

# Cyclic Incrementality in Competitive Coevolution: Evolvability through Pseudo-Baldwinian Switching-Genes

---

Rick Janssen<sup>\*,\*\*</sup>

Max Planck Institute  
for Psycholinguistics

Stefano Nolfi<sup>†</sup>

Institute of Cognitive Sciences  
and Technologies

Pim Haselager<sup>‡</sup>

Ida Sprinkhuizen-Kuyper<sup>‡</sup>

Donders Institute for Brain,  
Cognition, and Behaviour

**Abstract** Coevolving systems are notoriously difficult to understand. This is largely due to the Red Queen effect that dictates heterospecific fitness interdependence. In simulation studies of coevolving systems, master tournaments are often used to obtain more informed fitness measures by testing evolved individuals against past *and* future opponents. However, such tournaments still contain certain ambiguities. We introduce the use of a phenotypic cluster analysis to examine the distribution of opponent categories throughout an evolutionary sequence. This analysis, adopted from widespread usage in the bioinformatics community, can be applied to master tournament data. This allows us to construct behavior-based category trees, obtaining a hierarchical classification of phenotypes that are suspected to interleave during cyclic evolution. We use the cluster data to establish the existence of *switching-genes* that control opponent specialization, suggesting the retention of dormant genetic adaptations, that is, genetic memory. Our overarching goal is to reiterate how computer simulations may have importance to the broader understanding of evolutionary dynamics in general. We emphasize a further shift from a component-driven to an interaction-driven perspective in understanding coevolving systems. As yet, it is unclear how the sudden development of switching-genes relates to the gradual emergence of genetic adaptability. Likely, *context genes* gradually provide the appropriate genetic environment wherein the switching-gene effect can be exploited.

---

## Keywords

Coevolution, cyclic evolution,  
computational modeling, cluster analysis,  
switching-genes

---

## 1 Introduction

The study of the coevolution of competitive pursuit and escape behavior has attracted the interest of several researchers in the area of robotics and artificial life [29]. This can be explained by considering

---

\* Contact author.

\*\* Language and Genetics Department, Max Planck Institute for Psycholinguistics, The Netherlands. Affiliated with Radboud University Nijmegen when this study was conducted. E-mail: rick.janssen@mpi.nl

† Laboratory of Autonomous Robotics and Artificial Life, Institute of Cognitive Sciences and Technologies, National Research Council, Italy. E-mail: stefano.nolfi@istc.cnr.it

‡ Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, The Netherlands, and Department of Artificial Intelligence, Radboud University Nijmegen, The Netherlands. E-mail: w.haselager@donders.ru.nl (P.H.); i.kuyper@donders.ru.nl (I.S.-K.)

that pursuit and evasion behaviors do not only represent one of the most common and challenging problems for natural organisms but also provide an ideal challenge for robotics and embodied cognition. Indeed, the need to face highly dynamic, largely unpredictable, and hostile environments requires the development of fast, robust, and adaptable solutions [27, 9].

In this article, we report a series of coevolutionary computer simulations in which two populations of predator and prey robots were evolved for their ability to pursue and escape each other. We use a phenotypic cluster analysis to characterize coevolutionary progress.<sup>1</sup> The obtained results confirm previous findings that the coevolutionary dynamics converge toward cycles in which specific classes of behavioral strategies alternate. However, compared to previous studies, our results were obtained in a more automated and formal fashion. The obtained results indicate that evolving robots converge toward solutions that are effective against strategies displayed by the competitors. In addition, the evolved robots display a certain readiness for change in order to cope with variant strategies that competitors are likely to exhibit in future generations. This is achieved through the synthesis of genetic organizations characterized by *switching-genes*, that is, genes that enable rapid shifts between different strategies through a few single point mutations. Finally, we demonstrate that such switching-genes gradually evolve when species are exposed to an environment wherein challenges are periodically recurring.

For our study, we make use of a class of evolutionary algorithms (EAs) known as *coevolutionary algorithms* (CoEAs). These algorithms evolve multiple species in a shared environment [9]. Thus, they can be used to approximate the natural situation more accurately than regular EAs. However, their application necessitates addressing two important issues. First, since a species' fitness in a CoEA is interdependent with other species, measuring progress is more difficult than when using a regular EA [8]. Previously, *master tournaments* (MTs; Section 2.2) were used to obtain more reliable fitness estimates. But, while they certainly add informativeness, MTs also exhibit some ambiguities [7] (cf. Section 2.2). Secondly, in nature, competitively coevolving systems are hypothesized to lead to *evolutionary arms races* [11] that represent an important drive for change and innovation in evolution [16]. However, simulations with CoEAs do not a priori lead to the emergence of these arms races in all cases. Instead, *cycling dynamics* are often observed (Section 2.2) [13, 30].

When discussing complex systems such as those showing cycling dynamics, meta-analyses are invaluable. A concept that is particularly useful is that of the *phenotype*. In our study, we describe the phenotype as the set of behavioral traits that lead to the MT fitness score of a single individual (Sections 2.3 and 3.2). If a species (hereafter the *subject species*) is confronted with an opponent species that is cycling between behavioral categories (hereafter *phenotypic categories*), it could be beneficial for the subject species to be able to quickly adapt to such category alternations, for example, by cycling between phenotypes as well. If such rapid adaptations were to be purely genetic in nature (excluding, e.g., developmental effects such as neural plasticity), a reliance on merely a few switching-genes that express a specialization from one opponent to the next would be more efficient than having a dependence on numerous ones. This is because such genes would enable quick specialization shifts.

Even when switching-genes might enable a species to increase its evolvability (i.e., increase adaptive genetic predispositions) as evolution progresses, it does not immediately follow that they would evolve abruptly. In fact, the emergence of any *pseudo-Baldwinian* genetic predisposition to adaptability expressed by switching-genes might require long-term exposure to a considerable number of opponent cycles (Section 2.4). Unfortunately, since any genetic change might only be expressed in the phenotype (i.e., become behaviorally differentiable) against some specific opponent categories, identifying switching gene evolution is not a simple task. As a solution, we propose to use a phenotypic cluster analysis (Section 3.3) to examine the distribution of opponent categories throughout an evolutionary sequence. Following this, we can investigate which particular subject genes are expressed against which particular opponent categories (Section 3.4). In other words: If temporal fluctuations in certain genes and in opponent specialization in the phenotype strongly correlate, we have to consider the existence of switching-genes.

---

<sup>1</sup> We emphasize the use of the term "phenotypic" here, that is, relating to observable traits. Thus, we do not use the cluster analysis to reconstruct phylogenies. Instead, we use it as a tool to group individuals that share similar behavioral (phenotypic) traits.

Somewhat problematically, even with the clustering technique we propose, it remains challenging to exercise enough control over evolutionary dynamics to answer the questions we just posed. This is because fitness interdependence between coevolving species can give rise to chaotic interactions [1]. Here, the cluster analysis offers additional utility. If we extract an opponent's categories from a previously coevolved scenario, we can present these in alternating intervals and let a new subject species evolve against them. Effectively, we use the extracted categories to serve as an *emulated cycling* opponent (Section 3.5) so we can more precisely control which opponent categories are encountered by the subject species, and at what time intervals. Thus, we can systematically manipulate a fundamental variable that codetermines the selective pressure to evolve switching-genes.

The remainder of this article is organized as follows: Section 2 will give a more thorough theoretical background on the analytical techniques used with CoEAs, the biological plausibility of Baldwinian evolution and cyclic evolution, and the concept of the phenotype. Section 3 discusses a novel interpretation of CoEA analytics and introduces the formal concept of the *functionally extended phenotype*. It also examines how this provides a conceptual justification for the cluster analysis and its derivative techniques. Section 4 demonstrates the results obtained from applying these techniques to a case study (detailed in Appendix 1) and discusses multiple interpretations. Finally, Section 5 provides a general conclusion regarding the findings.

## 2 Background

### 2.1 Evolutionary Robotics

Evolutionary robotics (ER) is an engineering approach where autonomous agents are developed following a *proximal* (i.e., from the agent's) point of view [40]. This perspective describes agents in terms of internal states (e.g., neural activation), and thus ER is very much an approach that takes a dynamical systems outlook on agent design [4]. More specifically, while an agent can have very simple internals, it is able to give rise to complex behavior through a sensorimotor loop. The agent's behavior thus has to be understood as its being situated in its environment, possibly interacting with other agents. The behaviors that are then displayed are said to emerge out of these interactions. In contrast, an agent can also be regarded from a *distal* perspective [40], which involves the use of behavioral terms such as "approaching" or "evading," or even of psychological ones such as having "intentions" or "desires" [23]. The aim is to understand the system by interpreting it as a collection of components [4]. However, while useful in their own right, distal descriptors can impose unwarranted constraints or predispositions on robot design.

A similar multiplicity of epistemological descriptive levels is also seen in many other fields, like animal ethology. Animals can, for instance, be described on a suborganism (e.g., genes, neurons, organs), organism (e.g., behaviors, intentions) or superorganism level (e.g., populations, societies, ecosystems) level [34]. As we will emphasize in this study, ER not only offers a promising approach to study natural evolution that highlights the proximal perspective, but also draws a parallel with the suborganism and, through our analysis (Section 3), the superorganism level. Surprisingly, then, while there is thus significant overlap between the field of ER and the biological sciences, ER makes little use of tools developed by the bioinformatics community. For instance, cluster algorithms have found widespread application, with demonstrated usefulness when addressing phylogenetics, grouping organisms on genetic similarity, and thereby deriving ancestral lineages [47]. Analyses like these might prove valuable to ER and might be used to overcome certain ambiguities (Section 2.2).

