Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling (ICAPS 2022)

OFFICERS: Operational Framework for Intelligent Crime-and-Emergency Response Scheduling

Jonathan Chase, Siong Thye Goh, Phong Tran, Hoong Chuin Lau*

Singapore Management University
{jdchase, stgoh, bptran, hclau}@smu.edu.sg

Abstract

In the quest to achieve better response times in dense urban environments, law enforcement agencies are seeking AI-driven planning systems to inform their patrol strategies. In this paper, we present a framework, OFFICERS, for deployment planning that learns from historical data to generate deployment schedules on a daily basis. We accurately predict incidents using ST-ResNet, a deep learning technique that captures wide-ranging spatio-temporal dependencies, and solve a large-scale optimization problem to schedule deployment, significantly improving its scalability through a simulated annealing solver. Methodologically, our approach outperforms our previous works where prediction was done using Generative Adversarial Networks, and optimization was performed with the CPLEX solver. Furthermore, we show that our proposed framework is designed to be readily transferable between use cases, handling a wide range of both criminal and non-criminal incidents, with the use of deep learning and a general-purpose efficient solver, reducing dependence on context-specific details. We demonstrate the value of our approach on a police patrol case study, and discuss both the ethical considerations, and operational requirements, for deployment of a lightweight and responsive planning system.

Introduction

Faced with dense urban environments and budgetary constraints, modern law enforcement is turning to computer science for efficient patrol strategies. Combinatorial optimization problems can be tailored to the requirements of any given agency or police force, balancing incident response times, manpower utilization rates, and police visibility. In the rush to develop patrol strategies that offer excellent performance, however, the human element must not be forgotten. Police officers, with years of established practice behind them, can be resistant to overly radical modifications to their working patterns and shift structures, and intangible factors such as community relations and local knowledge can be lost in overly-prescriptive and dynamic deployment plans. Bringing together ideas from prior work that provided law enforcement deployment planning for a specific use case, we are proposing a framework called OFFICERS (Operational Framework For Intelligent Crime-and-Emergency Response Scheduling) to provide a generalized pattern for tackling law enforcement patrol planning, and discuss the steps required to achieve a deployed system and the ethical ramifications of predictive law enforcement. Additionally we present solution methods for two key components: the incident predictor and the manpower scheduler, and generate experimental results using police patrol as a use case.

*Corresponding Author

Related Work

Incident prediction is a vital prerequisite to effective deployment planning, as the knowledge of where emergencies occur will inform the positioning of response agents. Examples of prediction approaches used in the past include statistical aggregation of historical demands (Malleson and Andresen 2015), spatial hotspot identification with temporal regression (Butt et al. 2021), risk terrain modeling (Caplan, Kennedy, and Miller 2011), continuous time modelling of burglary (Mukhopadhyay et al. 2016), Gaussian Process incident generation for discrete space-time intervals (Chase et al. 2019), deep learning (Chase et al. 2021), (Wang et al. 2017), and network analytics (Dash, Safro, and Srinivasamurthy 2018). Where many of these methods fall short is in capturing the spatio-temporal dependencies in crime data, where incidents appearing in one location or time period implies they will not also appear in the immediately adjacent locations or periods. In this paper we employ ST-ResNet, a method designed to predict urban crowd flows (Zhang, Zheng, and Qi 2017), that captures spatial and temporal dependencies combined with external factors to learn longer cyclic patterns. A similar notion was considered in (Wang et al. 2017), but we apply the model to a wider range of emergency incident types and employ a novel performance metric that reflects the role of the prediction within our framework. In so doing, we show that this approach can be applied successfully to a range of use cases within law enforcement and emergency response.

Much has been done on the problem of optimizing emergency response, including law enforcement deployment (e.g. (Lau, Yuan, and Gunawan 2016), (Mukhopadhyay et al. 2016), (Saisubramanian, Varakantham, and Lau 2015)). The scheduling problem introduced in this paper draws on the concept of police 'presence', an increasingly popular objective for deployment planning (Wang et al. 2021). We apply this concept through scheduling with a patrol diversity con-

straint. However, there are scalability challenges for Mixed Integer Programs (MIPs) that must optimize at the scale of urban law enforcement on a daily basis. A relatively new and exciting contender in the field of heuristic search is quantum annealing, where combinatorial problems are formulated as Ising, or equivalently QUBO (Quadratic Unconstrained Binary Optimization), models and solved by an annealing machine, such as D-Wave (Date et al. 2019). It is therefore necessary to efficiently convert a constrained MIP to an unconstrained problem, such as in (Crispin and Syrichas 2013) for quantum vehicle scheduling and (Ikeda, Nakamura, and Humble 2019) which applies D-Wave's quantum annealer to nurse scheduling, but can only solve small problems. Due to the scale limitations of quantum devices, hybrid methods or decomposition have been proposed (Ajagekar, Humble, and You 2020), (Stollenwerk, Lobe, and Jung 2019). Quantum computing is still in its nascent stage, so solvers such as Fujitsu's Digital Annealer (DA) chip (Aramon et al. 2019) have been developed to solve QUBO models on conventional CMOS computers, implementing a meta-heuristic inspired by quantum annealing. While these solvers still have model size limitations due to hardware constraints, DA's fully connected 8192 binary variables (in contrast with D-Wave's 2048 qubits on a Chimera graph) is capable of modeling and solving larger scale problems, and without the intermediate step of embedding the QUBO model onto the underlying machine architecture. In this paper, we employ DA to implement our model so as to support meaningful problem sizes that could be deployed using current technology.

