

Designing Decision Support Systems for Emergency Response: Challenges and Opportunities

Geoffrey Prettet*, Hunter Baxter*, Sayyed Mohsen Vazirizade*,

Hemant Purohit†, Meiyi Ma*, Ayan Mukhopadhyay*, Abhishek Dubey*

*Department of Computer Science, Vanderbilt University, Nashville, TN 37235, USA

†Information Sciences and Technology Department, George Mason University, Fairfax, VA 22030, USA

Abstract—Designing effective emergency response management (ERM) systems to respond to incidents such as road accidents is a major problem faced by communities. In addition to responding to frequent incidents each day (about 240 million emergency medical services calls and over 5 million road accidents in the US each year), these systems also support response during natural hazards. Recently, there has been a consistent interest in building decision support and optimization tools that can help emergency responders provide more efficient and effective response. This includes a number of principled subsystems that implement early incident detection, incident likelihood forecasting and strategic resource allocation and dispatch policies. In this paper, we highlight the key challenges and provide an overview of the approach developed by our team in collaboration with our community partners.

Index Terms—emergency response, decision making under uncertainty, cyber-physical systems

I. INTRODUCTION

Designing effective emergency response management (ERM) systems to oversee incidents such as road accidents and fires is a challenge faced by communities across the globe. These systems must efficiently manage resources such as ambulances to respond quickly to incidents so that human and financial losses are minimized [1], [2]. Over the last several decades, much attention has been given to studying emergency incidents and response, but emergency incidents still cause thousands of deaths and injuries and result in losses worth billions of dollars each year [3].

A fundamental shift in how these systems is operated has begun in the last decade. Critically, the use of deep learning methods and statistical methods that provide principled uncertainty-aware decision support systems are now being applied to this domain [4], [5]. These methods typically include five major components: (1) data curation components that clean and merge environmental features like traffic and weather, (2) statistical models that provide fine-grained incident likelihood forecasting, (3) resource allocation algorithms that dynamically optimize the spatial locations of responders and depots to improve response effectiveness, (4) algorithms that can provide early incident information by extracting insights from crowd-sourced data, and (5) algorithms for providing dispatch recommendations across the region, including support for coordination across multiple agencies in the community.

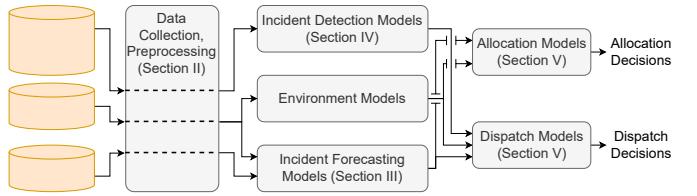


Fig. 1. ERM System Pipeline: historical data from different sources are used to design predictive models for incidents and the environment, which in turn are used to create allocation and response models. Events can be extracted using text and video data to expedite reporting and aid response.

Together, the goal of these components is to enable the agencies to (1) coordinate information and decisions between the many agencies involved with emergency response, (2) Collect, analyze and understand historical inefficiencies by creating data models using the diverse and high volume data generated from these events, (3) learn incident forecasting models that can generalize across large geographic areas, have high spatial-temporal resolution, and handle high data sparsity, (4) detect and report incidents automatically as quickly as possible, even from uncertain data-streams such as Waze and Twitter, and (5) dynamically adapt resource allocation and dispatch approaches, even if the environment in which the emergencies occur changes.

However, these goals are difficult to achieve, especially considering the following challenges. Principled emergency response systems need to account for the differing dynamics and variability of incidents. For example, the agencies have to respond to both daily incidents and large-scale disasters (e.g. natural hazards and man-made attacks) despite great variability in scales of impact for different categories of incidents. Disasters and security attacks can hinder operations and communications, which make centralized planning and cloud-based deployments infeasible. In fact, the emergency response pipelines must be robust to any communication outage. Also, to improve allocation and response, agencies require incident forecasts over fine spatial and temporal resolutions. However, learning incident prediction models at high resolutions is extremely difficult due to data sparsity [8].

Further, situational awareness during response requires information from heterogeneous data sources, e.g. weather, traffic, Twitter, Waze, 911 calls, historical emergency data, and real-time crowdsourced data (e.g. near real-time road flooding). However, integrating varying forms of data into real-

TABLE I
SUMMARY OF CHALLENGES WHEN DESIGNING AND DEPLOYING AN EMERGENCY RESPONSE MANAGEMENT FRAMEWORK.

