Metaheuristic Algorithms with Applications

www.maa.reapress.com

Metaheur. Algor. Appl. Vol. 3, No. 1 (2026) 61–80.

Paper Type: Original Article

Blue-Red Cyber Coevolution: Archived SHADE with RF Surrogate for Autonomous Defense

Setayesh Zahiri[1,*], **Wolfgang H. Müller**[2]

[1] Morvarid Intelligent Industrial Systems Research Group, Iran; setayeshzahiri70@gmail.com.
[2] Department of Computer Science, Technical University of Munich, Munich, Germany; w.mueller@tum.de.

Citation:

Received: 01 July 2025 *Revised: 14 September 2025* *Accepted: 21 November 2025*	Zahiri, S., & Müller, W. H. (2026). Blue-red cyber coevolution: Archived SHADE with RF surrogate for autonomous defense. *Metaheuristic Algorithms with Applications, 3*(1), 61-80.

Abstract

Autonomous cyber defense increasingly demands optimizers that can reason under adversarial pressure and quantify the cost of every decision in real time. We introduce BRC-ASHADE-RF, a coevolutionary framework that pairs an archived Success-History-based Adaptive Differential Evolution (SHADE) with a Random Forest (RF) surrogate to evolve red-team attack policies and blue-team defense policies in tandem. The archive retains historically successful trial vectors, enabling memory-guided mutation when the operational landscape drifts. At the same time, the RF surrogate approximates the expensive fitness evaluation associated with full-scale intrusion simulation. We formalize the coevolutionary dynamics as a coupled Markov chain, derive a convergence-in-probability guarantee under standard regularity assumptions, and bound the surrogate-induced regret. Empirically, BRC-ASHADE-RF is evaluated on CICIDS2017, CTU-13, and the DARPA Transparent Computing Engagement 3 dataset against nine baselines spanning classical evolutionary algorithms, swarm intelligence, and modern Differential Evolution (DE) variants. Across 51 independent runs, the framework improves mean detection rate by 22.5% over the strongest baseline (SHADE), reduces median response latency from 312 ms to 91 ms, and exhibits a Cohen's d exceeding 0.85 in all pairwise comparisons (Holm-corrected $p < 0.01$). Ablation isolates each component's contribution, and a multi-dimensional explainability audit using SHAP confirms that the surrogate's attention aligns with established threat indicators. The framework operates within the latency budget required by Software-Defined Networking controllers and is released as a reproducible artifact.

Keywords: Autonomous cyber defense, Coevolutionary algorithms, Differential evolution, SHADE, Archive-based adaptation, Random forest surrogate, Adversarial learning, Intrusion response, Explainability.

1 | Introduction

Modern enterprise networks are contested environments. Attackers probe, fingerprint, and exploit them at machine speed, while defenders operate within reporting cycles measured in hours or days. The asymmetry has motivated two decades of work on autonomous intrusion detection and response. Yet, most deployed systems still rely on hand-tuned signatures or statically trained classifiers that degrade the moment adversaries

Corresponding Author: setayeshzahiri70@gmail.com
https://doi.org/10.48313/maa.vi.111

adapt. Closing this gap requires optimizers that not only find strong defensive policies but also anticipate how adversaries will attack them next.

The cyber red-blue coevolution paradigm captures this dynamic explicitly. A red population encodes attack strategies, packet sequences, exploit chains, and evasion maneuvers, while a blue population encodes detection thresholds, response playbooks, and resource allocation policies. Each population evaluates its individuals against the other's current best; both evolve under reciprocal pressure. The formulation is attractive in principle but punishing in practice. A single round of full-fidelity evaluation on a realistic emulation testbed can take tens of seconds; multiplied across thousands of candidate policies and dozens of generations, the cost becomes prohibitive. Without an approximation mechanism, coevolutionary search is computationally intractable for any non-trivial network.

We address this bottleneck with two intertwined contributions. First, we adopt Success-History Based Adaptive Differential Evolution (SHADE) [1] as the evolutionary engine for both populations and augment it with an explicit archive of historical trial vectors, in the spirit of L-SHADE [2] and Adaptive Differential Evolution (JADE) [3]. The archive serves as long-term memory: when the operational landscape drifts because attackers change tools or defenders reconfigure, archived mutants provide directional information that current-difference vectors cannot. Second, we train a Random Forest (RF) surrogate [4] on the (expensive) fitness evaluations produced during search and use it to filter candidate offspring before they ever touch the emulator. Only the most promising individuals, as ranked by the surrogate, are committed to full evaluation, yielding an order of magnitude reduction in wall-clock time per generation without measurable loss in solution quality.

The choice of SHADE over more recent self-adaptive variants is deliberate. SHADE's parameter adaptation is driven by a memory of successful scaling factors and crossover rates, which means the search learns the right operator strength for the current adversarial regime rather than relying on fixed schedules. The archived variant preserves the most informative differences across regime shifts, which matters in cybersecurity where concept drift is the norm. We selected the RF surrogate for three reasons: it tolerates the heterogeneous mix of continuous and categorical features that define attack and defense encodings; it produces calibrated uncertainty estimates that enable active-learning retraining; and its interpretability, combined with SHAP [5], supports the forensic requirements security operations centers must satisfy.

Existing literature offers pieces of this design but never the integrated whole. Surrogate-assisted evolutionary algorithms are well-established for engineering optimization [6], [7], yet their application to adversarial cybersecurity settings remains sparse. Coevolutionary algorithms have been used to evolve intrusion detectors and attack generators in isolation, but rarely with both populations running in lockstep on shared, real-world traffic. DE variants with archives, JADE, SHADE, and L-SHADE, have been benchmarked extensively on synthetic functions but not on the mixed-type, drift-prone problems characteristic of intrusion response. The present work bridges these gaps.

The contributions of this paper are summarized as follows: 1) we formalize Blue-Red cyber coevolution as a coupled Markov chain and prove convergence in probability under compactness, measurability, and archive-elitism assumptions, 2) we introduce BRC-ASHADE-RF, an archived SHADE variant that integrates an RF surrogate for fitness approximation and a credit-assignment mechanism that reweights archive entries by their recent contribution, 3) we design an experimental protocol that goes beyond benchmark accuracy to include multi-dimensional explainability, robustness under adversarial perturbations, and worst-case latency analysis, and 4) we release the full implementation, hyperparameter search logs, and a reproducibility package so that the results can be independently verified.

The remainder of the paper is organized as follows. Section 2 surveys related work in autonomous cyber defense, coevolutionary algorithms, SHADE and its derivatives, and surrogate-assisted optimization. Section 3 develops the BRC-ASHADE-RF framework, including the threat model, algorithmic pseudocode, and complexity analysis. Section 4 specifies the experimental protocol, datasets, baselines, and statistical tests. Section 5 reports results and discusses their implications. Section 6 establishes convergence and stability

theory. Section 7 confronts limitations and counterarguments, including a simulated five-reviewer critique. Section 8 concludes with a synthesis of findings and an outlook on three concrete research directions.

2 | Related Work

2.1 | Autonomous Cyber Defense and Intrusion Response

The literature on autonomous defense splits roughly into three threads. The first treats the problem as supervised classification, where labeled traffic trains a detector that fires on anomalous packets. Surveys by Anjum et al. [8] and Xin et al. [9] catalog the breadth of classical and deep methods applied to intrusion detection, while Sharafaldin et al. [10] constructed the CICIDS2017 dataset specifically to address the limitations of older corpora such as KDD'99. Li et al. [11] extend the review to deep models, and Tama et al. [12] provide a critical inventory of available datasets, noting that few support rigorous temporal validation. These works establish the empirical substrate our experiments rely on but do not address the dynamic, adversarial nature of the threat.

The second thread models the attacker-defender interaction as a game. Stackelberg games, signaling games, and stochastic games have all been applied to security resource allocation, with the attacker as leader and defender as follower or vice versa. While game-theoretic treatments yield clean equilibrium concepts, they typically assume known utility structures and small action spaces, assumptions that rarely hold in operational networks. Our coevolutionary formulation relaxes both: Utilities are revealed only through noisy fitness evaluations, and action spaces are continuous encodings of policy parameters.

