

Optimizing Traffic Enforcement: From the Lab to the Roads

Ariel Rosenfeld^{1,†}, Oleg Maksimov² and Sarit Kraus²

¹ Dept. of Computer Science and Applied Mathematics,
Weizmann Institute of Science, Rehovot, Israel.

² Dept. of Computer Science, Bar-Ilan University, Ramat-Gan, Israel.

[†] Corresponding author, arielros1@gmail.com

Abstract. Road accidents are the leading causes of death of youths and young adults worldwide. Efficient traffic enforcement has been conclusively shown to reduce high-risk driving behaviors and thus reduce accidents. Today, traffic police departments use simplified methods for their resource allocation (heuristics, accident hotspots, etc.). To address this potential shortcoming, in [23], we introduced a novel algorithmic solution, based on efficient optimization of the allocation of police resources, which relies on the prediction of accidents. This prediction can also be used for raising public awareness regarding road accidents. However, significant challenges arise when instantiating the proposed solution in real-world security settings. This paper reports on three main challenges: 1) Data-centric challenges; 2) Police-deployment challenges; and 3) Challenges in raising public awareness. We mainly focus on the data-centric challenge, highlighting the data collection and analysis, and provide a detailed description of how we tackled the challenge of predicting the likelihood of road accidents. We further outline the other two challenges, providing appropriate technical and methodological solutions including an open-access application for making our prediction model accessible to the public.

1 Introduction

Every year the lives of approximately 1.25 million people are cut short and between 20 and 50 million people suffer disability or other severe injuries as a result of severe road accidents (accidents that cause death or injury) [29]. Efficient traffic enforcement can reduce the number and severity of severe road accidents by giving drivers the feeling that they are likely to be caught and sanctioned when breaking the law [11]. Road safety agencies have already identified the need for improvement in traffic enforcement and it is now an integral part of many countries' road safety policies [12]. Traffic police resources cannot cover the entire road network given the limited number of police cars and officers [7], and therefore some allocation mechanism is needed.

To address this challenge, we introduced the Traffic Enforcement Allocation Problem (TEAP) in a previous work [23]. TEAP is represented as an optimization problem which is shown to be NPH for approximation within any constant factor. TEAP relies on a pre-defined road network which is associated with two functions: 1) a function for measuring the likelihood that a severe traffic accident will occur on any road segment at

any time; 2) a function for measuring the effect that the police allocation (both past and present) has on the risk of accidents occurring on any road at any time. Despite its computational complexity, realistically sized TEAPs can be solved efficiently using a newly proposed relaxed optimization technique named ROSE which leverages the TEAPs' characteristics. Nevertheless, despite its appealing theoretical properties and its practical promise, several challenges need to be addressed before the proposed model and solution can be applied in real-world traffic enforcement settings.

This paper presents the challenges that arise from our current effort to apply the TEAP model and its proposed optimization solution for the Israeli Traffic Police (ITP). We focus on three main challenges: First, the TEAP assumes that the road network and all required models which are associated with the road network are pre-defined and given. Namely, the TEAP assumes that the road network has already been divided into equally sized road segments, that the risk that a severe traffic accident will occur on any road at any time is known and that the effect of police enforcement on the latter is given. We refer to the challenge of defining and estimating the above as the **Data-centric challenge**. Second, several logistical/technical issues need to be addressed. For example, during the morning shift, the traffic allocation needs to accommodate a lunch break for the police officers. We refer to these challenges as the **Police-deployment challenges**. Lastly, data was gathered during the design and development process. However, most of the gathered data is confidential and thus cannot be released to the public. The public may benefit from making some of this data available in some format which will not jeopardize confidential materiel. We refer to these challenges as **Challenges in raising public awareness**.

In addition to our discussion on the above challenges and the provided technical and methodological solutions, this paper also provides a “behind the scenes” view of the process of moving from a theoretical model, tested in lab-settings, to an actual deployed system in field trials on the roads.

2 Related Work and Background

Recent studies suggest that drivers respect traffic laws mainly due to enforcement concerns, rather than safety concerns (e.g., [24]). 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., [1,3,25]).

Within the Security Games (SG) field, optimal security resource allocation mechanisms and applications for mitigating various types of crimes have been developed. The generic SG framework consists of a defender (traffic police) which has a limited number of resources (police cars) to protect a large set of targets (road segments) from an adversary (reckless drivers). This approach has led to a variety of successfully deployed applications for the security of infrastructure and wildlife [26].

Other than our previous work [23], Brown et al. [4] 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 that drivers can first observe and then use in order to best respond. However, several practical issues make existing SG-based solutions seem unsuitable for the task of mitigating severe car accidents. First, traffic enforcement seeks to reduce road accidents (and not necessarily to apprehend reckless drivers) [12]. Second, over 4 decades of traffic enforcement literature shows that drivers are acting in a less strategic manner, *reacting* to observed police presence in the past and the present as well as on the current road and other roads in the vicinity [9]. While non-strategical adversaries in SG settings have recently been addressed (e.g., modeled as opportunistic criminals [32]), existing solutions do not account for the above challenges when combined. If we translate these seemingly technical issues into theoretical ones, we can identify that traffic enforcement is in fact a *non-Markovian* and *coupled* allocation problem. Namely, past police actions and states influence future states (non-Markovian), and police cars should coordinate their actions (coupled).