One other alternative approach for security deployment is game theory (security games), particularly the application of Stackelberg games. Stackelberg games can be applied effectively to many problems in security, as there are two players, a leader and a follower, who operate as adversaries. The leader chooses strategy first, and can represent the security force, while the follower responds to their choice, representing the various criminal elements who attempt to circumvent the security provision. Uncertainty can be incorporated through the use of mixed strategies, with players choosing actions according to expected payoff as determined by a probability distribution. Stackelberg security games (SSG) can be broadly divided into three categories: infrastructure security games, green security games, and opportunistic crime security games (An, Tambe, and Sinha 2017). The first category are focused on protecting infrastructure, such as ARMOR deployed at LAX (Jain et al. 2010b). These games end after a single attack with the defender only updating strategies after attacks, due to the rarity but seriousness of attacks. The attacker is assumed to employ detailed surveillance. The second category, of green games, focus on protecting vulnerable ecology, such as PAWS for anti-poaching operations (Fang, Stone, and Tambe 2015). Attacks are repeatedly executed, with less surveillance and planning than for infrastructure attacks. The aim is to cover a large sparsely populated geographic area, and machine learning can be incorporated to help with strategy creation. Finally, opportunistic crime games aim to tackle crimes with a low level of planning but repeated operation. A typical example would be pickpockets on public transport (Della Fave 2014). Note that game models hinge on the ability to design an accurate payoff matrix over a finite discrete set of actions, but such a setup is not possible nor suitable for our problem in this paper, which is about scheduling patrol units across sectors in space and time under various complex constraints.

Necessity of the Framework

One of the primary challenges faced by the law enforcement agency that forms the basis of our case study is the heterogeneity of the incidents that must be responded to. Much prior work in this field focuses on accurately predicting specific crime types or running interference against motivated attacks that seek to circumvent the defence. This is particularly true of SSG methods, which assume a degree of concerted rationality on the part of the follower, such that with varying degrees of surveillance and prior planning, they will attempt to adapt to the leader's strategy. In our use case, with incidents reported by the public via emergency call, the type of incident can range from trivial noise complaints, through to medical emergencies, up to assault, fighting, and even in rare cases, bomb threats. With this diversity of incidents, the Stackelberg game assumption of a rational intent behind the attackers' actions is less applicable, as only a subset of the incident instigators will respond to police behaviour, which we address by adding a 'police presence' constraint to our scheduler. Our problem domain also considers heterogeneous urban terrain, incorporating residential, industrial, and commercial areas on a daily basis. Therefore, we are not attempting to protect a particular set of targets, but to robustly cover a wide area in a varied urban environment. To handle the heterogeneity and uncertainty in the spatio-temporal incident occurrence, we incorporate a machine learning-based incident generator to drive the input of the model, minimizing the burden on the user that can occur in infrastructure-based systems (Jain et al. 2010b).

Scalability is a significant challenge in urban security. Deployed methods such as ARMOR and TRUSTS (Luber et al. 2013) have handled games with a similar number of agents and locations to those used in our experimental results, but for larger scale implementations, such as IRIS, a branch-and-price algorithm, ASPEN, was required (Jain et al. 2010a). Given that our framework must handle daily execution of deployments at the city level, we have designed a three-stage structure, to decompose the shift scheduling process. The first stage generates incident samples for training the deployment, capturing the range of incident possibilities. The second stage, as presented in previous work (Chase et al. 2021), finds the number of police cars required to patrol a set of locations throughout a shift with a time granularity of 2 hours. This demand is found by optimizing the possible car deployments against generated incidents, which means that the occurrence of an incident is not disruptive to the patrol schedule, but is built into the schedule creation, avoiding the need for real-time updates or plan re-creation delivered by mobile app to patrol agents. The final step, for which we present a scheduling model and scalability approach in this paper, assigns the patrol cars to the determined demand. We present a QUBO-based reformulation of the optimization problem that can be solved more

efficiently on a general purpose QUBO solver, than on the industry standard, branch-and-bound based CPLEX solver. Demonstrating the efficiency of a QUBO approach suitable for execution on a general purpose solver, we can be confident that our approach will scale to larger problems while meeting the requirement for daily execution.

OFFICERS Overview

The ability to respond rapidly to dynamically occurring incidents is a key requirement for law enforcement and emergency medical service agents of all stripes. A planner seeking to guarantee this requirement can be met must know two things: when and where incidents occur, and how agents should be stationed/deployed. We propose a general solution framework, and as a case study, consider specifically the problem of deploying police cars to patrol sectors. We take a proactive planning approach in determining robust schedules based on historical data. Our approach can also be applied to station ambulances and firefighting apparatus for Emergency Medical Services (EMS), as the framework accounts for a variety of incident types without assuming an antagonistic motive. We consider a scenario in which response agents are assigned to patrol deployment sectors so as to effectively respond to incidents. Sector boundaries are static and define a predetermined geographic area. Agents patrol their assigned sectors until an emergency call is made and a central dispatcher assigns the agent to respond.

Our framework, outlined along with the problem scenario in Fig. 1, requires three components, the first of which is the prediction of incident occurrence, which we present in this paper. Incident prediction methods were presented in earlier works, but we demonstrate that this method can achieve superior accuracy in precision and recall. Additionally, the method presented in this paper is the most general, identifying the spatio-temporal connections between incidents without relying on explicitly provided terrain data. The second is an optimization on the incident generation output to determine the number of agents required in each sector. For the purposes of deployment, we employ the method presented in (Chase et al. 2021) for the second component. This component is also a refinement on previous designs, offering improved scalability. The model regards incident types agnostically, considering only the priority, duration, and number of agents required, enabling a range of incidents to be considered on an equal basis. This component identifies the number of agents required to patrol each sector, ensuring that sufficient agents are able to satisfy the response times of incidents even while earlier incidents are being addressed. At this stage of the deployment planning process, other factors such as pre-allocation can also be considered, to allow the injection of human insight to provide coverage for factors such as special events that can't be predicted from historical data. The third, and critical, component is the scheduler, an optimization problem that determines which agents should patrol a sector in each time period of a shift. The goal of the scheduler is to control the amount of movement between locations required by agents, limiting the disruption to their patrolling and response activities. We separate the scheduler from the second component so that the disruption minimization decision does not compromise incident response, which is of greater priority. We present a model for scheduling and significantly enhance its scalability through a quantum-inspired simulated annealing method. The solution method is general purpose, and therefore the specific constraints of the scheduling model can be tweaked to fit individual application contexts (as shown by the introduction of a 'police presence' provision) without losing the benefits of the solver. The output of the scheduler provides the deployment plan for supervisors to put into action.

