A Multi-Objective Competitive Co-Evolutionary Framework with Progressive Shrinking for Wargame Scenarios

Ritam Guha[1], Ryan McKendrick[2], Bradley Feest[1], and Kalyanmoy Deb[1]

[1] Computational Optimization and Innovation (COIN) Laboratory
Michigan State University, East Lansing, Michigan 48864, USA
{guharita,kdeb}@msu.edu
[2] Northrop Grumman Corporation
2980 Fairview Park Drive Falls Church, VA 22042, USA
{ryan.mckendrick,bradley.feest}@ngc.com

COIN Report 2024010

Abstract. Dealing with multiple conflicting objectives in a multi-agent system is challenging, as agents' interactions complicate decision-making, especially when managing multiple Pareto-optimal fronts. In competitive co-evolutionary frameworks, not only do the objectives of each agent conflict with one another, but the agents' goals are also at odds. One such application domain is wargame strategy optimization, where the strategies of one agent must adapt based on the moves of opposing agents. Despite advancements in modern warfare, strategy analysis and decision-making are still largely manual, leaving room for great application of computational methods to automate different parts of the system. To address this, we propose a co-evolutionary optimization algorithm that integrates strategy search with interactive decision-making, allowing co-evolving populations to collaboratively identify their respective Pareto-optimal strategies. Central to this approach is a progressive-shrinking method that aligns feasible moves with those previously taken, ensuring smoother transitions. Our framework introduces a novel decision-making strategy using opposition front hypervolume improvement, particularly suited for competitive co-evolutionary contexts, combined with Penalty-based Boundary Intersection selection, to optimize strategy selections. We also examine the influence of various decision-making approaches, shrinking techniques, and parameter settings on the final results. This co-evolutionary framework, combining multi-agent interaction, evolutionary multi-objective optimization, and progressive shrinking, is not only effective for wargame strategy optimization but is also adaptable to other multi-agent conflicting systems.

Keywords: Attacker-defender system · Competitive co-evolution · Decision-making · Multi-objective games · Multi-agent systems · Progressive Shrinking

1 Introduction

In co-evolutionary optimization scenarios, where multiple agents interact, each agent's strategy evolves in response to the other. Unlike standard optimization problems, where objectives can be considered in isolation, co-evolutionary optimization requires both agents to continuously adapt to each other's strategies. This challenge is amplified when the agents have conflicting objectives, transforming the problem into a competitive co-evolutionary scenario. Each agent, while optimizing its own strategy, must also account for the actions of its opponent, making the process highly dynamic and intertwined.

One such practical application of competitive co-evolution is wargame strategy optimization, a critical aspect of military preparation and execution. Wargame simulations [1–3] have gained significant traction in the gaming community and military decision-making processes (MDMPs) [4]. Wargames simulate interactions between offense and defense agents, with each agent optimizing its resources to counter the opponent's moves. While our discussion highlights this scenario, we do not support or glorify warfare; we use wargames as an example of a challenging multi-agent optimization problem. In military decision-making, much of the strategy analysis and decision-making is done manually, but this paper introduces an automated framework to assist decision-makers in optimizing and selecting strategies.

In a typical wargame, the offense agent attacks, while the defense agent protects its assets. The success of any strategy relies heavily on the opponent's response, making it essential for both agents to co-evolve their strategies instead of optimizing in isolation. To address this need, we propose an iterative competitive Multi-objective Co-Evolutionary (MoCoEv) optimization framework, coupled with a novel decision-making process, specifically tailored for competitive co-evolutionary scenarios. This process utilizes opposition front hypervolume improvement, combined with a penalty-based boundary intersection (PBI) [5] method, customized for wargame situations. In this optimization problem, every strategy taken by the offense or defense can be represented in terms of multiple variables. The goal of the optimization is to co-evolve this set of variables for both offense and defense.

In the current wargaming context, if an agent needs to shift from one strategy to another, certain "switch variables" must remain fixed across the strategies under consideration. To facilitate these smooth transitions among strategies and further reduce problem dimensions based on variable convergence, we introduce the concept of progressive shrinking. This concept uses multiple rounds of MoCoEv and decision-making to reduce the dimension of the problem in each round. As we progressively reduce the variable dimension of the problem, we have named this process progressive shrinking. Progressive shrinking serves two key purposes: it enables agents to seamlessly transition between optimized strategies, and it significantly reduces the problem's dimensionality, making time-sensitive optimization tasks like online optimization more feasible and efficient. However, it is important to note that we do not address online optimization in this paper. The method iteratively refines the search space and decision-making process until a

stopping criterion is reached, ultimately delivering more focused and adaptable strategic solutions.
In summary, the key contributions of this paper are:

- We propose a novel progressive-shrinking-based multi-objective competitive co-evolutionary optimization framework. This framework narrows the search by gradually reducing the problem's dimensionality, enabling smooth strategy shifts during execution while simplifying decision-making for planners.
- We introduce a new and generically applicable decision-making approach for multi-agent systems involving more than one Pareto-optimal front.
- The proposed system offers some flexibility in execution, allowing for a wide range of possibilities. For instance, we can choose to progressively shrink strategies on only one side, instruct the system to prioritize more defensive decisions or adjust the reference point (origin) used in the PBI process. We explore the impact of these tunable parameters to gain deeper insights into the behavior of the final optimal strategies.

