

Building Large-Scale Drone Defenses from Small-Team Strategies

Grant Douglas
Adelaide University
Adelaide, Australia
grant.douglas@adelaide.edu.au

Claudia Szabo
Adelaide University
Adelaide, Australia
claudia.szabo@adelaide.edu.au

Stephen Franklin
Saab Australia Pty Ltd
Adelaide, Australia

Mingyu Guo
Adelaide University
Adelaide, Australia
mingyu.guo@adelaide.edu.au

ABSTRACT

Defending against large adversarial drone swarms requires coordination methods that scale effectively beyond conventional multi-agent optimisation. In this paper, we propose to scale strategies proven effective in small defender teams by integrating them as modular components of larger forces using our proposed framework. A dynamic programming (DP) decomposition assembles these components into large teams in polynomial time, enabling efficient construction of scalable defenses without exhaustive evaluation. Because a unit that is strong in isolation may not remain strong when combined, we sample across multiple small-team candidates. Our framework iterates between evaluating large-team outcomes and refining the pool of modular components, allowing convergence on increasingly effective strategies. Experiments demonstrate that this partitioning approach scales to substantially larger scenarios while preserving effectiveness and revealing cooperative behaviours that direct optimisation cannot reliably discover.

KEYWORDS

Genetic Algorithms; Drone Swarms; Scalability

ACM Reference Format:

Grant Douglas, Stephen Franklin, Claudia Szabo, and Mingyu Guo. 2026. Building Large-Scale Drone Defenses from Small-Team Strategies. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. <https://doi.org/10.65109/GYJS4496>

1 INTRODUCTION

Coordinating large teams of autonomous drones in adversarial environments is a formidable challenge. Even with modest numbers of agents, the combinatorial explosion of possible strategies quickly renders exhaustive search and naive optimisation infeasible. Reinforcement learning (RL) and skill-discovery methods have shown promise for small teams (e.g., [4, 15, 23, 27]), but they often scale poorly due to high sample complexity, non-stationarity, sparse rewards or computational overhead [15].

In this paper, we present a staged pipeline for building large-scale drone defenses from small team strategies. The key idea is to first identify and evolve promising behaviours in small defender teams, and then systematically reuse these strategies as modular components for larger forces. Dynamic programming (DP) is used to assemble these components in polynomial time, transforming what would otherwise be a super-exponential search into a tractable optimisation [19]. The resulting solution enables the efficient construction of coordinated swarm defenses that scale to realistic adversarial scenarios. While prior UAV swarm defense studies generally evaluate on no more than 12 agents (e.g., 10v10 in [24] and 5v5 in [18]), we demonstrate scalability to attacker swarms up to thirty agents and defender swarms up to 45 agents with stable execution, substantially exceeding the tested ranges of existing learning-based methods.

Our approach deliberately uses heuristics rather than direct policy learning. Heuristics offer several advantages critical for scalability: (i) they require limited manual scripting, enabling construction of diverse policies relatively quickly, (ii) can be expressed in compact mathematical forms for efficient JAX implementation [3], (iii) they can be simulated at scale without the overhead of RL training, and (iv) at runtime, execution only requires lightweight CPU-level evaluation of heuristic code - in contrast to neural network inference - enabling deployment on smaller onboard processors with modest memory. Prior work has explored heuristic evolution in multi-agent systems, including modular evolution [16, 22] and messy genetic algorithms [7], but we adapt these ideas to swarm defense by framing heuristics as sub-chromosomes in evolutionary search.

The contributions of this work are both methodological and engineering-oriented. We present a hybrid GA-DP framework leveraging chromosome factorisation for efficient evolutionary search and demonstrate its scalability on large-scale drone defense simulations. Implemented fully in JAX and deployed on high-performance computing infrastructure, the framework integrates algorithmic structuring that exploits priors and hierarchical decomposition to avoid exhaustive search while remaining tractable, even across billions of simulation time steps. To expand the behavioural search space, we additionally use large language models to generate and diversify candidate heuristics. Together, these elements enable experiments at scales that would otherwise be computationally prohibitive, positioning the framework as a practical foundation for future swarm defense research and applications.

The primary contributions of this work are fourfold:



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/GYJS4496>

- (1) A hybrid GA–DP framework for scalable defense. We propose a framework that combines genetic algorithms to evolve effective chromosomes in small-scale engagements with dynamic programming to assemble them into strategies for larger swarms. This allows tractable analysis at scales that naive evolutionary methods cannot reach [16, 19].
- (2) DP integration with a hierarchical defense policy. By embedding DP-guided allocations within a hierarchy, our method coordinates defenders at the sub-team level while retaining heterogeneity at the agent level, enabling effective pursuit–evasion dynamics in complex swarm encounters.
- (3) A chromosome factorisation perspective on evolutionary search. We show that the effective building blocks in this domain are entire chromosomes (combinations of heuristics) rather than individual genes, highlighting the importance of preserving high-order dependencies [7].
- (4) Iterative refinement through large-scale evaluation. By feeding back the outcomes of large-scale evaluations into the small-team heuristics, our approach progressively improves the accuracy of heuristic estimates and converges toward high-performing defender strategies at scale. This iterative design provides solutions to the problem, as well as a proven methodology for their systematic improvement.

We demonstrate that this combined GA-DP approach scales to large defensive swarms, uncovering cooperative strategies that are difficult to evolve with conventional representations. Collectively, these contributions provide a practical and scalable pathway for optimizing coordinated behavior in adversarial multi-agent systems.

2 RELATED WORK

We survey prior research on multi-agent coordination and strategy optimisation, highlighting approaches that inform our hierarchical heuristic framework.

