020), pp. 127–138.

[187] Geoffrey Pettet et al. "On algorithmic decision procedures in emergency response systems in smart and connected communities". In: *arXiv preprint arXiv:2001.07362* (2020).

[188] Constantine Toregas et al. "The Location of Emergency Service Facilities". In: *Operations Research*

19.6 (Oct. 1, 1971). Publisher: INFORMS, pp. 1363–1373.

[189] Richard Church and Charles ReVelle. "The maximal covering location problem". In: *Papers of the regional science association*. Vol. 32. 1. Springer-Verlag. 1974, pp. 101–118.

[190] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. "Solving an ambulance location model by tabu search". In: *Location Science* 5.2 (Aug. 1, 1997), pp. 75–88.

[191] Erhan Erkut, Armann Ingolfsson, and Güneş Erdoğan. "Ambulance location for maximum survival". In: *Naval Research Logistics (NRL)* 55.1 (2008). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/nav.20267, pp. 42–58.

[192] V. A. Knight, P. R. Harper, and L. Smith. "Ambulance allocation for maximal survival with heterogeneous outcome measures". In: *Omega*. Special Issue on Forecasting in Management Science 40.6 (Dec. 1, 2012), pp. 918–926.

[193] Richard McCormack and Graham Coates. "A simulation model to enable the optimization of ambulance fleet allocation and base station location for increased patient survival". In: *European Journal of Operational Research* 247.1 (Nov. 16, 2015), pp. 294–309.

[194] C Swoveland et al. "A simulation-based methodology for optimization of ambulance service policies". In: *Socio-Economic Planning Sciences* 7.6 (1973), pp. 697–703.

[195] Vladimir Marianov and Charles Revelle. "The queuing probabilistic location set covering problem and some extensions". In: *Socio-Economic Planning Sciences* 28.3 (Jan. 1, 1994), pp. 167–178.

[196] Hongzhong Jia, Fernando Ordóñez, and Maged Dessouky. "A modeling framework for facility location of medical services for large-scale emergencies". In: *IIE Transactions* 39.1 (Jan. 1, 2007). Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/07408170500539113, pp. 41–55.

[197] Hari K. Rajagopalan et al. "Developing effective meta-heuristics for a probabilistic location model via experimental design". In: *European Journal of Operational Research* 177.1 (Feb. 16, 2007), pp. 83–101.

[198] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. "A dynamic model and parallel tabu search heuristic for real-time ambulance relocation". In: *Parallel Computing*. Applications of parallel computing in transportation 27.12 (Nov. 1, 2001), pp. 1641–1653.

[199] Hari K. Rajagopalan, Cem Saydam, and Jing Xiao. "A multiperiod set covering location model for dynamic redeployment of ambulances". In: *Computers & Operations Research*. Part Special Issue: New Trends in Locational Analysis 35.3 (Mar. 1, 2008), pp. 814–826.

[200] Matthew S. Maxwell, Shane G. Henderson, and Huseyin Topaloglu. "Ambulance redeployment: An approximate dynamic programming approach". In: *Proceedings of the 2009 Winter Simulation Conference (WSC)*. Proceedings of the 2009 Winter Simulation Conference (WSC). ISSN: 1558-4305. Dec. 2009, pp. 1850–1860.

[201] Francisco Silva and Daniel Serra. "Locating emergency services with different priorities: the priority queuing covering location problem". In: *Journal of the Operational Research Society* 59.9 (2008), pp. 1229–1238.

[202] Adel A. Aly and John A. White. "Probabilistic Formulation of the Emergency Service Location Problem". In: *Journal of the Operational Research Society* 29.12 (Dec. 1, 1978), pp. 1167–1179.

[203] Mark S. Daskin. "A Maximum Expected Covering Location Model: Formulation, Properties and Heuristic Solution". In: *Transportation Science* 17.1 (Feb. 1, 1983). Publisher: INFORMS, pp. 48–70.