Altogether, further formalization and automation by incorporating techniques from bioinformatics into ER is a logical step in the attempt to understand dynamical systems through a proximal, interaction-driven perspective. More specifically, as we will demonstrate, autonomous agents can develop through the process of self-organization that the CoEA provides (Appendix 1), while they are analyzed by a cluster analysis and derivative techniques (Section 3). Thus, both agent development and analysis depend less on subjective interpretations.

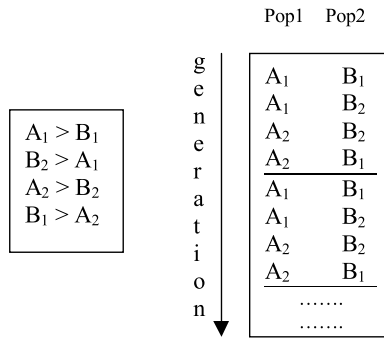


Figure 1. Systematization of limit cycle dynamics. The same strategies ( $A_1$  and  $A_2$  in population A, and  $B_1$  and  $B_2$  in population B) may be selected over and over again throughout generations. Strategy  $A_1$  enables population A to outperform competitors displaying strategy  $B_1$  but not  $B_2$ .  $B_2$  on the other hand enables population B to outperform competitors displaying strategy  $A_1$  but not  $A_2$ . (From [30].)

### 2.2 Long-Term Coevolutionary Dynamics

Competitive coevolution could spontaneously lead to a form of self-sustained incremental process that in turn might cause the coevolutionary process to outgrow standard evolutionary ones. Indeed, according to some authors, the establishment of *arms races* between and within species could be one of the main sources of evolutionary innovation in nature [45, 11, 37].

However, a continuous increase in complexity is not guaranteed. Computer simulations show that coevolving populations may in fact drive one another along twisting pathways where each new solution is just good enough to counterbalance the current strategies discovered by the opponent species, but is not necessarily more and possibly even less effective than solutions discovered some generations earlier. Thus, species are often found to be evolving cyclically: continuously resorting to previously discovered and then discarded strategies, without any apparent long-term incrementality arising (Figure 1) [13, 30].<sup>2</sup>

Another difficulty in competitive coevolution concerns the intrinsic complexity of the evolutionary dynamics that makes the analysis challenging. In coevolving populations, changes in one species affect the reproductive value of specific trait combinations in the other species. This effectively corresponds to a modification of the fitness landscape. It might thus happen that progress achieved by one lineage is reduced or eliminated by the competing species. Such bilaterally exerted selective pressure between species might lead to a reciprocal feedback loop, which is referred to as the *Red Queen effect* [45, 37]. This fitness interdependence makes it hard to monitor progress by using conventional indicators: Apparent oscillations of fitness values throughout generations might hide true progress, and periods of stasis might correspond to tightly coupled coevolutionary changes in both species.

Due to their inherently complex interaction, simulations on coevolving systems are often analyzed using MTs (Section 3). However, MTs are mainly used to *demonstrate* cycling or incrementality (e.g., [14]). To *understand* these dynamics, on the other hand, one has to address the variables influencing evolutionary progress, such as sensorimotor constraints, environmental richness, and ontogenetics (developmental effects) [29], as well as the feedback loop between genes and the phenotype.

### 2.3 Phenetics

Our focus on the phenotype and the use of the concept itself (Sections 1 and 3.2) might appear unconventional. Historically however, the term “phenotype” has had different interpretations, depending on the field of study as well as on the time frame. In the context of ER, phenotypes

<sup>2</sup> The term “cyclic evolution” does not necessarily imply a strictly circular, linear recurrence of strategies. The precise order in which strategies (or variants thereof) are revisited is not of prime importance for our study. What matters is that the rediscovery of historical phenotypes can be a valid alternative to a progressive refinement or accumulation of traits. Thus, we use the term “cyclic evolution” primarily for the sake of consistency with the existing literature.

are associated with the robot's body and/or controller (i.e., the actual instantiation of those parameters subjected to the EA). This parallels the use of the term in its originating, biological domain where it was used to distinguish it from the Mendelian genotype [21]. As such, the genotype was historically considered the unit that would be directly transmitted in ancestral heritage, while the phenotype was understood as the mere expression of the genotype.

However, the binary classification between genotype and phenotype has long been recognized to be rather limited and arbitrary, and the concept of the phenotype is nowadays considered from different (epistemological) viewpoints. For example, a phenotype that is closer to actual genetics would be the *endophenotype* that describes, for example, markers for heritable diseases not necessarily apparent as salient, conspicuous symptoms [22]. A simple example like this illustrates that the concept of a well-defined, fundamental level at which "the" phenotype resides is an oversimplification, as the term is very much context-dependent and should be considered to follow a continuum. After all, genes can be expressed in different domains, such as the (sub)cellular (e.g., RNA, proteins, cells), functionally descriptive (e.g., aerobic capacity, cardiac rhythm), or behaviorally characteristic [32].

As a product of genes and their interaction with their environment (cf. epigenetics), phenotypes can even be said to extend beyond a single organism. For example, a (hypothetical) termite's gene that would result in the behavioral inclination to build termite mounds would improve the survivability of that gene, regardless of the individual termite it resides in. Thus, the realization of a termite mound should be considered an expression of an *extended* phenotype of the whole collection of homologous, mound-building genes in an entire termite colony [10]. We emphasize that, in order to understand cyclic evolution, it is vital to pay specific attention to such extended phenotypic features. More specifically, we will regard the phenotype from a functionally extended perspective, instead of restricting ourselves to mere physicality. For this reason, we use the term "phenotype" to indicate a species' capability to evolve individuals to cope with rapidly alternating opponent behavior (Section 3.2).

## 2.4 Genetic Predispositions to Environmental Variation

Ontogenetics (i.e., developmental effects within an organism's lifetime) add a layer of complexity that is hard to dismiss when evaluating computer models on their biological plausibility. While there are varying interpretations of the specifics of the *Baldwin effect* [2], it can be, simply put, bisected into two phases [44]: First, species might evolve *phenotypic plasticity* (e.g., callosity, ectothermism, or intelligence), enhancing their adaptability and increasing their survivability. This would provide a degree of flexibility that genetics alone cannot, allowing exploration of ecological niches that would take generations for genetics. Following this, traits that are ontogenetically acquired and that have demonstrated utility might be genetically *assimilated* (i.e., a species might become so predisposed to acquire a particular trait that it comes to appear innate). Thus, fitness is increased at the cost of reduced flexibility.

We, instead, hypothesize a conceptual variant of the Baldwin effect that operates at the genetic level only and that consists in the predisposition to cope with environmental variations through minimal genetic change. This might be realized through the selection of genotypes that include switching-genes (cf. Section 1) that can enable significant phenotype reorganization. In this case, no phenotypes are assimilated, because all change is, fundamentally, genetic in the first place.

It is promising that the assimilation component of the (ontogenetic) Baldwin effect (cf. [46]) has been reproduced in computer simulations by using regular (i.e., not coevolving) EAs [18]. More recent studies, however, stress a dependence of assimilation on the rate of genetic as well as environmental change [3]. Finally, and more concretely, the development of ontogenetics in coevolving predator-prey simulations has been demonstrated in [15]. Thus, altogether, we are confident in hypothesizing that a purely genetic variant of the Baldwin effect should be reproducible in CoEAs as well. If so, this would provide an opportunity to explore the existence of switching-genes and how they might relate to cyclic evolution in a computational model. Models such as these have already received considerable academic interest, starting with the Lotka-Volterra equations [26].

In nature, cycling can be observed on what we classify as three distinct levels: A well-established example of (conspecific) *population size* cycling is found in the common side-blotched lizard,

*Uta stansburiana* [41, 42], structurally comparable to heterospecific cycling in CoEAs: *ultra-dominants* seduce females from *dominants*, dominants are adept at defending their females from *sneakers*, and sneakers in turn steal mates from ultra-dominants. A similar example of (heterospecific) population size cycling can be found in antibiotic production by different strains of the enteric bacterium *Escherichia coli* [25]. Interestingly, these patterns were first predicted by computer models before being confirmed in vitro [24].

*Hybrid cycling* (having a genetic component in addition to population size cycling) was already acknowledged as a plausible outcome of arms races when they were first popularized [11]. For example, species of the *Anolis* lizard (found dispersed over the Caribbean islands) that are bigger possess higher fitness, since they can eat larger insects, but at the cost of having to sustain a large body mass [39, 6]. Consequently, on islands that are only inhabited by one species, the lizards would only have a modest, solitary size. However, when a larger species invades such an island, the solitary-sized species is outcompeted by the invaders, being forced to move into a different ecological niche by reducing its size, followed by size reduction in the invaders. Eventually, the original, native species goes extinct and the invaders take its place, to await a similar turn of events.

Finally, an example of *genetic cycling*, mostly grounded on computational modeling and less established in vitro, can be found in parent-offspring conflict [36, 35]. In this scenario, offspring might evolve *conflictor genes* that exploit parental attention (e.g., birds that display begging behavior). Parents would in response evolve *suppressor genes* that enable them to conserve energy by ignoring conflictor gene expressions (and divide attention equally among offspring). The conflictor genes would then disappear (since they are costly and no longer of use), followed by the disappearance of the suppressor genes, after which the cycle repeats.

