Incident Prediction and Response Optimization

Doctoral Consortium

Ayan Mukhopadhyay
Vanderbilt University
Advisor : Yevgeniy Vorobeychik
ayan.mukhopadhyay@vanderbilt.edu

ABSTRACT

In urban areas across the globe, incidents like crime, fire and accidents often result in massive losses of life and property. In such scenarios, quick response can minimize or prevent damage. Emergency responder services are eager to adopt mechanisms that aid real-time decision making in the field. With limited resources however, it is imperative that resource allocation and dispatch decision are taken in a systematic and principled manner. This problem gives rise to challenges that are both technically intriguing and have immense applicability in practice. In this paper, I describe how my research seeks to identify problems in this space and addresses them. I also describe some of the existing challenges in the domain and potential approaches that could be used to tackle such problems.

KEYWORDS

Survival Analysis; Decision Theory; Stochastic Optimization; Markov Decision Processes; Machine Learning

ACM Reference Format:
Ayan Mukhopadhyay. 2018. Incident Prediction and Response Optimization. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

An increase in urban population density has led to a number of problems across the world, with traffic congestion, accidents, and crime being on the forefront of urban issues. To manage incidents, cities resort to diverse groups of emergency responder organizations, including fire and police departments. The broader problem of ensuring effective and fast response subsumes three implicit sub-problems - 1) forecasting incidents like crime, fire and traffic accidents in space and time, 2) placing responders in anticipation of such incidents and 3) deploying responders when incidents happen. Although the first two sub-problems have been extensively studied [4, 6–9]), there are several crucial gaps in literature which limit the practical applicability of such approaches. Surprisingly, the problem of dynamically dispatching responders has received considerably less attention. A principled approach for dispatch is however, extremely important. I seek to systematically address these three problems in my research. In this paper, I describe our work in this area, and highlight current challenges and scope for future work. Specifically, in section 2, I describe my work on designing generative models for spatial-temporal incident prediction. Then, based on such an incident prediction approach, in sections 3 and 4, I describe my work on responder placement and dispatch respectively. I touch upon how my research is being implemented in practice in section 5. Finally, in section 6, I highlight some of the current challenges in this domain and how my ongoing work attempts to solve them.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

2 INCIDENT PREDICTION

Incident prediction models can be broadly classified into three categories: spatial models, which predict future incidents by identifying common spatial patterns in historical incident data, spatial-temporal models, which capture the "attractiveness" of a discrete set of spatial locations in a geographic area, and risk-terrain models which capture environmental factors affecting incidents. However, these approaches lack the presence of a spatial-temporal generative framework that can capture arbitrary correlates like weather, population density, time of day and so on. In order to address this issue, we propose the use of Survival Analysis [1] to predict incidents by creating a stochastic generative model which is continuous in time but discrete in space (the total geographic area is assumed to be discretized into a set of grids) [6]. A natural fit to the problem, parametric survival models let us capture arbitrary covariates while explicitly allowing to model time to arrival of incidents, thereby bridging an important gap in prior literature. We also propose learning a joint distribution over incident arrival time and severity, which aids responder dispatch significantly (as dispatching decisions are directly influenced by severity). An important concern in using such a model learned over a discrete set of grids is the trade off between over-fitting (learning a different model for each grid) and losing spatial heterogeneity not explicitly captured in the feature space (creating a single *universal* model). In order to tackle this, we introduce a hierarchical clustering approach by iteratively merging grids that are *similar*, with similarity defined over a set of carefully identified features. We compare the hierarchical-survival prediction mechanism to other state-of-the-art approaches [3, 8] and show that this approach performs at least as well as the others, while creating a drastic reduction in running-time (which is crucial for real-time prediction and dispatching decisions).

3 RESPONDER PLACEMENT

Armed with a generative model to predict and simulate incidents in space and time, I now describe our work on placing emergency responders to minimize response times in anticipation of such incidents. It is important to point out that a major consideration in designing such mechanisms is whether placement of responders affects future incidents or not - while police patrols responding to crimes affect future crimes, ambulances responding to accidents do not affect future incidents. This makes it imperative to create specialized models for different incident types. First, I describe how our work tackles the problem of optimally responding to crimes. We propose a two-stage optimization framework to minimize the expected response time to a set of incidents that are distributed according to a known arrival distribution [6]. In the first-stage, the program considers that a set of police patrols are placed in space under uncertainty about incidents; and in the second stage, patrols respond to crimes with the uncertainty resolved. The resulting stochastic program is solved by using Bender's decomposition and sample average approximation. A key consideration in this approach is that the distribution over incident arrival is actually dependent upon police placement, which makes the top level decisions affect the seconds-level response optimization. In order to tackle this, we propose a novel iterative scheme, that intuitively gives police repeated chances to respond to changing crime distributions. Simulations comparing our approach to existing placement strategies in the metropolitan area of Nashville demonstrate significant reduction in response times.