Multi-Agent Defense and Pursuit–Evasion. Defending a target from coordinated attackers is a classic multi-agent control problem, often studied through pursuit–evasion games and adversarial swarm interactions. Algebraic and control-theoretic formulations yield insights for small engagements [10], but rapidly become intractable as team sizes increase [9]. Reinforcement learning approaches [15] can scale somewhat further, but face challenges of non-stationarity, sparse rewards, and high sample complexity, especially in competitive–cooperative domains such as drone defense [13, 22]. These limitations motivate methods that exploit structure or heuristics for tractability.

Evolutionary Algorithms in Multi-Agent Systems. Evolutionary computation offers a population-based search mechanism suited to the combinatorial complexity of multi-agent coordination. Genetic algorithms (GAs) have been applied to tasks such as role assignment, cooperative foraging, communication, and swarm robotics [16, 22]. However, scalability remains an issue, as the search space grows combinatorially with the number of agents. The messy genetic algorithm (mGA) [7] sought to preserve high-quality building blocks under deceptive landscapes, but its stochastic recombination remains expensive at scale. Our approach takes inspiration from these ideas but emphasises explicit chromosome factorisation and systematic reuse rather than probabilistic recombination.

Dynamic Programming (DP) for Team Optimisation. DP has been widely studied for coalition formation, task allocation, and sensor coverage in MAS [19]. DP leverages problem structure to find exact or near-optimal decompositions, but faces exponential blow-up when applied directly to large swarms. Hybrid approaches that combine DP with other mechanisms are therefore attractive, as they can preserve structure while avoiding combinatorial explosion.

Dynamic Target Assignment. Classical weapon–target assignment (WTA) focuses on the discrete assignment of weapons to static targets to maximise expected damage or minimise surviving threats, typically solved via integer programming or dynamic programming [8, 11]. In contrast, our framework addresses continuously moving attackers with kinematic constraints, where assignments must account for multi-agent interactions. Furthermore, whereas DP in classical WTA computes exact or relaxed optimal solutions [2], we employ DP as a combinatorial decomposition to guide evolutionary search, structuring chromosome sampling without requiring full optimality. This approach bridges discrete assignment principles with continuous, dynamic multi-agent engagements, enabling scalable strategy optimisation.

Hierarchical Policies in Swarm Control. Hierarchical organisation is a natural way to manage complexity in MAS. High-level policies assign roles or allocate agents to sub-problems, while low-level policies execute primitive behaviours such as pursuit or patrol. This paradigm has been explored both in reinforcement learning [5, 12, 17] and in hybrid heuristic–learning frameworks. Our approach falls into this class but differs in two respects: (i) we employ heuristics as the low-level policies, chosen for their largely automatic generation enabling diversity and efficient JAX implementation, and (ii) we embed DP-based allocation within the hierarchy. Consequently, they facilitate coordination to emerge from structured assignments, enabling scalability without requiring each agent to learn full joint coordination from scratch.

To our knowledge, no prior work has combined evolutionary search, DP, and heuristic factorisation within a hierarchical framework for drone defense. This integration enables tractable scaling from small evolved strategies to large swarm engagements, filling a gap left by existing MARL, GA, and DP-only approaches.

3 APPROACH

3.1 Problem Formulation and Setting

We formalise the drone defense scenario as a multi-agent coordination problem, detailing agent dynamics, objectives, and interactions, and introduce our hierarchical heuristics solved via a hybrid genetic algorithm and dynamic programming. The environment consists of an open two-dimensional battlefield with a fixed critical asset (target) that the Blue team must defend (see Figure 1). A Red team of attackers attempts to reach the target from randomised starting positions (radius and angle to target) following sinusoidal paths (randomised amplitudes and frequencies), introducing nonlinear and unpredictable motion that prevents defenders from relying on simple linear extrapolation. While not fully realistic, this design creates an informative testbed for evaluating hierarchical and heuristic coordination strategies. The objective of the Blue team is to maximise win rate - by preventing Red agents from reaching the

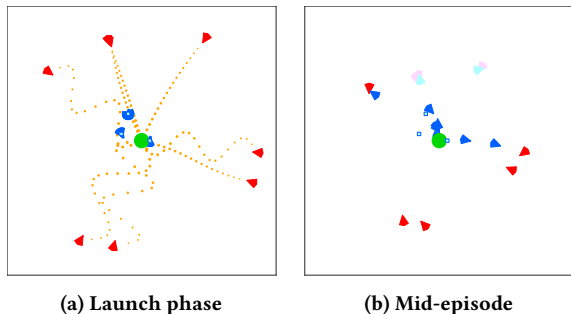


Figure 1: Scenario with 10 defenders (blue) protecting a target (green) from 8 attackers (red). In (a) attackers are shown with their offensive paths (orange) toward the target. In (b) collided drones that no longer participate have muted coloring.

asset - in contrast to much of the prior work, which often focuses on agent rewards [24] or attrition [18, 21].

Each agent’s state comprises its position (x, y) , heading θ , and speed v , with actions defined by turn and acceleration commands constrained within fixed bounds. Agents have full visibility of all teammates and opponents. Blue team coordination emerges through a hierarchical heuristic framework: each agent executes a selected heuristic, while a hybrid genetic algorithm combined with dynamic programming optimises the heuristic assignments and parameters to maximise team-level performance.

3.2 Overview

Each Blue agent is controlled by a selected heuristic, which determines its low-level actions in the environment. Team-level coordination emerges through the assignment of heuristics to agents, rather than through explicit joint planning at every time step. To optimise these assignments and any associated heuristic parameters, we employ a hybrid genetic algorithm combined with dynamic programming, enabling efficient exploration of the combinatorial space of possible team configurations. The blue agents have no prior knowledge of attacker trajectories. Defensive strategies are optimised via the GA based on estimating strategies and evaluating them through interaction. In this way our approach is agnostic to the attacker model.