We postulate that conflictor and suppressor genes should be regarded as examples of genetic dispositions to evolvability, as an instance of the genetic pseudo-Baldwinism we discussed earlier. More specifically, conflictor and suppressor genes would evolve in order to cope with parent and offspring expressions, respectively. Since quick adaptations to opponent changes are beneficial, this suggests that a small number of conflictor or suppressor genes is advantageous compared to having to rely on a large number of them. Thus, conflictor and suppressor genes can be considered instances of the more general class of switching-genes (Section 1). In the case of hybrid cycling, however, switching-genes would of course be less likely to emerge, since there is no genetic continuity in the cycling dynamic (since species go extinct).

### 3 Experiment and Analysis

#### 3.1 Overview

We evolved two populations of predator and prey robots (Section A1.1) for the ability to pursue and escape each other, respectively. The robots were controlled by artificial neural networks (NNs; Section A1.2) that allowed them to map input from the visual camera and infrared sensors to wheel speed activation. We used a CoEA (Section A1.3) to adjust the connection strengths of this NN, thereby changing the robot's behavioral response to environmental stimuli. The fitness of each individual was evaluated while it was allowed to interact with all opponents of the current generation, one at a time. We replicated the evolutionary process several times, starting from different randomly initialized genotypes (hereafter *seeds*). We realized the simulations using the Evorobot\* tool [31]. More details on the implementation can be found in Appendix 1.

We conducted two sequential series of experiments. The first one is the *genuinely coevolving experiment* (GENU), where both predator and prey were subjected to the evolutionary process as simulated by the CoEA. Subsequently, in a series of *emulated coevolving experiments* (EMUs), we evolved only a population of predators against fixed prey displaying selected behavioral strategies. In the EMUs, we extracted the prey individuals from the GENU run earlier.

Parameters were determined in a pilot study so that the evolutionary process converged to an equilibrium in terms of fitness values obtained [20]. Certain parameters were fixed on account of

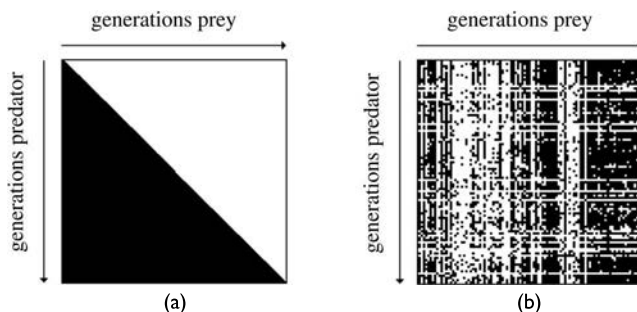


Figure 2. MT data. Each pixel represents the outcome of a post-evaluation test between the best predator and prey of two corresponding generations (black and white pixels indicate trials won by the predator or the prey, respectively) (from [30]). (a) In a prototypical arms race, an MT plot would show a diagonal bisection. This should be interpreted as each subject elite from generation  $g$  outperforming all opponent elites from previous generations  $\{g' : g' < g\}$ . (b) An actual example of a limit cycle dynamic, visible as distinct rectilinear banding, obtained from running a coevolution computer simulation.

theoretical considerations. For instance, in [9] the morphology of a robot’s visual sensors’ position and angle of view was evolved. This showed that predators evolved anterior eyes, with narrow angle of view, while prey evolved lateral eyes, with wide angle of view. This parallels what we often see in nature. Because morphological evolution was not the focus of our study, we fixed the robots’ visual sensors in a configuration that we deemed realistic.

In the context of artificial coevolution, where all the data are available to the researcher, specific measures have been proposed to evaluate the coevolutionary dynamics. In particular, [27] proposed to monitor progress by testing the performance of the best individual at each generation against all the best competing ancestors. The measurements obtained in this way have been called *CLAO* (current individual versus ancestral opponents) data. An extended version of this measure, in which the best individual of each generation is tested against the best competitors of past *and* future generations, is known as a *master tournament* (MT) [13].

MT data provide hints on whether variations occurring during a certain evolutionary phase represent real progress (i.e., enable individuals of successive generations to successfully cope with a larger number of competitors) or they enable individuals to gain advantage over the current competitors only. In this way, they provide indications of whether the coevolutionary process is characterized by a continuous refinement of traits (Figure 2a) or by a cyclic dynamic (Figure 2b).

CLAO and MT analyses, however, can only give indirect indications on evolutionary dynamics, since they do not provide a way to identify qualitatively different classes of strategies and counterstrategies. Thus, we propose to use a phenotypic cluster analysis (Section 3.3) to partition MT data so that the strategies displayed by the two species can be categorized into functionally different classes (where “functionally” refers to the ability of these strategies to cope with counterstrategies of specific classes). The cluster analysis can be used to identify the number and type of classes and subclasses in which the evolved behaviors are organized. In doing so, we can construct behavior-based “family trees” (hereafter, *category trees*<sup>3</sup>). We can then examine which (if any) genes are in control of alternating between these categories (Section 3.4).

In Section 3.2, we introduce a definition of MT data that demonstrates a species-dependent perspective on the MT (first informally raised as “subject” and “opponent” species in Section 1). We define the notion of a phenotype in the context of MT data. We also introduce the concepts of *genosequences* and *phenosequences*. We then use these definitions to measure phenotypic distances to be used in the cluster analysis (Section 3.3). Subsequently, we explain how to use the acquired clusters

<sup>3</sup> We use the term “category tree” to avoid confusion with phylogenetic family trees, which our cluster analysis does not address.

to investigate correlations between genosequences and phenosequences in order to examine the existence of switching-genes (Section 3.4). Finally, we demonstrate how to use the same clusters to set up the EMUs, where we simulate a coevolutionary situation but where we are able to exert more control over long-term evolutionary dynamics (Section 3.5).

### 3.2 Perspectives on Master Tournament Data

We introduce the notation

$$\langle S, O \rangle = \begin{cases} \langle PD, PY \rangle & \text{if predator is subject} \\ \langle PY, PD \rangle & \text{if prey is subject} \end{cases} \quad (1)$$

as a perspective-dependent but species-neutral reference (cf. Section 1). Following this, we denote nonspecific individuals as  $s \in S$  and  $o \in O$ .

We ran the GENU for 500 generations (Section A1.3), thus producing  $|S| = 500$  elites with which to organize a MT (Section 2.2). We can consider the genetic data representing  $S$ 's elites a sequence of  $|S|$  genotypes, each instantiating 28 genes (Section A1.3), denoted as an  $|S| \times |\mathbf{g}|$  matrix  $G^S$  (Equation 2). Thus,  $\mathbf{g}_{s,\gamma}^S$  (where  $\gamma = 1, 2, \dots, |\mathbf{g}|$ ) denotes the allele corresponding to the  $\gamma$ th gene of the  $s$ th individual of species  $S$ , and we have

$$G^S = \begin{pmatrix} g_{1,1}^S & g_{1,2}^S & \dots & g_{1,|\mathbf{g}|}^S \\ g_{2,1}^S & g_{2,2}^S & \dots & g_{2,|\mathbf{g}|}^S \\ \vdots & \vdots & \ddots & \vdots \\ g_{|S|,1}^S & g_{|S|,2}^S & \dots & g_{|S|,|\mathbf{g}|}^S \end{pmatrix} \quad (2)$$

Here we genetically describe a single individual  $s \in S$  with a *genotype*: a vector containing  $|\mathbf{g}|$  alleles:

**Definition 1.** *Genotype*:  $\mathbf{g}_{s,*}^S = \langle g_{s,1}^S, g_{s,2}^S, \dots, g_{s,|\mathbf{g}|}^S \rangle$ .

Over the course of evolution, a particular gene can mutate from one generation to the next, in what we can consider a genosequence: a vector denoting a gene  $\gamma \in \mathbf{g}$ 's temporal sequence of alleles over  $|S|$  generations:

**Definition 2.** *Genosequence*:  $\mathbf{g}_{s,\gamma}^S = \langle g_{1,\gamma}^S, g_{2,\gamma}^S, \dots, g_{|S|,\gamma}^S \rangle$ .

Similarly, we represent the data of a MT between two competing species by an  $|S| \times |O|$  matrix  $M^S$ :

$$M^S = \begin{pmatrix} m_{1,1}^S & m_{1,2}^S & \dots & m_{1,|O|}^S \\ m_{2,1}^S & m_{2,2}^S & \dots & m_{2,|O|}^S \\ \vdots & \vdots & \ddots & \vdots \\ m_{|S|,1}^S & m_{|S|,2}^S & \dots & m_{|S|,|O|}^S \end{pmatrix} \quad (3)$$

When we would shift perspective (i.e., when switching subject and opponent assignment as declared in Equation 1), we transpose and modify values as in  $M_{ji}^O = 1 - M_{ij}^S$ .

Comparably to how we describe an individual  $s \in S$  as a genotype (Definition 1), we can describe it phenetically as a *phenotype*, consisting of  $|O|$  *phenes* that each correspond to a certain fitness score against an opponent in the MT:

**Definition 3.** *Phenotype*:  $\mathbf{m}_{s,*}^S = \langle m_{s,1}^S, m_{s,2}^S, \dots, m_{s,|O|}^S \rangle$ .



Just as we describe temporal variations in genes as a genosequence (Definition 2), we can likewise describe phenes as a *phenosequence*.<sup>4</sup> Thus, whereas we consider an individual  $s \in S$ 's phenotype (Definition 3) as a sequence of different phene values against all MT opponents, we consider its phenosequence as the complete sequence of temporal changes over  $|S|$  generations in one particular phene (i.e., against one particular opponent individual  $o \in O$ ):

**Definition 4.** *Phenosequence:*  $\mathbf{m}_{s,o}^S = \langle m_{1,o}^S, m_{2,o}^S, \dots, m_{|S|,o}^S \rangle$ .

To reiterate, the definitions of phenotypes and phenosequences vary according to which species' (i.e., the subject's or opponent's) perspective the MT is interpreted from. The reason for this is that, when considering genotypes and genosequences, no ambiguities in perspective arise, since each species has a unique genetic code, whereas both species share a single phenetic "code" (i.e., MT data). For instance, what would be a predator individual's phenotype would be the prey species' phenosequence of one particular phene (and vice versa). This consideration thus provides the conceptual justification of correlating between subject genosequences and opponent phenotypes (i.e., subject phenosequences), and subsequently averaging such correlations over opponent category distributions (Section 3.4).