The TEAP model, which we propose in [23], addresses the above issues. TEAP leverages on the fact that, according to police enforcement experts, if police cars are stationed at the same place and time, their effectiveness in reducing road accidents cannot be assumed to be greater than the effectiveness of a single police car at the same point and time. In a TEAP, the interaction between drivers and police is modeled 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. The graph is then extended into a *transition graph* such that each vertex v (edge e) is replicated T times, one for each round, denoted v_t (e_t), assuming that it takes one unit of time to traverse each road (see [31] for an extended discussion of the use of transition graphs). 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} . The b_{v_t} values are assumed to be known in advance and cannot be changed by the police. Every road segment and time, e_t , is associated with an indicator $H[e_t]$ which assumes the value of 1 if a police resource is allocated to e_t . The allocation history of the police resources until time t (including) is denoted as H_t . The *risk of accidents* occurring at e_t is denoted as $\text{risk}(e_t)$. The *risk* function measures the likelihood that a severe traffic accident will occur at e_t in the absence of police enforcement (in the $[0,1]$ range). The *effectiveness of enforcement* is denoted as $\text{eff}(H_t, e_t)$. eff measures the effect that the police allocation history has on the risk of accidents occurring at e_t . Consequently, the TEAP is formulated as the following mathematical program:

$$\min_{H_T} \sum_t \sum_{e_t} \text{risk}(e_t) \cdot (1 - \text{eff}(e_t, H_t)) \quad (1)$$

$$\text{s.t.} \sum_{v'_{t-1}} H_t[(v'_{t-1}, v_t)_{t-1}] - \sum_{v'_{t+1}} H_{t+1}[(v_t, v'_{t+1})_{t+1}] = b_{v_t} \quad \forall v_t \in G_T \quad (2)$$

$$H_T[e_t] \in \{0, 1\} \quad \forall e, t \quad (3)$$

Complete details and source code, including a master-slave optimization technique for solving the TEAP, are presented and evaluated in [23]. We refer the reader to the original paper for a thorough discussion about the TEAP’s benefits and limitations.

Generally speaking, security settings are often very dynamic and complex, which may make it impractical to capture all of the necessary characteristics of the designated domain in a general game model built in a lab [21]. As a result, when moving from a theoretical model, tested in lab-settings, to an actual deployed system in field trials, different challenges arise. Therefore, despite the evidence showing the benefits of the above TEAP formulation, several data and logistical challenges prevent it from being implemented in the field as an “off-the-shelf” product. Furthermore, some of the assumptions made in the original formulation do not hold in practice, which necessitates the modification of the proposed formulation.

It is common to consider ARMOR as the first deployed SG system [22]. The system was deployed at the Los Angeles International Airport (LAX) in 2007 in order to randomize checkpoints on the roadways entering the airport as well as canine patrol routes within the airport terminals. For its deployment, the authors faced different challenges, mainly in instantiating their model to the LAX environment and increasing organizational acceptance of the proposed solution [16]. For example, the defender’s payoffs were hard-coded after a series of interviews with airport security personnel. In addition, in order to allow security personnel the needed flexibility to adjust and change the provided allocation, an “override” option was added to the system. These insights are integrated in this work as well. Note that ARMOR provides a static allocation of security resources, which does not account for the spatio-temporal aspects of traffic enforcement.

A more similar system to ours, which also requires transition-based scheduling, is the TRUSTS system which is designed for fare-evasion deterrence in urban transit systems [31]. Similar to traffic enforcement settings, the TRUSTS system also allows defender resources to move across a graph structure over time. In deploying TRUSTS, several issues had to be addressed, with the prominent one being execution uncertainty [8]. In real world trials, a significant fraction of the executions of the pre-generated schedules got interrupted for a variety of reasons (e.g., writing citations, felony arrests, etc.), causing the officers to miss the train that they were supposed to take. Despite the resemblances between the two domains, exact timing is of far lesser significance in traffic enforcement as temporal constraints are more flexible than in transit system enforcement. Specifically, traffic enforcement officers are not bound to a fixed train schedule and thus delays have less far-reaching consequences on a planned schedule. According to ITP experts, traffic police schedules are usually macro-managed, for example, enforcement of a road segment is scheduled in hours and not minutes, thus officers can better plan their actions and adapt to changes in real time. Therefore, delays do not pose a big concern in our setting. Several parameters of the TRUSTS formulation are estimated from available data such as ridership of different trains. Other parameters, which are more complex to estimate given available data, are estimated using non-data-driven methods (e.g., the uncertainty parameter). Similarly, in this work, most parameters are learned directly from Israeli-based data such as the risk of accidents occurring at each road segment and time which are estimated based on 11 years of collected data. Other

parameters, such as the effectiveness of enforcement efforts which are much harder to learn, are estimated using other sources (in our case, past literature).