Figure 2 illustrates the pipeline within the framework, showing how heuristic selection, environment interaction, and GA-DP optimisation are integrated to maximise the Blue team’s win rate. The simulation itself follows a standard Markov Decision Process (MDP) structure: at each time step, all agents submit actions, the environment transitions accordingly, and new observations are returned (along with potential rewards, though these are unused here). Unlike conventional reinforcement learning approaches, all decision-making in our framework is purely heuristic-driven, with coordination emerging from GA-DP optimisation rather than from policy-gradient or value-based learning.

This modular approach enables scaling by reusing and combining optimised heuristics without retraining.

3.3 Hierarchical Policy as Chromosomes

In our framework, the hierarchical policy for the Blue team is encoded as a chromosome suitable for GA optimisation. Each chromosome represents a full team configuration of N agents, where each agent’s gene encodes its heuristic, spawn location, and heuristic-specific parameter. Formally, the chromosome C can be written:

$$C = (g_1, g_2, \dots, g_N), g_i = (H_i, W_i, P_i)$$

where g_i is the gene of agent i , H_i is the selected heuristic, W_i is the spawn location, and P_i is a heuristic-specific parameter. This representation allows the GA to explore combinations of agent strategies while maintaining the hierarchy of decision-making.

The hierarchical policy separates low-level behaviour, determined by the agent’s heuristic H_i , from high-level coordination, determined by the assignment of heuristics across the team C . During optimisation, crossover and mutation operate at the gene level, allowing recombination of heuristics and parameters between team configurations while preserving feasible individual behaviours. Dynamic programming complements the GA by efficiently evaluating the combinatorial space of heuristic assignments, enabling scalable optimisation of large teams. This chromosome-based representation ensures the search respects the hierarchy of decisions and allows the hybrid GA-DP to maximise team-level win rate effectively.

3.4 Heuristics and Generation

Each gene in a chromosome encodes a heuristic H_i that governs the low-level behaviour of the corresponding Blue agent. The heuristics define action selection based on the current environment state, including the positions of Red agents, teammates, and the target. Heuristics are parameterised to allow flexibility in behaviour. By varying these parameters, a single heuristic template can generate multiple distinct strategies, increasing the diversity of candidate actions available for the GA-DP to optimise.

The use of heuristics is a deliberate design decision to avoid the complexity of skill discovery approaches such as feudal networks [23], variational auto-encoders [1], or general value functions [4], while still allowing explicit incorporation of action-masking and death-masking without additional training or the complexity of attention mechanisms.

We design a set of modular heuristics that capture a range of tactical behaviours, such as intercepting attackers, maintaining defensive formations, and prioritising high-threat targets. Each heuristic encodes a distinct strategy, allowing the framework to combine them into larger teams while preserving interpretability and explicit action control. While one heuristic may be more familiar or conventional, focusing solely on it risks missing emergent cooperative behaviours. Accordingly, the heuristic set is curated for both diversity and effectiveness.

We utilize a combination of hand-crafted and LLM-generated heuristics. LLMs are employed to generate candidate heuristics due to their capacity to generate diverse strategies from concise textual prompts. The LLMs were provided a one-shot attempt at coding against our Python API, and made only minor coding errors. This allows us to substantially accelerate exploration of the heuristic design space and facilitate discovery of novel, and potentially effective, tactical behaviours with relatively less coding effort

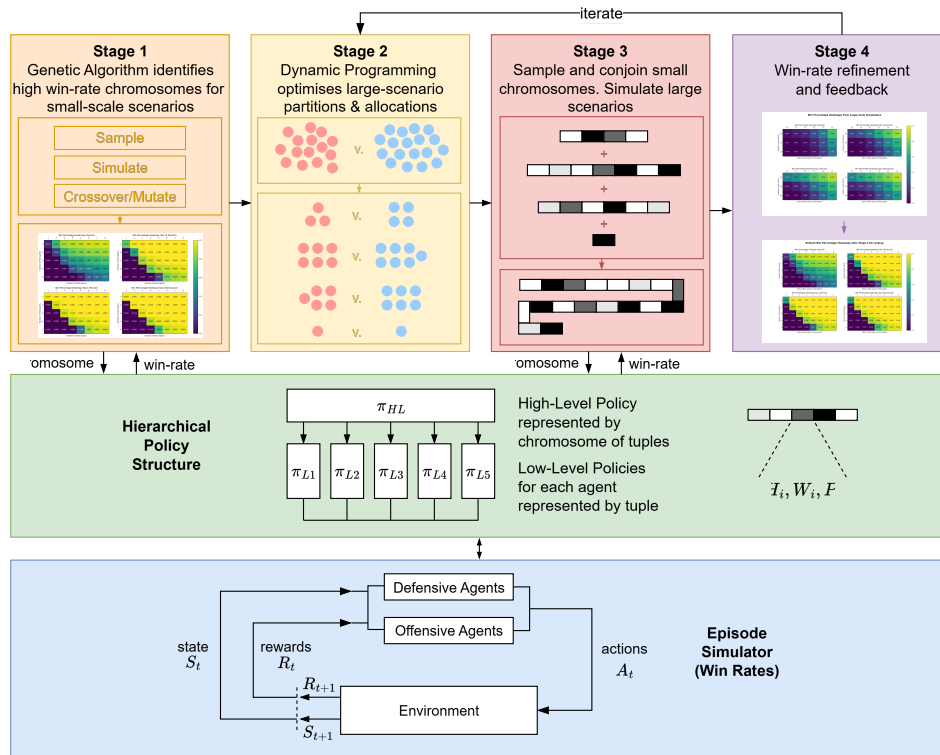


Figure 2: Overall hybrid GA–DP framework. The four-stage pipeline (Stage 1: GA evolution, Stage 2: dynamic programming allocation, Stage 3: chromosome sampling, Stage 4: refinement) operates over a hierarchical policy structure assigning heuristics to agents. These heuristics drive all decision-making. The simulator follows a standard MDP cycle, but no reinforcement learning is employed—coordination arises entirely from heuristic assignment and GA–DP optimisation.