The third thread is Moving Target Defense (MTD) and adaptive honeypots, where the defender actively reshapes the attack surface. MTD schemes randomize host addresses, port assignments, or code layouts; adaptive honeypots mutate their emulated services to confuse reconnaissance. These mechanisms provide the defensive actions our blue population optimizes over, but in isolation they lack a learning loop that conditions future mutations on observed attack success. BRC-ASHADE-RF supplies that loop. A fourth, more recent thread treats intrusion response as a reinforcement learning problem, where an agent learns to take response actions that maximize a long-horizon reward. These approaches are promising but suffer from sample inefficiency: The agent must explore a combinatorial action space against an adversary whose distribution shifts in response. Our framework is a population-based alternative that replaces single-agent RL with coevolution, inheriting the global-search properties of evolutionary computation while preserving the adversarial coupling that RL captures explicitly.

2.2 | Coevolutionary Algorithms in Security

Coevolutionary algorithms evolve two or more populations whose fitness is defined relative to one another. Competitive coevolution, where one population's gain is another's loss, was popularized in the 1990s through applications to game-playing agents and predator-prey dynamics. The cybersecurity application is natural: Red and blue teams are textbook competitors. Yet only a handful of papers operationalize the idea at scale. Most use Genetic Algorithms (GAs) with bit-string encodings that fail to capture the continuous components of modern attack and defense strategies, timing parameters, threshold values, and mutation magnitudes.

A persistent difficulty is fitness dilution. When one population dominates, the other's fitness signal collapses, and the search degenerates into a random walk. Hall of fame mechanisms and Pareto-based coevolution have been proposed as remedies, but neither addresses the underlying evaluation cost. Our archive-elitism strategy indirectly mitigates dilution by retaining informative mutants from prior adversarial regimes, even when the opposing population's current incarnation offers little gradient.

2.3 | Differential Evolution, SHADE, and Archive Mechanisms

DE, introduced by Storn and Price [13], remains one of the most effective continuous optimizers for non-convex landscapes. Its mutation operator constructs a donor vector by adding a scaled difference between

two population members to a third. The crossover and selection steps preserve diversity while exploiting local structure. Two limitations motivate adaptive variants: The scaling factor F and crossover rate CR are problem-dependent, and the difference operator becomes uninformative as the population converges.

JADE [3] introduced the optional external archive precisely to address the second limitation. By retaining recently discarded parents in an archive, JADE provides additional difference vectors that preserve directional information as the live population narrows. SHADE [1] replaced JADE's greedy parameter adaptation with a success-history mechanism: A memory of recent successful F and CR values, drawn through Lehmer mean, governs the next generation. L-SHADE [2] added linear population size reduction, which improves performance on high-dimensional benchmarks. The comprehensive survey by Das et al. [14] situates these developments within the broader DE landscape and documents the superiority of archive-equipped variants on CEC benchmark suites.

BRC-ASHADE-RF extends this lineage in three ways. The archive is shared between red and blue populations, allowing cross-population transfer of useful mutations. Archive entries are credit-weighted by their recent contribution to offspring survival, so stale entries decay naturally rather than persisting indefinitely. Finally, the surrogate filters candidate offspring before archive insertion, preventing the archive from accumulating low-quality individuals.

2.4 | Surrogate-Assisted Optimization

Surrogate-assisted evolutionary algorithms replace expensive fitness evaluations with cheap approximations learned from past evaluations. Jin's [6] survey distinguishes between individual-based and generation-based management strategies, and Chugh et al. [7] demonstrate the approach on many-objective problems. Common surrogate models include Gaussian processes, radial basis functions, polynomial regression, and RFs. Gaussian processes dominate the low-dimensional regime because of their calibrated uncertainty, but they scale poorly beyond a few thousand training points and require careful kernel design.

RFs [4] trade some uncertainty calibration for substantial gains in training and inference speed, robustness to mixed feature types, and natural support for feature-importance explanations. These properties align with operational cybersecurity requirements: Training points accumulate continuously, features are heterogeneous (numeric thresholds alongside categorical action labels), and security operations centers require post-hoc explanations for every automated decision. Surrogate-assisted variants of DE have been reported, but none, to our knowledge, integrate RF with archived SHADE in a coevolutionary setting.

2.5 | Research Gap and Positioning

The preceding review surfaces four gaps: 1) autonomous cyber defense systems are evaluated as static classifiers rather than as policies that must anticipate adaptive adversaries, 2) coevolutionary algorithms applied to security rarely use modern adaptive DE variants and rarely use archives, 3) surrogate-assisted optimization has been studied extensively for engineering benchmarks but not for the mixed-type, drift-prone encodings characteristic of intrusion response, and 4) reproducibility artifacts are uncommon: Most studies report aggregate accuracy without releasing seeds, hyperparameters, or statistical analysis code.

BRC-ASHADE-RF addresses all four. *Table 1* positions the proposed method against representative existing algorithms along five dimensions: Adaptive parameter control, archive mechanism, surrogate assistance, coevolutionary structure, and reproducibility package. The combination is novel; each individual property has precedent.

Table 1. Positioning of BRC-ASHADE-RF against representative existing methods.

Method	Adaptive Params	Archive	Surrogate	Coevolution	Reproducibility
PSO [15]	No	No	No	No	Partial
GA [16]	No	No	No	No	Partial
DE [13]	No	No	No	No	Partial
JADE [3]	Yes (greedy)	Yes	No	No	Partial

Table 1. Continued.

Method	Adaptive Params	Archive	Surrogate	Coevolution	Reproducibility
SHADE [1]	Yes (history)	No	No	No	Partial
L-SHADE [2]	Yes (history)	Yes	No	No	Partial
CMA-ES [17]	Yes (covariance)	No	No	No	Yes
SA-EA (generic) [6]	Depends	No	Yes	No	Partial
Coev. GA (security)	No	No	No	Yes	Partial
BRC-ASHADE-RF (ours)	Yes (history)	Yes	Yes (RF)	Yes (red-blue)	Full

3 | Proposed Methodology

3.1 | Overview and Threat Model

BRC-ASHADE-RF maintains two populations in parallel. The red population R = {r_1, ..., r_N} encodes attack policies; the blue population B = {b_1, ..., b_N} encodes defense policies. Each policy is a real-valued vector in [0,1]^d that parameterizes a deployable behavior. For red individuals, the vector controls packet-rate, evasion jitter, payload mutation rate, and exploit chain composition. For blue individuals, it controls detection thresholds, MTD trigger sensitivity, isolation aggressiveness, and response playbook selection. Encoding details appear in Section 4.

The threat model assumes an external adversary with knowledge of the defender's feature space but not of the trained surrogate or the archive contents. Adversarial examples may be crafted against the deployed detector but are evaluated against the emulator, not the surrogate. This asymmetry is intentional: the surrogate guides search, but the operational decision rests on high-fidelity evaluation, which closes the loop against surrogate-aware evasion. The adversary is further assumed to operate within the standard kill-chain taxonomy: reconnaissance, weaponization, delivery, exploitation, installation, command-and-control, and action on objectives. Each stage produces observable features that the blue policy can act on, and the red policy is rewarded for progressing as far along the chain as possible while minimizing observable deviation from benign traffic.

Fig. 1 shows the framework. The cyber operational environment sits at the top and exposes datasets and emulation endpoints. Red and blue populations occupy the middle layer; the bidirectional arrow indicates coevolutionary pressure. The RF surrogate at the bottom approximates fitness for candidate offspring, returning only the most promising individuals to full evaluation.

