# When Security Games Hit Traffic: A Deployed Optimal Traffic Enforcement System

Ariel Rosenfeld, Oleg Maksimov, Sarit Kraus

*Bar-Ilan University*
*Israel*

**Abstract**

Road accidents are the leading causes of death among youths and young adults worldwide. Efficient traffic enforcement is an essential, yet complex, component in preventing road accidents. In this article, we present a novel model, an optimizing algorithm and a deployed system which together mitigate many of the computational and real-world challenges of traffic enforcement allocation in large road networks. Our approach allows for scalable, coupled and non-Markovian optimization of multiple police units and guarantees optimality. Our deployed system, which utilizes the proposed approach, is used by the Israeli traffic police and is shown to provide meaningful benefits compared to existing standard traffic police enforcement practices.

*Keywords:* Security, Traffic Enforcement, Deployed System

## 1. Introduction

About 1.35 million people die worldwide each year as a result of road accidents and between 20 and 50 million people suffer disability or other severe injuries [1]. An essential component in mitigating serious traffic accidents (accidents that cause death or injury) is efficient traffic enforcement, which is based on giving drivers the feeling that they are likely to be caught and sanctioned when breaking the law [2]. In fact, recent studies suggest that drivers respect traffic laws mainly due to enforcement concerns, rather than safety concerns (e.g., [3]). As a result, efficient traffic enforcement has been shown to reduce a wide range of high-risk, illegal driving behaviors, including driving while under the influence of drugs/alcohol, speeding, lack of seatbelt use and red-light running, and thus reduces road accidents (e.g., [4, 5, 6, 7] to name a few). Unfortunately, traffic police cannot cover the entire road network given its limited number of police cars and officers [8].

Within the Security Games (SG) field, optimal police allocation mechanisms for mitigating various types of crimes

*Email address:* ariel.rosenfeld@biu.ac.il (Ariel Rosenfeld)

have been developed. The generic SG framework consists of a defender (traffic police) who has a limited number of resources (police cars) to protect a large set of targets (road segments) from an adversary (reckless drivers) [9, 10]. To the best of our knowledge, [11] is the only work in the scope of SG which addresses traffic enforcement. The authors model the problem as a Stackelberg Security Game (SSG) where traffic police seek to apprehend reckless drivers who in turn seek to avoid apprehension. In a SSG, the traffic police commit to a mixed strategy where drivers can first observe and then respond as best as possible. In practice, traffic enforcement seeks to reduce traffic accidents (and not necessarily to apprehend reckless drivers) [12]. Furthermore, due to the dynamic environmental factors which influence driving behavior (weather, traffic jams, etc.), drivers have been shown to act in a less strategical manner, responding to changes in their environment, including the observed police presence in current and past rounds [13]. Therefore, SSGs seem unsuitable to the task of preventing serious road accidents.

Non-strategical adversaries in SG settings have recently been modeled as opportunistic criminals which choose where and when to commit a crime in real-time based on police presence and the attractiveness of the potential targets [14]. Opportunistic criminals are *reactive* to police actions and do not consider their behaviors' effect on future police actions. We adopt this approach here, modeling the drivers, and consequently accidents, as reactive to police allocations. However, unlike [14], drivers may react to *both present and past police enforcement allocations*, making the authors' Markovian assumption unsuitable. For example, it has been shown that drivers continue to react to police presence long after the enforcement operation has ceased (a phenomena also known as time and distance halo) [13]. Basilico et al. [15] have investigated non-Markovian strategies for robotic patrols. However, the authors assume that the attacker is strategic, and therefore the approach is inapplicable. To our knowledge, no work has efficiently addressed the non-Markovian property in SG.

Most allocation mechanisms in SG simplify the computational task by assuming that planning for each police unit *separately* will bring about a (near-)optimum solution [16]. However, this is not the case in traffic enforcement. For example, experts from the Israeli Traffic Police (ITP) claim that if police cars are stationed at the same place and time, their effectiveness in reducing traffic accidents cannot be assumed to be greater than the effectiveness of a single police car at the same point and time, *a fact we leverage in this work*. Furthermore, significant benefits may accrue from coordination across multiple police units, e.g., allocating two police cars in *adjacent* road segments could have a stronger impact than allocating a single police car. This notion relates to the coordinated actions notion in [17] which captures the combined effects of multiple defenders guarding the same target simultaneously. As a result, the computational task of deriving optimal traffic enforcement allocation in order to prevent serious road accidents is both *coupled* and *non-Markovian*, which makes it computationally intractable. Namely, the optimal allocation of traffic enforcement at time $t$ could depend on the trajectories of *all police cars* (i.e., coupled) up to time $t$ (i.e., non-Markovian).

In order to address these shortcomings, we first formulate the TRAFFIC ENFORCEMENT ALLOCATION PROBLEM (TEAP). We prove that deriving or approximating the optimal solution to a TEAP is NP-Hard, and remedy this hardness by introducing an optimal novel algorithm called the RELAXED OPTIMIZATION SOLVER ENHANCER, or ROSE for short. ROSE uses a master/slave optimization approach, aimed at reducing the computational burden of directly solving the TEAP, and leverages common characteristics of TEAPs that have not been investigated in previous works. In an extensive set of lab-based empirical evaluations, we show that ROSE favorably compares to several baseline approaches, achieving a significant speed-up, using both synthetic and real-world road networks.

Based on the promising results obtained in our lab-based evaluation (Section 4.1), we extended our model to make it suitable for real-world deployment. Through a four month-long controlled field study with the ITP, we show that the results obtained in our lab-based evaluation translate well to the field, resulting in an improved traffic enforcement policy as depicted by the number of accidents and average driving speeds in the deployment district compared to the control condition.

As such, the contribution of this article is twofold: First, the theoretical TEAP modeling and the novel ROSE solution technique. Second, the adaptation of these lab-based developments into a beneficial real-world application which can be adopted and extended by researchers and practitioners in the traffic safety and security games fields.

All code and procedures used in this study are available at `http://www.biu-ai.com/trafficPolice` in order to encourage other researchers to tackle the important and challenging task of preventing serious traffic accidents.[1]

The remainder of this article is structured as follows: Section 2 presents the TEAP formulation and analyzes its properties. In Section 3 we present our solution approach, followed by its evaluation in the lab and in the field (Section 4). Then, in Section 5, we discuss the obtained results. Lastly, we summarize the work and highlight future research directions in Section 7.

## 2. Traffic Enforcement Allocation

We model the interaction between drivers and police as a repeated game over $T(< \infty)$ rounds, which takes place on a road network, represented as a graph $G = \langle V, E \rangle$ where $V = \{v\}$ is the set of intersections and $E = \{e = (u, v)\}$ is the set of road segments. We assume that no accidents occur off-road, and therefore $E$ is the set of enforcement targets in this work (intersection $v$ is considered part of the road segments that share $v$, thus there is no need to consider $v$ as a different

[1] Unfortunately, we are unable to release confidential data provided by the ITP such as the number of police resources used, where and when traffic citations where given, etc.

target). Without losing generality, we assume that the time it takes to travel through each road segment is 1 round; this assumption can be relaxed by including dummy vertices.

The traffic police has $k(<< |E|)$ police cars at its disposal. At each round $t$, the police places enforcement on a subset of size $k$ from $E$, which we refer to as the allocation at round $t$ denoted $a_t$, such that the allocation respects the graph's connectivity constraints and *no more than a single police car is assigned to any edge*. Namely, at round $t$, each police car can either stay in its current road segment (enforcing for a longer period of time) or move to an adjacent edge given $a_{t-1}$. $a_1$ can assume any subset of size $k$ of $E$. We denote the traffic police *allocation history* at round $t$ as $H_t = \langle a_1, \dots, a_t \rangle$. We use the notation $e_t$ to denote road $e$ at round $t$, and $H[e_t]$ as an indicator of whether a police car is assigned to road $e$ at round $t$. Simultaneously, drivers choose whether to obey the law (drive safely) or not at each road segment $e \in E$.

We assume that drivers' actions at round $t$ are visible to the police. For example, the ITP, as with many other police departments, uses anonymous cellular reports provided by commercial companies to evaluate the distribution of speeds on each road in real-time. Other technological aids such as speed cameras are also in use. Note that while the police does not consider the behavior of *each driver individually*, they do obtain aggregated statistics on traffic behavior for the entire road network. On the other hand, drivers are only exposed to a noisy signal regarding the police allocation. For example, common applications such as WAZE and other technological instruments such as police scanners allow drivers to have an indicator of police presence at $e_t$. However, these indicators are not completely accurate (police presence in a road segment is not always reported in WAZE, an indicator of police presence may not be up-to-date, a police car may be covert, etc.). As a result, the game is conducted *under one-sided uncertainty*. Due to this uncertainty, the drivers base their actions at $e_t$ according to $a_t$ (although not completely visible) and the police's past allocations ($H_{t-1}$), which together constitute $H_t$.

Following recent advancements in predictive policing, including the prediction model constructed in the course of this study, and in the same spirit as done in previous works such as [18], we define the *risk of accidents* occurring at $e_t$ as $\texttt{risk}(e_t)$. The risk function measures the likelihood that a serious traffic accident will occur at $e_t$ *in the absence of police enforcement* (in the [0,1] range). We further define the *effectiveness of enforcement* as $\texttt{eff}(H_t, e_t)$. $\texttt{eff}$ measures the effect that the police allocation history has on the risk of accidents occurring at $e_t$.