To conclude, it follows from Equation 3 and Definitions 3 and 4 that the MT data contain a description of all the phenotypes and phenosequences that were displayed during the tournament. This, first of all, allows us to cluster phenotypes on the basis of MT data (Section 3.3). Secondly, if switching-genes exist, there should be patterns in the relations between genosequences of certain genes (i.e., temporal changes in alleles) and phenosequences of certain phenes (i.e., temporal changes in terms of fitness scores against particular opponent individuals). More concretely, a change in switching gene should be accompanied by a marked change in phenotype. Thus, correlating between genosequences and phenosequences (Section 3.4) should enable identification of switching-genes.

### 3.3 Cluster Analysis

The cluster analysis we applied in this study (Algorithm 1) is based on the unweighted pair group method with arithmetic mean (UPGMA) algorithm, an agglomerative (i.e., bottom-up) hierarchical clustering algorithm [43] that allows building well-informed, phenotypic category trees. Agglomerative cluster algorithms construct hierarchies by iteratively grouping the two most similar clusters in a collection until all clusters are grouped under one big root cluster, thereby forming a binary tree. In it, each node in the phenotypic tree represents an (abstract) category of individuals, and each leaf (i.e., singleton cluster) represents a (concrete) individual, that is, a phenotype (Definition 3).

The cluster algorithm (Algorithm 1) used a distance matrix to keep track of newly formed clusters and their similarity to the other known clusters. The algorithm works by iteratively merging the two closest clusters into a new cluster. We calculated any newly formed cluster's unweighted average (Equation 5 below) and updated the distance matrix. We simultaneously stored cluster representations in a binary tree data structure. We continued this process until every cluster is grouped under one, all-encompassing root cluster.

Species  $S$ 's categories are initially represented by singleton clusters (i.e., concrete individuals). The dissimilarity between two such singleton clusters  $\mathbf{m}_{r,*}^S$  and  $\mathbf{m}_{r',*}^S$  (Definition 3) is equivalent to their Euclidean distance to each other,

$$\text{dist} \left( \mathbf{m}_{r,*}^S, \mathbf{m}_{r',*}^S \right) = \sqrt{\sum_{o=1}^{|O|} \left( \mathbf{m}_{r,o}^S - \mathbf{m}_{r',o}^S \right)^2} \tag{4}$$

---

<sup>4</sup> While the phenes in the simulation are not independent entities, they require independent measurements in view of the dynamical systems nature of the experiment. More specifically, while the simulation is a fully deterministic system, it would be practically impossible to predict what phenotypes emerge from what genotypes.

---

**Algorithm 1.** The UPGMA-based cluster algorithm.

---

```

MasterData ← loadMasterData();
foreach Species do
    Tree ← newEmptyTree();
    DistanceMatrix ← initDistanceMatrix(MasterData,Species);
    while sizeof(DistanceMatrix > 1) do
        // Euclidean metric used
        ClusterA,ClusterB ← removeClosestClusters(DistanceMatrix);
        // averaging measure used
        NewCluster ← mergeClusters(ClusterA,ClusterB);
        addCluster(NewCluster,DistanceMatrix);
        addCluster(NewCluster,Tree);
    
```

---

Distance computation between compound clusters necessitated the formulation of an appropriate averaging measure that specifies how nested clusters can be represented by a single vector to be used in Equation 4. Let  $(S' \cup S'') \subseteq S$  be a compound cluster (i.e., a phenotypic category), where  $S'$  or  $S''$  can either represent another (nested) compound cluster or a singleton cluster  $\mathbf{m}_{s,*}^S$  (i.e., a concrete individual). The averaging measure for a compound cluster is now formulated as follows:

$$\mu_u(S' \cup S'') = \frac{1}{2} \left( \frac{1}{|S'|} \sum_{s \in S'} \mathbf{m}_{s,*}^S + \frac{1}{|S''|} \sum_{s \in S''} \mathbf{m}_{s,*}^S \right) \tag{5}$$

Conceptually, this unweighted average reflects the idea that categories’ phenotypic profiles are equally important when combined to form a new supercategory, regardless of the subcategories’ size.

To visually draw the phenogram, we converted the binary tree to a Newick formatted string [12], which we then visualized using the Environment for Tree Exploration (ETE) toolkit [19]. We inserted small average-fitness histograms at the tree’s nodes. This allows for quick observation of the overall performance of the category corresponding to that node against the opponent. For this visualization, we are however interested in the weighted performance rather than the unweighted (since a category’s fitness performance is an average over all its members, not solely a binary one of its two subcategories):

$$\mu_w(S' \cup S'') = \frac{1}{|S'| + |S''|} \left( \sum_{s \in S'} \mathbf{m}_{s,*}^S + \sum_{s \in S''} \mathbf{m}_{s,*}^S \right) \tag{6}$$

### 3.4 Geno-Pheno Correlations

#### 3.4.1 Identification

We denote the Pearson product-moment correlation coefficient (PCC; cf. [38]) between a subject’s genosequence and phenosequence (i.e., opponent phenotype) as  $\rho(\mathbf{g}_{*,\gamma}^S, \mathbf{m}_{o,*}^O)$ . To investigate whether switching-genes might exist, we average the correlations between a single genosequence and a whole set of subject phenosequences (i.e., category of opponent phenotypes)  $O'$ , following

$$\rho_\mu(\mathbf{g}_{*,\gamma}^S, O') = \frac{1}{|O'|} \left( \sum_{o \in O'} \rho(\mathbf{g}_{*,\gamma}^S, \mathbf{m}_{o,*}^O) \right) \tag{7}$$

Importantly, note that while the aim is to measure a correlation between subject genosequences and phenosequences, the averaged PCC calculation in Equation 7 uses the opponent matrix  $M^O$ . This is

because, since a subject phenosequence  $\mathbf{m}_{*o}^S$  is equivalent to the opponent phenotype  $\mathbf{m}_{o*}^O$  (Section 3.2), cluster data on the opponent species (Section 3.3) can be justifiably used to determine which subject genes got strongly expressed as which category of subject phenosequences (i.e., category of opponent phenotypes). Since we use the averaging measure (7) for exactly this purpose (Section 4.3.1), the notation is of pragmatic, explicit origin.

### 3.4.2 Transplantation

Once we suspect a (set of) switching gene(s)  $\mathbf{g}' \subseteq \mathbf{g}^S$  to control the subject category  $S'$ 's specialization, we can extract the corresponding alleles  $\{\mathbf{g}_{s,\gamma}^S : s \in S', \gamma \in \mathbf{g}'\}$  and then randomly distribute those over the complete conspecific population  $S$  (i.e., stochastically varying the originating individual  $s$ , each time copying  $s$ 's complete set of alleles).

If switching-genes indeed control opponent specialization, the associated allele's transplantation should force the entirety of  $S$  to assume the category  $S'$ 's specialization, regardless of pre-transplantation category membership. Conversely, if we distribute non-switching alleles, this should not result in any switch in specialization.

### 3.5 Opponent Cycling Emulation

Results show that genetic adaptability, possibly controlled by switching-genes, appears to emerge slowly as the subject species is exposed to an increasing number of opponent category cycles (cf. Sections 4.1.2 and 4.3.2). To investigate, we extracted the opponent categories against which switching-genes got strongly expressed in the preceding GENU. We randomly sampled opponents from these categories and presented those in alternating fashion to the subject species to evolve against *de novo*. We switched the opponent categories every 75 generations (an approximation of the cycling interval in the GENU) for a total of 1000 generations (Section 4.1.1).

For example, two extracted opponent categories  $O' \subseteq O$  and  $O'' \subseteq O$  could be presented in the alternation sequence  $\langle O', O'', O', O'', \dots \rangle$ . Each cycle (i.e., every 75 generations), the opponent category is switched. If adaptability in the subject species increases throughout evolution, there should be an increase in fitness performance against both categories. More specifically, the difference in either of the opponent categories' consecutive cycles' fitness maxima and minima should increase as evolution progresses (Figure 3). Having fixed, alternating categories served to emulate cycling dynamics in the opponent species, thereby gaining more control over selective pressure on the subject species. Moreover, we now used the Gaussian mutator ( $\mu = 0, \sigma = 4$ ) for the subject mutations instead of the random bitwise one used in the GENU (Section A1.3). Again, this was done in order to attain a higher degree of control over evolutionary progress and subsequent exploration of the geno space.

## 4 Results and Discussion

We will demonstrate the results following the methods elaborated on in Appendix 1 and Section 3. Our focus is mainly on one experiment replication (seed 6), since it shows promising initial results using the cluster analysis and provides a clear example of cycling dynamics.<sup>5</sup> First, we briefly discuss the classical, online fitness measurements and MT data in Section 4.1.1. These raw data are then used to obtain the cluster analysis' category classifications in Section 4.2.1. In Section 4.3.1, we examine the correlations between subject (predator) genosequences and phenosequences, mapped on the previously obtained opponent (prey) categories, and their influence on opponent specialization in situ. Finally, in Section 4.4.1, we will address the emergence of genetic adaptability by emulating opponent cycling.