To the best of our knowledge, the most recently deployed application is PAWS [30]. PAWS is designed to combat illegal poaching through the optimization of human patrol resources. However, initial field tests of PAWS have revealed several data-centric and deployment-centric challenges [13]. For example, it turns out that the PAWS grid-based model does not capture important factors which may hinder the quality of the provided allocation and even prevents patrollers from completing their tasks, such as terrain elevation and accessibility. Furthermore, in the original model, PAWS assumed that many parameters are fixed and known, however, due to animal movement and seasonal changes, this assumption does not hold. Similarly, in this work, in order to deploy the proposed traffic enforcement solution, important factors are integrated within our modified model and certain assumptions are relaxed.

Data gathering always poses a great challenge in the development and deployment of security-based applications. Unfortunately, this data is usually withheld from the public. In traffic enforcement, one usually encounters aggregated statistics on the number of accidents and their severity on a monthly or even yearly basis. Releasing some of the gathered data and its analysis to the public in an efficient and natural way could potentially help raise public awareness to road dangers, help drivers trade off travel time, distance and safety, and help in achieving our main goal – reducing the number of road accidents. In this paper, we provide such a publicly available system without jeopardizing the police’s interests.

3 Data-Centric Challenges

Obtaining and leveraging the “right” data is a challenge for most SG-based systems (e.g., [19,17]). To instantiate the TEAP formulation in Israel, we had to address four cardinal data-related challenges: 1) building a road network (Section 3.1); 2) extending the road network into a transition graph (Section 3.2); 3+4) deriving *risk* and *eff* models which are associated with the constructed graph (Sections 3.3 and 3.4). The solution for each of the above four challenges constitutes a component in our solution architecture as illustrated in Figure 1.

3.1 Building a Road Network

The TEAP formulation, similar to many other transition-based formulations, builds on a graph G . In traffic enforcement settings, G represents a road network where the vertices indicate intersections and edges indicate road segments for enforcement. While it may seem trivial to obtain a network graph from open source systems such as OpenStreetMap¹, the translation of this *raw map* into a useful graph raises a number of practical issues and fundamental questions. First, the ITP does not *enforce* traffic laws on local roads (e.g., inner-city roads) which are in the jurisdiction of the local police departments. A naïve solution would be to omit the roads and intersections in question

¹ <https://www.openstreetmap.org>

from the graph. However, this omission may disconnect the graph, which is undesired. Second, some of a road’s features such as the speed limit and number of lanes may change over different parts of the road. In order to construct a suitable prediction model for `risk`, a challenge we address in Section 3.3, one should be able to identify these features for each road segment e . Unfortunately, these features are hard to obtain and are usually aggregated according to some road segmentation by the data collector. In Israel, this data is collected by both the ITP and the Central Bureau of Statistics (denoted CBS, www.cbs.org.il). Finally, it is well known that achieving organizational acceptance of SG-based solutions is a highly complex task (see Section 4). Therefore, using a police-defined road segmentation, when available, is preferable. Police-based segmentation can also encapsulate other domain specific knowledge and constraints, for example, on a narrow road segment without shoulders it may be impossible or highly undesirable to perform enforcement from the police’s perspective. Therefore, using police-planned enforcement segments has significant benefits.

Based on the discussion above, we sought to use the police-based segmentation. However, in the initial stage of the research, the ITP did not allow us access to their segmentation (see Section 4). Therefore, we used the CBS’s road network. In a much later phase of the research, after the publication of [23] which relied on the CBS’s road network, more than a year and a half after the project was initiated, the police segmentation was made available to us. Unfortunately, the police segmentation is far from perfect. Despite its resemblance to the CBS’s segmentation, ample effort had to be invested to transform the provided segmentation into a complete, connected and valid road network. Namely, in many cases, adjunct roads were not marked as such; some road segments were overlapping by mistake while others were disjointed despite being physically connected, etc. Therefore, we manually processed the police segmentation on the basis of the CBS’s network (about 40 human hours). We denote the resulting road network graph as G from this point onwards. G consists of several hundred road segments.

The above is illustrated as **Component A** in our solution architecture, Figure 1.

3.2 Extending G into a transition graph

According to the TEAP’s original formulation, G is extended into a transition graph by replicating each vertex and edge T times. Thus, the formulation assumes that at each time step a police car is assigned to enforce a different road segment. Namely, each police car can perform a single action at each road e given that it is present at e . The formulation also assumes that the enforcement action takes exactly 1 time unit. While this is reasonable in theory, in practice a police car may either engage in various types of enforcement (e.g., using speed-guns to catch speeding cars or randomly choosing drivers on whom to perform breathalyzer tests in order to catch drunk drivers) or in transit (moving across the road network without enforcement so as to reach the intended enforcement site). As a result, the solution to the original TEAP does not prescribe which enforcement should be conducted but only where and when. Furthermore, the solution is limited to only enforcement actions, which according to ITP experts take significantly longer than simple transit actions.

