

# Adversarial Genetic Programming for Cyber Security: A Rising Application Domain Where GP Matters

Una-May O'Reilly · Jamal Toutouh ·  
Marcos Pertierra · Daniel Prado  
Sanchez · Dennis Garcia · Anthony Erb  
Luogo · Jonathan Kelly · Erik Hemberg

Received: 18 October 2018 / Revised: 18 September 2019 / Published: 02 April 2020

**Abstract** Cyber security adversaries and engagements are ubiquitous and ceaseless. We delineate *Adversarial Genetic Programming for Cyber Security*, a research topic that, by means of genetic programming (GP), replicates and studies the behavior of cyber adversaries and the dynamics of their engagements. Adversarial Genetic Programming for Cyber Security encompasses extant and immediate research efforts in a vital problem domain, arguably occupying a position at the frontier where GP matters. Additionally, it prompts research questions around evolving complex behavior by expressing different abstractions with GP and opportunities to reconnect to the Machine Learning, Artificial Life, Agent-Based Modeling and Cyber Security communities. We present a framework called *RIVALS* which supports the study of network security arms races. Its goal is to elucidate the dynamics of cyber networks under attack by computationally modeling and simulating them.

**Keywords** Genetic programming · Coevolutionary algorithms · Cyber Security

**Copyright notice:** This article has been published in the *Genetic Programming and Evolvable Machines* journal. Cite as: O'Reilly, U.M., Toutouh, J., Pertierra, M. et al. Adversarial genetic programming for cyber security: a rising application domain where GP matters. *Genet Program Evolvable Mach* 21, 219250 (2020). <https://doi.org/10.1007/s10710-020-09389-y>

---

Una-May O'Reilly  
MIT CSAIL  
Tel.: +617-253-6437  
Fax: +123-45-678910  
E-mail: unamay@csail.mit.edu

Erik Hemberg  
MIT CSAIL  
Tel.: +123-45-678910  
Fax: +123-45-678910  
E-mail: hembergerik@gmail.com

## 1 Introduction

The natural world yields many examples of adversarial advantages arising through evolution. These are testaments to biological arms races that have played out with complex dynamics, typically over macroscopic time scales. There is a wide range of advantages evolved by predators and prey including examples of animal coloration – camouflage that is static such as the yellow crab spider (*Misumena vatia*) and dynamic such as the cuttlefish (see Fig. 1a,1b), bioluminescence (e.g. mid-water crustaceans and fireflies, Fig 1e,1f), and mimicry (e.g. the Monarch and Viceroy butterflies, Fig 1c,1d).

Contemporary human-made ecosystems also yield arms races that arise from adversarial competitions. In the U.S. automotive industry, firms engage in an arms race to develop innovative new products ahead of the competition. The relationship between innovations and firm performance has been



(a) A spider(*Misumena vatia*) [138]



(b) Cuttlefish [137]




(c) A Monarch butterfly [140]



(d) A Viceroy Butterfly [139]



(e) Fireflies [43]



(f) Firefly squid [141], Appendix A.1.

Fig. 1: Adversarial defense and attack adaptations. Coloration adaptation – camouflage (a), (b) and biomimicry (c), (d). Bioluminescence defensive (e), (f).

analyzed from a coevolutionary perspective [133] (See Appendix A.2). The *circular economy* and its related *cradle-to-cradle* idea describe manufacturer-customer relationships that cooperatively evolve while business ecosystems within the economy also evolve, though competitively [105]. In an example from taxation, tax shelters prey upon unintended lapses of the tax code until defensive mediations address them. On occasion, however, these mediations either introduce new ambiguity or signal evaders to target a different weakness or inefficiency and the arms race continues [125, 147].

One of the contemporary crucibles of adversarial activity and arms races is cyberspace. It is a complex, human-made ecosystem that spans networks (including the Internet), edge devices, servers, clouds, supercomputers. Each sub-ecosystem has sets of resources that collectively are a source of contention. Each includes human actors in circumstances where their software is proxy for their part in cyber adversarial engagements. Actors can have conflicting objectives. From a security perspective one population of actors have benign intentions and the other has malicious intentions. Benign “cyber actors” are conforming within social and legal norms. They may control data which they need to keep confidential or private. They may operate networks, or websites in order to conduct their enterprise’s business – digital commerce. Different devices e.g. cellphones and laptops support everyday social and economic functions. Underpinning all these activities are computational hardware and software. Through attacks on these resources, malicious actors seek to disrupt, infiltrate and exfiltrate to steal, poison, extort so they can profit financially and advance their causes.

Focusing on network security, one pernicious kind of network attack is named Denial of Service (DOS). The goal of a DOS attack is to consume resources of a target so the target has none left to serve legitimate clients. At the extreme end of the range of volume are Distributed DOS (DDOS) attacks. These are extremely aggressive but more rare, likely due to the cost of setting them up. At a technical level, DDOS attacks are composed from a software toolkit. One set of software comprises multiple tactics that allow a network of compromised servers, i.e. a botnet, to be stealthily established. Another has tactics that enable, from a command and control location, a (human) controller to direct the bots to launch DOS attacks. These tactics can even react to defensive measures being deployed during an attack. The controller strategically selects the tactics and aims them at a target.

It is arguable that DDOS attacks are analogous to population members of a species, with tactical variation akin to biological variation. This population arguably undergoes evolutionary adaptation because ineffective attacks are not reused and effective ones undergo adaptation in order to circumvent defensive measures (i.e. reproductive selection and variation are present). With each new generation of attacks, the defensive measures themselves adapt or evolve then they face the next round of adaptation based on the toolkit. The attackers general goal is to minimize its use of resources while maximizing its denial of service and the defenders general goal is to minimize the impact of the attack.




Fig. 2: Time series of size of bots (number of unique IPs of the compromised servers) in the MIRAI net. The size escalates prior to an attack. It plummets as the attack is repelled but adaptations allow it to sustain its size then regrow [6], modified for legibility [4].

In terms of dynamics, an escalating, adversarial arms race around engagements based on the conflicting objectives consequently emerges.

There is reliable and thorough documented evidence of DDOS attacks and counter measures that support an evolutionary interpretation, at a high level of abstraction [10]. MIRAI [130] is a notorious DDOS, receiving publicity for two high profile attacks: Krebs [24] and DYN [118]. From its inception to the present, MIRAI has adapted its tactics and targets to thwart defensive counter-measures, see Fig. 2. Fig. 3 shows a Red Queen like DDOS dynamics – while evolution is churning beneath the surface, the number of attacks and their diversity of volume barely changes. Defensive gains are relative and transient.

DOS attacks sit amongst an alarming array of other cyber attack types, such as Advanced Persistent Threats (APT) and Ransomware (see Appendix A.3) that exhibit intelligent and agile attacks as well as attack versus defense arms race dynamics. Across the entire cyber ecosystem there is ceaseless evolution of kinds, arguably species, of attacks and co-evolving responses. Beyond their inconvenience, attacks risk lives, have financial costs, threaten businesses, impinge upon privacy and disrupt legitimate activity.

One way to improve defenses requires understanding better how an attacker behaves and how competing attacks and defenses react to one another. Feeding back information on *what could happen* is extremely valuable. It elucidates potential or past behaviors. For defensive design, knowing what could happen can help by *identifying unforeseen or particularly critical attacks*. Knowledge of arms race dynamics can inform the *adversarially hardening* of defensive designs *ahead of time*.

Understanding what could happen requires examining, actually playing out, or simulating, the dynamics of ongoing engagements between multiple parties with varying behaviors. Genetic programming (GP) can meet these




Fig. 3: Evolution is churning underneath the surface where we see DDOS attack densities over time, modified for legibility [5].

requirements. It is ideally suited to represent the behavior of an adversary when it engages an opponent. Within a competitive evolutionary algorithm, GP representations can serve as the units of its selection and variation and GP operators can serve as the means of generating new solutions from existing ones.

Cyber security needs a way of exploring and executing attack versus defense engagements. GP is a paradigm that handles the performance selection and variation of units that **behave!** It can explore by directly manipulating executable functions without syntactic destruction or explicit program knowledge. It can express behavioral solutions drawn from a search space of hierarchical, varying length, executable structures and it has genetic operators that are able to blindly, at the representation level, generate novel syntactically correct solutions. Cyber security needs an expressively powerful way to describe abstracted attacks and defenses in a way that they are also able to actually compete against each other. GP can fulfill this because it can accommodate any abstraction that can be described as a set of *functions* and *terminals* or a computational language expressed by a grammar.

Combinations of adversarial evolutionary algorithms and GP matched with the rich problem domain of cyber security thus meld into an increasingly critical intersection with an agenda of compelling and still untapped scope, see Fig. 4. We name this topic *Adversarial Genetic Programming for Cyber Security*. It is *not* a new topic; it is the delineation of extant and current research approaches. We call attention to them because cyber security is urgent and escalating in challenge, and despite contributions to date, many novel approaches are required. At the 20 year mark of GP, when it is no longer necessary to show that GP works, it is timely to direct mature and nascent GP methods plus GP researchers in the direction of this critical domain with real world problems. Can search-based software engineering techniques gain traction on




Fig. 4: Domain intersections of *Adversarial Genetic Programming for Cyber Security*.

security exploits by considering adaptation at actual code level? Can GP enable the exploration of cyber space actors' goals and resulting plans, thereby addressing their intentions and cognitive reasoning? These questions and ones in between represent an opportunity for GP research that matters.

*Paper Roadmap* Having motivated *Adversarial Genetic Programming for Cyber Security*, in Section 2, we first succinctly cover GP and evolutionary algorithms (EAs). To consider adversarial evolution, we then describe competitive coevolutionary algorithms to remind readers of their complexity. We end the section by showing how GP and a coevolutionary algorithm can be combined. In Section 3 we survey prior evolutionary computation research into Artificial Intelligence (AI) and games (relevant because games are also competitive contexts), and software engineering (related because tests have been evolved to compete with program solutions to spur correctness). We then survey research around and within our central topic with Table 1 summarizing. Sections 2 and 3 inform the design motivations underlying a framework we present in Section 4. Named *RIVALS*, it combines competitive coevolutionary algorithms and GP for network security design. We describe three systems drawing on the framework and a component that extracts useful solutions to support the design decisions of a network defense designer. We draw to a close with a summary and a broad discussion of possible paths forward within this exciting, important paradigm.

## 2 Basis Algorithms for Adversarial Evolution

Variable length genotypes, such as the hierarchical tree structures introduced by Koza [69], define large search spaces and improve the flexibility and power of Evolutionary Algorithms conducted by executable structure search [67, 100]. GP is arguably defined by its distinctive structures that are variable length and executable plus its operators that preserve syntactic correctness while changing genotype size and structure. An Evolutionary Algorithm (EA) can evolve individual solutions in the form of executable structures, fixed length

genotypes such as the bit strings used by Genetic Algorithms (GAs) [51] are inflexible.

Biological coevolution refers to the influences two or more species exert on each other’s evolution [112]. A seminal paper on reciprocal relationships between insects and plants coined the word “coevolution” [40]. Coevolution can occur in the context of cooperation, where the species experience mutual benefit and adaptation, or in the context of competition, where the species negatively interact due to constrained resources that are under shared contention or a predator-prey relationship.

