

OPRADI: Applying Security Game to Fight Drive under the Influence in Real-World

Luzhan Yuan, Wei Wang, Gaowei Zhang, Yi Wang

Beijing University of Posts and Telecommunications
{johnyuan, weiwang, zhanggaowei, yiwang}@bupt.edu.cn

Abstract

Driving under the influence (DUI) is one of the main causes of traffic accidents, often leading to severe life and property losses. Setting up sobriety checkpoints on certain roads is the most commonly used practice to identify DUI-drivers in many countries worldwide. However, setting up checkpoints according to the police's experiences may not be effective for ignoring the strategic interactions between the police and DUI-drivers, particularly when inspecting resources are limited. To remedy this situation, we adapt the classic Stackelberg security game (SSG) to a new SSG-DUI game to describe the strategic interactions in catching DUI-drivers. SSG-DUI features drivers' bounded rationality and social knowledge sharing among them, thus realizing improved real-world fidelity. With SSG-DUI, we propose OPRADI, a systematic approach for advising better strategies in setting up checkpoints. We perform extensive experiments to evaluate it in simulated environments, contexts reconstructed from real-world data, and preliminary deployments, in collaborating with a Chinese city's police bureau. The results reveal its effectiveness in improving police's real-world operations, thus having significant practical potentials.

Introduction

Traffic accidents, particularly those caused by driving under the influence (DUI) of alcohol or drugs, often lead to significant life and property losses. According to the National Highway Traffic Safety Administration (NHTSA)'s statistics, more than 7,000 people died from DUI in the United States in 2020 (NHTSA 2022). Currently, identifying DUI cases mostly relies on the traffic police officers' efforts in setting up checkpoints on certain roads to examine the sobriety of drivers passing through (Erke, Goldenbeld, and Vaa 2009). Once identified as a DUI-driver, one faces severe punishments, including monetary penalties, driving license revoking, or even imprisonment (Yao, Johnson, and Beck 2014). Thus, a DUI-driver, if still having any sense left, would try to avoid being caught by the police, while the police hope to catch DUI-drivers as many as possible after setting up checkpoints.

Therefore, the police and DUI-drivers are opponents of each other. Due to limited resources, the police have to de-

termine where to set up checkpoints to maximize the effectiveness of catching DUI-drivers (Elder et al. 2002). Meanwhile, DUI-drivers may deliberately choose some routes to reduce the chance of being caught. Hence, the two parties form a game where they are the players, and their strategies are selecting the locations of checkpoints and driving routes, respectively. Here, such a game has several practical restrictions, (1) the checkpoints can only be set at a very limited number of locations; (2) even a small city's road system is often complex enough to allow a DUI-driver find an *escape* route. Therefore, determining the potential checkpoints is a strategic resource allocation optimization problem for the police under certain practical constraints; (3) a checkpoint cannot be changed in a period, resulting in some drivers may gradually learn its locations by certain means; and (4) some people may help their friends escape by sharing the locations of checkpoints.

Conventionally, such a problem is approached with the Stackelberg security game (Yin et al. 2010) where DUI-drivers are attackers and the police are the defender. For instance, assuming that DUI-drivers could observe the defender's strategy, Jie et al. (2017, 2020) use SSGs to model DUI-drivers and the police's strategic interactions and compute the Strong Stackelberg Equilibrium (SSE) to offer an optimal strategy to the police. However, these conventional models in literature often fail to consider the above four restrictions faithfully in describing real-world DUIs inspections, making the SSE (if exists) less desirable.

Thus, we propose a novel extension (**SSG-DUI**) to the classic SSG to address the police's problem of finding an effective strategy for setting up sobriety checkpoints in real-world contexts. In our model, the attackers are restricted by their bounded rationality (Simon 1990), i.e., they have inaccurate prior knowledge about the police (the defender)'s inspection strategy but may learn it via certain ways. In addition, we acknowledge the existence of social knowledge sharing among drivers (DUI-drivers and sober drivers); they may share the current locations of checkpoints for some reasons, e.g., social reputation (Sanfey 2007). Thus, both shall be viewed as attackers though having distinct payoffs.

Based on SSG-DUI, we propose a systemic approach named **OPRADI** (Optimization for Police Resources Allocation in Drive under the influence Inspection) to advise effective strategies to the police. With the help from a small

Chinese city (Chiping)’s police bureau, we do not only evaluate the OPRADI’s effectiveness in simulated environments, but also test its usefulness in complex real-world contexts. Particularly, our empirical evaluation reveals that OPRADI can offer a much more effective (**2.72 times higher** regarding the police’s payoff) strategy to the police than what is currently being used. Specifically, our contributions are tri-fold:

1. SSG-DUI, a novel extension to the SSG, which features DUI-drivers’ bounded rationality and knowledge sharing, thereby achieving high fidelity in reflecting reality;
2. OPRADI, a systematic approach advising the police effective strategies for setting up checkpoints;
3. Extensive evaluations of OPRADI, highlighting its effectiveness in real-world police operations with the assistance from the police bureau of Chiping.

Moreover, per the agreement with the police bureau of Chiping, OPRADI is currently deployed to its operations, and the preliminary outcomes are encouraging.

Background and Related Work

DUI inspection is a typical scenario involving both attackers (the drivers) and defenders (the police) who act sequentially. Thus, it could be abstracted as Stackelberg Security Games (SSGs) (Paruchuri et al. 2007; Pita et al. 2008; Yin et al. 2010). In an SSG, a defender must determine a limited number of targets to put their limited resources on, to fight against an attacker’s potential attacks. Meanwhile, an attacker can observe and learn the defender’s strategy, i.e., how the defender’s resources are deployed, and choose actions of attack after some deliberations.

