An extensive study of Dynamic Bayesian Network for patrol allocation against adaptive opportunistic criminals

Shahrzad Gholami*, **Chao Zhang***, **Arunesh Sinha, Milind Tambe**
University of Southern California, Los Angeles, CA, 90089
{sgholami,zhan661,aruneshs,tambe}@usc.edu

Abstract

Police patrols are used ubiquitously to deter crimes in urban areas. A distinctive feature of urban crimes is that criminals react opportunistically to patrol officers' assignments. Different models of adversary behavior have been proposed but their exact form remains uncertain. Recent work [Zhang *et al.*, 2015] has explored learning the model from real-world criminal activity data. To that end, criminal behavior and the interaction with the patrol officers is represented as parameters of a Dynamic Bayesian Network (DBN), enabling application of standard algorithms such as EM to learn the parameters. More specifically, the EMC^2 algorithm is a sequence of modifications to the DBN representation, that allows for a compact representation resulting in better learning accuracy and increased speed of learning. In this paper, we perform additional experiments showing the efficacy of the EMC^2 algorithm. Furthermore, we explore different variations of Markov model. Unlike DBNs, the Markov models do not have hidden states, which indicate distribution of criminals, and are therefore easier to learn using standard MLE techniques. We compare all the approaches by learning from a real data set of criminal activity obtained from the police department of University of Southern California (USC) situated in Los Angeles, USA. We demonstrate a significant better accuracy of predicting the crime using the EMC^2 algorithm compared to other approaches. This work was done in collaboration with the police department of USC.

1 Introduction

Crime in urban areas plagues every city in all countries. A notable characteristic of urban crime, distinct from organized terrorist attacks, is that most urban crimes are opportunistic in nature, i.e., criminals do not plan their attacks in detail, rather they seek opportunities for committing crime and are agile in their execution of the crime [Zhang *et al.*, 2014; Short *et al.*, 2008]. In order to deter such crimes, police officers conduct patrols with the aim of preventing crime. However, by observing on the spot the actual presence of patrol units, the criminals can adapt their strategy by seeking crime opportunity in less effectively patrolled location. The problem of where and how much to patrol is therefore important.

There are two approaches to solve this problem. The first approach is to determine patrol schedules manually by human planners, which is followed in various police departments including police in USC. However, it has been demonstrated that manual planning of patrols is not only time-consuming but it is also highly ineffective in many related scenarios of protecting airport terminals [Jain *et al.*, 2010] and ships in ports [Shieh *et al.*, 2012]. The second approach is to use automated planners to plan patrols against urban crime. This approach has either focused on modeling the criminal explicitly [Zhang *et al.*, 2014; Short *et al.*, 2008] (rational, bounded rational, limited surveillance, etc.) in a game model or to learn the adversary behavior using machine learning [Chen *et al.*, 2004]. However, the proposed mathematical models of criminal behavior have not been validated with real data. Also, prior machine learning approaches have either only focused on the adversary actions ignoring their adaptation to the defenders' actions [Chen *et al.*, 2004].

In recent work we proposed a novel approach to tackle the problem of generating patrol strategies against opportunistic criminals in [Zhang *et al.*, 2015]: learn the criminal behavior from real data. We did so by modeling the interaction between the criminal and patrol officers as a Dynamic Bayesian Network (DBN). As far as we know, we are the first to use a DBN model that considers the temporal interaction between defender and adversary in the learning phase.

Given a DBN model, we use the well-known Expectation Maximization (EM) algorithm to learn unknown parameters in the DBN from given learning data. However, using EM with the basic DBN model has two drawbacks: (1) the number of unknown parameters scales exponentially with the number of patrol areas and in our case is much larger than the available data itself; this results in over-fitting (2) EM cannot scale up due to the exponential growth of runtime in the number of patrol areas. We demonstrate these two drawbacks both theoretically and empirically.

The second algorithm in [Zhang *et al.*, 2015] utilizes is a sequence of modifications of the initial DBN model resulting in a compact representation of the model, that leads to better learning accuracy and increased speed of learning of the EM algorithm when used for the compact model. This sequence of modifications involve marginalizing states in the DBN using approximation technique from the Boyen-Koller algorithm [Boyen and Koller, 1998] and exploiting structure of this problem. In the compact model, the parameters scale polynomially with the number of patrol areas, and EM applied to this compact model runs in polynomial time.

Finally, we propose different variations of Markov model for prediction of the crime based on available data for crime and/or defender. The main purpose for this modeling is to evaluate other possible bayesian network structures. The Markov models do not have hidden states, which is distribution of criminals.

*S. Gholami and C. Zhang are joint first-authors for this paper

With regard to that, we explored three broad structures: (1) current step crimes depends on previous step crimes (2) current step crimes depends on previous step police patrol and (3) current step crimes depends on previous step crimes and police patrol. Our finding was that the compact DBN model outperformed all the Markov models.