The traffic police is interested in minimizing the total expected number of accidents occurring throughout the game. Formally, it seeks to minimize the following objective of the optimization problem we denote as the TRAFFIC ENFORCEMENT ALLOCATION PROBLEM (TEAP):

$$min_{H_T} \sum_{t=1,\dots,T} \sum_{e \in E} \texttt{risk}(e_t)(1 - \texttt{eff}(e_t, H_t)) \tag{1}$$

$\texttt{risk}(e_t)$ cannot be influenced by police *enforcement* but rather through modification of the road's characteristics (e.g.,

| notation | meaning |
|---|---|
| $t \leq T$ | Game round index. |
| $e_t$ | Road segment $e$ at round $t$. |
| $a_t$ | Defender's allocation in round $t$. |
| $H_t$ | Defender's allocation history at round $t$. |
| $H[e_t]$ | Indicator whether police is present at $e_t$. |
| $\texttt{risk}(e_t)$ | Likelihood of a car accident occurring at $e_t$ in the absence of police enforcement. |
| $\texttt{eff}(e_t, H_t)$ | The effectiveness of police enforcement on $e_t$. |

Table 1: Summary of key notations.

number of lanes), traffic (e.g., reducing speed-limit), etc. On the other hand, eff heavily depends on police enforcement, $H_t$. We assume both $\texttt{risk}(e_t)$ and $\texttt{eff}(e_t, H_t)$ are known to the police and can be computed in polynomial time.

A summary of the notations used in this article is available in Table 1.

The solution to Eq. (1) prescribes a pure strategy for the traffic police. The police could optimize over all rounds simultaneously, however this approach is computationally expensive; it needs to solve a possibly non-convex optimization problem as the police must consider drivers' responses (modeled within eff). Unfortunately, approximating the optimal solution to a TEAP, within any constant factor, is hard even for a single driver and a single police car.

**Theorem 1.** TEAP *cannot be approximated within any factor of* $c \geq 1$ *in polynomial time, unless* $P = NP$.

*Proof.* In order to prove the theorem, we give a reduction from SAT to TEAP with one driver and one police car: On input $\Phi(x_1, \ldots, x_n)$, construct $n+1$ nodes $V = \{v_i\}$, $i = 1, \ldots, n+1$. Then connect node $i$ with node $i+1$ ($i = 1, \ldots, n$) using 2 directed edges, one for $x_i = True$ and one for $x_i = False$ and a single directed edge from $v_{n+1}$ to $v_1$ representing $Satisfiable$ (S). Consider the resulting graph $G = (V, E)$ as the road network for a TEAP with $T = n + 1$. A single police car starts at $v_0$. Let risk assume 0 for all edges at all rounds except for edge $S$ at round $t + 1$, which assumes the value of 1. Let eff assume 0 for all edges, rounds and allocation histories except for $\texttt{eff}(S, H_t)$, which assumes the value of 1 if the police trajectory ($H_t$) corresponds to a satisfying assignment for $\Phi(x_1, \ldots, x_n)$ and 0 otherwise. Clearly, the driver's action (causing an accident at edge S or not) can be decided in polynomial time.

The above construction takes polynomial time. Assume to the contrary that such an approximation polynomial time algorithm $App(G)$ exists. If there is no satisfying assignment to $\Phi$, then every trajectory the police car may take will bring about an objective value of 1, thus $App(G) \geq c$. If there is a satisfying assignment, then the defender can take the respective trajectory and receive a value of 0, hence $App(G) = 0$. □

Two key computational challenges arise from the TEAP formulation. First, the arbitrary `risk` and `eff`, which can take any polynomial time computable form and depend on an unbounded history of police actions (`eff`), pose a significant optimization challenge. Second, the space of possible police strategies (joint schedules for all police cars) grows exponentially in the number of resources and the number of time steps which make the computation even more challenging.

## 3. Optimizing Police Strategy

In this work we derive an optimal *pure strategy* for traffic enforcement for $T$ steps. Our goal is to find the pure strategy that would minimize the total expected number of serious accidents. In our framing, any randomized mixed-strategy, which is the combination of pure strategies, results in a greater or equal number of accidents than the optimal pure strategy, as in [14].

Given Theorem 1, we resort to remedying the hardness of solving the TEAP by introducing an optimal novel algorithm called the RELAXED OPTIMIZATION SOLVER ENHANCER, or ROSE for short. ROSE uses a master/slave optimization approach, aimed at reducing the computational burden of directly solving the TEAP. It exploits the fact that no two police cars are allowed to enforce the same road segment at the same time. ROSE is guaranteed to return an *optimal solution*, hence, in the worst case, ROSE will run in exponential time. Nevertheless, experimental results (see Section 4) on both synthetic and the Israeli road networks demonstrate that ROSE is able to derive an *optimal solution* significantly faster than competing approaches under various real-world conditions.

Before introducing ROSE, we first cast the TEAP as a binary graph flow problem and present an exponential sized Binary Integer Program (BIP) for solving it.

### *3.1. TEAP as Graph Flow*

Similar to other transition-based security formulations, we model the TEAP using a *transition graph* [19]. The transition graph is a compact representation which captures the spatio-temporal structure of the road network and allows us to handle the exponential strategy space by avoiding the enumeration of all pure strategies. Technically, given a road network $G$, we transform it into a $T$ time-expanded graph $G_T$ such that each vertex $v$ (edge $e$) is replicated $T$ times, one for each round, denoted $v_t$ ($e_t$).

Each $v_t$ in the transition graph is associated with the number of police cars that start their trajectories in it minus the number of police cars that end their trajectory in it, denoted $b_{v_t}$. For example, $b_{v_t} = 0$ means that either no police cruiser starts or ends its route in $v$ at time $t$, or, more generally, the same number of police cruisers start and end their route in $v$ at time $t$. $b_{v_t}$ is assumed to be known in advance and cannot be changed by the police.[2] The resulting flow problem can be formulated as the following mathematical program:

$$\min_{H_T} \sum_t \sum_{e_t} \texttt{risk}(e_t) \cdot (1 - \texttt{eff}(e_t, H_t)) \tag{2}$$

$$\text{s.t} \sum_{v'_{t-1}} H_t[(v'_{t-1}, v_t)_{t-1}] - \sum_{v'_{t+1}} H_{t+1}[(v_t, v'_{t+1})_{t+1}] = b_{v_t} \quad \forall v_t \in G_T \tag{3}$$

$$H_T[e_t] \in \{0, 1\} \quad \forall e, t \tag{4}$$

Constraints (3) and (4) are standard binary flow constraints. Let $Sol = \{e_t | H_T[e_t] = 1\}$ denote the set of $e_t$s that were assigned a unit of flow (a police car) in the optimal assignment.

We transform the above mathematical program into a 0-1 integer linear program (or Binary Integer Problem, BIP for short) of exponential size, using the following non-standard procedure. The procedure is intended to cast the original optimization problem formulation into a unique structure that will then be exploited by our novel optimization technique – ROSE (Section 3.2). $\texttt{risk}(e_t)$ and $\texttt{eff}(e_t, H_t)$ are enumerable; for every $e_t$ and possible $H_t$ (which is bounded in size by $2^{|V||E||T|}$) one can *conceptually* calculate the value of $\texttt{risk}(e_t) \cdot (1 - \texttt{eff}(e_t, H_t))$ offline and store it in a table. For every entry $i$ in the table, let $Value_i$ denote the value of $\texttt{risk}(e_t) \cdot (1 - \texttt{eff}(e_t, H_t))$ for a given $e_t$ and assumed allocation history $H_t$. We denote $Visited_i = \{e_t | H_t[e_t] = 1\}$ as the set of $e_t$s that assumed the value of 1 under the allocation history of table entry $i$. For every entry $i$ we create a new binary variable $p_i$ which takes the value of 1 if $Sol \cap Visited_i = Visited_i$. Mathematically, we add the constraint:

$$p_i = \prod_{e_t \in \{Visited_i\}} H_T[e_t] \tag{5}$$

Equation (5) might seem non-linear at first. However, it is rather easy to linearize it using a fix-sized set of linear constraints that will force the indicator $p_i$ to assume the correct value (an explanation of the procedure is available in Appendix B).

For all table entries $i$, other than the empty allocation $H_T = \emptyset$, there exists at least one table entry $j$ such that $Visited_j$ is a *strict* (proper) subset of $Visited_i$. As such, if $p_i = 1$ then $p_j = 1$. In order to "isolate" the true value of some $H_T$,

[2] This formulation allows police cars to start and finish their paths at different times and locations.

we modify the optimization objective (2) using the inclusion-exclusion principle as follows: Let $Pow(i)$ be the set of all *strict* (proper) subsets of $Visited_i$.

$$min_{H_T} \sum_t \sum_{e_t} \sum_i p_i(Value_i + \sum_{Visited_j \in Pow(i)} (-1)^{|Visited_j \cap Visited_i|+1} Value_j) \tag{6}$$

Intuitively, given an $H_T$, we derive the true value of $H_T$ by appropriately adding and deducting the values associated with all proper subsets of the visited $e_t$s according to the inclusion-exclusion rule. For interperability purposes, for a given $e_t$ and $i$, we shall refer to a summed term as *penalty* if the summed term is *positive*, and *reward* otherwise.