Deep Learning for Incident Prediction

The capacity to accurately predict the patterns of incident occurrence is essential to a law enforcement deployment system. Inaccurate prediction will result in a 'garbage in, garbage out' situation. Since we perform deployment to discrete patrol sectors, an initial attempt to predict incidents might discretize the prediction space accordingly. However, the probability of an incident occurring in one sector in a particular time period is not independent of the sectors or time periods around it, thus our prediction method must consider both spatial and temporal dependencies, as well as external factors such as public holidays. To capture these dependencies we adopt a deep-learning method called ST-ResNet (Zhang, Zheng, and Qi 2017) which tackled crowd flow prediction in large urban environments, which makes it suitable for adaptation to our context. The data input for ST-ResNet is a set of time-discretized images showing the spatial distribution of data points in that time interval. The ST-ResNet architecture, illustrated in Fig. 2, combines four major components. The first three are similarly structured with alternating convolutional neural networks and residual units, which allows deep neutral networks without compromising training effectiveness, and therefore large city problems can be considered. The three components model three temporal properties: closeness, period, and trend, which represent increasingly large time granularities to capture temporal dependencies over both the short and long term. The output of the first three components are fused by weighted linear matrix combination with learnable weights. The fourth component models external factors such as weather, weekends, and public holidays using a fully connected shallow network. The external data component output is combined with the fused components' output by matrix addition followed by application of a hyperbolic tangent, *tanh*, to normalize the combined output. This normalized output forms the model's prediction of the next 1 hour of data. To measure the accuracy of the training we reserve the final 1 hour's worth of input data for testing, with the goal of training being to minimize the mean squared error (MSE) between the predicted data and the test data.

Applying ST-ResNet to Law Enforcement

In our use case we start with a set of emergency call records that record the time and lat-long coordinates of incidents. To train ST-ResNet we divide the geographic map into a grid corresponding to a set of pixels according to a chosen *resolution* parameter, and each incident is represented by a pixel

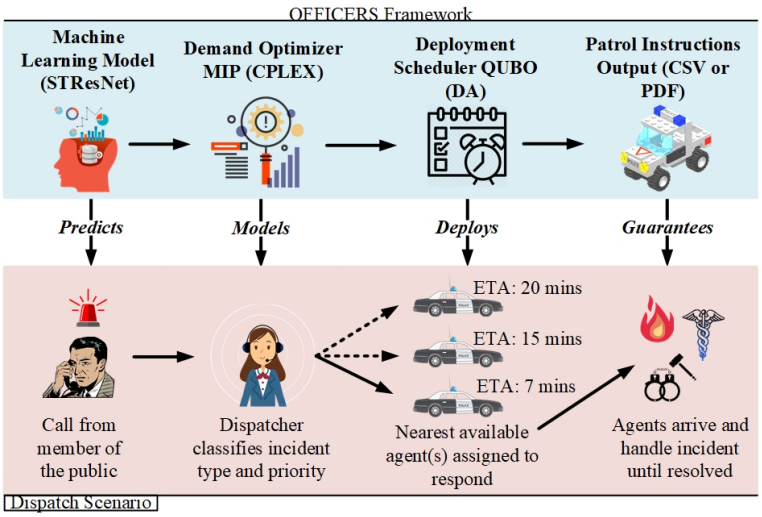

Figure 1: Components of the OFFICERS Framework and how they address the challenges of our problem scenario.

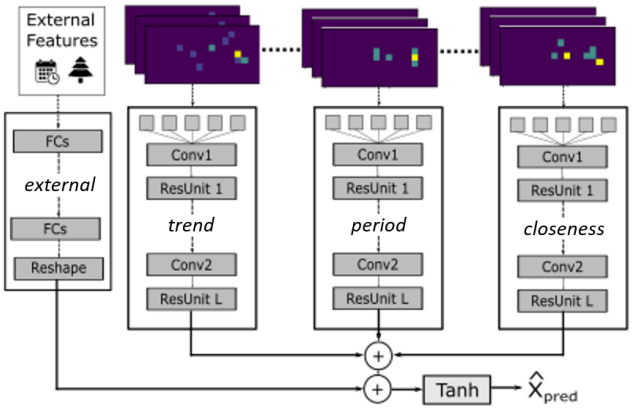

Figure 2: The four components of the ST-ResNet predictor.

indicating the grid location. Initially, we generated an image for each incident and sorted them by time for input to ST-ResNet. However, the data we use is too sparse to derive good dependencies using this conversion method. Therefore, we also choose a *sampling rate* that defines a time frame for each image, allowing multiple incidents to be included in a single image, resulting in a richer heatmap that can identify hotspot areas and the dependencies around them. The results in Table 1 illustrate the importance of well-chosen resolution and sampling rate parameters (see next section for description of F1 metric calculation). The data is sequenced according to the three levels of temporal granularity. The *closeness* data has a length of 24 and sampling rate of 1 - meaning images snapshot incidents hourly for the previous 24 hours. The *period* data has a sampling rate of 24, meaning images snapshot incidents daily, while the *trend* data has a sampling rate of 168 with each image snapshotting a week's worth of incidents. This data sequence is used to predict the next hour's incidents, with a sliding window

Distance (m)	Time (mins)	Resolution 0.01	0.015	0.02
1000	10	0.124	0.143	0.168
2000	12	0.202	0.278	0.313
2000	60	0.375	0.674	0.750
Predicted-Actual Counts		66	19	-2

Table 1: F1 Scores (larger is better) for varying success criteria and resolutions

approach using a one year data set enabling us to predict the following year's incidents hour by hour. External factors consist of one-hot encoded data indicating the various public holidays, a weekday/weekend indicator, and six cyclical features identifying the sine and cosine values for the hour-to-be-predicted's point in the year, week, and day.

