

TOMASZ P. MICHALAK
MICHAŁ T. GODZISZEWSKI
ANDRZEJ NAGÓRKO

Protecting critical infrastructure with game theory, optimization techniques, and AI algorithms

Abstract

In light of recent geopolitical developments, Europe and Poland are acutely aware of the urgent importance of infrastructure security. Despite heightened interest and increased investments, our security resources remain severely limited, rendering continuous protection for every potential target unattainable. Consequently, the strategic allocation of security resources becomes an ongoing imperative. This paper presents a short introduction to the core principles behind advanced methods that facilitate automated decision-making in security resource allocation. These methods leverage artificial intelligence (AI), game theory, and optimization techniques, and have demonstrated their effectiveness through multiple real-life deployments in the USA. We also provide a concise overview of this exciting body of research and discuss the solutions and software developed by our team, “AI for Security” at the IDEAS NCBR research institute to protect critical infrastructure in Poland and in Europe.

Keywords:

optimization,
security games,
artificial
intelligence,
critical
infrastructure

Introduction

Los Angeles International Airport (LAX) is one of the busiest and largest airports in the world, serving as a vital transportation hub for the city of Los Angeles and the surrounding region. In terms of the number of passengers, it is about four times larger than the Frederic Chopin Airport, Warsaw, which is the biggest Polish airport. The LAX airport encompasses a vast area and features four parallel runways and nine terminals each serving different airlines and destinations. The largest one is the Tom Bradley International Terminal dedicated to international flights. The Central Terminal Area serves as the focal point of the airport connecting all the terminals. It includes a complex network of roadways, parking structures, and transportation services, such as shuttles and taxis, to facilitate passenger movement around the airport.

Given its prominence and size, LAX is one of the prime targets on the West Coast of potential attacks. Safeguarding such a complex and sprawling facility requires a delicate balance between security measures and operational efficiency. Unfortunately, available security resources are limited, making it impossible to provide round-the-clock security for each and every place of interest. For instance, while the number of canines exceeds the number of terminals, there are only a handful of canines of particular expertise, such as explosive detection ones. It means that it is simply impossible to provide constant coverage of a single explosive detection canine patrol per each terminal at LAX. The same hold for other types of canine patrols, such as those specialised in drug detection.

Interestingly, however, in 2008, the drug sniffing canine patrols at the LAX airport turned out to be much more effective than it was previously believed possible. While in the 15 month period between April 2006 to July 2007, only 4 drug-related offences were recorded, in the 12 months of 2008 the number of cases grew to 30. The reason behind this significant improvement was ARMOR which stands for the Assistant for Randomized Monitoring over Routes. The ARMOR system was an innovative software tool developed by Milind Tambe and his colleagues at the University of Southern California (USC)¹, within the Department of Homeland Security's first University Center of Excellence.

¹ J. Pita et al., *Using game theory for Los Angeles Airport security*, "AI Magazine" 2009, vol. 30, no. 1, pp. 43–57.

The primary objective of the ARMOR system is to assist security personnel in making better and more efficient decisions by weighing risks against available resources. To this end, ARMOR leverages artificial intelligence (AI), game theory, and optimization techniques. The system makes it difficult for adversaries to plan how to avoid security forces during an attack. Even more importantly, it allows security forces to deploy their limited resources in the most efficient manner and achieve maximal effectiveness.

The implementation of ARMOR at LAX resulted in the improvement in security coverage, resource allocation efficiency, and deterrence against potential attackers. For example, Chart shows that the introduction of the ARMOR system at the LAX Los Angeles Airport resulted in more than threefold increase in the number of detected offences. The system serves as an exemplar of how advanced technology and AI-driven approaches can contribute to strengthening security measures and enhancing the safety of airports and their passengers.




Chart. The number of offences in the period of 15 months before introducing ARMOR (dark blue bars) vs. the number of offences in the period of 12 months after introducing the system (light blue bars).

Source: own elaboration based on: Pita et al., *Using game theory for Los Angeles Airport security*, “AI Magazine” 2009, vol. 30, no. 1, pp. 43–57.

The success of ARMOR received significant attention in the context of security applications. A number of systems based on similar principles were deployed in the USA to protect other critical infrastructure sites. These include, in particular:

- IRIS system² – to optimize the routes and schedule of the security agents in the U.S. Air Marshals program;
- PROTECT³ – to optimize the security of the Boston and New York ports;
- TRUSTS system⁴ – to prevent fare evasion created for the railway transport system in Los Angeles.

Furthermore, it has been advocated in the context of cybersecurity⁵. There are also a growing number of civilian applications such as protecting endangered species in national parks (systems PAWS⁶ and MIDAS⁷). In all these cases, it was possible to significantly improve security, not by adding many additional security resources but by better deployment of the available ones.

This is a very important lesson for Europe and Poland in particular. Given recent geopolitical developments, and the on-going Russian full-scale invasion in Ukraine that started in 2022, we are all well aware of the pressing concern of infrastructure security. Europe has already witnessed a few such attacks⁸. The question now is not if the next attacks will happen but when.

² J. Tsai et al., *Iris - a tool for strategic security allocation in transportation networks*, in: *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2009)*, pp. 37–44 (the proceedings of the AAMAS conference series are available at: <https://dl.acm.org/conference/aamas/proceedings> – editor’s note).

³ E. Shieh et al., *Protect: A deployed game theoretic system to protect the ports of the United States*, in: *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2012)*, vol. 1, pp. 13-20.

⁴ Z. Yin et al., *Trusts: Scheduling randomized patrols for fare inspection in transit systems*, in: *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, vol. 26, no. 2, pp. 2348–2355 (the proceedings from the AAAI conferences and symposia are available at: <https://aaai.org/aaai-publications/aaai-conference-proceedings/> – editor’s note).

⁵ Y. Zhang, P. Malacaria, *Bayesian Stackelberg games for cyber-security decision support*, “Decision Support Systems” 2021, vol. 148, art. 113599. <https://doi.org/10.1016/j.dss.2021.113599>.

⁶ R. Yang et al., *Adaptive resource allocation for wildlife protection against illegal poachers*, in: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pp. 453–460.

⁷ W. Haskell et al., *Robust protection of fisheries with COMPASS*, in: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2014)*, vol. 28, no. 2, pp. 2978-2983.

⁸ An example is the deliberate cutting of two optical fibers of the Deutsche Bahn communication system on October 8, 2022, which stopped rail traffic in northern Germany for approximately 3 hours.

Unfortunately, the problem is exacerbated by the growing technological sophistication of critical infrastructure. While modern communication, computation, and control technologies enhance the efficiency, they also make modern critical infrastructure systems increasingly complex and vulnerable to deliberate attacks and random failures. These attacks can manifest in various forms, severities, and magnitudes, ranging from acts of terrorism on local infrastructure to major kinetic strikes during times of war, such as the ongoing Russian invasion on Ukraine. New technologies, such as drones, also enhance the capabilities of the potential attackers.

Facing the evolving and expanding landscape of threats, despite increased interest and investments in infrastructure security, our security resources will remain limited, making it impossible to provide constant protection for everything. Thus, the need to strategically allocate security resources becomes a perpetual necessity. The example of the LAX airport as well as other aforementioned examples from the USA show that such well-designed strategic decision making is beneficial and delivers significant improvement in security.

In this paper, we discuss the fundamentals behind these advanced methods of protecting critical infrastructure, briefly review this body of research, and present the solutions and software which is developed by our team “AI for security” at the IDEAS NCBR research institute.