Clearly, the result is a BIP. Furthermore, the resulting BIP is *not sensitive to the number of police cars*. The correctness of the above procedure easily follows that of the inclusion-exclusion principle. In order to understand the procedure better, consider the following example:

**Example 1.** *Assume a time-expanded graph with 2 vertices $(v, u)$ expanded over 3 time steps $(v_1, u_1, v_2, u_2, v_3, u_3)$ such that $v_1$ and $u_1$ are connected to $v_2$ and $u_2$, and $v_2$ and $u_2$ are connected to $v_3$ and $u_3$. There are 2 guards, starting at nodes $v_1$ and $u_1$, and they finish their trajectories at $v_3$ and $u_3$. Overall, the problem induces $8$ binary decision variables, written in short as $I_{v_1,v_2}, I_{v_1,u_2}, I_{u_1,v_2}, I_{u_1,u_2}$, etc.* `risk` *is set to 1 for all edges.* `eff` *is set to 1 for all edges and strategies except for $(v_1, v_2)$ which is set to 0.6 if $I_{v_1,v_2} = 1$, 0.8 if $I_{u_1,v_2} = 1$ and to 0.5 if both $I_{v_1,v_2} = 1$ and $I_{u_1,v_2} = 1$. Technically, one can define $2^8 - 1$ new variables, each corresponding to a non-empty subset of the $8$ binary decision variables defined above. However, given that only $I_{v_1,v_2}$ and $I_{u_1,v_2} = 1$ bear an effect on the objective, it is sufficient to consider only three new variables, $p_1, p_2$ and $p_3$, and the following three new constraints: $p_1 = I_{v_1,v_2}$, $p_2 = I_{u_1,v_2}$ and $p_3 = I_{v_1,v_2} \cdot I_{u_1,v_2}$. In simple words, $p_1 = 1$ iff a guard traverses the $(v_1, v_2)$ edge, $p_2 = 1$ iff a guard traverses the $(u_1, v_2)$ edge and $p_3 = 1$ iff $p_1 = p_2 = 1$. Considering the above 3 new variables, the modified optimization objective is:* $\min_I 8 + (0.6 - 1)p_1 + (0.8 - 1)p_2 + (0.5 - (0.6 + 0.8) + 1)p_3$. *Note that the terms associated with $p_1$ and $p_2$ are* ***rewards*** *as placing a guard on the corresponding edges helps the optimizer lower the objective. The term associated with $p_3$ is a* ***penalty*** *as it lowers the effectiveness of the guards obtained though $p_1$ and $p_2$, thus obstructing the optimizer from lowering the objective.*

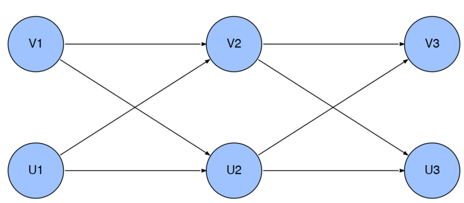

### 3.2. Linear Optimization Using ROSE

The resulting BIP of the procedure above cannot scale up due to the exponential number of variables and constraints (see evaluation in Section 4). To overcome this limitation we introduce a novel master/slave-based optimization algorithm, ROSE, Algorithm 1. The Master program consists of two levels: At the high level the Master program maintains a subset of penalty terms, denoted $P$. At the low level a BIP solver is used to solve a **relaxed BIP** in which only a subset of penalty terms are introduced along with their associated binary variables, $p_i$. At the beginning of the execution, $P$ contains all penalty terms extracted from Eq. (6) (Line 3) and the low level solver generates a solution, $Sol$ (Line 6), while considering only $p_i$ variables associated with reward terms (Line 2) . Simply put, the solver is executed on a smaller scale optimization problem which consists of only a subset of penalty terms (at the initial execution, an empty set). Given $Sol$, the Slave program is used to examine whether any penalty term $p \in P$ is *triggered* (Line 7), that is, the Slave program (Lines 12-14) checks whether any binary variable $p_i$ associated with a penalty term in $P$ should assume the value of 1 given $Sol$. If no penalty terms from $P$ are triggered, the Slave returns an empty set, indicating that an optimal solution has been found and ROSE terminates (Lines 8 and 9); otherwise, a set of penalty terms $P' \subseteq P$ is returned. The returned $P'$ is injected into the relaxed BIP, $P'$ is removed from $P$ by the Master and the process is repeated (Lines 10 and 11).

The Slave program can return any subset $P' \subseteq P$ as long as it obeys the following two rules: 1) $P' = \emptyset$ if no penalty terms from $P$ are triggered under $Sol$; and 2) $P'$ contains at least one penalty term (if such exists). We use an elementary implementation of the Slave program, returning all *triggered* penalty terms from $P$. Namely, we enumerate all (relevant) subsets of $Sol$ and return the triggered penalty terms. The investigation of more elaborate Slave programs which *predict* which penalty terms are most beneficial to introduce, in terms of minimizing ROSE 's run-time, is left for future work.

We demonstrate the use of ROSE in the following example:

**Example 2.** *Using the same setup of Example 1, the set of penalty terms extracted from the original formulation consists of only $p_3$ and its associated value. Therefore, the relaxed BIP's objective is:* $\min_I 8 + (0.6 - 1)p_1 + (0.8 - 1)p_2$. *With the constraints of Example 1 in mind, it is clear that the generated solution $Sol$ will result in $p_1 = p_2 = 1$. The Slave program will then detect that $p_3 \in P$ assumes the value of 1 and therefore injects it into the relaxed BIP and the process is repeated. In the second iteration, no new penalty terms are triggered (as there are none), and the process is concluded.*

It is well known that BIP solvers are sensitive to the number of constraints. Therefore, ROSE's computational performance depends on the number of penalty terms in $P$ which can be avoided in the iterative penalty generation process. it is hard to guarantee the computational benefit of the approach in the general case. While ROSE may be inefficient in some cases (e.g., no penalty term can be avoided regardless of Slave implementation, as is the case in Example 2), in several

**Algorithm 1** ROSE

**Require:** Time-expanded graph $G_T$, BIP Solver $Solver$.

**function** MASTER
    $BIP \leftarrow$ Initialize BIP with reward terms
    $P \leftarrow$ Penalty terms
    $Sol \leftarrow \emptyset$
    **repeat**
        $Sol \leftarrow Solver(BIP)$
        $P' \leftarrow Slave(Sol, P)$
        **if** $P' = \emptyset$ **then**
            **return** $Sol$
        $P = P \setminus P'$
        Introduce $P'$ into $BIP$
**function** SLAVE(Sol,P)
    $P' = \{p | p \in P \wedge$ p is triggered by Sol$\}$
    **return** $P'$

settings, including realistic and real-world traffic enforcement settings, it can bring about a significant improvement in runtime without jeopardizing the solution quality (see Theorem 1).

**Proposition 2.** ROSE *always terminates and returns an optimal solution.*

*Proof.* The Slave program introduces at least one penalty term to the relaxed BIP at each non-terminal iteration. Due to the finite number of penalty terms, ROSE terminates after a finite number of steps. At each iteration, the value of each feasible solution cannot decrease as ROSE only introduces penalty terms to the objective function. When ROSE terminates, all penalty terms triggered by $Sol$ have been injected into the relaxed BIP, therefore the relaxed BIP's objective value under $Sol$ is the optimal value under both the relaxed BIP and the original BIP, and its objective value would not change if any additional penalties from $P$ were to be added to the objective function. □

ROSE bears some similarity to the classic column generation approach [20]. The premise of both approaches is that most of the variables will assume a value of zero in the optimal solution. As such, only a subset of variables needs to be considered when solving the problem. However, the two differ in that column generation leverages the above idea to generate only the variables which have the *potential to improve* the objective function while ROSE introduces variables which *definitely deteriorate* the objective function yet must be taken into account in order to assert optimality.

Similar to other security settings, `eff` is assumed to be *submodular* [17]. Namely,

**Definition 3.** `eff` *is submodular if for every* $e_t, e'_t$ *and* $H_t \subseteq H'_t$, $\texttt{eff}(e_t, H_t \cup \{e'_t\}) - \texttt{eff}(e_t, H_t) \geq \texttt{eff}(e_t, H'_t \cup \{e'_t\}) - \texttt{eff}(e_t, H'_t)$

A submodular `eff` means that performing an additional enforcement activity (allocating a police car at $e'_t$) has diminishing gains in effectiveness. In Section 4 we show the significant runtime benefits that can be generated by ROSE when this property holds.

## 4. Evaluation

We evaluate our approach in two experimental settings: First, we perform a *lab-based evaluation* using a newly built realistic simulation environment consisting of both synthetic and the Israeli road networks. Through this examination, we demonstrate the strengths and limitations of ROSE . Second, following a year and a half-long process, we perform an *in-the-field evaluation* of our approach in a controlled study with the ITP. For reproducibility purposes and to facilitate future research on traffic enforcement, we release our simulation environment. Complete details, source code and data are provided in `http://www.biu-ai.com/trafficPolice`.

### *4.1. Lab-based Evaluation*

We first describe our simulation environment followed by the competing approaches evaluated in this part of the study. Then, we report the results obtained for the synthetic and the Israeli road networks.

#### *4.1.1. Simulation Environment*

Our simulation environment consists of 3 components: 1) Synthetic and Israeli road networks; 2) A state-of-the-art prediction model for modeling `risk`; and 3) A submodular `eff` function. `risk` and `eff` are derived from 11 years of accident data, extensive literature review on accident prevention and analysis and human expert knowledge from the ITP. We will describe the main components below.

`risk`. We obtained a record of 11 years of accident reports from the Israeli Central Bureau of Statistics (2005-2015). By cross-referencing these reports with additional sources such as the Israeli GIS database and weather reports, we were able to characterize each accident using 117 features, including infrastructure characteristics (e.g., number of lanes), date and time characteristics (e.g., weekend/weekday), weather (e.g., precipitation), traffic (e.g., average speed), etc. The full list of features is available in Appendix A. To the best of our knowledge, this is the largest set of features ever to be used to predict serious car accidents. For comparison, the Indiana traffic police use an intelligent accident prediction tool `http://www.in.gov/isp/ispCrashApp/main.html` which is based on approximately 90 features which we also use here. Experts in traffic enforcement claim that only the Indiana and Tennessee State traffic police use accident prediction tools but we were only able to obtain the latter's features. Using more than 30,000 accident records and under sampling the "non-accident" class (see [21]), we trained a deep neural network model that, given 110 features representing $e_t$, returns a value in the [0,1] range, *acting as a proxy* to the likelihood of an accident occurring at $e_t$.[3] We compared our prediction model to several baseline prediction models such as logistic regression, SVM and XGBoost (which is currently in use by the Indiana traffic police). Our model achieves an AUC of 0.87, outperforming logistic regression, SVM and XGBoost which recorded 0.78, 0.77 and 0.82, respectively.