In order to adequately extend G into a practical transition graph, we amend the TEAP formulation in the following way: We denote $A = \{\alpha_1, \alpha_2, \dots, \alpha_n\}$ as the set of actions that a police car can take at every road segment and time. Jain et al. [16] found that providing a very detailed allocation and micro-managing resources does not get as positive a reception from users. Instead, the authors suggest using more abstract instructions, which they found to be better received. Therefore, we simplify the model by assuming two actions, *Enforce* and *Transit*, and leave the investigation of different enforcement options for future work. Let $l(e, t, \alpha_i)$ denote the time it takes for a police car to perform action α_i in road segment e at time t . Then, G is extended such that each vertex v is replicated T times, denoted $\{v_t\}$. If an edge $e = (u, v)$ exists in G , then two types of edges are added for each $t < T$ to the transition graph. We use the following procedure: With accordance to ITP standards, enforcement actions are set to 1 hour (regardless of e and t), therefore an **enforcement** edge is drawn from each u_t to v_{t+60} . On the other hand, transit edges depend on the travel duration of road segment e at time t . Therefore, for each u_t an edge is drawn 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>).² Unfortunately, the above procedure does not suffice when representing the real world. The TEAP formulation relies on the assumption that no two police cars should enforce the same road segment at the same time. However, although not common in practice, 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, we first duplicated each transit edge by the number of police cars available. Namely, for each edge $e = \langle u, v \rangle$ in G and t , we created multiple transit edges to connect u_t and $v_{t+l(e,t)}$. 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. Therefore, while theoretically justified, we avoid replicating transit edges for our deployment. We denote the resulting transition graph as G_T from this point onwards. Note that the notation v_t is still used in its usual meaning. However, notation e_t is not well defined since G_T has multiple edges (e_t may denote an enforcement edge or a transit edge), and therefore will not be used from this point onwards.

Per the ITP's request, we allow all police cars to start and end their routes at any intersection by introducing a dummy source vertex which is connected to all intersections at time 0 and a dummy sink vertex, accessible from all intersections at time $T + 1$. Thus, we assign each $b_{v_t} = 0$. Note that in practical deployment (Section 4) b_{v_t} may assume a value other than zero in cases where a police car is scheduled to visit a specific intersection at a specific time (e.g., due to road work) or plans to start its route at a certain intersection (e.g., once the road work is completed). In such cases b_{v_t} is set appropriately.

The above is illustrated as **Component B** in our solution architecture, Figure 1.

² Time was discretized in 10 minute time-frames.

3.3 Modeling `risk`

The `risk` function captures the risk of accidents occurring at each road segment and time. Recently, traffic police forces began implementing the predictive policing paradigm [20] through which police officers can identify people and locations which are at increased risk. From a methodological standpoint, the effort of predicting road accidents has mainly focused on aggregative analysis, most commonly the prediction of the *annual* number of severe accidents per road segment, using statistical methods such as empirical Bayes, Poisson or negative binomial regression models [15,5]. 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. According to experts in traffic enforcement, the state-of-the-art prediction model is the one used by the Indiana traffic police <https://www.in.gov/isp/ispCrashApp/main.html>. The Indiana system does not consider each road segment separately but instead covers the Indiana state map (including residential areas) with a grid of 1km by 1km squares and provides a prediction as to the risk of a severe accident occurring at each square in three-hour time-frames. According to our discussion with the Indiana traffic police, the prediction model uses approximately 90 features and achieves an Area Under the Curve (AUC) of approximately 0.8. In this paper, we were able to construct a prediction model that provides beneficial predictions for each road segment on *one hour time-frames* by using a unique set of 122 features and 11 years of collected data that achieves an AUC of 0.89. Our model is available at <http://www.biu-ai.com/trafficPolice/> in order to encourage other researchers to tackle the important and challenging task of preventing severe road accidents.

We obtained the records of 11 years of accident reports from the Israeli CBS (2005-2015). By cross-referencing these reports with additional sources such as the Israeli GIS database and the Israeli Meteorological Service (IMS, www.ims.gov.il) weather reports, we were able to characterize each accident using a unique set of 122 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 Israeli 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 consists of rural low-density residential areas for the most part [3 features]. The ITP further divides Israeli 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] and average speed [5 features]. 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].

Training a Deep Neural Network Using more than 30,000 accident records (accidents that took place between 2005 and 2015 in Israel) and undersampling the “non-accident” class (see [6]), we trained a deep neural network model. The model receives, as input, vectors of 122 features (as described above) representing a road segment and time. The model returns a value in the [0,1] range, *acting as a proxy* to the likelihood of an accident occurring on that road at that time. Note that severe accidents are sporadic events in both time and space. Therefore, directly estimating the probability of accidents occurring on any road segment at any time is extremely challenging.

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.

We compared our prediction model to several baseline prediction models, such as logistic regression, SVM and XGBoost. The latter is currently in use by the Indiana traffic police for the same task. Our model achieves an AUC of 0.89, outperforming logistic regression, SVM and XGBoost, which recorded 0.78, 0.77 and 0.82, respectively, using 10-cross validation.

Using entropy-based ranking feature selection [14], we identify 5 groups of high ranking features in the following order of importance: 1) number of past accidents; 2) traffic volume; 3) road type; 4) speed limit; and 5) weather. Contrary to what the authors initially expected, among the lowest ranking features one would find: 1) day of the week; 2) month of the year; and 3) enforcement district. We plan to further analyze these results in order to provide additional practical suggestions for the ITP.