Defender-Attacker Security Games

Game theory studies interactions between intelligent entities like individuals, companies, or countries. In the context of security, these entities can represent “the defenders”, e.g., security forces, police, military, and “the attackers”, e.g., criminals, terrorists, and state actors. Game-theoretic approaches help us understand how these intelligent actors interact, assuming they can anticipate and respond to each other’s actions. By using game theory, we can develop a strategy to efficiently distribute scarce security resources for infrastructure protection. This approach considers the importance of different targets and how adversaries may react to specific protection strategies.

A non-cooperative game is defined by the set of players, the set of strategies for each player and the payoff function that assigns each

player the utility for any possible combination of strategies. Each game is also associated with some set of rules, e.g., we may require the players to move simultaneously or sequentially.

Table 1 presents a sample game from Pita et al.⁹ In this case, we have two players, each having two strategies: $\{A, B\}$ and $\{C, D\}$, respectively. The values of the payoff function are given by the pairs of numbers in the matrix, where each cell corresponds to a given combination of strategies, also known as a 'strategy profile'. For instance, if Player 1 plays strategy A and Player 2 plays strategy C then the Player 1 earns the payoff of 2 while Player 2 the payoff of 1.

Sometimes it is possible to stipulate how the rational players (where the notion of rationality is explicated with a precise mathematical formula) actually would play the game given its rules. Such a collection of players' strategies is called an equilibrium of the game. Perhaps the most well-known equilibrium concept is the Nash equilibrium. A combination of strategies is a Nash equilibrium if any player would not like to change their strategy, given the strategies chosen by the opponents. For instance, in Table 1, the combination of strategies (A, B) is not a Nash equilibrium, because Player 2 would like to change their strategy from D to C , assuming that Player 1 sticks to strategy A . Conversely, the combination of strategies (A, C) is a Nash equilibrium, because, for Player 1, A is the best strategy if Player 2 plays C , and, for Player 2, C is the best strategy if Player 1 plays A .

Table 1. A payoff matrix for a sample game.

		Player 2	
		C	D
Player 1	A	(2,1)	(4,0)
	B	(1,0)	(3,2)

Source: J. Pita et al., *Using game theory for Los Angeles Airport security*, "AI Magazine" 2009, vol. 30, no. 1, pp. 43–57.

Players in a non-cooperative game do not have to focus on particular strategies. Instead of choosing a single strategy with certainty, a player can choose one strategy with some probability, other strategy with other probability, and so on and so forth. That is, a player can assign

⁹ Pita et al., *Using game theory for Los Angeles Airport...*

probabilities to each available strategy. For instance, Player 1 may choose to play strategy A with a certain probability, denoted by p , and strategy B with a probability of $1 - p$. Similarly, Player 2 can assign probabilities to strategies C and D. By using mixed strategies, the players introduce randomness into their decision-making process. The concept of the Nash equilibrium also extends to the mixed strategies.

Consider the game with the payoff matrix defined in Table 2.

Table 2. An example of payoff matrix for a game with no Nash equilibrium in pure strategies, but with a Nash Equilibrium in mixed strategies.

		Player 2	
		C	D
Player 1	A	(2,1)	(1,2)
	B	(1,2)	(3,1)

In this game, there is no pure strategy Nash equilibrium, since for every strategy profile, it is beneficial for one of the players to switch their strategy, if the other player's strategy is fixed, but there exists a mixed strategy Nash equilibrium, namely the following:

- Player 1's mixed strategy: A with probability $\frac{1}{2}$, B with probability $\frac{1}{2}$;
- Player 2's mixed strategy: C with probability $\frac{2}{3}$, D with probability $\frac{1}{3}$.

The expected payoff for Player 1 in the equilibrium is:

$$\frac{1}{2} \times \frac{2}{3} \times 2 + \frac{1}{2} \times \frac{1}{3} \times 1 + \frac{1}{2} \times \frac{2}{3} \times 1 + \frac{1}{2} \times \frac{1}{3} \times 3 = \frac{5}{3}$$

and the expected payoff for Player 2 is:

$$\frac{1}{2} \times \frac{2}{3} \times 1 + \frac{1}{2} \times \frac{1}{3} \times 2 + \frac{1}{2} \times \frac{2}{3} \times 2 + \frac{1}{2} \times \frac{1}{3} \times 1 = \frac{3}{2}$$

The above model is, of course, simplified. In particular, in the case of protecting critical infrastructure, one may argue that the players do not move simultaneously. This is because, an attacker may be able to observe the defensive measures (strategies) employed by a defender. To address this problem, let us consider a seminal economic model proposed by Stackelberg¹⁰, a game is played between two players: a leader and a follower.

¹⁰ H. von Stackelberg, *Marktform und Gleichgewicht*, J. Springer 1934.

That is, in contrast to the previous example, the Stackelberg game is played sequentially rather than simultaneously. Specifically, the leader selects their strategy first, and this choice is observed by the follower, who subsequently determines their own move accordingly.

The Stackelberg model has garnered substantial attention within the realm of security applications due to its inherent ability to capture the dynamics of defender-attacker interactions. In this context, Stackelberg games are often called security games.

The model encompasses the following aspects:

- the defender, who assumes the role of the leader in the Stackelberg game, allocates limited security resources to protect a designated set of targets. Recognizing that adversaries possess the capability to observe defense strategies and exploit discernible patterns, the defender naturally opts for a mixed (randomized) strategy. For instance, in the case of LAX, the management of the canine unit determines the frequency of visits to each terminal by a specific type of patrol within a given week. In other words, they establish the probability distribution for each patrol type across all terminals;
- the attacker, acting as the follower in the Stackelberg game, observes the defender's chosen strategy, i.e., these probability distributions. This assumption embodies a prudent and realistic scenario, presupposing an intelligent attacker who thoroughly surveys critical infrastructure before devising and executing an attack;
- lastly, having acquired knowledge of the probabilities selected by the defender, the attacker strategically selects the optimal course of action for themselves and subsequently executes their move accordingly.

It is crucial to emphasize that the attacker has the capability to observe the probability distribution chosen by the defender but not the defender's actual move. As an illustration, let us consider an scenario involving the United States Coast Guard (USCG) responsible for patrolling the Mexican Bay to combat drug trafficking via boats. The smugglers can observe the frequency of patrols in specific sea areas and how frequently patrol boats alter their course, e.g., by changing the patrol direction to a different one. In other words, the attackers have knowledge of the probability distribution. Nonetheless, they lack the capacity to predict whether a patrol boat will change its course at a given moment or not. So they cannot simply wait until a patrol boat goes away as there is

a non-zero probability that it may immediately return. Interestingly, as we already mentioned in the introduction, the system called PROTECT, based on the Stackelberg game was introduced by the USCG to enhance port/coastal security.

One may say that the defenders of a critical infrastructure are at disadvantage as they move first (decide on the allocation of defense resource and the probability distributions) and their move is observed by the attacker. However, a more careful analysis reveals that the defender, as the first mover, may have a significant influence on the choices made by the attacker. Intuitively, the defender may push the attacker to choose one strategy not the others.

Table 3. A payoff matrix for a sample game.

		Player 2	
		C	D
Player 1	A	(1,1)	(3,0)
	B	(0,0)	(2,1)

Source: D. Korzhyk et al., *Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness*, "Journal of Artificial Intelligence Research" 2011, vol. 41, no. 2, pp. 297–327.