Based on the the algorithms we proposed in [Zhang *et al.*, 2015], we do an extensive experiments to analyze the performance of those algorithms in this paper. We also compare that with the performance of the various Markov models we explored. As part of our collaboration with the police department of USC, we obtained criminal activity and patrol data for three years. Given the results showing the good predictive output with the DBN model, we expect our algorithm to be tested and eventually deployed in USC. The main focus of this paper is on the estimation and prediction of the number of crimes. But as it is addressed in our previous work, these techniques can be used as a basis for planning purposes. More broadly, by introducing a novel framework to reason about urban crimes along with efficient learning and planning algorithms, we open the door to a new set of research challenges.

2 Related Work

We categorize the related work into five main areas. First, recent research has made inroads in applying machine learning and data mining in criminology domain to analyze crime patterns and support police in making decisions. A general framework for crime data mining is introduced in [Chen *et al.*, 2004]. In [Nath, 2006], data mining is used to model crime detection problems and cluster crime patterns; in [De Bruin *et al.*, 2006], data mining approaches are applied in criminal career analysis; in [Oatley *et al.*, 2006], the authors apply machine learning techniques to soft forensic evidence and build decision support systems for police. However, this area of research considers only crime data and does not model the interaction between patrol officers and criminals.

The second line of work we compare is Pursuit-Evasion Games(PEG). PEG models a pursuer(s) attempting to capture an evader, often where their movement is based on a graph[Hespanha *et al.*, 2000]. However, in common settings of Pursuit Evasion Games, evader's goal is to avoid capture and not to seek opportunities to commit crimes and a pursuer's goal is to capture the evader and not to deter the crime; thus common PEG settings are different from the setting in this work.

The third area of work we compare with is Stackelberg Security Games (SSG) [Tambe, 2011], which models the interaction between defender and attacker as a game and recommends patrol strategies for defenders against attackers. SSG has been successfully applied in security domains to generate randomized patrol strategies, e.g., to protect flights [Tambe, 2011], for counter-terrorism and fare evasion checks on trains [Jiang *et al.*, 2013]. While the early work on SSG assumed a perfectly rational attacker, recent work has focused on attackers with bounded rationality and learning the parameters of the bounded rationality model using machine learning methods such as maximum-likelihood estimation. An example of this approach is the PAWS model [Yang *et al.*, 2014]. PAWS addresses the problem of learning criminals' behavior in the domain of wildlife crime such as illegal poaching, within a game-theoretic interaction between defenders and criminals. Recent research has also made progress in designing leader-follower patrol strategies against adversaries in graph settings [Basilico *et al.*, 2009a]. In [Basilico *et al.*, 2009b], patrol strategies against various types of adversaries are designed.

However, including various extensions, security games include an explicit model of the adversary such as bounded rationality models and limited observation models. Distinct from these approaches, we do not model the adversary explicitly, rather we aim to learn the adversary interaction with defender using real world data. In our case these are how the adversary moves from one patrol area to another, and the probability of his committing a crime given some patrol officers presence.

In addition, in SSG randomized strategies are generated assuming attackers learn the strategy through long-term observation. However, in our work, we plan in a more dynamic environment where we keep pace with criminals in a real time fashion. We update their behavior model and change our strategy accordingly from time to time. In this dynamic environment, criminals have little time to observe and exploit current strategy before we switch to another one, which means pure strategy suffices for our purpose. Therefore, we show (Section 6) that the optimal strategy in each small period is a pure one.

A fourth thread of recent research combines machine learning with game theory. In [Blum *et al.*, 2014], the defender's optimal strategy is generated in a SSG by learning the payoffs of potential attackers from their best responses to defender's deployments. An inherent problem with such an approach is that the defender strategy is geared towards learning the adversary payoff, and not exploiting the improved knowledge of the adversary payoff as the game progresses. Adversarial machine learning or adversarial classification [Vorobeychik and Li, 2014] is another technique that uses concepts from game theory to learn a separator in a classification problem in which the adversary that can modify labels arbitrarily. We restrict our attention to adversaries that do not aim to attack the machine learning algorithm itself.

The last area of work we compare against works modeling opportunistic criminals. In [Short *et al.*, 2008] burglars' movement is modeled as a random walk, and in [Zhang *et al.*, 2014], a more general model of opportunistic criminals was propose with algorithms for optimal strategy against such criminals. Again, these works include explicit models of the criminals and lack real world data to learn the interactions.

3 Motivating Example

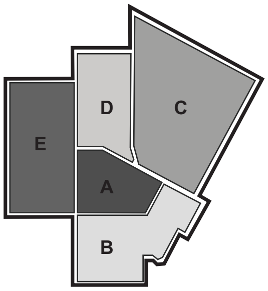

Figure 1: Campus map

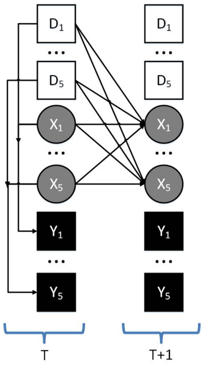

Figure 2: DBN for games

Domain Description: The motivating example for this study is the problem of controlling crime on a university campus. Our case study is about USC in USA. USC has a Department of Public Safety (DPS) that conducts regular patrols, similar to police patrols in urban settings. As part of our collaboration with USC DPS, we have access to the crime report as well as patrol schedule on campus for the last three years (2011-2013). USC is a large enough university that allows us to claim that our methods are applicable to large urban neighborhoods.

