

Designing Emergency Response Pipelines : Lessons and Challenges

Ayan Mukhopadhyay
Stanford University
ayanmukh@stanford.edu

Geoffrey Pettet
Vanderbilt University
geoffrey.a.pettet@vanderbilt.edu

Mykel Kochenderfer
Stanford University
mykel@stanford.edu

Abhishek Dubey
Vanderbilt University
abhishek.dubey@vanderbilt.edu

Abstract

Emergency response to incidents such as accidents, crimes, and fires is a major problem faced by communities. Emergency response management comprises of several stages and sub-problems like forecasting, resource allocation, and dispatch. The design of principled approaches to tackle each problem is necessary to create efficient emergency response management (ERM) pipelines. Over the last six years, we have worked with several first responder organizations to design ERM pipelines. In this paper, we highlight some of the challenges that we have identified and lessons that we have learned through our experience in this domain. Such challenges are particularly relevant for practitioners and researchers, and are important considerations even in the design of response strategies to mitigate disasters like floods and earthquakes.

Introduction

Emergency response to incidents such as accidents, medical calls, urban crimes, poaching, and fires is one of the most pressing problems faced by communities across the globe. Emergency response must be fast and efficient in order to minimize loss of lives (Jaldell 2017; Jaldell, Lebnak, and Amornpetsathaporn 2014). Significant attention in the last several decades has been devoted to studying emergency incidents and response. The broader goal of designing ERM systems is to enable communities to deal with emergency response in a manner that is principled, proactive, and efficient. ERM systems based on data-driven models can help reduce both human and financial losses. Insights from principled approaches can also be used to improve policies and safety measures. Although such models are increasingly being adopted by government agencies, emergency incidents still cause thousands of deaths and injuries and also result in losses worth more than billions of dollars directly or indirectly each year (Hattis 2015).

ERM can be divided into five major components: 1) mitigation, 2) preparedness, 3) detection, 4) response, and 5) recovery. The stages of ERM systems are heavily inter-linked (Mukhopadhyay et al. 2020). Mitigation involves sustained efforts to reduce long-term risks to people and prop-

erty. It also involves creating forecasting models to understand the spatial and temporal characteristics of incidents. Preparedness involves creating policy and allocating resources that enable emergency response management. The third phase seeks to use automated techniques to detect incidents as they happen in order to expedite response. The dispatch phase, the most critical phase in the field, involves responding to incidents when they occur. Finally, the recovery phase seeks to provide support and sustenance to communities and individuals affected by the emergency. While most of the 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 such that intricate interdependencies are considered. In this paper, we highlight some of the challenges that we have faced while designing ERM pipelines and lessons that we have learned through the process. A large portion of the insights have come from our colleagues at the Nashville Fire Department and the Tennessee Department of Transportation, who have provided invaluable domain expertise to us.

Challenges and Lessons

While ERM forms a crucial component of the cities and governments, designing and deploying principled approaches to ERM is challenging. We consider the following to be the main challenges in the design and deployment of ERM pipelines.

1. **How to forecast incident occurrence?** It has been noted that emergency incidents are generally difficult to predict due to the inherent random nature of such incidents and spatially varying factors (Shankar, Mannering, and Barfield 1995; Mukhopadhyay et al. 2017). Incidents are highly sporadic, making it particularly difficult to learn models of incident occurrence. For example, well-known regression models such as Poisson regression and negative binomial regression have been shown to perform poorly on accident data due to the prevalence of zero counts (Mukhopadhyay et al. 2020). Incident data has also been shown to be particularly sensitive to the scale of spatial and temporal resolutions, thereby making it difficult to perform meaningful inference. There are several

approaches that have shown to alleviate these concerns. First, identifying clusters of incidents (both spatial and non-spatial) has shown to balance model variance and spatial heterogeneity particularly well (Sasidharan, Wu, and Menendez 2015; Mukhopadhyay et al. 2017). Second, dual state models like zero-inflated Poisson models can be used to address the issue of a high number of zero counts in data (Qin, Ivan, and Ravishanker 2004).

2. **When to optimize?** Arguably, the most important component of ERM pipelines is to dispatch responders when incidents occur. While resource allocation and dispatch to emergency incidents evolve in highly uncertain and dynamic environments, the expectation is that response is very timely (Felder and Brinkmann 2002; Mukhopadhyay, Wang, and Vorobeychik 2018). Approaches to optimize dispatch typically focus on decision-making *after* an incident occurs (Toro-DíAz et al. 2013; Mukhopadhyay, Wang, and Vorobeychik 2018; Keneally, Robbins, and Lunday 2016). However, our conversations and collaborations with first responders revealed that there is limited applicability of such approaches in practice due to two important reasons. First, response to incidents occurs almost instantaneously after a report is received. Although optimizing dispatch can minimize response times in the long run, time spent to optimize dispatch *after* incidents occur is perceived as costly in the field. Second, it is almost impossible to judge the severity of an incident from a call for service. Consequently, it is imperative for first responders to follow a *greedy* strategy and dispatch the closest available responder to incidents. An alternative approach is to periodically optimize the spatial distribution of responders *between* incidents. Pettet et al. (Pettet et al. 2020) introduced this idea recently for emergency response. While there are challenges with respect to scalability of such an approach, planning between incidents is much more applicable in the field, as it does not violate constraints under which first responder operate.
3. **How to model communication?** Approaches to tackle emergency response typically assume that the agents can observe the world completely and communicate with centralized servers and each other. This assumption is usually satisfied in practice. However, in scenarios that involve natural disasters (like floods, earthquakes, wildfires, etc.) communication mechanisms can break down and power failures are common. In such scenarios, it is important that agents can optimize response based on information gathered locally. One way to approach such a problem is to design distributed approaches for ERM, in which agents can optimize their own decisions (Pettet et al. 2020). This is feasible since modern agents (ambulances, for example) are equipped with laptops. While distributed approaches perform worse than their centralized counterparts, they provide the benefit of performance in scenarios where communication is challenging.
4. **How to model the environment?** A problem in designing approaches to ERM is that environmental conditions under which response takes place is dynamic. Consider a decision-theoretic model for dispatching responders. For