As an example, let us assume that Player 1 in Table 3 is the leader in the Stackelberg game. Observe that if the players move simultaneously, then actually the only Nash equilibrium (in pure strategies - observe that, trivially, every Nash equilibrium in pure strategies is also a Nash equilibrium in mixed strategies) of this game is for Player 1 to play A and Player 2 to play C, which gives Player 1 the expected payoff equal to 1. Now, by the power of moving first, Player 1 can choose a uniform mixed strategy of playing A and B with equal probability of $1/2$, instead of A with probability 1 and B with probability 0 as in the case of the Nash equilibrium for a simultaneous game. The choice of the leader pushes the follower (Player 2) to choose strategy D instead of C¹¹. In result, by being the leader, Player 1 can secure the expected payoff of $5/2$ instead of 1, which is quite a significant difference.

¹¹ In Stackelberg Games, the assumption is that if the follower remains indifferent, the tie is resolved in favor of the leader, since otherwise the optimal solution is ill-defined.

In Appendix A we present a more formal introduction to security games. Let us now comment on the computational challenges posed by security games and then move on to the overview of approaches in the literature.

Challenges and approaches

The game-theoretic approach described in the previous section has been well-established in the literature for many years. However, it is only in the last two decades that these concepts have been effectively deployed in practice to protect critical infrastructure. The reason for this delayed implementation can be attributed to the computational challenges associated with security games. In this section, we will first discuss these challenges and explore how optimization and AI techniques have emerged as effective tools to address them. Furthermore, we will briefly review the existing lines of research on security games to shed light on the challenges involved in developing practical and feasible solutions.

Computational challenges

In real-life deployments, Stackelberg games pose significant computational challenges due to the following key factors:

- first and foremost, decision spaces in complex and large-scale environments of critical infrastructure are immense. The number of possible strategies and actions that can be taken by the players, such as defenders and attackers, can be truly enormous. For example, the New York City Subway system is one of the largest and busiest public transportation networks in the world, serving millions of commuters and visitors daily. The subway system comprises a vast network of tracks, stations, and interconnected lines, covering a total of 472 stations and over 800 miles (1,287 kilometers) of tracks. There are not only various layers of protective measures (i.e., enormous strategy space of the defender) but also very many attack options (i.e., enormous strategy space of the attacker). This requires efficient algorithms to explore and optimize such a vast decision spaces;
- this difficulty is further exacerbated by uncertainty and incomplete information regarding the intentions, capabilities, and actions

of adversaries. Bayesian security games (see the appendix) explicitly model this uncertainty using probability distributions. This, however, adds computational complexity to the problem;

- next, real-life situations are often dynamic and they constantly evolve. Adversaries may adapt their strategies, and defenders need to respond accordingly. Modeling and optimizing strategies in such dynamic environments require solving repeated or sequential games, which further increase the computational challenge.

The scientific literature took a few routes to deal with the computational challenge posed by security games. One approach is to employ mathematical optimization methods to solve these games efficiently. Researchers have developed algorithms and optimization techniques that can handle large-scale game models and provide solutions within reasonable time frames. These optimization methods exploit the structure of the game to reduce the computational burden and improve computational efficiency. They leverage mathematical programming, linear programming, integer programming, and other optimization frameworks to find optimal strategies and resource allocations.

Due to the inherent complexity of these games, finding exact solutions for large-scale scenarios is often infeasible. Therefore, researchers and practitioners often resort to developing approximation algorithms and heuristics to tackle computational challenges while maintaining a reasonable level of accuracy.

Furthermore, AI techniques may play an important role in enhancing the performance of the optimization algorithms. Using AI to optimize algorithms in general and optimization solvers in particular has led to improvements of the state of the art in the solving of hard computational problems for many years. AI allows existing approaches to scale better, and can be applied in many different contexts, for example the one we consider here, i.e., games and optimization¹². On the other hand, AI models can approximate the result of expensive computational processes quickly and reliably, allowing to do more with the same amount of resources. For example, when deciding what intervention to use to improve resilience and security, some alternatives will be obviously inferior. AI models can help to identify such inferior interventions quickly and cheaply, even accounting

¹² F. Hutter et al., *Boosting Verification by Automatic Tuning of Decision Procedures*, in: *Proceedings of the 19th International Conference on Computer Aided Verification (CAV 2007)*, pp. 27–34.

for the uncertainty of the approximation. The same kind of techniques allow systems like AlphaGO to explore a vast space of possible actions in seconds, used for example to deter wildlife poachers¹³.

Another option is to resort to parallel computing and distributed computing techniques. By distributing the computational workload across multiple processors or machines, much larger-scale game models can be handled. Here, advancements in hardware technology that improve parallel computing capabilities, are especially important.

Approaches in the Literature

A brief overview of the Stackelberg games with a couple of illustrative examples can be found in the work by Sinha et al.¹⁴ A major and very recent literature review concerning security games can be found in the paper by Hunt and Zhuang¹⁵. The review examines the present state-of-the-art in game-theoretic modeling for attacker-defender scenarios and analyzes the literature based on common application areas, modeling approaches, and solution methods, additionally addressing significant gaps in the existing body of research and providing a comprehensive discussion on future directions. Other extensive surveys concerning security games can be found in the work by Fang and Nguyen¹⁶ and in the one by Nguyen et al.¹⁷, where it is demonstrated that security agencies regularly employ decision aids based on game theory to optimize the allocation of limited security resources against strategic adversaries, and that the unique characteristics of these applications demand innovative solutions from AI systems.

¹³ S. Gholami et al., *Adversary models account for imperfect crime data: Forecasting and planning against real-world poachers*, in: *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, pp. 823–831.

¹⁴ A. Sinha et al., *Stackelberg security games: Looking beyond a decade of success*, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pp. 5494–5501.

¹⁵ K. Hunt, J. Zhuang, *A review of attacker-defender games: Current state and paths forward*, “*European Journal of Operational Research*” 2023, in press. <https://doi.org/10.1016/j.ejor.2023.04.009>.

¹⁶ F. Fang, T.H. Nguyen, *Green security games: Apply game theory to addressing green security challenges*, “*ACM SIGecom Exchanges*” 2016, vol. 15, no. 1, pp. 78–83. <https://doi.org/10.1145/2994501.2994507>.

¹⁷ T.H. Nguyen et al., *Towards a science of security games*, in: *Mathematical Sciences with Multidisciplinary Applications*, B. Toni (ed.), Springer Cham 2016, pp. 347–381.

Two already classical monographs have emerged as prominent references in the intersection of game theory and security. Tambe's book¹⁸ centers around algorithmic advancements and the adoption of game-theoretic software by government stakeholders. On the other hand, Bier and Azaiez's monograph¹⁹ presents a compilation of works that combine game theory and risk analysis within the realm of security.

The Stackelberg games have been increasingly employed to examine a wide array of security issues, spanning from scenarios like missile defense systems²⁰, terrorism²¹, policing²², to computer network security²³.

Sometimes the models of security games are being referred to as attacker-defender games. They have been a subject of extensive research for the past years and there is a large body of literature concentrating on variety of different problems. An example, again, might be a resource allocation model where e.g., a government wishes to allocate defensive resources among a set targets (e.g., airports or train stations) in an optimal way, and an adversary seeks to attack some of these targets. In the paper by An et al.²⁴, the authors present an overview of the aforementioned game-theoretic system PROTECT, utilized by the USCG for scheduling patrols in the Port of Boston and New York (see Image). Importantly, the successful evaluation of PROTECT in the Port of Boston has led to its further deployment in the Port of New York. The foundation of PROTECT lies exactly in the attacker-defender Stackelberg game model. However, the development and implementation of the system involved significant contributions in theory as well as comprehensive evaluations. What

¹⁸ M. Tambe, *Security and game theory: algorithms, deployed systems, lessons learned*, Cambridge 2011.