The SSGs framework naturally fits many real-world scenarios and has been applied to solving many practical problems, resulting in dozens of successful international deployments. For instance, in cybersecurity, researchers often use it to deal with problems such as information privacy (Alvim et al. 2017; Shokri 2015), intrusion detection systems (Sengupta et al. 2018; Zhu and Başar 2013; Roy et al. 2021), obfuscation & deception (Clark et al. 2012; Pawlick and Zhu 2017; Zhu et al. 2012), etc. Besides, in tackling social problems, Fang et al. developed and deployed SSG applications for poaching and illegal fishing prevention (Fang, Stone, and Tambe 2015). Meanwhile, many factors were introduced into the classic SSGs to better describe real-world scenarios. Such factors include the existence of informant (Shen et al. 2020), attackers’ memory limitations (Nguyen et al. 2013), and so on. Interested readers may refer to some quite comprehensive recent surveys (Kar et al. 2018; Sinha et al. 2018).

In an SSG, the defender’s mixed strategy is the probability distribution over deploying a set of resources on targets to be protected, while an attacker’s mixed strategy is the probability distribution of attacking a target (Yin et al. 2010). If a target i is attacked and it is protected by the defender, the defender gets a reward R_i^d and the attack receives a penalty P_i^a . But if target i is not protected, the defender gets a penalty P_i^d and the attacker receives a reward R_i^a . For a defender

strategy $S = \langle s_i \rangle$, where s_i is the probability of protecting target i , the players’ expected utilities when target i is attacked are:

$$U_i^a = s_i P_i^a + (1 - s_i) R_i^a \quad (1)$$

$$U_i^d = s_i R_i^d + (1 - s_i) P_i^d \quad (2)$$

SSGs’ characteristics make them particularly suitable for modeling, exploring, and deriving solutions for certain traffic violations involving police and offenders (Yoo and Langari 2013; Rosenfeld and Kraus 2017; Guo et al. 2016). Jie et al. (2017) suggested that catching DUI-drivers could be modeled as a zero-sum SSG and proposed an OISDD algorithm to compute the optimal strategy for allocating scarce police resources. Guo et al. (2016) proposed the Network Flow Interdiction Game (NIFG) model, a variant of Stackelberg game model, to detect illegal network flow. Later, Jie et al. (2020) modeled the DUI-drivers inspecting scenario as a network interdiction problem and considered factors such as the traffic network’s scalability and the attackers’ diversity, aiming at improving the expressiveness of their models (Jie et al. 2020). However, these approaches had several key shortcomings: (1) they assumed the DUI-drivers can fully observe the defender’s strategy ahead of their actions; (2) they omitted the knowledge sharing among drivers; (3) sober drivers might also engage in the game by sharing police’s strategies they observed with DUI-drivers, thus are attackers too. These shortcomings thus limited the potentials in real-world applications.

A Motivating Scenario

The police’s goal is to catch DUI-drivers as many as possible by setting up checkpoints on some selected roads based on their experiences. Meanwhile, if DUI-drivers decide to drive, they want to minimize the chance of being caught. For the police, when and where a specific DUI driver will appear are hard to predict, though they may have some rough estimates. For a DUI-driver, learning the police’s strategy is non-trivial since they can only directly observe a couple of blocks in the city. However, they might indirectly get some partial knowledge about checkpoints from their friends via a distributed knowledge sharing process, e.g., one might send their friends an SMS saying “*Be careful, police are checking drunk drivers on the xxx street.*”, which effectively results that more and more DUI drivers may (partially) figure out the police’s strategy. Consider the following scenario.

Two drivers d_1 and d_2 who just had a decent dinner were going to drive home (from A to B on the map shown in Fig. 1). While at dinner, d_1 did not drink alcohol for some reason and left earlier; d_2 had a few shots and decided to play some billiards before going home. On the way (e_1) to B, d_1 saw that a few cops were setting up a checkpoint. “Oh, my gosh! d_2 will get a heavy ticket if he drives through here. I shall warn him,” d_1 thought. So, he pulled over his car and sent his current location to d_2 , saying, “Hey dude, avoid this place; otherwise cops will catch you!” d_2 heard his cell phone beeping and then read d_1 ’s message. He immediately decided to take another route (e_2). Thirty minutes later, he arrived home (B), took a long breath, “what a lucky night!”




Figure 1: The simple map with two routes e_1, e_2 from A to B . The red X is the checkpoint set by the police.

The above story can be abstracted as a typical SSG. d_1 and d_2 were attackers, and the police were defenders. In the beginning, both had zero knowledge about the strategy of defenders (checkpoints on e_1) because they could only observe a limited spatial scope¹. Let us shamelessly assume d_1 and d_2 could only choose routes from either e_1 or e_2 , which were basically indifferent to them. In case d_1 did not share the information to d_2 , the police would have 50% probability to catch d_2 . But once d_1 shared the information to d_2 , the drunk driver could update his prior knowledge about the defender's strategy and therefore chose e_2 , voiding the police's efforts and increasing the potential public safety risks. This story demonstrates the importance of considering the limited scope of attackers and knowledge sharing among the attackers (drunk and sober drivers) in shaping the game's dynamics. However, such factors have been long ignored in the literature.

SSG-DUI Game: Definitions and Notations

To better reflect the real-world DUI inspections, we define a new game structure **SSG-DUI** as follows.