(only doing the JAX conversion manually) or other expert training. Unlike heuristics discovered through MARL, LLM-generated heuristics are immediately executable and generally interpretable. As all generated heuristics are assessed solely through quantitative simulation under identical conditions, any biases inherent in the language model are neutralised by the evolutionary selection process, ensuring only empirically effective strategies persist.

3.5 Stage 1 - GA for Chromosome Evolution

In Stage 1 of our hybrid approach, a genetic algorithm (GA) is applied to small-scale team configurations—Red teams of 1–5 agents and Blue teams of 1–8 agents—to explore the space of hierarchical policies represented as chromosomes. Each chromosome encodes a full Blue team configuration, specifying the heuristic, spawn location, and parameterisation for each agent. Chromosome fitness is evaluated as the mean win rate over 128 simulated episodes.

The GA operates with a random initial population of 1024 chromosomes. Testing of Sobol [20] and Latin Hypercube [14] sampling population initialisation methods showed immaterial difference in performance. Evolution proceeds using 20% elitism, tournament selection, and single-point crossover (at 70%). Mutation introduces stochastic changes to genes to maintain population diversity. These


operators allow the GA to efficiently explore a wide range of team-level strategies, exploiting high-performing chromosomes while exploring novel configurations.

Figure 3 illustrates the GAs impact on team performance: heatmaps show win rates as a function of Blue and Red agent counts, comparing the random initial population with the GA-optimised population. The results demonstrate that GA evolution substantially improves team performance across small-scale configurations.

By restricting GA application to small-scale teams, we enable computationally feasible evaluation while producing a diverse set of high-quality chromosomes. These chromosomes serve as the candidate pool for Stage 2, where dynamic programming is applied to optimally assign heuristics and scale strategies to larger teams.

3.6 Stage 2 - DP for Chromosome Allocation

In the second stage, we address the core challenge of scaling defender policies from small scenarios to engagements with a larger number of attackers. Directly evolving effective chromosomes for scenarios with ten or more attackers is computationally prohibitive, as the search space grows combinatorially with the number of agents. Instead, we leverage the observation that outcomes of small-scale encounters, such as 2 vs. 3 or 3 vs. 4, can be represented compactly as empirical success probabilities obtained in Stage 1.



(a) Heatmap without GA application. (b) Heatmap following 128 GA generations.

Figure 3: Heatmaps showing mean population win rates for the top-performing 10 percentile and the top-performing individual chromosome, both before (a) and after (b) application of GA generations. The latter, demonstrating significantly improved win rates for small scale swarm-defense scenarios, are utilised in subsequent stages of the approach.

These outcomes provide a lookup table of defender effectiveness against subgroups of attackers under varying team sizes.

Stage 2 uses a DP search to compose these small-scale results into effective strategies for larger engagements. The central idea is to represent a set of R attackers as a partition into subgroups, for example $R = 10 \rightarrow \{3, 3, 4\}$. Each subgroup is assumed to be engaged independently by a sub-team of defenders. For any given partition, the DP algorithm explores feasible allocations of the available defenders across the subgroups, consulting the Stage 1 heatmaps to evaluate the probability of success for each allocation.

Let $P(r, b)$ denote the probability of successfully defending against r attackers with b defenders, as estimated from Stage 1. For a partition $\pi = \{r_1, r_2, \dots, r_k\}$ of R , and an allocation $\alpha = \{b_1, b_2, \dots, b_k\}$ of B , the joint success probability is given by:

$$V(\pi, \alpha) = \prod_{i=1}^k P(r_i, b_i) \quad (1)$$

The DP search enumerates partitions of R , and for each partition computes the allocation α that maximises $V(\pi, \alpha)$. By caching intermediate results, the DP approach avoids redundant computation, enabling efficient evaluation of all possible subgroup structures and defender splits.

This stage therefore produces a set of high potential candidate defender configurations for larger-scale swarm-defense scenarios, which are then passed forward to Stage 3. In addition, this stage also enables exploration by trade-offs between success and resource expenditure by varying the number of defenders B , supporting systematic investigation of efficiency, for example through determining whether an 11-defender team achieves significantly higher success than a 10-defender team against the same set of attackers.