In USC, the campus map is divided into five patrol areas, which is shown in Fig 1. DPS patrols in three shifts per day. In the crime data all crimes are local, i.e., no crime happens across two patrol areas or patrol shifts. At the beginning of each patrol shift, DPS assigns each available patrol officer to a patrol area and the officer patrols this area in this shift. At the same time, the criminal is seeking for crime opportunities by deciding which target they want to visit. Discussions with DPS reveals that criminals act opportunistically, i.e., crime is not planned in detail, but occurs when opportunity arise and there is insufficient presence of DPS officers.

There are two separate reports that DPS shared with us. The first is about crime activity that includes details of each reported crime during the last three years, including the type of crime and the location and time information about the crime. In this paper, we do not distinguish between different types of crime and hence we consider only the number of crimes in each patrol area during each shift. Therefore, we summarize the three year crime report into $365 \times 3 \times 3 = 3285$ crime data points, one for each of the 8-hour patrol shift. Each crime data point contains five crime numbers, one for each patrol area.

Shift	A	B	C	D	E
1	1	1	2	2	2
2	1	1	1	2	1
3	2	1	1	3	1

Table 1: Crime data for 3 shifts.

The second data-set contains the DPS patrol allocation schedule. Every officer is allocated to patrolling within one patrol area. We assume that all patrol officers are homogeneous, i.e., each officer has the same effect on criminals' behavior. As a result, when generating a summary of officer patrol allocation data, we record only the number of officers allocated to each patrol area in each shift.

Table 1 shows a sample of the summarized crime data, where the row corresponds to a shift, the columns correspond to a patrol area and the numbers in each cell is the number of crimes. Table 2 shows a sample of the summarized officer patrol allocation data, where the numbers in each cell is the number of patrol officers. For example, from Table 1 we know that in Shift 2, there were 2 crimes in area D and 1 crime in the other areas, while from Table 2, we know that in Shift 2, the number of officers in area A and B is 1 while the number of officers in area C, D and E is 2. However, we do not know the number of criminals in any patrol area in any patrol shift. We call the patrol area as targets, and each patrol shift a time-step.

Problem Statement: Given data such as the real-world data from USC, our goal is to build a general learning and planning framework that can be used to design optimal defender patrol allocations in any comparable urban crime setting.

Shift	A	B	C	D	E
1	2	1	1	1	1
2	1	1	2	2	2
3	2	1	1	3	1

Table 2: Patrol data for 3 shifts.

Due to the lack of the data for criminals, we model the learning problem as a DBN which contains hidden states for criminals. Hidden states are unobserved data that would be estimated. Next section deals with the basic model, the EM algorithm and a compact form of our model that leads to improved learning performance. After that, we present Markov model that considers observed data of crime and defender as the input for modeling and also accuracy of the prediction in different variations of the models are evaluated. The main focus of this study is on prediction and estimation of the number of crime; however, this prediction can be used for planning techniques.

4 Learning with Hidden States

We propose to learn the criminals' behavior, i.e, how the criminals pick targets and how likely are they to commit crime at that target. This behavior is in part affected by the defenders' patrol allocation. Due to the fact that we do not have detailed information about the criminals, we assume that criminals are homogeneous in this paper, i.e., all criminals behave in the same manner. Further, as stated earlier, the patrol officers are also homogeneous. It means that they are given similar tasks and they do their job with the same quality. Thus, crime is affected only by the number of criminals and patrol officers, and not by which criminal or patrol officer is involved.

We propose a DBN model for learning the criminals' behavior. In every time-step of the DBN we capture the following actions: the defender assigns patrol officers to protect N patrol areas and criminals react to the defenders' allocation strategy by committing crimes opportunistically. Across time-steps the criminal can move from any target to any target, since a time-step is long enough to allow such a move. The criminals' payoff is influenced by the attractiveness of targets and the number of officers that are present. These payoffs drive the behavior of the criminals. However, rather than model the payoffs and potential bounded rationality of the criminals, we directly learn the criminal behavior as modeled in the DBN.

The DBN is shown in Fig 2: squares are observed states, where N white squares represent input states (number of defenders at each target) and N black squares represent output states (number of crime at each target) while N circles (number of criminals at each target) are hidden states. For ease of exposition, we use C to denote the largest value that any state can take. Next, we introduce the various parameters of this DBN.

4.1 DBN Representation

First, we introduce parameters that measure the size of the problem

- N: Total number of targets in the graph.
- T: Total time steps of the training data.