The Game: The SSG-DUI game is a special form of SSG between one defender (the police) and multiple attackers A indexed by a_i (DUI-drivers and sober drivers sharing knowledge²) on a traffic network G with nodes V (each node represents a crossing) and E edges (each edge represents a road section between any two adjacent crossings). Once the checkpoint is set up, it could not move since moving checkpoints could be costly for interrupting traffic flows which might cause potential traffic dangers and accidents.

The Strategies: (i) The defender can only set N ($N \ll E$) stationary checkpoints and chooses a mixed strategy at the beginning of the game. (ii) DUI-drivers choose routes on G based on their knowledge about the defender's strategy and share knowledge with a limited number (up to F) of friends;

¹In a general sense, this is resulting from one's bounded rationality (Simon 1990) in practical decision-making.

²In the rest of the paper, we simply assume that sober drivers who do not share knowledge do not participate in the game.

sober drivers participate in the game only by sharing knowledge. (iii) Attackers would update their limited knowledge once receiving the message from their friends containing knowledge about the checkpoints.

The Defender's Mixed Strategy: The mixed strategy of the defender is defined as the probability distribution of setting checkpoints on E , denoted by S which is an E -dimensional vector. We use C_i to represent the chance of inspecting an attacker a_i regardless of one's sobriety.

The Payoff Structure: (i) The attacker a_i would receive either a reward R_i^a if not being caught by the police or a penalty P_i^a if being caught by the police. Attackers will also share a message about the defender's strategy and get a small fixed payoff p_s . For sober drivers, R_i^a and P_i^a are equal to 0. (ii) The defender will get a payoff R_i^d for catching DUI-drivers.

With the mixed strategy of the defender, the defender aim to maximize his expected utility:

$$U^d = \sum_i R_i^d * C_i \quad (3)$$

And the expected utility of the attacker A_i is:

$$U_i^a = P_i^a * C_i + R_i^a * (1 - C_i) + p_s \quad (4)$$

Therefore, the defender's target is to maximize the expected total utility satisfying the vNM theory (Von Neumann and Morgenstern 1947).

An Attacker's Knowledge: The knowledge represents an attacker's dynamic beliefs about the defender's mixed strategy. In the early stage, there shall be a substantial divergence between an attacker's belief and the defender's actual mixed strategy. With the progress of the game, the divergence would decrease due to the knowledge sharing among attackers. We use a_i^K to denote a_i 's knowledge.

The above defines the SSG-DUI game, we summarize all notations in Tab. 1.

Notations	Meanings
G	The traffic map.
E	The number of the roads.
V	The set of the nodes.
a_i	The i^{th} attacker (maybe sober or drunk).
a_i^K	The a_i 's prior knowledge.
F	The maximum number of attacker's friends.
S	The defender's mixed strategy.
C_i	The probability of inspecting a_i .
R_i^a	The a_i 's reward of escaping from inspecting.
P_i^a	The a_i 's penalty of being caught by the police.
R_i^d	The defender's reward of inspecting driver a_i .
p_s	The attacker's reward of sharing message with friends.
U_i^a	The attacker a_i 's expected utility.
U^d	The defender's expected utility.

Table 1: Notations and their meanings in SSG-DUI.

OPRADI: Our Approach in A Nutshell

With the SSG-DUI game, this section introduces a systematic approach to help the police (the defender) form an effective strategy in setting up DUI checkpoints. We name the approach as **OPRADI**.

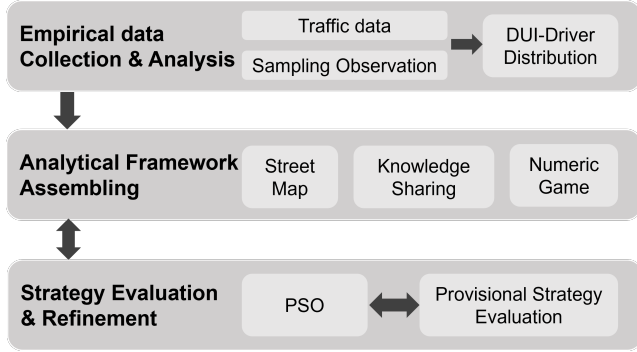


Figure 2: An overview of OPRADI.

Fig. 2 overviews OPRADI consisting of three stages. Stage 1 collects the empirical traffic data and observes a sample of drinking venues, which yields some estimations of the distribution of DUI-drivers. Using the empirical insights derived in Stage 1, Stage 2 assembles the abstract street map, knowledge sharing mechanism among attackers, and SSG-DUI’s numeric form to constitute the analytical framework. Such an analytical framework brings empirical context and the abstract game model together. Then, in Stage 3, the defender’s strategy is optimized using the Particle Swarm Optimization (PSO) algorithm within the analytical framework in an iterative fashion, which finally presents actionable strategic advice to the police. The rest of this section is going to introduce the main operationalizations in each stage.

Stage 1: Empirical Data Collection & Analysis

The empirical data collection has two key operations, **collecting traffic flow data** of interested time frames and **performing sampling observations**. While traffic flow data could be easily extracted from the videos recorded by road surveillance cameras, sampling observations have to be done through human efforts. The police often sample a few roads which are the gathering places of restaurants and pubs. They often sit in an unmarked car wearing civilian clothes, watch people exiting from these drinking venues, judge if they are potential DUI-drivers. Doing so allows them to get rough estimations of the proportion of potential DUI-drivers joining the traffic flow. Thus, with the traffic flow data, they could have some estimations of the distribution of DUI-drivers in the population of all drivers in specific periods.

Stage 2: Analytical Framework Assembling

The insights gained in Stage 1 are integrated into a street map (e.g., Fig. 5c) abstracted from the corresponding city’s real-world road system (e.g., Fig. 5b). Then, we need to define attackers’ knowledge sharing mechanism (DUI-drivers and participating sober drivers).