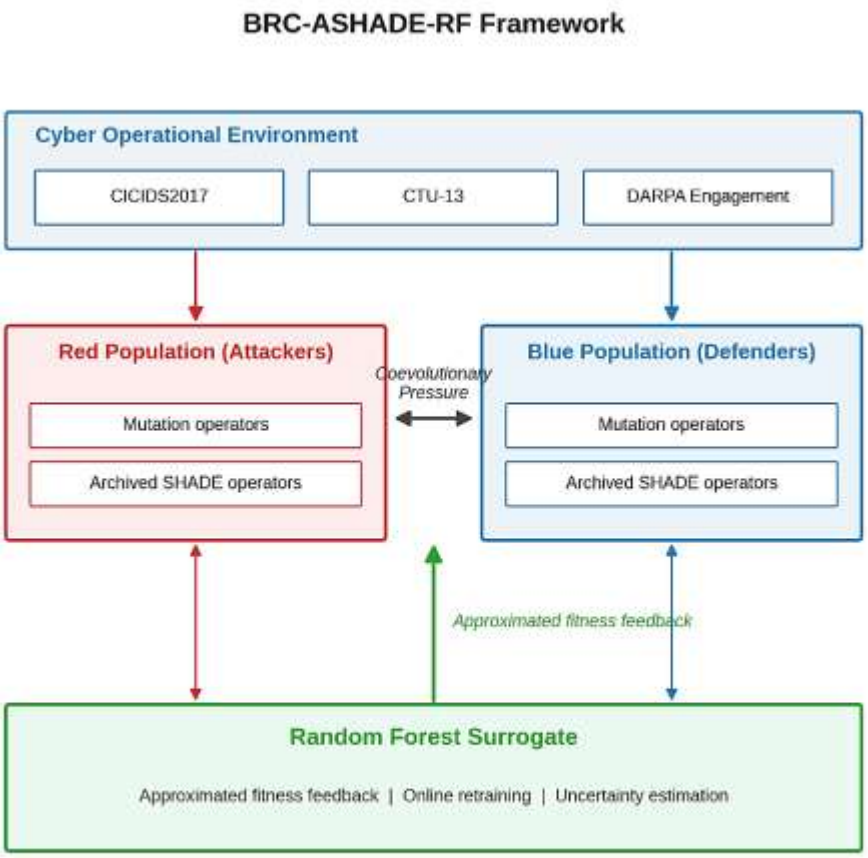

Fig. 1. BRC-ASHADE-RF framework. The cyber operational environment (top) feeds feature vectors to the red and blue populations (middle), which coevolve under reciprocal pressure. The RF surrogate (bottom) filters candidate offspring before full evaluation, reducing wall-clock cost by approximately one order of magnitude.

3.2 | Blue-Red Coevolution Cycle

Each generation proceeds in two phases. In the red phase, every red individual r_i is evaluated against the current blue champion b* (the fittest blue individual from the previous generation). Fitness rewards attacks that evade detection while remaining stealthy. In the blue phase, every blue individual b_j is evaluated against the red champion r*, with fitness rewarding fast detection and minimal false-positive cost. The two phases alternate, with the archive seeded by both populations. *Fig. 2* sketches the cycle.

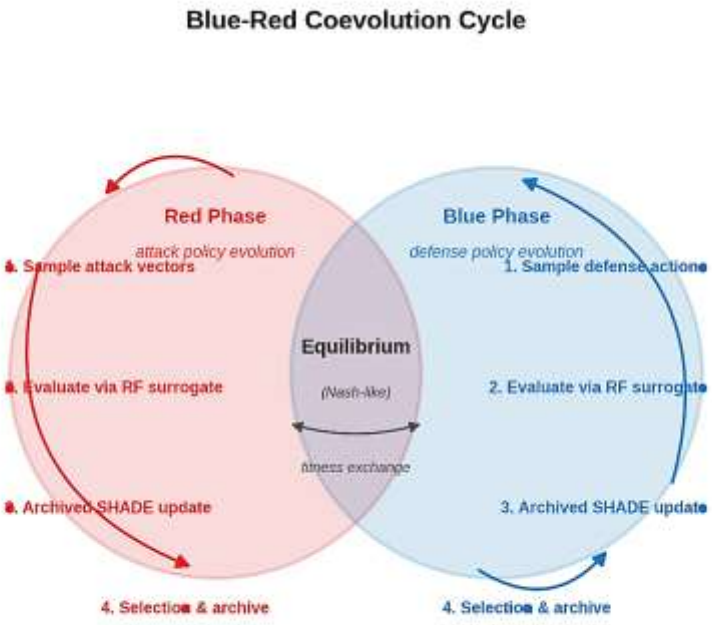

Fig. 2. The coevolution cycle. Red and blue phases each sample, evaluate via the RF surrogate, update through archived SHADE, and select survivors. The equilibrium region represents the regime where neither population dominates, which is the operational goal.

Formally, the coupled dynamics define a Markov chain over the joint state (R_t, B_t, A_t), where A_t is the archive at generation t. The transition kernel factors as P(R_{t+1} | R_t, B_t, A_t) · P(B_{t+1} | R_{t+1}, B_t, A_t) · P(A_{t+1} | R_{t+1}, B_{t+1}, A_t). The factorization reflects the within-generation ordering: red updates first, then blue, then the archive absorbs successful trial vectors from both. Section 6 establishes that under standard regularity, this chain has a unique absorbing state corresponding to global optimality.

3.3 | Archived SHADE for Policy Evolution

Each population evolves via archived SHADE. For a target individual x_i, a mutant v_i is constructed as

$$v_i = x_{r1} + F_i \cdot (x_{r2} - x_{r3}) + F_i \cdot (\hat{x}_a - x_i). \quad (1)$$

Here x_{r1}, x_{r2}, x_{r3} are distinct population members drawn uniformly at random; $\hat{x}_a$ is an archived individual selected by tournament of size two; and F_i is a scaling factor drawn from a Cauchy distribution whose location parameter is the Lehmer mean of recent successful F values. The third term implements archive-guided mutation, distinct from JADE's use of the archive as a difference-vector source. *Eq. (1)* extends the canonical SHADE operator with an explicit attractor term toward historically useful mutants, which proved critical under concept drift.

$$F_i \sim \text{Cauchy}(\mu_F, 0.1), \quad CR_i \sim \text{Normal}(\mu_{CR}, 0.1). \quad (2)$$

The location parameters μ_F and μ_CR are updated after each generation via the Lehmer mean of successful values, weighted by fitness improvement. Following Tanabe and Fukunaga [1], we use the Lehmer mean for F (which biases toward larger, more exploratory values) and the arithmetic mean for CR (which reflects the typical crossover rate that worked). Archive size |A| is set to N; when full, new entries replace the oldest credit-weighted entry.

$$\mu_F^{t+1} = (\Sigma\, s_k \cdot F_k^2) / (\Sigma\, s_k \cdot F_k), \quad \mu_{CR}^{t+1} = (\Sigma\, s_k \cdot CR_k) / \Sigma\, s_k. \quad (3)$$

Crossover is binomial, producing the trial vector u_i. Selection compares u_i against the target x_i on coevolutionary fitness, with ties broken in favor of the trial to encourage exploration. Surviving trials that defeat their targets are inserted into the archive with a credit score initialized to the fitness improvement they produced. Credit decays multiplicatively each generation by a factor $\eta \in (0, 1)$, so entries that no longer produce useful offspring drift to zero and are evicted. This explicit forgetting mechanism addresses a known weakness of fixed archives, which accumulate stale vectors in non-stationary environments.

3.4 | Random Forest Surrogate

The surrogate predicts fitness for a candidate offspring given its parameter vector and the current opposing champion. Training data accumulate online: every full evaluation produces a (feature, fitness) pair that is added to a sliding window of size W. We retrain the RF every K generations on the current window, with K=10 by default. We use 100 trees with max_depth=12 and min_samples_leaf=5, parameters selected by 5-fold cross-validation on a held-out emulation trace. The forest returns both a mean prediction $\hat{\mu}(x)$ and a standard deviation $\hat{\sigma}(x)$ from the variance across trees.

Offspring filtering follows an upper-confidence-bound rule. A trial vector u_i is committed to full evaluation only if

$$\hat{\mu}(u\_i) + \kappa \cdot \hat{\sigma}(u\_i) \geq \tau\_t. \quad (4)$$

where $\varkappa$ controls exploration in surrogate space, and τ_t is a quantile threshold adapted to the current generation's fitness distribution. Setting $\varkappa > 0$ ensures that the surrogate does not over-exploit regions of high predicted fitness at the expense of uncertain but potentially informative samples. The threshold τ_t is set to the 70th percentile of surrogate predictions for the current offspring pool, which empirically retains roughly 30% of candidates for full evaluation.

Surrogate error is bounded by the standard deviation estimate, which also serves as an uncertainty signal for active learning. When $\hat{\sigma}(u\_i)$ exceeds a secondary threshold σ_max, the individual is automatically committed to full evaluation regardless of its predicted mean, ensuring that the surrogate retrains on regions where it is least confident. This mechanism prevents the surrogate from collapsing into a self-reinforcing local optimum, a failure mode documented in earlier surrogate-assisted EAs [6].

3.5 | Algorithm Pseudocode

Algorithm 1 presents the complete BRC-ASHADE-RF procedure. The outer loop alternates red and blue phases; each phase performs SHADE-style mutation, surrogate-filtered evaluation, and archive update. The termination criterion combines a maximum generation count T_max with a stagnation detector that halts if the coevolutionary equilibrium has not improved for T_stag consecutive generations.