¹⁹ V.M. Bier, M.N. Azaiez, *Game Theoretic Risk Analysis of Security Threats*, Springer 2008. <https://doi.org/10.1007/978-0-387-87767-9>.

²⁰ G. Brown et al., *A Two-Sided Optimization for Theater Ballistic Missile Defense*, "Operations Research" 2005, vol. 53, no. 5, pp. 745–763. <https://doi.org/10.1287/opre.1050.0231>.

²¹ T. Sandler, *Terrorism & Game Theory*, "Simulation & Gaming" 2003, vol. 34, no. 3, pp. 319–337. <https://doi.org/10.1177/1046878103255492>.

²² N. Gatti et al., *Game theoretical insights in strategic patrolling: Model and algorithm in normal-form*, in: *Proceedings of the 2008 conference on ECAI 2008: 18th European Conference on Artificial Intelligence (ECAI 2008)*, pp. 403–407.

²³ K-w. Lye, J. Wing, *Game Strategies in Network Security*, "International Journal of Information Security" 2005, vol. 4, pp. 71–86. <https://doi.org/10.1007/s10207-004-0060-x>.

²⁴ B. An et al., *A Deployed Quantal Response-Based Patrol Planning System for the U.S. Coast Guard*, "Interfaces" 2013, vol. 43, no. 5, pp. 400–420. <https://doi.org/10.1287/inte.2013.0700>.

is crucial about the system, it does not assume that adversaries possess perfect rationality, allowing for more realistic and robust scenarios.



Image. The PROTECT system was deployed by the United States Coast Guards to protect the Staten Island Ferry route operated by the New York City Department of Transportation. The picture shows the United States Coast Guards boat protecting one of the ferry vessels.

It is noteworthy that numerous Stochastic Stackelberg Games exist, wherein the decision-making abilities of the adversary are limited by bounded rationality. The majority of systems based on Stackelberg games have traditionally relied on the conventional game-theoretical assumption of adversaries being perfectly rational, which aligns with the standard in game theory literature. However, this assumption may not accurately reflect the behavior of real-world adversaries, as human adversaries often exhibit bounded rationality. Therefore, taking inspiration from psychological and behavioral economics models, researchers (e.g. Yang et al.²⁵) have delved into studying parametrized models of bounded rationality in these games. The models in question offer versatile approaches for incorporating bounded rationality into game settings, making them applicable to a wide range of game interactions beyond Stackelberg Security Games. An example of such an approach is an instance studied in the work by

²⁵ R. Yang et al., *Improving Resource Allocation Strategy Against Human Adversaries in Security Games*, in: *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pp. 458–464.

Nguyen et al.²⁶, where rather than selecting a single target as the optimal response to the induced coverage C of targets by defense resources, the adversary's response $h(C)$ entails probabilistically choosing a target t based on a probability q_t associated with that target.

Let us conclude by noting that there are various further potential applications of Stackelberg Games in modelling adversarial security scenarios. They do include – among others – the following:

- Patrolling games by Vorobeychik et al.²⁷, designed to simulate situations where environments need to be patrolled to deter intruders. These games draw inspiration from the widely recognized pursuit-evasion game model but have been expanded in diverse ways, including the incorporation of alarm systems;
- In plan interdiction games by Vorobeychik and Pritchard²⁸, the defender is tasked with selecting a mitigation strategy to intercept potential attack actions, while the attacker, in response, devises an optimal attack plan that bypasses the implemented mitigations. This model finds relevance in the context of adversaries operating in cybersecurity;
- Audit games (J. Blocki et al.²⁹) investigate the economic aspects involved in designing audit mechanisms, with a specific emphasis on efficient resource allocation and suitable punishment schemes. The audit game model expands upon the security game model by introducing an extra parameter related to punishment. These models find practical application in audits aimed at ensuring compliance with privacy policies within diverse institutions, including medical hospitals;

²⁶ T.H. Nguyen et al., *Analyzing the effectiveness of adversary modeling in security games*, in: *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, no. 1, pp. 718–724.

²⁷ Y. Vorobeychik, B. An, M. Tambe, *Adversarial Patrolling Games*, in: *Papers from the 2012 AAAI Spring Symposium*, vol. 3, pp. 91–98.

²⁸ Y. Vorobeychik, M. Pritchard, *Plan interdiction games*, in: *Adaptive Autonomous Secure Cyber Systems*, S. Jajodia et al. (ed.), Springer Cham 2020, pp. 159–182. https://doi.org/10.1007/978-3-030-33432-1_8.

²⁹ J. Blocki et al., *Audit games with multiple defender resources*, in: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*, vol. 29, no. 1, pp. 791–797.

- Coalitional security games (Guo et al.³⁰) tackle the issue of optimizing the prevention of attacker coalitions, where attackers have the ability to form alliances. This concept is particularly relevant in domains such as disrupting terrorist networks, dismantling cells of these networks, or preventing collusion among multiple attackers.

“AI for Security” Team at the IDEAS NCBR

At the IDEAS NCBR research institute, our team “AI for Security” builds Stackelberg models for various types of critical infrastructure. Currently, we have been focusing on developing software for protecting ports, LNG terminals, railways, and power grids. Figure presents the interface of our basic software.




Figure. The snapshot from the Interface of the security-game software that is under development at IDEAS NCBR by the team “AI for Security”. Here, we can see the map of the LNG Terminal in Świnoujście. The red circles represent targets (the size of the circle corresponds to the importance of the target). The blue circles represent the positioning of the patrols and their field of view. The heatmap on the right shows the relative probability of which parts of the site should be patrolled (the optimal strategy of the defender). Note that the positioning of targets and patrols on this snapshot is for demonstration purposes only.

³⁰ Q. Guo et al., *Coalitional security games*, in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, pp. 159–167.

Our goal is to create a system with the following features:

- **Risk Assessment:** our system should conduct continuous risk assessments through the analysis of diverse data sources, current and historical ones, and intelligence reports. In order to optimize the allocation of security resources, it should take into account elements such as threat levels, target desirability, and vulnerabilities.
- **Randomized Patrol Strategy:** Our system should employ randomized strategies to determine the optimal patrol routes for security personnel coupled with optimal deployment of security devices. By randomizing routes, the system increases the difficulty for potential adversaries to predict security patterns, thereby strengthening the element of surprise and deterring potential threats.
- **Dynamic Adaptation:** our system should accommodate evolving threats and adapt to changing security scenarios. It should possess the capability to dynamically modify patrol routes and allocate resources in response to real-time information, such as emerging intelligence that updates the knowledge of the actual risk situation. This goal is to ensure optimal coverage and enhance response abilities.
- **Collaboration and Coordination:** the system should facilitate collaboration among diverse security teams and agencies operating within the protected area. It should enable the sharing of information, coordination of efforts, and real-time intelligence exchange, all aimed at enhancing situational awareness and achieving improved security outcomes.
- **Performance Evaluation and Feedback:** finally, it should incorporate mechanisms for evaluating performance, enabling security personnel to analyze the effectiveness of the system and adapt strategies accordingly. The system should offer feedback, identifying areas for improvement and recognizing patterns that may warrant attention.