Consider the case $R = 8$ attackers and $B = 10$ defenders. The DP routine computes $f(8, 10)$, the best defense probability for this scenario. One possible split is partition to $\{5, 3\}$ attackers and allocate $\{7, 3\}$ defenders. The value for this split is computed as:

$$f(8, 10) = f(5, 7) \times f(3, 3)$$

Algorithm 1 Dynamic Program for Partitioning and Allocation

Require:

- R (number of attackers)
- B (number of defenders)
- $P(r, b)$ (Stage 1 success probabilities)

Ensure: Best partition π^* and allocation α^*

```

function BESTDEFENSE( $R, B$ )
  if  $R = 0$  then
    return (success = 1, allocation = [])
  end if
  if cache[ $R, B$ ] exists then
    return cache[ $R, B$ ]
  end if
   $bestValue \leftarrow 0$ 
   $bestSolution \leftarrow \text{None}$ 
  for each partition  $\pi = \{r_1, \dots, r_k\}$  of  $R$  do
    for each allocation  $\alpha = \{b_1, \dots, b_k\}$  with  $\sum b_i = B$  do
       $value \leftarrow \prod_{i=1}^k P(r_i, b_i)$ 
      if  $value > bestValue$  then
         $bestValue \leftarrow value$ 
         $bestSolution \leftarrow (\pi, \alpha)$ 
      end if
    end for
  end for
  cache[ $R, B$ ]  $\leftarrow (bestValue, bestSolution)$ 
  return cache[ $R, B$ ]
end function

```

Another split might be $\{4, 3, 1\}$ attackers with defender allocation $\{5, 4, 1\}$. By considering all such splits and retaining the maximum value, the algorithm builds $f(8, 10)$ from small sub-problems. Each of these is computed recursively, terminating in base cases where outcomes are directly available from the Stage 1 heatmaps. Once a value is computed, it can be used in all higher-level states, thus avoiding redundant recomputation.

By way of contrast, a brute force approach would require enumerating all integer partitions of R , and for each partition considering all feasible allocations of defenders. The number of (unordered) partitions of an integer R is given by the Bell number B_R which approaches super-exponential growth. For a partition of size k , the number of possible allocations of defenders (including an empty allocation) is $\binom{B+k-1}{k-1}$. Combining these, it becomes clear that overall complexity of brute force means it is intractable for all but small values of R . The DP approach on the other hand has polynomial complexity.

PROPOSITION 1. *Let R and B denote the number of red attackers and blue defenders, respectively, and let k denote the maximum subgroup size of red agents for which precomputed outcomes are available. Then Algorithm 1 has worst-case time complexity*

$$O(R^3 B^3),$$

which is polynomial in R and B for fixed k .

3.7 Stage 3 - Chromosome Sampling with Priors

The role of Stage 3 is to transition from small-scale empirical knowledge (Stage 1) and dynamic allocations (Stage 2) to full-sized configurations that can be tested in the target scenario. Rather than relying on evolutionary operators, Stage 3 systematically samples and combines high-performing small-scale chromosomes.

Stage 3 commences by sampling from the distribution of promising sub-chromosomes discovered in Stage 1, then assembling them into full chromosomes consistent with the allocations identified in Stage 2. For example, if the DP procedure recommends that a sub-team of three defenders be allocated to a partition of five attackers, Stage 3 samples multiple candidate sub-teams from the high-performing small-scale chromosomes identified in Stage 1. This sampling introduces diversity, recognising that a strong sub-team in isolation may not retain its effectiveness when combined with other sub-teams. Repeating this process across all partitions generates a population of candidate full-team chromosomes for evaluation in the full simulation.

These assembled chromosomes are then evaluated directly in the full simulation environment, where the stochastic dynamics and interdependencies between sub-teams are faithfully represented. This step is crucial, as the independence assumptions underpinning Stage 2's DP analysis can only provide an approximation of real-world multi-agent interactions.

By systematically joining and testing factorised chromosome components, Stage 3 achieves two goals: (i) it leverages prior knowledge to drastically reduce the search space compared to naïve sampling, and (ii) it ensures that the retained strategies are robust to the complexities of large-scale multi-agent coordination.

3.8 Stage 4 - Iterative Refinement

During Stage 4, win rate performance of the full-team chromosomes are attributed back to the sub-chromosomes that were utilised. For each sub-chromosome, we accumulate a record of its performance across all simulations in which it appeared (it may be utilised in several full-team chromosomes), enabling the calculation of a mean win rate that reflects its effectiveness in different contexts. This

process establishes an empirical performance profile for each sub-chromosome, which serves as the basis for the weighted refinement applied in Stage 4. By grounding sub-chromosome estimates in repeated observations, Stage 3 provides the link between small-scale evaluations and the more comprehensive large-scale simulations.

For each sub-chromosome utilised in the large-scale evaluations, its win rate estimate is refined by combining priors from Stage 1 with empirical evidence accumulated in Stage 3. Specifically, the refined estimate, $w_{new}^{(1)}$, is obtained through a weighted combination:

$$w_{new}^{(1)} = \frac{K}{n+K} w_{old}^{(1)} + \frac{n}{n+K} w^{(3)}$$

where $w_{old}^{(1)}$ is the current Stage 1 estimate, $w^{(3)}$ is the average performance observed across Stage 3 simulations, and n is the number of times the sub-chromosome has featured in those simulations. Here K is a hyper-parameter that controls the relative influence of the prior versus empirical data.

The value of K was tuned via sensitivity analysis and trial-and-error, balancing the need for meaningful updates against the risk of overreacting to noisy outcomes. A value of $K = 200$ was ultimately selected, providing smooth yet efficient refinement of the Stage 1 estimates.

This weighted refinement ensures that sub-chromosomes with consistent large-scale success are promoted, while those that underperform are down-weighted, even if they appeared promising in small-scale settings. By iteratively applying this procedure across all evaluated chromosomes, Stage 4 produces more reliable win rate estimates for each building block. As evaluations accumulate, the refined win rates converge toward values that accurately reflect performance under large-scale conditions. Stage 4 thus validates composite strategies at scale while iteratively sharpening the estimated effectiveness of the underlying building blocks.