<sup>5</sup> For some suggestions on generality, see Appendix 2.

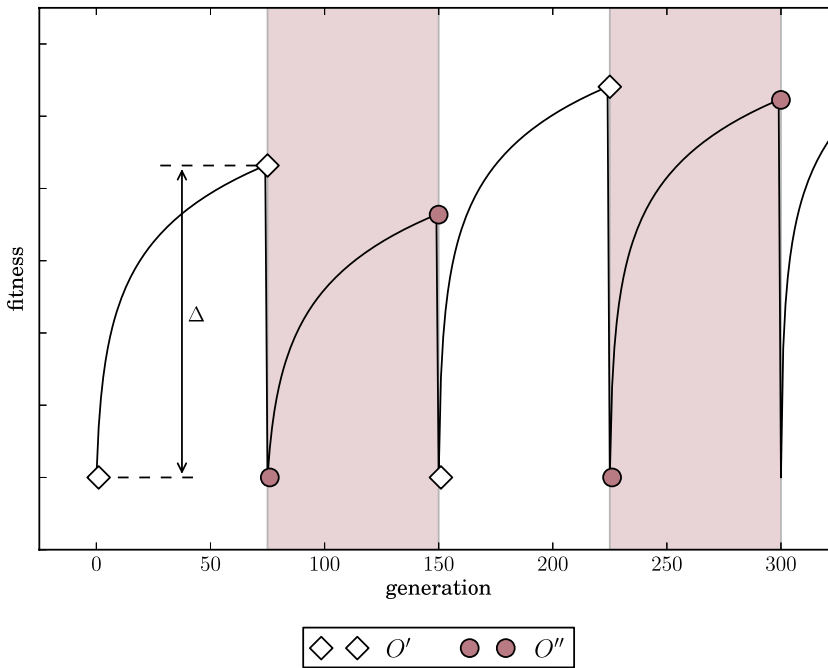


Figure 3. The hypothesized subject fitness progression in the EMU. The opponent category is manually alternated between  $O'$  and  $O''$  every 75 generations (x-axis; shaded areas). Opponent cycling is visible as sudden drops in subject fitness, followed by a period of readaptation against another opponent category. The rate of adaptability against a particular opponent category in a particular cycle is denoted by  $\Delta$ ; it becomes larger as evolution progresses. Note that fitness between different opponent categories need not per definition increase consecutively (i.e., performance against one category (e.g.,  $O''$ ) can be worse than against the other (e.g.,  $O'$ ) in a preceding cycle, while overall adaptability is still increasing).

## 4.1 Classical Analysis

### 4.1.1 Results

Figure 4a shows the online fitness measures in the GENU of seed 6, which displays large fluctuations that would be expected in cycling scenarios. The fitness averages (over 10 seeds; Figure 4b) are shown to be approximating an equilibrium of  $\approx 0.5$  (fine-tuned in a pilot study [20]). The MT shows a more informed representation of evolutionary progress (Figure 5).

### 4.1.2 Discussion

Not only is the Red Queen effect hypothesized to lead to arms races [45]; online fitness measures are also not particularly informative because of it. Neither the results from seed 6 (Figure 4a) nor the fitness averages over 10 seeds (Figure 4b) show any increase in fitness values as evolution progresses, while the results from seed 6 also show considerable noise. The other seed-specific results, not shown here, all display so much variation such that any general interpretations regarding cycling or incrementality are hard to justify.

The results from seed 6 (Figure 5a) show the rectilinear patterning that is often interpreted as indicative of cyclic progress (Section 2.2). The seed-averaged MT results (Figure 5b), on the other hand, display a subtle, yet clear diagonal bisection that resembles the arms race scenario (cf. Figure 2a), preliminarily confirming that, on the most general level, coevolution can indeed lead to long-term

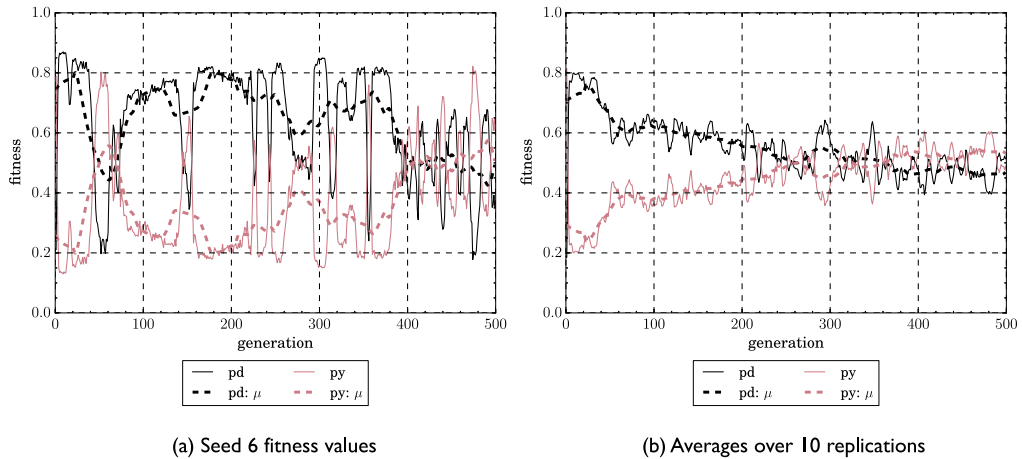


Figure 4. Fitness values of the predator and prey elites throughout the generations. Solid and dashed lines plot elite fitness and elite weighted average over 50 generations ( $\mu$ ), respectively. Dark lines plot predator fitness; light ones plot prey fitness.

progress (in both predator and prey species). Taken both the seed-specific and averaged results into account, one could thus suggest that evolution is characterized by a subtle degree of incrementality while at the same time cycling among phenotypic categories. In other words, evolution can be said to display *cyclic incrementality*.

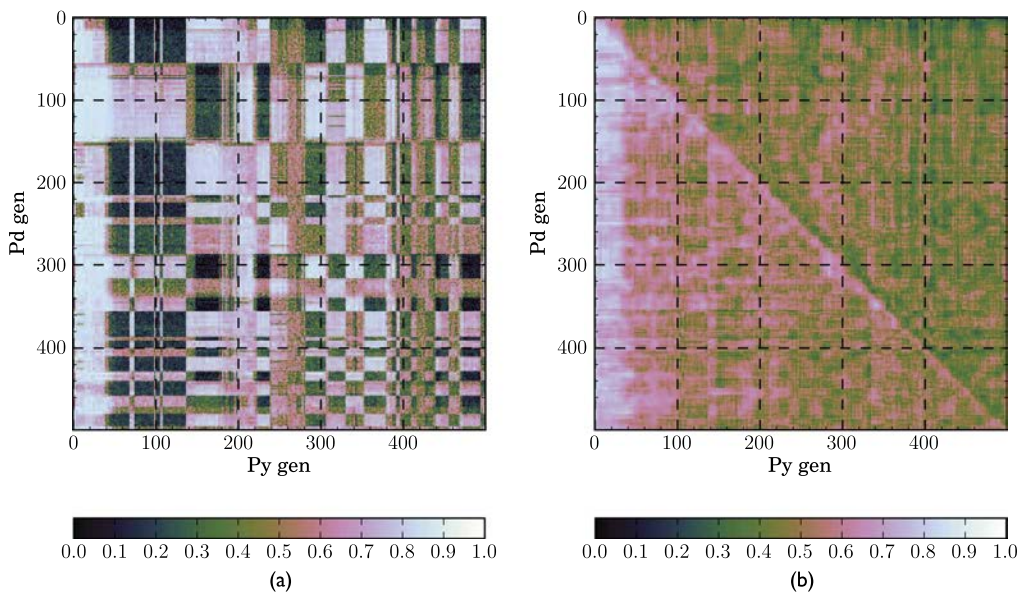


Figure 5. Visualization of MT data (i.e., performance of the predator elite of each generation against the prey elite of each generation). Each pixel represents the fitness value in terms of predator performance. The higher the value, the faster the predator caught the prey. The lower the value, the longer the prey managed to escape the predator. (a) Seed 6 fitness values. Each outcome (pixel) is an average over 25 trials. (b) Averages over 10 replications. Note the resemblance to Figure 2a.

## 4.2 Cluster Analysis

### 4.2.1 Results

We apply the cluster analysis (Section 3.3) to the phenotypes (Definition 3) expressed during the MT of seed 6 (Figure 5a). Figure 6 shows how the predator species of that seed is clustered. Each node in the predator dendrogram shows a histogram displaying the weighted average (Equation 6) of the fitness score of that predator category against all 500 prey elites in the MT (the prey’s originating generation is thus denoted on the histograms’ x-axes).

The fitness characteristics of the different predator and prey categories are shown in vertical alignment in Figure 7. The fitness progression displays a clear alternation between the two sets of categories. For the predator species, the alternation is visible in all categories (Figure 7a). For the prey species, differences between individual nodes (Figure 7b) are larger. Still, an alternating pattern is visible here as well, notably between the categories  $PY_1$ ,  $PY_3$ , and  $PY_4$  on one side, and the categories  $PY_2$  and  $PY_6$  on the other. Not only do we see an alternating pattern in the averaged performances of categories; they also appear to be distributed over time accordingly (Figure 8). In particular, the distribution of the categories  $PY_1$  and  $PY_6$  (Figure 8a) aligns with the performance of  $PD_1$  and  $PD_2$  (Figure 7a), while the distribution of  $PD_1$  and  $PD_2$  (Figure 8b) corresponds symmetrically to performance changes in  $PY_1$  and  $PY_6$  (Figure 7b). Averaged hetero-specific fitness values are shown in Table 1.