`eff`. We base `eff` on [22], which used a unique database to track the exact location of the Dallas Police Department's patrol cars throughout 2009 and cross-referenced it with the car accidents of that year. To the best of our knowledge, this is the most recent investigation of the topic. The author found that if $e_t$ is enforced, `eff` should assume a value of 36%. However, enforcement effects are not restricted to the specific time and space in which the enforcement is performed. For example, *Time halo* is the time and the intensity to which the effects of enforcement on drivers' behavior continue after the enforcement operations have been concluded. It has been recorded that longer enforcement efforts cause more intense time halo effects that can last for hours and influence the next day(s) or even week(s) during the same time of day as the enforcement. *Distance halo* is defined as the distance over which the effects of an enforcement operation last after a driver passes the enforcement site. The most frequent distance halo effects are in the range of 1.5 - 3.5 kilometers from the enforcement site (see [13] for a review). Currently, there is no consensus on a mathematical modeling of these two halo effects, separately or combined. In the absence of real-world data such as was used in [22], we resort to an expert-based approach [23] and define `eff` in accordance with the ITP's estimations. We define time halo effects in the exponential diminishing form $\frac{36}{2^k}\%$ where $k \geq 0$ is the number of time-steps that have passed since the enforcement effort. To avoid negligible effects, we prune the effect at $k = 3$. The Distance halo effect is defined to be 5%, given that the two road segments are adjutant. Given the police allocation, `eff` assumes a simple submodular form where `eff` takes the largest applicable effect and adds half of each of the smaller appropriate effects to it. For example, if both $e_t$ and $e_{t+1}$ are enforced (and no other time or distance halo effects are appropriate), `eff` assumes 45% $(= 36\% + \frac{18}{2}\%)$.

[3]Note that serious accidents are sporadic events in both time and space. Therefore, directly estimating the probability of accidents occurring at $e_t$ is extremely challenging.

We are currently investigating a more data-driven approach for modeling `eff` in Israel.

#### 4.1.2. Competing Approaches

ROSE is compared with 4 baseline solutions: First, a **Naïve** solver which solves the entire BIP (Eq. (6)) in its general form. Second, a **Random** solver which for each police unit selects an action at random at each time step, resolving conflicts locally. Third, a **Greedy** solver, which computes a greedy path for each individual police car *iteratively*, capturing a (wrongly) assumed additivity in individual police car gains. Greedy considers a simplified version of `eff` which only accounts for the marginal gains that an enforcement in a road segment will generate given the current allocation of other police cars. Given the calculated path, Greedy updates the simplified `eff` given the visited road segments and continues to the next police car. Finally, we compare ROSE with **Domain Expert** allocations from the ITP. We could not evaluate Cartesian product solutions, which capture the joint effects of all police units, such as the ones presented in [24, 14], due to their lack of scalability in the number of road segments (we were unable to solve road networks larger than 5 road segments, which are unrealistic).

The resulting allocations are evaluated on the basis of two criteria: 1) **Accidents**, the expected number of accidents (i.e., the objective value of Eq. (6)) normalized by the no-police enforcement condition; 2) **Runtime and Scalability** of the deployed algorithm with respect to the number of police cars, road segments and the density of the road network [25].

The evaluation was done on a personal computer with 16 GB RAM and a CPU with 4 cores each operating at 4 GHz. The BIP solver was GUROBI [26].

#### 4.1.3. Synthetic Road Networks

We evaluate ROSE, Naïve, Random and Greedy on a series of synthetic road networks. We used 2 sets of synthetic road networks: Small networks (each consists of between 40 and 100 road segments in intervals of 10) and realistic networks (each consists of between 200 and 400 road segments in intervals of 100). Connectivity between road segments (i.e., the network density) is randomized such that each two road segments are connected by an intersection with a probability ranging between 0.05 and 0.15 (in intervals of 0.05), allowing for different topologies. `risk` uniformly samples a value in the $[0, 1]$ interval for each road segment and round and `eff` is defined as in our simulation environment. The number of police cars is set to either 5, 10 or 15 and $T$ is set for either 8, 16 or 24. Overall, 270 networks were evaluated. A 30 minute timeout was set for all conditions and networks.

*Accidents.* As expected, ROSE and Naïve return optimal allocations. On average, they reduce 22.7% and 5.3% of the no-enforcement objective value (the expected number of accidents) in small and realistic networks, respectively. On the other hand, on average, Random and Greedy reduce 1% of the no-enforcement objective value in both small and realistic

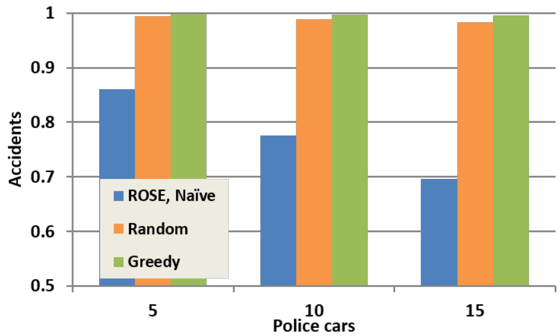


(a) Accidents/Police cars ($T$=8)

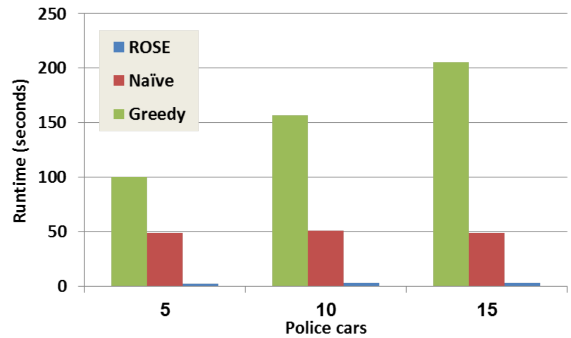


(b) Runtime/Police cars ($T$=24)

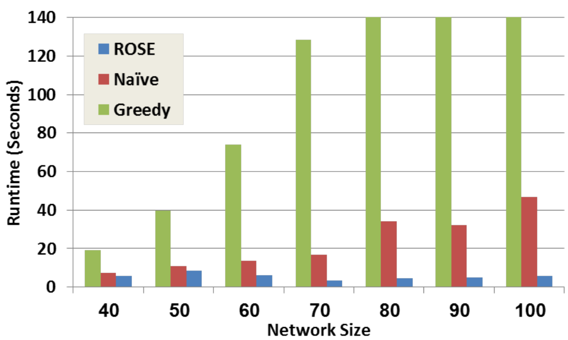


(c) Runtime/Size ($T$=24)

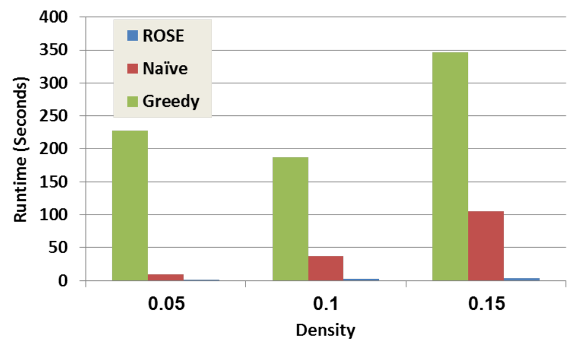


(d) Runtime/Density ($T$=24)

Figure 1: Synthetic road networks: results for the small networks set. In all Figures, the lower - the better.

networks. In realistic networks Greedy exceeded the timeout for all networks of size 300 and 400 and thus its quality cannot be evaluated properly. In our trials, Random and Greedy did not come up with an optimal allocation in any of the cases. Figures 1a and 2a present the results.

*Runtime and Scalability.* We begin by analyzing the non-optimal algorithms, aimed at reducing runtime. Random takes negligible time under all settings ($<$ 3 seconds). Greedy is linear in the number of police cars (it iteratively solves the problem for each police car separately) but exponential in the size of the network. For example, for a network of size 100 with a density of 0.1, 10 police cars and $T = 16$, ROSE takes exactly 1 second to derive an optimal solution while Greedy takes 289 seconds, and produces a suboptimal solution. Greedy reached the timeout for all realistic networks.