Next, we introduce random variables for the observed state (input defender distribution and output crime distribution in our case) and the hidden state. We use three random variables to represent the global state for defenders, criminals and crimes at all targets.

- d_t: Defender's allocation strategy at step t: number of defenders at each target in step t. d_t can take C^N possible values.
- x_t: Criminals' distribution at step t with C^N possible values
- y_t: Crime distribution at step t with C^N possible values.

Next, we introduce the unknown parameters that we wish to learn.

- π: Initial criminal distribution: probability distribution of x_1.
- A: The transition matrix that decides how the criminals distribution evolves from one time step to other. We call this the movement matrix of the DBN. Formally, $A(d_t, x_t, x_{t+1}) = P(x_{t+1}|d_t, x_t)$. Given the C^N values for each argument of A, representing A requires $C^N \times C^N \times C^N$ parameters.
- B: The transition matrix that decides how the criminals commits crime at any target. We call this the crime matrix of the DBN. Formally, $B(d_t, x_t, y_t) = P(y_t|d_t, x_t)$. Given the C^N values for each argument of B, representing B requires $C^N \times C^N \times C^N$ parameters.

We can apply the EM algorithm to learn the unknown initial criminal distribution π, movement matrix A and output matrix B. However, EM applied to the basic DBN model above results in practical problems that we discuss in the next section.

4.2 Expectation Maximization

We start with a brief overview of EM. EM is a class of algorithms for finding maximum likelihood estimation for unknown parameters in DBN [Dempster *et al.*, 1977]. The EM algorithm has an initialization step, expectation (E) step and maximization (M) step. The initialization step chooses initial estimates for unknown parameters (π, A, B). The E step computes some intermediate parameters using these estimates. The M step updates the estimates of π, A, B using values from E step. By iteratively performing E and M step, the EM algorithm converges to a local maxima of the likelihood function for parameters in the DBN. The particular mathematical equations used in E and M depends on the underlying model. For our basic model these equations are not hard to derive and are delegated to the online appendix[1].

In EM algorithm, the size of movement matrix A is $C^N \times C^N \times C^N$ and the size of crime matrix B is also $C^N \times C^N \times C^N$. The number of unknown variables is $O(C^{3N})$. The exponentially many parameters make the model complex, and hence results in over-fitting given limited data. In addition, the time complexity as well as the space complexity of EM depends on the number of parameters, hence the problem scales exponentially with N. In practice, we can reduce C by categorizing the number of defenders, criminals and crimes. For example, we can partition the number of defenders, criminals and crimes into two categories each: the number of officers at each station is 1 (meaning ≤ 1) or 2 (meaning ≥ 2); the number of criminals/crimes is 0 (no criminal/ crime) or 1 (≥ 1 criminal/crime). However, the number of unknown parameters is still exponential in N. As a concrete example, in USC, $N = 5$ and the number of unknown parameters are more than 32768, even when we set $C = 2$. As we have daily data for three years, which is $365 \times 3 \times 3 = 3285$ data points, the number of parameters is much more than the number of data points. Therefore, we aim to reduce the number of parameters to avoid over-fitting and accelerate the computing process.

4.3 Compact model and EMC^2 procedure

In this section, we introduce our second contribution, which is to modify the basic DBN model to reduce the number of parameters. In the resultant compact model, the EM learning process runs faster and avoids over-fitting to the given data. The improvement may be attributed to the well-established machine learning principle of Occam's Razor [Blumer *et al.*, 1987], and our experimental results support our claims. We use three modifications to make our model compact. (1) We infer from the available crime data that crimes are local, i.e., crime at a particular target depends only on the criminals present at that target. Using this inference, we constructed a factored crime matrix B that eliminates parameters that capture non-local crimes. (2) Next, we rely on intuition from the Boyen-Koller [Boyen and Koller, 1998] (BK) algorithm to decompose the joint distribution of criminals over all targets into a product of independent distributions for each target. (3) Finally, our consultations with the DPS in USC and prior literature on criminology [Short *et al.*, 2008] led us to conclude that opportunistic criminals by and large work independently. Using this independence of behavior of each criminal, we reduce the size of the movement matrix. After these steps, the number of parameters is only $O(N \cdot C^3)$.

EM on CompaCt model (EMC^2) procedure applies the EM algorithm to the compact DBN model to find maximum likelihood estimation for unknown parameters.

4.4 Additional Experiments

We conduct experiments in addition to the ones conducted in [Zhang *et al.*, 2015]. The purpose of such experiments is to better analyze the results obtained in that paper.

Experimental setup. We following the settings in [Zhang *et al.*, 2015]. All our experiments were performed on a machine with 2.4GHz and 16GB RAM. MATLAB was our choice of programming language. To avoid leaking confidential information of USC Department of Public Safety, all the crime numbers shown in the results are normalized.