4 RESULTS AND DISCUSSION

4.1 Baseline

Performance in all cases is measured as the defender win rate, defined as the proportion of simulations in which defenders successfully prevent attackers from reaching the target. Since prior work in multi-agent attack–defense scenarios has focused almost exclusively on small-scale settings (typically involving fewer than ten agents) or algebraic formulations that do not scale, no directly comparable baselines exist in the literature. Accordingly, we construct two internal baselines to anchor our large-scale evaluation.

Random population. In this baseline, defender chromosomes are generated by assigning heuristics and parameters at random. This provides a lower-bound performance level and establishes how well defenders can perform without any optimisation.

Genetic algorithm. In this baseline, we apply the same GA used to optimise the small-scale scenarios in Stage 1 directly to large-scale scenarios, starting from a random initial population. This represents the natural alternative to our staged approach, in which a conventional GA is tasked with discovering effective defender strategies without factorisation or staged refinement.

Together, these baselines capture the extremes of unstructured search: a purely random allocation and a flat GA applied directly to the full problem. By comparing against these, we can demonstrate

the benefits of staged factorisation, DP, and large-scale refinement in producing scalable and effective defender strategies.

As the GA baseline only exhibits minor performance enhancement over the random baseline, we show only the former in Figure 4a. Overall win rates remain close to zero across most scenarios, and even the best-performing chromosomes achieve only modest outcomes (e.g., a maximum win rate of 0.328 for 10 Red at a 1.5 Blue/Red ratio). Performance collapses almost entirely for larger attacker swarms, with even elite win rates rarely exceeding 0.1.

These results clearly establish the difficulty of the large-scale defense problem, setting the stage for Section 4.2, where we demonstrate our hybrid GA–DP approach substantially outperforms baselines, achieving high win rates even against dense attacker swarms.

4.2 Comparative Results

We perform an ablation study to quantify the contribution of each stage in our hybrid GA–DP pipeline. Figure 4 present heatmaps of defender win rates for three configurations: (a) GA baseline, (b) Stage 3: DP-informed chromosome sampling, and (c) Stage 4: Iterative Refinement. This allows a clear evaluation of the incremental benefit of each component.

Stage 3 Ablation (DP-informed Chromosome Sampling). Stage 3 (Figure 4b) combines the dynamic programming allocation from Stage 2 with sampling of promising small-scale sub-chromosomes to assemble candidate full-team strategies. This stage alone substantially improves performance relative to the GA baseline (Figure 4a), producing nontrivial defender strategies across swarm sizes. For example, at $R = 20$ and a 1.5 Blue/Red ratio, the best chromosome achieves a win rate of 0.69 compared to just 0.13 in the baseline. However, Stage 3 alone struggles to optimise: while performance improves, win rates for some smaller Blue/Red agent ratios remain modest (e.g., 0.17 at $R = 30, B/R = 1.3$). These results highlight that DP-informed sampling establishes a strong foundation but is insufficient for fully robust large-scale strategies.

Stage 4 Ablation (Iterative Refinement). Incorporating Stage 4 refinement (Figure 4c) significantly amplifies performance across all scenarios. Iterative re-weighting and evaluation of sub-chromosomes consistently produces chromosomes that outperform Stage 3, particularly at lower Blue/Red ratios where coordination is critical. For instance, at $R = 20$ and $B/R = 1.3$, the top 10 percent improves from 0.22 (Stage 3) to 0.45 after Stage 4 refinement. For smaller swarms, defenders approach near-complete reliability, with win rates exceeding 0.80 at $R = 10$ and $B/R = 1.3$. Even the most challenging scenarios with $R = 30$ achieve a top win rate of 0.52, far above the negligible performance of the baseline.

Overall, this ablation study demonstrates that Stage 3 establishes a foundation of promising strategies, but Stage 4 is essential for convergence toward robust, high-performing defender allocations. This refinement allows the framework to scale efficiently, preserving effectiveness even as the size and complexity of the attacking swarm increase.

4.3 Discussion

Scalable swarm defense poses an inherently hard optimisation problem due to the explosive combinatorial growth of possible defender

allocations. The staged approach presented here substantially outperforms the baseline (random initialisation + GA) and achieves robust performance across increasingly challenging scenarios.

Breaking complexity into solvable pieces. The central advantage of the staged pipeline lies in factorisation of the search space. By first evolving small chromosomes in Stage 1, we obtain candidate building blocks under low-dimensional constraints. Stage 3 then leverages these factors through dynamic programming, assembling them into large-scale defender allocations. Finally, Stage 4 corrects biases introduced in earlier stages by iteratively refining chromosomes under direct evaluation in large scenarios. This decomposition dramatically reduces the effective search complexity compared to tackling large-scale swarm defense directly.

Refinement turns small-scale wins into large-scale victories. The comparison of Stages 3 and 4 highlights that some heuristics are strongly context-dependent. Certain heuristics perform well in small battles but fail to generalise when scaled, while others remain robust across sizes. Pruning experiments show that removing low-frequency heuristics has little effect, whereas removing the most frequent, high-impact heuristic substantially reduces performance, indicating that robustness depends on preserving key strategies. Iterative refinement in Stage 4 exposes these differences and converges on strategies that persist under large swarm conditions. This underscores the importance of multi-scale evaluation: assessing only at small scales risks over-committing to heuristics that collapse when scaled up.

Synergy, not parts: why chromosomes matter. An additional experiment explored gene-level sampling in Stage 3, where individual tuples were recombined independently rather than preserving entire chromosomes. Results were markedly poor. High-performing strategies appear to depend on synergistic combinations of tuples rather than on the marginal quality of individual components. Tuples that appear promising in isolation may be harmful if overused, while seemingly weak tuples can be critical in the right context. This resonates with the idea of deceptive landscapes [6, 26] and with findings in modular evolution [25], where interactions between submodules complicate search. In our setting, the effective unit of selection is the chromosome: the landscape is factorisable at the chromosome level but becomes brittle at the gene level.