Analyzing the Naïve and ROSE conditions head-to-head provides interesting insights. First, in *all* tested networks, ROSE performed faster than Naïve. On average, for small networks, ROSE requires only 19% of the runtime needed by Naïve. We were able to manually engineer circumstances in which Naïve outperforms ROSE, mainly in very small net-

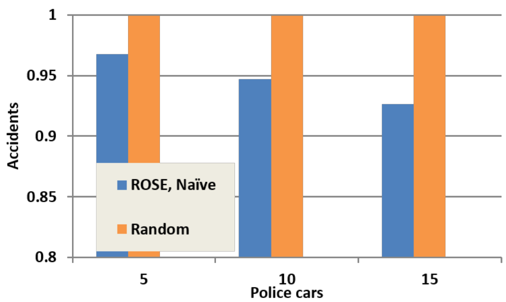


(a) Accidents/Police cars ($T$=8)

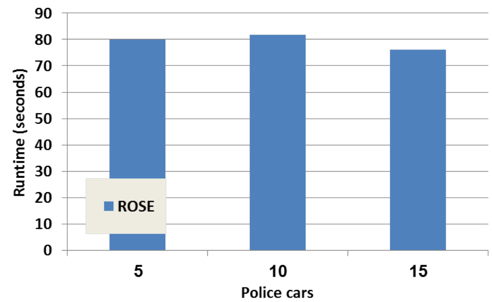


(b) Runtime/Police cars ($T$=24)

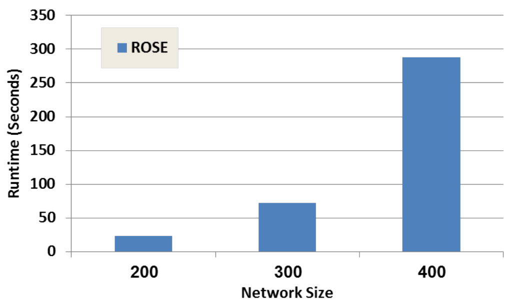


(c) Runtime/Size ($T$=24)

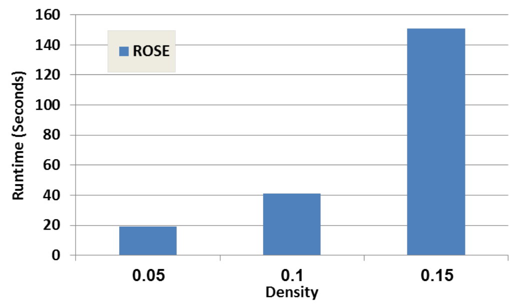


(d) Runtime/Density ($T$=24)

Figure 2: Synthetic road networks: results for the realistic networks. In all Figures, the lower - the better. Note that Naïve and Greedy exceeded the timeout and thus do not appear.

works (size $< 40$) or in networks with a high number of police cars ($> 25$). The runtime difference increases significantly depending on the network's size and density but *slightly* decreases in the number of police cars and the network's density. Similar to Greedy, Naïve was unable to solve most networks of size 200 and all networks of size 300 (and above) in 30 minutes time. See Figures 1b,1c, 1d, 2b, 2c and 2d.

### *4.1.4. Real-World Road Network*

We evaluate ROSE using of the Israeli road network. Unlike for synthetic networks, for the Israeli road network we used the `risk` prediction model available in our simulation environment. $T$ was set to $8, 16$ and $24$, and the number of police cars varied between $5$ and $40$ (in intervals of 5), with 24 settings in total. We also evaluate a Domain Expert condition in which we asked an experienced ITP superintendent who specialized in traffic enforcement to provide an allocation.

The Israeli road network is much larger than the synthetic networks analyzed previously, consisting of 715 road segments, but with a very low density (on average, each intersection connects between 3 and 4 road segments). Therefore, the results display slightly different patterns.

The results show that for cases with 30 or fewer police cars, both in terms of quality and runtime, ROSE outperforms the Naïve, Greedy and Domain Expert conditions by a large margin. Specifically, in these cases, Naïve achieves the same solution quality as ROSE (a 5.5% decrease in the number of expected accidents), but requires up to **6** times longer for runtime. For example, under $T = 16$ and 10 police cars, ROSE requires only 45 seconds compared to almost 4 minutes required by Naïve. However, a transition occurs between 30 and 35 police cars. Specifically, while Naïve and ROSE achieve the same solution quality (averaging an 8% decrease in the number of expected accidents), Naïve favorably compares to ROSE . See Table 2 for the results.

Greedy and Random produced extremely poor solutions across the conditions, averaging less then $1\%$ improvement over the no-enforcement condition. Greedy required a significantly longer runtime than ROSE and reached our timeout of 30 minutes in most cases. As expected, Random required negligible runtime under all settings ($< 2$ seconds). The *Domain Expert* produced allocations where police cars were allocated permanently at the most risky road segments. The quality of the proposed allocation was about $1\%$.

Note that Table 2 further shows that the runtime benefits of ROSE are diminished as the number of police units increases. The reason is simple: with the increase in the number of police cars, penalty terms are more likely to be triggered by feasible solutions and thus more iterations are needed. Next, in the online evaluation of our approach, we demonstrate that the real-world benefits are, in fact, substantial.

### *4.2. Online Evaluation*

Based on the promising results obtained through our lab-based settings (Section 4.1), we began the process of making our approach suitable for real-world deployment. First, several technical issues had to be addressed in order to make our approach suitable for real-world deployment. These include extending the TEAP formulation to include the scheduling of lunch breaks for the officers such that a minimum number of officers are always on duty, allowing officers to "transit" through road segments without enforcing the traffic laws (i.e., bypassing our restriction to a single police cruiser at each road segment and time), etc. As these modifications are mostly technical in nature and do not alter our underling approach, solution technique and theoretical properties, they are reported in Appendix C. Next, we discuss the design of our controlled field experiment and its evaluation metrics. We then present and analyze the results.

#### *4.2.1. Setup*

After attaining the approval of the ITP's commander of the traffic police forces (in the rank of a Major General), we began a controlled field experiment for a period of four months from January up to and including April 2019. The ITP has

| Police Cars | $T = 8$ ROSE | $T = 8$ Naïve | $T = 16$ ROSE | $T = 16$ Naïve | $T = 24$ ROSE | $T = 24$ Naïve |
|---|---|---|---|---|---|---|
| 5 | 5 | 33 | 31 | 153 | 58 | 352 |
| 10 | 7 | 36 | 45 | 191 | 212 | 402 |
| 15 | 11 | 36 | 219 | 301 | 384 | 875 |
| 20 | 12 | 40 | 119 | 263 | 471 | 695 |
| 25 | 21 | 53 | 394 | 487 | 1432 | 1520 |
| 30 | 36 | 40 | 479 | 520 | N/A | N/A |
| 35 | 53 | 51 | 611 | 591 | N/A | N/A |
| 40 | 85 | 53 | 1072 | 836 | N/A | N/A |

Table 2: Runtime of ROSE and Naïve for the Israeli road network with varying numbers of police cars and $T$. Runtime is measured in seconds. N/A means that a timeout of 30 minutes was reached.

chosen the "Shfela"[4] district to test out our system. The Shfela district includes approximately 100 road segments.

During the evaluation period, a 24-hour schedule ($T = 24$) was automatically provided to the ITP every morning around 6am[5] given the number of police units at our disposal and the operational constraints imposed by the ITP on that day[6]. The use of a single schedule for 24 hours is naturally flawed as unexpected delays, operational constraints and changes in manpower may occur during that time, making the pre-calculated schedule infeasible or sub-optimal. In full-scale deployment, one may need to recalculate an allocation for all police cars when such unexpected events occur. However, for this evaluation, the ITP had used ad-hoc local adjustments to follow the schedule "to the best of their ability".

[4]`https://en.wikipedia.org/wiki/Shfela`

[5]At this point in time, the ITP preferred not to integrate our system in their computer network due to security considerations.

[6]The exact number of police cars available on each day and time is withheld at the ITP's request.

Only sporadically, when these local adjustments were too complex to perform, had the ITP requested a new schedule to replace an existing one. These adjustments were not made available to us.

Recall that several technical changes were made to the original TEAP formulation (e.g., the inclusion of lunch breaks for the officers such that a minimum number of officers are always on duty. See Appendix C). As such, we report an additional comparison of ROSE and the Naïve algorithm next in Section 4.2.2.

As in many security settings, it is hard to expect that a comparison over a relatively short time frame will yield a statistically significant difference [27]. In our case, it is unlikely that we will observe a statistically significant decrease *in the number of serious accidents*, due to the fact that road accidents are very sporadic. Therefore, in addition to the evaluation of the number of severe accidents, we define a more subtle metric which is well-known to correlate with our system's main objective - *average speed*. Average speed is strongly associated with traffic safety, influencing both the risk of a traffic accident and the severity of the injury that results from an accident [28, 29]. In order to evaluate the average speed of drivers, we used anonymized cellular reports purchased from Decell Technologies[7] which have been shown to match drivers' average speeds as measured by other (more conventional) instruments in Israel [30].

The use of alternative metrics which are not directly optimized by the system is not unique to this study. For example, in PAWS [31], the authors faced a similar challenge in quantifying the number of saved wildlife due to their provided ranger patrols. The authors used human and animal signs as indicators that PAWS patrols prioritize areas with higher animal and poacher activity. In the same spirit, in this study, we use drivers' average speed as an additional metric.

If our model is capable of successfully prioritizing dangerous road segments and times, it is reasonable to expect it to prioritize road segments and times in which the average speed is high. It is important to stress once more that both the TEAP formulation and the existing ITP practices do not explicitly optimize for police presence in roads segments and times with high average speeds.

#### *4.2.2. Results*

*Runtime.* Recall that some technical changes were made to our initial model (see Appendix C). As such, we compared the ROSE approach to the Naïve algorithm once more in order to quantify its runtime benefit. As was the case in the offline evaluation, ROSE performs more than twice as fast as the Naïve algorithm, averaging less than 5 minutes compared to 11.5 minutes. However, a closer examination of the results demonstrates different patterns compared to the offline evaluation. As shown in Table 3, when the number of police cars is small (i.e., 5 or 10), ROSE outperforms the Naive algorithm by a significant margin. For example, for the case of 5 police cars and $T = 24$, ROSE demonstrates a runtime of 43 seconds

[7] `http://www.decell.com/`

compared to the more than 10 times slower Naïve algorithm (8.5 minutes). However, for cases with many police cars (i.e., 25 and 30), we see that Naïve consistently outperforms ROSE . It is important to recall that the Shfela district consists of *only 100 road segments* (about one seventh of the entire Israeli road network). As such, the bottom half of the table is considered very unrealistic as there is no traffic police force that can cover up to 30% of its road network at any given moment. According to the ITP, most police forces do not have enough police cars and officers to cover more than 10% of any district at any given moment. Averaging over these cases alone in Table 3 shows that ROSE outperforms the Naïve algorithm by averaging 36 seconds compared to about 5 minutes, respectively.

| | $T = 8$ | | $T = 16$ | | $T = 24$ | |
|---|---|---|---|---|---|---|
| Police Cars | ROSE | Naïve | ROSE | Naïve | ROSE | Naïve |
| 5 | 4 | 42 | 25 | 218 | 43 | 509 |
| 10 | 18 | 57 | 31 | 287 | 97 | 619 |
| 15 | 68 | 59 | 328 | 364 | 452 | 615 |
| 20 | 126 | 67 | 359 | 531 | 636 | 931 |
| 25 | 410 | 104 | 625 | 664 | 1831 | 1373 |
| 30 | 861 | 193 | 1707 | 752 | 3805 | 2026 |

Table 3: Runtime of ROSE and Naïve for the Shfela road network with varying numbers of police cars and $T$. Runtime is measured in seconds. N/A means that a timeout of 30 minutes was reached.