Learning(Settings): In this paper, we do extensive study on evaluating performance of EMC^2 algorithm in learning criminals' behavior. We use the case study of USC in our experiments. We obtained three years of crime report and corresponding patrol schedule followed in USC. We divide the three year data into four equal parts of nine months each. For each part we train on the first eight months data and test on the ninth month data. Since EMC^2 algorithm and EM algorithm only reach locally optimal solution, we run the algorithms for 30 different randomly chosen start points and choose the best solution from among these runs. These start points, i.e., values of A, B and π, are generated by sampling values from a uniform random distribution over $[0, 1]$ for all the elements and then normalizing the probabilities so that they satisfy the initial conditions. C is set to 2 following [Zhang *et al.*, 2015].

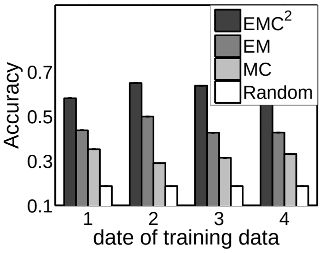

Figure 3: Accuracy (4 datasets)

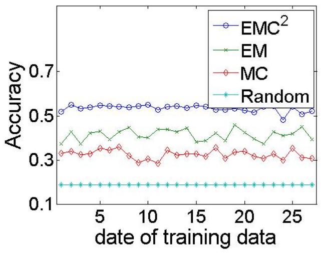

Figure 4: Accuracy (27 datasets)

Result: To begin with, we extend the result of Figure 3 in [Zhang *et al.*, 2015] into Figure 4, which expands the number of datasets. We measure learning performance for each individual target using a metric that we call accuracy. To define this metric, let n_{it} be the actual number of crimes at target i for time step t, let n'_{it} be the predicted number of crimes at target i at time step t. Then, accuracy at step t is the probability of the event $\sum_{i=1}^{N} |n_{it} - n'_{it}| \leq 1$. In other words, it is the probability that we make less than one mistake in predicting crimes for all N targets. The reported accuracy is the average accuracy over all t. In Figure 4, the y-axis represents the accuracy. The higher accuracy is, the more accurate our prediction is. We compare four different algorithm: MC, EM, EMC^2 algorithm and the uniform random algorithm, which sets equal probability for all possible numbers of crimes at each target. As expected, EMC^2 outperforms all other algorithms in all training groups. In addition, even though the accuracy of the algorithms varies in different training groups, which we attribute to the noisy nature of the data in the field, the largest difference is within 15% in all 27 datasets. This indicates accuracy of the algorithms are data-independent.

As introduced in motivating example, DPS conducts three shifts per day and each shift is eight hours by default. How-

[1] `http://keep-pace-with-criminals.weebly.com/`

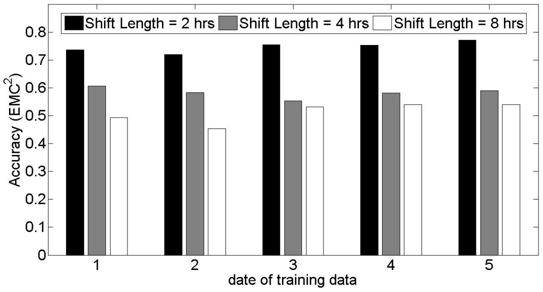

Figure 5: Varying length of shift

ever, the criminals may react in a shorter period. Therefore, we evaluate EMC^2 algorithm by varying the length of the shift in Figure 5. The new length of the shift should be the divisor of the original length. This is because we 'generate' the new strategy by duplicate the original strategy within one shift. In Figure 5, the y-axis shows the accuracy . The accuracy increases as the length of the shift decreases. This is because when the length of the shift decreases, we capture the criminals' adaptive behavior better. Thus, the accuracy increases. This indicates we can improve the performance of the algorithm by decreasing the length of the shift.

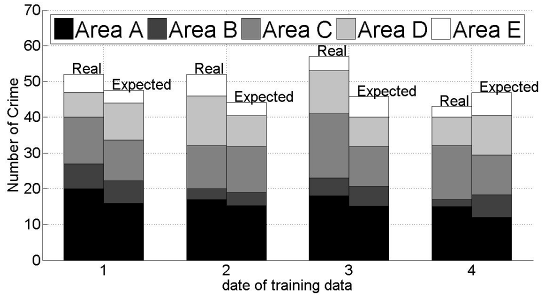

Figure 6: Target by target comparison

In Figure 5 in [Zhang *et al.*, 2015], we compare the expected number of crime that is calculated by EMC^2 with the real number of crime. However, in that experiment, we only consider the total number of crimes over all 5 targets. Even the expectation of the sum is close to the real number of crimes, the expectation for each target may be quite different from the real number of crimes at that target. Therefore, in Figure 6, we compare the expected number of crimes at each target that is calculated by EMC^2 algorithm with the real number of crimes. The y-axis shows the accumulated (expected) number of crimes. As we can see in Figure 6, the expected number of crimes that is calculated by EMC^2 algorithm at each target is closed to the real number of crimes at that target. This indicates that the prediction of EMC^2 algorithm is closed to the reality not by coincidence.

