Hall of Fame in Coevolutionary Algorithm for Stackelberg Security Games

Adam Żychowski[1][0000−0003−0026−5183],
Jacek Mańdziuk[1][0000−0003−0947−028X],
Andrew Perrault[2][0000−0002−5062−7958]

[1]*Warsaw University of Technology, Faculty of Mathematics and Information Science*
adam.zychowski@pw.edu.pl jacek.mandziuk@pw.edu.pl
[2]*The Ohio State University, Department of Computer Science and Engineering*

Abstract. *Stackelberg Security Games (SSGs) is a popular game-theoretic model for strategic interactions between a Defender and an Attacker. The computational challenges of identifying optimal strategies for larger games led to the development of the CoEvoSG coevolutionary method. This paper introduces an extension to CoEvoSG, incorporating a Hall of Fame (HoF) mechanism with Mixed Nash Equilibrium. The HoF stores successful strategic configurations, enhancing algorithm robustness. Results across distinct game types demonstrate that the proposed method consistently outperforms the baseline CoEvoSG algorithm with standard HoF approach. This improvement is achieved with minimal computation time increase.*
Keywords: *Stackelberg Security Games, Evolutionary Computation*

1. Introduction

Stackelberg Security Games (SSGs) represent a widely applicable game model effectively deployed across various real-world security domains, including security patrolling [1, 2], poaching prevention [3], terrorist deterrence [4], and cybersecurity [5]. The challenge of identifying the optimal pair of players' strategies (Stackelberg Equilibrium) for SSGs is proven to be NP-hard, prompting the development of numerous approximation methods in the literature. This study builds upon the

CoEvoSG coevolutionary method, which has demonstrated superior performance in terms of computation time and the ability to address larger games that were previously intractable. In this paper we extend the CoEvoSG method by the Hall of Fame (HoF), containing strategies of the Nash Equilibrium for the current state of players' populations. This addition enhances the proposed method, yielding improved results with only a marginal increase in computation time.

1.1. Problem definition

SSGs involve two players: the Defender (D) and the Attacker (A). Each game consists of m time steps during which both players simultaneously choose actions. A player's *pure strategy* σ_P ($P \in \{D, A\}$) is a sequence of their actions in consecutive time steps: $\sigma_P = (a_1, a_2, \ldots, a_m)$. The set of all possible pure strategies of player P is denoted by Σ_P. Furthermore, the player's *mixed strategy* is a probability distribution $\pi_P \in \Pi_p$ over Σ_P, where Π_p is the set of all mixed strategies for player P. For any pair of strategies (π_D, π_A) the expected payoffs for the players are denoted by $U_D(\pi_D, \pi_A)$ and $U_A(\pi_D, \pi_A)$, resp. The goal of the game is to find the *Strong Stackelberg Equilibrium* (SSE), i.e. a pair of strategies (π_D, π_A) satisfying the following conditions:

$$\pi_D = \arg\max_{\bar{\pi}_D \in \Pi_D} U_D(\bar{\pi}_D, BR(\bar{\pi}_D)), \qquad BR(\pi_D) = \arg\max_{\pi_A \in \Pi_A} U_A(\pi_D, \pi_A).$$

The first one optimizes the selection of the Defender's strategy π_D under the assumption that the Attacker always opts for the best response strategy ($BR(\pi_D)$) to the Defender's committed strategy. Both players determine their strategies at the beginning of a game (first the Defender and then the Attacker). Once the strategies are chosen, they remain fixed throughout the duration of the game.

2. Proposed method - Mixed Nash Equilibrium Hall of Fame

The CoEvoSG algorithm [6] is a coevolutionary approach designed to address the computational challenges associated with evaluating strategies in SSGs. In standard evaluations, the Attacker's pure strategies are exhaustively explored to determine the best response to the Defender's strategy. This process becomes impractical for larger games or continuous strategy spaces. CoEvoSG maintains two populations, one for the Defender's mixed strategies and the other for the Attacker's pure strategies. The algorithm operates by alternating modifications of

the Attacker's and Defender's populations, and evolving them through a specified number of generations. The Attacker's population represents pure strategies which are developed with evolutionary operators such as crossover and mutation. **Crossover** combines strategies by swapping actions, while **mutation** introduces new actions, both contributing to the exploration of the strategy space.

Evaluation process involves the Defender's population being assessed against a subset of the Attacker's strategies. The best Attacker's response is determined, and the Defender's payoff against this response is used as the fitness value for the evaluated Defender's strategy. The evaluation of the Attacker's population is more complex, considering the adaptability of the best response to different Defender strategies. Namely, in order to ensure the effectiveness of the Attacker's strategies against multiple Defender's strategies (not just the best one) and at the same time to avoid oscillations, $N_{top} = 20$ highest-fitness Defender's strategies are used for each Attacker's strategy evaluation. **Selection** determines which individuals advance to the next generation. It includes elite preservation and binary tournaments. The algorithm concludes when a predefined number of generations is reached or when no improvement in the best-found solution (Defender's payoff) is observed over a specified number of consecutive generations.