A Metric for Assessing Performance

We propose a novel matching metric to compare the performance of the model at different tolerances. The output of the ST-ResNet predictor is an image showing a heatmap for the hour and date to be predicted. Each pixel of the output image represents the number of incidents that occur in the spatial grid space in that hour. To generate incidents, then, for each pixel we generate a number of incidents based on its colour intensity. Each incident is assigned a lat-long value randomized within the geographic area represented by the grid cell, and assigned a start time based on the time division (in our problem case, we use 1 hour steps), with the precise start minute randomly chosen within the hour. This element of randomization allows us to generate multiple incident scenarios from the model. We then compare the incidents to a set of corresponding real test data, and run a matching algorithm that pairs generated and test incidents based on their spatio-temporal distance, as illustrated in Fig. 3. We mark a prediction as a success if it can be paired with a test incident

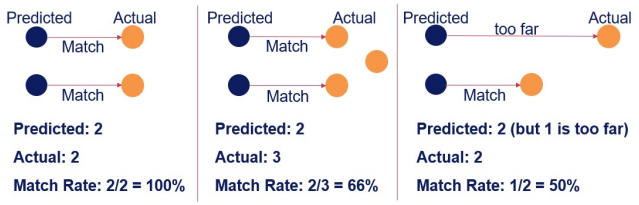

Figure 3: Illustration of the matching metric used to assess model prediction. Matching between two incidents is based on satisfying both spatial and temporal distance requirements.

within the time and distance threshold (e.g. incidents occur within 2km of each other with start times no more than 60 minutes apart). We define *precision* as the proportion of generated incidents that can be matched to a test incident, and *recall* as the proportion of test incidents that can be matched to a generated incident. If the number of generated incidents is too high, the excess incidents will be marked as failures on the precision score as they cannot be matched to a test incident, while if the number is too low, the excess test incidents will be marked as failures on the recall score. However, since it would be possible to achieve high recall by a surfeit of incidents, and a high precision by generating a deficit, we adopt the F1 score (Chinchor 1992), given in Eq. (1), as our metric to balance precision and recall. Both the accuracy of the prediction positions and the predicted counts are important, and therefore we consider both metrics equally, rather than prioritizing one or the other.

$$F_1 = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall} \tag{1}$$

Evaluating ST-ResNet for Law Enforcement

Precisely predicting the time and location of general crimes is difficult due to randomness in incident occurrence. However, our prediction functions as a component within our framework, and if it can accurately identify the time and position of incidents to the granularity of our schedule (1-hour, within a sector), we can be confident that agents can respond in time. Generation of multiple prediction sets further alleviates this issue by using Sample Average Approximation (SAA) in the second component to be robust to random variation. The sector boundaries in our use case are based on historically defined boundaries but are sufficiently small to guarantee intra-sector response times for available cars. Predicting incidents directly according to sector limits can introduce significant bias if boundaries cut across the distribution space(Chase et al. 2021), and from an ethical perspective, care must also be taken when applying this framework to other contexts that sectors are not defined by underlying human bias about the nature of criminal behaviour.

We train our model on a year's data and test over the following half year. The resolution and sampling rate hyperparameters are selected by identifying the smallest value that gives a strong performance at the required F1 distance and time settings. The results in Table 1 indicate that a resolution of 0.02 gives the best result without having a geographic granularity larger than the patrol sectors the agents patrol. Training with a sampling rate of 1 hour also provides the best balance of results, as we found that shortening the sampling rate to 30 minutes resulted in a model built on very sparse data that significantly underestimated the incident counts. 1 hour provides a reasonable time window as we will schedule patrols on an hourly basis. To evaluate the value of ST-ResNet against other methods, we compare it to a prediction approach based on Generative Adversarial Networks (GAN) that was itself proven superior to a Gaussian Process-based approach in (Chase et al. 2021). Since the GAN model also predicts the lat-long coordinates of incidents, we can apply the F1 metric to directly compare on an equal footing with the ST-ResNet model, as shown in Table 2. We find that ST-ResNet is able to attain a much higher F1 score, which is likely due to the rich spatio-temporal dependencies that can compensate for the minor random variations in the data.

Distance (m)	Time (mins)	Method	
		ST-ResNet	GAN
1000	10	0.17	0.15
2000	10	0.31	0.40
2000	30	0.60	0.43
2000	60	0.75	0.43

Table 2: F1 Scores for ST-ResNet vs GAN Models

Given the results shown, we conclude that ST-ResNet is a strong candidate for crime and emergency incident prediction with a high degree of accuracy, due to its capacity to model incident occurrence with spatio-temporal dependencies. The dependency modelling is a key feature, as the patrol sectors create artificial boundaries that do not reflect independence in the incident distribution. Therefore, it is important to treat the underlying space as interconnected so as to provide accurate information to the patrol scheduler.