In the recent work, members of our team “AI for Security” published a paper at one of the key computer science conferences: the Conference on Uncertainty in Artificial Intelligence (UAI 2023, Pittsburgh, USA)³¹ that analyzed a situation of an attack that has two phases. Typically, the attack

³¹ A. Nagórko, P. Ciosmak, T. Michalak, *Two-phase security games*, in: *Proceedings of the Thirty-Nine Conference on Uncertainty in Artificial Intelligence (UAI 2023)*, pp. 1489–1498.

in security games is modeled as a one-off assault during which the attacker has no chance to update their strategy even if new valuable information is gained in the process. This, however, does not cover certain tactics that can be applied by ever more agile covert organizations.

To address this, we propose a model in which, in the first phase, the attacker makes a preliminary move designed to gain extra information on the defender's activities in this particular instance of the game. Next, in the second phase, this insight is used to choose an optimal concluding move.

A recent real-world example of the tactics that are explicitly modeled in our two-phase game are the actions of Lukashenko's regime in Belarus which exploits immigrants to probe the border with Ukraine³². This callous behaviour puts the lives of the immigrants in extreme danger both due to very difficult terrain and the ongoing war. In more details, Ukraine's northwestern border of nearly 900 km is a heavily forested area full of forbidding wetlands and the Chernobyl Exclusion Zone. On top of that, the border – that was crossed by the Russian army in February 2022 and then subsequently restored by the Ukrainian counteroffensive – is now heavily fortified with trenches, walls and mine fields.

Regrettably, despite the fact that the border has now become one of the most perilous in the world, the Belarusian border guards are actively organizing and coordinating groups of immigrants in an attempt to breach it. Their objective is to expose and disrupt the Ukrainian defenses, which are obligated to respond to such attempts due to the threat posed by Russian saboteurs.

Given that some sophisticated electronic security measures are in place, most of these border crossings are detected. However, it should be noted that detection does not necessarily guarantee the presence of a patrol close enough to pre-vent unauthorized entry, meaning that the border is not entirely impenetrable. Nevertheless, even in cases where a specific section of the border is unguarded at the moment of entry, the Ukrainian headquarters promptly dispatch a team to the area.

Consequently, subsequent attempts to enter the same section of the border are highly unlikely to succeed, given the swift response and reinforcement measures taken by the Ukrainian authorities.

³² V. Romanenko, *Belarus uses migrants for intelligence on the border with Ukraine*, <https://www.pravda.com.ua/eng/news/2022/12/6/7379514/> [accessed: 25 VI 2023].

Let us consider a scaled-down version of the problem, with four sections of the Belarus-Ukraine border ($S_1, S_2, S_3,$ and S_4) and two patrol units. This setting can be modelled as a standard security game in the spirit of the one used at the Los Angeles World Airport³³. Pure strategies (moves) of the Ukrainian defenders are possible assignments of patrols to the sections of the border:

$$I = \{S_1S_2, S_1S_3, S_1S_4, S_2S_3, S_2S_4, S_3S_4\}$$

We assume there are two possible types of the attacker: low- and high-profile human traffickers (type 1 and 2, respectively). The high-profile type of the attacker inflicts a much larger loss upon the defender as they organize much bigger groups. Both types have the same strategy space, i.e., an attacker of each type can either choose one of the four sections of the border or back off, i.e., $J_1 = J_2 = \{S_1, S_2, S_3, S_4, \emptyset\}$. The payoffs of both parties, depending on the attacker type, increase linearly with S_i : for a high-profile attackers payoffs are 50, 100, 150 and 200 respectively and for a low-profile attacker the payoffs are five times smaller. The defender payoffs are opposite, with small random noise added uniformly from interval $[-5, 5]$.

Assuming that probabilities of attacks by these two types are $p_1 = 0,8$ for the low-profile attacker and $p_2 = 0,2$ for the high-profile one, an optimal strategy for the defender is:

$$(x_{S_1S_2}, x_{S_1S_3}, x_{S_1S_4}, x_{S_2S_4}, x_{S_2S_4}, x_{S_3S_4}) = (0\%, 50\%, 0\%, 0\%, 50\%, 0\%)$$

According to this strategy, border sections S_1 and S_2 are never protected simultaneously. Such a situation is typical for Stackelberg equilibria in one-phase games and can be easily exploited by performing a two-phase attack.

Now, let us discuss the concept of a two-phase attack. Assume that, unbeknownst to the defender, the attacker possesses the necessary resources and capabilities of both a low-profile human trafficker and a high-profile one. Consequently, the attacker can attempt to breach two sections of the border sequentially, in distinct phases.

Based on the optimal strategy derived previously, let's consider the scenario where, in the first phase, a low-profile human trafficker makes an attempt to breach the border at section S_1 . This initial

³³ J. Pita et al., *Using game theory for Los Angeles Airport...*

phase provides the attacker with valuable information, regardless of the defender’s positioning. This is due to the fact that the attacker now possesses knowledge of a conditional probability distribution pertaining to the defender’s resources.

Let $t \in \{0\%, 17\%, 33\%, 50\%, 67\%, 83\%, 100\%\}$ be a chance of encountering a two-phase attacker, $(1 - t) \times 80\%$ be a probability of encountering a low-profile attacker and $(1 - t) \times 20\%$ be a likelihood of encountering a high-profile attacker. For $t = 0\%$ this is the standard one-phase model, while $t = 100\%$ describes a pure two-phase attack.

Table 4 shows that presence of two-phase attackers significantly alters the Stackelberg equilibrium of the game. For example, for 33% probability of a two-phase attack (with 53% chance of a single-phase low-profile attack and 13% chance of a single-phase high-profile attack, keeping the 4 : 1 low-to high-profile ratio), the optimal defender strategy becomes

$$(x_{S_1S_2}, x_{S_1S_3}, x_{S_1S_4}, x_{S_2S_3}, x_{S_2S_4}, x_{S_3S_4}) = (12\%, 15\%, 17\%, 17\%, 18\%, 21\%)$$

As we see in Table 4, two-phase Stackelberg equilibria are much more robust against changes of attacker profiles.

Table 4. Each row presents an optimal mixed strategy of the defender against a group of attackers with a given chance of encountering a two-phase attack. As we can see in the last row, without presence of two-phase attackers the Stackelberg equilibrium heavily over-fits to the random noise in payoff matrices.

0.085	0.11	0.12	0.2	0.25	0.23	100%	Chance of a two-phase attack
0.085	0.11	0.12	0.2	0.25	0.23	83%	
0.12	0.11	0.12	0.2	0.25	0.23	67%	
0.12	0.15	0.17	0.17	0.18	0.21	50%	
0.12	0.15	0.17	0.17	0.18	0.21	33%	
0.15	0.15	0.17	0.16	0.18	0.18	17%	
0	0.5	0	0	0.5	0	0%	
S_1S_2	S_1S_3	S_1S_4	S_2S_3	S_2S_4	S_3S_4		Moves of the defender (patrol placements)

Table 5 shows how defender payoffs change against different compositions of attacker groups. For example, the expected payoff of the defender against a single-phase attack drops to -175 when single-phase strategy is pitted against a two-phase attacker.

Table 5. Expected defender payoff when playing a strategy from Table 4 against a given chance of a two-phase attack. As we can see in the last column, the loss incurred by playing a strategy that ignores the possibility of a two-phase attack is an order of magnitude larger than over-cautious protection against such attacks.