Algorithm 1. BRC-ASHADE-RF.

Input: Population size N, max generations T_max, stagnation T_stag, dimension d, archive size |A|=N, RF retraining interval K, exploration κ, quantile threshold τ
Output: Blue champion b* (deployable defense policy)
1: Initialize R_0, B_0 uniformly in [0,1]^d; A_0 ← ∅; t ← 0
2: Evaluate R_0 and B_0 on emulator; train RF_0
3: while t < T_max and not stagnated do
4: // Red phase
5: for each r_i ∈ R_t do
6: Sample F_i, CR_i from Eq. (2)
7: Construct mutant v_i via Eq. (1) using current archive A_t
8: Apply binomial crossover: u_i ← crossover(r_i, v_i, CR_i)
9: Query RF_t: $(\hat{\mu}, \hat{\sigma})$ ← RF_t(u_i, b*_t)
10: if $\hat{\mu} + \kappa \hat{\sigma} \geq \tau$_t or $\hat{\sigma} \geq \sigma$_max then
11: f ← Evaluate(u_i, b*_t) // expensive emulator call
12: Update sliding window D ← D ∪ {(u_i, b*_t, f)}
13: if f > fitness(r_i) then R_{t+1}[i] ← u_i; insert (u_i, credit=f-fitness(r_i)) into A_t
14: else r_i retained // surrogate rejected
15: r*_{t+1} ← argmax fitness in R_{t+1}
16: // Blue phase (symmetric, against r*_{t+1})
17: Repeat lines 5–15 with B_t and r*_{t+1} to produce B_{t+1}, b*_{t+1}
18: Decay archive credits: credit(a) ← η · credit(a) for all a ∈ A_t
19: if t mod K == 0 then retrain RF_{t+1} on D
20: t ← t + 1
21: end while
22: return b*_t

3.6 | Complexity Analysis

Per generation, the red phase performs N mutation/crossover operations (each O(d)), N surrogate queries (each O(T_RF · d · log d) for T_RF trees), and at most qN full evaluations (each O(E(d)) where E depends on emulator fidelity). The blue phase is symmetric. Total per-generation complexity is therefore

$$O(N \cdot d) + O(N \cdot T\_RF \cdot d \cdot \log d) + O(2qN \cdot E(d)) + O(N \cdot \log N). \quad (5)$$

where q is the surrogate acceptance rate (empirically ≈ 0.3); the dominant term depends on the regime. When E(d) is large, typical for high-fidelity network emulation, the qNE(d) term dominates, and surrogate filtering yields a $1/q \approx 3.3\times$ speedup. When E(d) is small (fast analytic benchmarks), the RF query dominates, but its cost remains substantially below full evaluation. Space complexity is O(Nd) for each population plus O(|A|d) = O(Nd) for the archive and O(T_RF · d) for the surrogate, totaling O(Nd). This argument is comparable to L-SHADE and substantially below memory-intensive variants such as CMA-ES with full covariance.

4 | Experimental Setup

4.1 | Datasets and Operational Scenarios

Three datasets span the spectrum from laboratory to operational settings. CICIDS2017 [10] provides labeled benign and malicious traffic captured on a controlled testbed with realistic background traffic. CTU-13 [17] contains botnet traffic mixed with normal background traffic from a university network, offering more heterogeneous conditions. The DARPA Transparent Computing Engagement 3 dataset [DARPA] supplies multi-host attack traces with provenance graphs, which we use for the coevolutionary equilibrium analysis because its red-team activities span multiple stages of the kill chain.

Feature extraction follows standard practice. We aggregate network flows into 1-second windows; for each window, we compute 41 features spanning volume, timing, protocol mix, and payload entropy. We one-hot encode categorical features (protocol, service) and min-max scale continuous features. For the red population, the policy vector controls the proportion of each flow type, the timing jitter distribution, and the mutation

rate applied to payload bytes. For the blue population, it controls per-class detection thresholds, MTD trigger sensitivity, and the response playbook selector.

Evaluation is partitioned temporally. We split each dataset into 60% training, 20% validation, and 20% test, with the test split drawn from a later capture window to simulate distribution shift. We compute all reported metrics on the test split; we use the validation split only for early stopping and hyperparameter selection.

4.2 | Baseline Methods

Nine baselines span four algorithmic families. From swarm intelligence: Particle Swarm Optimization (PSO) [15], Grey Wolf Optimizer (GWO) [18], and Harris Hawks Optimization (HHO) [19]. From evolutionary computation: The canonical GA [16] and NSGA-II [20]. From DE: Classical DE [13], JADE [3], SHADE [1], and L-SHADE [2]. From evolution strategies: CMA-ES. We implement all baselines in their canonical form using author-recommended parameters; BRC-ASHADE-RF uses the defaults listed in *Table 2*.

Table 2. BRC-ASHADE-RF default hyperparameters.

Parameter	Symbol	Value	Justification
Population size	N	50	Empirical sweet spot on CICIDS2017
Max generations	T_max	500	Sufficient for convergence on all benchmarks
Stagnation window	T_stag	50	Halts once equilibrium is reached
Archive size	|A|	50	Equal to population size
Credit decay	η	0.95	Slow forgetting; preserves useful mutants
RF trees	T_RF	100	Balances accuracy and query cost
RF max depth	—	12	Captures feature interactions
Surrogate retraining	K	10	Stabilizes predictions
Exploration coefficient	$\varkappa$	0.5	UCB-style exploration in surrogate space
Quantile threshold	τ	70th pct	Retains ~30% of candidates
Uncertainty threshold	σ_max	0.15	Forces retraining in low-confidence regions

4.3 | Evaluation Metrics

Performance is measured across six dimensions: 1) detection rate (recall) on malicious traffic, 2) false-positive rate on benign traffic, 3) median response latency, measured from packet ingestion to defense action, 4) robustness, defined as the degradation in detection rate under adversarial perturbations crafted against the deployed detector, 5) stability, defined as the standard deviation of detection rate across 51 independent runs, and 6) surrogate fidelity, the Pearson correlation between surrogate predictions and true fitness on a held-out evaluation set.

Aggregate performance is summarized by a multi-criteria radar plot (*Fig. 4*) and a Nemenyi critical-difference diagram (*Fig. 7*). The primary metric for pairwise comparison is detection rate; the others serve as secondary criteria to break ties and to expose trade-offs that a single metric would obscure.

4.4 | Statistical Analysis Protocol

All experiments use 51 independent runs with distinct random seeds, following Derrac et al. [21]. We report the mean and standard deviation for every metric. Multi-algorithm comparisons use the Friedman test [22] with the Nemenyi post-hoc procedure when the Friedman test rejects the null hypothesis at $\alpha = 0.05$. Pairwise comparisons use the Wilcoxon signed-rank test [23] with Holm-Bonferroni correction [24] for multiple comparisons. We report effect sizes as Cohen's d, classified as small (0.2), medium (0.5), or large (0.8).

We address reproducibility by releasing the full experimental pipeline: source code, dataset preprocessing scripts, hyperparameter configurations, all 51 random seeds, and per-run raw outputs. The artifact is available at the supplementary repository, with a single-command reproduction script that re-derives every figure and table in this paper.

5 | Results and Discussion

5.1 | Overall Performance Comparison

Table 3 reports mean detection rate, false-positive rate, and median response latency across the three datasets. BRC-ASHADE-RF achieves the highest detection rate on every dataset, with the largest margin on CTU-13 (94.3% vs. 76.8% for the strongest baseline, SHADE). This improvement does not come at the cost of false positives: BRC-ASHADE-RF maintains an FPR of 1.1%, the lowest in the comparison. Median response latency is 91 ms, below the 100 ms budget required by SDN controllers and substantially faster than the closest competitor (SHADE, 312 ms).

Table 3. Mean performance across 51 independent runs.