5 Learning from Observed States

There are different techniques to model the relation between the number of crime and patrol allocation and predict number of the crime based on the history. As we discussed earlier, in our Markov model based approach there are no hidden states and we explore three variations of the Bayesian network structure. The main purpose of this section is to study the performance of simpler models that need less computation in comparison with model including hidden states in the previous section. The three cases are when only crime predicts crime, only defender predicts crime and both crime and defender predicts crime. We discuss these in details in the sub-sections below.

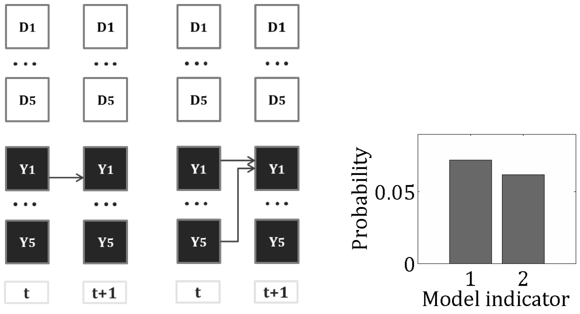

(a) Model 1 (b) Model 2 (c) plot of probability of correct prediction

Figure 7: Model Structure: Crime predicts crime

5.1 Crime predicts crime

Before describing these modifications in details, we introduce some notations that aid in describing the different quantities at each target: $Y_t = [Y_{t,1}, Y_{t,2}, ..., Y_{t,N}]$ is a N by 1 random vector indicating the number of crimes $Y_{t,i}$ at each target i at step t. D_t is a N by 1 random vector indicating the number of defenders $D_{t,i}$ at each target i at step t. X_t is a N by 1 random vector indicating the number of criminals $X_{t,i}$ at each target i at step t.

In the first model shown in Fig 7(a) the effect of the crime in the previous steps at the same target is investigated. This correlation can be defined with the following mathematical function:

$$Y_{t+1,n} = f(Y_{t,n})$$

For this mathematical modeling, probability for a sequence of events can be calculated as below:

$$\begin{aligned}
P(Y_n; A) &= P(Y_{t,n}, ..., Y_{0,n}; A) \\
&= P(Y_{t,n}|Y_{t-1,n}, ..., Y_{0,n}; A) \times ... \times P(Y_{1,n}|Y_{0,n}; A) \\
&= P(Y_{t,n}|Y_{t-1,n}; A)...P(Y_{1,n}|Y_{0,n}; A) \\
&= P(Y_{t,n}|Y_{t-1,n}; A)...P(Y_{1,n}|Y_{0,n}; A) \\
&= \prod_{1 \leq t \leq T} P(Y_{t,n}|Y_{t-1,n}; A)
\end{aligned}$$

Log likelihood for this model can be written as following:

$$\begin{aligned}
l(A) = logP(Y_n; A) &= log \prod_{1 \leq t \leq T} P(Y_{t,n}|Y_{t-1,n}; A) \\
&= log \prod_{1 \leq t \leq T} A_{Y_{t,n}Y_{t-1,n}} \\
&= \sum_{1 \leq t \leq T} logA_{Y_{t,n}Y_{t-1,n}} \\
&= \sum_{i=1}^{|S_Y|} \sum_{j=1}^{|S_Y|} \sum_{t=1}^{T} 1\{Y_{t,n} = S_i \wedge Y_{t-1,n} = S_j\} logA_{ij}
\end{aligned}$$

In the above equation $|S_Y|$ indicates the total possible number

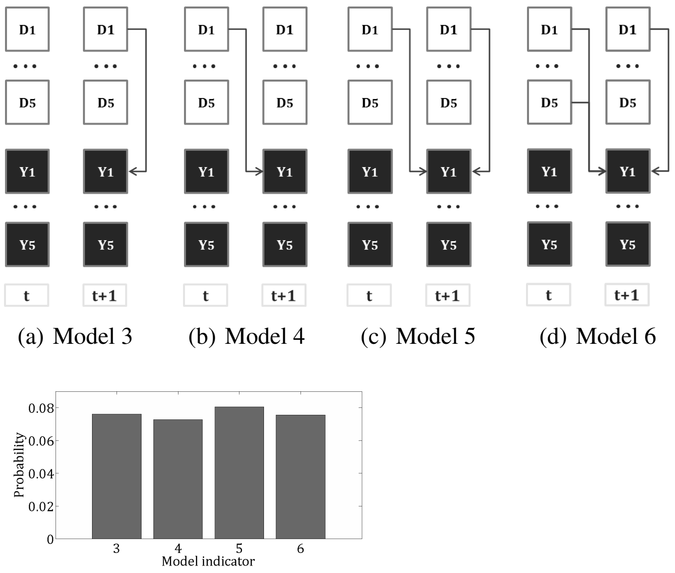

(a) Model 3 (b) Model 4 (c) Model 5 (d) Model 6

(e) plot of prob. of correct prediction

Figure 8: Model Structure: Defender allocation predicts crime

of values that Y can take.
Also the expression shown as $1\{Y_{t,n} = S_i \wedge Y_{t-1,n} = S_j\}$ would be equal to one when the internal expression is satisfied. To find the parameters of the transition matrix we need to maximize the log likelihood defined in the previous step.