EAs usually abstract the evolutionary process while ignoring coevolution. To evaluate the quality of an individual, they apply an *a-priori* defined fitness function and this function does not reflect the possible interactions between another individual or a dynamic environment. Optimality can be formulated by an absolute rather than relative objective. Coevolutionary algorithms extend EAs by basing the fitness of an individual on its interactions with other individuals or the environment. They mimic the coupled species-to-species interaction mechanisms of natural coevolution in order to solve a niche of search, optimization, design, and modeling problems [11, 108, 112, 120].

Similar to the biology, coevolutionary algorithms are categorized as cooperative or competitive. Cooperative coevolutionary algorithms abstractly model the beneficial interaction between populations or in environments with time dependent factors [70]. They are frequently set up with multiple populations each solving a distinct sub-problem of a larger problem. We continue this section by describing competitive coevolutionary algorithms in more detail as it is central to adversarial behavior.

## 2.1 Competitive Coevolutionary Algorithms


A basic competitive coevolutionary algorithm evolves two coupled populations, each with conventional selection and representation-aligned variation (crossover and mutation) operators which suit the representation of the population’s genotype. One population comprises what is commonly called *tests* and the other *solutions*. We refer to individuals in a competitive coevolutionary algorithm generally as *adversaries*. In each generation, different competitions are formed by pairing up a test and a solution drawn from their respective populations. This couples the two population as they share a fitness evaluation component. We shall shortly describe different ways in which these competitions can be set up.

A test competes to demonstrate the solution as incorrect; this is its objective. The solution competes to solve the test correctly; this is its objective. In a security setting, it may be more apt to translate from *test* and *solution* to *attack* and *defense* as well as to refer to *engagements* rather than *competitions*. Fitness is calculated over all of an adversary’s engagements. The dynamic of the algorithm, driven by conflicting objectives and guided by performance-based selection and random variation can gradually produce better and more

robust solutions (i.e defenses) [112, 120]. Generally, competitive coevolutionary algorithms suit domains in which there is no exogenous objective measure of performance but wherein performance is relative to others. These have been called *interactive* domains in [108] and include games and software engineering. In most domains a competition is often computationally expensive because it involves simulation or a complex testbed. We next describe how competition structuring addresses this challenge.

*Competition Structures* For efficiency, the algorithm designer tries to minimize the number of competitions per generation while maximizing the accuracy of its fitness estimate of each adversary. The designer is able to control how competitions are structured and how many competitions are used to estimate the fitness of an adversary. Assuming one or both populations are of size  $N$ , two extreme structures are: *one-vs-one*, each adversary competes only once against a member of the opposing population, see Fig. 5a and 5d, and *all-vs-all*, each adversary competes against all members of the opposing population, see Fig. 5b and 5e. *One-vs-one* has minimal fitness evaluations<sup>1</sup> ( $O(N)$ ) and strong competition bias. In contrast, *all-vs-all* has a quadratic number of fitness evaluations, yielding a high computational cost,  $O(N^2)$  but weaker competition bias [120]. Other structures provide intermediate trade-offs between computation costs and competition bias, e.g. a *tournament* structure ranks individuals based on different rounds of peer competitions, see Fig. 5c.

In [89], adversaries termed hosts and parasites are placed on a  $M \times M$  grid with a fixed neighborhood (size  $c$ ) and one host and parasite per cell. The structure of the competitions is competition among all competitors in the neighborhood. Fitness evaluations are reduced to  $O(Mc^2)$  by this. An adversary has an outcome for each competition. We next discuss how these outcomes can be united to assign it a fitness score.



(a) One-vs-one, (b) All-vs-all, for (c) Tournament, (d) One-vs-one, (e) All-vs-all, for  
for 1 population. 1 population. for 1 population. for 2 populations. 2 populations.

Fig. 5: Competition structures between adversaries/competitors.

*Fitness Assignment* At each generation an adversary needs to be assigned a fitness score derived from some function of its performance outcomes in its

<sup>1</sup> computational cost is shown for two populations.



competitions. For example, the function can average the performance outcomes or use the maximum, minimum, or median outcome [15]. Other approaches have been defined depending on the specific problem domain, e.g. [44, 9, 27, 76, 81, 135]. A more formal approach, using the *solution* and *test* perspective, describes fitness assignments as *solution concepts* [108]. Solution concepts include: best worst case, maximization of expected utility, Nash equilibrium, and Pareto optimality.

Coevolutionary algorithms are more complex and challenging to direct toward some expected outcome than a single population EA. We next explain the pathologies they exhibit and describe accepted remedies for them.

*Pathologies and Remedies* The interactive aspect of fitness implies that a competitive coevolutionary algorithm lacks an *exact* fitness measurement. While EAs use a fitness function that allows an individual to be compared and globally ranked, in coevolutionary algorithms two members of the same population, during selection, may be imprecisely compared when they did not compete against the same opponents. Adding more imprecision, regardless of the function that computes a fitness score, an adversary’s score may change if it later competes against different opponents. These properties imply that any ranking of individuals is a noisy estimate. This estimation can lead to competitive coevolutionary algorithms *pathologies* that limit the effectiveness of their arms race modeling [108]. Furthermore, pathologies can arise from competitive imbalance, particularly when the behavior space of one population is not the same as that of the other. This asymmetry is considered and addressed in [98]. Pathologies include: *disengagement* – occurring when opponents are not challenging enough, eliminating an incentive to adapt, or when opponents are too challenging and progress becomes impossible [28], *cyclic dynamics* – generally appearing in transitive domains (A beats B, B beats C and C beats A) as oscillation of the evaluation metric [58], *focusing* or *overspecialization* – arising when an adversary evolves to beat only some of its opponents, while neglecting the others [26], and *coevolutionary forgetting* – occurs when an adversary’s ability to beat an opponent evolves away [23].

In order to remediate these pathologies, assorted methods have been proposed [22, 23, 26, 28, 42, 58, 70, 144]. These center on archives or memory mechanisms, maintaining suboptimal individuals in the population, using a neighborhood spatial structure, and/or monitoring the progress of the algorithm.

With the basics of GP and competitive coevolutionary algorithms introduced, we are now able to present a combined competitive coevolutionary and GP algorithm.

## 2.2 GP and Adversarial Evolution

Combining GP and competitive coevolutionary algorithms enables *cyber security* arms races to be replicated by evolving executable adversaries. An example of an alternating GP competitive coevolutionary algorithm [14] is shown

---

**Algorithm 1** Example of alternating GP Competitive coevolutionary algorithm

**Input:**

$T$ : number of generations    $\mathcal{L}$ : Fitness function,    $\mathcal{F}$ : Functions,    $\mathcal{T}$ : Terminals  
 $\mu$ : Mutation probability,    $\xi$ : Crossover probability,    $N$ : Population size

---

```

1:  $\mathbf{A}_0 \leftarrow [\mathbf{a}_{1,0}, \dots, \mathbf{a}_{N,0}] \sim \mathcal{U}(\{\mathcal{F}, \mathcal{T}\})$  ▷ Initialize minimizer(attacker) population
2:  $\mathbf{D}_0 \leftarrow [\mathbf{d}_{1,0}, \dots, \mathbf{d}_{N,0}] \sim \mathcal{U}(\{\mathcal{F}, \mathcal{T}\})$  ▷ Initialize maximizer(defender) population
3:  $t \leftarrow 0$  ▷ Initialize iteration counter
4: repeat
5:    $t \leftarrow t + 1$  ▷ Increase counter
6:    $\mathbf{A}_t \leftarrow \text{select}(\mathbf{A}_{t-1})$  ▷ Selection
7:    $\mathbf{A}_t \leftarrow \text{mutate}(\mathbf{A}_t, \mu)$  ▷ Mutation
8:    $\mathbf{A}_t \leftarrow \text{crossover}(\mathbf{A}_t, \xi)$  ▷ Crossover
9:    $\mathbf{a}'_*, \mathbf{d}'_* \leftarrow \arg \min_{\mathbf{a} \in \mathbf{A}_t} \arg \max_{\mathbf{d} \in \mathbf{D}_{t-1}} \mathcal{L}(\mathbf{a}, \mathbf{d})$  ▷ Best minimizer
10:  if  $\mathcal{L}(\mathbf{a}'_*, \mathbf{d}'_*) < \mathcal{L}(\mathbf{a}_{N,t-1}, \mathbf{d}_{N,t-1})$  then ▷ Replace worst minimizer
11:     $\mathbf{a}_{N,t-1} \leftarrow \mathbf{a}'_*$  ▷ Update population
12:  end if
13:   $\mathbf{A}_t \leftarrow \mathbf{A}_{t-1}$  ▷ Copy population
14:   $t \leftarrow t + 1$  ▷ Increase counter before alternating to maximizer
15:   $\mathbf{D}_t \leftarrow \text{select}(\mathbf{D}_{t-1})$  ▷ Selection
16:   $\mathbf{D}_t \leftarrow \text{mutate}(\mathbf{D}_t, \mu)$  ▷ Mutation
17:   $\mathbf{D}_t \leftarrow \text{crossover}(\mathbf{D}_t, \xi)$  ▷ Crossover
18:   $\mathbf{a}'_*, \mathbf{d}'_* \leftarrow \arg \min_{\mathbf{a} \in \mathbf{A}_t} \arg \max_{\mathbf{d} \in \mathbf{D}_t} \mathcal{L}(\mathbf{a}, \mathbf{d})$  ▷ Best maximizer
19:  if  $\mathcal{L}(\mathbf{a}'_*, \mathbf{d}'_*) > \mathcal{L}(\mathbf{a}_{N,t}, \mathbf{d}_{N,t-1})$  then ▷ Replace worst maximizer
20:     $\mathbf{d}_{N,t-1} \leftarrow \mathbf{d}'_*$  ▷ Update population
21:  end if
22:   $\mathbf{D}_t \leftarrow \mathbf{D}_{t-1}$  ▷ Copy population
23: until  $t \geq T$ 
24:  $\mathbf{a}_*, \mathbf{d}_* \leftarrow \arg \min_{\mathbf{a} \in \mathbf{A}_T} \arg \max_{\mathbf{d} \in \mathbf{D}_T} \mathcal{L}(\mathbf{a}, \mathbf{d})$  ▷ Best minimizer
25: return  $\mathbf{a}_*, \mathbf{d}_*$ 

```

---

in Algorithm 1. The adversaries (attacker/defender) are coupled and evolve in an alternating manner. First, the adversaries are initialized. Then, at each generation, a new attacker population,  $\mathbf{A}_t$ , is selected, mutated, recombined and evaluated against the current defense population,  $\mathbf{D}_{t-1}$ . Based on the evaluation the attackers are replaced. Then, control reverts to the defender population so it can evolve to conclude the generation. This algorithm allows more attack evolution than defensive evolution, reflecting what can happen in real cyber systems. A note of caution about this approach in the context of algorithms, rather than reality, is offered by [87] where it is observed that moderate environmental variation across generations is sufficient to promote the evolution of robust solutions.