Knowledge Sharing: Current OPRADI initializes attackers’ prior knowledge using the simplest uniform distribution, i.e., the probability of the police appearing on every road is the same at the beginning of the game for all attackers. Each attacker shares current knowledge with up to F friends; those who received the message will update their prior knowledge accordingly before joining the traffic. An attacker a_i ’s knowledge updating process is approximated as a combination of one’s prior knowledge and the knowledge shared by other attackers, see Eq. 5.

$$a_i^K \leftarrow (1 - \lambda)a_i^K + \lambda \sum_j m_j \quad (5)$$

where λ is the parameter referring to a_i ’s confidence in the authenticity of the knowledge shared by other attackers and m_j represents the knowledge shared by a_j . The knowledge updating would gradually reduce the divergence between an attacker’s knowledge and the defender’s actual mixed strategy. Note that OPRADI may use other sophisticated knowledge updating mechanisms, e.g., stochastic learning dynamics (Arieli and Young 2016), to improve its expressiveness of real-world phenomena.

The next step is assigning numeric values to the SSG-DUI game. In the scope of this paper, we simply use the numeric values shown in Tab. 2. In each pair of numbers, the first is the driver’s payoff and the second is the police’s payoff.

	Inspected	Uninspected
DUI-driver	(-0.9, 1)	(1.1, 0)
Sober driver	(0.1, 0)	(0.1, 0)

Table 2: Players’ numerical payoffs used in OPRADI.

By assembling the street map and the knowledge sharing mechanism with the numeric form of the SSG-DUI game, we establish an analytical framework that not only keeps the analytical power of game-theoretic models but also reflects a substantial level of empirical reality.

Stage 3: Strategy Evaluation & Refinement

Evaluating Strategy: Recall that we define the defender’s strategy as S in SSG-DUI game, which yields a utility in the

Algorithm 1: Calculating the defender’s utility.

- 1: Initialize the map $G = (V, E)$ and players including payoff structure R_i^d, P_i^a, R_i^s, p_s and attackers’ prior knowledge a_i^K .
 - 2: The defender adopts the mixed strategy S .
 - 3: **for** a_i in attackers **do**
 - 4: Obtain a_i ’s strategy based on the prior knowledge a_i^K .
 - 5: Calculate C_i i.e., the probability of inspecting a_i .
 - 6: Randomly select F friends and update their knowledge according to Eq. 5.
 - 7: **end for**
 - 8: Calculate defender’s utility according to Eq. 3.
-

form of Eq. 3. Thus, by calculating the utility associated with a specific S , and comparing it with other strategies' utility, we could evaluate each strategy's effectiveness to identify ones that fulfill our requirement. How to calculate the defender's utility is described in Algo. 1.

Refining Strategy: The refining process starts from a randomly chosen provisional strategy. Then, it will be iteratively refined. OPRADI aims to find a strategy maximizing the defender's utility, i.e., as many as possible DUI-drivers could be caught. We select the Particle Swarm Optimization(PSO) algorithm to fulfill this task due to its independence to the gradient or any differential form of the objective function (Kennedy and Eberhart 1995). Moreover, OPRADI does not exclude using other derivative-free optimization algorithms like Bacterial Foraging Optimization method (Passino 2002). PSO algorithm optimizes the objective function by simulating the behavioral dynamics of birds in a flock. Algo. 2 lists the pseudo-code of applying the PSO algorithm in our task.

Note that OPRADI does not target finding the Strong Stackelberg Equilibrium (SSE). In SSG-DUI, the stochastic knowledge sharing process leads to uncontrollable uncertainties in attackers' knowledge and strategies. Thus, the equilibrium would be extremely difficult to be computationally tackled because the state space of the game explodes in a drastic pattern due to the knowledge sharing.

Algorithm 2: PSO algorithm (Kennedy and Eberhart 1995).

```

1: Initialize the particles  $k$  which represents a provisional
   strategy  $S$ ,  $p_{best}^k$  and  $g_{best}$ .
2: for  $iter = 1$  to  $I$  do
3:   for each particle  $k$  do
4:     Obtain defender's utility  $U_k^d$  according to Algo. 1.
5:     if defender's utility  $U_k^d > p_{best}^k$  then
6:        $p_{best}^k = U_k^d$ .
7:     end if
8:   end for
9:   Set  $g_{best} = \max\{p_{best}^k\}$ .
10:  for each particle  $k$  do
11:    Calculate the velocity according to
        $v = \omega v + c_1 r_1 (p_{best} - k) + c_2 r_2 (g_{best} - k)$ .
12:    Update the particle according to  $k = k + v$ .
13:  end for
14: end for

```

Evaluating OPRADI

OPRADI's effectiveness is evaluated in both simulated environments and real-world contexts. The results suggest that it can provide improved strategic advice to the police, thus having significant potential in real-world applications. All computations were running on a machine equipped with a 2.8GHz AMD EPYC7543 32-core CPU and 128G memory.

Evaluating OPRADI in Simulated Environments

OPRADI was first evaluated with a simple 3×3 grid map. We simulated 10000 drivers (sober drivers and DUI-drivers)

Parameter	Description	Value
I	Number of iterations	100
ω	Inertia coefficient	0.8
c_1	Acceleration coefficient of p_{best}	0.5
c_2	Acceleration coefficient of g_{best}	0.5

Table 3: The PSO algorithm's parameters.

as attackers with randomly determined sources and destinations. We assessed OPRADI's effectiveness by comparing the defender's payoffs from the following three strategies and the OPRADI strategy.