Challenge	Description	Contribution
Coordination	ERM requires coordination between multiple agencies and decision makers, each with their own objectives. Each decision maker often has access to only incomplete information, and coordination must happen quickly while a situation is unfolding.	ERM requires coordination across jurisdictions and agencies. We have shown that response times to incidents decrease when resources dedicated to a sub-region of a city “help” other regions in response [6]. The strategy for such coordinated response is an open challenge.
Data Collection	It is difficult to collect, integrate, and preprocess the eclectic data that forms the foundation for emergency response systems. Much of the data has high volume and velocity, and is from diverse sources. This large set of data must then be narrowed down to a set of useful features.	We have compiled and released WildfireDB, the first comprehensive and open-source dataset that relates fires with relevant covariates [7]. We have also designed data collection and feature engineering pipelines for road accident data [8].
Incident Forecasting	Incident occurrence is difficult to model due to incidents’ inherent randomness and high sparsity. We have also shown that Incident models are sensitive to spatial-temporal resolution, which makes high-fidelity models challenging to learn.	We have developed an incident prediction pipeline that combines sparsity mitigation techniques with various statistical and learning based forecasting models [8]. We integrate this pipeline directly with decision making approaches to evaluate the end-to-end performance of each model.
Incident Detection	Fast incident detection is critical for timely response, but traditional reporting methods have time delays. Crowdsourced datastreams (e.g. Waze) provide an opportunity for early identification, but are noisy and uncertain.	We have developed CROME, a novel practitioner-centered incident detection approach using crowdsourced data [9].
Dynamic Environments	The environments in which ERM systems operate change over both long and short time scales. ERM systems must adapt to this nonstationarity.	We have developed scalable online planning approaches that can adapt to nonstationary environments, such as our hierarchical ambulance allocation framework [6].
Communications	Many emergency incidents cause failures in communication networks. ERM systems must be robust to communication loss to maintain service in such situations.	We have developed a partially decentralized approach to ambulance allocation that allows ambulances to operate with limited communication [10].

time incident detection models is highly non-trivial, especially when data exhibits noise and uncertainty [9], [11], [12]. Lastly, it is critical that any learning procedure designed for emergency response be adaptive. That is, it should be able to provide recommendations even if the traffic patterns, housing density and the distribution of large events in the community change.

In this paper, we describe some of these challenges and describe the approaches developed by our team to address them (see Table I)^{1,2}. Figure 1 presents our overall ERM pipeline. In particular, section II discusses the challenges involved in collecting and processing data for ERM systems. Section III presents our framework for using this sparse incident data to create useful incident forecasting models. Section IV discusses how to use crowdsourced data to detect incidents quickly. Section V discusses the challenges in designing adaptive resource allocation and dispatch techniques that are robust to communication failures. Finally, section VI discusses open research questions that require further study.

II. DATA COLLECTION AND INTEGRATION

A. Challenges in data processing

The collection of diverse, high-volume geospatial data can pose challenges for designing effective emergency response management systems. Traditional desktop computers can be quickly overwhelmed by large volumes of geospatial data, necessitating the use of server computers with large amounts

¹Our work has primarily focused on motor vehicle crashes, however, the approach we take is generalizable and apply to other incidents (except crime as it has different dynamics [13]) that are responded to by our partners.

²For a comprehensive review of the state-of-the-art, please see our survey on incident prediction, resource allocation, and dispatch models [4].

of main memory, a cluster of desktop computers, or cloud computing resources. These resources can be financially prohibitive to smaller government organizations. Thus, precise programs are essential to minimize costs. These programs will inevitably contain spatial operations or queries to a spatial database. Spatial Operations and queries to spatial databases are computationally expensive, most notably spatial range, join, and k-nearest neighbor operations [14]. In an emergency response system, a variety of geospatial data formats (vector and raster) may arise, necessitating complex solutions: NoSQL data-stores or complex data structures, which can add additional processing time for the CPU [15].