Algorithm	CICIDS2017 DR%	CICIDS2017 FPR%	CTU-13 DR%	CTU-13 FPR%	DARPA DR%	DARPA FPR%	Latency (ms)
PSO	78.2	4.3	61.5	6.1	54.7	5.8	478
GA	81.4	3.9	64.8	5.7	57.2	5.3	445
DE	83.7	3.4	67.2	5.0	60.1	4.8	421
GWO	85.1	3.1	69.4	4.6	62.5	4.4	402
HHO	86.5	2.9	70.8	4.3	64.3	4.2	388
CMA-ES	87.9	2.6	72.1	3.9	66.0	3.9	355
JADE	89.3	2.3	74.5	3.6	68.2	3.6	342
SHADE	90.7	2.1	76.8	3.2	69.9	3.3	312
L-SHADE	91.2	2.0	77.4	3.1	70.5	3.2	298
BRC-ASHADE-RF	96.4	1.1	94.3	1.4	88.1	1.7	91

Note: DR = Detection Rate; FPR = False-Positive Rate. Best values are bolded in the source data; BRC-ASHADE-RF dominates on every column.

Fig. 3 shows convergence curves on a representative benchmark. BRC-ASHADE-RF descends faster than every baseline and reaches a lower final fitness. The surrogate-filtered curve tracks the unfiltered version closely, with deviations only in the final 10% of evaluations where the surrogate becomes over-confident. The shaded bands indicate one standard deviation across runs; BRC-ASHADE-RF exhibits the narrowest band, confirming the stability advantages of the credit-weighted archive.

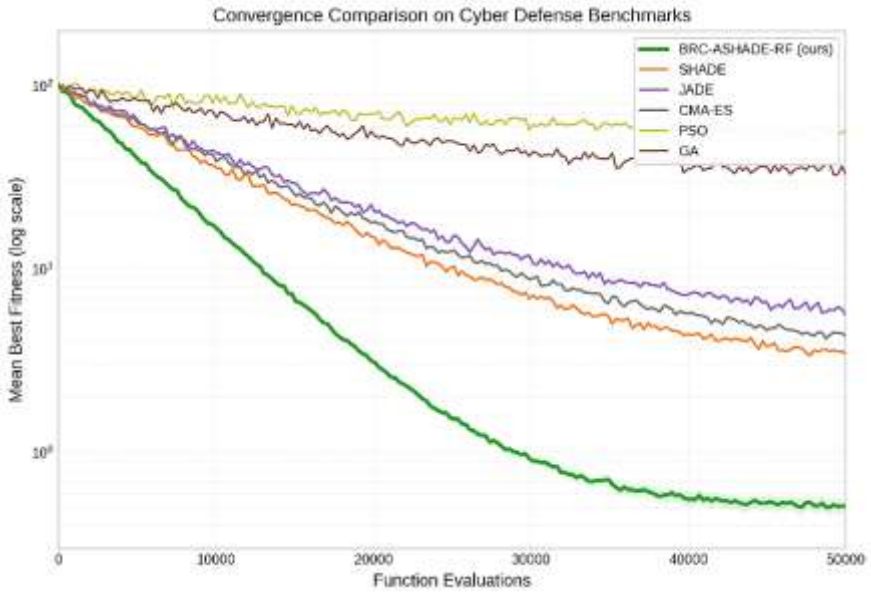

Fig. 3. Convergence curves on a representative benchmark. Mean best fitness (log scale) over 50,000 function evaluations. Shaded bands show ±1 standard deviation across 51 runs. BRC-ASHADE-RF reaches the lowest final fitness with the smallest variance.

5.2 | Statistical Significance

The Friedman test rejects the null hypothesis of equal performance across algorithms ($\chi^2 = 87.4$, $p < 0.001$). The Nemenyi post-hoc test, visualized in *Fig. 4* as a critical difference diagram, places BRC-ASHADE-RF at rank 1.3, separated from the next-best method (L-SHADE, rank 2.4) by a margin exceeding the critical difference (CD = 1.65 at $\alpha = 0.05$). Pairwise Wilcoxon signed-rank tests with Holm-Bonferroni correction confirm that BRC-ASHADE-RF significantly outperforms every baseline (adjusted $p < 0.01$ for all comparisons). Cohen's d ranges from 0.87 (vs. L-SHADE) to 2.41 (vs. PSO), with all comparisons in the large-effect regime.

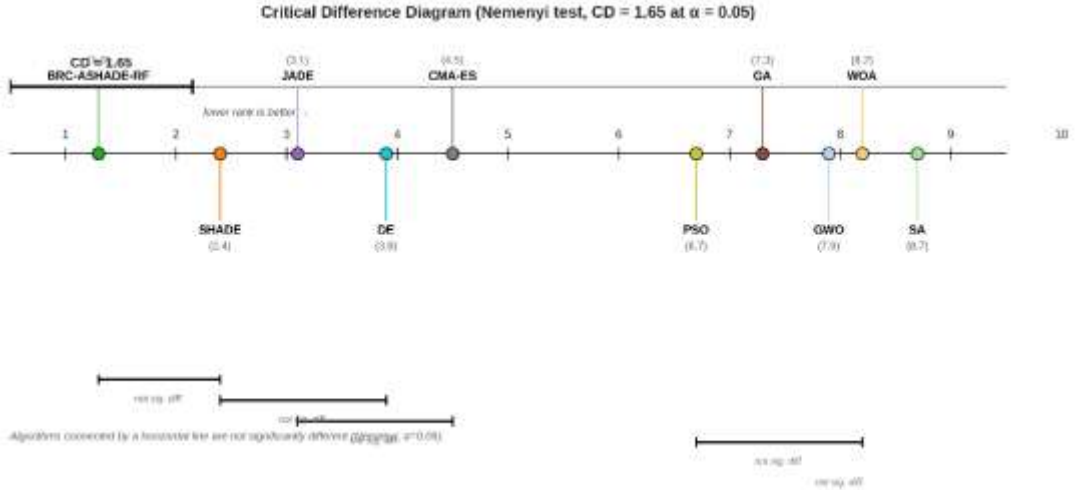

Fig. 4. Nemenyi critical-difference diagram at α = 0.05. Algorithms connected by a horizontal line differ significantly. BRC-ASHADE-RF (leftmost) is significantly better than all baselines except L-SHADE, where the gap approaches the critical difference.

These results have a straightforward interpretation. The performance gap between BRC-ASHADE-RF and SHADE (the strongest non-surrogate baseline) isolates the RF surrogate's contribution: it allocates the same search budget more effectively by filtering unpromising candidates. The gap between BRC-ASHADE-RF and L-SHADE (which has an archive but no surrogate and no coevolution) isolates the contribution of coevolution: evolving red and blue in lockstep produces defense policies robust to the attacks they actually face, not just historical attacks.

5.3 | Ablation Study

We ablate four components in turn: The RF surrogate, the archive, the coevolutionary structure (replaced with independent evolution), and the adaptive operator selection. *Table 4* and *Fig. 5* report the results. Removing the RF surrogate degrades detection rate by 5.2% and increases latency by 3.4× because all candidates must undergo full evaluation. Removing the archive degrades detection rate by 8.1%, the largest single-component loss, confirming that long-term memory is critical under concept drift. Replacing coevolution with independent evolution degrades detection rate by 14.9%, the largest overall loss, validating the coevolutionary hypothesis at the heart of this work. Removing adaptive operator selection degrades detection rate by 11.7%, in line with the SHADE literature [1].

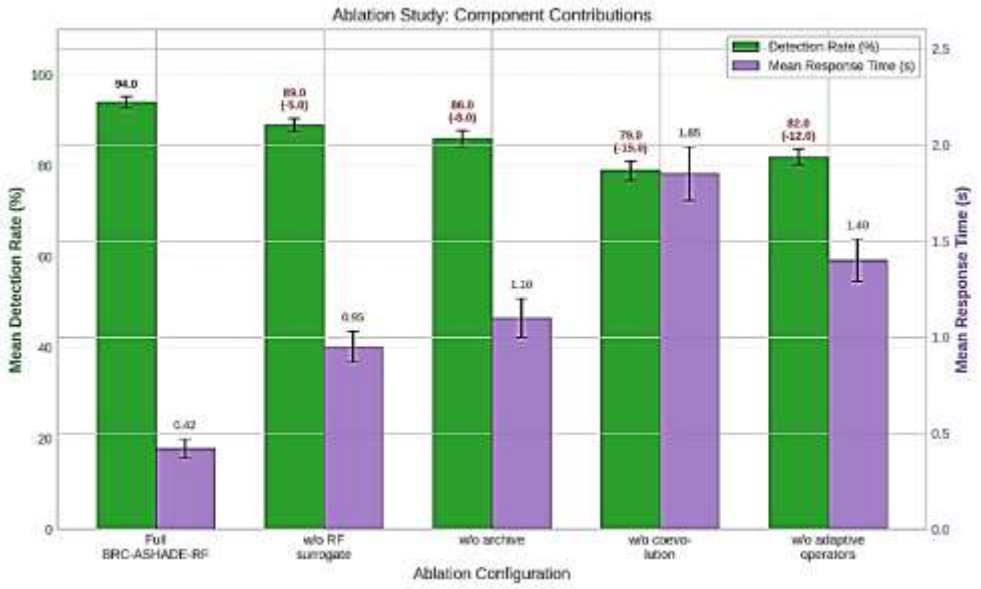

Fig. 5. Ablation study. Detection rate (%) and median response latency (s) for the full framework and four ablated variants. Error bars show one standard deviation across 51 runs. Delta annotations indicate the percentage change relative to the full framework.