As described before, the problem of optimally placing ambulances to respond to accidents is fundamentally different. Working in life-saving scenarios in the field, such responder services need to provide guarantees of arrival times. Ambulance dispatch is also heavily affected by the reported severity of incidents. Also, unlike police patrols, ambulances do not continuously patrol urban areas and are often stationed together in depots. Leveraging on our work on predicting incident arrival and severity, we propose a novel optimization problem to maximize incident density coverage with restrictions on waiting times [4]. Any such problem formulation suffers from a major concern - constraints that bound arrival times are non-convex and non-linear. While there are existing heuristic approaches to solve such optimization problems, such algorithms fail when multiple responders need to be placed together in depots, as the search space becomes significantly more complex. We propose a novel heuristic algorithm based on greedy random adaptive search to solve this problem. Centered on locally searching around promising candidate solutions, the approach is both computationally feasible and provides solutions that outperform existing dispatch strategies used in practice.

4 RESPONDER DISPATCH

The final sub-problem in effective emergency responder management is to dynamically dispatch responders as incidents occur. We propose a principled decision theoretic framework for continuous-time resource dispatch [5]. We formulate the problem of responder dispatch as a Semi-Markov Decision Process which evolves in continuous time, and derive an equivalent discrete-time process for the formulation. We create a novel algorithm, that accesses a simulator to simultaneously estimate values of states, as well as estimate transition probabilities for the decision process. This work is inspired by existing work on sparsely sampling a simulator to estimate *values* of states [2]. Our work shows drastic reduction in response times over state-of-the-art dispatch strategies.

5 IMPLEMENTATION

The broad goal of our work in incident prediction is to provide emergency responder services with algorithms and tools that can aid real-time decision making. As part of this venture, a part of my research has been showcased in various global smart-city summits. Collaborating with the Police and the Fire departments of Nashville has immensely helped in getting expert-opinions and understanding the problem domain better. Currently, a web-based visualization and dispatch tool developed by us is being deployed at the Nashville Fire Department.

6 ONGOING WORK AND CHALLENGES

Having described our work on incident prediction and responder dispatch, I now draw attention to some of the existing challenges that this domain faces, and how our current work tries to address these problems. A major challenge that plagues this domain is the computational overhead involved in making real-time dispatch decisions. Most decision-theoretic formulations that capture the dynamics of an entire urban area result in complicated and large state spaces. While an optimal dispatch policy could be calculated offline for constant-time access, the problem definition can change at any moment in practice. Consider that an ambulance designated to serve a particular area breaks down. Emergency responders would need access to an immediate solution in such a scenario. Recalculating an optimal policy however, could potentially take days. A plausible solution to this is clustering the states of the decision process based on an appropriate similarity measure, such that when the problem structure changes, the policy for only a subset of the states needs to be re-computed. This would greatly enhance the practical applicability of our work in the field. Another problem faced by incident prediction models is the inherent bias in the input data due to responder placement. Consider a criminal who wishes to commit a crime at a specific location. Imagine now that a police patrol passes by, thereby causing the criminal to change his location for committing the crime. Since crime data only shows reported crime, the final location where the crime is committed is observed, but the original intended location is not. We are working to retrieve this underlying latent behavior model by creating an adversarial framework for learning about occurrences of crimes. This would help police plan their patrol routes and identify new areas that could be potential locations for crimes. I plan to investigate these two problems as well as work to deploy my research about crime prediction in practice.

REFERENCES

[1] David Roxbee Cox and David Oakes. 1984. *Analysis of survival data*. Vol. 21. CRC Press.

[2] Michael Kearns, Yishay Mansour, and Andrew Y Ng. 2002. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine learning* 49, 2 (2002), 193–208.

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

[4] Ayan Mukhopadhyay, Yevgeniy Vorobeychik, Abhishek Dubey, and Gautam Biswas. 2017. Prioritized Allocation of Emergency Responders based on a Continuous-Time Incident Prediction Model. In *International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 168–177.

[5] Ayan Mukhopadhyay, Zilin Wang, and Yevgeniy Vorobeychik. 2018. A Decision Theoretic Framework for Emergency Responder Dispatch. In *International Conference on Autonomous Agents and MultiAgent Systems (to appear)*. International Foundation for Autonomous Agents and Multiagent Systems.

[6] Ayan Mukhopadhyay, Chao Zhang, Yevgeniy Vorobeychik, Milind Tambe, Kenneth Pence, and Paul Speer. 2016. Optimal Allocation of Police Patrol Resources Using a Continuous-Time Crime Model. In *Conference on Decision and Game Theory for Security*.

[7] Francisco Silva and Daniel Serra. 2008. Locating emergency services with different priorities: the priority queuing covering location problem. *Journal of the Operational Research Society* 59, 9 (2008), 1229–1238.

[8] Chao Zhang, Victor Bucarey, Ayan Mukhopadhyay, Arunesh Sinha, Yundi Qian, Yevgeniy Vorobeychik, and Milind Tambe. 2016. Using abstractions to solve opportunistic crime security games at scale. In *International Conference on Autonomous Agents and Multiagent Systems*.

[9] Chao Zhang, Arunesh Sinha, and Milind Tambe. 2015. Keeping pace with criminals: Designing patrol allocation against adaptive opportunistic criminals. In *International Conference on Autonomous Agents and Multiagent Systems*. 1351–1359.