### 4.2.2 Discussion

When observing results of the cluster analysis of the predator species (Figure 6), note that the higher in the tree’s hierarchy, the larger the difference in fitness performance between two sibling categories is. This is to be expected, since the tree is constructed bottom-up, that is, the most similar individuals and categories were clustered first, the most dissimilar ones last (Section 3.3).



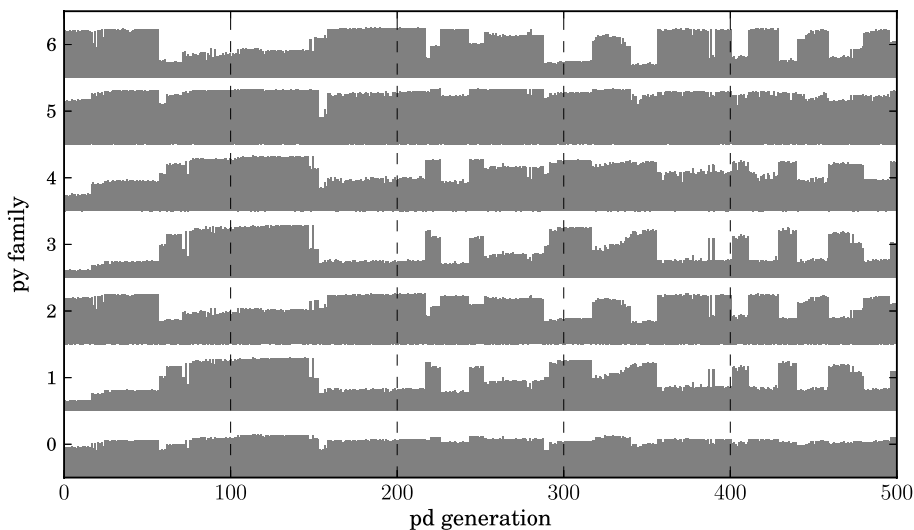


Figure 6. Visualization of the categorization tree resulting from the cluster analysis of predator strategies. Nodes show the average performance of a category’s members against all opponent strategies (cf. Figure 7). For each category, its size (i.e., the number of strategies belonging to the category) is shown in italics, and the category id (which corresponds to the labels on the y-axis of Figures 7 and 8) in upright text. Branch length is proportional to the Euclidean distance between categories. The tree has been visualized up to the fourth level only. At this level, some categories include only a single individual (e.g.,  $PD_7$ ,  $PD_{13}$ , and  $PD_{14}$ ).



(a) Predator category performance against prey opponent



(b) Predator category performance against predator opponent

Figure 7. Average performance of selected predator and prey categories against opponents. The labels (y-axis) correspond to the labels used in Figure 6. Only nodes up until a depth of level 3 are shown.

Seed 6 seems to have produced two very large, well-demarcated predator categories that are hard to further differentiate. This follows from a number of observations: First, the two most dissimilar sibling categories (at level 1 in the tree) are equally large in terms of the individuals they represent. Furthermore, the tree is notably unbalanced down from level 1 (Figure 6). Finally, both species' category performance (Figure 7) and distributions (Figure 8) show a clear alternating pattern. Overall, these observations strongly suggest that the two level 1 nodes ( $PD_1$  and  $PD_2$ ) capture the predator's cycling categories to a reasonable extent.

When we inspect numerical fitness averages of  $PD_1$  and  $PD_2$  on one hand, and  $PY_1$  and  $PY_6$  on the other (Table 1), we see a clear heterospecific cross-specialization. More specifically, fitness averages indicate that categories are circularly related as  $PD_1 > PY_1 > PD_2 > PY_6 > PD_1$ . These

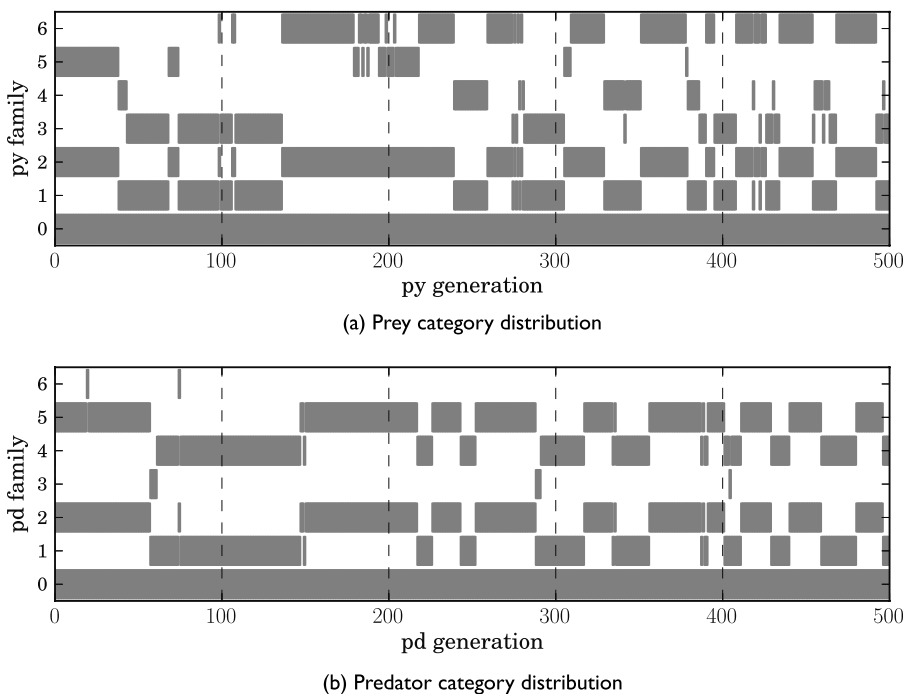


Figure 8. Distribution of category membership throughout generations. For example, the prey individual from generation 300 belongs to categories 0, 1, and 3. The labels (y-axis) correspond to the labels used in Figures 6 and 7. Only nodes up until a depth of level 3 are shown.

Table 1. Heterospecific fitness performance of selected (cycling) categories.

|        | $PD_1$ | $PD_2$ |
|--------|--------|--------|
| $PY_1$ | 0.71   | 0.34   |
| $PY_6$ | 0.34   | 0.70   |

results are in line with previously established conceptualizations of cyclic evolution, both in simulated (Section 2.2 and Figure 1) and in natural settings (Section 2.4).

We can now draw a number of conclusions. First, the elites representing distinct categories are alternatingly distributed in the MT. This indicates a similar distribution in their originating population during the actual evolutionary run.<sup>6</sup> Secondly, the alternating performance of categories corresponds to alternations in the opponent category distribution. Finally, on average, there exists a heterospecific cross-specialization between categories. In conclusion, the results demonstrate that the cluster analysis is able to capture fitness dynamics such as cyclicity, paving the way for further exploration (cf. Sections 4.3.1 and 4.4.1).

On a more conceptual note, the definition of a phenotype as a sequence of behavior-derived fitness values (Section 3.2), and the following empirical operationalization through the cluster analysis (Section 3.3), fit well with the underlying philosophy of doing ER. Since we evolved agents using the “algorithm” of natural selection, while we determined their phenotypic categorization through a (comparably automated) cluster algorithm, we increase our distance from subjective, distal interpretations

<sup>6</sup> In order to make this inductive inference, we are (fairly) assuming that the elites are phenotypically representative of their originating generation.



(Section 2.1). This observation, however, requires two issues to be addressed. First, the definition of a phenotype we propose (Section 3.2) might appear unconventional. In Section 2.3, however, it was argued that constraining an organism's phenotype to, for example, its physical body can restrict one's understanding to that of the individual as the sole level of selection. As the classical extended phenotype considers a gene's effects on the environment as part of the phenotype, the functionally extended phenotype (formalized in Section 3.2) more abstractly considers the (hypothetical) performance of an organism against opponents from different time periods to be part of the phenotype.

Secondly, we should emphasize that any definition of behavior, or indeed of a phenotype, is in effect an arbitrary one and might not per definition correspond to human intuition (cf. Section 2.3). For instance, we might regard two robot strategies (e.g., encircling the opponent either clockwise or counterclockwise) as different from each other from a distal perspective, but we would consider them similar from the proximal one (Section 2.1). The strength of using the cluster analysis thus lies in the fact that it lessens ambiguities that may arise from an external observer trying to interpret an agent's situated behavior. Furthermore, as in our case, the behaviors displayed might actually be very hard to clearly differentiate, because they differ from each other in ways so subtle it would be hard to justify any interpretation. Moreover, note that for each generation, to arrive at a full understanding of the behaviors displayed, there would be 400 combinations to check, times 500 generations. It is for precisely these reasons that the cluster analysis is useful in abstracting away from the finer details of the behavioral dynamics displayed. Thus altogether, the question exactly what kind of cluster analysis is used is subservient to the usability of the results. More specifically, we used a basic hierarchical cluster algorithm and obtained workable findings. Future studies might explore more informed clustering variants, but this decision should be mainly pragmatic and well motivated (e.g., if a hierarchical algorithm proved to be inadequate).

In conclusion, we have proposed a usage of the term "phenotype" that extends the classical concept further (Sections 2.3 and 3.2). More specifically, we made the observation that, in cyclic evolution, a subject individual would obtain a particular fitness performance against a particular opponent. This performance is based on the behavior displayed by the individual, which in turn is based on the individual's genetically determined "brain." Of course, it could be that an individual displays similar performance against a number of different opponents, but generally speaking, every subject individual should be definable by a very distinct sequence of fitness scores.