$$\max_{A} l(A) \qquad s.t. \begin{cases} \sum_{i=1}^{|S_Y|} A_{i,j} = 1, & j = 1...|S_Y| \\ A_{ij} \geq 0, & i,j = 1...|S_Y| \end{cases}$$

This optimization problem can be solved in the closed form using Lagrangian multipliers. So the the parameters of the system can be found from the following equation:

$$\hat{A_{ij}} = \frac{\sum_{t=1}^{T} 1\{Y_t = S_i \wedge Y_{t-1} = S_j\}}{\sum_{t=1}^{T} 1\{Y_{t-1} = S_j\}}$$

So for each target we can find the transition matrix like the above formulation. The same procedure for deriving the transition matrix can be done for all other models in this section and also the other more complicated variation for modeling of this problem which consist crime information from other targets. Model shown in 7(b) includes number of the crime at all other targets.

$$Y_{t+1,n} = f(Y_{t,1:n})$$

To compare the accuracy of these two models with each other, we learned from two month of data and then found the average probability of predicting the number of the crime over all the targets correctly for one week. Fig 7(c) shows that the first model is performing better and the number of the crime in an area is much more related to the number of the crime at that specific area.

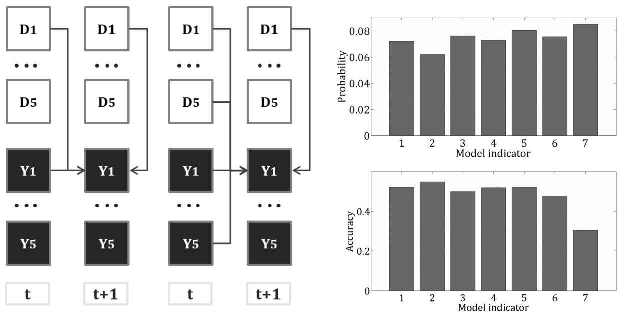

(a) Model 7 (b) Model 8 (c) plot of prob. of correct prediction and one mistake in prediction

Figure 9: Model Structure: Defender allocation predicts crime

5.2 Defender allocation predicts crime

To study the effect of the defender allocation on the number of the crime, the first model in Fig 8(a) is presented which can be defined by the following mathematical function.

$$Y_{t+1,n} = f(D_{t+1,n})$$

In Fig 8(b), we present the relation between the defender strategy at each target in previous time step and crime at that specific area to evaluate the effect of the defender strategy on the prediction.

$$Y_{t+1,n} = f(D_{t,n})$$

Additionally, the effect of the defender allocation in the current time step can also be included, this model is shown in Fig 8(c).

$$Y_{t+1,n} = f(D_{t,n}, D_{t+1,n})$$

Another variation is shown in Fig 8(d) which includes number of the defenders at all other targets from previous time and the number of the defender from the current time.

$$Y_{t+1,n} = f(D_{t,1:n}, D_{t+1,n})$$

The same procedure as the previous subsection can be used to find the transition matrix for the above models. Fig8(e) shows the average probability of correct prediction over all targets for all of the models described in this subsection. This figure demonstrates that number of the defender at the current step has more effect on the prediction. Also, adding information from previous step helps to increase the accuracy for exact prediction. However, in comparison of the model 6 and 5, it is observed that model 5 outperforms model 6. Similar to the previous subsection we can see that crime prediction at each target is mostly affected by the information from that target rather that all other targets.

5.3 Crime and Defender allocation predicts crime

One more variation to the above model is including the observed state of the crime at that specific target in addition to the number of defender at the previous and current step. This model is shown in Fig 9(a).

$$Y_{t+1,n} = f(D_{t,n}, D_{t+1,n}, Y_{t,n})$$

In order to compare the accuracy of the exact prediction in model 7, Fig9(c) (upper plot) illustrates the probability for all the models. From this figure, it can be concluded that model 7 outperforms all other models described in this section and has the best prediction potential. It is worth noting that based on the plot of probabilities for exact prediction, the most complicated and comprehensive model is one that considers all of the observed states for defender and crime from the previous time at all targets in addition to the number of defender in the current time and the specific target. However due to the large number of parameters for this model and overfitting issues, we ignore this item in our study. This model is illustrated in Fig9(b). Fig9(c) (lower plot) presents the accuracy based on the probability of one mistake in the prediction, according to the metric used in the previous section. From this plot and results of previous sections, it can be concluded that in all of the cases, DBN with hidden states outperforms Markov models.