There are some open challenges when applying *Adversarial Genetic Programming for Cyber Security*. Some of them are inherent in the use of EAs and adversarial evolution, such as the **subjective solution evaluation**, i.e., individuals interact with different opponents and the fitness is based on these, so it only provides a relative ordering [109, 110, 107]. Other amplified challenges are the **evaluation cost**, since the individual's fitness is calculated based on its interaction with other individuals which might require high computational costs [17, 66, 78, 89, 79]; Finally, combining GP and adversarial evolution generates complex models that complicate the **algorithm operator and parameter selection**.

This section has described algorithms that form the foundation of computational adversarial evolution. In the next section we describe prior Evolutionary Computation work that use GP, competitive coevolutionary algorithms or a combination of them in the domains of AI and Games, security, and Software Engineering.

### 3 Related Work

We now turn our attention to prior work relating to the theme of adversarial contexts. We organize by application domain, starting with domains which are competitive in nature but not cyber security: AI and games (3.1) and Software Engineering (3.2). Most examples take technical approaches that use competitive coevolutionary algorithms or GP. They provide essential illustrations of how the algorithms can be used as building blocks. Table 1 catalogs the related work by domain, algorithm, representation, competition structure, and fitness scoring. Fig. 6 illustrates how they intersect thematically.




Fig. 6: Venn diagram showing the intersections of domains in *Adversarial Genetic Programming for Cyber Security* related work.

#### 3.1 AI and Games

We find related work in AI and games because many games are competitive and game playing is a fertile research ground for investigating adversarial learning. We embrace a broad definition of **game** – encompassing board games, e.g. [106], video game playing, e.g. [63, 121] and social science games e.g. [15].

One of the first competitive coevolutionary algorithms was used to find efficient strategies for the Iterated Prisoners Dilemma [15, 16]. This seminal project used a single population and it structured its strategy competition by running a tournament among the population members. A single population is typical for symmetric games, i.e. games where each player has the

Table 1: Some representative related work to *cyber security*. It includes three different domains: games, security, and software, and three algorithmic solutions: GP (genetic programming), Comp Coev (Competitive coevolution), and combination of Comp Coev and GP. The columns relate to properties described in Section 2.1, algorithm class, number of populations, representation of the individuals, competition structure and how fitness is assigned to an individual.

Reference	Algorithm	# of Populations	Representation	Competition Structure	Fitness Scoring
<i>Games</i>					
Axelrod, et al. [15,16]	Comp Coev GA	One	Bit strings	Tournament	Maximum expected utility
Crawford-Marks et al. [35]	Comp Coev GP	Two: players and smart-balls	Trees	One vs subset	Maximum expected utility
Harper [55]	Comp Coev GP	One	Variable integer vector	One vs a subset	Maximum expected utility
Keaveney & O'Riordan [63]	Comp Coev GP	One	Trees	One vs subset	Maximum expected utility
Lam et al. [76]	GP	One	Trees	Playing games	Pareto optimality
Luke et al. [81]	Comp Coev GP	One	Trees	One vs one	Maximum expected utility
Miles et al. [88]	Comp Coev GA	One	Bit strings	One vs one	Maximum expected utility
<i>Security</i>					
Garcia et al. [48]	Comp Coev GP	Two: Defender & Attacker	Integers & BNF Grammar	NA	Multiple comparison
Henberg et al. [56]	Comp Coev GP	Two: Defender & Attacker	Integers & BNF Grammar	All vs All	Pareto optimality
Hingston & Preuss [57]	Comp Coev GA	Four: two of learners and two of tests	Pairs of real numbers	One vs all	Best-worst case
Kewley & Embrechts [64]	Comp Coev GA	Eight: four friendly and four enemy forces	Vector of four parameters	One vs a subset	Maximum expected utility
McDonald & Upton [86]	Comp Coev GA	Two: red and blue teams	Vector of parameters	One vs all	Force Exchange Ratio
Ostaszewski et al. [102]	Comp Coev AIS	Two: detectors and anomalies	Vector of parameters	One vs subset	Maximum expected utility
Rush et al. [113]	Comp Coev EA	Two: defenders and attackers	<i>Not specified</i>	One vs subset	Maximum expected utility
Service & Tauritz [119]	Comp Coev GA	Two: system hardenings and system faults	Bit strings: unified power flow controller installation	One vs a subset	Nash equilibrium
Sharaz-Tangil et al. [129]	GP	One	Trees: intrusion detection rules	Risk assessment	Maximum expected utility
Wintterrose & Carter [145]	GA	One: strategies	Bit strings: finite states machine	Fixed scenarios	Success against the defender
<i>Software engineering</i>					
Arcuri & Yao [13]	Comp Coev GP	Two: programs and tests	Trees	One vs all	Nash equilibrium
Adamopoulos et al. [3]	Comp Coev GA	Two: programs and tests	Sequences of mutant programs and tests	One vs all	Maximum expected utility
Oliveira et al. [97]	Comp Coev GA	Two: programs and tests	Sequences of mutant programs and tests	One vs all	Pareto optimality
Wilkerson & Tauritz [142]	Comp Coev GP	Two: programs and tests	Trees and lists	One vs a subset	Maximum expected utility

same set of moves (behavioral repertoire) and objectives. We find symmetric board game projects that use single population competitive coevolutionary algorithms to evolve Tic-Tac-Toe, Othello and Backgammon players. The Tic-Tac-Toe project represented players with GP [9] while both Othello and Backgammon projects used temporal difference learning, which is a gradient-based local search method [71, 106, 131, 132, 123]. A subsequent work used symmetric competitive coevolution in games and demonstrated impressive results through self-play [30]. A noteworthy system at the intersection of competitive coevolutionary algorithms and GP used a single population of agents cooperating as a team to compete for the “RoboCup97” soccer competition [81]. Soccer, being symmetric, would allow the system to potentially train against itself. In [122] the authors designed an evolutionary strategizing machine for game playing and other contexts. In [124], the authors coevolved strategies to solve Rubik’s Cube where the test population evolved the Rubik’s Cube configurations while the learner population evolved GP individuals to solve the Cube configurations. An ablation study demonstrated the contribution of competitive coevolution over random replacement of Cube configurations.

Cyber security adversaries are arguably asymmetric. In asymmetric games where opponents have different objectives and move sets, we find multi-population competitive coevolutionary algorithms. These include a competitive coevolutionary algorithm with evolutionary strategies that evolves Pac-Man versus Ghost Team players [27] and a PushGP system coevolving player versus smart-ball for the game “Quidditch”. The latter is noteworthy for its intersection of a competitive coevolutionary algorithm and GP.

In some asymmetric games one adversary is external to the learning system and quite complex. This drives single or multi-population coevolutionary algorithm systems. A competitive coevolutionary algorithm and GP study compared using a single population to nine populations when evolving controllers for a car racing game [135]. The multi-population approach generally produced better controllers. Similarly, real time strategy games (which provide challenging AI problems [72, 95]) have applied a single-population competitive coevolutionary algorithm and GP to develop automated players that use a progressive refinement planning technique [63]. Other work [88] coevolved real time strategy players by representing the AI agents with Influence Maps [38]. A real-time Neuroevolution of Augmenting Topologies (rtNEAT) approach evolved neural networks in real time while a game was played [128]. Behavior Trees (BT) were introduced to encode formal system specifications [7], but have also represented game AI behavior in commercial games [94]. GP evolved BT controllers for “Mario” [103] and “DEFCON” [76] games. Another study used Grammatical Evolution to generate Java programs to compete in “Robocode”. The players were represented by programs [55] and a spatial and temporal competitive coevolutionary algorithm and GP were used.

### 3.2 Software Engineering and Software Testing

Software testing called *fuzzing* arguably started with [54] which used grammars to generate program inputs for tests. This transitioned to direct automated random testing. One relevant security example is [50] that searched for bugs in a security protocol. Later, constraint based automatic test data generation was used to generate tests to verify if software variants were safe from malicious manipulation [39].

For similar purposes, the search based software engineering community uses evolutionary computation. Coevolution is used in SBSE with GA representations, see e.g. [3, 8, 97]. The community has used competitive coevolutionary algorithms and GP to coevolve programs and unit tests from their specification, e.g. [13, 142, 14]. Another study, [18], used coevolution to distinguish correct behavior from incorrect.

### 3.3 Cyber Security

We now turn to the domain of cyber security. There are a number of examples where, without any computational modeling or simulation, cyber security dynamics are described as evolutionary or as arms races. According to [143], *cyber security* is the task of minimizing an *attack surface* over time. The attack surface is the portion of a system which has vulnerabilities. Attackers attempt to influence the system's nominal state and operation by varying their interactions with the attack surface by stealth and non-compliance (see Appendix A.4). *Cyber security* perimeter protection (e.g. a firewall) is a battle fought and lost. Attackers are able to now gain access and defenders must resort to strategies that assume the attacker is present but hidden. One tactic to detect their presence is deception. For this a "honeypot" that entraps an attacker by imitating its target is commonly deployed. This arms race escalates *ad infinitum* as attackers then anticipate what the defenders have anticipated (e.g. the honeypots) and so on.

Retrospective security event reports also document coevolution. For example, [53] is an empirical study of malware evolution. Arguments for employing nature-inspired technologies for *cyber security* that mention how biological and ecological systems use information to adapt in an unpredictable world include [34, 45, 60, 85, 114].

A selected set of cyber security research that takes technical approaches is described in the Security subsection of Table 1. This set include works that we would not consider Adversarial Genetic Programming for Cyber Security: one work uses a genetic algorithm and neither coevolution or GP[145] and six others use are adversarial in nature, i.e. they use coevolution, but not GP[57, 64, 86, 102, 113, 119]. The remaining three projects fit within the topic of Adversarial GP[129, 48, 56]. The research within this set, along with other related work, can also be distinguished by what specific application it addresses in the domain of cyber security:

**Moving target defense (MTD)** techniques seek to randomize components to reduce the likelihood of a successful attack, reduce the attack lifetime, and diversify systems to limit the damage [41,96]. A MTD study investigates an adaptable adversarial strategy based on Prisoners Dilemma in [145]. Strategies are encoded as binary chromosomes representing finite state machines that evolve according to GA. The study has one adaptive defender population in GA and few fixed scenarios.

**Network Defense Investigation** is studied with the coevolutionary agent-based network defense lightweight event system (CANDLES) [113] a framework designed to coevolve attacker and defender agent strategies with a custom, abstract computer network defense simulation. The *RIVALS* network security framework, elaborated in Section 4 supports three studies into respectively DDOS, deceptive and isolation defense.

**Self-Adapting Cyber Defenses** are for example the Helix self-regenerative architecture [74]. Helix shifts the attack surface by automatically repairing programs that fail test cases using Software Dynamic Translation, e.g. Strata [117]. Another example of automated fault analysis and filter generation within a system is named “FUZZBUSTER” [91,92]. Like Helix, FUZZBUSTER is designed to automatically identify software vulnerabilities and create adaptations that protect those vulnerabilities before attackers can exploit them. Both FUZZBUSTER and Helix use a GP system called GenProg [136] for automatically fixing code.