A more detailed description of the CoEvoSG algorithm can be found in [6].

2.1. Mixed Nash Equilibrium (MNE) Hall of Fame

We expanded the CoEvoSG algorithm by incorporating the Hall of Fame (HoF) archive mechanism featuring Nash Equilibrium. HoF is a well-established concept in evolutionary computation which typically consists in the inclusion of the best individual in each generation to contribute to subsequent evaluations.

Please note, that HoF differs from the well-known elitism mechanism, as the latter consists in preserving the best individuals between generations (in the selection process), so as not to lose the overall best solution, while HoF extends this concept by storing the best solutions found throughout the entire evolutionary process for their further usage or for reference.

In contrast, our method deviates by integrating the MNE strategy instead of a single best individual. A similar approach was previously proposed in [7] for optimizing robust decision trees.

Two distinct HoFs are maintained, one for the Defender's population and another for the Attacker's population. Following each generation, the MNE strategy is computed based on the current players' populations. The resultant mixture com-

prises the Attacker's pure strategies and the Defender's mixed strategies, which are subsequently appended to the respective HoFs. Another modification to CoEvoSG is introduced in the evaluation procedure. Individuals (strategies) from a given population (Defender's or Attacker's) undergo evaluation not only against individuals from the adversarial population but also with each corresponding element within their respective HoF.

The HoF acts as a repository of historical knowledge containing strategic interactions between the Defender and the Attacker over multiple generations. By storing MNE in the HoF, the algorithm preserves strategic configurations that have demonstrated success in dealing with a variety of opponent strategies. It improves algorithm robustness since strategies in the next generations are evaluated also against effective strategic configurations from the past rather than relying on a single best individual from the current adversarial population.

3. Results

The proposed method has been tested on 2 distinct types of Security Games: Search Games (SEG) [8] and FlipIt Games (FIG) [9]. The same game instances were also used for CoEvoSG evaluation [6]. Please refer to [10] for a detailed description of the rules and characteristics of the games. To maintain consistency and ensure a fair comparison, the baseline CoEvoSG was configured with the same set of parameters as reported in [6].

	C2016	O2UCT	EASG	CoEvoSG	CoEvoSG+HoF	CoEvoSG+NEHoF
5	0.890	0.887	0.886	0.886	0.886	0.887
10	0.854	0.848	0.847	0.845	0.845	0.849
15	0.811	0.805	0.802	0.798	0.801	0.806
20	-	0.779	0.780	0.772	0.775	0.776
25	-	-	0.754	0.746	0.751	0.754
30	-	-	-	0.730	0.732	0.735
40	-	-	-	0.722	0.726	0.733

Table 1. Averaged Defender's payoff with respect to game nodes for FIG.

Tables 1 and 2 present the results obtained from five distinct methods. *C2016* [11] represents an exact method based on Mixed Integer Linear Programming, ensuring the return of optimal Defender's strategies. However, its applicability is constrained by extensive computation times, which makes it suitable only

	C2016	O2UCT	EASG	CoEvoSG	CoEvoSG+HoF	CoEvoSG+NEHoF
15	0.122	0.116	0.115	0.115	0.115	0.116
20	0.117	0.107	0.106	0.101	0.104	0.106
25	-	0.119	0.117	0.115	0.116	0.119
30	-	-	0.136	0.135	0.135	0.135
40	-	-	-	0.150	0.152	0.156
50	-	-	-	0.139	0.144	0.146

Table 2. Averaged Defender's payoff with respect to game nodes for SEG.

for small-scale game instances. O2UCT [12] is a heuristic method which utilizes the Upper Confidence Bounds applied to Trees (UCT) algorithm - a variant of the Monte Carlo Tree Search method (MCTS). *EASG* [10] is an evolutionary algorithm maintaining a single population of Defender's strategies. While utilizing the same evolutionary operators as the CoEvoSG Defender's population, EASG evaluates candidate Defender's strategies against all possible pure Attacker's strategies, ensuring precise payoff calculations. However, this precision results in a significantly longer computation time compared to CoEvoSG. *CoEvoSG+HoF* is a baseline CoEvoSG algorithm augmented with standard HoF - the best individual from the current population is added to the respective HoF after each generation. *CoEvoSG+NEHoF* represents the proposed method incorporating MNE strategies into the respective HoFs.