Recall that in full-scale deployment, one may need to execute ROSE again and again given unexpected events. As such, the runtime differences between Naïve and ROSE may prove very substantial in practice.

*Accidents.* We first analyze the number of serious accidents that occurred in the Shfela district during the evaluation period. For comparison, we consider the adjacent "HaSharon"[8] district which is highly similar in its size and land use, with many

[8] https://en.wikipedia.org/wiki/Sharon_plain

roads spanning across the two districts and the general trend in all other parts of Israel. To that end, we use official data published by the Israeli CBS. Prior to the deployment of our system, the number of serious accidents in the Shfela district decreased from 84 in January-April of 2017 to 77 in the same period of 2018 (8% decrease). In the same period in 2019, while using our system, 65 serious accidents were reported (15.5% decrease from 2018). During the same time frame, the number of serious accidents in the HaSharon district and in all other parts of Israel decreased by 8% and 7% from 2017 to 2018, and by an additional 9.5% and 11% from 2018 to 2019, respectively. The Shfela district has demonstrated the sharpest decrease in the number of serious accidents among all Israeli districts. See Figure 3 for a month-by-month analysis. Part of the decrease in the number of accidents in Israel (and in the Shfela and HaSharon districts in particular) is naturally attributed to the increased prevalence of automotive safety measures such as MobileEye[9] and to the increased number of vehicles on the roads which translate into slower average driving speeds. To our knowledge, no other significant traffic-related changes have occurred during the above-mentioned period, specifically enforcement-related changes.

A month-by-month analysis reveals some sharp changes in the number of accidents in January and February 2019 in the Shfela district (a decrease of more than 30%) while only a modest decrease is recorded for the other Israeli districts (less than 10%). Inconsistency in the number of serious accidents is also demonstrated in the HaSharon district. On the other hand, the change in the number of serious accidents in all other Israeli districts displays low variance (maximum of 13% and minimum of 7.5%). The high variance in both districts is mostly attributed to the sporadic nature of accidents while the low variance in the Israeli measurement is due to the aggregative calculation of accidents outside the Shfela distinct which "balances" the differences encountered in each individual district.

*Average Speed.* To provide some additional insights into the evaluation, we take a closer look at the average speed on road segments and times which were enforced by officers using our system and compare it to the average speed on roads enforced by officers not using our system in the HaSharon district. Using a standard $t$-test, we cannot reject the null hypothesis that the average speed limits of the two districts are the same ($p = 0.99$), nor can we reject the null hypothesis that the average of speeds across the two districts are the same ($p = 0.96$), thus supporting our hypothesis that the two districts are indeed similar.

On average, police officers in the Shfela district were allocated to roads and times with an average speed 8% higher as compared to the locations and times that police officers in the adjutant district were allocated. The difference is statistically significant using a standard $t$-test ($p < 0.05$). Namely, police officers in the Shfela district (who use our system) were allocated to road segments and times which are assumed to be more dangerous than the ones in which the officers from

[9] `http://www.mobileye.com`

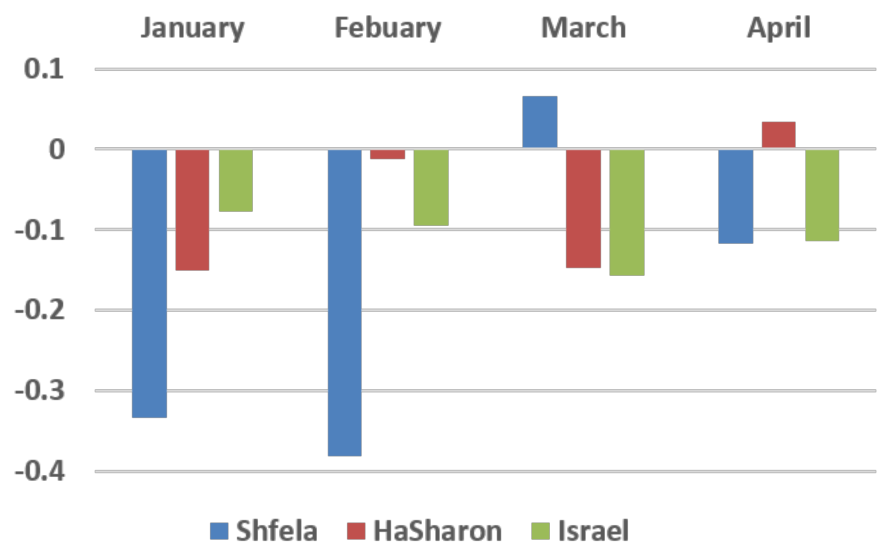


Figure 3: A month-by-month comparison of the number of serious accidents between 2018 and 2019.

the adjutant district were allocated (using standard police practices), assuming that police presence does not increase the average speed. Under the assumption that the presence of a police cruiser reduces the average speed on a road segment at a certain time to a similar extent (e.g., reduce it by 20%) this would mean that the absolute reduction in average speed is greater in the Shfela district. The above assumption seems reasonable as the two districts are very similar and drivers cannot distinguish between officers who use our system and those who do not.

It is, however, important to note that the two districts do not significantly differ in the average speed *across all segments* even after using our system ($p = 0.94$). The reason is simple: the ITP's limited resources can affect the average speed for only a (very) small subset of road segments. In turn, these changes do not translate into statistically significant differences at the entire district level.

## 5. Discussion

The results of our evaluation (Section 4), both in the lab and in real-world deployment settings, show that our approach can successfully address the computational challenges associated with traffic enforcement (compared to the competing optimization approaches) and outperform existing police practices. The latter is measured in terms of both the number of accidents and the average speeds in the deployment site compared to the control condition.

When presented with the results, ITP officials stated that they were "very happy" with the system and the outcomes. They are currently considering deployment in additional districts.

However, when presenting a new formulation such as TEAP, new solution techniques such as ROSE, and an in-the-field evaluation of a deployed system, it is worth discussing limitations.

ROSE allows us to optimally solve large TEAPs with significant runtime improvement compared to baseline approaches. This improvement is most significant for large, dense networks. However, ROSE's runtime is impaired with the increase in the number of police units. The reason is simple: with the increase in the number of police cars, penalty terms are more likely to be triggered by feasible solutions. Therefore, in a "congested" TEAP (i.e., a small network with many police cars), ROSE could be counterproductive. According to ITP experts, traffic police worldwide use a network size – police car ratio similar to the one deployed in Israel. Therefore, in real-world deployment in other countries, one is most likely to encounter large networks with a relatively low number of police resources, like the settings investigated in this article.

The TEAP solution is a pure strategy for the police, which makes *predictability* an issue. Unlike various other security models such as adversarial robotic patrolling (e.g., [32, 33, 34]), in this article TEAP assumes that the drivers are reactive to police presence and essentially do not learn the police's actual policy. This assumption may lead to repetitive police allocations which drivers may (eventually) understand and anticipate. A possible indication for the occurrence of such a phenomena in practice can be seen in Figure 3. The benefit from our system in the first half of the evaluation period was much larger than the benefit in the second half. It is, however, possible that factors of which we are unaware have caused this discordance, such as operational constraints or natural noise in the data. All the same, a practical solution to this concern is to (periodically) define additional allocation constraints that impose or restrict the enforcement of a specific road segment, similar to the entropy-based approach suggested in [11]. Today, police forces occasionally define road segments that must or must not be visited during a shift due to special enforcement needs (e.g., road work). The injection of these constraints in the TEAP formulation is straightforward, yet the injection itself was not performed in our online evaluation following the ITP's request. An automatic process may also randomly select which road segments must/must not be enforced in a given allocation such that every road segment has at least a user-defined $\epsilon$ probability of being enforced at every time step. Note that the adoption of SSG for preventing serious road accidents seems unsuitable as it can be argued that drivers follow an opportunistic behavioral model rather than a strategic one.

A common challenge to many human-centered problems and systems, such as TEAP, is the efficient adaption to human-driven changes in the environment. For example, adapting to a human's changing preferences or abilities (e.g., [35]). In traffic enforcement, this challenge may manifest itself as a police cruiser being delayed, which might make the proposed allocation undesirable or infeasible. An efficient way to resolve this issue is for central command to allocate the police cars, assuming perfect execution. Only after a non-default transition occurs does the central command resolve the TEAP, starting from the current state [16]. Given the positive runtime results of ROSE, especially the ones demonstrated in our online evaluation (Section 4.2.2), such reallocation should not pose a significant computational concern and the former's runtime advantage accrues over time.