**Physical Infrastructure defense** is studied in terms of how network components can be made resilient in [119].

**Anomaly detection** attempts to discern network or other activity behavior that is out of the ordinary or that differs from normal. In the auto-immune system computational paradigm, normal has been characterized as “self” [46]. Anomaly detection has been studied as a one class learning problem by [102] who use the artificial immune systems paradigm and coevolution for search. Attempts to build GP anomaly detectors have mostly assumed labeled data sets and they evolve a classifier that labels outputs as normal or not. They frequently encounter a class imbalance issue. This is explicitly addressed in [126]. Later work adopted an explicitly Pareto archive formulation of competitive coevolution[75].

**Vulnerability Testing** Vulnerability testing involves testing a system with modifications to known exploits or attack vectors. Examples of developing mimicry attacks to test for vulnerability are found in [61,62]. The approach used the alarm signal for coevolution of a GP exploit generator. Moreover, GP was limited to instructions that were in legitimate applications, forcing GP to search for exploits described in terms of instructions used by legitimate applications. Others have also used GP to coevolve port scans against the ‘Snort’ intrusion detector with the objective of evolving port scans that demonstrate holes in the IDS. See for example [73]. Vulnerability testing for malware in mobile applications using coevolution appears in [25].

**Malware detection** has seen GP used to evolve novel PDF malware to automatically evade classifiers [148] and to study malware in the form of return-oriented program evolution [47].

**Intrusion detection** Most examples of intrusion detection are addressed as a multi-class classification problem where activity such as network traffic must be labeled as normal, denial of service, botnet or something else. Multi-class classification has been studied using GP without coevolution. For example, [129] optimizes intrusion detection rules. Others study botnet detection with GP under streaming data label budgets and class imbalance [65]. With both intrusion detection and anomaly detection, one challenge is obtaining a dataset that truly reflects the network or any other cyber environment. For example, a widely used dataset known at KDD'99 has been criticized for its artificially generated normal data, which did not encompass the diversity in real normal behavior, see [77,84] for more details. It is also challenging to label and maintain datasets, see [49] as an example. The persistent nature of these challenges provides additional motivation for *RIVALS*.

**Battle management** has historically noteworthy work by [12] which focuses on security in battle management. It was novel in considering semi-autonomous control of several intelligent agents; plan adjustment based on developments during plan execution; and the influence of the presence of an adversary in devising plans. Similar but contemporary work on a computational military tactical planning system uses fuzzy sets and competitive coevolutionary algorithm with a battlefield tactics simulator for decision support [64].

**Red teaming** is a technique utilized by the military operational analysis community to check defensive operational plans by pitting actors posing as attackers (a red team) against the defense (a blue team). This activity has evolved from human teams to teams of programmers overseeing automated attacks tactics and defensive measures. See [37] and studies [1,93] trying to automate the exercise to use less manpower, or use it more efficiently with agent-based models [19,86,146,149,150,57].

*The Next Section* Prior work in AI and games, cyber security and software engineering domains is helpful in showing how competitive coevolutionary algorithms and GP can be used in adversarial contexts. It also provides a modest number of examples of GP combined with a competitive coevolutionary algorithm. We now proceed to present a detailed example of research within *Adversarial Genetic Programming for Cyber Security*. Named *RIVALS*, it is a software framework for studying network security. *RIVALS* addresses an aspect of central importance to the 2016 DARPA grand challenge in cyber security named “The World’s First All-Machine Hacking Tournament” [37]. This aspect that adaptation of both sides of a cyber security battle ground must be anticipated and that eventually posture reconfiguration needs to be fully automated.





Fig. 7: High level view of *RIVALS* framework. The left hand side component is a competitive coevolutionary algorithm that evolves two competing populations: attacks and defenses. The right hand side component from a computational perspective is a modeling or simulation environment. The environment is initialized with a mission and can be reset each time it is passed an attack-defense pairing to run an engagement. It first installs the defense, then it starts the mission and the attack. It evaluates the performance of the adversaries relative to their objectives and returns appropriate measurements.

## 4 RIVALS

Consistent with the *Adversarial Genetic Programming for Cyber Security* paradigm, our goal is to study the dynamics of cyber networks under attack by computationally modeling and simulating them. Ultimately we aim to provide defenders with information that allows them to anticipate attacks and design for resilience and resistance before their attack surfaces are attacked. We exploit competitive coevolutionary algorithms and GP for these purposes within a framework named *RIVALS* [101,104,111].

Conceptually, the framework consists of three elements: *adversaries*, *engagements and their environments*, and *competitive coevolution*. These fit together as two connected modules – one executing the coevolutionary algorithm and executing the engagements, and the other executing the engagements with the competitive coevolutionary algorithm directing attack and defense engagement pairings, see Fig. 7,.

We now proceed to describe each of these elements and *RIVALS*’ use cases.

### 4.1 Elements of *RIVALS*

*Adversaries:* The system consists of two adversarial populations – *attacks* and *defenses*. The primary objective of an *attack* is to impair a mission by maximizing disruption of some resource. The primary objective of a *defense* is to complete a mission by minimizing disruptions. Both attack and defense can have secondary objectives based on the cost of their behavior. We design attack and defense behavior by defining grammars that can express the different variations of attacks and defenses. An attack or defense is a “sentence” in the grammar’s language. We implement a rewriting process (“generator”) to form sentences from the grammar [99], see Fig. 8. The sentences (i.e. attacks and




Fig. 8: Grammatical evolution string rewriting. A context free grammar rewrites an integer sequence (genotype) to an output sentence (phenotype).

defenses) are functions that are executable in the layers of the engagement environment.

The grammar has Backus Naur Form (BNF) and is parsed to a context free grammar structure. The (rewrite) rules of the generator express how a sentence, i.e. attack or defense, can be composed by rewriting a start symbol. The adversaries are represented with variable length integer vectors. The generator decodes these vectors to control its rewriting. As a result, different vectors generate different attacks or defenses. For different use cases, it is only necessary to change the BNF grammar, engagement environment and fitness function of the adversaries. This modularity, and reusability of the parser and rewriter are efficient software engineering and problem solving advantages. The grammar additionally helps communicate to the designer how the system works and its assumptions, i.e. threat model. This enables conversations and validation at the domain level with experts and increases the confidence in solutions from the system.

*Engagements and the Engagement Environment* An engagement is an attack on a defense. Engagements take place in an *engagement environment* which is initialized with a mission to complete and a set of resources such as network services that support the mission. The defense is first installed in the environment and then, while the mission runs, the attack is launched. The scenario (mission and resources) and attacks are then executed. Engagements have outcomes that match up to objectives; they are phrased in terms of mission completion (primary objective) and resource usage (second objective). Implementation-wise, the engagement environment component can support a problem-specific network testbed, simulator or model. Mod-sim is appropriate when testbeds incur long experimental cycle times or do not abstract away irrelevant detail.

*Coevolutionary Algorithms* *RIVALS* maintains two *populations* of competing attackers and defenders. It calculates the fitness of each population member (in both attack and defense populations) by assessing its ability to successfully engage one or more members from the adversarial population, given its objective(s). It also directs selection and variation. *RIVALS* [111,104] utilizes different coevolutionary algorithms to generate diverse behavior. The algorithms, for further diversity, use different “solution concepts”, i.e. measures of



Fig. 9: Overview of the **ESTABLO** framework used by *RIVALS* for decision support through selection and visualization by using a compendium of solutions from coevolutionary algorithms.

adversarial success. However, engagements are often computationally expensive and have to be pairwise sampled from two populations at each generation, a number of enhancements enables efficient use of a fixed budget of computation or time.

*RIVALS' Compendium* Solely emulating or simulating cyber arms races is not sufficient to practically inform the design of better, anticipatory defenses. In fact, competitive coevolution poses general challenges when used for design optimization. The following ones in *RIVALS* make it difficult to present a designer with clear information derived solely from multiple simulation runs [115, 111]:

1. Attacks (and defenses) of different generations are not comparable because fitness is based solely on the composition of the defense (attack) population at each generation. So no clear champion emerges from running the algorithm.
2. From multiple runs, with one or more algorithms, it is unclear how to automatically select a “best” attack or design.

To this end, *RIVALS* provides an additional decision support component, named **ESTABLO** [111, 115], see Fig. 9. At the implementation level, the engagements and results of every run of any of the system’s coevolutionary algorithms are cached. Later, offline, **ESTABLO** filters these results and moves a subset to its *compendium*. To prepare for the decision support analysis, it then competes all the attacks in the compendium against all defenses and ranks them according to multiple criteria, e.g. maximum-expected utility, best-worst case. For the defensive designer, it also provides visualizations and comparisons of adversarial behaviors to inform the design process.

*Use Cases* The *RIVALS* framework supports use cases – investigations into arms races of a particular cyber network context. They each interface with the framework through a set of domain specific information. For each use case, grammars define the behavioral space of the adversaries, objectives define the goals of the adversaries for scoring fitness and the threat environment describes its engagement environment. For each use case, different parameters to control the coevolutionary algorithm are also available. Fig. 10 depicts the high level decomposition of the framework and shows how a use case interfaces with it.

Table 2 presents the domain specific information for three use cases. By name, these are:

- *DIFFRACTED*: Defending a peer-to-peer network against Distributed Denial of Service (DDOS) attacks [48] (Section 4.2)
- *AVAIL*: Defenses against device compromise contagion in a segmented enterprise network [56] (Section 4.3), and
- *DARK-HORSE*: Deceptive defense against the internal reconnaissance of an adversary within a software defined network [104] (Section 4.4)

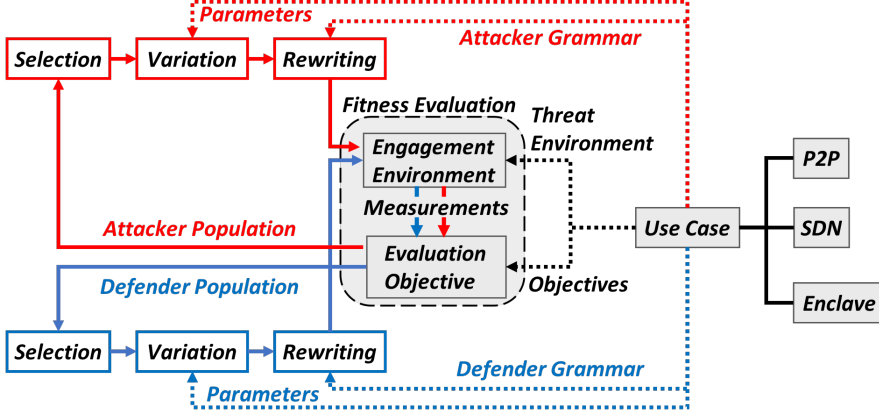


Fig. 10: RIVALS framework component decomposition showing interface with use cases.

The following sections elaborate on these use cases.

#### 4.2 DIFFRACTED– DOS Attacks on Peer-to-Peer Networks

A peer-to-peer (P2P) network is a robust and resilient means of securing mission reliability in the face of extreme distributed denial of service (DDOS) attacks. *DIFFRACTED* [48], assists in developing P2P network defense strategies against DDOS attacks. Attack completion and resource cost minimization serve as attacker objectives. Mission completion and resource cost minimization are the reciprocal defender objectives. DDOS attack strategies are modeled with a variety of behavioral languages.