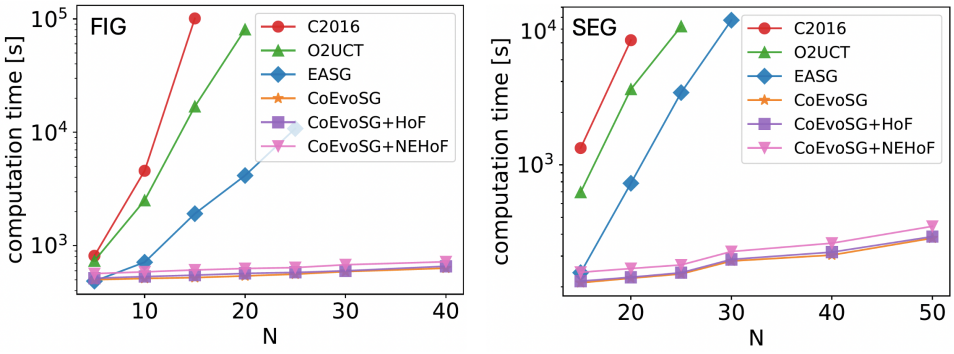

Figure 1. Computation times with respect to game nodes (N) for FIG and SEG.

Results across all datasets consistently indicate that the inclusion of an additional HoF element has a positive impact on the outcomes. In the majority of

games, CoEvoSG+NEHoF achieved higher Defender's payoffs than the baseline CoEvoSG algorithm. Furthermore, it outperforms CoEvoSG+HoF, supporting the claim that utilizing MNE results for HoF, as opposed to a simple best individual, is advantageous. Comparisons with EASG, which is a reference point for the optimal Defender's strategy evaluation, reveal that CoEvoSG+NEHoF achieves competitive results while exhibiting significantly lower computation times.

Figure 1 depicts the averaged computation times for each method. Notably, for all three coevolutionary algorithms, these times are significantly lower than those for other methods. The introduced overhead caused by HoF maintenance and MNE computation is insignificant. On average, CoEvoSG+HoF and CoEvo+NEHoF computation times are respectively 4% and 13% higher than plain CoEvoSG method that does not use the HoF mechanism.

4. Summary

In this study, we introduced an enhanced coevolutionary algorithm designed to address SSGs. It incorporates a HoF mechanism based on MNE derived from players' population strategies. The proposed method outperforms the baseline algorithm (absent of HoF), with only slight increase in computation time.

Acknowledgment

Adam Żychowski was funded by the Warsaw University of Technology within the Excellence Initiative: Research University (IDUB) programme.

References

[1] Sinha, A., Fang, F., An, B., Kiekintveld, C., and Tambe, M. Stackelberg Security Games: Looking Beyond a Decade of Success. In *27th IJCAI Conference*, pages 5494–5501. 2018.

[2] Karwowski, J. and Mańdziuk, J. A Monte Carlo Tree Search approach to finding efficient patrolling schemes on graphs. *European Journal of Operational Research*, 277:255–268, 2019.

[3] Żychowski, A., Mańdziuk, J., Bondi, E., Venugopal, A., Tambe, M., and Ravindran, B. Evolutionary approach to Security Games with signaling. *31st IJCAI Conference*, pages 620–627, 2022.

[4] Karwowski, J., Mańdziuk, J., Żychowski, A., Grajek, F., and An, B. A memetic approach for sequential security games on a plane with moving targets. In *Proceedings of the 33rd AAAI conference*, volume 33, pages 970–977. 2019.

[5] Mańdziuk, J. and Żychowski, A. Duel-based neuroevolutionary method for stackelberg security games with boundedly rational attacker. *Applied Soft Computing*, 146:110673, 2023.

[6] Żychowski, A. and Mańdziuk, J. Coevolution of players strategies in Security Games. *Journal of Computational Science*, 68:101980, 2023.

[7] Żychowski, A., Perrault, A., and Mańdziuk, J. Coevolutionary Algorithm for Building Robust Decision Trees under Minimax Regret . In *38th AAAI Conference*. 2024.

[8] Bošanský, B. and Čermak, J. Sequence-Form Algorithm for Computing Stackelberg Equilibria in Extensive-Form Games. In *29th AAAI Conference*, pages 805–811. 2015.

[9] Van Dijk, M., Juels, A., Oprea, A., and Rivest, R. L. Flipit: The game of "stealthy takeover". *Journal of Cryptology*, 26(4):655–713, 2013.

[10] Żychowski, A. and Mańdziuk, J. Evolution of Strategies in Sequential Security Games. In *20th AAMAS Conference*, pages 1434–1442. 2021.

[11] Čermák, J., Bošanský, B., Durkota, K., Lisý, V., and Kiekintveld, C. Using correlated strategies for computing stackelberg equilibria in extensive-form games. In *30th AAAI Conference*, pages 439–445. 2016.

[12] Karwowski, J. and Mańdziuk, J. Double-oracle sampling method for Stackelberg Equilibrium approximation in general-sum extensive-form games. In *34th AAAI Conference*, volume 34, pages 2054–2061. 2020.