Deployment Scheduler

Given a set of response agents and a set of sectors, the scheduler allocates each agent to patrol a sector for a discrete time period (e.g. 1 hour for our use case), with allocations permitted to change throughout the shift. The allocation must satisfy the individual sectors' agent demand with the goal of minimizing the total migration time across the shift. Migration time is the time taken for an agent to move from one sector to the next at the end of each time period. Patrol behaviour within each sector is handled autonomously by the agents, allowing them to draw on their local knowledge and training in a range of crime prevention activities.

For the police patrol use case, given n patrol cars to assign, m deployment sectors, and T number of time steps, our goal is to minimize the total traveling time between sectors at the end of each time step. At each time step, we are given the demands that are required for each location by the second framework component, and we require that the total demand across all locations is equal to the number of cars, n. We assume that the travel time to move between each lo-

Notation	Represents
n	Number of cars
m	Number of sectors
T	Number of time steps to consider
i	Car index
j	Location index
t	Time index, member of set $\{1, \ldots, T\}$
$y_{i,j}^t$	Variable: Indicates assignment of car i to location j in time t
x_{i,j_1,j_2}^{t-1}	Variable: indicates i moves from location j_1 to j_2 from time $t-1$ to time t
T_{j_1,j_2}	Parameter: Time to travel from location j_1 to j_2
d_j^t	Parameter: number of cars required in location j in time t.

Table 3: Key notations used in the Scheduling model

cation is also known. To evaluate the travel time we define a variable, x_{i,j_1,j_2}^t, which takes value 1 when y_{i,j_1}^{t-1} and y_{i,j_2}^t both take value 1, indicating that the car i goes from location j_1 to location j_2 at the end of time $t-1$. This can be defined using the following conditional expression:

$$x_{i,j_1,j_2}^{t-1} = \begin{cases} 1, & \text{if } y_{i,j_1}^{t-1}, y_{i,j_2}^t = 1 \\ 0, & \text{otherwise} \end{cases} \quad \forall i, j_1, j_2, t > 0 \quad (2)$$

The notation used in the model is given in Table **??**.

Classical Optimization Formulation

The classical constrained optimization model is as follows:

$$\min \sum_{i=1}^{n} \sum_{j_1=1}^{m} \sum_{j_2 \neq j_1}^{m} \sum_{t=1}^{T} T_{j_1,j_2} x_{i,j_1,j_2}^{t-1} \quad (3)$$

s.t.

$$\sum_{i=1}^{n} y_{i,j}^t = d_j^t, \ \forall j \in \{1, \ldots, m\} \quad \forall t \in \{1, \ldots, T\}, \quad (4)$$

$$\sum_{j=1}^{m} y_{i,j}^t = 1, \ \forall i \in \{1, \ldots, n\} \quad \forall t \in \{1, \ldots, T\}, \quad (5)$$

$$y_{i,j_1}^{t-1} + y_{i,j_2}^t \leq 1 + 2x_{i,j_1,j_2}^{t-1} \quad \forall i, j_1, j_2, t > 0, \quad (6)$$

$$y_{i,j_1}^{t-1} + y_{i,j_2}^t \geq 2x_{i,j_1,j_2}^{t-1} \quad \forall i, j_1, j_2, t > 0, \quad (7)$$

The objective function (3) minimizes the total cost across all cars, sectors and time steps, excluding only cases where the origin and destination sectors are the same, since this incurs no travel time. Alternative formulations could consider a min-max approach to limit the maximum travel, but this can be enforced through fairness constraints if required. Equation (4) ensures the demand is satisfied for each sector in each time step. (5) ensures that each car is assigned to exactly 1 location in each time step. Constraints (6) and (7) linearize the conditional definition of x_{i,j_1,j_2}^{t-1} in (2) to satisfy the requirements of the MIP construct, enumerating all possible combinations of y_{i,j_1}^{t-1} and y_{i,j_2}^t.

Due to the linearized formulation of the travel time constraints, we require $\mathcal{O}(nm^2T)$ variables and $\mathcal{O}(nm^2T)$ equations for this formulation. Thus, while conceptually straightforward, the resulting MIP is large and intractable to solve by exact methods.

Enforcing Police Presence We enhance our model based on the concept of police 'presence' - the effect of deploying patrols not simply to respond to incidents, but to project a psychological effect - reassurance to honest citizens, and intimidation to criminals. We inject unpredictability into the movement of cars within their patrol sectors, by introducing a 'diversity' requirement, that stipulates a maximum number of consecutive time steps that each car may remain in the same sector for. When a car is at a particular location at time t, it must be at a different location after a certain duration after t. This requirement means that sectors will be patrolled by different officers at different times. Since the activities undertaken within the sector are at the autonomous discretion of the individual officers, the local patrol behaviour will change with the manpower allocation, making it harder to plan nefarious endeavours. For this problem we set a limit of 2 consecutive time steps on the amount of time a car can spend in a sector. This does not preclude returning to the sector later, but will ensure a reasonable degree of mobility among the team members. To implement this, the following constraint can be added to the MIP in (3)-(7):

$$y_{i,j}^{t-1} + y_{i,j}^{t+1} \leq 1 \quad \forall i, j, t. \quad (8)$$

Comparison with Constraint Programming An alternative approach is to formulate the problem as a CP model. One concern with the application of constraint programming is over the ability to solve the CP model (in our case, to minimize the overall travel time) efficiently, with say, the CP Optimizer, with the global cardinality constraint replacing the demand constraint. This would work but the time taken could be large relative to our solution approach.