Figure 5 illustrates this effect. Each network graph shows gene co-occurrence within the top 10% of high-performing chromosomes in the final population (simplified to heuristics only for clarity), with edge weights proportional to co-occurrence frequency. Panel (a) shows that GA alone can distil effective combinations in a small setting (4R v. 7B). In contrast, panel (b) shows that the same process fails in the larger 10R v. 15B case, where the search space overwhelms direct GA optimisation. The full staged pipeline in panel (c) restores structure, producing simplified chromosomes dominated by heuristics #3 Predictive Interception and #8 Threat-Level Assessment, sometimes supported by #2 Circle Target or #4 Rotation Patrol in low-ratio settings. These patterns show how staged evolution isolates a compact set of repeatable heuristics, balancing proactive interception with minimal passive guarding, and sustaining effectiveness even in complex scenarios.

In summary, the staged pipeline succeeds because it balances tractable search with large-scale fidelity. Early stages discover and preserve useful strategic structures, while later stages refine them




Figure 4: Ablation study of defender performance across Red swarm sizes and Blue-to-Red ratios. Columns show Blue/Red ratios and rows show the number of Red attackers. Shading indicates win rate, with lighter colors representing higher performance. For each subfigure, the top row shows the mean win rate of the top 10% of chromosomes and the bottom row the best-performing chromosome. Stage 3 demonstrates the benefit of combining DP allocation with sampling, while Stage 4 illustrates the additional gains from iterative refinement, achieving the highest and most consistent win rates across all scenarios.




Figure 5: Network graphs of heuristic co-occurrence (edge thickness proportional to co-occurrence frequency). GA alone can distill effective structures in small-scale cases (a), but fails to simplify when scaled up (b). Only the full pipeline (c) restores parsimony, converging on a compact set of synergistic heuristics.

under the full complexity of large engagements. Reducing the problem further to the gene level destroy this synergy, confirming that in this MAS defense setting the meaningful building blocks are entire chromosomes, not individual genes. These findings highlight both the limitations of naive factorisation and the promise of staged, chromosome-level approaches for large-scale multi-agent defense.

5 CONCLUSIONS

This study introduces a staged factorization approach for optimizing defender allocation in large-scale multi-agent drone defense scenarios. The key innovation lies in decomposing the problem into manageable stages: small-scale genetic evaluation of heuristics, factor refinement, and large-scale evaluation of combined strategies. This staged approach reduces search complexity, mitigates early evaluation biases, and enables the identification of high-performing defender allocations that generalize across swarm sizes — outcomes that conventional methods struggle to achieve.

Our results highlight that some heuristics are highly effective only in small battles, whereas others provide robust defense in large-scale engagements. The method’s ability to combine these heuristics

into effective large-scale strategies demonstrates the practical value of staged factorization in multi-agent systems. In real-world drone defense contexts, this implies that a principled, factorized evaluation can guide the deployment of autonomous agents to protect high-value assets against diverse adversary strategies.

An additional insight is the explicit mapping between defensive effectiveness and available resources. By evaluating defender allocations across a range of blue-to-red agent ratios, the approach reveals diminishing returns for certain heuristics and identifies allocation strategies that achieve high win rates with minimal resources. This provides actionable guidance for resource-constrained operations, allowing decision-makers to balance effectiveness against the number of deployed agents and optimize overall system efficiency.

A limitation of our approach is the assumption of full observability and deterministic agent dynamics. Real-world environments may feature partial information, communication constraints, and stochastic disturbances, which could affect generalizability of our findings. Additionally, the current heuristics and allocation schemes do not incorporate learning-based adaptation during deployment.

Future work could address these limitations by integrating partially observable decision-making, adaptive heuristics, and reinforcement learning at lower levels. Expanding to dynamic adversary behaviors, environmental obstacles, and more heterogeneous agent capabilities could further enhance the applicability of the approach to operational drone defense scenarios.

ACKNOWLEDGMENTS

G. Douglas has received support through the Australian Government Research Training Program Scholarship. This work has also been supported in part by Saab Australia Pty Ltd through the Defence Trailblazer program. This work was supported with supercomputing resources provided by the Phoenix HPC service at Adelaide University.