The above is illustrated as **Component C** in our solution architecture, Figure 1.

3.4 Modeling eff

The eff function measures the effect that the police allocation history has on the risk of accidents occurring on any road segment at any time. Unfortunately, as discussed in Section 4, getting the ITP's allocation history was very difficult. We only gained access to a single month's allocation in 2017. Unfortunately, the allocation was recorded in the format of GPS coordinates which police cars report once every few minutes while on duty. Translating these GPS coordinates into a usable format which will allow us to understand when and where a police car was enforcing the law, driving through a road segment or handling other events (e.g., a traffic police car may be temporarily assigned to assist police patrol cars) is extremely complex. For example, a police car may stay for some time at a gas station. The police car may be refueling, there may be a technical problem with the police car that needs fixing, or the police may be having a lunch break or enforcing the law on a nearby road. We are currently working with the ITP on a methodology to address the above issues in the future.

Isolating the effect that the police allocation has on the likelihood of road accidents is extremely complex regardless of the above mentioned issues. First, endogeneity is a big concern. Naturally, police cars are likely to be stationed on dangerous roads. Naïve statistical inference may conclude that the presence of a police car *increases* the likelihood of accidents. The endogeneity problem is particularly relevant in the context of time series analysis of causal processes, which is the case in traffic enforcement. Even if we assume that eff is independent of all other factors within a given period, but is influenced by the average speed of traffic in the preceding period, then eff is exogenous within the period but endogenous over time, which poses an additional statistical challenge.

As a result of the above discussion, we base our estimation of eff on [28]. The author 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 on a given road at a given time, if 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 [9] for a review). In accordance with the ITP's expert estimations, we define time halo effects in the exponential diminishing form $\frac{36}{2^k}\%$ where $k \geq 0$ is the hours 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 a road segment is enforced for two hours straight (and no other time or distance halo effects are appropriate), eff assumes 45% ($= 36\% + \frac{18}{2}\%$).

As discussed before, we are currently working towards a more data-driven approach for modeling eff in Israel.

The above is illustrated as **Component D** in our solution architecture, Figure 1.

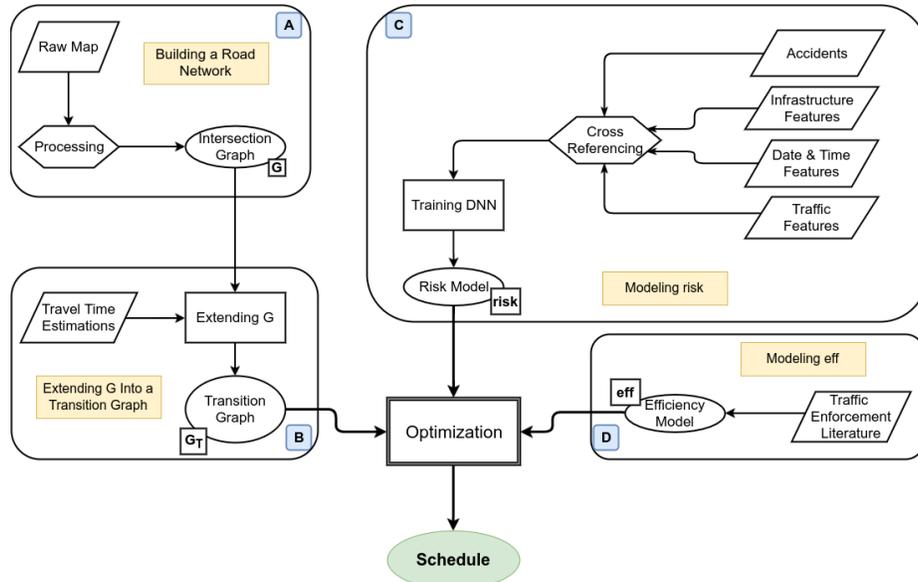


Fig. 1: High-Level Solution Architecture

4 Police-Deployment Challenges

In order to deploy our model and solution in Israel, several logistical/technical issues had to be addressed. In this section we discuss challenges which arise from our interaction with the ITP and do not focus on data.

4.1 Security Clearance

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

The clearance that we obtained allowed us full access to ITP information. However, due to bureaucratic reasons, obtaining each piece of information required a long approval process. As a result, at the moment, we do not have access in *real time* to

important pieces of information such as average speeds. Note that the ITP does have accurate estimations thereof using anonymized cellular reports [2].