Simulated Annealing and QUBO

A general QUBO formulation can be written as follows:

$$\min_{x \in \{0,1\}^n} x^T Q x \quad (9)$$

where Q is an $n \times n$ matrix and x are binary decision variables. Penalty methods and slack variables are used to transform the constraints (Lucas 2014). QUBO is closely related to the Ising model in quantum mechanics, and there has been great interest in using quantum adiabatic optimization techniques to solve NP-hard optimization problems via QUBO solvers. Here, we study if the deployment scheduling problem, where solution quality and execution speed are both important, can be solved more efficiently and effectively by a metaheuristic QUBO solver than a classical MIP solver (such as IBM's CPLEX solver). QUBO formulations permit quadratic terms and thus we are able to express products of binary variables in the objective function directly. We can take advantage of this non-linearity and avoid the inefficient

need to express $x^{t-1}_{i,j_1,j_2} = y^{t-1}_{i,j_1} \cdot y^{t}_{i,j_2}$ explicitly in the formulation, and more importantly, remove constraints (6) and (7), rewriting the optimization problem as:

$$\min \sum_{i=1}^{n} \sum_{j_1=1}^{m} \sum_{j_2 \neq j_1}^{m} \sum_{t=1}^{T} T_{j_1,j_2} y^{t-1}_{i,j_1} y^{t}_{i,j_2}, \tag{10}$$

subject to (4) and (5). This rewritten formulation is more efficient than the original MIP in (3)-(7), but it is not in true QUBO format, not being 'unconstrained'. To complete the transformation, we apply the penalty method, introducing penalty parameters A and B to move the constraints, (4) and (5) into the objective expression. We can also include the 'presence' constraint with a third penalty term and parameter, C, to generate the following expression:

$$\min \sum_{i,j_1,j_2,t} T_{j_1,j_2,} y^{t-1}_{i,j_1} y^{t}_{i,j_2} + A \sum_{j,t} \left(\sum_{i=1}^{n} y^{t}_{i,j} - d^{t}_{j} \right)^2 + B \sum_{i,t} \left(\sum_{j=1}^{m} y^{t}_{i,j} - 1 \right)^2 + C \sum_{i,j,t} y^{t-1}_{i,j} y^{t+1}_{i,j}. \tag{11}$$

This QUBO formulation can be passed to a QUBO solver and in our case we use Fujitsu's DA. This formulation only requires $\mathcal{O}(nmT)$ variables, thus we have saved a multiplicative factor of m variables. Hence, the QUBO representation is more compact than the MIP formulation. In order to minimize the original unconstrained problem, we have to tune the penalty parameters, A, B, and C. We employ Hyperopt (Parzen Tree Estimator) (Bergstra et al. 2011) to do this. At each Hyperopt parameter, we check the feasiblity of the problem and record when the parameter is suitable.

Annealing Results

We base our case study input on the real world problem in (Chase et al. 2021), scheduling police patrols for 3 different teams, X, Y, and Z, serving respectively 9, 13, and 17 patrol sectors, with manpower equal to the number of sectors served by each team. Time steps are hourly, and we consider both 6-hour and 12-hour shifts. We use two solvers to generate our experimental results - the IBM CPLEX solver provides the performance figures for the MIPs, and the Fujitsu DA solver provides the figures for the QUBO formulations.

Setting a Baseline First, we examine the scheduler without the diversity constraint. Whilst we would prefer to implement the presence feature in a practical setting, this provides a good baseline to establish whether a QUBO solver can outperform an industry standard MIP solver on a general problem and the results are given in Table 4. We set the DA to run for 10^8 iterations in Parallel Tempering mode, and for 40 Hyperopt iterations. With similar execution times, DA outperforms CPLEX, especially when the problem size increases. In area X, the small problem size enables CPLEX to prove optimality quickly, whereas DA does not prove optimality, and therefore continues to run despite finding a feasible solution quickly. As the instance size increases, the search space for the MIP grows and the advantage for DA and, by implication, quantum annealing, becomes evident. In Fig. 4, we present a representative excerpt from the schedules provided by each solver for team Z. Each column corresponds to a car, and each colour denotes a different patrol sector. As the colour changes down the column, the car patrols different sectors. The DA schedule is more stable with better individual travel times having only 3 cars with over 50 mins total travel time, while CPLEX has only one car with under 50 mins travel time.

	n, m, T	CPLEX Obj	DA Obj	CPLEX Time (s)	DA Time (s)
X	9, 9 , 6	572.5	572.5	7.7	1418
Y	13, 13, 12	301	254.9	2716	2714
Z	17, 17, 12	1504.7	538.6	4789	4785

Table 4: Comparison of DA and CPLEX. Objective measures the total migration time of the schedule in minutes.

DA	Car							
Time	1	2	3	4	5	6	7	8
T1	11	1	3	13	2	6	7	15
T2	11	1	3	13	2	6	7	15
T3	11	1	3	13	2	6	7	15
T4	11	1	3	13	2	6	7	15
T5	11	1	5	17	2	6	7	15
T6	11	1	5	17	2	6	7	15
T7	11	1	5	16	1	6	7	15
T8	11	1	5	16	1	6	7	15
T9	11	1	5	16	2	6	7	15
T10	11	1	5	16	2	6	7	15
T11	11	1	5	13	2	6	7	15
T12	11	1	5	13	2	6	7	15
Total:	0	0	18	128	39	0	0	0

CPLEX	Car							
Time	1	2	3	4	5	6	7	8
T1	3	2	13	2	9	13	11	4
T2	3	2	13	2	9	13	11	4
T3	3	2	13	2	9	13	11	4
T4	3	2	13	2	9	13	11	4
T5	5	2	17	3	9	13	11	4
T6	5	4	2	11	9	14	11	3
T7	5	16	1	7	10	13	10	3
T8	7	13	1	15	10	16	10	7
T9	7	16	1	15	10	14	11	7
T10	7	14	1	15	10	16	11	7
T11	4	2	1	15	4	11	11	7
T12	2	4	1	15	4	11	11	7
Total:	80	223	144	101	62	158	32	52

Figure 4: DA and CPLEX solutions for team Z. Each colour represents a different patrol sector.