To illustrate the problem, consider the context of wildfire response. Data regarding fire occurrence and relevant features are often collected and stored by different agencies and sources. Importantly, such data is usually stored in different data models. For example, while fire occurrence data is usually stored in the vector form, information about vegetation, fuel, and topographic features is available in a raster model. These two data models use different storage mechanisms and computational methods that make it difficult to combine them. Also, collecting information about all relevant features through high spatial and temporal periods is challenging and mining such large-scale feature data is a massive computational bottleneck. To tackle such issues, we designed *WildfireDB*, a collection of over 17 million data points that capture fire spread in continental United States in the last decade [7]. Our data generation, to a large extent, is based on work on large-scale vector and raster analysis [16], and uses a principled algorithmic approach to merge large-scale raster and vector data. *WildfireDB* can be used to model the spread of wildfires and inform response strategies [17]. We manage incident data

collection and joins in a similar manner.

B. Data Sources

Relevant and accurate data is essential for performant statistical models. In our work in incident forecasting, we have relied on four primary types of data: automobile incidents, roadway geometry, traffic, and weather data. Past automobile accidents are relevant to the forecasting of future incidents because segments with a high recent incident rate are more likely to have incidents in the future [18]–[20]. Spatial roadway features like lanes, speed limit, and road intersections often signify different driver behavior and, along with segment shape, are effective predictors of automobile incidents [21]–[23]. Weather data is valuable as precipitation, visibility, snowfall, and light levels can potentially alter driver behavior and is a valuable predictor of accident rates [18], [24], [25]. Traffic data is critical because it gives information about how drivers behave on roadways through speed and free-flow speed. Traffic data also adds congestion, a measurement we have found to be one of the most predictive features.

III. INCIDENT FORECASTING

Incident forecasting is necessary to understand the future demand of emergency resources in a given region, which is critical for a proactive ERM system. Road accident prediction is well studied in the literature, and a variety of modeling approaches having been applied in the domain. These include statistical models such as hierarchical Poisson Models [26], [27] and zero-inflated models [28], [29]. In recent years, data mining models such as neural networks [30]–[32] and support vector machines [33], [34] have also been explored. We refer interested readers to our survey on emergency response for a detailed analysis of prior work in forecasting road accidents (as well as allocation and dispatch techniques) [4].

A shortcoming of prior accident forecasting approaches is their inability to deal with large data sparsity. Roadway accident data is extremely sparse – we have observed more than 99% sparsity on interstate highway segments in Tennessee (a state in the USA with a total area of over 100,000 sq. km.) when discretized into four hour time windows (more details can be found in paper [8]). Even zero-inflated models, the only class of statistical models that have been shown to work fairly well on sparse spatial-temporal data, fail to perform well when trained on these sparse roadway incidents. Prior approaches are therefore ill-equipped to make accurate short-term predictions, since as the temporal resolution increases, the sparsity increases as well. While long-term predictions can be useful to analyze policies (optimize road construction, for example), short-term predictions are desirable for allocating and dispatching responders.

To address these challenges, we have developed an incident prediction pipeline that integrates various sparsity mitigation and modeling techniques, as shown in figure 2. The pipeline consists of five major stages: (1) clustering, (2) resampling, (3) learning, (4) inference, and (5) integration with allocation and dispatch algorithms. First, we use clustering to identify

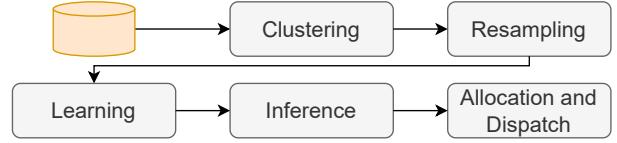


Fig. 2. Overview of our incident prediction approach. We extract spatial temporal information from a variety of data sources, focus on heterogeneity not explicitly modeled in the feature space by identifying clusters, perform resampling to address sparsity, and learn models on incident occurrence.

spatial segments with similar patterns of incident occurrence. This aggregates the data for each cluster to avoid overfitting on individual segments while capturing heterogeneity that is not explicitly modeled in the feature space.

Next, we perform synthetic under- and over-sampling to address sparsity by “balancing” the data. Naive synthetic sampling performs poorly, however, as the relative frequencies of incident occurrence vary significantly across clusters. We address this by performing resampling such that the ratio of accidents between each cluster is the same as in the original dataset. Specifically, we start with the cluster having the highest frequency of incident occurrence (cluster A, say) and perform synthetic resampling such that the number of positive data points (spatial segments in temporal windows that *have* accidents) is the same as the number of negative data points (spatial segments in temporal windows that *do not* have accidents). We then apply resampling to the other clusters such that the ratio of accident frequency for any given cluster compared to cluster A is the same as in the original dataset.