The rest of the paper is structured as follows: Section 2 sheds light on some of the scarce existing literature related to competitive multi-objective co-evolutionary algorithms. A brief description of the wargame strategy optimization problem is provided in Section 3. The proposed process of intertwined multi-objective co-evolutionary optimization and decision-making is described in Section 4. The outcome of the application of the proposed process to the WSOP under consideration is shown in Section 5 and finally we conclude the paper in Section 6.

2 Related Studies

In recent years, co-evolutionary algorithms have gained tremendous interest in various applications of multi-agent systems [6–8]. The idea of co-evolution is inspired by similar problem-solving tasks that nature addresses in many different forms. The survival and success of different species in nature have been achieved through co-evolution among multiple species [9]. There are three types of co-evolutionary processes: cooperative [10–13], competitive [14–17] and coopetitive [18]. When two or more species evolve by cooperating to achieve their own goals, the process is called *cooperative* co-evolution. On the other hand, when multiple species evolve by competing against each other to achieve their individual goals, it is called *competitive* co-evolution. Another form of co-evolution involves species that collaborate in certain areas while competing in others. This paradigm is referred to as *coopetition* [18]. In the case of wargame strategy optimization, generally, two agents' goals are in conflict, thereby making the task a competitive co-evolution. Moreover, each agent in the wargame problem usually optimizes multiple conflicting objectives, making the overall wargame problem a lot more complicated, challenging, and interesting.

3 Wargame Strategy Optimization Problem

Wargame strategy optimization problems (WSOPs) are highly idiosyncratic, and there is no single computing model that can fully represent all types of wargames. In the Military Decision-Making Process (MDMP), wargaming is a crucial step

[4], as shown in the flowchart in Figure 1. Currently, many parts of steps 4, 5, and 6 are conducted manually by wargame specialists or analysts, creating bottlenecks. As a result, only a small fraction of the strategies or Courses of Actions (COAs) are evaluated during wargaming.

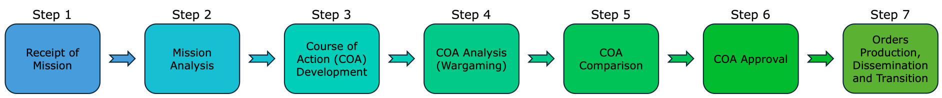

Fig. 1: Military Decision-making Process (MDMP) as outlined in [4].

This paper proposes an end-to-end automated WSOP framework that integrates steps 3, 4, and 5 using competitive co-evolutionary optimization and intertwined decision-making. We focus on a wargame involving two agents: offense and defense. The offense's goal is to maximize damage to an airbase (measured in air-to-ground missile hits) while minimizing its own costs (monetary expenses and losses of offensive lives) and imposing the highest possible costs on the defense (monetary expenses and losses of defensive lives). Conversely, the defense aims to counteract these goals, creating a direct competition between the two agents. Mathematically, the overall problem can be defined as follows:

Offense Goals:

$$\begin{aligned} &\text{Minimize}_{\mathbf{x}} \ (\text{-OH}(\mathbf{x},\mathbf{y}), \text{OE}(\mathbf{x},\mathbf{y}) - \text{DE}(\mathbf{x},\mathbf{y})), \\ &\text{subject to } x_i \in \{x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(c_i)}\},\ i = 1, \ldots, n_1. \end{aligned} \tag{1}$$

Defense Goals:

$$\begin{aligned} &\text{Minimize}_{\mathbf{y}} \ (\text{OH}(\mathbf{x},\mathbf{y}), \text{DE}(\mathbf{x},\mathbf{y}) - \text{OE}(\mathbf{x},\mathbf{y})), \\ &\text{subject to } y_j \in \{y_j^{(1)}, y_j^{(1)}, \ldots, y_j^{(c_j)}\},\ j = 1, \ldots, n_2. \end{aligned} \tag{2}$$

where OH, OE and DE represent `OffenseHits`, `OffenseExpenditures` and `DefenseExpenditures`, respectively; x_i and y_j denote the values of the i^{th} offense and j^{th} defense variable, respectively; c_i and c_j represent the number of possible categorical values for x_i and y_j, respectively. The values for hits and expenses are typically derived from a high-fidelity simulation of the wargame scenario. However, in this project, we have utilized surrogate models to approximate these indicators during the optimization process. The entire process of building the surrogate models from the training data collected from the simulator is mentioned in the supplementary material.

4 Proposed Wargame Strategy Optimization Framework

The proposed framework consists of two interconnected components that work together to solve the WSOP. The first component is the optimization process, which interacts with the wargame to explore the search space strategies and identify optimized strategies for both offense and defense. The second component is the decision-maker, responsible for selecting a strategy for each side.

As mentioned before, we introduce a progressive shrinking approach where after each optimization cycle, some of the variables get fixed based on the decision-making and the optimization is re-run with reduced dimension. This ensures that the agents can freely switch from one strategy to another in the final optimized space and at the same time, the dimension reduction process narrows down the focus to smaller variations which helps in the final decision-making process. The entire workflow of this MoCoEv and decision-making is described in Figure 2.