Police Presence Constraint In Table 5, we compare the results and pure annealing time of DA with the CPLEX performance with a 10 minute solving time limit, which is comparable to the DA time. We find that setting the number of iterations to be 1.5×10^8 and the Hyperopt iterations to be 6 offers a good trade off to get a total annealing time of approximately 10 minutes, which is good for a daily deployment plan generator in a real life scenario. Currently, we have excluded the QUBO construction time which can be expensive. However, techniques such as precomputation can be adopted to reduce this time.

From Framework to Fieldwork

Deployment of the framework for law enforcement planning presented in this paper must satisfy two sets of requirements: the technical, and the human.

Firstly, a deployable system must be implementable at reasonable cost. The methods we have presented for both the incident prediction and the scheduler are usable on a single, albeit powerful, machine. We developed our approach within a service-based system architecture, where a central server process manages a web frontend for users to re-

m, n, T	CPLEX time (s)	CPLEX Obj	DA annealing time (s)	DA Obj
13,11,12	602	420	269	330
17,17, 10	603	2686	358	1761
18,18,20	608	9964	608	8350
20,20,14	607	8509	586	5434
20,20,16	609	9437	610	6566
20,20,18	610	11113	634	8880

Table 5: Solution times and objectives for police presence.

quest deployment plans and upload incident data. A worker process generates the patrol schedules, with a modular approach taken to the framework components, and a Python API allowing plug-and-play integration of external methods. Our ST-ResNet model is trained on a Tesla V100 GPU, but only making use of 4 cores for TensorFlow. Hyperparameter tuning is provided by the Ray Tune library integrated with the worker by a Flask app. The optimization model that determines the car demand per sector uses CPLEX, but achieves scalability by partitioning and can be executed on desktop hardware, in this case an Intel Core i7 3.40GHz with 16GB RAM. The annealing results were generated with a local datacenter-installed DA chip, but equivalent performance can be achieved using the Fujitsu DA cloud service, accessible from any system running Python. This approach also permits future-readiness, allowing a genuine quantum annealing method to be substituted when services such as DWave can operate at sufficient scale. Thus, all parts of the framework we propose have been implemented without the need for investment in an expensive computing cluster, in line with the needs of the case study agency, removing a significant barrier to wider adoption.

Secondly, the deployed system must fit within the existing working practices and operating procedures of the deploying agency so as to gain approval from both end users and management. We have developed our modeling assumptions in consultation with a real law enforcement agency. Our system user interface was developed through a collaborative design process to ensure that the tools, visuals, and terminology were in line with current operating procedures. The scheduler operates in line with current shift structures and deployment timings as agents on the ground have structured their lives around the existing shifts and are familiar with a degree of latitude in patrolling their sectors. Shift timings are easily adjustable system parameters, with users able to select the start and end times to a granularity of 30 minutes. The system passed User Acceptance Testing (UAT) with both the transformation team and the end users, although plans for a live field trial were disrupted by the COVID-19 pandemic. The live field trial should focus on a set of sectors that are active enough to provide insight into performance improvements. The trial should go through phases, first guiding actual patrols on a smaller region (such as Area X from the annealing results) and finally on a larger region, such as Area Z. If performance meets targets set by the agency, the path would be open for tender to developers to implement the system prototype as commercial calibre software.

Conclusion

We have presented a general framework for planning daily law enforcement deployments, consisting of three core components, of which we have presented viable generally applicable solution methods for two, using deep learning and simulated annealing. We have demonstrated the effectiveness of these methods using police patrol as a case study, outlining the necessary criteria for deployment of a system implementation, and discussed the ethical implications of such a system. Future work should move to field trials, while continuing to refine the methods, and examining the potential for AI to tackle ethical safeguards.

Acknowledgements

This research is funded by the National Research Foundation Singapore under its Corp Lab @ University scheme and Fujitsu Limited as part of the A*STAR-Fujitsu-SMU Urban Computing and Engineering Centre of Excellence.

Ethical Impact

The ethics of law enforcement has been a hot topic over the last year, and concerns about police bias arising primarily in the US, but also in other countries such as the UK, have highlighted the potential for misuse in an automated patrol system. In (Richardson, Schultz, and Crawford 2019), three case studies are presented highlighting cases where corrupt police practices and manipulated crime records, such as racially biased stop-and-frisk activity, was either directly used in prediction, or evidence of corrupt practices show a high likelihood of 'dirty' data being used, with a lack of transparency and accountability preventing assessment of the extent of the data corruption. Dirty data results in a feedback loop where biased input encourages biased behaviour. Our case study uses call records generated automatically by the public and therefore the distribution of incidents is not based on police decision-making, although the assignment of types and the extent of the response are. Additionally, in our case study there exists a high degree of trust between public and police, and a lack of racially segregated housing areas, both of which help to reduce the potential for bias. However, this may not hold true if our framework is applied to another case with different data, and a key part of the future development of this work to general law enforcement problems must lie in ethical safeguards. Safeguards include limiting use of data based on subjective police judgements, and disregarding data from periods of time when corrupt practices have been discovered, as well as crime types that are known to have strong prejudice connected (e.g. drugs-related offences in the US). Future work should focus on the development of intelligent tools to identify and compensate for bias, such as correlation with data from similar but less corrupt areas, analysis of demographics, not as predictors of crime but as predictors of bias, and machine learning models to learn patterns of bias that can be applied to temper prediction models. Predictive policing has the potential for generating powerful insights that can do great social good, but can similarly do great harm if misused.

References

Ajagekar, A.; Humble, T.; and You, F. 2020. Quantum computing based hybrid solution strategies for large-scale discrete-continuous optimization problems. *Computers & Chemical Engineering*, 132: 106630.

An, B.; Tambe, M.; and Sinha, A. 2017. *Stackelberg Security Games (SSG) Basics and Application Overview*, 485–507. Cambridge University Press.