-16.2	-16.2	-16.2	-20.3	-20.3	-24.9	-175	100%
-14.8	-14.8	-14.8	-17.3	-17.3	-20.9	-146	83%
-13.4	-13.4	-13.4	-14.3	-14.3	-16.9	-116	67%
-12	-12	-12	-11.3	-11.3	-12.8	-87.1	50%
-10.7	-10.7	-10.7	-8.36	-8.36	-8.84	-57.9	33%
-9.27	-9.27	-9.27	-5.38	-5.38	-4.83	-28.6	17%
-7.89	-7.89	-7.89	-2.41	-2.41	-0.816	0.7	0%
100%	83%	67%	50%	33%	17%	0%	

Defender strategy

Chance of a two-phase attack

In order to fix this flaw, we propose a new model which allows for considering one-phase and two-phase attackers simultaneously. With our security model, the expected payoff against coordinated attackers jumps from -175 to -16.2 (the defender is still at a disadvantage). The optimal strategy:

$$(x_{S_1S_2}, x_{S_1S_3}, x_{S_1S_4}, x_{S_2S_4}, x_{S_2S_4}, x_{S_3S_4}) = (8,5\%, 11\%, 12\%, 20\%, 25\%, 23\%)$$

forces the low-profile attacker to attack S_1 and the high-profile attacker to back off if S_1 was not patrolled. Note that this comes at a cost: for the uncoordinated (one-phase) attack, when low- and high-profile attackers act independently, this strategy brings payoff -7.89 to the defender (a drop from 0.7).

Summary

In this paper, we have provided an exposition of advanced methodologies to ensure the safety of critical infrastructure that incorporate the combination of game theory, optimization techniques, and AI algorithms. The effectiveness of these methods has been demonstrated through their successful deployment in practice in multiple location in the USA. It is crucial to underline that these improvements were achieved not by adding on security resources but rather by optimally deploying the available ones. The work of our team “AI for Security” at the IDEAS NCBR research institute is focused on extending these results and make them applicable to various types of critical infrastructure and to the security threats that have recently reappeared in Europe. We aspire to have them implemented soon to optimize in practice the protection of the Polish critical infrastructure sites and systems.

A Formal Description

We start with a formal description of security games that follows modern treatment by Xu³⁴. Then we describe a broader class of Bayesian Stackelberg games that forms a basis for the two-phase model discussed in the previous part of article and we derive a quadratic optimization problem that can be used to solve these games.

A.1 Security Games

A security game, once again, is a two-player game between a defender and an attacker. The defender possesses multiple security resources and aims to allocate these resources to protect n targets from the set $[n] = \{1, 2, \dots, n\}$. A defender pure strategy is a subset of targets that is protected (covered) in a feasible allocation of these resources. A representation of a pure strategy is a binary vector $e \in \{0, 1\}^n$ where the entries of value 1 specify the covered targets. Let $E \subseteq \{0, 1\}^n$ denote the set of all defender pure strategies. A defender mixed strategy is a probability distribution x over the elements in E . An attacker pure strategy is a target $i \in [n]$. An attacker mixed strategy is denoted by $y \in \Delta_n$, where Δ_n is an n -dimensional simplex. We will use y_i to denote the probability of attacking target i .

³⁴ H. Xu, *The Mysteries of Security Games: Equilibrium Computation Be-Comes Combinatorial Algorithm Design*, in: *Proceedings of the 2016 ACM Conference on Economics and Computation (ACM EC 2016)*, pp. 497–514.

In a most general phrasing, security games are a form of a bilinear game. A bilinear game is given by a pair of matrices (A,B) and polytopes (P,Q) . Given that player 1 plays $x \in P$ and player 2 plays $y \in Q$, the utilities for player 1 and 2 are $x^T Ay$ and $x^T By$ respectively.

We can now give different notions of equilibria for security games. A strategy profile (x, y) is a Nash equilibrium, if:

$$\forall x' \in P \quad \forall y' \in Q \quad x^T Ay \geq x'^T Ay \quad \& \quad x^T By \geq x'^T By'$$

By The Nash Theorem, there exists at least one NE, possibly multiple NEs, in any bilinear game.

When one player moves before another player, the Stackelberg equilibrium serves as a more appropriate solution concept. A two-player Stackelberg game is played between a leader and a follower. The leader moves first, or equivalently, commits to a mixed strategy. The follower observes the leader's strategy and best responds. The leader's optimal strategy, together with the follower's best response, forms an equilibrium.

Let

$$y_x = \arg \max_{y' \in Q} x^T By'$$

denote the follower's best response to a leader strategy $x \in P$. A strategy profile (x, y) is a strong Stackelberg equilibrium if:

$$x = \arg \max_{x' \in P} x'^T Ay_{x'} \quad \text{and} \quad y = y_x$$

When $B = -A$, the bilinear game is zero-sum. In such games, both NE and SSE, are equivalent to the minimax equilibrium.

A strategy profile (x, y) is a **minimax equilibrium** if

$$\forall x' \in P \quad \forall y' \in Q \quad x^T Ay \geq x'^T Ay \quad \& \quad x^T Ay \leq x^T Ay'$$

If (x, y) is a minimax equilibrium, the strategy x is the player 1's maximin strategy, and y is the player 2's minimax strategy.

The value of the game is:

$$V = x^T Ay = \max_{x' \in P} \min_{y' \in Q} x'^T Ay'$$

We will now describe the payoff structure of the game – given that the attacker attacks target i :

- the defender gets a reward r_i if target i is covered or a cost c_i if i is uncovered,
- the attacker gets a cost ξ_i if target i is covered or a reward ρ_i if i is uncovered,
- both players have utility 0 on the other $n - 1$ unattacked targets.

A crucial assumption here is the following: for all $i \in [n]$ we have:

$$r_i > c_i \text{ and } \rho_i > \xi_i.$$

This means that:

- covering a target is strictly beneficial to the defender than uncovering it,
- the attacker prefers to attack a target when it is uncovered.

Definition 1 (Security Game). A security game G with n targets is a tuple (r, c, ρ, ξ, E) that satisfies $r_i > c_i$ and $\rho_i > \xi_i$ for all $i \in [n]$.

The defender's utility can be defined as follows:

$$U^d(e, i) = r_i e_i + c_i(1 - e_i)$$

Given $p \in \Delta_{|E|}$ and $y \in \Delta_n$, the defender's expected utility is:

$$\begin{aligned} U^d(p, y) &= \sum_{e \in E} \sum_{i \in [n]} p_e y_i U^d(e, i) = \\ &= \sum_{e \in E} \sum_{i \in [n]} p_e y_i (r_i e_i + c_i(1 - e_i)) = \\ &= \sum_{i \in [n]} y_i \sum_{e \in E} p_e (r_i e_i + c_i(1 - e_i)) = \\ &= \sum_{i \in [n]} y_i \left(r_i \sum_{e \in E} p_e e_i + c_i \left(1 - \sum_{e \in E} p_e e_i \right) \right) \end{aligned}$$

Given $p \in \Delta_{|E|}$ and $y \in \Delta_n$ the defender's expected utility is:

$$U^d(p, y) = \sum_{i \in [n]} y_i \left(r_i \sum_{e \in E} p_e e_i + c_i \left(1 - \sum_{e \in E} p_e e_i \right) \right)$$

If we follow the convention of using:

$$x_i := \sum_{e \in \mathcal{E}} p_e e_i$$

where x_i is the marginal coverage probability of target i , then we have:

$$U^d(p, y) = \sum_{i \in [n]} y_i (r_i x_i + c_i (1 - x_i))$$

Let $x = (x_1, \dots, x_n)^T$ denote the marginal probability for all targets induced by the mixed strategy p . Then the equation above shows that the defender's expected utility can be compactly expressed as the bilinear form:

$$U^d(x, y) = \sum_{i \in [n]} y_i (r_i x_i + c_i (1 - x_i))$$

Note that $U^d(x, y)$ has the bilinear form

$$x^T A y + a x$$

for some non-negative diagonal matrix A .