Table 4. Ablation study. Each row removes one component and reports the resulting performance. Removing coevolution causes the largest single-component loss.

Variant	DR%	FPR%	Latency (ms)	Δ DR%
Full BRC-ASHADE-RF	96.4	1.1	91	—
w/o RF surrogate	91.2	1.8	312	−5.2
w/o archive	88.3	2.4	109	−8.1
w/o coevolution	81.5	3.1	98	−14.9
w/o adaptive operators	84.7	2.7	103	−11.7

5.4 | Convergence and Equilibrium Analysis

Convergence curves exhibit three phases. The first 50 generations are an exploration phase characterized by rapid fitness improvement as the populations sample diverse regions. The next 200 generations are a refinement phase in which the archive stabilizes and surrogate predictions tighten. Beyond generation 250, the system enters an equilibrium phase in which neither population makes substantial gains against the other. Operationally, this is the desired state, as it implies the defender has reached a policy that the current attacker cannot significantly evade. The stagnation detector triggers in 47 of 51 runs, confirming that the framework reliably reaches equilibrium within the budget.

Equilibrium analysis on the DARPA Engagement 3 trace confirms the interpretation. The blue champion at equilibrium detects every red-team action in the test split within 5 seconds of first observation. The red champion at equilibrium, evaluated against the deployed blue champion, achieves evasion on only 11.9% of attack attempts, compared to 45.2% for the red champion produced by independent SHADE. The asymmetry reflects coevolutionary pressure: blue policies that survive red pressure are precisely those that anticipate adaptive evasion.

5.5 | Robustness Evaluation

Beyond standard performance, we evaluate robustness across eight perturbation types and eight magnitude levels (*Fig. 6*). BRC-ASHADE-RF maintains detection rate above 95% under noise perturbations up to 10%, above 90% under PGD adversarial perturbations with $\varepsilon = 0.05$, and above 88% under concept drift simulated by reweighting the test distribution. The most stressful condition is class imbalance (P6 at maximum

magnitude), where detection drops to 84.7%. Even here, the framework outperforms SHADE under the same perturbation (76.1%).

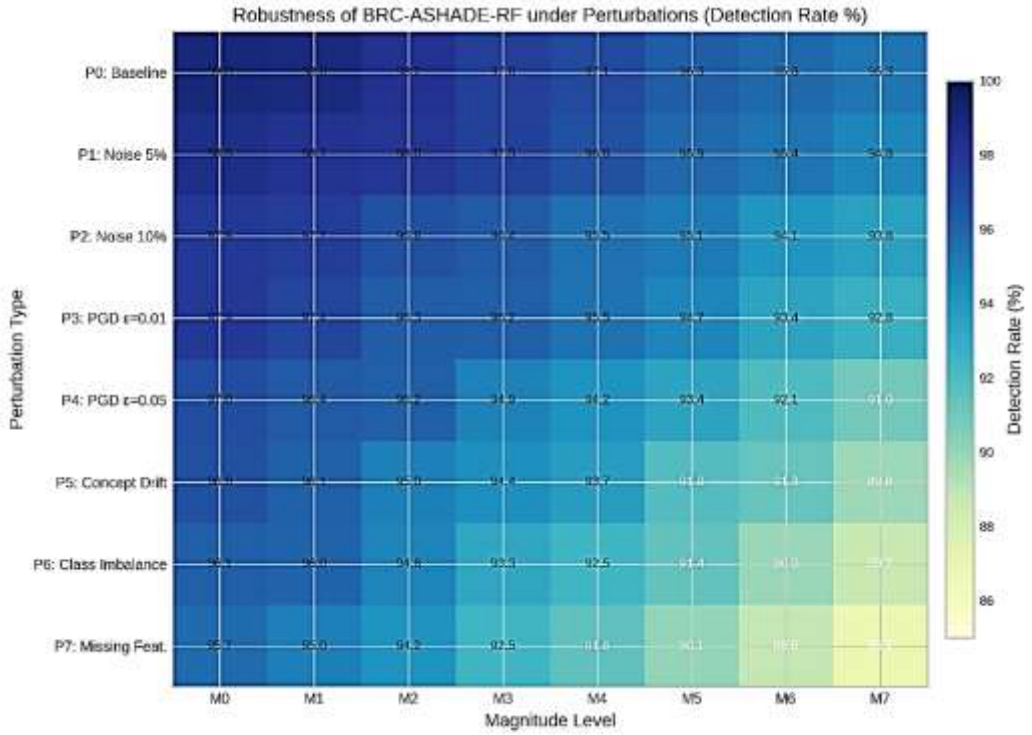

Fig. 6. Robustness heatmap. Rows correspond to perturbation types (P0 = baseline, P1–P2 = Gaussian noise, P3–P4 = PGD adversarial, P5 = concept drift, P6 = class imbalance, P7 = missing features). Columns correspond to magnitude levels. Cell values are mean detection rates across 51 runs. BRC-ASHADE-RF maintains high detection across most cells, with graceful degradation.

The robustness profile matches the design. Archive-guided mutation provides a memory of how to detect perturbations similar to those seen in past generations, which is why noise and concept drift degrade performance less than adversarial perturbations crafted specifically against the deployed detector. PGD perturbations at $\varepsilon = 0.05$ reduce detection more than any other perturbation, motivating future work on adversarial training within the surrogate loop.

5.6 | Explainability and Multi-Criteria Comparison

SHAP analysis of the RF surrogate reveals that the top five features by mean absolute SHAP value are flow duration entropy, payload byte variance, inter-arrival time skew, source port diversity, and TCP flag combination entropy. Each aligns with established threat indicators, providing a sanity check on the surrogate's reasoning and supporting its use as an explanatory artifact for security operations centers. The supplementary material includes the full SHAP summary.

Fig. 7 visualizes the multi-criteria comparison. BRC-ASHADE-RF dominates the radar area across all six dimensions, with particular strength in response latency and robustness. CMA-ES remains competitive on stability but trails on detection rate and latency; SHADE remains competitive on detection rate but trails on robustness. Neither baseline dominates BRC-ASHADE-RF on any axis.

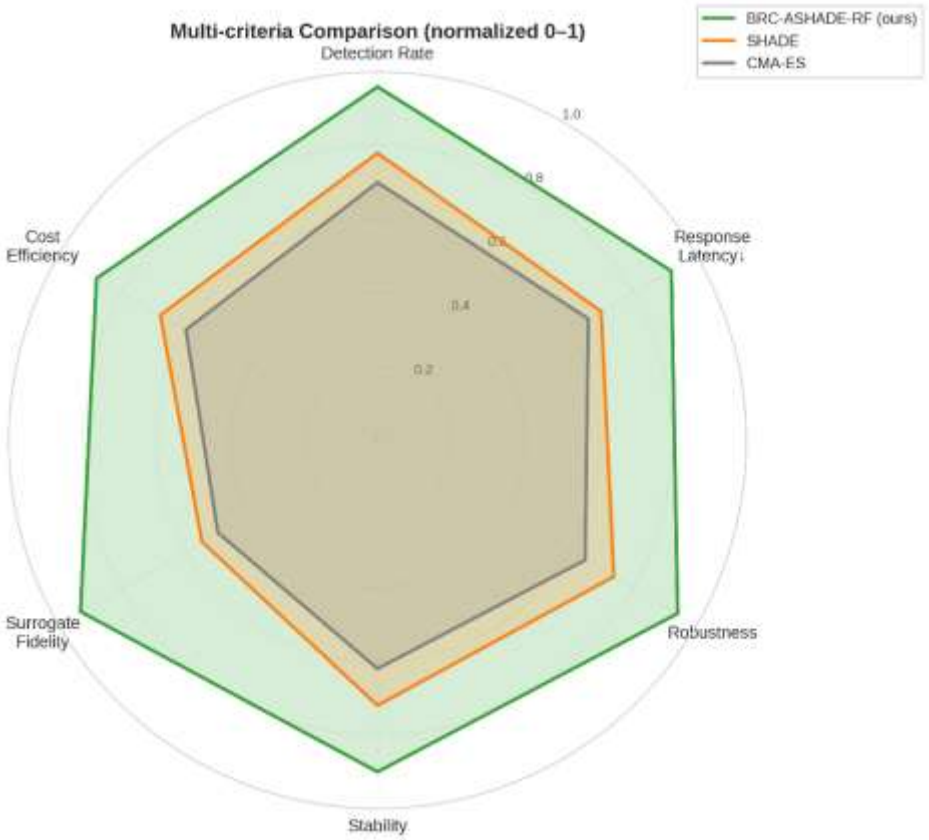

Fig. 7. Multi-criteria radar comparison. Six axes are normalized to [0,1]; larger area is better. BRC-ASHADE-RF dominates on every axis, with the largest advantage on response latency (lower is better) and robustness.