4.2 Logistics

In order to deploy our model and solution in the real world, logistical constraints need to be addressed. There are three main constraints: First, some police resources may have specific schedule constraints of the form “Officer X must arrive at traffic court at 2pm and stay there until 3pm to testify in a trial”. Such constraints are easily integrated within our model by setting the b values of the intersections in the transition graph appropriately. Second, 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 4th 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 4th 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 4th hour. These vertices are accessible from any vertex during the 4th 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 4th hour were duplicated such that every node had two copies – “pre-break” and “post-break”. Then, pre-break nodes were disconnected from 5th hour 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 car can only reach the 5th hour of the shift if it goes through 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. Third, an unexpected event may cause a police car to deviate from its schedule. For example, a police car may be sent to clear an unexpected road block, making its schedule infeasible. The ITP claims that there is no easy way to determine the likelihood of these unexpected events in the real world, making the MDP-based approach used by TRUSTS inapplicable. We resolve this issue as proposed in [8], by *assuming* perfect execution and only after a non-default transition occurs does the central command resolve the TEAP starting from the current state. Yet this requires a quick solution, as the ITP would not accept a long wait time. Given that the original TEAP formulation (as presented in [23]) was modified in this work, we reevaluated the runtime results of the proposed solution given the new formulation. Similar to our previous findings, the runtime of the amended solution does not impose a significant concern. Specifically, given the modified formulation for allocating 10 police cars for 96 hours (4 days, 12 shifts) in a designated district, the proposed solution runs in approximately 1 minute, compared to more than 30 minutes by a naïve solver as described in the original paper. Namely, the modifications proposed in this paper do not jeopardize the solution method’s superiority over baseline methods.

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.

4.3 Deployment and Evaluation

As a first step, the ITP wanted to know how different our provided allocations are from their current hand-crafted, time-consuming allocations. Unfortunately, as mentioned above, usable records of past police allocations are currently unavailable. Instead, we were given a list of road segments which the ITP considers to be of “special enforcement interest”. The list, which consists of approximately 5% of the road segments, was constructed by the ITP’s researchers and acts as a guideline as to which road segments are the most important to enforce during a given month. Generally speaking, the ITP focuses on these road segments and their surroundings. We generated a schedule for a whole month using the modified TEAP formulation as described above, and identified the number of times a police car was assigned to one of the designated road segments. On average, a road segment on the list is enforced 25 times more often than other road segments. Note that the entire list was associated with very high `risk` scores, specifically, if road segments were to be sorted according to their average `risk` score on any day or at any time, the entire list would be at the top 15% of all road segments.

There are many challenges when attempting to evaluate deployed SG-based systems in the field [27]. Specifically, unlike conceptual ideas, where we can run thousands of careful simulations under controlled conditions, we cannot conduct such experiments in the real world [18]. We are currently working towards a head-to-head comparison of our proposed solution against an expert-generated allocation. Our controlled experiment would take place in two very similar enforcement districts. During a period of at least one month, one district will use our system while the other will use an expert-generated allocation. It is hard to believe that such a comparison will yield a statistically significant difference, due to the fact that road accidents are very sporadic. A similar problem was encountered in PAWS, where the authors faced a similar issue of how to quantify the number of saved wildlife due to their provided solution. The authors instead use human and animal signs as indicators that PAWS patrols prioritize areas with higher animal and poacher activity. In the same spirit, we would also record other statistics such as the average speeds and the number of citations issued by the police officers, as well as other statistics, as proxies to the allocation’s quality. Speed has been identified as a key risk factor in road traffic safety, influencing both the risk of a road crash and the severity of the injury that results from crashes [10]. Furthermore, a higher number of issued citations can suggest that the provided solution can avoid the human-generated predictability. Note that the benefit of our solution, as well as other deployed solutions such as PAWS, should be expected in the long-term.

5 Raising Public Awareness

Recently, the World Health Organization (WHO) has released a report on road traffic injuries and how they can be reduced [29]. The WHO mentioned that governments

should take a more holistic approach to mitigating road accidents, which includes not only better enforcement but also the modification of infrastructure and the raising of public awareness. The WHO further mentions the latter as one of its own tasks, “sharing information with the public on (road) risks. . .”.

As discussed in Section 3.3, the Indiana traffic police has provided a visual tool that uses color shading to show a low, moderate or high probability of a crash occurring in each 1-square-kilometer area in the state. This interactive map predicts where crashes are likely to occur across the state of Indiana, so citizens and law enforcement can be more proactive in avoiding or preventing accidents. Despite its publication in the general media, according to the developers, only a few citizens use the system. We speculate that there are two main reasons for this disappointing feedback: First, the use of a grid-cell instead of road segments and the relatively large time frames (3-hours long) make the system less practical for drivers. Consider a driver who tries to decide which route to take at 8am from one point to the other. The driver is interested in the risk associated with the different route options at 8am rather than grid-cells between 8am and 11am. Second, the system uses a newly designed interface. However, in creating solutions for people, we must be cognizant to how difficult it will be for a user to adopt our solution. Each deviation from existing methodology is a step away from the familiar that we must convince the user to accept [26].