As with any in-the-field experiment, one is limited by the data he or she is exposed to and by the natural noise it contains. In our case, as discussed before, the number of accidents and average speeds recorded during our evaluation period should be partially attributed to additional factors which we cannot control and observe. For example, unexpected operational constraints may have influenced the ITP's ability to follow our schedule (e.g., a police officer took an unexpected sick day). Similarly, special circumstances such as sports events, road work, political events[10] and others may have influenced both the occurrence of accidents as well as the recorded average speeds. Unfortunately, the ITP does not record most of these events in their systems and they are unable to report them to us (both in real time and in retrospect). In addition, considering 2018 as a baseline introduces additional noise as we have little information on traffic-related circumstances other than the fact that the number of police officers has not changed. For example, In March 2019, a slight increase in the number of accidents was recorded.

## 6. Related Notions in Traffic Enforcement

It has been established that a significant reduction in the occurrence of serious traffic accidents can be achieved by efficient traffic police allocation [2]. Specifically, efficient traffic enforcement has been shown to reduce a wide range of high-risk, illegal driving behaviors, including driving while under the influence of drugs/alcohol, speeding, lack of seat belt use and red-light running, and thus reduces traffic accidents (e.g., [5, 6]). Therefore, recently, traffic police forces have begun implementing the predictive policing paradigm [36] through which police officers can identify people and locations at increased risk. From a methodological standpoint, the effort of predicting traffic accidents has mainly focused on aggregative analysis, specifically on the prediction of the *annual* number of serious accidents per road segment using statistical methods such as Poisson or negative binomial regression models [37]. Such aggregation is limited in its use to police forces as the allocation of traffic police enforcement requires a prediction on a much more finely-grained level. To the best of our knowledge, the state-of-the-art prediction models provide prediction for three hour time-frames. Overall, despite its promise and successful implementation, predictive policing does not provide police officers with a means to derive optimal enforcement allocations. In this study, we were able to construct a prediction model that provides beneficial

[10] General elections took place on April 4$^{\text{th}}$.

predictions for *one hour time-frames* by using a unique set of features and 11 years of collected data.

The Gambler's Fallacy is the phenomenon where people tend to put ample weight on previous events, believing that they influence future outcomes. This phenomenon manifests itself in the context of traffic enforcement in the form of *halo effects*. For over 4 decades traffic halo effects have been validated repeatedly, showing that enforcement effects are not restricted to the specific time and space in which the enforcement is performed. Two such effects are called time-halo and distance-halo [13]. To our knowledge, this is the first work to formulate and integrate halo effects in enforcement optimization. Existing works on modeling human behavior in SG settings such as [38, 39, 40] consider the adaptive nature of human behavior to successes and failures in past rounds. However, the integration of halo effects in such models is not straightforward.

## 7. Conclusions

This article introduces a novel framework for optimizing traffic police allocation in real-world settings. First, we model the interaction between drivers and traffic police as a Traffic Enforcement Allocation Problem (TEAP) and prove that accurately solving or approximating the optimal solution of a TEAP is hard. Next, we cast the TEAP as a binary graph flow problem, which in turn is translated into a unique binary optimization problem, and we show how to solve it efficiently and optimally by a new algorithm called the RELAXED OPTIMIZATION SOLVER ENHANCER, ROSE. Extensive empirical evaluation, both in lab-based settings and in a controlled field experiment in Israel, demonstrates the benefits of our approach and its applicability.

We hope that this study will encourage other researchers to tackle the important and challenging task of preventing serious traffic accidents. To assist others with this challenge, we also provide a realistic simulation environment, which we name *SECURE*, that includes a state-of-the-art accident prediction model along with useful road networks and data.

In future work, we intend to extend our investigation in two ways: First, we wish to study the use of police resources which vary in their capabilities. For example, breathalyzers are often used by police officers to detect the blood alcohol content of drivers. The decision of which officers should be equipped with breathalyzers and when has yet to be captured in our proposed model and, to the best of our knowledge, has yet to be captured by any optimization-based model. A similar challenge of integrating traffic enforcement drones within our framework was recently addressed in [41]. Second, we seek to investigate additional security settings in which the submodularity assumption does not hold, yet other structural properties, such as supermodularity, may prove useful.

## Acknowledgments

This article extends two previous reports: 1) [42] from the 2017 International Joint Conference on Artificial Intelligence (IJCAI); and 2) [43] from the 2017 International Conference on Decision and Game Theory for Security (GameSec). The former paper discusses the main components of our model and solution technique while focusing on lab-based synthetic instances. The latter paper then extends the initial model and solution technique to properly account for some of the (then expected) real-world deployment challenges. In this article we elaborate and extend the two papers in two major aspects: First, we examine the real-world deployment of our approach through an in-the-field four-month long study with the ITP (see Section 4). The results clearly demonstrate the potential impact of our approach in the real-world as well as support and strengthen our results, which were obtained in prior lab-based settings. Second, we provide full details of both the development and deployment process of our model, solution technique and deployed system which will, hopefully, allow traffic enforcement researchers and practitioners to replicate and adapt our approach to additional deployment sites across the globe. Additional explanations, more detailed descriptions of our assumptions, model, solution techniques and evaluation and more thorough examples are now also provided.

This work was supported in part by the LAW-TRAIN project which received funding from the European Union Horizon 2020 research and innovation program under grant agreement 653587 and in part by a grant from the Ministry of Science & Technology, Israel & the Japan Science and Technology Agency (jst), Japan.

## Appendix A. Modeling `risk`

We characterize each road segment in time using a unique set of 117 features. The features are divided into 3 categories: 1) infrastructure features; 2) date and time features; and 3) traffic features. To the best of our knowledge, this is the largest set of features ever to be used to predict severe car accidents.

*Infrastructure Features*

The geography of Israel is very diverse, with desert conditions in the south and snow-capped mountains in the north. It is customary to divide Israel into 3 regions: North, South and Center. These three regions differ significantly in their population and land use. For example, the central region is a metropolitan area (e.g, the Tel-Aviv metropolis) characterized by dense urban building and high-tech land use, whereas the southern region is mostly a desert which for the most part consists of rural low-density residential areas [3 features]. The ITP further divides Israel into 15 districts according to geographic criteria [15 features].

Each road segment is characterized according to its type (e.g., highway) [7 features], its length in KM [1 feature], the number of lanes [7 features], the posted speed limit [5 features], road signals [2 features], road width [5 features], whether a traffic light is present on the road segment [2 feature], road surface conditions (e.g., gravel/paved) [6 features] and whether the road is lit up at night [5 features]. Unfortunately, to date, we were unable to obtain additional features that have been shown to affect the prevalence of road accidents in past literature. These features include the existence of road shoulders, the road segment's curvature, incline/decline etc.

*Date and Time Characteristics*

We characterize the date using the month of the year [12 features], day of the week [7 features] and an indicator whether it is a weekday, weekend, holiday, holiday evening or another type of special day [5 features]. Time is characterized on an hourly scale [24 features] and by an indicator of whether it is daytime or nighttime [2 features].

In addition, we characterize the weather in the vicinity of the road segment at the given time using the publicly available IMS reports and forecasts [4 features].

*Traffic Characteristics*

While the infrastructure characteristics do not change frequently, the traffic that goes through the road segments changes rapidly over time. We characterize the traffic by its volume [1 feature]. Traffic volume is provided by the CBS and average speeds are provided by the ITP. We further identify the number of severe accidents which have occurred on that road segment in the prior 30, 90, 180 and 365 days [8 features].

### *Appendix A.1. Training a Deep Neural Network*

Our network consists of 3 layers, 1024x512x1, where the hidden layer uses the common RelU activation function. Several other architectures were tested and found to be of lower quality in terms of AUC.

## Appendix B. Linearization Technique

Let us assume a product term of $n$ binary variables denoted $x_1 \cdots x_n$.

Define a new binary variable $z$ which will represent the product term using the following $n+1$ constraints:

$$z \leq x_i \text{ for } i = 1, \ldots, n$$

$$z \geq \sum_{1}^{n} x_i - (n-1)$$

It is easy to verify that $z$ will be forced to take the value of 0 if at least one $x_i = 0$ (first $n$ constraints) and that it must take the value of 1 otherwise (last constraint).

## Appendix C. From the Lab to the Roads

A few steps needed to be taken before our approach could be deployed in the field.

*Security Clearance.* Before any meaningful intersection with the ITP could take place (e.g., allowing us access to their confidential data), the first two authors had to obtain security clearance, including a 2-hour background check and an interview at the ITP headquarters. The clearance came through about 6 months into the process.

*Adding Transit Edges.* The ITP has requested the addition of transit commands to their schedule. Namely, in addition to directing police cruisers to enforce different road segments, they have requested that we explicitly model the option of a police car traveling through a road segment without enforcing the law. To that end, when time-extending the road network $G$ such that each vertex $v$ is replicated $T$ times, two types of edges are added for each $t < T$ to the transition graph. Specifically, transit edges are added from each vertex $u_t$ to $v_{t+l(e,t)}$ where $l(e, t)$ is the estimated travel time to cross $e$ at time $t$ according to Google Maps (`https://maps.google.com`).[11] Unfortunately, the above does not suffice. Specifically, the TEAP's formulation relies on the assumption that no two police cars should enforce the same road segment at the same time. However, this rule does not necessarily apply to *transit* actions, where more than one police car can be present on the same road and at the same time. We investigated this issue empirically; first we duplicated each transit edge by the number of police cars available. Practically, under various conditions, we did not encounter any realistic settings in which more than a single police car was present on the same road segment at the same time in Israel.