After clustering and resampling, we apply and compare various well-known spatial-temporal forecasting models to the processed dataset, including logistic regression, zero-inflated Poisson regression, random forests, and artificial neural networks. We use a suite of metrics to compare the performance of each model. Conversations with first responders revealed that *ranking* roadway segments based on relative risk is useful for resource allocation. Therefore, we report the correlation of each model’s marginal accident likelihood distribution over space with the real accident distribution in addition to standard statistical metrics like F1-scores and accuracy. Another important consideration is that these models are part of a pipeline that informs allocation decision making algorithms. Rather than only assess the models in isolation, we integrate them with an allocation decision agent and report each model’s final impact on response times to simulated incidents.

We evaluated the forecasting pipeline on three years of road accident data from the interstate highway network of Tennessee and found it can significantly reduce incident response times compared to current approaches followed by first responders [8]. An important result of our work is the finding that standard statistical metrics like F1-score do not give an accurate representation of each model’s performance when integrated with the full ERM pipeline: when used for resource allocation, the models which performed best in terms of these metrics ended up reducing response times less than models with ‘worse’ scores. This result highlights the need to

treat ERM pipelines holistically during both development and deployment.

IV. INCIDENT DETECTION

Emergency response is highly dependent on the time of incident reporting. However, the traditional approach of receiving incident reports (e.g., calling 911 in the USA) has time delays. Crowdsourcing platforms such as Waze provide an opportunity for early identification of incidents. However, detecting incidents from crowdsourced data streams is difficult due to the noise and uncertainty associated with such data. Further, simply optimizing over detection accuracy can compromise spatial-temporal localization of the inference, thereby making such approaches infeasible for real-world deployment.

We designed *CROME* (Crowdsourced Multi-objective Event Detection), that quantifies the relationship between the performance metrics of incident classification (e.g., F1 score) and the requirements of model practitioners (e.g., 1 km. radius for incident detection) [9]. First, we show how crowdsourced reports, ground-truth historical data, and other relevant determinants such as traffic and weather can be used together in a Convolutional Neural Network (CNN) architecture for early detection of emergency incidents. However, naively maximizing the accuracy of detection can harm emergency response. For example, consider an approach that raises an alert every time a report is received on the crowdsourcing platform. While such an approach will correctly detect many (if not most) ground-truth incidents, it will lead to wastage of resources if a responder is dispatched on every alert. To tackle this challenge, we use a Pareto optimization-based approach to optimize the output of the CNN in tandem with practitioner-centric parameters to balance detection accuracy and spatial-temporal localization. We evaluated the efficacy of CROME using crowdsourced data from Waze and traffic accident reports from Nashville, TN, USA [9]. Our findings suggest that CROME can detect more than 40% of the incidents earlier than traditional mechanisms.

V. RESOURCE ALLOCATION AND DISPATCH

Once we have an understanding of when and where incidents are likely to occur, the final stage of the ERM pipeline is using forecasting models to make *allocation* and *dispatch* decisions. Resource allocation (also referred to as the stationing problem [10], [13]) is the problem of spatially distributing resources such as ambulances in anticipation of stochastic incidents, while the dispatch problem is determining which resource to send to an incident once a call for service is received. The distinction between these problems is not always clear, as allocation policies create the implicit dispatching policy of sending the nearest available responder to an incident (greedy dispatch). In fact, conversations with emergency response agencies reveal that there are also legal constraints on dispatching models – it is very difficult to judge the severity of an incident from a call for service [10], so it is imperative to follow this greedy strategy to ensure that if the incident is severe, a responder arrives as quickly as possible.