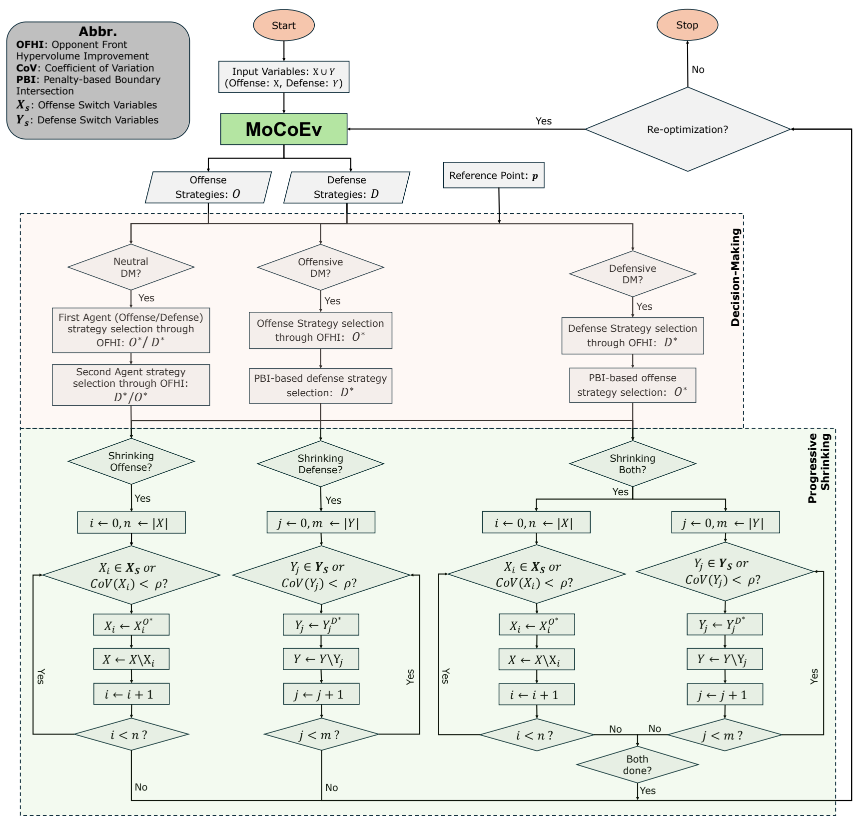

Fig. 2: The entire framework consisting of PS-MoCoEv and Co-evolutionary decision-making developed for WSOP.

4.1 Progressive Shrinking based Multi-obj. Co-Evol. (PS-MoCoEv)

The first phase of the process focuses on optimizing wargame strategies using MoCoEv. In this wargame, there are two agents: offense and defense. MoCoEv optimizes both agents simultaneously, considering their competitive objectives. The process works as follows: for a certain number of iterations (τ_1), MoCoEv optimizes the population of one agent based on its objectives while keeping the other agent's population fixed. During this time, the objective values for any strategy are computed by pairing it against every strategy from the fixed opponent population and computing the objective values for each pair, followed by an aggregation. In this work, we have used average-based aggregation, but median or worst-case aggregation measures can also be used. After τ_1 iterations are completed, the same process is applied to the other agent for a different set of iterations (τ_2). These cycles of improvement are applied to each agent in turn until a stopping criterion (e.g., total number of generations) is met. The optimization for each agent involves applying genetic operators to the parent population to generate offspring, followed by selecting the best solutions from the combined pool of parents and offspring. As the algorithm progresses, both agents evolve in turns, each being optimized alternately. We have used a co-evolutionary version of NSGA-II [19] as the underlying MoCoEv algorithm. A detailed description of the algorithm is provided in [20].

After the MoCoEv run is completed, an additional analysis step reduces the dimensionality of the search space by continuing with previously-taken moves leading to more focused optimization runs. In the current WSOP, if an agent wishes to transition to a new strategy, some variables (called switch variables) must remain consistent with the previously committed strategy. This step is crucial for applications like WSOP, where rapid responses are needed for strategic planning. For example, while the initial optimization and decision-making occur offline, we may need to make quick adjustments once the game is in progress. The goal is to ensure that, in the final set of optimal strategies, agents can practically and effectively make strategic shifts. At the same time, as a by-product, reduced dimensions help to narrow down the search to fewer variables, which might allow us to perform online optimization runs as well (online optimization is currently out of the scope of this paper).

To achieve this reduction, we use the decision-making process outlined in Section 4.2 to fix the switch variables based on automated decisions. After the initial MoCoEv run, we select one offense and one defense strategy and use their values of the switch variables to fix them. With the reduced problem dimensions, we run another MoCoEv to find optimal strategies for the simplified problem, as described in [21]. We then analyze the solutions, fix variables with a Coefficient of Variation (CoV) below a certain threshold ($\rho = 0.3$) [22], and do not use them in the next optimization phase. After each dimensionality reduction, we re-optimize using MoCoEv, progressively shrinking the search space and simultaneously reducing the number of generations needed for optimization. This method, termed Progressive Shrinking-based MoCoEv (PS-MoCoEv), can be applied to only offense, only defense, or both agents. The results of these different optimization approaches are discussed in Section 5.