Thus, we define a phenotype as a sequence of behaviors in different contexts, identifying the corresponding individual. When clustering, individuals are therefore grouped on (functionally extended) phenotypic qualities.

### 4.3 Geno-Pheno Correlation

#### 4.3.1 Results

**Identification** If switching-genes exist, there should be a correlation between genotypes and the phenotypes shown during the MT. To investigate, we calculated the PCCs between every geno-sequence and every phenosequence for the predator species (Section 3.4.1). Results show that some of seed 6's predator geno-sequences, particularly  $\mathbf{g}_{*,25}^{PD}$  and  $\mathbf{g}_{*,26}^{PD}$ , correlate up to  $\rho \approx |0.7|$  with predator phenosequences, when averaged (Equation 7) over prey category distributions (Figure 9). The correlations are notably positively expressed against prey categories  $PY_3$ ,  $PY_8$ ,  $PY_{17}$ , and  $PY_{18}$  ( $\rho_\mu \geq 0.7$ ; Table 2a), but notably negatively against  $PY_6$ ,  $PY_{14}$ ,  $PY_{27}$ , and  $PY_{28}$  ( $\rho_\mu \leq -0.5$ ; Table 2b).

Of these categories, we selected two sets ( $\{PY_3, PY_6\}$  and  $\{PY_8, PY_{14}\}$ ) for further investigation by EMUs (Section 4.4.1). We selected categories based on the criteria that (a) they are highly correlated and represent large categories, following the weighted correlation coefficient

$$\theta \stackrel{\text{def}}{=} \frac{|PY'|}{|PY|} \left( \frac{1}{2} \left( \rho \left( \mathbf{g}_{*,25}^{PD}, PY' \right) + \rho \left( \mathbf{g}_{*,25}^{PD}, PY' \right) \right) \right) \quad (8)$$




Figure 9. Correlation coefficients between seed 6’s 28 predator genosequences and the 29 highest-level (i.e., up until level 5 in the cluster hierarchy) prey categories.

—implying that the aforementioned correlations are not trivial—and (b) each set contains categories showing inverse correlation signs, which suggests opposing specializations.

**Transplantation** Looking at the data in Section 4.3.1, we might now suggest that predator genes  $\mathbf{g}_{*,25}^{PD}$  and  $\mathbf{g}_{*,26}^{PD}$  control opponent specialization. To test that, we sampled pairs of those genes’ alleles from either  $PD_1$  or  $PD_2$  and distributed those over the entire predator population. We evaluated the effect on predator performance in two sets of two MTs each (Section 3.4.2).

In the first set (Figure 10a,b and Table 3a), we extracted alleles from the genes  $\mathbf{g}_{*,25}^{PD}$  and  $\mathbf{g}_{*,26}^{PD}$ , while in the second set (Figure 10c,d and Table 3d), we extracted alleles from all genes *except*  $\mathbf{g}_{*,25}^{PD}$  and  $\mathbf{g}_{*,26}^{PD}$ . For each set, the alleles’ originating category was varied ( $PD_1$  or  $PD_2$ ). We selected alleles from random members of a category, and then distributed them over the entirety of  $PD$ .

Figure 10a,b show that allele transplantation leads to an approximation of the fitness characteristics of the genes’ originating category. In other words, the entire population now becomes specialized in coping with one particular opponent category (Table 3a). Correspondingly, transplanting less-correlating genes results in little change in the MT: Although opponent specialization becomes less pronounced, no actual switch is apparent (Figure 10c,d and Table 3d). Altogether, this indicates

Table 2. Notably strong PCCs between predator genosequences and predator phenosequences expressed against various prey categories (following  $\rho_{\mu}(\mathbf{g}_{*,\gamma}^{PD}, PY')$ ; Equation 7).

|                       | (a) Strongly positive PCC values<br>(i.e., $\rho \geq 0.7$ ) |        |           |           | (b) Strongly negative PCC values<br>(i.e., $\rho \leq -0.5$ ) |           |           |           |        |
|-----------------------|--|--------|-----------|-----------|---|-----------|-----------|-----------|--------|
|                       | $PY_3$   | $PY_8$ | $PY_{17}$ | $PY_{18}$ | $PY_6$  | $PY_{14}$ | $PY_{27}$ | $PY_{28}$ |        |
| $ PY' $               | 149  | 144    | 75        | 39        | $ PY' $   | 209       | 176       | 104       | 72     |
| $\mathbf{g}_{*,25}^S$ | 0.70   | 0.72   | 0.71      | 0.73      | $\mathbf{g}_{*,25}^S$   | -0.50     | -0.59     | -0.66     | -0.50  |
| $\mathbf{g}_{*,26}^S$ | 0.55   | 0.58   | 0.59      | 0.56      | $\mathbf{g}_{*,26}^S$   | -0.48     | 0.55      | -0.56     | -0.55  |
| $\theta$              | 0.186  | 0.187  | 0.098     | 0.050     | $\theta$  | -0.205    | -0.201    | -0.127    | -0.076 |

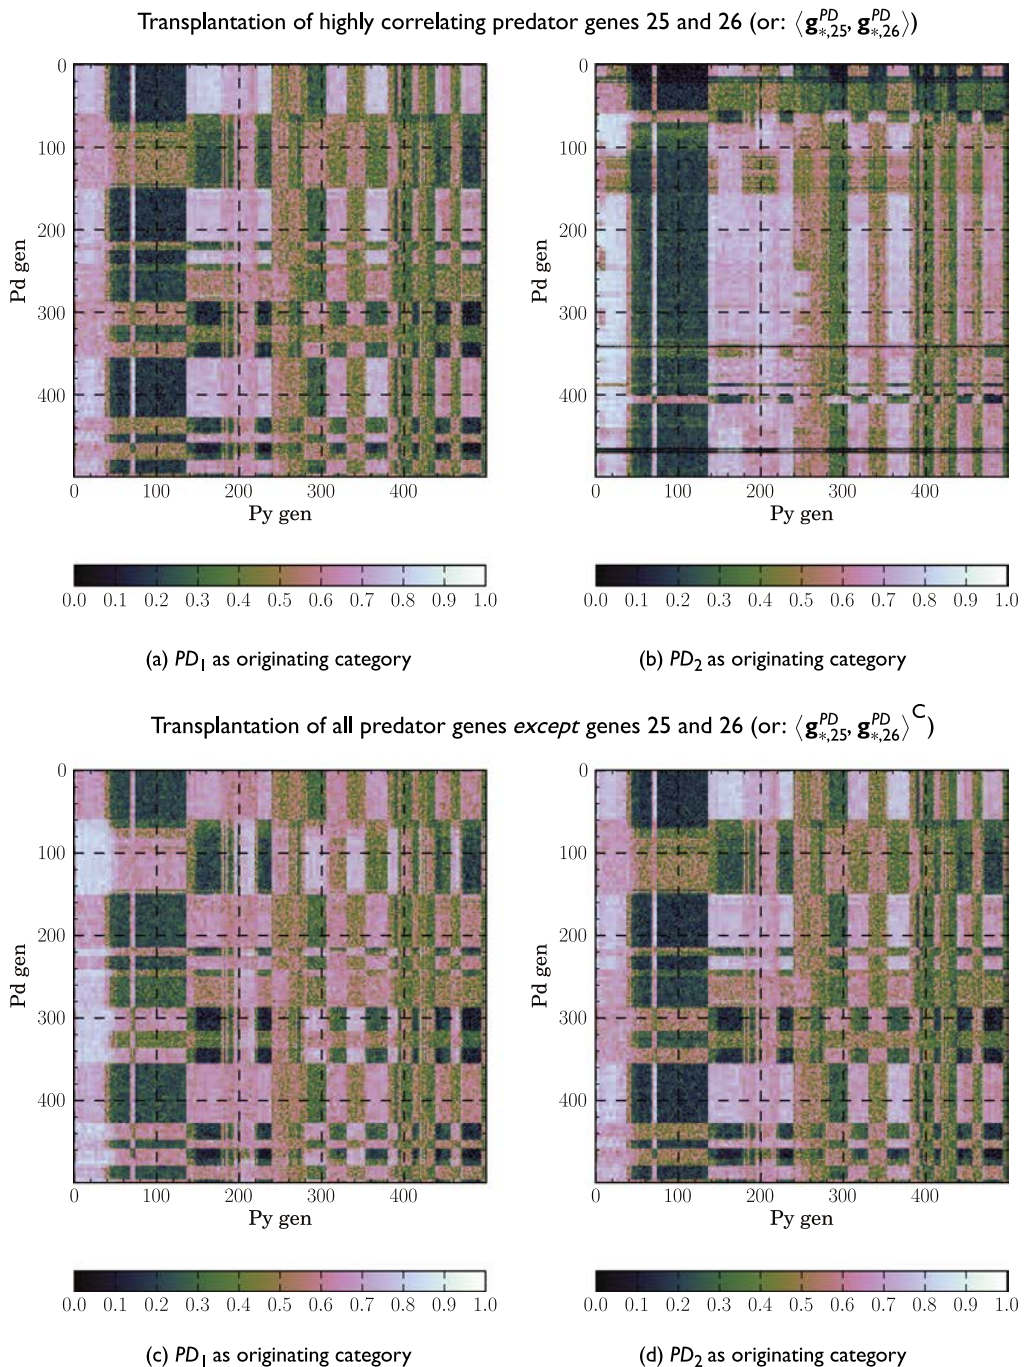


Figure 10. MT plots after predator allele transplantation (compare with Figure 5a).

that (the combination of) predator genes  $\mathbf{g}_{*,25}^{PD}$  and  $\mathbf{g}_{*,26}^{PD}$  can indeed be considered to fulfill the role-switching-genes that we first hypothesized in Section 1. Possible conflictor and suppressor genes that might emerge in conspecific parent-offspring conflict (Section 2.4) could control phenotypic expressions through a similar mechanism.