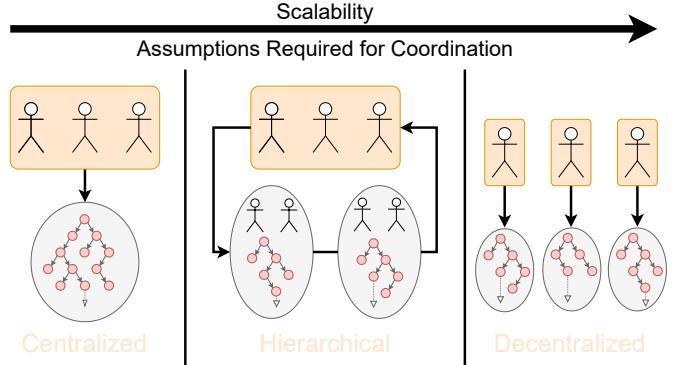


Fig. 3. A spectrum of planning approaches for online ERM planning, each with inherent tradeoffs. A centralized approach uses a monolithic state representation. A decentralized approach segments the state-action space for each individual agent, and each agent performs their own decision making. A hierarchical approach segments the planning problem into sub-problems to improve scalability without agents estimating other agents' actions.

Since dispatching policies are constrained by policy, we focus on solving the allocation problem.

A popular strategy used to plan under uncertainty is to learn a *policy*, which is a general mapping from states of the environment to actions that should be taken. Reinforcement Learning and Approximate Dynamic Programming are examples of this approach, and have recently been applied to ambulance allocation and similar resource allocation problems such as ride-pooling [35]–[38]. Unfortunately, these methods are slow to converge when applied to large-scale, practical problems. This certainly applies to ERM – for example, a simple case of just 20 responders and 30 potential waiting locations has $\frac{|R|!}{(|R|-|V|)!} = \frac{30!}{10!} = 7.31e+25$ possible assignments at each decision point. Since urban environments are nonstationary, a policy might be out of date by the time it is deployed. Policies can also have issues adapting to failures that were not considered during training.

An alternative to learning a policy is online planning: given a model of the environment, heuristic search approaches such as Monte-Carlo tree search (MCTS) can find promising actions for the current system state. The environmental models can be updated to reflect any detected changes, and the planner can incorporate these updates into the decision making process to adapt to a dynamic environment. Such an approach has been applied to real-time control problems in domains such as the smart grid [39], game playing [40], and autonomous driving [41]. However, state-of-the-art MCTS approaches have difficulty converging in a reasonable amount of time for problems with large state-action spaces such as responder allocation. Another consideration is that many emergency incidents can cause failures in communication networks, so an allocation strategy should be able to operate with limited communication. Therefore, scalable online planning approaches that are robust to communication outages are needed for ERM.

We have developed two approaches to tackle the problem of scalable online resource allocation, as shown in figure 3. The first is decentralized planning, where each resource builds its

own decision tree to independently determine its actions [10]. By limiting the state-action space for each resource to only information relevant to the agent, the overall size of the state-action space is significantly decreased and the approach scales notably better than centralized approaches. A decentralized approach also addresses partial communication losses: since most mobile resources like ambulances are now equipped with compute units, they can perform their individual planning locally, only needing to communicate sporadically for coordination. Unfortunately, there is also a downside to decentralization – during individual planning, resources must make assumptions regarding the behavior of other resources, which can lead to sub-optimal decisions when those assumptions do not hold. This tradeoff is worthwhile for problems like ERM that are difficult to solve tractably with centralized approaches.

The second approach we developed is hierarchical planning, which leverages the spatial structure of the ERM environment to split a large resource allocation system into smaller sub-problems [6]. This decomposition drastically decreases the overall complexity of the planning problem compared to a centralized approach, but preserves the dependencies between nearby resources unlike a decentralized approach. We created a principled framework for solving the sub-problems and tackling the interaction between them. First, the overall spatial area is divided into regions using k-means clustering on historical data of incident occurrence. A high-level planner assigns responders to each region based on estimated waiting times using a queueing approximation. Then, a low-level planner solves the allocation problem within each sub-area in parallel using MCTS. Compared to a decentralized framework, a hierarchical approach is more reliant on communication channels within a region and is less scalable, but better models inter-resource dependencies.

We applied our decentralized and hierarchical planning frameworks to the problem of dynamic ambulance allocation in the city of Nashville, TN. We found that the decentralized approach easily scaled to the full ERM system with reasonable computation times, and reduced mean incident response times compared to centralized planning [10]. The hierarchical framework was also able to scale to the full system, and further improved average response times compared to the decentralized approach due to making fewer assumptions regarding the agents' interactions [6]. We also found that both approaches are able to adapt to changes in the environment, such as large shifts in the incident distribution or ambulance failures.