4.2 Co-evolutionary Decision-Making

Co-evolutionary optimization problems introduce unique challenges, particularly in decision-making, where strategies for both agents must be considered simultaneously. This interdependence adds complexity to the process, making it essential to develop a decision-making approach tailored to the co-evolutionary context. We have created a novel decision-making process based on opponent front hypervolume improvement and PBI-oriented selection.

The process unfolds in two steps: selecting a strategy for the first agent and then selecting a strategy for the second agent. Suppose that we have N_1 strategies for the first agent and N_2 strategies for the second agent. For decision-making, we need to consider $N_1 \times N_2$ possible strategy combinations. When choosing a strategy for the first agent, we must evaluate it against all N_2 opposing strategies, forming a collection of N_2 objective value pairs (pairs as the problem is two-objective) for each candidate strategy. To select one strategy from the candidates, we use the concept of opponent front hypervolume improvement, which considers the entire fleet of solutions to ensure a comprehensive evaluation.

Once the first agent's strategy is selected, we focus on the N_2 strategies for the second agent. Here, the objective is to choose one of these opposing strategies based on the specific goals of the decision-making process. In this work, we explore three types of decision-making: Offensive, Defensive, and Neutral. The details of opponent front hypervolume improvement and PBI-based final selection are discussed in the following section.

Opponent Front Hypervolume Improvement In a co-evolutionary setting, when an agent is choosing a strategy, it must evaluate how that strategy performs against all possible strategies of its opponent. This means it cannot just look at a single objective score pair, it needs to consider a range of score pairs. To address this, we propose a new approach called "opponent front hypervolume improvement" which is specifically designed for competitive co-evolutionary optimization.

Here is how it works: For the first agent, we examine all its strategies. For each of these strategies, there are multiple corresponding opponent strategies. From these strategy combinations, we determine the non-dominated front (NDF) based on the opponent's objectives, which is different from the traditional approach. We then calculate the hypervolume of this NDF and aim to improve it. Essentially, this hypervolume pushes the worst-performing points with respect to the first agent's objectives closer to the first agent's goals. A higher hypervolume indicates that the first agent performs well against most of the opponent's strategies.

Let us clarify this concept with some examples. Suppose that we are focusing on offensive decision-making. In Figure 3, the orange circles represent all combinations of offense and defense strategies plotted in the objective space defined by `OffenseExpenditures-DefenseExpenditures` vs `-OffenseHits`. The offense aims to minimize both objectives, while the defense aims to maximize them.

Figures 3a and 3b show the distribution of objective scores for two specific offense strategies: `Offense_7` and `Offense_1`, respectively. The blue squares represent the objective distribution of the selected offense strategy against the defense's population of strategies. From these distributions, we can then identify the Pareto front based on the defense's objectives (maximizing both), and then calculate its hypervolume based on a reference point. According to the opponent front hypervolume improvement approach, we prefer the offense strategy that maximizes this hypervolume.

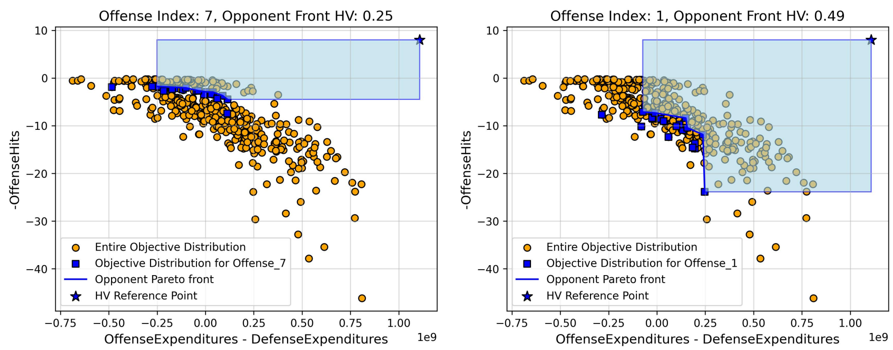

(a) The opponent front for `Offense_7` (b) The opponent front for `Offense_1`

Fig. 3: Opponent fronts for different offense strategies.

This approach essentially means that the distribution of the objectives is increasingly pushed towards meeting offensive goals, which is ultimately advantageous for the offense. Simply put, it indicates that the chosen offense strategy is more robust against a wide range of defense strategies. In this case, `Offense_1` in Figure 3b is preferred over `Offense_7` in Figure 3a because it achieves a higher hypervolume.

The same idea can be applied to defensive decision-making in the opposing sense (the front should be identified based on offense objectives). In neutral decision-making, we first apply the opponent front hypervolume improvement for offense and then for defense. The intersection of those two fronts represents the final strategies for offense and defense. For offensive and defensive decision-making, an extra step is performed through a PBI-based selection which is described next.

PBI-based selection In the case of non-neutral decision-making goals, after the first agent selects its strategy, the next step is to choose a strategy for the second agent. This process involves using a reference point ($\mathbf{p}$) and a reference angle (θ^*).