- *Random Strategy*: Randomly choosing roads as checkpoints.
- *Common Strategy*: Always setting up checkpoints in some fixed places.
- *Stackelberg Strategy*: The defender's optimal strategy under the assumption that attackers have perfect knowledge of the defender's strategy.

Experiment Procedures: For the above three strategies, we ran simulations to determine the defender's payoffs. For OPRADI, we first calculated the defender's strategy and then the payoff corresponding to the strategy. Since we were performing evaluations in simulated environments, no empirical data was necessary. The key parameters of OPRADI were set as follows: (1) an attacker's number of friends (with who would share knowledge) is sampled from a uniform distribution $U[0, F]$ where $F = 30$; (2) confidence in the knowledge received $\lambda = 0.4$; (3) a DUI-driver's reward and penalty were the default values as shown in Tab. 2. Besides, when using PSO in Stackelberg Strategy and OPRADI related calculations, its parameters were summarized in Tab. 3.

Experiment Results: Fig. 3 and Fig. 4 depict the main results we obtained. We ran 100 trials in the simulated environment and the cumulative payoff of OPRADI is 565.81 ± 8.64 , outperforming the others (*Stackelberg*: 392.12 ± 5.66 , *Random*: 337.59 ± 31.59 , *Common*: 207.62 ± 1.01). Fig. 3 shows that the slope of OPRADI's payoff curve remains stable as the number of DUI-drivers increases, indicating OPRADI is more resistant to knowledge sharing than other strategies. In contrast, for *Common Strategy*, attackers could easily learn the defender's strategy through social knowledge sharing and bypass these roads since the defender is always locked in fixed places.

Fig. 4 reports the evaluation of OPRADI's reliability and scalability in different settings of the number of defender resources and the number of DUI-drivers. The x-axis is the number of checkpoints and the y-axis is the utility of the defender in different settings. The colors of the blocks represents the number of DUI-drivers in all simulated drivers. We can conclude that OPRADI is reliable and scalable in multiple settings since (1) increasing police resources constantly yields higher payoffs for the defender, and (2) the number of DUI-drivers being caught is proportional to the corresponding population sizes.

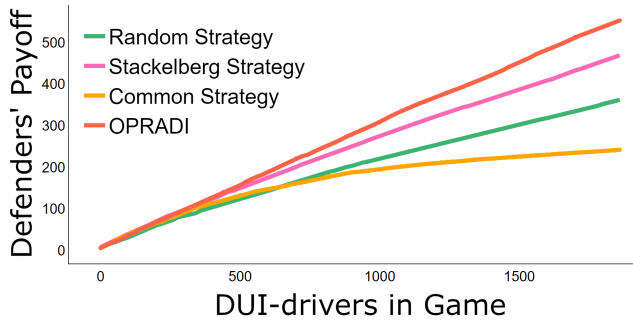


Figure 3: The dynamics of defender’s cumulative payoffs with different strategies.

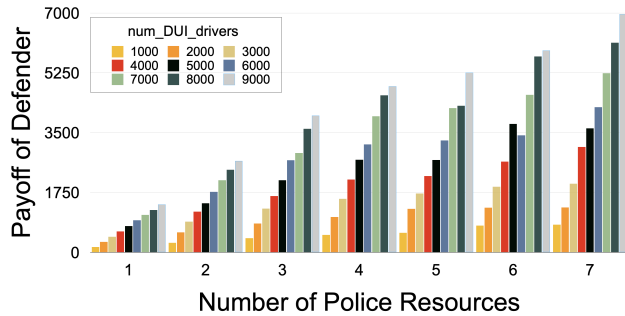


Figure 4: OPRADI’s performances with different police resources and DUI-drivers.

Further Analysis: In the above experiment with the small 3×3 grid map, we noticed that the differences in performances between OPRADI and other strategies became much more distinguishable as the increase of the number of DUI drivers. In such a small road network, the information about the police’s strategy could be learned pretty fast when there are many DUI drivers because the police, as defenders, have very limited strategic choices. To further test such a proposition, we ran another experiment by reusing the settings of the 3×3 grid map experiment but on a larger, more complex Chiping’s map. Note that we did not bring any empirical data from Chiping into this experiment. The experiment results of 100 trials are shown in Tab. 4. Obviously, OPRADI’s advantages shrank in a larger and more complex map though still outperformed the other three strategies. When maps become larger and more complex, it would take DUI drivers more

Strategies	Expectations	SDs
<i>Common Strategy</i>	86.62	10.18
<i>Random Strategy</i>	100.59	1.07
<i>Stackelberg Strategy</i>	103.14	1.74
OPRADI	108.57	2.40

Table 4: The defender’s expected payoffs with different strategies on Chiping Map without empirical data.

time to learn the police’s strategies no matter which strategy the police used.

Evaluating OPRADI in Contexts Reconstructed with Real-world Data

To further evaluate OPRADI’s in the environment similar to real-world, we conducted experiments with real-world data from Chiping District, a small city in Shandong Province, China, with assistance from the local police bureau. An author paid several field visits to the city during late 2021 and early 2022, and negotiated a collaboration agreement with the local police bureau. Then, the research team and a few police officers executed OPRADI’s three stages together.

Experiment Procedures: In stage 1, the local police bureau provided summary traffic flow statistics extracted from the traffic videos collected at the city’s three key crossings between 5 P.M. and 1 A.M. in 2022/06 (see attached data file). The research team also joined a team of plainclothes police officers in sampling observations to identify potential DUI-drivers in some gathering places, and we obtained the percentage of DUI drivers is about 10%. With these two types of data, we estimated the distributions of DUI-drivers over all drivers. Then, we consulted the police officers if these estimations fit their impressions and received affirmative replies. In addition, our police colleagues also offered four recent DUI inspections’ data, which serves as an ideal real-world benchmark to evaluate OPRADI.