A simple language e.g. allows a strategy to select one or more network servers to disable for some duration. Defenders choose one of three different network routing protocols: shortest path, flooding and a peer-to-peer ring overlay to try to maintain their performance. A more complex language allows a varying number of steps over which the attack is modulated in duration, strength and targets. It can even support an attack learning a parameterized




Fig. 11: The x axis shows a sorted subsample of attackers (note, the top 10 are shown and then every tenth) and the y axis shows the ranking score. The ranking is done on the scores from the compendium. The values for the same run and unseen test sets are shown on separate lines. The algorithm used to evolve the attacker is shown by the marker and the color. The attacker in the box with the solid line is the top ranked solution from the Combined Score ranking schemes. The solution in the dashed box is the top ranked solution from the Minimum Fitness ranking scheme.

condition that controls how it adapts during an attack, i.e. “online”, based on feedback it collects from the network on its impact. Defenders have simple languages related to parameterizations of P2P networks that influence the degree to which resilience is traded off with service costs. A more complex language allows the P2P network to adapt during an attack based on local or global observations of network conditions.

An example of attackers from *ESTABLO* on a mobile resource allocation defense used in *DIFRACTED* [115] is shown in Fig. 11. The mobile asset placement defense challenge is to optimize the strategic placement of assets in the network. While under the threat of node-level DDOS attack, the defense must enable a set of tasks. It does this by fielding feasible paths between the nodes that host the assets which support the tasks. A mobile asset is, for example, mobile personnel or a software application that can be served by any number of nodes. A task is, for example, the connection that allows personnel to use a software application.

#### 4.3 *AVAIL*- Availability Attacks on Segmented Networks

Attackers often introduce malware into networks. Once an attacker has compromised a device on a network, they spread to connected devices, akin to contagion. *AVAIL* considers *network segmentation*, a widely recommended defensive strategy, deployed against the threat of serial network security attacks that delay the mission of the network’s operator [56] in the context of malware spread.

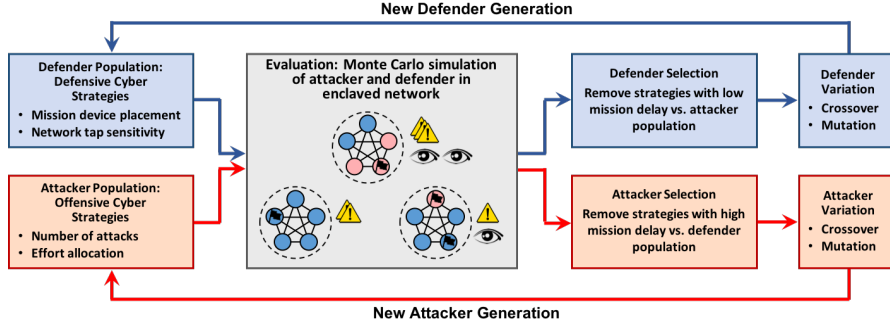


Fig. 12: Flow of the competitive coevolutionary algorithm used to evaluate defensive network enclave configurations (blue boxes) and contagion attacks (red boxes) in *AVAIL* [56]. A Monte Carlo simulation of device compromise contagion in a segmented network is used to assign the fitness of the adversaries. *AVAIL* uses Maximum Expected Utility as a fitness score based on the mission delay

*Network segmentation* divides the network topologically into *enclaves* that serve as isolation units to deter inter-enclave contagion. How much network segmentation is helpful is a tradeoff. On the one hand, a more segmented network provides lower mission efficiency because of increased overhead in inter-enclave communication. On the other hand, smaller enclaves contain compromise by limiting the spread rate, and their cleansing incurs fewer mission delays. Adding complexity, given some segmentation, a network operator can monitor threats and utilize cleansing policies to detect and dislodge attackers, with the caveat of cost versus efficacy.

*AVAIL* assumes an enterprise network in carrying out a *mission*, and that an adversary employs *availability attacks* against the network to disrupt it. Specifically, the attacker starts by using an exploit to compromise a vulnerable device on the network. This inflicts a mission delay when a mission critical device is infected. The attacker moves laterally to compromise additional devices to further delay the mission. Fig 12 shows *AVAIL*.

*AVAIL* employs a Monte Carlo simulation model as its engagement environment. Malware contagion of a specific spread rate is assumed. The defender decides placement of mission devices and tap sensitivities in the pre-determined enclave segmentation. The attacker decides the strength, duration and number of attacks in an attack plan targeting all enclaves. For a network with a set of four enclave topologies, the framework is able to generate strong availability attack patterns that were not identified *a priori*. It also identifies effective configurations that minimize mission delay when facing these attacks.




Fig. 13: *DARK-HORSE* fitness evaluation(left side) has a defensive SDN system based on the POX SDN Controller that creates different network views and manipulates traffic, and an attacker that performs NMAP scans in a *mininet* simulation. The scan results are passed into a fitness function that assigns fitness values for the adversaries' configurations. The search in the competitive coevolutionary algorithm uses Maximum Expected Utility of the detection time to assign a fitness value.

#### 4.4 *DARK-HORSE*– Internal Reconnaissance in Software Defined Networks

Once an adversary has compromised a network endpoint, they can perform network reconnaissance [127]. After reconnaissance provides a view of the network and an understanding of where vulnerable nodes are located, attackers are able to execute a plan of attack. One way to protect against reconnaissance is by obfuscating the network to delay the attacker. This approach is well suited to software defined networks (SDN) such as those deployed in cloud server settings because it requires programmability that they support [68]. The SDN controller knows which machines are actually on the network and can superficially alter (without function loss) the network view of each node, as well as place decoys (honeypots) on the network to mislead, trap and slow down reconnaissance. *DARK-HORSE* is shown in Fig. 13.

One such multi-component deceptive defense system [2] foils scanning by generating “camouflaged” versions of the actual network and providing them to hosts when they renew their DHCP leases. We use this deception system and *mininet* [134] within the framework as an engagement environment. This allows us to explore the dynamics between attacker and defender on a network where the deception and reconnaissance strategies can be adapted in response to each other [104].

A deception strategy is executed through a modified POX SDN controller. A reconnaissance strategy is executed by an NMAP scan[82]. The attacker strategy includes choices of: which IP addresses to scan, how many IP addresses to scan, which subnets to scan, the percent of the subnets to scan, the scanning speed, and the type of scan. The defender strategy includes choices of: the number of subnets to set up, the number of honeypots, the distribution of the real hosts throughout the subnets, and the number of real hosts that exist on the network. Fitness is comprised of four components: how fast the defender detects that there is a scan taking place, the total time it takes to run the scan, the number of times that the defender detects the scanner, and the number of

Table 2: How *RIVALS* components are configured to express specific use cases.

Component <i>NAME</i>	<i>Use Case</i>		
	<b>Robustness vs Denial</b> <i>DIFFRACTED</i>	<b>Deception vs Reconnaissance</b> <i>DARK-HORSE</i>	<b>Isolation vs Contagion</b> <i>AVAIL</i>
Threat Environment Evaluation Objective	P2P vs DDOS ALFA sim Mission disruption	Pox SDN vs Nmap Mininet Detection speed	Enclave vs Malware High level Sim Mission delay
Behavior: Attacker	Node/Edge impairment	Scanning parameters	Strength & Duration
Behavior: Defender	Network settings	Honeypots	Network taps & device placement
Adaptivity	Yes	No	No
Text description	Section 4.2	Section 4.4	Section 4.3

real hosts that the scanner discovers. Through experimentation and analysis, the framework is able to discover configurations that the defender can use to significantly increase its ability to detect scans. Similarly, there are specific reconnaissance configurations that have a better chance of being undetected.

#### 4.5 *RIVALS* Summary

We summarize *RIVALS* use cases with Table 2 and its architecture with Fig. 10. The *RIVALS* framework is one example of Adversarial Genetic Programming for Cyber Security, where the domain of network security is one of many possibilities in *cyber security*. Possible future steps are elaborated in the next section.

### 5 Taking Stock and Moving Forward

*Summary* Considering the paper from bottom to top, we described three use cases: *DIFFRACTED*, *AVAIL*, and *DARK-HORSE*. For 3 specific network security contexts, they investigate unique threat environments, objectives and behaviors for their adversary species, and varying capacities to adapt during an attack on a mission. They draw upon the *RIVALS* framework in which adversarial arms races are the activity of interest. *RIVALS* is an example of *Adversarial Genetic Programming for Cyber Security* and is motivated by a call to arms to work on automated defensive reaction to waves of attacks. *RIVALS* pushes towards this goal by considering how *both attack and defense* adapt to each other. While we identified a breadth of applications to cyber security in the set of EAs we surveyed, we found just a modest number of extant *adversarial GP* that represented approaches for behavior investigation and two population dynamics, the latter combining competitive coevolution



with GP. The body of prior work relevant to cyber security from AI and games and software engineering more generally informs the reader as to how competitive coevolutionary algorithms and EAs can be used as building blocks in different adversarial contexts. Finally, at the top, we argue that *Adversarial Genetic Programming for Cyber Security* is compellingly worthy of future community attention because of its potential to solve a critical and growing set of challenges in cyber security and because of how well GP and adversarial evolution match the technical nature of the challenges.

*Future directions* To end, we consider what future research questions and directions of study, stemming from the domain of cyber security, are prompted by *Adversarial GP*.

We forecast explorations driven by the wide array of cyber security scenarios will arise. There is a plethora of attack surfaces. We have mentioned some in Section 3.3, however, there are others such as networked cyber physical systems, ransomware, gadgets and malware. For example, how can GP help examine the adversarial cyber security of a power grid or self-driving cars? How could ransomware's next mutation be predictable? Can GP leverage code repositories and gadgets to discover new malware? Social engineering is currently an Achilles heel for security measures because it preys on human nature and once privileged access has been obtained, it is much harder to discern an attacker. Insider threats similarly disguise attacks. How can GP inform insider threats and social engineering?

One challenge for the GP community is to become adequately conversant in cyber security topics (though no member needs to be conversant in all of them). Collaborations with cyber security researchers seem advisable, offering additive and synergistic benefits. It will be important to identify the most appropriate level of abstractions from which to design GP languages or function and terminal sets. It will be imperative to develop metrics of success and paths to technology transfer and solution deployment. At a practical level, undoubtedly new modeling and simulation modeling platforms will be needed. They will need to be general purpose (somewhat like *RIVALS*) and specific. Undoubtedly some will need to be scaled due to the computational cost of executing a mission.

We forecast the opportunity for extensions and innovations in GP and coevolutionary algorithms for cyber security. Different patterns and rates of evolution will be observed across different security scenarios. The opening to develop new algorithmic models of these phenomena is exciting. GP evolves behavior but many different expressions of behavior remain unexplored. Developing the new behavioral evolutionary capabilities that will be needed is both challenging and motivating.

We also forecast that research into *Adversarial Genetic Programming for Cyber Security* will push bridge building to other research areas. One impetus will be the quest to describe at a finer granularity what is happening during cyber evolution. For this, it may be advisable to consider evolution through the lens of individuals, such as how Artificial Life (ALife) [83] models individuals.