The reference point is specified by the user, and the reference angle is set to 45° for defensive decision-making and 225° for offensive decision-making, with the origin at $\mathbf{p}$. For defense, which aims to maximize both objectives ($-$*OffenseHits* vs *OffenseExpenditures* $-$ *DefenseExpenditures*), an angle of 45° is ideal. Conversely, for offense, which aims to minimize objectives, an angle of 225° is ideal. It is important to normalize both objectives before considering these angles.

We then examine the available strategies within the ranges of [0°, 90°] for defense and [180°, 270°] for offense. If no strategies fall within these ranges, the closest strategy (angle-wise) outside the range is selected. For each strategy within the specified range, we calculate a PBI (Penalized Boundary Intersection) metric as follows:

$$\text{PBI}(\mathbf{x}) = w_1 \times \hat{d\theta} - w_2 \times \hat{dx}. \quad (3)$$

Here, $\hat{d\theta} = \frac{|\theta - \theta^*|}{360}$ measures the normalized deviation of the current angle θ from the reference angle θ^*, and $\hat{dx}$ represents the normalized distance from the reference point $\mathbf{p}$ to the strategy point $\mathbf{x}$ along the direction of the reference angle, as shown in Figure 4. The goal is to minimize $\hat{d\theta}$ and maximize $\hat{dx}$. Ultimately, the strategy with the lowest PBI score is chosen as the strategy for the second agent. In our work, we have used $w_1 = 0.9$ and $w_2 = 0.1$.

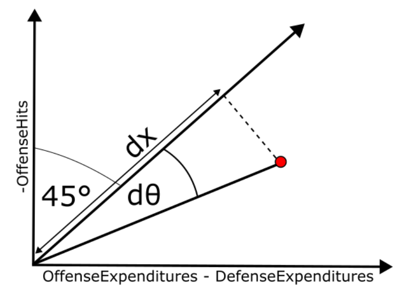

Fig. 4: illustration of the angular deviation ($d\theta$) and distance (dx) along the reference angle used for the calculation of the PBI metric. The reference point ($\mathbf{p}$) can be considered to be at the origin.

5 Results and Discussion

In this section, we begin by showing the results of the PS-MoCoEv run along with the decision-making process associated with it. The overall outcome depends on several key factors: the reference point used for PBI, the shrinking mechanism, and the decision-making strategy itself. To understand how changes in each of these factors affect the results, we conducted a scenario exploration study, which we describe at the end of this section.

5.1 Objectives

The objectives of the problem and the goals of the participating entities are specified in Table 1. Each objective is calculated using a surrogate model trained using grammatical evolution on high-fidelity simulation data.

Table 1: WSOP Objectives and Entity Goals.

i	**Objective (f_i)**	**Offense Goal**	**Defense Goal**
1	*OffenseExpenditures - DefenseExpenditures*	↓	↑
2	*OffenseHits*	↑	↓

5.2 Hyperparameters

As mentioned before, each PS-MoCoEv and decision-making run requires certain hyperparameters, which can be broadly classified into four groups: *DM strategy*, *Shrinking side*, *PBI Reference Point* and *MoCoEv algorithm paramters*. The *DM strategy* has three alternatives: offensive, defensive or neutral. The *Shrinking side* also has three options: shrinking offense only, defense only, or shrinking both. *Reference points* can be any point in the 2-D objective space suggesting the area of interest for the user. Finally, the *MoCoEv algorithm parameters* are the parameters used by the underlying NSGA-II approach like population size or number of generations.

5.3 PS-MoCoEv Round

First, we want to show the results of a PS-MoCoEv run. Any PS-MoCoEv run consists of multiple rounds. Every round is a combination of a MoCoEv run and a post-optimization co-evolutionary decision-making. Then the next round is run with a reduced variable dimension. For illustration purposes, we select a single hyperparameter configuration and discuss the outcome of the process. In this example, we have used offensive decision-making and dual shrinking strategy (shrinking both offense and defense variables) with a fixed PBI reference point of $(-5 \times 10^8, 0.25)$.

MoCoEv Run Results After the initial MoCoEv run, the evolved strategy combinations are displayed in Figure 5a. The black star represents the PBI reference point. In addition, Figure 5c and Figure 5e represent the coefficient of variation (CoV) of the offense and defense variables in the initial and final population of strategies, respectively. This analysis shows the degree of variation of different variables across the population of strategies. But, in the first round, we always fix the "switch variables" which are specified by our industrial collaborator. From the next round onwards, the variables with a CoV less than ρ are fixed.

Post-MoCoEv Decision-Making In order to reduce the dimension, we need to select the values of the variables to be fixed. To achieve this, the automated decision-making process chooses one offense strategy and one defense strategy. The combination of these strategies is represented by a green star in Figure 5a. During the preliminary round, we fix the values of the switch variables to those corresponding to the point marked by the green star. In subsequent steps, we fix the values of the variables with a Coefficient of Variation (CoV) less than ρ. This completes one round of the progressive shrinking process.