example, see the semi-Markov decision process formulation to minimize expected response times to accidents (Mukhopadhyay, Wang, and Vorobeychik 2018). An approach to solve such large-scale decision problems is to use a simulator to find a policy that picks the optimal action given the state of the world. However, in urban areas, events like road closures, constructions, or increased traffic due to a public gathering can drastically alter the distribution of incidents. Further, ambulances can be unavailable due to breakdowns or maintenance. In such cases, it is crucial that the actual state of the environment is taken into account while creating allocation and dispatch decisions. One approach is to create high-fidelity models of covariates such as traffic and weather. Such models can then be used in online forecasting models, that can accommodate incoming streams of updated information (Mukhopadhyay et al. 2019). Similarly, decision-theoretic approaches that can quickly compute promising actions for the current state of the world can be more valuable in emergency response than approaches that find policies for the entire state space of the problem (Pettet et al. 2020).

The principled design of ERM systems is an important problem faced by communities. As smart and connected communities evolve, they present both opportunities and challenges to manage ERM systems. In this short paper, we highlight common problems and lessons learned through our experience in designing ERM systems over the last few years.

Acknowledgement

We would like to acknowledge the National Science Foundation and the Center of Automotive Research at Stanford for funding this research. We would also like to thank the Nashville Fire Department (NFD) and the Tennessee Department of Transportation (TDOT) for collaborating with us and providing invaluable knowledge about the intricacies of emergency response.

References

- [Felder and Brinkmann 2002] Felder, S., and Brinkmann, H. 2002. Spatial allocation of emergency medical services: minimising the death rate or providing equal access? *Regional Science and Urban Economics* 32(1):27–45.
- [Hattis 2015] Hattis, S. H. 2015. *Crime in the United States*. Lanham Bernan Press.
- [Jaldell, Lebnak, and Amornpetsathaporn 2014] Jaldell, H.; Lebnak, P.; and Amornpetsathaporn, A. 2014. Time is money, but how much? the monetary value of response time for thai ambulance emergency services. *Value in Health* 17(5):555–560.
- [Jaldell 2017] Jaldell, H. 2017. How important is the time factor? saving lives using fire and rescue services. *Fire Technology* 53(2):695–708.
- [Keneally, Robbins, and Lunday 2016] Keneally, S. K.; Robbins, M. J.; and Lunday, B. J. 2016. A markov decision process model for the optimal dispatch of military medical evacuation assets. *Health Care Management Science* 19(2):111–129.

- [Mukhopadhyay et al. 2017] Mukhopadhyay, A.; Vorobeychik, Y.; Dubey, A.; and Biswas, G. 2017. Prioritized allocation of emergency responders based on a continuous-time incident prediction model. In *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 168–177.
- [Mukhopadhyay et al. 2019] Mukhopadhyay, A.; Pettet, G.; Samal, C.; Dubey, A.; and Vorobeychik, Y. 2019. An online decision-theoretic pipeline for responder dispatch. In *International Conference on Cyber-Physical Systems (ICCPs)*, 185–196. ACM.
- [Mukhopadhyay et al. 2020] Mukhopadhyay, A.; Pettet, G.; Vazirizade, S.; Lu, D.; Baroud, H.; Jaimes, A.; Vorobeychik, Y.; Kochenderfer, M.; and Dubey, A. 2020. A review of emergency incident prediction, resource allocation and dispatch models.
- [Mukhopadhyay, Wang, and Vorobeychik 2018] Mukhopadhyay, A.; Wang, Z.; and Vorobeychik, Y. 2018. A decision theoretic framework for emergency responder dispatch. In *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 588–596.
- [Pettet et al. 2020] Pettet, G.; Mukhopadhyay, A.; Kochenderfer, M.; Vorobeychik, Y.; and Dubey, A. 2020. On algorithmic decision procedures in emergency response systems in smart and connected communities. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [Qin, Ivan, and Ravishanker 2004] Qin, X.; Ivan, J. N.; and Ravishanker, N. 2004. Selecting exposure measures in crash rate prediction for two-lane highway segments. *Accident Analysis & Prevention* 36(2):183–191.
- [Sasidharan, Wu, and Menendez 2015] Sasidharan, L.; Wu, K.-F.; and Menendez, M. 2015. Exploring the application of latent class cluster analysis for investigating pedestrian crash injury severities in switzerland. *Accident Analysis & Prevention* 85:219–228.
- [Shankar, Mannering, and Barfield 1995] Shankar, V.; Mannering, F.; and Barfield, W. 1995. Effect of roadway geometrics and environmental factors on rural freeway accident frequencies. *Accident Analysis & Prevention* 27(3):371–389.
- [Toro-DíAz et al. 2013] Toro-DíAz, H.; Mayorga, M. E.; Chanta, S.; and Mclay, L. A. 2013. Joint location and dispatching decisions for emergency medical services. *Computers & Industrial Engineering* 64(4):917–928.