[204] Charles Revelle and Kathleen Hogan. "The maximum reliability location problem and α-reliablep-center problem: Derivatives of the probabilistic location set covering problem". In: *Annals of Operations Research* 18.1 (Dec. 1, 1989), pp. 155–173.

[205] Charles ReVelle and Kathleen Hogan. "The Maximum Availability Location Problem". In: *Transportation Science* 23.3 (Aug. 1, 1989). Publisher: INFORMS, pp. 192–200.

[206] John F. Repede and John J. Bernardo. "Developing and validating a decision support system for locating emergency medical vehicles in Louisville, Kentucky". In: *European Journal of Operational Research* 75.3 (June 30, 1994), pp. 567–581.

[207] Rajan Batta, June M. Dolan, and Nirup N. Krishnamurthy. "The Maximal Expected Covering Location Problem: Revisited". In: *Transportation Science* 23.4 (Nov. 1, 1989). Publisher: INFORMS, pp. 277–287.

[208] Richard C. Larson. "A hypercube queuing model for facility location and redistricting in urban emergency services". In: *Computers & Operations Research* 1.1 (Mar. 1, 1974), pp. 67–95.

[209] Sean K Keneally, Matthew J Robbins, and Brian J Lunday. "A markov decision process model for the optimal dispatch of military medical evacuation assets". In: *Health care management science* 19.2 (2016), pp. 111–129.

[210] David G Olson and Gordon P Wright. "Models for allocating police preventive patrol effort". In: *Journal of the Operational Research Society* 26.4 (1975), pp. 703–715.

[211] Maite Dewinter et al. "Analysing the Police Patrol Routing Problem: A Review". In: *ISPRS International Journal of Geo-Information* 9.3 (2020), p. 157.

[212] George L Kelling et al. "The Kansas City preventive patrol experiment". In: *Washington, DC: Police Foundation* (1974).

[213] George L Kelling et al. "The Newark foot patrol experiment". In: *Washington, DC: Police Foundation* (1981), pp. 94–96.

[214] Toby Davies and Kate Bowers. “Quantifying the deterrent effect of police patrol via GPS analysis”. In: *GIS research UK (GISRUK)* (2015).

[215] Private Communication. Metropolitan Nashville Police Department. USA, 2017.

[216] Jan M Chaiken and Peter Dormont. “A patrol car allocation model: Capabilities and algorithms”. In: *Management Science* 24.12 (1978), pp. 1291–1300.

[217] Hope Corman and H Naci Mocan. “A time-series analysis of crime, deterrence, and drug abuse in New York City”. In: *American Economic Review* 90.3 (2000), pp. 584–604.

[218] Thomas B Marvell and Carlisle E Moody. “Specification problems, police levels, and crime rates”. In: *Criminology* 34.4 (1996), pp. 609–646.

[219] Kenneth Chelst. “An algorithm for deploying a crime directed (tactical) patrol force”. In: *Management Science* 24.12 (1978), pp. 1314–1327.

[220] Anthony A Braga et al. “Hot spots policing of small geographic areas effects on crime”. In: *Campbell Systematic Reviews* 15.3 (2019), e1046.

[221] David Weisburd and John E Eck. “What can police do to reduce crime, disorder, and fear?” In: *The Annals of the American Academy of Political and Social Science* 593.1 (2004), pp. 42–65.

[222] Police Executive Research Forum. “Violent crime in America: What we know about hot spots enforcement”. In: Police Executive Research Forum Washington, DC. 2008.

[223] Tae Jin Chung. *GEOGRAPHIC INFORMATION SYSTEMS: AN EFFECTIVE TECHNOLOGICAL INTERVENTION OF PROACTIVE POLICING*. LAP Lambert Academic Publishing, 2009.

[224] Tegan Herchenrader and Steven Myhill-Jones. “GIS supporting intelligence-led policing”. In: *Police Practice and Research* 16.2 (2015), pp. 136–147.

[225] Abish Malik et al. “Proactive spatiotemporal resource allocation and predictive visual analytics for community policing and law enforcement”. In: *IEEE transactions on visualization and computer graphics* 20.12 (2014), pp. 1863–1872.