To address these two issues we provide the *www.SafeRoad.today* open-access system, which is mounted on the popular Google maps interface. Users can access our website and review the risk in each road segment in Israel in 1-hour time frames discretized into 5 risk levels – “very low”, “low”, “average”, “high” and “very high”. However, unlike the ITP, drivers may not be interested in the risk of an accident on a road segment but rather **their own risk** of being involved in an accident. Therefore we provide 2 layers, one illustrating the *risk* function as described in Section 3.3, and the other illustrating the *risk* function after the normalization by the expected traffic volume. Users can query the system for risks associated with any road at any time. Another type of query is a route query. The user can query the system for routes to take her from one point to another using the regular Google maps interface. Once a query for a route is made, in addition to the travel time and distance for each possible route provided by Google maps, the system provides a *risk* estimation for each route, enabling the driver to consider the safety factor in her route selection. An additional layer allows the user to query the system for past road accidents. Given dates and locations, provided by the user, all severe accidents that occurred during the designated time and at the designated locations will appear on the map, each according to the place of the crash. Each accident appears alongside some basic information regarding the crash such as the date and time, type of crash (e.g., a car and a motorcycle), number of injuries, etc. Note that our system does not jeopardize the confidentiality of police data – it simply does not contain any restricted data. To the best of our knowledge, our system is the first of its kind.

The *SafeRoad.today* system joins existing publicly accessible systems designed for raising public awareness of other road dangers. For example, WAZE³ provides road

³ <https://www.waze.com/livemap>

danger alerts for drivers such as road work, Sustrans⁴ provides safe cycle routes, factoring in bike lanes and traffic free routes and Rudder⁵ provides safe pedestrian routes, factoring in street lighting.

A snapshot of our system is provided in Figure 2.

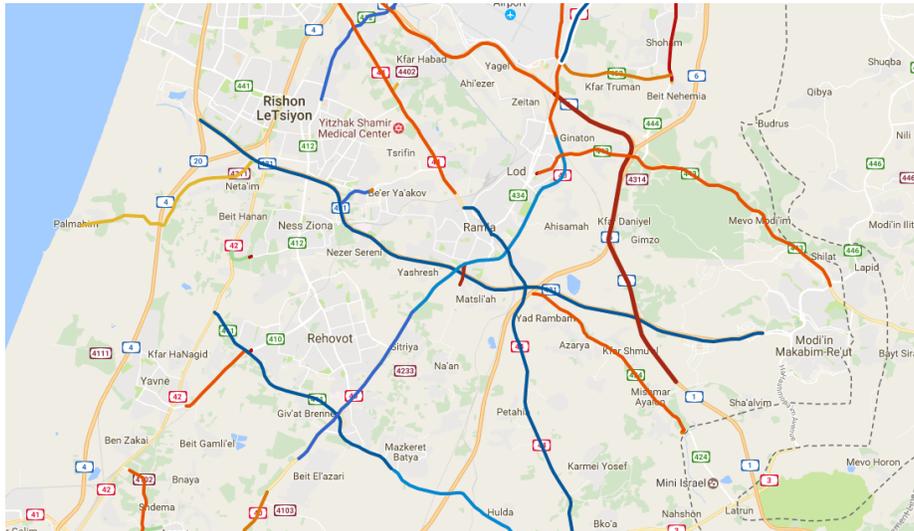


Fig. 2: A snapshot from *SafeRoad.today* system. Risk illustration is as follows: Red = “Very High”, Orange = “High”, Azure = “Low” and Dark Blue = “Very Low”. Average risks are not depicted.

6 Conclusions

In this paper we present key challenges and solutions in transforming a lab-based theoretical model for mitigating road accidents through efficient enforcement to an operational system in field trials. We focus on three main challenges: data, deployment and raising public awareness. These challenges, and specifically the data-centric challenges, are very common in security-based applications. Two important components of our provided solution include a novel traffic accident prediction model, available at <http://www.biu-ai.com/trafficPolice/>, and an open-access risk visualization system, *SafeRoad.today*, which is available for public use. Our prediction model provides a state-of-the-art prediction tool, based on 122 features and 11 years of collected data, that we hope will encourage other researchers and practitioners to tackle the important and challenging task of preventing severe road accidents. In addition, our

⁴ <http://www.sustrans.org.uk/ncn/map>

⁵ <https://walkrudder.com/>

SafeRoad.today system is designed for and targeted at raising public awareness and allowing drivers to make better decisions. These components, which amend and extend our police allocation mechanism from [23], combine to provide a viable tool for mitigating road accidents in Israel and can be adapted to other countries as well.

This “behind the scenes” paper also provides a unique look into the considerations and decisions that developers of deployed security-based applications have to face. Since the challenges discussed and addressed in this work are not unique to traffic enforcement, we hope that the provided discussion and insights will assist others in the process of deploying their security-based systems in the real world.

Acknowledgement

This work was supported in part by the LAW-TRAIN project that has 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 and the Japan Science and Technology Agency (jst), Japan .