MoCoEv Re-Run After fixing the variables, the updated solutions might not be optimal based on these fixing constraints. So, following the concept of regularity in optimization as mentioned in [21, 23], we perform a re-optimization using a new MoCoEv. This cycle continues until all variables have CoV beyond the ρ threshold. In this particular scenario, there were 14 rounds of MoCoEv, Decision-Making, and Re-optimization cycles. Figure 5b shows the final optimized strategy distribution in the objective space. Interestingly, what variables

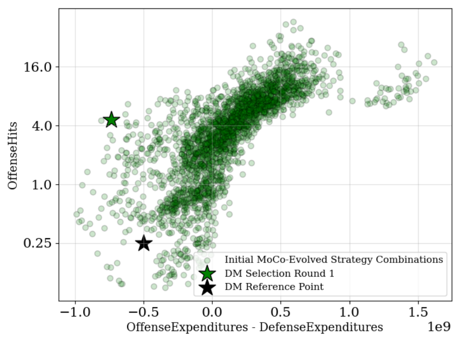

(a) Preliminary MoCo-Evolved optimized strategy combinations

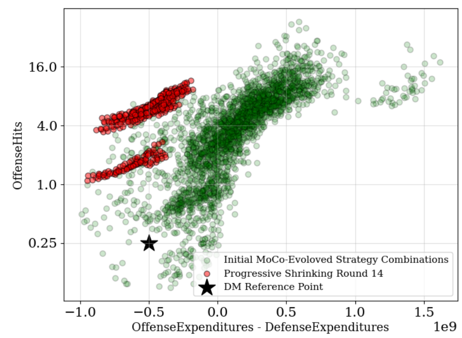

(b) Final round of MoCo-Evolved optimized strategy combinations

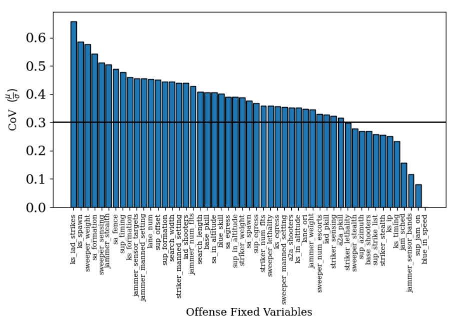

(c) The CoV for the offense variables across the population of preliminary MoCo-Evolved strategies

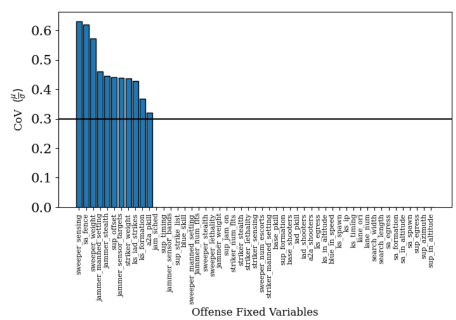

(d) The CoV for the offense variables across the population of the final round of MoCo-Evolved strategies

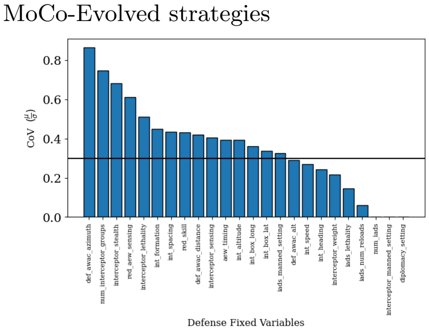

(e) The CoV for the defense variables across the population of preliminary MoCo-Evolved strategies

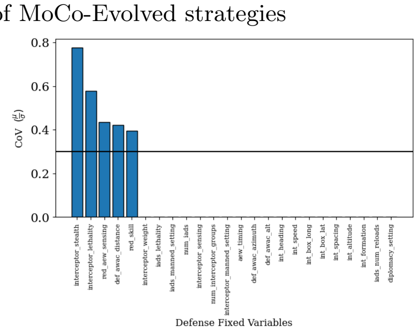

(f) The CoV for the defense variables across the population of the final round of MoCo-Evolved strategies

Fig. 5: Initial and Final round of MoCoEv run results for offensive decision-making, dual progressive shrinking and a reference point of $(-5 \times 10^8, 0.25)$.

have been reduced in the process can be observed for both offense (from Figure 5c to Figure 5d) and defense (from Figure 5e to Figure 5f). In the scenario under consideration, we have reduced the variable dimension from 74 to 16 through 14 progressive shrinking rounds.

5.4 Effect of PBI Reference Points, Shrinking Side and DM Strategy on Final Decision-Making

The outcome of progressive shrinking is clearly influenced by the chosen hyperparameters. To explore how these settings impact the final strategy, we keep NSGA-II algorithm parameters fixed and vary other problem-specific hyperparameters. Specifically, we test different decision-making strategies—defensive, offensive, or neutral—as well as shrinking sides, focusing on offense, defense, or both. Additionally, we evaluate four distinct PBI reference points. This is not intended to determine the best set of hyperparameter values. Instead, it is an exploratory analysis to observe how varying system configurations influence the outcomes of the approach.