VI. RESEARCH OPPORTUNITIES AND DISCUSSIONS

While tremendous progress has been made in designing emergency response management systems, the research area of building and deploying efficient and effective ERM in practice is still in its infancy. In this paper, we presented our recent efforts as the first step in developing key ERM components in collaboration with our community partners. There are many open research problems and real-world challenges in this exciting new area that need further studies.

Uncertainty and guarantee: Many current approaches have been designed to be robust to uncertainty and noise in the existing data. However, they barely deal with the uncertainty from other services/systems or human behavior in the real world, which could significantly affect the performance of incidents prediction and resource allocation. It is an open question of how to model the behaviors of a complex and integrated system considering this real-world uncertainty and, more importantly, providing run-time guarantees.

Robustness: Handling non-stationarity and environmental disruptions that will very likely reduce the effectiveness of the decision support algorithms are a key concern. In this context, it is clear that the online tree search algorithms are superior to a pure offline learning-based strategy that relies on reinforcement learning or a neural network approach. However, it is not yet clear if a combined approach that bootstraps the tree search can lead to improved outcomes. While preliminary experiments conducted by the team do indicate that such an approach can help with non-stationarity, further investigation is required.

Interaction with domain experts: Despite the increasing intelligence in ERM, there is still a high demand for inputs and interaction from the domain experts, e.g., emergency responders and city decision-makers. Two critical questions are (1) how to support experts to use our systems (e.g., providing requirements and domain knowledge), and then convert the information to machine-understandable language, and (2) how to help decision-makers interpret and trust the outcomes and decisions made by the ERM systems?

Fairness and equity: As data-driven approaches, machine learning models rely significantly on existing data. If unfairness exists in historical data, it is very challenging or impossible to learn fair models without exogenous intervention. Therefore, measuring and validating group fairness and addressing equity when designing ERM pipelines is still an important yet challenging future direction.

Practical Constraints: Another issue that is an opportunity for further research is the integration of practical constraints into the decision support procedures. Such dynamic considerations include reduction in the number of ambulances available, constraints on the maximum mileage that can be put on a vehicle, road closures and reduced manpower. Recent experiments by Ferber et.al. in [42] have shown that including the decision outcomes in the prediction pipeline can lead to improved forecasting models. Such considerations may also be better in integrating various practical constraints that can be encountered.

ACKNOWLEDGEMENTS

This work is sponsored by the National Science Foundation under award numbers CNS1640624, CNS1818901, IIS1814958, IIS-1815459 and a grant from the Tennessee Department of Transportation.

REFERENCES

- [1] H. Jaldell, "How important is the time factor? saving lives using fire and rescue services," *Fire Technology*, vol. 53, no. 2, pp. 695–708, 2017.

[2] H. Jaldell, P. Lebnak, and A. Amornpetchsathaporn, "Time is money, but how much? the monetary value of response time for thai ambulance emergency services," *Value in Health*, vol. 17, no. 5, pp. 555–560, 2014.

[3] S. H. Hattis, *Crime in the United States*. Lanham Berman Press, 2015.

[4] A. Mukhopadhyay, G. Pettet, S. M. Vazirizade, D. Lu, A. Jaimes, S. E. Said, H. Baroud, Y. Vorobeychik, M. Kochenderfer, and A. Dubey, "A review of incident prediction, resource allocation, and dispatch models for emergency management," *Accident Analysis & Prevention*, vol. 165, p. 106501, 2022.

[5] A. Mukhopadhyay, "Robust Incident Prediction, Resource Allocation and Dynamic Dispatch," Ph.D. dissertation, Vanderbilt University, 2019.

[6] G. Pettet, A. Mukhopadhyay, M. J. Kochenderfer, and A. Dubey, "Hierarchical planning for resource allocation in emergency response systems," in *Proceedings of the ACM/IEEE 12th International Conference on Cyber-Physical Systems*, May 2021, pp. 155–166.

[7] S. Singla, A. Mukhopadhyay, M. Wilbur, T. Diao, V. Gajjewar, A. Eldawy, M. Kochenderfer, R. Shachter, and A. Dubey, "WildfireDB: An Open-Source Dataset Connecting Wildfire Spread with Relevant Determinants," in *Conference on Neural Information Processing Systems Track on Datasets and Benchmarks*, 2021.