References

1. Adler, N., Hakkert, A.S., Kornbluth, J., Raviv, T., Sher, M.: Location-allocation models for traffic police patrol vehicles on an interurban network. *Annals of Operations Research* 221, 9–31 (2014)
2. Bar-Gera, H.: 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)
3. Bates, L., Soole, D., Watson, B.: The effectiveness of traffic policing in reducing traffic crashes. In: *Policing and security in practice*, pp. 90–109. Springer (2012)
4. Brown, M., Saisubramanian, S., Varakantham, P.R., Tambe, M.: Streets: game-theoretic traffic patrolling with exploration and exploitation. In: *Proceedings of the 2014 AAAI Conference on Artificial Intelligence* (2014)
5. Chang, L.Y.: Analysis of freeway accident frequencies: negative binomial regression versus artificial neural network. *Safety science* 43(8), 541–557 (2005)
6. Chawla, N.V.: Data mining for imbalanced datasets: An overview. In: *Data mining and knowledge discovery handbook*, pp. 853–867. Springer (2005)
7. DeAngelo, G., Hansen, B.: Life and death in the fast lane: Police enforcement and traffic fatalities. *American Economic Journal: Economic Policy* 6(2), 231–257 (2014)
8. Delle Fave, F.M., Jiang, A.X., Yin, Z., Zhang, C., Tambe, M., Kraus, S., Sullivan, J.P.: 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)
9. Elliott, M., Broughton, J., et al.: How methods and levels of policing affect road casualty rates. *Transport Research Laboratory* (2005)
10. Elvik, R.: 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)
11. Elvik, R., Vaa, T., Erke, A., Sorensen, M.: *The handbook of road safety measures*. Emerald Group Publishing (2009)

12. European Transport Safety Council: How traffic law enforcement can contribute to safer roads: Pin flash report 31 (June 2016), http://etsc.eu/wp-content/uploads/NEW_PIN_FLASH31_final.pdf/
13. Fang, F., Nguyen, T.H., Pickles, R., Lam, W.Y., Clements, G.R., An, B., Singh, A., Tambe, M., Lemieux, A.: Deploying paws: field optimization of the protection assistant for wildlife security. In: Twenty-Eighth IAAI Conference (2016)
14. Guyon, I., Elisseeff, A.: An introduction to variable and feature selection. *Journal of machine learning research* 3(Mar), 1157–1182 (2003)
15. Hauer, E.: Empirical bayes approach to the estimation of unsafety: the multivariate regression method. *Accident Analysis & Prevention* 24(5), 457–477 (1992)
16. Jain, M., Tsai, J., Pita, J., Kiekintveld, C., Rathi, S., Tambe, M., Ordóñez, F.: Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces* 40(4), 267–290 (2010)
17. Kar, D., Ford, B., Gholami, S., Fang, F., Plumtre, A., Tambe, M., Driciru, M., Wanyama, F., Rwetsiba, A.: Cloudy with a chance of poaching: Adversary behavior modeling and forecasting with real-world poaching data. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems. pp. 159–167 (2017)
18. Nguyen, T.H., Kar, D., Brown, M., Sinha, A., Jiang, A.X., Tambe, M.: Towards a science of security games. In: *Mathematical Sciences with Multidisciplinary Applications*, pp. 347–381. Springer (2016)
19. Park, N., Serra, E., Subrahmanian, V.: Saving rhinos with predictive analytics. *IEEE Intelligent Systems* 30(4), 86–88 (2015)
20. Perry, W.L.: Predictive policing: The role of crime forecasting in law enforcement operations. Rand Corporation (2013)
21. Pita, J., Bellamane, H., Jain, M., Kiekintveld, C., Tsai, J., Ordóñez, F., Tambe, M.: Security applications: Lessons of real-world deployment. *ACM SIGecom Exchanges* 8(2), 5 (2009)
22. Pita, J., Jain, M., Marecki, J., Ordóñez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., Kraus, S.: Deployed armor protection: the application of a game theoretic model for security at the los angeles international airport. In: Proceedings of the 2008 International Conference on Autonomous Agents & Multiagent Systems. pp. 125–132 (2008)
23. Rosenfeld, A., Kraus, S.: When security games hit traffic: Optimal traffic enforcement under one sided uncertainty. In: Proceedings of the 26th International Conference on Artificial Intelligence, IJCAI (2017)
24. Schechtman, E., Bar-Gera, H., Musicant, O.: Driver views on speed and enforcement. *Accident Analysis & Prevention* 89, 9–21 (2016)
25. Simandl, J.K., Graettinger, A.J., Smith, R.K., Jones, S., Barnett, T.E.: 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)
26. Tambe, M.: *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press (2011)
27. Taylor, M.E., Kiekintveld, C., Western, C., Tambe, M.: A framework for evaluating deployed security systems: Is there a chink in your ARMOR? *Informatica* 34(2) (2010)
28. Weisburd, S.: Does police presence reduce car accidents? Tech. rep. (August 2016)
29. World Health Organization: Road traffic injuries fact sheet (May 2017), <http://www.who.int/mediacentre/factsheets/fs358/en/>
30. Yang, R., Ford, B., Tambe, M., Lemieux, A.: Adaptive resource allocation for wildlife protection against illegal poachers. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems. pp. 453–460 (2014)
31. Yin, Z., Jiang, A.X., Johnson, M.P., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Tambe, M., Sullivan, J.P.: Trusts: Scheduling randomized patrols for fare inspection in tran-

- sit systems. In: Proceedings of the 2012 Innovative Applications of Artificial Intelligence (2012)
32. Zhang, C., Gholami, S., Kar, D., Sinha, A., Jain, M., Goyal, R., Tambe, M.: Keeping pace with criminals: An extended study of designing patrol allocation against adaptive opportunistic criminals. *Games* 7(3), 15 (2016)