*Logistics.* According to the ITP, during an 8-hour shift, each police car should have a break of about 1 hour to eat and reach its next destination. The rationale is that the ITP has arranged various different places for police officers to eat and therefore no special requirements should be implemented as to where a police car should have its break. This break is scheduled for different times, for example, interleaving during the $4^{th}$ hour of work so as to avoid having all officers on break at the same time. Specifically, officers are interleaved as to when they would go on a break during the $4^{th}$ hour of work such that at least $k$ police cars are not on break at any given moment ($k$ is a police defined constant). We amend our model by adding designated "break" vertices during the $4^{th}$ hour. These vertices are accessible from any vertex during the $4^{th}$ hour and are connected to all vertices which are one hour later. For example, a police car can go on a break from any location at 12:00, and continue its schedule from any vertex at 13:00. This formulation was specifically tailored at the request of the ITP. To make sure each police car goes on a single break, nodes during the $4^{th}$ hour were duplicated such that every node had two copies – "pre-break" and "post-break". Then, pre-break nodes were disconnected from $5^{th}$ hour

[11] Time was discretized in 10 minute time-frames.

nodes and post-break nodes were disconnected such that they are only accessible from break nodes or other post-break nodes. Simply put, a police crusier can only reach the $5^{th}$ hour of the shift if it goes though a post-break node. Naturally, the post-break nodes do not allow re-access to a break node, ensuring that each police car visits only a single break node on its path.

*Non-default transition.* Note that, given a non-default transition, we recalculate the allocation for all police cars, as local adjustments may produce suboptimal allocations. We plan to investigate local methods for adjusting infeasible or undesired allocations in future work.

## References


[1] World Health Organization. Road traffic injuries fact sheet. Technical report, December 2018. URL `https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries`.

[2] Rune Elvik, Truls Vaa, Alena Erke, and Michael Sorensen. *The handbook of road safety measures*. Emerald Group Publishing, 2009.

[3] Edna Schechtman, Hillel Bar-Gera, and Oren Musicant. Driver views on speed and enforcement. *Accident Analysis & Prevention*, 89:9–21, 2016.

[4] Nicole Adler, Alfred Shalom Hakkert, Jonathan Kornbluth, Tal Raviv, and Mali Sher. Location-allocation models for traffic police patrol vehicles on an interurban network. *Annals of Operations Research*, 221:9–31, 2014.

[5] Lyndel Bates, David Soole, and Barry Watson. The effectiveness of traffic policing in reducing traffic crashes. In *Policing and security in practice*, pages 90–109. Springer, 2012.

[6] Jenna K Simandl, Andrew J Graettinger, Randy K Smith, Steven Jones, and Timothy E Barnett. Making use of big data to evaluate the effectiveness of selective law enforcement in reducing crashes. *Transportation Research Record: Journal of the Transportation Research Board*, (2584):8–15, 2016.

[7] José I Castillo-Manzano, Mercedes Castro-Nuño, Lourdes López-Valpuesta, and Diego J Pedregal. From legislation to compliance: The power of traffic law enforcement for the case study of spain. *Transport policy*, 75:1–9, 2019.

[8] Gregory DeAngelo and Benjamin Hansen. Life and death in the fast lane: Police enforcement and traffic fatalities. *American Economic Journal: Economic Policy*, 6(2):231–257, 2014.

[9] Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, and Sarit Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *Proceedings of the 2008 International Conference on Autonomous Agents & Multiagent Systems*, pages 895–902, 2008.

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

[11] Matthew Brown, Sandhya Saisubramanian, Pradeep Reddy Varakantham, and Milind Tambe. Streets: game-theoretic traffic patrolling with exploration and exploitation. In *Proceedings of the 2014 AAAI Conference on Artificial Intelligence*, 2014.

[12] European Transport Safety Council. How traffic law enforcement can contribute to safer roads: Pin flash report 31. Technical report, June 2016. URL `http://etsc.eu/wp-content/uploads/NEW_PIN_FLASH31_final.pdf/`.

[13] MA Elliott, Jeremy Broughton, et al. *How methods and levels of policing affect road casualty rates*. Transport Research Laboratory, 2005.

[14] Chao Zhang, Shahrzad Gholami, Debarun Kar, Arunesh Sinha, Manish Jain, Ripple Goyal, and Milind Tambe. Keeping pace with criminals: An extended study of designing patrol allocation against adaptive opportunistic criminals. *Games*, 7(3):15, 2016.

[15] Nicola Basilico, Nicola Gatti, and Francesco Amigoni. Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proceedings of the 2009 International Conference on Autonomous Agents & Multiagent Systems*, pages 57–64, 2009.

[16] Francesco Maria Delle Fave, Albert Xin Jiang, Zhengyu Yin, Chao Zhang, Milind Tambe, Sarit Kraus, and John P Sullivan. Game-theoretic patrolling with dynamic execution uncertainty and a case study on a real transit system. *Journal of Artificial Intelligence Research*, 50:321–367, 2014.

[17] Francesco Maria Delle Fave, Eric Shieh, Manish Jain, Albert Xin Jiang, Heather Rosoff, Milind Tambe, and John P Sullivan. Efficient solutions for joint activity based security games: fast algorithms, results and a field experiment on a transit system. *Autonomous Agents and Multi-Agent Systems*, 29(5):787–820, 2015.

[18] 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 2012 International Conference on Autonomous Agents & Multiagent Systems*, pages 13–20, 2012.

[19] Zhengyu Yin, Albert Xin Jiang, Matthew Paul Johnson, Christopher Kiekintveld, Kevin Leyton-Brown, Tuomas Sandholm, Milind Tambe, and John P Sullivan. Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *Proceedings of the 2012 Innovative Applications of Artificial Intelligence*, 2012.

[20] Marco E Lübbecke and Jacques Desrosiers. Selected topics in column generation. *Operations research*, 53(6): 1007–1023, 2005.

[21] Nitesh V Chawla. Data mining for imbalanced datasets: An overview. In *Data mining and knowledge discovery handbook*, pages 853–867. Springer, 2005.

[22] Sarit Weisburd. Does police presence reduce car accidents? Technical report, Auguest 2016.

[23] Ariel Rosenfeld and Sarit Kraus. Predicting human decision-making: From prediction to action. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(1):1–150, 2018.

[24] Rong Yang, Albert Xin Jiang, Milind Tambe, and Fernando Ordonez. Scaling-up security games with boundedly rational adversaries: A cutting-plane approach. In *Proceedings of the 2013 International Joint Conference on Artificial Intelligence*, 2013.

[25] Paul Black. Dictionary of Algorithms and Data Structures, U.S. National Institute of Standards and Technology, `https://xlinux.nist.gov/dads/`, 2005.

[26] Inc. Gurobi Optimization. Gurobi optimizer manual, 2016. URL `http://www.gurobi.com`.

[27] Matthew E Taylor, Christopher Kiekintveld, Craig Western, and Milind Tambe. A framework for evaluating deployed security systems: Is there a chink in your ARMOR? *Informatica*, 34(2), 2010.

[28] Göran Nilsson. *Traffic safety dimensions and the power model to describe the effect of speed on safety*. PhD thesis, Lund Institute of Technology, Department of Technology and Society, Traffic Engineering, Lund, Sweden., 2004.

[29] Rune Elvik. A re-parameterisation of the power model of the relationship between the speed of traffic and the number of accidents and accident victims. *Accident Analysis & Prevention*, 50:854–860, 2013.

[30] Hillel Bar-Gera. Evaluation of a cellular phone-based system for measurements of traffic speeds and travel times: A case study from israel. *Transportation Research Part C: Emerging Technologies*, 15(6):380–391, 2007.

[31] 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, 2014.

[32] Noa Agmon, Gal A Kaminka, and Sarit Kraus. Multi-robot adversarial patrolling: facing a full-knowledge opponent. *Journal of Artificial Intelligence Research*, 42:887–916, 2011.

[33] Efrat Sless, Noa Agmon, and Sarit Kraus. Multi-robot adversarial patrolling: facing coordinated attacks. In *Proceedings of the 2014 International Conference on Autonomous Agents & Multiagent Systems*, pages 1093–1100, 2014.

[34] Noa Agmon, Sarit Kraus, Gal A Kaminka, and Vladimir Sadov. Adversarial uncertainty in multi-robot patrol. In *Proceedings of the 2009 International Joint Conference on Artificial Intelligencee*, pages 1811–1817, 2009.

[35] Ariel Rosenfeld, Moshe Cohen, Sarit Kraus, and Joseph Keshet. Online prediction of time series with assumed behavior. *Engineering Applications of Artificial Intelligence*, 88, 2020.

[36] Walt L Perry. *Predictive policing: The role of crime forecasting in law enforcement operations*. Rand Corporation, 2013.

[37] Li-Yen Chang. Analysis of freeway accident frequencies: negative binomial regression versus artificial neural network. *Safety science*, 43(8):541–557, 2005.

[38] Raz Lin, Sarit Kraus, Noa Agmon, Samuel Barrett, and Peter Stone. Comparing agents' success against people in security domains. In *Proceedings of the 2011 AAAI Conference on Artificial Intelligence*, pages 809–814, 2011.

[39] Debarun Kar, Fei Fang, Francesco Delle Fave, Nicole Sintov, and Milind Tambe. A game of thrones: when human behavior models compete in repeated stackelberg security games. In *Proceedings of the 2015 International Conference on Autonomous Agents & Multiagent Systems*, pages 1381–1390, 2015.

[40] Fei Fang, Peter Stone, and Milind Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *Proceedings of the 2015 international joint conference on Artificial Intelligence*, 2015.

[41] Ariel Rosenfeld and Oleg Maksimov. Optimal cruiser-drone traffic enforcement under energy limitation. *Artificial Intelligence*, 277:103–166, 2019.

[42] Ariel Rosenfeld and Sarit Kraus. When security games hit traffic: Optimal traffic enforcement under one sided uncertainty. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3814–3822, 2017. doi: 10.24963/ijcai.2017/533.

[43] Ariel Rosenfeld, Oleg Maksimov, and Sarit Kraus. Optimizing traffic enforcement: From the lab to the roads. In *International Conference on Decision and Game Theory for Security*, pages 3–20. Springer, 2017.