REFERENCES

- [1] Joshua Achiam, Harrison Edwards, Dario Amodei, and Pieter Abbeel. 2018. Variational option discovery algorithms. *arXiv preprint arXiv:1807.10299* abs/1807.10299 (2018).
- [2] Darryl Ahner and Carl Parson. 2013. Weapon tradeoff analysis using dynamic programming for a dynamic weapon target assignment problem within a simulation. In *2013 Winter Simulations Conference (WSC)*. IEEE Press, Washington, D.C., 2831–2841. <https://doi.org/10.1109/WSC.2013.6721653>
- [3] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. *JAX: composable transformations of Python+NumPy programs*. Google Research. <http://github.com/google/jax>
- [4] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. 2018. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070* abs/1802.06070 (2018).
- [5] Mohammad Ghavamzadeh, Sridhar Mahadevan, and Rajbala Makar. 2006. Hierarchical multi-agent reinforcement learning. *Autonomous Agents and Multi-Agent Systems* 13, 2 (2006), 197–229. <https://doi.org/10.1007/s10458-006-7035-4>
- [6] David Goldberg, Kalyanmoy Deb, and Jeffrey Horn. 1994. Massive Multimodality, Deception, and Genetic Algorithms. *Parallel Problem Solving from Nature 2* (08 1994).
- [7] David E. Goldberg, Bradley Korb, and Kalyanmoy Deb. 1989. Messy genetic algorithms: Motivation, analysis, and first results. *Complex systems* 3, 5 (1989), 493–530.
- [8] Cai Huaiping, Liu Jingxu, Chen Yingwu, and Wang Hao. 2006. Survey of the research on dynamic weapon-target assignment problem. *Journal of Systems Engineering and Electronics* 17, 3 (2006), 559–565. [https://doi.org/10.1016/S1004-4132\(06\)60097-2](https://doi.org/10.1016/S1004-4132(06)60097-2)
- [9] Rufus Isaacs. 1951. *Games of Pursuit*. RAND Corporation, Santa Monica, CA.
- [10] Rufus Isaacs. 1965. *Differential games: a mathematical theory with applications to warfare and pursuit, control and optimization*. Wiley, New York. Bibliography: p. 379–380.
- [11] Alexander Kline, Darryl Ahner, and Raymond Hill. 2019. The Weapon-Target Assignment Problem. *Computers & Operations Research* 105 (2019), 226–236. <https://doi.org/10.1016/j.cor.2018.10.015>
- [12] Tejas D. Kulkarni, Karthik R. Narasimhan, Ardavan Saeedi, and Joshua B. Tenenbaum. 2016. Hierarchical deep reinforcement learning: integrating temporal abstraction and intrinsic motivation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (Barcelona, Spain) (NIPS'16)*. Curran Associates Inc., Red Hook, NY, USA, 3682–3690.
- [13] Arthur Louette, Pascal Leroy, Yanis Geurts, and Damien Ernst. 2025. Existing Gaps In Reinforcement Learning For Drone Warfare. (2025). Author preprint.
- [14] Michael McKay, Richard Beckman, and William Conover. 1979. A Comparison of Three Methods for Selecting Vales of Input Variables in the Analysis of Output From a Computer Code. *Technometrics* 21 (05 1979), 239–245. <https://doi.org/10.1080/00401706.1979.10489755>
- [15] Zepeng Ning and Lihua Xie. 2024. A survey on multi-agent reinforcement learning and its application. *Journal of Automation and Intelligence* 3, 2 (2024), 73–91. <https://doi.org/10.1016/j.jjai.2024.02.003>
- [16] Liviu Panait and Sean Luke. 2005. Cooperative Multi-Agent Learning: The State of the Art. *Autonomous Agents and Multi-Agent Systems* 11, 3 (2005), 387–434. <https://doi.org/10.1007/s10458-005-2631-2>
- [17] Shubham Pateria, Budhitama Subagdja, Ah-hwee Tan, and Chai Quek. 2022. Hierarchical Reinforcement Learning. *ACM computing surveys*. 54, 5 (2022), 1–35. <https://doi.org/10.1145/3453160>
- [18] Adrian Selmonaj, Szehr Szeher, Giacomo Del Rio, Alessandro Antonucci, Adrian Schneider, and Michael Rügsegger. 2023. Hierarchical Multi-Agent Reinforcement Learning for Air Combat Maneuvering. In *2023 International Conference on Machine Learning and Applications (ICMLA)*. ICMLA, Jacksonville, 1031–1038. <https://doi.org/10.1109/ICMLA58977.2023.00153>
- [19] Onn Shehory and Sarit Kraus. 1998. Methods for task allocation via agent coalition formation. *Artificial Intelligence* 101, 1 (1998), 165–200. [https://doi.org/10.1016/S0004-3702\(98\)00045-9](https://doi.org/10.1016/S0004-3702(98)00045-9)
- [20] Ilya Sobol'. 1967. On the distribution of points in a cube and the approximate evaluation of integrals. *U. S. S. R. Comput. Math. and Math. Phys.* 7, 4 (1967), 86–112. [https://doi.org/10.1016/0041-5553\(67\)90144-9](https://doi.org/10.1016/0041-5553(67)90144-9)
- [21] Laura Gail Strickland. 2022. *Coordinating Team Tactics for Swarm-vs.-Swarm Adversarial Games*. Thesis. School of Interactive Computing, College of Computing.
- [22] Vito Trianni. 2008. *Evolutionary Robotics for Self-Organising Behaviours*. Springer Berlin Heidelberg, Berlin, Heidelberg, 47–59. https://doi.org/10.1007/978-3-540-77612-3_4
- [23] Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. 2017. Feudal networks for hierarchical reinforcement learning. In *International conference on machine learning*. PMLR, Sydney, 3540–3549.
- [24] Baolai Wang, Shengang Li, Xianzhong Gao, and Tao Xie. 2021. UAV Swarm Confrontation Using Hierarchical Multiagent Reinforcement Learning. *International Journal of Aerospace Engineering* 2021 (2021), 1–12. <https://doi.org/10.1155/2021/3360116>
- [25] Richard A. Watson and Jordan B. Pollack. 2005. Modular Interdependency in Complex Dynamical Systems. *Artificial Life* 11, 4 (2005), 445–457. <https://doi.org/10.1162/106454605774270589>
- [26] L. Darrell Whitley. 1991. Fundamental Principles of Deception in Genetic Search. In *Foundations of Genetic Algorithms*, Gregory J.E. Rawlins (Ed.). Vol. 1. Elsevier, San Mateo, 221–241. <https://doi.org/10.1016/B978-0-08-050684-5.50017-3>
- [27] Jiachen Yang, Igor Borovikov, and Hongyuan Zha. 2020. Hierarchical Cooperative Multi-Agent Reinforcement Learning with Skill Discovery. In *International Foundation for Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Auckland, 1566–1574. <https://doi.org/10.5555/3398761.3398941>