6 Conclusion

This paper evaluates a novel framework to design patrol allocation against adaptive opportunistic criminals. Such framework models the interaction between officers and adaptive opportunistic criminals as a DBN. Next, a sequence of modifications to the basic DBN resulting in a compact model that enables better learning accuracy is proposed. Finally, by investigating different Markov models, we can conclude that not only the number of the crime but also, number of the defender in each area can affect the prediction of the crime. Including both information increase the accuracy of the prediction. Further, considering hidden states as done in the DBN improves prediction accuracy. By extensive experimental validation with real data, we show that our choice of model and assumptions is supported. Further, our modeling assumptions were informed by inputs from our collaborators in the DPS at USC. These promising results has opened up the possibility of deploying our method for policing in USC.

7 Acknowledgement

This research is supported by MURI grant W911NF-11-1-0332.

References

[Basilico *et al.*, 2009a] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 57–64. International Foundation for Autonomous Agents and Multiagent Systems, 2009.

[Basilico *et al.*, 2009b] Nicola Basilico, Nicola Gatti, Thomas Rossi, Sofia Ceppi, and Francesco Amigoni. Extending algorithms for mobile robot patrolling in the presence of adversaries to more realistic settings. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 557–564. IEEE Computer Society, 2009.

[Blum *et al.*, 2014] Avrim Blum, Nika Haghtalab, and Ariel D Procaccia. Learning optimal commitment to overcome insecurity. In *Proceedings of the 28th Annual Conference on Neural Information Processing Systems (NIPS). Forthcoming*, 2014.

[Blumer *et al.*, 1987] Alselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Occam's razor. *Inf. Process. Lett.*, 24(6):377–380, April 1987.

[Boyen and Koller, 1998] Xavier Boyen and Daphne Koller. Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 33–42. Morgan Kaufmann Publishers Inc., 1998.

[Chen *et al.*, 2004] Hsinchun Chen, Wingyan Chung, Jennifer Jie Xu, Gang Wang, Yi Qin, and Michael Chau. Crime data mining: a general framework and some examples. *Computer*, 37(4):50–56, 2004.

[De Bruin *et al.*, 2006] Jeroen S De Bruin, Tim K Cocx, Walter A Kosters, Jeroen FJ Laros, and Joost N Kok. Data mining approaches to criminal career analysis. In *Data Mining, 2006. ICDM'06. Sixth International Conference on*, pages 171–177. IEEE, 2006.

[Dempster *et al.*, 1977] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.

[Hespanha *et al.*, 2000] Joao P Hespanha, Maria Prandini, and Shankar Sastry. Probabilistic pursuit-evasion games: A one-step nash approach. In *Decision and Control, 2000. Proceedings of the 39th IEEE Conference on*, volume 3, pages 2272–2277. IEEE, 2000.

[Jain *et al.*, 2010] Manish Jain, Jason Tsai, James Pita, Christopher Kiekintveld, Shyamsunder Rathi, Milind Tambe, and Fernando Ordóñez. Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40(4):267–290, 2010.

[Jiang *et al.*, 2013] Albert Xin Jiang, Zhengyu Yin, Chao Zhang, Milind Tambe, and Sarit Kraus. Game-theoretic randomization for security patrolling with dynamic execution uncertainty. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 207–214. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

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

[Oatley *et al.*, 2006] Giles Oatley, Brian Ewart, and John Zeleznikow. Decision support systems for police: Lessons from the application of data mining techniques to soft forensic evidence. *Artificial Intelligence and Law*, 14(1-2):35–100, 2006.

[Shieh *et al.*, 2012] Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. Protect: A deployed game theoretic system to protect the ports of the united states. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 13–20. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

[Short *et al.*, 2008] Martin B Short, Maria R D'ORSOGNA, Virginia B Pasour, George E Tita, Paul J Brantingham, Andrea L Bertozzi, and Lincoln B Chayes. A statistical model of criminal behavior. *Mathematical Models and Methods in Applied Sciences*, 18(supp01):1249–1267, 2008.

[Tambe, 2011] Milind Tambe. *Security and game theory: algorithms, deployed systems, lessons learned.* Cambridge University Press, 2011.

[Vorobeychik and Li, 2014] Yevgeniy Vorobeychik and Bo Li. Optimal randomized classification in adversarial settings. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 485–492. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[Yang *et al.*, 2014] Rong Yang, Benjamin Ford, Milind Tambe, and Andrew Lemieux. Adaptive resource allocation for wildlife protection against illegal poachers. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 453–460. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[Zhang *et al.*, 2014] Chao Zhang, Albert Xin Jiang, Martin B Short, P Jeffrey Brantingham, and Milind Tambe. Defending against opportunistic criminals: New game-theoretic frameworks and algorithms. In *Decision and Game Theory for Security*, pages 3–22. Springer, 2014.

[Zhang *et al.*, 2015] Chao Zhang, Arunesh Sinha, and Milind Tambe. Keeping pace with criminals: Designing patrol allocation against adaptive opportunistic criminals(to appear). In *Proceedings of the 2015 international conference on Autonomous agents and multi-agent systems*. 2015.