Stage 2’s first task was to build an abstract street map from the city’s urban area’s road system. An author manually implemented this transition. The original road system and the abstract map are shown in Fig. 5b & 5c. The insights obtained from the stage 1 were also added to the map. Then, we reused the default knowledge sharing mechanism presented before. According to our observation, drivers often used the group chat function in their social network apps (e.g., WeChat) to share DUI inspection information to multiple people simultaneously rather than sending messages to everyone. So, the cost of sharing such knowledge was minimal. We conveniently set $F = 30$ (upper bounds of the number of friends receiving messages) and $\lambda = 0.4$ (confidence in the knowledge received). The numeric parameters for the SSG-DUI game were set to the default values shown in Tab. 2. We explained the numeric parameters and meanings to our police colleagues, who later confirmed that the values were reasonable. After this stage, the analytical framework was successfully assembled.

Stage 3 was responsible for deriving actionable advice to the police. As we mentioned above, this task was fulfilled by the PSO algorithm, which kept optimizing provisional strategies by interacting with the provisional strategy evaluation module. The provisional strategy evaluation module took charge of calculating the defender’s utility with a specific provisional strategy. The initial provisional strategy was randomly generated. The PSO algorithm’s parameters reused that in Tab. 3.

Experiment Results: OPRADI was evaluated against two benchmarks in the simulated environment similar to real-world. The first is “*Random Strategy*” which the defender randomly decided where to set up checkpoints. The expected




Figure 5: (a): A police is inspecting DUI-drivers at a checkpoint; (b): The satellite map of Chiping, whose major roads in its urban area are marked by yellow lines; (c): The abstract street map and the brown nodes refer to the places where DUI-drivers frequently haunt.

real-world utilities of these two strategies were calculated in the simulated environment reconstructed from the data collected in Chiping. The second is “*Current Strategy*” actually used by the police in Chiping. As mentioned above, we obtained “*Current Strategy*” and its payoffs from our police colleagues’ recent four DUI inspections. The utility of *Current Strategy* is calculated based on the number of DUI-drivers who were actually caught by Chiping police in these four DUI inspections.

Tab. 5 reports the expected performances of strategies in real-world contexts. The expectations and the standard deviations of random strategy and OPRADI were obtained from 100 trials in the simulated environment. Apparently, OPRADI brought significantly higher payoffs than the other two did. Compared with the *Random Strategy*, its payoff achieved a 27.46% improvement (112.59 vs. 88.33). Compared with the *Current Strategy*, its payoff was 2.75 times higher (112.59 vs. 30), suggesting that deploying OPRADI in the police’s practices would enable them to make much better use of their limited resources. It is not surprising that the *Current Strategy* recorded the lowest payoff. To some degree, it is very similar to the common strategy in the simulated experiment, which almost has no resistance to knowledge sharing among attackers. We went back to the empirical data shared by our police colleagues, and found that the *Current Strategy* did perform poorly in practice, often could only catch a very small proportion (<5%) of potential DUI drivers.

Strategies	Expectations	SDs
<i>Random Strategy</i>	88.33	1.52
<i>Current Strategy</i>	30 ³	
OPRADI	112.59	1.93

Table 5: The defender’s expected performances with different strategies.

³According to Chiping police’s historical inspection activities in June and July 2022, catching 30 DUI drivers in a single inspection activity was the best performance of *Current Strategy*.

	checkpoints	Caught DUI-drivers
2022/06-07	7	58
2022/08	6	133

Table 6: The results of OPRADI’s Preliminary Deployment.

Preliminary Deployments

Our police colleagues recently started testing OPRADI in their practices. In 2022/08, our police colleagues conducted two DUI inspection activities. During each inspection activity, they set up three checkpoints according to the recommended strategies obtained from the simulated experiments. Tab. 6 summarizes the performances of these two inspection activities which evidence significant improvements, as compared to the inspection activities in June and July.

As we mentioned above, local police employed *Common Strategy* (usually setting up checkpoints at some fixed locations) in inspections happened in 06 and 07/2022. They conducted four inspection activities setting up seven checkpoints, catching 58 DUI drivers in total. Thus, a single checkpoint only caught 8.14 (58/7) DUI-drivers on average. In 2022/08, with the help of OPRADI, a single checkpoints caught 22.17 (133/6) DUI drivers on average. While these results are preliminary, they provide some support to the effectiveness of OPRADI in real-world applications. The 2.72-time improvement (22.17 vs. 8.14) in each checkpoint’s payoff is non-trivial, also consistent with the performances in evaluating OPRADI in contexts reconstructed with real-world data.

Reflections

In the above section, both the second evaluation study and the preliminary deployment provided solid support to OPRADI’s effectiveness. With our police colleagues, the rationales behind OPRADI’s outstanding real-world performances were briefly analyzed in a qualitative fashion. We found that OPRADI’s advantages were likely to result from the following two considerations.

Making the Knowledge Sharing Explicit

First, OPRADI made the knowledge sharing among attackers explicit in the underlying game model (SSG-DUI). Therefore, when computing the police’s strategies, knowledge sharing must be taken into consideration. Compared with the other three strategies, explicit knowledge sharing forced OPRADI to yield strategies to have some capabilities to resist knowledge sharing. Our police colleagues told us that the strategy provided by OPRADI for DUI inspections, often coincided with some unavoidable paths for a driver leaving drinking avenues, or some positions critical to the traffic network but had been neglected before. Let us have a look at the following intuitive example.