[8] S. M. Vazirizade, A. Mukhopadhyay, G. Pettet, S. El Said, H. Baroud, and A. Dubey, "Learning incident prediction models over large geographical areas for emergency response," in *IEEE International Conference on Smart Computing*, 2021, pp. 424–429.

[9] Y. Senarath, A. Mukhopadhyay, S. Vazirizade, H. Purohit, S. Nannapaneni, and A. Dubey, "Practitioner-centric approach for early incident detection using crowdsourced data for emergency services," in *IEEE International Conference on Data Mining*, 2021, pp. 1318–1323.

[10] G. Pettet, A. Mukhopadhyay, M. Kochenderfer, Y. Vorobeychik, and A. Dubey, "On algorithmic decision procedures in emergency response systems in smart and connected communities," in *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 2020, pp. 1046–1054.

[11] Y. Senarath, S. Nannapaneni, H. Purohit, and A. Dubey, "Emergency Incident Detection from Crowdsourced Waze Data using Bayesian Information Fusion," in *The 2020 IEEE/WIC/ACM International Joint Conference On Web Intelligence And Intelligent Agent Technology*, 2020.

[12] H. Purohit, S. Nannapaneni, A. Dubey, P. Karuna, and G. Biswas, "Structured summarization of social web for smart emergency services by uncertain concept graph," in *2018 IEEE International Science of Smart City Operations and Platforms Engineering in Partnership with Global City Teams Challenge*, 2018, pp. 30–35.

[13] A. Mukhopadhyay, C. Zhang, Y. Vorobeychik, M. Tambe, K. Pence, and P. Speer, "Optimal allocation of police patrol resources using a continuous-time crime model," in *International conference on decision and game theory for security*. Springer, 2016, pp. 139–158.

[14] J. Yu, J. Wu, and M. Sarwat, "Geospark: A cluster computing framework for processing large-scale spatial data," in *International Conference on Advances in Geographic Information Systems*, 2015, pp. 1–4.

[15] S. Li, S. Dragicevic, F. A. Castro, M. Sester, S. Winter *et al.*, "Geospatial big data handling theory and methods: A review and research challenges," *ISPRS journal of Photogrammetry and Remote Sensing*, vol. 115, pp. 119–133, 2016.

[16] S. Singla, A. Eldawy, T. Diao, A. Mukhopadhyay, and E. Scudiero, "The Raptor Join Operator for Processing Big Raster + Vector Data," in *International Conference on Advances in Geographic Information Systems*, 2021, p. 324–335.

[17] T. Diao, S. Singla, A. Mukhopadhyay, A. Eldawy, R. Shachter, and M. Kochenderfer, "Uncertainty aware wildfire management," *arXiv preprint arXiv:2010.07915*, 2020.

[18] A. Mukhopadhyay, Y. Vorobeychik, A. Dubey, and G. Biswas, "Prioritized allocation of emergency responders based on a continuous-time incident prediction model," in *International Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 168–177.

[19] A. Mukhopadhyay, G. Pettet, C. Samal, A. Dubey, and Y. Vorobeychik, "An online decision-theoretic pipeline for responder dispatch," in *International Conference on Cyber-Physical Systems*, 2019, pp. 185–196.

[20] A. Mukhopadhyay, Z. Wang, and Y. Vorobeychik, "A decision theoretic framework for emergency responder dispatch," in *International Conference on Autonomous Agents and MultiAgent Systems*, 2018, pp. 588–596.

[21] V. Shankar, F. Mannering, and W. Barfield, "Effect of roadway geometries and environmental factors on rural freeway accident frequencies," *Accident Analysis & Prevention*, vol. 27, no. 3, pp. 371–389, 1995.

[22] M. Poch and F. Mannering, "Negative binomial analysis of intersection-accident frequencies," *Journal of Transportation Engineering*, vol. 122, no. 2, pp. 105–113, 1996.

[23] H. C. Chin and M. A. Quddus, "Applying the random effect negative binomial model to examine traffic accident occurrence at signalized intersections," *accident analysis & prevention*, vol. 35, no. 2, pp. 253–259, 2003.