The final points selected after the progressive shrinking run and successive decision-making processes are presented in Figure 6. Each plot illustrates the PBI reference point used in the study, along with the direction of the final selected points relative to the PBI reference points for various combinations of decision-making and progressive shrinking strategies. Additionally, the angles formed by the final points with respect to the PBI reference points (assuming a 2D origin at these points) are shown. The black outline in each plot represents the convex hull of the solutions found after the first MoCoEv run. In this angular representation, movements within the range of [90°, 180°] are categorized as offensive, while movements within [270°, 360°] are considered defensive. The remaining angular movements represent trade-off movements.

One key observation is that the final outcomes are highly dependent on the PBI reference points. For instance, from certain reference points, it is easier to find offensive strategy combinations (Figure 6c) using progressive shrinking, whereas other reference points tend to facilitate defensive solutions (Figure 6d). As expected, we generally observe more offensive movements when offensive decision-making is employed, and more defensive movements when defensive decision-making is applied. However, in some cases, offensive or defensive decision-making does not lead to the corresponding offensive or defensive movements, especially when both offense and defense are involved in progressive shrinking. This suggests that the progressive shrinking process can sometimes restrict movement in the desired direction, which is intuitive, as the degrees of freedom are reduced due to the shrinking process.

6 Conclusions

In this study, we have developed a novel multi-objective co-evolutionary framework based on progressive shrinking (PS), PS-MoCoEv, for optimizing strategic wargame scenarios. The framework incorporates PS techniques alongside decision-making strategies, allowing for the co-evolution of offensive and defensive strategies in tandem for competitive wargames. Through iterative optimiza-

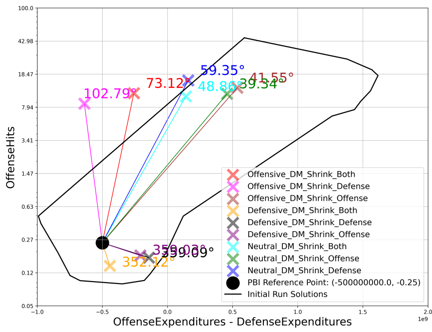

(a) Decision-making movement of selected strategy combination from a reference point of (-5e8, -0.25).

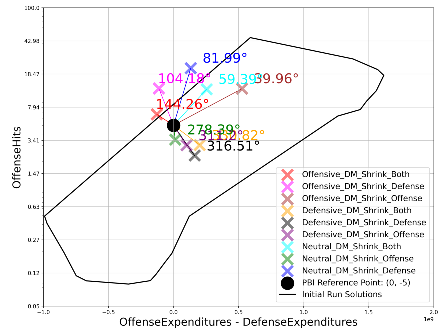

(b) Decision-making movement of selected strategy combination from a reference point of (0, -5).

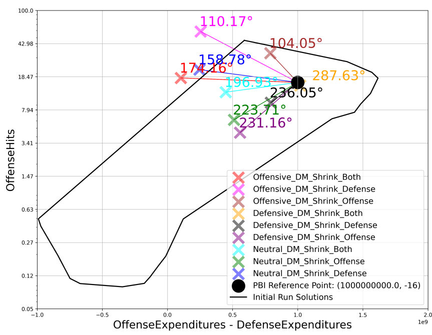

(c) Decision-making movement of selected strategy combination from a reference point of (1e9, -16).

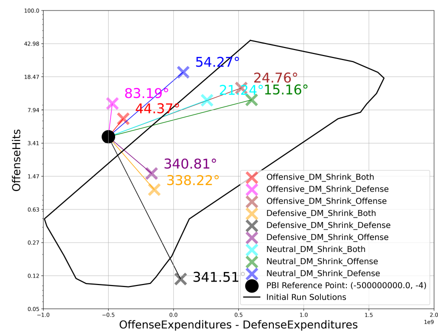

(d) Decision-making movement of selected strategy combination from a reference point of (-5e8, -4).

Fig. 6: Movement of the selected strategy combination after the entire PS-MoCoEv workflow using different PBI reference points.

tion and decision-making, our approach reduces the dimensionality of the problem space, maintaining the practicality of already taken moves to be continued in later moves.

The results of our various optimization studies have highlighted the significance of the PBI reference points, decision-making strategies, and shrinking mechanisms in determining the final strategy outcomes. In our specific scenario, better offensive strategies have emerged more often, while in limited initial scenarios, better defensive solutions have emerged. Moreover, our results have shown the sensitivity of the progressive shrinking approach in restricting latter moves in the desired direction due to reduced degrees of freedom yielded from previous moves. The proposed PS-MoCoEv stays as a systematic, algorithmic, and automated system for wargame strategy optimization, enhancing the efficiency of decision-making in complex, multi-agent military or other similar scenarios.

Future work may explore additional refinements to the decision-making process and further investigate the role of different intermediate steps for more efficient performance of PS-MoCoEV.

Acknowledgement

This research was supported by Northrop Grumman Corporation as a part of a collaboration between Computational Optimization and INnovation (COIN) Lab at Michigan State University and Northrop Grumman.