Aramon, M.; Rosenberg, G.; Valiante, E.; Miyazawa, T.; Tamura, H.; and Katzgraber, H. G. 2019. Physics-Inspired Optimization for Quadratic Unconstrained Problems Using a Digital Annealer. *Frontiers in Physics*, 7: 48.

Bergstra, J. S.; Bardenet, R.; Bengio, Y.; and Kégl, B. 2011. Algorithms for hyper-parameter optimization. In *Advances in neural information processing systems*, 2546–2554.

Butt, U. M.; Letchmunan, S.; Hassan, F. H.; Ali, M.; Baqir, A.; Koh, T. W.; and Sherazi, H. H. R. 2021. Spatio-Temporal Crime Predictions by Leveraging Artificial Intelligence for Citizens Security in Smart Cities. *IEEE Access*, 9.

Caplan, J. M.; Kennedy, L. W.; and Miller, J. 2011. Risk Terrain Modeling: Brokering Criminological Theory and GIS Methods for Crime Forecasting. *Justice Quarterly*, 28(2): 360–381.

Chase, J.; Phong, T.; Long, K.; Le, T.; and Lau, H. C. 2021. GRAND-VISION: An Intelligent System for Optimized Deployment Scheduling of Law Enforcement Agents. *Proceedings of the International Conference on Automated Planning and Scheduling*, 31(1): 459–467.

Chase, J. D.; Nguyen, D. T.; Sun, H.; and Lau, H. C. 2019. Improving law enforcement daily deployment through machine learning-informed optimization under uncertainty. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-2019)*.

Chinchor, N. 1992. MUC-4 Evaluation Metrics. In *Fourth Message Uunderstanding Conference (MUC-4)*.

Crispin, A.; and Syrichas, A. 2013. Quantum Annealing Algorithm for Vehicle Scheduling. In *2013 IEEE International Conference on Systems, Man, and Cybernetics*, 3523–3528.

Dash, S. K.; Safro, I.; and Srinivasamurthy, R. S. 2018. Spatio-temporal prediction of crimes using network analytic approach. In *2018 IEEE International Conference on Big Data (Big Data)*, 1912–1917.

Date, P.; Patton, R.; Schuman, C.; and Potok, T. 2019. Efficiently embedding QUBO problems on adiabatic quantum computers. *Quantum Information Processing*, 18(4).

Della Fave, F. M. e. a. 2014. Security Games in the Field: Deployments on a Transit System. In *Lecture Notes in Artificial Intelligence*, 103–126. Springer.

Fang, F.; Stone, P.; and Tambe, M. 2015. When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing.

Ikeda, K.; Nakamura, Y.; and Humble, T. S. 2019. Application of quantum annealing to nurse scheduling problem. *Scientific reports*, 9(1): 1–10.

Jain, M.; Kardes, E.; Kiekintveld, C.; Tambe, M.; and Ordonez, F. 2010a. Security Games with Arbitrary Schedules: A Branch and Price Approach. In *National Conference on Artificial Intelligence (AAAI)*.

Jain, M.; Tsai, J.; Pita, J.; Kiekintveld, C.; Rathi, S.; Tambe, M.; and Ordóòez, F. 2010b. Software Assistants for Randomized Patrol Planning for the LAX Airport Police and the Federal Air Marshal Service. *Interfaces*, 40(4): 267–290.

Lau, H. C.; Yuan, Z.; and Gunawan, A. 2016. Patrol scheduling in urban rail network. *Annals of Operations Research*, 239(1): 317–342.

Luber, S.; Yin, Z.; Delle Fave, F.; Xin Jiang, A.; Tambe, M.; and Sullivan, J. P. 2013. Game-theoretic Patrol Strategies for Transit Systems: the TRUSTS System and its Mobile App (Demonstration). In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)[Demonstrations Track]*.

Lucas, A. 2014. Ising formulations of many NP problems. *Frontiers in Physics*, 2: 5.

Malleson, N.; and Andresen, M. A. 2015. Spatio-temporal crime hotspots and the ambient population. *Crime Science*, 4(1): 1–8.

Mukhopadhyay, A.; Zhang, C.; Vorobeychik, Y.; Tambe, M.; Pence, K.; and Speer, P. 2016. Optimal Allocation of Police Patrol Resources Using a Continuous-Time Crime Model. In *Decision and Game Theory for Security - 7th International Conference, GameSec 2016*, 139–158.

Richardson, R.; Schultz, J.; and Crawford, K. 2019. Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice. *New York University Law Review Online*, 94(192).

Saisubramanian, S.; Varakantham, P.; and Lau, H. C. 2015. Risk Based Optimization for Improving Emergency Medical Systems. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, 702–708.

Stollenwerk, T.; Lobe, E.; and Jung, M. 2019. Flight Gate Assignment with a Quantum Annealer. In Feld, S.; and Linnhoff-Popien, C., eds., *Quantum Technology and Optimization Problems*, 99–110. Cham: Springer International Publishing.

Wang, B.; Zhang, D.; Zhang, D.; Brantingham, P. J.; and Bertozzi, A. L. 2017. Deep Learning for Real Time Crime Forecasting. *arXiv:*, 1707.03340.

Wang, W.; Dong, Z.; An, B.; and Jiang, Y. 2021. Toward Efficient City-Scale Patrol Planning Using Decomposition and Grafting. *IEEE Transactions on Intelligent Transportation Systems*, 22(2): 747–757.

Zhang, J.; Zheng, Y.; and Qi, D. 2017. Deep Spatio-Temporal Residual Networks for Citywide Crowd Flows Prediction. In *Thirty-First AAAI Conference on Artificial Intelligence*, AAAI'17, 1655–1661. AAAI Press.