[226] Arunesh Sinha et al. “Stackelberg security games: Looking beyond a decade of success”. In: IJCAI. 2018.

[227] Victor Bucarey et al. “Building real stackelberg security games for border patrols”. In: *International Conference on Decision and Game Theory for Security*. Springer. 2017, pp. 193–212.

[228] Yevgeniy Vorobeychik, Bo An, and Milind Tambe. “Adversarial patrolling games”. In: *2012 AAAI Spring Symposium Series*. 2012.

[229] Nicola Basilico and Nicola Gatti. “Automated abstractions for patrolling security games”. In: *Twenty-Fifth AAAI Conference on Artificial Intelligence*. 2011.

[230] Fei Fang, Peter Stone, and Milind Tambe. “When security games go green: Designing defender strategies to prevent poaching and illegal fishing”. In: *Twenty-Fourth International Joint Conference on Artificial Intelligence*. 2015.

[231] Thanh H Nguyen et al. “Making the most of our regrets: Regret-based solutions to handle payoff uncertainty and elicitation in green security games”. In: *International Conference on Decision and Game Theory for Security*. Springer. 2015, pp. 170–191.

[232] Haifeng Xu et al. “Optimal patrol planning for green security games with black-box attackers”. In: *International Conference on Decision and Game Theory for Security*. Springer. 2017, pp. 458–477.

[233] Chao Zhang et al. “Keeping pace with criminals: An extended study of designing patrol allocation against adaptive opportunistic criminals”. In: *Games* 7.3 (2016), p. 15.

[234] Chao Zhang et al. “Using abstractions to solve opportunistic crime security games at scale”. In: *International Conference on Autonomous Agents and Multiagent Systems*. 2016.

[235] Engineering National Academies of Sciences, Medicine, et al. *Proactive policing: Effects on crime and communities*. National Academies Press, 2018.

[236] Dennis P Rosenbaum. “The limits of hot spots policing”. In: *Police innovation: Contrasting perspectives* (2006), pp. 245–263.

[237] Juan Del Toro et al. “The criminogenic and psychological effects of police stops on adolescent black and Latino boys”. In: *Proceedings of the National Academy of Sciences* 116.17 (2019), pp. 8261–8268.

[238] David Weisburd. “Does hot spots policing inevitably lead to unfair and abusive police practices, or can we maximize both fairness and effectiveness in the new proactive policing”. In: *U. Chi. Legal F.* (2016), p. 661.

[239] Rod K Brunson and Jody Miller. “Young black men and urban policing in the United States”. In: *British journal of criminology* 46.4 (2006), pp. 613–640.

[240] Charles F Manski and Daniel S Nagin. “Assessing benefits, costs, and disparate racial impacts of confrontational proactive policing”. In: *Proceedings of the National Academy of Sciences* 114.35 (2017), pp. 9308–9313.

[241] Richard R Johnson. “Citizen expectations of police traffic stop behavior”. In: *Policing: An International Journal of Police Strategies & Management* (2004).

[242] Linda Green. “A multiple dispatch queueing model of police patrol operations”. In: *Management Science* 30.6 (1984), pp. 653–664.

[243] Jonathan Chase et al. “Law enforcement resource optimization with response time guarantees”. In: *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE. 2017, pp. 1–7.

[244] Jonathan David CHASE et al. “Improving law enforcement daily deployment through machine learning-informed optimization under uncertainty”. In: IJCAI. 2019.

[245] Todd Litman. *Autonomous vehicle implementation predictions*. Victoria Transport Policy Institute Victoria, Canada, 2017.

[246] Bentolhoda Jafary et al. "A Survey on Autonomous Vehicles Interactions with Human and other Vehicles". In: Sept. 2018.

[247] Alex Campolo et al. "AI now 2017 report". In: *AI Now Institute at New York University* (2017).

[248] Cynthia Rudin. "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead". In: *Nature Machine Intelligence* 1.5 (2019), pp. 206–215.