[24] Y. Qi, B. L. Smith, and J. Guo, "Freeway accident likelihood prediction using a panel data analysis approach," *Journal of Transportation Engineering*, vol. 133, no. 3, pp. 149–156, 2007.

[25] P. Songchitruksa and K. Balke, "Assessing weather, environment, and loop data for real-time freeway incident prediction," *Transportation Research Record: Journal of the Transportation Research Board*, no. 1959, pp. 105–113, 2006.

[26] M. Deublein, M. Schubert, B. T. Adey, J. Köhler, and M. H. Faber, "Prediction of road accidents: A Bayesian hierarchical approach," *Accident Analysis & Prevention*, vol. 51, pp. 274–291, 2013.

[27] M. A. Quddus, "Modelling Area-Wide Count Outcomes with Spatial Correlation and Heterogeneity: An Analysis of London Crash Data," *Accident Analysis and Prevention*, vol. 40, pp. 1486–1497, 2008.

[28] X. Qin, J. N. Ivan, and N. Ravishanker, "Selecting exposure measures in crash rate prediction for two-lane highway segments," *Accident Analysis & Prevention*, vol. 36, no. 2, pp. 183–191, 2004.

[29] H. Huang and H. C. Chin, "Modeling road traffic crashes with zero-inflation and site-specific random effects," *Statistical Methods & Applications*, vol. 19, no. 3, pp. 445–462, 2010.

[30] A. Pande and M. Abdel-Aty, "Assessment of freeway traffic parameters leading to lane-change related collisions," *Accident Analysis & Prevention*, vol. 38, no. 5, pp. 936–948, 2006.

[31] L. Zhu, F. Guo, R. Krishnan, and J. W. Polak, "The use of convolutional neural networks for traffic incident detection at a network level," in *Transportation Research Board Annual Meeting*, 2018.

[32] J. Bao, P. Liu, and S. V. Ukkusuri, "A spatiotemporal deep learning approach for citywide short-term crash risk prediction with multi-source data," *Accident Analysis & Prevention*, vol. 122, pp. 239–254, 2019.

[33] Y. Zhang and Y. Xie, "Forecasting of short-term freeway volume with v-support vector machines," *Transportation Research Record*, vol. 2024, no. 1, pp. 92–99, 2007.

[34] R. Yu and M. Abdel-Aty, "Utilizing support vector machine in real-time crash risk evaluation," *Accident Analysis & Prevention*, vol. 51, pp. 252–259, 2013.

[35] M. S. Maxwell, S. G. Henderson, and H. Topaloglu, "Ambulance redeployment: An approximate dynamic programming approach," in *2009 Winter Simulation Conference*, Dec. 2009, pp. 1850–1860.

[36] A. A. Nasrollahzadeh, A. Khademi, and M. E. Mayorga, "Real-Time Ambulance Dispatching and Relocation," *Manufacturing & Service Operations Management*, vol. 20, no. 3, pp. 467–480, Apr. 2018.

[37] X. Yu and S. Shen, "An Integrated Decomposition and Approximate Dynamic Programming Approach for On-Demand Ride Pooling," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 9, pp. 3811–3820, Sep. 2020.

[38] S. Shah, M. Lowalekar, and P. Varakantham, "Neural Approximate Dynamic Programming for On-Demand Ride-Pooling," *AAAI Conference on Artificial Intelligence*, vol. 34, no. 01, pp. 507–515, Apr. 2020.

[39] H. Shuai and H. He, "Online Scheduling of a Residential Microgrid via Monte-Carlo Tree Search and a Learned Model," *IEEE Transactions on Smart Grid*, vol. 12, no. 2, pp. 1073–1087, Mar. 2021.

[40] D. Silver, A. Huang, C. J. Maddison, *et al.*, "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.

[41] C. Hoel, K. Driggs-Campbell, K. Wolff, L. Laine, and M. J. Kochenderfer, "Combining Planning and Deep Reinforcement Learning in Tactical Decision Making for Autonomous Driving," *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 2, pp. 294–305, Jun. 2020.

[42] A. Ferber, B. Wilder, B. Dilksina, and M. Tambe, "Mipa: Mixed integer program as a layer," in *AAAI Conference on Artificial Intelligence*, vol. 34, no. 02, 2020, pp. 1504–1511.