A note is in order here: the convex hull of E is a polytope of all the feasible (i.e., implementable by a defender mixed strategy) marginal probabilities:

$$\mathcal{P} = \{x = \sum_{e \in \mathcal{E}} p_e e : p \in \Delta_{|\mathcal{E}|}\}$$

so we can simply interpret a point $x \in \mathcal{P}$ as a mixed strategy and denote the defender's utility by: $U^d(x, y)$.

Similarly, the attacker's expected utility can be compactly represented in the following form:

$$U^a(x, y) = \sum_{i \in [n]} y_i (\rho_i (1 - x_i) + \xi_i x_i)$$

Note that $U^a(x, y)$ also has the bilinear form

$$x^T B y + \beta y$$

for some non-positive diagonal matrix B .

In zero-sum games, all standard equilibrium concepts are payoff-equivalent to the minimax equilibrium, and our goal is to compute the minimax equilibrium in polynomial time.

When the game is not zero-sum, the main solution concept is the strong Stackelberg equilibrium (SSE): the defender plays the role of the leader and can commit to a mixed strategy before the attacker moves. The attacker observes the defender's mixed strategy and best responds. In this case, the goal is to compute the optimal mixed strategy for the defender to commit to (note that the attacker is not able to observe the defender's real-time deployment, i.e., the sampled pure strategy, since he has to plan the attack before the defender's real-time pure strategy is sampled).

A.2 Bayesian Security Games

The solution deployed at the LAX Airport was based on a broader class of games, called Bayesian security games or Bayesian Stackelberg games. We follow³⁵ to describe this class and we use the Belarus-Ukraine border protection problem discussed in Section 4 as a running example.

In a Bayesian Stackelberg game, the defender plays against a group of attackers of n distinct types. In each round, the defender plays against a single attacker and encounters the attacker of type $1 \leq t \leq n$ randomly, with probability p_t . Attackers may have different sets of moves at their disposal that inflict different damage to the defender. In the running example, we have a low-profile attacker ($t = 1$) and a high-profile attacker ($t = 2$), with

$$p_1 = \frac{4}{5} \text{ and } p_2 = \frac{1}{5}.$$

We let I denote the set of defender's moves. In the running example, the border patrol assigns two patrolling units to four segments of the border, hence $I = \{S_1S_2, S_1S_3, S_1S_4, S_2S_3, S_2S_4, S_3S_4\}$.

In a Bayesian Stackelberg game, the defender picks his mixed strategy x first. Here $x = \{x_i\}_{i \in I}$ is a probability measure on I , which we denote by $x \in \text{Prob}(I)$ with

$$\text{Prob}(I) = \{x: I \rightarrow R: \sum_{i \in I} x_i = 1, x \geq 0\}$$

Strategy x does not depend on t as the defender doesn't know the type of attacker he will encounter.

³⁵ A. Nagórko, P. Ciosmak, T. Michalak, *Two-phase security games...*

Let J_t denote the set of moves of attacker of type t . In the running example, $J_1 = J_2 = \{S_1, S_2, S_3, S_4, \emptyset\}$, i.e. attackers may either attack one of the border segments or back off. Attacker picks his strategy second, with the knowledge of the defender's strategy x .

Although it may seem counter-intuitive at first, it is advantageous to the defender to disclose his mixed strategy to the attacker (but not his current defensive positions). It is quite common to disclose information in such scenarios to force the adversary to a favorable response, see e.g. action "ZNICZ" carried out each year by the Polish Police³⁶.

In each round of the game, both players move independently, according to strategies x and $y^t(x)$ they picked prior. Let $r_{i,t,j}$ denote the defender's payoff if she played move $i \in I$ against the attacker of type $1 \leq t \leq n$ who played a move $j \in J_t$. Let $c_{i,t,j}$ denote attacker's payoff (which may be different from $-r_{i,t,j}$) as we do not assume that the games are zero-sum in general.

Player payoffs may be compactly presented using payoff matrices. In the running example, the payoff matrices for the high-profile attack are:

		S_1		S_2		S_3		S_4		\emptyset
S_1, S_2	51,	-50	102,	-100	-152,	150	-211,	200	0,	0
S_1, S_3	55,	-50	-123,	100	175,	-150	-221,	200	0,	0
S_1, S_4	59,	-50	-108,	100	-169,	150	206,	-200	0,	0
S_2, S_3	-69,	50	101,	-100	168,	-150	-221,	200	0,	0
S_2, S_4	-55,	50	113,	-100	-170,	150	212,	-200	0,	0
S_3, S_4	-75,	50	-123,	100	166,	-150	211,	-200	0,	0

The first number in row i , column j is the defender payoff $r_{i,t,j}$ (here 1 stands for high-profile attacker $t=1$). The second number is $c_{i,1,j}$. The payoffs for the low-profile attack are:

		S_1		S_2		S_3		S_4		\emptyset
S_1, S_2	14,	-10	23,	-20	-34,	30	-42,	40	0,	0
S_1, S_3	10,	-10	-20,	20	32,	-30	-43,	40	0,	0
S_1, S_4	12,	-10	-23,	20	-33,	30	44,	-40	0,	0
S_2, S_3	-11,	10	24,	-20	31,	-30	-41,	40	0,	0
S_2, S_4	-11,	10	20,	-20	-31,	30	42,	-40	0,	0
S_3, S_4	-11,	10	-21,	20	34,	-30	44,	-40	0,	0

³⁶ *Policyjne działania Znicz*, <https://policja.pl/pol/aktualnosci/210088,Policyjne-dzialania-ZNICZ.html> [accessed: 25 VI 2023].

Attacker t picks an optimal strategy $\bar{y}^t = \bar{y}^t(x)$ that depends on strategy x known by him and that maximizes his expected payoff

$$\bar{c} = \sum_{i \in I} \sum_{j \in J_t} x_i \bar{y}_j^t c_{i,t,j}$$

This payoff is maximized by a pure strategy, i.e., y^t is optimal if and only if

$$\bar{c} \geq \sum_{i \in I} x_i c_{i,t,j}$$

The defender acts to maximize his expected payoff against the optimal strategies of the attackers, i.e. he picks an optimal strategy x that maximizes his expected payoff:

$$\sum_{i \in I} \sum_{t=1}^n \sum_{j \in J_t} p_t x_i \bar{y}_j^t r_{i,t,j}$$

Hence the following quadratic optimization problem solves Bayesian Stackelberg games:

$$\max_{x, y^t} \sum_{i \in I} \sum_{t=1}^n \sum_{j \in J_t} p_t x_i y_j^t r_{i,t,j}$$

subject to:

$$\begin{aligned} \sum_{i \in I} x_i &= 1, \\ \sum_{j \in J_t} y_j^t &= 1 \text{ for each } 1 \leq t \leq n, \\ \sum_{i \in I} \sum_{j \in J_t} x_i y_j^t c_{i,t,j} &\geq \sum_{i \in I} x_i c_{i,t,j} \text{ for each } 1 \leq t \leq n, j \in J_t, \\ x &\geq 0, y^t \geq 0 \text{ for each } 1 \leq t \leq n. \end{aligned}$$

This formulation coupled with a linearization technique leads to a mixed integer linear programming formulation of Bayesian Stackelberg games published by Paruchuri et al.³⁷, as the celebrated DOBSS algorithm.

³⁷ P. Paruchuri et al., *Playing Games for Security: An Efficient Exact Algorithm for Solving Bayesian Stackelberg Games*, in: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, vol. 2, pp. 895–902.