References

1. Geraldo Mulato De Lima Filho, André Rossi Kuroswiski, Felipe Leonardo Lôbo Medeiros, Mark Voskuijl, Herman Monsuur, and Angelo Passaro. Optimization of unmanned air vehicle tactical formation in war games. *IEEE Access*, 10:21727–21741, 2022.
2. Zhi-Xiang Jia and Jean-Fu Kiang. War game between two matched fleets with goal options and tactical optimization. *AI*, 3(4):890–930, 2022.
3. Atsuo Ozaki, Masakazu Furuichi, Katsumi Takahashi, and Hitoshi Matsukawa. Design and implementation of parallel and distributed wargame simulation system and its evaluation. *IEICE TRANSACTIONS on Information and Systems*, 84(10):1376–1384, 2001.
4. Peter J Schwartz, Daniel V O'Neill, Meghan E Bentz, Adam Brown, Brian S Doyle, Olivia C Liepa, Robert Lawrence, and Richard D Hull. Ai-enabled wargaming in the military decision making process. In *Artificial intelligence and machine learning for multi-domain operations applications II*, volume 11413, pages 118–134. SPIE, 2020.
5. Qingfu Zhang and Hui Li. Moea/d: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on evolutionary computation*, 11(6):712–731, 2007.
6. Jie Luo, John Cooper, Chengyu Cao, and Khanh Pham. Cooperative adaptive control of a two-agent system. In *2012 American Control Conference (ACC)*, pages 2413–2418. IEEE, 2012.
7. Yingli Li, Jiahai Wang, and Zhengwei Liu. A simple two-agent system for multi-objective flexible job-shop scheduling. *Journal of Combinatorial Optimization*, 43(1):42–64, 2022.
8. Kunfeng Wang, Chao Gou, Yanjie Duan, Yilun Lin, Xinhu Zheng, and Fei-Yue Wang. Generative adversarial networks: introduction and outlook. *IEEE/CAA Journal of Automatica Sinica*, 4(4):588–598, 2017.
9. Edmund Russell. Coevolutionary history. *The American Historical Review*, 119(5):1514–1528, 2014.
10. Nattavut Keerativuttitumrong, Nachol Chaiyaratana, and Vara Varavithya. Multi-objective co-operative co-evolutionary genetic algorithm. In *International Conference on Parallel Problem Solving from Nature*, pages 288–297. Springer, 2002.
11. Arash Atashpendar, Bernabé Dorronsoro, Grégoire Danoy, and Pascal Bouvry. A scalable parallel cooperative coevolutionary pso algorithm for multi-objective optimization. *Journal of Parallel and Distributed Computing*, 112:111–125, 2018.
12. Bernabé Dorronsoro, Grégoire Danoy, Antonio J Nebro, and Pascal Bouvry. Achieving super-linear performance in parallel multi-objective evolutionary algorithms by means of cooperative coevolution. *Computers & Operations Research*, 40(6):1552–1563, 2013.

13. Nicolás Garcıa-Pedrajas, César Hervás-Martınez, and José Munoz-Pérez. Multi-objective cooperative coevolution of artificial neural networks (multi-objective cooperative networks). *Neural Networks*, 15(10):1259–1278, 2002.
14. Chi Keong Goh, Kay Chen Tan, DS Liu, and Swee Chiang Chiam. A competitive and cooperative co-evolutionary approach to multi-objective particle swarm optimization algorithm design. *European Journal of Operational Research*, 202(1):42–54, 2010.
15. Ivan Reinaldo Meneghini, Frederico Gadelha Guimaraes, and Antonio Gaspar-Cunha. Competitive coevolutionary algorithm for robust multi-objective optimization: The worst case minimization. In *2016 IEEE congress on evolutionary computation (CEC)*, pages 586–593. IEEE, 2016.
16. Andrew R McIntyre and Malcolm I Heywood. Multi-objective competitive coevolution for efficient gp classifier problem decomposition. In *2007 IEEE International Conference on Systems, Man and Cybernetics*, pages 1930–1937. IEEE, 2007.
17. Fanchao Zeng, James Decraene, Malcolm Yoke Hean Low, Wentong Cai, and Philip Hingston. Studies on pareto-based multi-objective competitive coevolutionary dynamics. In *2011 IEEE Congress of Evolutionary Computation (CEC)*, pages 2383–2390. IEEE, 2011.
18. Leonardo Corbo, Sascha Kraus, Božidar Vlačić, Marina Dabić, Andrea Caputo, and Massimiliano M Pellegrini. Coopetition and innovation: A review and research agenda. *Technovation*, 122:102624, 2023.
19. Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.
20. Ritam Guha, Ryan Mckendrick, Bradley Feest, and Kalyanmoy Deb. Attacker-defender strategy optimization using multi-objective competitive co-evolution. In *International Conference on Parallel Problem Solving from Nature*, pages 351–366. Springer, 2024.
21. Ritam Guha and Kalyanmoy Deb. Compromising pareto-optimality with regularity in platform-based multi-objective optimization. *IEEE Transactions on Evolutionary Computation*, 2023.
22. Coefficient of variation. Accessed on September 27, 2023.
23. Ritam Guha and Kalyanmoy Deb. RegEMO: Sacrificing pareto-optimality for regularity in multi-objective problem-solving. In *International Conference on Evolutionary Multi-Criterion Optimization*, pages 29–42. Springer, 2023.