What cyber security scenario should be studied with ALife? How can agents in a cyber security ALife system be represented by evolvable GP behaviors? What if a GP defender “died” after an attack and attackers starve as they fail to penetrate defenses? How would dynamics at this scale offer different insights? What would a simulation of the DDOS attack ecosystem reveal about the progression of botnet sizes if intra-species competition was examined? What can ALife studies into ecosystem dynamics around intra-species competition and cooperation reveal?

Arguably, adversarial GP is a form of agent-based modeling (ABM, [90]) where agents are GP executable structures. Some ABM systems model more complex domains while the agents have simpler strategy spaces. Are there ABM investigations such as those of financial markets, crowd control, infectious diseases or traffic simulation that suggest innovative transfer to Adversarial GP and cyber security? Both ALife and ABM have examples that consider the spatial dimension of a domain. Competitive coevolutionary algorithms have spatial models. Interesting new models lie at these intersections.

Another area with potentially valuable interaction is Machine Learning (ML). Statistical machine learning relies on retrospective training data to learn a model capable of generalizing to new unseen data drawn from the same distribution as the retrospective data. In contrast, *Adversarial Genetic Programming for Cyber Security* is not data driven and this conveys it a niche. However, what can be transferred from the studies that have started to bridge supervised learning and test-based co-optimization[107]? Adversarial classification [36] as a game between classifier and adversary also appears in ML. As well, Generative Adversarial Networks (GANs) [52] derive generative models from samples of a distribution using adversarial neural network training. Coevolution has been applied to improve the robustness and scalability of these systems [116]. Competitive coevolution has also been combined with multi-layer perceptron classifiers with floating point representation [29].

Studies of adversarial coevolution in cyber security based on data that classify attacks using machine learning methods exist, e.g. see a pro-active defense for evolving cyber threats and a predictive defense against evolving adversaries in [31,32,33]. In adversarial Machine Learning the attacker manipulates data in order to defeat the Machine Learning model [80,59,20]. Can coevolution be used to anticipate continual adversarial evasion and guide model hardening?

Game theory considers equilibria and paths to equilibria. What is an equilibrium in cyber security? It could be the complete wipeout of one side of the adversarial equation. Or, it could be the point where an attacker chooses not to target a defense because it is less expensive to go elsewhere. It could be a Nash equilibrium where neither attacker or defender has a better tactical position to go to unless the other simultaneously changes also. One direction should investigate how these game theoretic notions can inform the highly empirical, algorithmic systems of Adversarial GP.

We believe the future directions are not enumerable because the likelihood of new adaptations by adversaries will never be zero. To date, defenses expose far more attack surface than they can protect and attackers only need to find

one crack to penetrate. To this end, the emerging importance of *Adversarial Genetic Programming for Cyber Security* is not likely to fade. We trust this justifies our call to arms.

## Acknowledgements

This was supported by the CSAIL CyberSecurity Initiative. This material is based upon work supported by DARPA. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements. Either expressed or implied of Applied Communication Services, or the US Government. This project has received funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 799078.

## References

1. Abbass, H.A.: The art of red teaming. In: Computational Red Teaming, pp. 1–45. Springer (2015)
2. Achleitner, S., Laporta, T., McDaniel, P.: Cyber deception: Virtual networks to defend insider reconnaissance. In Proceedings of the 2016 International Workshop on Managing Insider Security Threats pp. 57–68 (2016)
3. Adamopoulos, K., Harman, M., Hierons, R.M.: How to overcome the equivalent mutant problem and achieve tailored selective mutation using co-evolution. In: Genetic and Evolutionary Computation–GECCO 2004, pp. 1338–1349. Springer (2004)
4. Akamai: Akamai's State of the Internet / Security Report - Q1 2017 report. Tech. rep., Akamai Technologies, Inc. (2017). URL <https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/q1-2017-state-of-the-internet-security-report.pdf>
5. Akamai: Akamai's State of the Internet / Security Report - Q3 2017 report. Tech. rep., Akamai Technologies, Inc. (2017). URL <https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/q3-2017-state-of-the-internet-security-report.pdf>
6. Akamai Technologies: State of the internet quarterly security reports (2017). URL <https://www.akamai.com/us/en/about/our-thinking/state-of-the-internet-report/global-state-of-the-internet-security-ddos-attack-reports.jsp>
7. Alex, J.C.: Behavior trees for next-gen game ai. In: Game Developers Conference, Lyon, France, pp. 3–4 (2007)
8. Anand, S., Burke, E.K., Chen, T.Y., Clark, J., Cohen, M.B., Grieskamp, W., Harman, M., Harrold, M.J., McMin, P., et al.: An orchestrated survey of methodologies for automated software test case generation. *Journal of Systems and Software* **86**(8), 1978–2001 (2013)
9. Angeline, P.J., Pollack, J.B.: Competitive environments evolve better solutions for complex tasks. In: Proceedings of the Fifth International Conference (GA93), Genetic Algorithms, pp. 264–270 (1993)
10. Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., Dumeric, Z., Halderman, J.A., Invernizzi, L., Kallitsis, M., et al.: Understanding the mirai botnet. In: 26th {USENIX} Security Symposium ({USENIX} Security 17), pp. 1093–1110 (2017)

11. Antonio, L.M., Coello, C.A.C.: Coevolutionary multi-objective evolutionary algorithms: A survey of the state-of-the-art. *IEEE Transactions on Evolutionary Computation* pp. 1–16 (2018). DOI 10.1109/TEVC.2017.2767023
12. Applegate, C., Elsaesser, C., Sanborn, J.: An architecture for adversarial planning. *Systems, Man and Cybernetics, IEEE Transactions on* **20**(1), 186–194 (1990)
13. Arcuri, A., Yao, X.: Coevolving programs and unit tests from their specification. In: *Proceedings of the twenty-second IEEE/ACM international conference on Automated software engineering*, pp. 397–400. ACM (2007)
14. Arcuri, A., Yao, X.: Co-evolutionary automatic programming for software development. *Information Sciences* **259**, 412–432 (2014)
15. Axelrod, R.: *The Evolution of Cooperation*. 10. Basic, NY, New York (1984)
16. Axelrod, R., et al.: The evolution of strategies in the iterated prisoners dilemma. *The dynamics of norms* pp. 1–16 (1987)
17. Bari, A.G., Gaspar, A., Wiegand, R.P., Bucci, A.: Selection methods to relax strict acceptance condition in test-based coevolution. In: *2018 IEEE Congress on Evolutionary Computation (CEC)*, pp. 1–8. IEEE (2018)
18. Barr, E., Harman, M., McMinn, P., Shahbaz, M., Yoo, S.I.: The oracle problem in software testing: A survey (2015)
19. Beard, D.: *Enhancing Automated Red Teaming with Monte Carlo Tree Search* (2011)
20. Biggio, B., Roli, F.: Wild patterns: Ten years after the rise of adversarial machine learning. *arXiv preprint arXiv:1712.03141* (2017)
21. Bodeau, D., Graubart, R.: Characterizing effects on the cyber adversary: A vocabulary for analysis and assessment. The MITRE Corporation, Bedford, MA (2013)
22. Bongard, J.C., Lipson, H.: Nonlinear system identification using coevolution of models and tests. *IEEE Transactions on Evolutionary Computation* **9**(4), 361–384 (2005)
23. Boyd, R.: Mistakes allow evolutionary stability in the repeated prisoner's dilemma game. *Journal of theoretical Biology* **136**(1), 47–56 (1989)
24. Brian Krebs: Akamai on the Record KrebsOnSecurity Attack. <https://krebsonsecurity.com/2016/11/akamai-on-the-record-krebsonsecurity-attack/> (2016). (Accessed October 10,2018)
25. Bronfman-Nadas, R., Zincir-Heywood, N., Jacobs, J.T.: An artificial arms race: Could it improve mobile malware detectors? In: *2018 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 1–8. IEEE (2018)
26. Bucci, A.: *Emergent geometric organization and informative dimensions in coevolutionary algorithms*. Ph.D. thesis, Brandeis University (2007)
27. Cardona, A.B., Togelius, J., Nelson, M.J.: Competitive coevolution in ms. pac-man. In: *2013 IEEE Congress on Evolutionary Computation*, pp. 1403–1410 (2013)
28. Cartlidge, J., Bullock, S.: Combating coevolutionary disengagement by reducing parasite virulence. *Evolutionary Computation* **12**(2), 193–222 (2004)
29. Castellani, M.: Competitive co-evolution of multi-layer perceptron classifiers. *Soft Computing* **22**(10), 3417–3432 (2018)
30. Chellapilla, K., Fogel, D.B.: Evolution, neural networks, games, and intelligence. *Proceedings of the IEEE* **87**(9), 1471–1496 (1999)
31. Colbaugh, R., Glass, K.: Proactive defense for evolving cyber threats. In: *Intelligence and Security Informatics (ISI), 2011 IEEE International Conference on*, pp. 125–130. IEEE (2011)
32. Colbaugh, R., Glass, K.: Predictive defense against evolving adversaries. In: *Intelligence and Security Informatics (ISI), 2012 IEEE International Conference on*, pp. 18–23. IEEE (2012)
33. Colbaugh, R., Glass, K.: Moving target defense for adaptive adversaries. In: *Intelligence and Security Informatics (ISI), 2013 IEEE International Conference on*, pp. 50–55. IEEE (2013)
34. Crandall, J.R., Ensafi, R., Forrest, S., Ladau, J., Shebaro, B.: The ecology of malware. In: *Proceedings of the 2008 workshop on New security paradigms*, pp. 99–106. ACM (2009)
35. Crawford-Marks, R., Spector, L., Klein, J.: Virtual witches and warlocks: A quidditch simulator and quidditch-playing teams coevolved via genetic programming. In: *Late-Breaking Papers of GECCO-2004, the Genetic and Evolutionary Computation Conference*. Published by the International Society for Genetic and Evolutionary Computation (2004)