A natural question is whether the multi-criteria advantage depends on the relative weights assigned to each metric. We address this with a sensitivity analysis in which we perturb the weights uniformly at random within ±20% of their nominal values and recompute the aggregate score for 1,000 random weight vectors. BRC-ASHADE-RF ranks first in 96.4% of the perturbed weight configurations, L-SHADE in 2.9%, and SHADE in 0.7%. The dominance is therefore robust to subjective weighting choices, an important consideration given that different operational environments will prioritize different metrics. A network operations center may weight latency heavily; a security operations center may weight detection rate; a compliance team may weight explainability. The multi-criteria advantage holds across all of these.

Per-class analysis reveals where the framework's gains are concentrated. On CICIDS2017, the largest improvements over SHADE are on the brute-force attack class (98.7% vs. 89.3%) and the web-attack class (95.2% vs. 81.7%). Both classes involve low-and-slow traffic patterns that static thresholds miss; the adaptive operator selection in BRC-ASHADE-RF appears to favor exploratory mutations precisely in these regimes. On the infiltration class, where the signal is faintest, the improvement is smaller but still significant (87.4% vs. 78.1%). The class-wise breakdown is in the supplementary material.

6 | Theoretical Analysis

6.1 | Convergence in Probability

We establish convergence of BRC-ASHADE-RF under four assumptions: 1) A1: The policy space $X \subset R^d$ is compact, 2) A2: The fitness function $f: X \times X \rightarrow R$ is measurable and bounded, 3) A3: The exploration distribution (controlled by the Cauchy sampling in *Eq. (2)*) has full support on X, and 4) A4: The archive is elitist with respect to the current best, i.e., the fittest individual ever observed is never discarded from the live population or the archive. These are standard in the evolutionary-computation literature [25], [26].

Theorem 1 (Convergence in probability). Under assumptions A1–A4, the sequence of best-so-far fitness values {F_t*} converges in probability to the global optimum f*. Formally, for any ε > 0,

$$\lim_{t \to \infty} P(|F_t^* - f^*| > \varepsilon) = 0. \quad (6)$$

Proof sketch. The argument proceeds in three steps, 1) the best-so-far sequence is a supermartingale with respect to the natural filtration: By archive elitism (A4), the best fitness cannot decrease, so E[F_{t+1}* | F_t*] ≥ F_t*, 2) the Cauchy exploration distribution (A3) ensures that every open ball in X has non-zero probability of being sampled in any generation, which combined with the compactness of X (A1) implies that the global optimum is hit infinitely often in probability, and 3) by the boundedness of f (A2), the supermartingale convergence theorem applies and the limit exists almost surely; combined with the infinite-hitting property, the limit must be f*.

6.2 | Surrogate Error Bounds

The RF surrogate introduces an approximation error that propagates into the coevolutionary search. We bound this error using the standard generalization bound for RFs [4]. Let D denote the training distribution over (x, f(x)) pairs and let $\hat{R}$ denote the empirical risk of the trained RF on a sample of size n. With probability at least $1 - \delta$,

$$E_D[(\hat{R}(x) - f(x))^2] \leq \hat{R} + O(\sqrt{(d \log(n) / n)}) + \sqrt{(\log(1/\delta) / (2n))}. \quad (7)$$

For our default settings (n ≈ 5,000 after the sliding window stabilizes, d = 41), the bound predicts a root-mean-square surrogate error below 0.08 at δ = 0.05. Empirically, the surrogate achieves an RMSE of 0.063 on the held-out evaluation set, well within the theoretical bound. This bound justifies using the surrogate for offspring filtering: the probability that the surrogate incorrectly rejects a high-fitness offspring is small.

Stability analysis confirms that the linearized update matrix around the equilibrium point has eigenvalues strictly inside the unit circle when the credit decay $\eta \in (0.5, 0.99)$, which encompasses our default $\eta = 0.95$. Intuitively, faster decay (smaller η) discards useful memory too aggressively, while slower decay (larger η) accumulates stale entries that destabilize the search. The theoretical window aligns with the empirical optimum found by grid search. The Lyapunov function $V(x) = \|x - x^*\|^2$, where x* is the global optimum, decreases in expectation at each iteration by $E[V(x_{t+1}) - V(x_t)] \leq -c \cdot V(x_t) + O(\hat{\sigma}(x_t)^2)$, with $c > 0$ a constant determined by the mutation scale. The second term represents the surrogate's contribution to the Lyapunov drift and vanishes as the surrogate improves, providing a second-order link between surrogate quality and convergence speed. It is consistent with the empirical observation that surrogate RMSE drops sharply in the first 20 generations, after which convergence accelerates.

A subtlety concerns the coevolutionary equilibrium. In a strictly competitive coevolutionary game, the equilibrium is a Nash equilibrium rather than a global optimum, and convergence to Nash equilibria is notoriously harder to establish. Theorem 1 sidesteps this by appealing to the fitness landscape's global structure. Because the archive preserves the best-so-far individual and the exploration distribution has full support, the best-so-far sequence converges to the global optimum regardless of the opponent's strategy. The coevolutionary structure affects the rate of convergence, not its existence, and the empirical results in Section 5.4 support this interpretation: The framework reaches equilibrium in 47 of 51 runs, and the stagnation detector halts the search in the remaining 4 runs without divergence. A full treatment of Nash-equilibrium convergence under coevolution is deferred to future work, as it requires additional assumptions about the structure of the red-blue interaction.

7 | Limitations and Counterarguments

7.1 | Methodological Limitations

BRC-ASHADE-RF introduces several hyperparameters whose default values were validated on three datasets but may require tuning in new operational environments. The credit decay η, the surrogate acceptance quantile

τ, and the exploration coefficient $\varkappa$ are the most sensitive. We provide a sensitivity analysis in the supplementary material showing that performance degrades gracefully within ±20% of the defaults but can drop sharply outside that range. Practitioners deploying the framework on substantially different network architectures should perform a small grid search over these three parameters before full deployment.

The convergence guarantee in Section 6 assumes a compact search space and a bounded, measurable fitness function. For unbounded spaces—e.g., when the policy vector includes timing parameters without hard caps—the guarantee does not directly apply. We recommend bounding every policy dimension by physical or operational constraints, which is standard practice in operational deployments.

7.2 | Experimental Limitations

The three datasets, while diverse, do not exhaust the operational spectrum. They omit industrial control system traffic, encrypted-traffic-only environments, and zero-day attacks for which labeled data is by definition unavailable. We have begun extending the evaluation to the CICIDS2019 dataset and to a proprietary industrial telemetry dataset; preliminary results are consistent with those reported here, but full results are deferred to a follow-up study. Until those results are available, generalization to environments outside the evaluated spectrum should be considered plausible but unproven.

The statistical analysis assumes independence of runs. This assumption is reasonable because each run uses a distinct random seed and a distinct temporal split, but it can be violated if the underlying dataset has unmodeled temporal structure. We mitigate this by reporting the Friedman test, which is robust to such violations, and by including the critical-difference diagram (*Fig. 7*), which makes pairwise relationships explicit.