Bibliography

An B. et al., *A Deployed Quantal Response-Based Patrol Planning System for the U.S. Coast Guard*, “Interfaces” 2013, vol. 43, no. 5, pp. 400–420. <https://doi.org/10.1287/inte.2013.0700>.

Bier V.M., Azaiez M.N., *Game Theoretic Risk Analysis of Security Threats*, Springer 2008. <https://doi.org/10.1007/978-0-387-87767-9>.

Blocki J. et al., *Audit games with multiple defender resources*, in: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015)*, vol. 29, no. 1, pp. 791–797.

Brown G. et al., *A Two-Sided Optimization for Theater Ballistic Missile Defense*, “Operations Research” 2005, vol. 53, no. 5, pp. 745–763. <https://doi.org/10.1287/opre.1050.0231>.

Fang F., Nguyen T.H., *Green security games: Apply game theory to addressing green security challenges*, “ACM SIGecom Exchanges” 2016, vol. 15, no. 1, pp. 78–83. <https://doi.org/10.1145/2994501.2994507>.

Gatti N. et al., *Game Theoretical Insights in Strategic Patrolling: Model and Algorithm in Normal-Form*, in: *Proceedings of the 2008 conference on ECAI 2008: 18th European Conference on Artificial Intelligence (ECAI 2008)*, pp. 403–407. <https://doi.org/10.3233/978-1-58603-891-5-403>.

Gholami S. et al., *Adversary models account for imperfect crime data: Forecasting and planning against real-world poachers*, in: *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, pp. 823–831.

Guo Q. et al., *Coalitional security games*, in: *Proceedings of the 2016 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, pp. 159–167.

Haskell W. et al., *Robust protection of fisheries with COMPASS*, in: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2014)*, vol. 28, no. 2, pp. 2978–2983.

Hunt K., Zhuang J., *A review of attacker-defender games: Current state and paths forward*, “European Journal of Operational Research” 2023, in press. <https://doi.org/10.1016/j.ejor.2023.04.009>.

Hutter F. et al., *Boosting Verification by Automatic Tuning of Decision Procedures*, in: *Proceedings of the 19th International Conference on Computer Aided Verification (CAV 2007)*, pp. 27–34.

Korzhyk D. et al., *Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness*, “Journal of Artificial Intelligence Research” 2011, vol. 41, no. 2, pp. 297–327.

Lye K-w., Wing J., *Game Strategies in Network Security*, “International Journal of Information Security” 2005, vol. 4, pp. 71–86. <https://doi.org/10.1007/s10207-004-0060-x>.

Nagórko A., Ciosmak P., Michalak T., *Two-phase security games*, in: *Proceedings of the Thirty-Nine Conference on Uncertainty in Artificial Intelligence (UAI 2023)*, pp. 1489–1498.

Nguyen T.H. et al., *Analyzing the effectiveness of adversary modeling in security games*, in: *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI 2013)*, no. 1, pp. 718–724.

Nguyen T.H. et al., *Towards a science of security games*, in: *Mathematical Sciences with Multidisciplinary Applications*, B. Toni (ed.), Springer Cham 2016, pp. 347–381.

Paruchuri P. et al., *Playing Games for Security: An Efficient Exact Algorithm for Solving Bayesian Stackelberg Games*, in: *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, vol. 2, pp. 895–902.

Pita J. et al., *Using game theory for Los Angeles Airport security*, “AI Magazine” 2009, vol. 30, no. 1, pp. 43–57. <https://doi.org/10.1609/aimag.v30i1.2173>.

Sandler T., *Terrorism & Game Theory*, “Simulation & Gaming” 2003, vol. 34, no. 3, pp. 319337. <https://doi.org/10.1177/1046878103255492>.

Shieh E. et al., *Protect: A deployed game theoretic system to protect the ports of the United States*, in: *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2012)*, vol. 1, pp. 13–20.

Sinha A. et al., *Stackelberg security games: Looking beyond a decade of success*, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pp. 5494–5501.

Stackelberg H. von, *Marktform und Gleichgewicht*, J. Springer 1934.

Tambe M., *Security and game theory: algorithms, deployed systems, lessons learned*, Cambridge 2011.

Tsai J. et al., *Iris - a tool for strategic security allocation in transportation networks*, in: *Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2009, Industry Track)*, pp. 37–44

Vorobeychik Y., An B., Tambe M., *Adversarial Patrolling Games*, in: *Papers from the 2012 AAAI Spring Symposium*, vol. 3, pp. 91–98.

Vorobeychik Y., Pritchard M., *Plan interdiction games*, in: *Adaptive Autonomous Secure Cyber Systems*, S. Jajodia et al. (ed.), Springer Cham 2020, pp. 159–182. https://doi.org/10.1007/978-3-030-33432-1_8.

Xu H., *The Mysteries of Security Games: Equilibrium Computation Be-Comes Combinatorial Algorithm Design*, in: *Proceedings of the 2016 ACM Conference on Economics and Computation (ACM EC 2016)*, pp. 497–514.

Yang R. et al., *Adaptive resource allocation for wildlife protection against illegal poachers*, in: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pp. 453–460.

Yang R. et al., *Improving Resource Allocation Strategy Against Human Adversaries in Security Games*, in: *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pp. 458–464.

Yang R. et al., *Improving resource allocation strategies against human adversaries in security games: An extended study*, “Artificial Intelligence” 2013, vol. 195, pp. 440–469. <https://doi.org/10.1016/j.artint.2012.11.004>.

Yin Z. et al., *Trusts: Scheduling randomized patrols for fare inspection in transit systems*, in: *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, vol. 26, no. 2, pp. 2348–2355.

Zhang Y., Malacaria P., *Bayesian Stackelberg games for cyber-security decision support*, “Decision Support Systems” 2021, vol. 148, art. 113599. <https://doi.org/10.1016/j.dss.2021.113599>.

Internet sources

Policyjne dzialania Znicz, <https://policja.pl/pol/aktualnosci/210088,Policyjne-dzialania-ZNICZ.html> [accessed: 25 VI 2023].

Romanenko V., *Belarus uses migrants for intelligence on the border with Ukraine*, <https://www.pravda.com.ua/eng/news/2022/12/6/7379514/> [accessed: 25 VI 2023].

Tomasz P. Michalak, PhD

Leader of an independent research team at IDEAS NCBR and a lecturer at the Faculty of Mathematics, Informatics, and Mechanics at the University of Warsaw. Graduate of the Faculty of Economic Sciences at the University of Warsaw. During his academic career, he conducted research at the Department of Computer Science at the University of Oxford, the School of Electronics and Computer Science at the University of Southampton, the Department of Computer Science at the University of Liverpool, and the Faculty of Applied Economics at the University of Antwerp, where he obtained his Ph.D. in economics.

Michał T. Godziszewski, PhD

Specialist in logic and its applications (in mathematics, philosophy and computer science), artificial intelligence (specialisation - multi-agent systems theory: algorithmic game theory, network analysis, computational social choice theory) and theoretical computer science. His current research focuses on algorithmic analysis of social networks and Stackelberg games, computational complexity in game theory and their applications to security systems modelling.

Andrzej Nagórko, PhD

Adjunct at the Faculty of Mathematics, Informatics and Mechanics at the University of Warsaw. Former employee of the Mathematical Institute of the Polish Academy of Sciences and at universities in United States and Israel. For many years he applied mathematical optimization techniques in diverse fields of mathematics - from artificial intelligence, through game theory to geometric group theory. In IDEAS NCBR he works on applications of these methods to the protection of critical infrastructure.