36. Dalvi, N., Domingos, P., Sanghai, S., Verma, D., et al.: Adversarial classification. In: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 99–108. ACM (2004)
37. DARPA: The World's First All-Machine Hacking Tournament. <http://archive.darpa.mil/cybergrandchallenge/> (2016). (Accessed October 10, 2018)
38. DeLoura, M.A.: Game programming gems 2. Cengage learning (2001)
39. DeMilli, R., et al.: Constraint-based automatic test data generation. *Software Engineering, IEEE Transactions on* **17**(9), 900–910 (1991)
40. Ehrlich, P.R., Raven, P.H.: Butterflies and plants: a study in coevolution. *Evolution* **18**(4), 586–608 (1964)
41. Evans, D., Nguyen-Tuong, A., Knight, J.: Effectiveness of moving target defenses. In: *Moving Target Defense*, pp. 29–48. Springer (2011)
42. Ficici, S.G.: Solution concepts in coevolutionary algorithms. Ph.D. thesis, Brandeis University (2004)
43. Flickr: Fireflies brighter (2014). URL <https://www.flickr.com/photos/antoniseb/14325795079/in/gallery-flickr-72157645552049011flickr>. Picture taken by Jay Cross - License: CC-BY-SA-2.0
44. Fogel, D.: *Blondie24: Playing at the edge of artificial intelligence* (2001)
45. Ford, R., Bush, M., Bulatov, A.: Predation and the cost of replication: New approaches to malware prevention? *computers & security* **25**(4), 257–264 (2006)
46. Forrest, S., Hofmeyr, S.A., Somayaji, A., Longstaff, T.A.: A sense of self for unix processes. In: *Proceedings 1996 IEEE Symposium on Security and Privacy*, pp. 120–128. IEEE (1996)
47. Fraser, O.L., Zincir-Heywood, N., Heywood, M., Jacobs, J.T.: Return-oriented programme evolution with roper: a proof of concept. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 1447–1454. ACM (2017)
48. Garcia, D., Lugo, A.E., Hemberg, E., O'Reilly, U.M.: Investigating coevolutionary archive based genetic algorithms on cyber defense networks. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO '17*, pp. 1455–1462. ACM, New York, NY, USA (2017)
49. Garcia, S., Grill, M., Stiborek, J., Zunino, A.: An empirical comparison of botnet detection methods. *computers & security* **45**, 100–123 (2014)
50. Godefroid, P., Klarlund, N., Sen, K.: Dart: directed automated random testing. In: *ACM Sigplan Notices*, vol. 40, pp. 213–223. ACM (2005)
51. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization and Machine Learning*, 1st edn. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA (1989)
52. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014)
53. Gupta, A., Kuppili, P., Akella, A., Barford, P.: An empirical study of malware evolution. In: *Communication Systems and Networks and Workshops, 2009. COMSNETS 2009. First International*, pp. 1–10. IEEE (2009)
54. Hanford, K.V.: Automatic generation of test cases. *IBM Systems Journal* **9**(4), 242–257 (1970)
55. Harper, R.: Evolving robocode tanks for evo robocode. *Genetic Programming and Evolvable Machines* **15**(4), 403–431 (2014)
56. Hemberg, E., Zipkin, J.R., Skowrya, R.W., Wagner, N., O'Reilly, U.M.: Adversarial co-evolution of attack and defense in a segmented computer network environment. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 1648–1655. ACM (2018)
57. Hingston, P., Preuss, M.: Red teaming with coevolution. In: *Evolutionary Computation (CEC), 2011 IEEE Congress on*, pp. 1155–1163 (2011). DOI 10.1109/CEC.2011.5949747
58. Hornby, G.S., Mirtich, B.: Diffuse versus true coevolution in a physics-based world. In: *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-Volume 2*, pp. 1305–1312. Morgan Kaufmann Publishers Inc. (1999)
59. Huang, L., Joseph, A.D., Nelson, B., Rubinstein, B.I., Tygar, J.: Adversarial machine learning. In: *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, pp. 43–58. ACM (2011)

60. Iliopoulos, D., Adami, C., Szor, P.: Darwin inside the machines: malware evolution and the consequences for computer security. arXiv preprint arXiv:1111.2503 (2011)
61. Kayacik, H.G.: Can the best defense be a good offense? Evolving (MIMICRY) attacks for detector vulnerability testing under a black-box assumption. Ph.D. thesis, Dalhousie University Halifax (2009)
62. Kayacik, H.G., Zincir-Heywood, A.N., Heywood, M.I.: Can a good offense be a good defense? Vulnerability testing of anomaly detectors through an artificial arms race. *Applied Soft Computing* **11**(7), 4366–4383 (2011)
63. Keaveney, D., O’Riordan, C.: Evolving coordination for real-time strategy games. *IEEE Transactions on Computational Intelligence and AI in Games* **3**(2), 155–167 (2011)
64. Kewley, R., Embrechts, M.: Computational military tactical planning system. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* **32**(2), 161–171 (2002). DOI 10.1109/TSMCC.2002.801352
65. Khanchi, S., Vahdat, A., Heywood, M.I., Zincir-Heywood, A.N.: On botnet detection with genetic programming under streaming data label budgets and class imbalance. *Swarm and evolutionary computation* **39**, 123–140 (2018)
66. Kim, H.S., Cho, S.B.: An efficient genetic algorithm with less fitness evaluation by clustering. *Proceedings of the 2001 Congress on Evolutionary Computation* pp. 887–894 (2001)
67. Kinnear, K.E., Langdon, W.B., Spector, L., Angeline, P.J., O’Reilly, U.M.: *Advances in genetic programming*, vol. 3. MIT press (1999)
68. Kirkpatrick, K.: Software-defined networking. *Communications of the ACM* **56**(9) (2013)
69. Koza, J.R.: *Genetic programming II, automatic discovery of reusable subprograms*. MIT Press, Cambridge, MA (1992)
70. Krawiec, K., Heywood, M.: Solving complex problems with coevolutionary algorithms. In: *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference Companion*, pp. 687–713. ACM (2016)
71. Krawiec, K., Szubert, M.G.: Learning n-tuple networks for othello by coevolutionary gradient search. In: *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation, GECCO ’11*, pp. 355–362. ACM (2011)
72. Lara-Cabrera, R., Cotta, C., Fernandez-Leiva, A.J.: A review of computational intelligence in rts games. In: *2013 IEEE Symposium on Foundations of Computational Intelligence (FOCI)*, pp. 114–121 (2013)
73. LaRoche, P., Zincir-Heywood, N., Heywood, M.I.: Evolving tcp/ip packets: a case study of port scans. In: *2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications*, pp. 1–8. IEEE (2009)
74. Le Goues, C., Nguyen-Tuong, A., Chen, H., Davidson, J.W., Forrest, S., Hiser, J.D., Knight, J.C., Van Gundy, M.: Moving target defenses in the helix self-regenerative architecture. In: *Moving Target Defense II*, pp. 117–149. Springer (2013)
75. Lemczyk, M., Heywood, M.I.: Training binary gp classifiers efficiently: a pareto-coevolutionary approach. In: *European Conference on Genetic Programming*, pp. 229–240. Springer (2007)
76. Lim, C.U., Baumgarten, R., Colton, S.: Evolving behaviour trees for the commercial game DEFCON. In: *European Conference on the Applications of Evolutionary Computation*, pp. 100–110. Springer (2010)
77. Lippmann, R., Haines, J.W., Fried, D.J., Korba, J., Das, K.: The 1999 darpa off-line intrusion detection evaluation. *Computer networks* **34**(4), 579–595 (2000)
78. Liskowski, P., Krawiec, K.: Non-negative matrix factorization for unsupervised derivation of search objectives in genetic programming. In: *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference*, pp. 749–756. ACM (2016)
79. Liskowski, P., Krawiec, K.: Online discovery of search objectives for test-based problems. *Evolutionary computation* **25**(3), 375–406 (2017)
80. Lowd, D., Meek, C.: Adversarial learning. In: *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pp. 641–647. ACM (2005)
81. Luke, S., et al.: Genetic programming produced competitive soccer softbot teams for robocup97. *Genetic Programming* **1998**, 214–222 (1998)

82. Lyon, G.: Nmap network scanner. <https://nmap.org/> (2018). [Online; accessed 6-July-2018]
83. Macal, C.M., North, M.J.: Tutorial on agent-based modelling and simulation. *Journal of simulation* **4**(3), 151–162 (2010)
84. Mahoney, M.V., Chan, P.K.: An analysis of the 1999 darpa/lincoln laboratory evaluation data for network anomaly detection. In: G. Vigna, C. Kruegel, E. Jonsson (eds.) *Recent Advances in Intrusion Detection*, pp. 220–237. Springer Berlin Heidelberg, Berlin, Heidelberg (2003)
85. Mazurczyk, W., Drobniak, S., Moore, S.: Towards a systematic view on cybersecurity ecology. *arXiv preprint arXiv:1505.04207* (2015)
86. McDonald, M.L., Upton, S.C.: Investigating the dynamics of competition: coevolving red and blue simulation parameters. In: *Proceedings of the 37th conference on Winter simulation*, pp. 1008–1012. Winter Simulation Conference (2005)
87. Milano, N., Carvalho, J.T., Nolfi, S.: Moderate environmental variation across generations promotes the evolution of robust solutions. *Artificial life* **24**(4), 277–295 (2019)
88. Miles, C., Quiroz, J., Leigh, R., Louis, S.J.: Co-evolving influence map tree based strategy game players. In: *Computational Intelligence and Games, 2007. CIG 2007. IEEE Symposium on*, pp. 88–95. IEEE (2007)
89. Mitchell, M.: *Coevolutionary learning with spatially distributed populations. Computational intelligence: principles and practice* (2006)
90. Moran, N., Pollack, J.: Effects of cooperative and competitive coevolution on complexity in a linguistic prediction game. In: *Artificial Life Conference Proceedings 14*, pp. 298–205. MIT Press (2017)
91. Musliner, D.J., Friedman, S.E., Rye, J.M., Marble, T.: Meta-control for adaptive cybersecurity in fuzzbuster. In: *Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on*, pp. 219–226. IEEE (2013)
92. Musliner D. J. Friedman Scott E., R.J.M.: Automated fault analysis and filter generation for adaptive cybersecurity. In: *Proc. 6th Int'l Conf. on Adaptive and Self-Adaptive Systems and Applications* (2014)
93. Nettles, A.B.: The president has no clothes: the case for broader application of red teaming within homeland security. Tech. rep., DTIC Document (2010)
94. Nicolau, M., Perez-Liebana, D., O'Neill, M., Brabazon, A.: Evolutionary behavior tree approaches for navigating platform games. *IEEE Transactions on Computational Intelligence and AI in Games* **9**(3), 227–238 (2017)
95. Nogueira-Collazo, M., Porras, C.C., Fernandez-Leiva, A.J.: Competitive algorithms for coevolving both game content and ai. a case study:planet wars. *IEEE Transactions on Computational Intelligence and AI in Games* **8**(4), 325–337 (2016)
96. Okhravi, H., Hobson, T., Bigelow, D., Streilein, W.: Finding focus in the blur of moving-target techniques. *Security Privacy, IEEE* **12**(2), 16–26 (2014). DOI 10.1109/MSP.2013.137
97. de Oliveira, A.A.L., Camilo-Junior, C.G., Vincenzi, A.M.R.: A coevolutionary algorithm to automatic test case selection and mutant in mutation testing. In: *2013 IEEE Congress on Evolutionary Computation*, pp. 829–836 (2013)
98. Olsson, B.: Co-evolutionary search in asymmetric spaces. *Information Sciences* **133**(3–4), 103–125 (2001)
99. O'Neill, M., Ryan, C.: *Grammatical evolution: evolutionary automatic programming in an arbitrary language*, vol. 4. Springer (2003)
100. O'Reilly, U.M., Angeline, P.J.: Introduction to the special issue: Trends in evolutionary methods for program induction. *Evolutionary Computation* **5**(2), v–ix (1997)
101. O'Reilly, U.M., Erik, H.: An artificial coevolutionary framework for adversarial ai. In: *Adversary-aware Learning techniques and trends in Cybersecurity, AAAI Fall Symposium* (2018)
102. Ostaszewski, M., Seredynski, F., Bouvry, P.: Coevolutionary-based mechanisms for network anomaly detection. *Journal of Mathematical Modelling and Algorithms* **6**(3), 411–431 (2007)
103. Perez, D., Nicolau, M., O'Neill, M., Brabazon, A.: Evolving behaviour trees for the Mario AI competition using grammatical evolution. In: *European Conference on the Applications of Evolutionary Computation*, pp. 123–132. Springer (2011)