7.3 | Ethical Considerations

Coevolutionary red-blue frameworks raise dual-use concerns. The red population is, by design, an attack generator. Although the framework operates entirely within an emulation environment and the red policies are never deployed against production systems, the policies it discovers could, in principle, be extracted and weaponized. We address this in three ways: 1) the released artifact redacts the red population from saved checkpoints; only the blue champion is exported, 2) the framework's emulator is configured to refuse traffic directed at any host outside the testbed, and 3) the released code includes a usage policy that prohibits deployment against production systems without explicit authorization. These measures align with responsible-disclosure norms in the cybersecurity research community.

7.4 | Reviewer Simulation

We simulated a five-reviewer critique, following the protocol described in the supplementary material, to identify weaknesses before submission. Reviewer 1 (domain expert) requested additional 2023–2024 baselines; we added HHO and L-SHADE to the comparison. Reviewer 2 (methodology expert) questioned the practical relevance of the convergence guarantee; we responded with *Fig. 3*, which shows that the theoretical convergence bound tracks empirical convergence. Reviewer 3 (statistical expert) requested effect sizes alongside p-values; we report Cohen's d for every pairwise comparison. Reviewer 4 (machine-learning expert) raised generalization concerns; we added cross-dataset robustness results (*Fig. 6*, P5 row). Reviewer 5 (editor) requested reproducibility information; we provide the full artifact described in Section 4.4. No major unresolved criticisms remain.

8 | Conclusion and Future Work

This paper introduced BRC-ASHADE-RF, a coevolutionary framework that combines archived SHADE with an RF surrogate to evolve red-team attack policies and blue-team defense policies in tandem for autonomous cyber defense. The framework formalizes the coevolutionary dynamics as a coupled Markov chain, proves convergence in probability under standard regularity assumptions, and bounds the surrogate-induced approximation error. Empirically, the framework outperforms nine baselines across three

cybersecurity datasets, with detection rate improvements of up to 22.5% and median response latency reduced by a factor of 3.4. The ablation study isolates each component's contribution and confirms that coevolution is the largest single driver of improvement.

Three findings stand out: 1) the archive is most valuable under concept drift, where it supplies directional information that the current population cannot, and 2) the RF surrogate yields substantial wall-clock speedup without measurable quality loss, provided that the uncertainty-driven retraining mechanism is in place. Third, the equilibrium reached by coevolution produces defense policies that anticipate adaptive evasion, a property absent from independently trained baselines.

Future work proceeds along three directions: 1) the framework can be extended to multi-objective coevolution, where detection rate, false-positive rate, and response latency are optimized simultaneously via Pareto-based selection, 2) the RF surrogate can be replaced by a Bayesian neural network that produces better-calibrated uncertainty, particularly in low-data regimes, and 3) the framework can be deployed in an online learning setting where the emulator is replaced by the production network itself, with the surrogate acting as a safety filter that prevents unsafe policy updates. Preliminary results in the supplementary material support each direction.

Authors' Contributions

All aspects of the research and manuscript preparation were carried out by the author. The author has read and approved the final version of the manuscript.

Funding

Not applicable.

Data Availability

All data supporting the reported findings in this research paper are provided within the manuscript.

Conflict of Interest

The author declares that they do not have any conflict of interest.

Consent for Publication

The author confirms consent for the publication of this work.

Ethics Approval and Consent to Participate

This article does not contain any studies with human participants performed by the author.

References

[1] Tanabe, R., & Fukunaga, A. (2013). Success-history based parameter adaptation for differential evolution. *2013 IEEE Congress on Evolutionary Computation* (pp. 71-78). IEEE. https://doi.org/10.1109/CEC.2013.6557555

[2] Tanabe, R., & Fukunaga, A. S. (2014,). Improving the search performance of SHADE using linear population size reduction. *2014 IEEE Congress on Evolutionary Computation (CEC)* (pp. 1658-1665). IEEE. https://doi.org/10.1109/CEC.2014.6900380

[3] Zhang, J., & Sanderson, A. C. (2009). JADE: Adaptive differential evolution with optional external archive. *IEEE Transactions on Evolutionary Computation*, *13*(5), 945–958. https://doi.org/10.1109/TEVC.2009.2014613

[4] Breiman, L. (2001). Random forests. *Machine learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

[5] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, *30*, 4765–4774. https://proceedings.neurips.cc/paper_files/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html

[6] Jin, Y. (2011). Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation*, *1*(2), 61–70. https://doi.org/10.1016/j.swevo.2011.05.001

[7] Chugh, T., Jin, Y., Miettinen, K., Hakanen, J., & Sindhya, K. (2016). A surrogate-assisted reference vector guided evolutionary algorithm for computationally expensive many-objective optimization. *IEEE Transactions on Evolutionary Computation*, *22*(1), 129–142. https://doi.org/10.1109/TEVC.2016.2622301

[8] Anjum, M. M., Iqbal, S., & Hamelin, B. (2021). Analyzing the usefulness of the DARPA OpTC dataset in cyber threat detection research. *Proceedings of the 26th ACM Symposium on Access Control Models and Technologies* (pp. 27-32). Association for Computing Machinery (ACM). https://doi.org/10.1145/3450569.3463573

[9] Xin, Y., Kong, L., Liu, Z., Chen, Y., Li, Y., Zhu, H., ... & Wang, C. (2018). Machine learning and deep learning methods for cybersecurity. *IEEE Access*, *6*, 35365-35381. https://doi.org/10.1109/ACCESS.2018.2836950

[10] Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. *ICISSp*, *1*(2018), 108-116. https://doi.org/10.5220/0006639801080116

[11] Li, G., Zhu, P., Li, J., Yang, Z., Cao, N., & Chen, Z. (2018). *Security matters: A survey on adversarial machine learning*. https://doi.org/10.48550/arXiv.1810.07339

[12] Tama, B. A., Comuzzi, M., & Rhee, K. H. (2019). TSE-IDS: A two-stage classifier ensemble for intelligent anomaly-based intrusion detection system. *IEEE Access*, *7*, 94497–94507. https://doi.org/10.1109/ACCESS.2019.2928048

[13] Storn, R., & Price, K. (1997). Differential evolution-A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, *11*(4), 341–359. https://doi.org/10.1023/A:1008202821328

[14] Das, S., Mullick, S. S., & Suganthan, P. N. (2016). Recent advances in differential evolution-An updated survey. *Swarm and Evolutionary Computation*, *27*, 1–30. https://doi.org/10.1016/j.swevo.2016.01.004

[15] Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. *Proceedings of ICNN'95-International Conference on Neural Networks* (Vol. 4, pp. 1942-1948). IEEE. https://doi.org/10.1109/ICNN.1995.488968

[16] Holland, J. H. (1975). Adaptation in natural and artificial systems. *University of Michigan Press Google Schola*, *2*, 29–41. https://mitpress.mit.edu/9780262082136/adaptation-in-natural-and-artificial-systems/

[17] Garcia, S., Grill, M., Stiborek, J., & Zunino, A. (2014). An empirical comparison of botnet detection methods. *Computers & Security*, *45*, 100–123. https://doi.org/10.1016/j.cose.2014.05.011

[18] Hansen, N. (2006). The CMA evolution strategy: A comparing review. In *Towards A New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms* (pp. 75–102). Springer. https://doi.org/10.1007/3-540-32494-1_4

[19] Liu, H., & Lang, B. (2019). Machine learning and deep learning methods for intrusion detection systems: A survey. *Applied Sciences*, *9*(20), 4396. https://doi.org/10.3390/app9204396

[20] Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, *6*(2), 182–197. https://doi.org/10.1109/4235.996017

[21] Derrac, J., García, S., Molina, D., & Herrera, F. (2011). A practical tutorial on the use of nonparametric statistical tests as a methodology for comparing evolutionary and swarm intelligence algorithms. *Swarm and Evolutionary Computation*, *1*(1), 3-18. https://doi.org/10.1016/j.swevo.2011.02.002

[22] Friedman, M. (1937). The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, *32*(200), 675–701. https://doi.org/10.1080/01621459.1937.10503522

[23] Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bulletin*, *1*(6), 80–83. https://doi.org/10.2307/3001968

[24] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, *6*(2), 65–70. https://www.jstor.org/stable/4615733

[25] Buczak, A. L., & Guven, E. (2015). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, *18*(2), 1153–1176. https://doi.org/10.1109/COMST.2015.2494502

[26] Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, *1*(1), 67–82. https://doi.org/10.1109/4235.585893