Example: Two drivers d_1 (sober driver in the right green car) and d_2 (DUI driver in the left red car) were driving home (from A to B shown in Fig. 6). d_1 would share the location of the checkpoints on road y with d_2 after his observation. d_2 would choose another route $x - z$ to avoid being




Figure 6: A simple example showing how a strategy could be resistant to knowledge sharing among drivers.

inspected by the police. However, if the checkpoint was set up on road x , DUI driver d_2 would have no way to escape even after he got the checkpoints' location from d_1 . In other words, setting up checkpoint at road x was a better inspection strategy with more resistant to knowledge sharing.

The above simple example demonstrates how the strategies could resist knowledge sharing. In real-world contexts, the situations could be much more complex than the example. However, the rationale of making knowledge sharing explicit and considering it in identifying the strategies should bring better performances in catching DUI-drivers.

Integrating with Real-World Data and Knowledge

The second advantage lies in OPRADI's emphasis on integrating abstract models with real-world data and knowledge. As we have seen in the "Further Analysis" part of the first evaluation study, simply using a city's abstract map did not bring significant improvements in OPRADI's performance. Nevertheless, the extensive empirical data and knowledge introduced in the first two stages were the critical factors enabling OPRADI's success in its preliminary deployment. They gave the realistic contexts and the domain-specific hints to the algorithm part of OPRADI to figure out strategies that are resistant to attackers' knowledge-sharing.

Of course, such data and knowledge were not free or came from nowhere. They were indeed the product of collaboration between the research team and the local police bureau. Such collaboration cost multiple rounds of negotiations and months of efforts. Even establishing some common grounds took a lot of work. However, all efforts reaped the rewards. The police, who were experts of the local city, had rich data and knowledge. But they were often tacit and had never been put into use. With the OPRADI approach, researchers helped codify and integrate them into SSG-DUI models, realizing Michael Polanyi's knowledge transitions from the tacit dimension to the explicit dimension (Polanyi 2009). In return, the police's DUI inspection efforts become more effective. Thus, the close collaboration between artificial intelligence researchers (incl. ourselves) and domain experts in deploying any AI systems in the real-world shall be proactively pursued.

Conclusion and Future Work

The current practices of employing sobriety checkpoints for fighting DUIs are most dominated by the police's experience and arbitrary judgments, lacking an appropriate and rigorous game-theoretic model and analytical framework. This paper fills such a gap by proposing SSG-DUI game, a novel extension of the classic SSGs, and a corresponding analytical approach OPRADI. By incorporating DUI-drivers bounded rationality and social knowledge sharing, SSG-DUI achieves the balance between the game-theoretic model's analytical power and high fidelity in abstracting real-world phenomena. The OPRADI approach built on SSG-DUI provides effective, actionable solutions to improve police DUI-Inspection operations by identifying more effective strategies, confirmed by our extensive evaluations. OPRADI has been deployed in Chiping's police bureau's operations and obtained encouraging performances. We will continue to monitor its deployment closely and plan for applications in other cities.

For future work, SSG-DUI is flexible to accommodate other factors in real-world DUI, e.g., a dynamic driver population. Future work may also include exploring and integrating more sophisticated knowledge-sharing mechanisms (Bondi et al. 2020) and optimization algorithms (Liang et al. 2006) to improve OPRADI's performance.

Acknowledgements

We would like to thank the help from two mentors: Drs. Sean McGregor and Bradley J. Clement, as well as the anonymous reviewers' for their insightful suggestions. We are in debt to the Police Bureau of Chiping District for their assistance throughout the project. This work is partially supported by National Natural Science Foundation of China under grants 62076232 and 62172049. Corresponding author: Wei Wang.

References

- Alvim, M. S.; Chatzikokolakis, K.; Kawamoto, Y.; and Palamidessi, C. 2017. Information Leakage Games. In *Proc. of International Conference on Decision and Game Theory for Security*, 437–457.
- Arieli, I.; and Young, H. P. 2016. Stochastic Learning Dynamics and Speed of Convergence in Population Games. *Econometrica*, 84(2): 627–676.
- Bondi, E.; Oh, H.; Xu, H.; Fang, F.; Dilkina, B.; and Tambe, M. 2020. To Signal or Not to Signal: Exploiting Uncertain Real-Time Information in Signaling Games for Security and Sustainability. In *Proc. of the AAAI Conference on Artificial Intelligence*, 02, 1369–1377.
- Clark, A.; Zhu, Q.; Poovendran, R.; and Başar, T. 2012. Deceptive Routing in Relay Networks. In *Proc. of International Conference on Decision and Game Theory for Security*, 171–185.
- Elder, R. W.; Shults, R. A.; Sleet, D. A.; Nichols, J. L.; Zaza, S.; and Thompson, R. S. 2002. Effectiveness of Sobriety Checkpoints for Reducing Alcohol-Involved Crashes. *Traffic Injury Prevention*, 3(4): 266–274.