104. Pertier, M.: Investigating coevolutionary algorithms for expensive fitness evaluations in cybersecurity. Master's thesis, Massachusetts Institute of Technology (2018)
105. Petric, A.: Circular economy: a coevolutionary perspective on diversity. *uwf UmweltWirtschaftsForum* **24**(2), 253–260 (2016)
106. Pollack, J.B., Blair, A.D.: Co-evolution in the successful learning of backgammon strategy. *Machine Learning* **32**(3), 225–240 (1998)
107. Popovici, E.: Bridging supervised learning and test-based co-optimization. *Journal of Machine Learning Research* **18**(38), 1–39 (2017)
108. Popovici, E., Bucci, A., Wiegand, R.P., De Jong, E.D.: *Coevolutionary Principles*, pp. 987–1033. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
109. Popovici, E., Bucci, A., Wiegand, R.P., De Jong, E.D.: Coevolutionary principles. In: *Handbook of Natural Computing*, pp. 987–1033. Springer (2012)
110. Popovici, E., Winston, E.: A framework for co-optimization algorithm performance and its application to worst-case optimization. *Theoretical Computer Science* **567**, 46–73 (2015)
111. Prado Sanchez, D.: Visualizing adversaries - transparent pooling approaches for decision support in cybersecurity. Master's thesis, Massachusetts Institute of Technology (2018)
112. Rosin, C.D., Belew, R.K.: New methods for competitive coevolution. *Evolutionary Computation* **5**(1), 1–29 (1997)
113. Rush, G., Tauritz, D.R., Kent, A.D.: Coevolutionary agent-based network defense lightweight event system (CANDLES). In: *Proceedings of the Companion Publication of the 2015 on Genetic and Evolutionary Computation Conference*, pp. 859–866. ACM (2015)
114. Sagarin, R.D., Taylor, T.: Natural security: how biological systems use information to adapt in an unpredictable world. *Security Informatics* **1**(1), 14 (2012)
115. Sanchez, D.P., Pertier, M.A., Hemberg, E., O'Reilly, U.M.: Competitive coevolutionary algorithm decision support. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 300–301. ACM (2018)
116. Schmiedlechner, T., Al-Dujaili, A., Hemberg, E., O'Reilly, U.M.: Towards distributed coevolutionary gans. *arXiv preprint arXiv:1807.08194* (2018)
117. Scott, K., Davidson, J.: Strata: A software dynamic translation infrastructure. In: *IEEE Workshop on Binary Translation* (2001)
118. Scott Hilton: Dyn Analysis Summary Of Friday October 21 Attack. <https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/> (2016). (Accessed October 10, 2018)
119. Service, T., Tauritz, D.: Increasing infrastructure resilience through competitive coevolution. *New Mathematics and Natural Computation* **5**(02), 441–457 (2009)
120. Sims, K.: Evolving 3d morphology and behavior by competition. *Artificial life* **1**(4), 353–372 (1994)
121. Sipper, M.: *Evolved to Win*. Lulu. com (2011)
122. Sipper, M., Azaria, Y., Hauptman, A., Shichel, Y.: Designing an evolutionary strategizing machine for game playing and beyond. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* **37**(4), 583–593 (2007)
123. Smith, R.E.: Co-adaptive genetic algorithms: An example in othello strategy. In: *Proc. Florida Artificial Intelligence Research Symposium*, 1994 (1994)
124. Smith, R.J., Heywood, M.I.: Coevolving deep hierarchies of programs to solve complex tasks. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 1009–1016. ACM (2017)
125. Son of Boss: Son of boss — Wikipedia, the free encyclopedia. [https://en.wikipedia.org/wiki/Son\\_of\\_Boss](https://en.wikipedia.org/wiki/Son_of_Boss) (2018). (Accessed October 10, 2018)
126. Song, D., Heywood, M.I., Zincir-Heywood, A.N.: Training genetic programming on half a million patterns: an example from anomaly detection. *IEEE Transactions on Evolutionary Computation* **9**(3), 225–239 (2005)
127. Sood, A., Enbody, R.: Targeted cyberattacks: a superset of advanced persistent threats. *IEEE security & privacy* **11**(1), 54–61 (2013)
128. Stanley, K.O., Bryant, B.D., Miikkulainen, R.: Real-time neuroevolution in the nero video game. *IEEE Transactions on Evolutionary Computation* **9**(6), 653–668 (2005)



129. Suarez-Tangil, G., Palomar, E., de Fuentes, J.M., Blasco, J., Ribagorda, A.: Automatic rule generation based on genetic programming for event correlation. In: Computational Intelligence in Security for Information Systems, pp. 127–134. Springer (2009)
130. Symantec Security Response: Mirai: what you need to know about the botnet behind recent major DDoS attacks. <https://www.symantec.com/connect/blogs/mirai-what-you-need-know-about-botnet-behind-recent-major-ddos-attacks> (2016). (Accessed October 10, 2018)
131. Szubert, M., Jaskowski, W., Krawiec, K.: Coevolutionary temporal difference learning for othello. In: 2009 IEEE Symposium on Computational Intelligence and Games, pp. 104–111 (2009)
132. Szubert, M., Jakowski, W., Krawiec, K.: On scalability, generalization, and hybridization of coevolutionary learning: A case study for othello. *IEEE Transactions on Computational Intelligence and AI in Games* **5**(3), 214–226 (2013)
133. Talay, M.B., Calantone, R.J., Voorhees, C.M.: Coevolutionary dynamics of automotive competition: Product innovation, change, and marketplace survival. *Journal of Product Innovation Management* **31**(1), 61–78
134. Team, M.: Mininet - realistic virtual sdn network emulator. <http://mininet.org/> (2018). [Online; accessed 6-July-2018]
135. Togelius, J., Burrow, P., Lucas, S.M.: Multi-population competitive co-evolution of car racing controllers. In: 2007 IEEE Congress on Evolutionary Computation, pp. 4043–4050 (2007)
136. Weimer, W., Forrest, S., Le Goues, C., Nguyen, T.: Automatic program repair with evolutionary computation. *Communications of the ACM* **53**(5), 109–116 (2010)
137. Wikimedia Commons: Cuttlefish changing color. URL [https://upload.wikimedia.org/wikipedia/commons/thumb/1/1c/Cuttlefish\\_color.jpg/636px-Cuttlefish\\_color.jpg](https://upload.wikimedia.org/wikipedia/commons/thumb/1/1c/Cuttlefish_color.jpg/636px-Cuttlefish_color.jpg). Picture taken by Nick Hobgood - License: CC BY-SA 3.0
138. Wikimedia Commons: Misumena vatia with wasp (1998). URL <https://en.wikipedia.org/wiki/File:Misumena.vatia.beute.wespe.1771.jpg#filelinks>. Picture taken by Olaf Leillinger - License: CC-BY-SA-2.0/DE and GNU FDL
139. Wikimedia Commons: Viceroy butterfly (2005). URL [https://commons.wikimedia.org/wiki/File:Viceroy\\_Butterfly.jpg](https://commons.wikimedia.org/wiki/File:Viceroy_Butterfly.jpg). License: CC BY-SA 3.0. Subject to disclaimers.
140. Wikimedia Commons: Monarch in may (2007). URL [https://en.wikipedia.org/wiki/Monarch\\_butterfly#/media/File:Monarch\\_In\\_May.jpg](https://en.wikipedia.org/wiki/Monarch_butterfly#/media/File:Monarch_In_May.jpg). Created: 29 May 2007 By Kenneth Dwain Harrelson - License: CC BY-SA 3.0
141. Wikimedia Commons: Bioluminescence in ocean organisms (2014). URL [https://en.wikipedia.org/wiki/Bioluminescence#/media/File:Squid\\_Counterillumination.png](https://en.wikipedia.org/wiki/Bioluminescence#/media/File:Squid_Counterillumination.png). Picture taken by Chiswick Chap - License: CC BY-SA 4.0
142. Wilkerson, J.L., Tauritz, D.: Coevolutionary automated software correction. In: Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation, GECCO '10, pp. 1391–1392. ACM, New York, NY, USA (2010)
143. Willard, G.: Understanding the co-evolution of cyber defenses and attacks to achieve enhanced cybersecurity. *Warfare* **14**, 17–31 (2015)
144. Williams, N., Mitchell, M.: Investigating the success of spatial coevolution. In: Proceedings of the 7th annual conference on Genetic and evolutionary computation, pp. 523–530. ACM (2005)
145. Winterrose, M.L., Carter, K.M.: Strategic evolution of adversaries against temporal platform diversity active cyber defenses. In: Proceedings of the 2014 Symposium on Agent Directed Simulation, p. 9. Society for Computer Simulation International (2014)
146. Wood, B.J., Duggan, R., et al.: Red teaming of advanced information assurance concepts. In: DARPA Information Survivability Conference and Exposition, 2000. DIS-CEX'00. Proceedings, vol. 2, pp. 112–118. IEEE (2000)
147. Wright Jr, D.: Financial alchemy: How tax shelter promoters use financial products to bedevil the irs (and how the irs helps them). *Ariz. St. LJ* **45**, 611 (2013)
148. Xu, W., Qi, Y., Evans, D.: Automatically evading classifiers. In: Proceedings of the 2016 Network and Distributed Systems Symposium (2016)
149. Yuen, J.: Automated cyber red teaming. Tech. rep., DTIC Document (2015)

150. Zeng, F., Decraene, J., Low, M., Zhou, S., Cai, W.: Evolving optimal and diversified military operational plans for computational red teaming. *Systems Journal, IEEE* **6**(3), 499–509 (2012). DOI 10.1109/JSYST.2012.2190693

## Appendix A

### A.1 Firefly squid bioluminescence defense

When firefly squid are seen from below by a predator, the bioluminescence helps to match the squid's brightness and color to the sea surface above.

### A.2 A coevolutionary perspective of the U.S. automotive industry

Talay, Calantone, and Voorhees presented an study that explicitly terms competitive interactions between firms “Red Queen competition”, in which gains from innovations are relative and impermanent [133].

### A.3 Advanced Persistent Threats and Ransomware

Some DOS attacks includes Advanced Persistent Threats (APT) and Ransomware. The first ones have multiple stages starting at external reconnaissance then moving to intrusion (e.g. social engineering or use of zero day exploits), laterally moving malware, command and control direction to data exfiltration and, finally, self-erasure. Ransomware, which largely preys upon unpatched systems and which exploits anonymous payment channels enabled by Bitcoin, has also recently become more frequent.

### A.4 Cyber security attack categorization

Examples of attacks, classifications and taxonomies can be found at <https://cwe.mitre.org/index.html>. One categorization is: *A*) Advanced Persistent Threats, “lurking” threats from resourceful persevering adversaries *B*) Denial of Service Attack, defense resource limitation and exposure, means of penetrating systems *C*) Identity theft, e.g. impersonating users. Attacks are also characterized by their stages on a timeline. Another characterization is based on the attacker identity, from individuals to organized criminals and nation states, and what resources they access, see [21] for details.