- Erke, A.; Goldenbeld, C.; and Vaa, T. 2009. The Effects of Drink-Driving Checkpoints on Crashes—A Meta-Analysis. *Accident Analysis & Prevention*, 41(5): 914–923.
- Fang, F.; Stone, P.; and Tambe, M. 2015. When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing. In *Proc. of International Conference on Artificial Intelligence*, 2589–2595.
- Guo, Q.; An, B.; Zick, Y.; and Miao, C. 2016. Optimal Interdiction of Illegal Network Flow. In *Proc. of International Joint Conference on Artificial Intelligence*, 2507–2513.
- Jie, Y.; Li, M.; Tang, T.; and Guo, C. 2017. Optimal Allocation Astrategy Based on Stackelberg Game for Inspecting Drunk Driving on Traffic Network. *KSII Transactions on Internet and Information Systems (TIIS)*, 11(12): 5759–5779.
- Jie, Y.; Liu, C. Z.; Li, M.; Choo, K.-K. R.; Chen, L.; and Guo, C. 2020. Game Theoretic Resource Allocation Model for Designing Effective Traffic Safety Solution Against Drunk Driving. *Applied Mathematics and Computation*, 376: 125142.
- Kar, D.; Nguyen, T. H.; Fang, F.; Brown, M.; Sinha, A.; Tambe, M.; and Jiang, A. X. 2018. Trends and Applications in Stackelberg Security Games. In *Handbook of Dynamic Game Theory*, 1223–1269. Springer International Publishing.
- Kennedy, J.; and Eberhart, R. 1995. Particle Swarm Optimization. In *Proc. of International Conference on Neural Networks*, 1942–1948.
- Liang, J.; Qin, A.; Suganthan, P.; and Baskar, S. 2006. Comprehensive Learning Particle Swarm Optimizer for Global Optimization of Multimodal Functions. *IEEE Transactions on Evolutionary Computation*, 10(3): 281–295.
- Nguyen, T.; Yang, R.; Azaria, A.; Kraus, S.; and Tambe, M. 2013. Analyzing the Effectiveness of Adversary Modeling in Security Games. In *Proc. of AAAI Conference on Artificial Intelligence*, 718–724.
- NHTSA. 2022. Alcohol-Impaired Driving. <https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813294>. Accessed: 2022-8-06.
- Paruchuri, P.; Pearce, J. P.; Tambe, M.; Ordonez, F.; and Kraus, S. 2007. An Efficient Heuristic Approach for Security against Multiple Adversaries. In *Proc. of International Joint Conference on Autonomous Agents and Multiagent Systems*.
- Passino, K. 2002. Biomimicry of Bacterial Foraging for Distributed Optimization and Control. *IEEE Control Systems Magazine*, 22(3): 52–67.
- Pawlick, J.; and Zhu, Q. 2017. A Mean-Field Stackelberg Game Approach for Obfuscation Adoption in Empirical Risk Minimization. In *Proc. of IEEE Global Conference on Signal and Information Processing*, 518–522.
- Pita, J.; Jain, M.; Marecki, J.; Ordóñez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. Deployed ARMOR Protection: The Application of a Game Theoretic Model for Security at the Los Angeles International Airport. In *Proc. of International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track*, 125–132.
- Polanyi, M. 2009. *The Tacit Dimension*. The University of Chicago Press.
- Rosenfeld, A.; and Kraus, S. 2017. When Security Games Hit Traffic: Optimal Traffic Enforcement under One Sided Uncertainty. In *Proc. of International Joint Conference on Artificial Intelligence*, 3814–3822.
- Roy, S.; Kadir, S. U.; Vorobeychik, Y.; and Laszka, A. 2021. Strategic Remote Attestation: Testbed for Internet-of-Things Devices and Stackelberg Security Game for Optimal Strategies. In *Proc. of International Conference on Decision and Game Theory for Security*, 271–290.
- Sanfey, A. G. 2007. Social Decision-Making: Insights from Game Theory and Neuroscience. *Science*, 318(5850): 598–602.
- Sengupta, S.; Chowdhary, A.; Huang, D.; and Kambhampati, S. 2018. Moving Target Defense for the Placement of Intrusion Detection Systems in the Cloud. In *Proc. of International Conference on Decision and Game Theory for Security*, 326–345.
- Shen, W.; Chen, W.; Huang, T.; Singh, R.; and Fang, F. 2020. When to Follow the Tip: Security Games with Strategic Informants. In *Proc. of International Joint Conference on Artificial Intelligence*, 371–377. Main track.
- Shokri, R. 2015. Privacy Games: Optimal User-Centric Data Obfuscation. *Proc. of Privacy Enhancing Technologies*, 2015(2): 299–315.
- Simon, H. A. 1990. *Bounded Rationality*. Springer.
- Sinha, A.; Fang, F.; An, B.; Kiekintveld, C.; and Tambe, M. 2018. Stackelberg Security Games: Looking Beyond a Decade of Success. In *Proc. of International Joint Conference on Artificial Intelligence*, 5494–5501.
- Von Neumann, J.; and Morgenstern, O. 1947. *Theory of Games and Economic Behavior*, 2nd rev. Princeton university press.
- Yao, J.; Johnson, M. B.; and Beck, K. H. 2014. Predicting DUI Decisions in Different Legal Environments: Investigating Deterrence with a Conjoint Experiment. *Traffic Injury Prevention*, 15(3): 213–221.
- Yin, Z.; Korzhyk, D.; Kiekintveld, C.; Conitzer, V.; and Tambe, M. 2010. Stackelberg vs. Nash in Security Games: Interchangeability, Equivalence, and Uniqueness. In *Proc. of International Conference on Autonomous Agents and Multiagent Systems*, 1139–1146.
- Yoo, J. H.; and Langari, R. 2013. A Stackelberg Game Theoretic Driver Model for Merging. In *Proc. of Dynamic Systems and Control Conference*.
- Zhu, Q.; and Başar, T. 2013. Game-theoretic Approach to Feedback-Driven Multi-stage Moving Target Defense. In *Proc. of International conference on decision and game theory for security*, 246–263.
- Zhu, Q.; Clark, A.; Poovendran, R.; and Başar, T. 2012. Deceptive Routing Games. In *Proc. of IEEE Conference on Decision and Control*, 2704–2711.