Table 3. Heterospecific performance (compare with Table 1) after allele transplantation (Section 3.4).

| (a) Transplantation of highly correlating predator genes 25 and 26 (or: $\langle \mathbf{g}_{*,25}^{PD}, \mathbf{g}_{*,26}^{PD} \rangle$ ). |        |        |                                    |        |        | (d) Transplantation of all predator genes except genes 25 and 26 (or: $\langle \mathbf{g}_{*,25}^{PD}, \mathbf{g}_{*,26}^{PD} \rangle^C$ ). |        |        |                                    |        |        |
|---|--------|--------|------------------------------------|--------|--------|---|--------|--------|------------------------------------|--------|--------|
| (b) $PD_1$ as originating category  |        |        | (c) $PD_2$ as originating category |        |        | (e) $PD_1$ as originating category  |        |        | (f) $PD_2$ as originating category |        |        |
|   | $PD_1$ | $PD_2$ |                                    | $PD_1$ | $PD_2$ |   | $PD_1$ | $PD_2$ |                                    | $PD_1$ | $PD_2$ |
| $PY_1$  | 0.62   | 0.49   | $PY_1$                             | 0.39   | .33    | $PY_1$  | 0.62   | 0.38   | $PY_1$                             | 0.49   | 0.33   |
| $PY_6$  | 0.35   | 0.33   | $PY_6$                             | 0.58   | 0.63   | $PY_6$  | 0.35   | 0.58   | $PY_6$                             | 0.34   | 0.63   |

### 4.3.2 Discussion

Although we deliberately did not address the distal perspective (Sections 2.1 and 4.2.2), a brief elaboration of how switching-genes can influence behavior is appropriate. Upon closer inspection, genes  $\mathbf{g}_{*,25}^{PD}$  and  $\mathbf{g}_{*,26}^{PD}$  express the strength of the NN connections between the center-right and right camera sectors in the input layer and the output turning neuron (Sections A1.1 and A1.2). Thus, by modifying just these two weights, the EA is able to realize quick adaptations to different opponent categories. Of course, this setup only works if the rest of the NN provides the right background for such a mechanism to take effect. Therefore, we can say that the NN has evolved to a state that is predisposed to enable large change in the phenotype through small structural (i.e., genetic) change. More speculatively, more complex NNs might realize genetic adaptability in a more “exotic” fashion. For instance, one could imagine switching-genes controlling neural “gates” that could either block or facilitate downstream neural activity. This way, a neural network could encode the responses for multiple opponent strategies, while only one gets promoted as actual motor activation.

Adaptive mechanisms such as the ones observed and hypothesized are not unlike the functioning of some genes in nature. For example, while the current study used a simplified model where genes and neural connections correspond isomorphically, gene expression in actual organisms is dependent on the interactions between many genes (*epistasis*) that form complex regulatory networks (cf. [33]). Striking examples of transcription factors involved in these networks are *hox genes* that control morphogenesis in embryonic development [5]. Intriguingly, traits that are suppressed can still remain dormant in a species’ genome. For example, a single mutation in a particular *hox* gene might lead to examples of chickens growing teeth, or horses growing toes [17]. This reexpression of ancestral, dormant features is known as *atavism* and can be said to parallel the retention of dormant specializations in specific opponent categories, which forms the basis of cyclic incrementality (further discussed in Section 4.4.2). The difference from switching-genes emphasized in the current study is that *hox* genes do not directly encode phenotypic features, but instead promote or suppress gene sequences that do. Mechanistically, however, they appear to play comparable roles, and might be of value in enhancing a species’ evolvability.

## 4.4 Opponent Cycling Emulation

### 4.4.1 Results

To determine whether a genetic predisposition for adaptability might slowly evolve (and how switching-genes might play a role in this), we emulated cycling in the prey opponent (Section 3.5). So, instead of evolving both species (as in the GENU), we now only evolved the predator species. The prey individuals, on the other hand, were randomly sampled from the previously evolved categories  $\{PY_3, PY_6\}$  (Figures 11a and 11b) and  $\{PY_8, PY_{14}\}$  (Figures 11c and 11d). Every 75 generations, we manually alternated the prey category that was sampled from. This presented the predator with a series of opponents indistinguishable from a genuinely cyclically coevolving prey.

We ran each experiment 25 times with random initial conditions (prey genotypes were fixed after having been sampled). Then we ran a MT (switching prey categories following the same 75-generation interval) to measure fitness progress more accurately. So, in this MT, we tested the elites of the newly evolved predators against the prey elites from the same category samples manually presented during the EMU.

The data obtained show that the predator species displays an increase in fitness obtained against both prey categories as evolution progresses (Figure 12 and Table 4), except against  $PY_{14}$  during the prey sequence  $\langle PY_8, PY_{14}, \dots \rangle$  (Figures 11c and 12c). The increase is less pronounced against the first prey category when including the first cycle (i.e., generations 0–75), and in fact shows an apparent fitness drop.

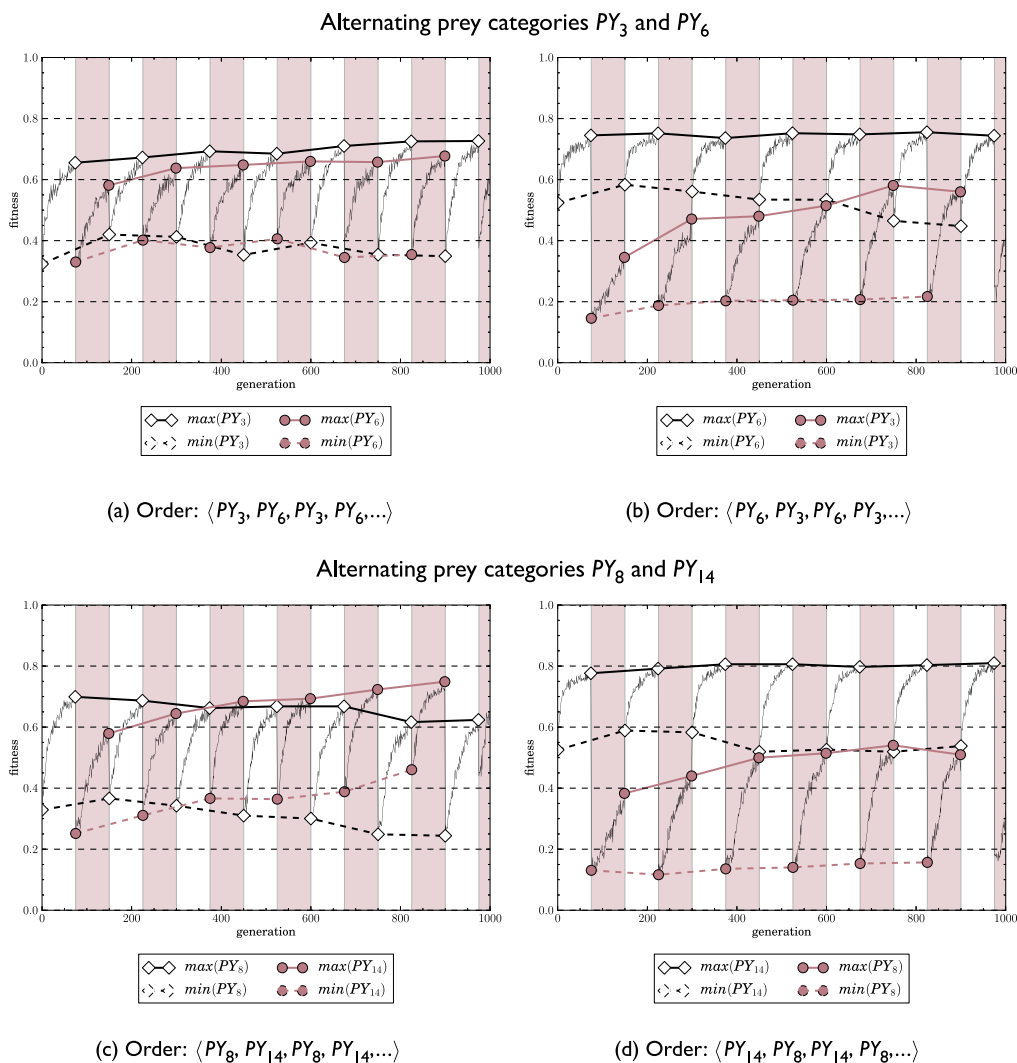


Figure 11. Master fitness of evolved predator species and evolved prey categories during the EMU. Solid and dashed lines mark cycle maxima and minima respectively. Vertically shaded areas mark the prey cycle interval. Each value represents an average over 500 trials. See also Figure 3.

### 4.4.2 Discussion

We might explain the initial fitness drop (Section 4.4.1) by observing that it is, crucially, due to an increase in the minimum fitness and not due to a decrease in maximum (Figure 11), and is therefore advantageous. This is a plausible progression because, according to the observations, the predator gets primed against the first prey category it encounters, as seen in a consistent performance decrease from the first cycle to the second (when it encounters the second category for the first time). When encountering the third cycle, the predator is adapted to the second category. While then having to readapt to the first prey category from the second is of course a challenge, it is better than having to adapt from a randomly initialized state; hence the increase in the minimum fitness (and the apparent fitness drop). This explanation provides a justification for disregarding the first cycle when determining the overall increase in adaptability (Table 4); only from the second cycle onwards is the predator species able to “realize” it has to cope with multiple opponent categories.

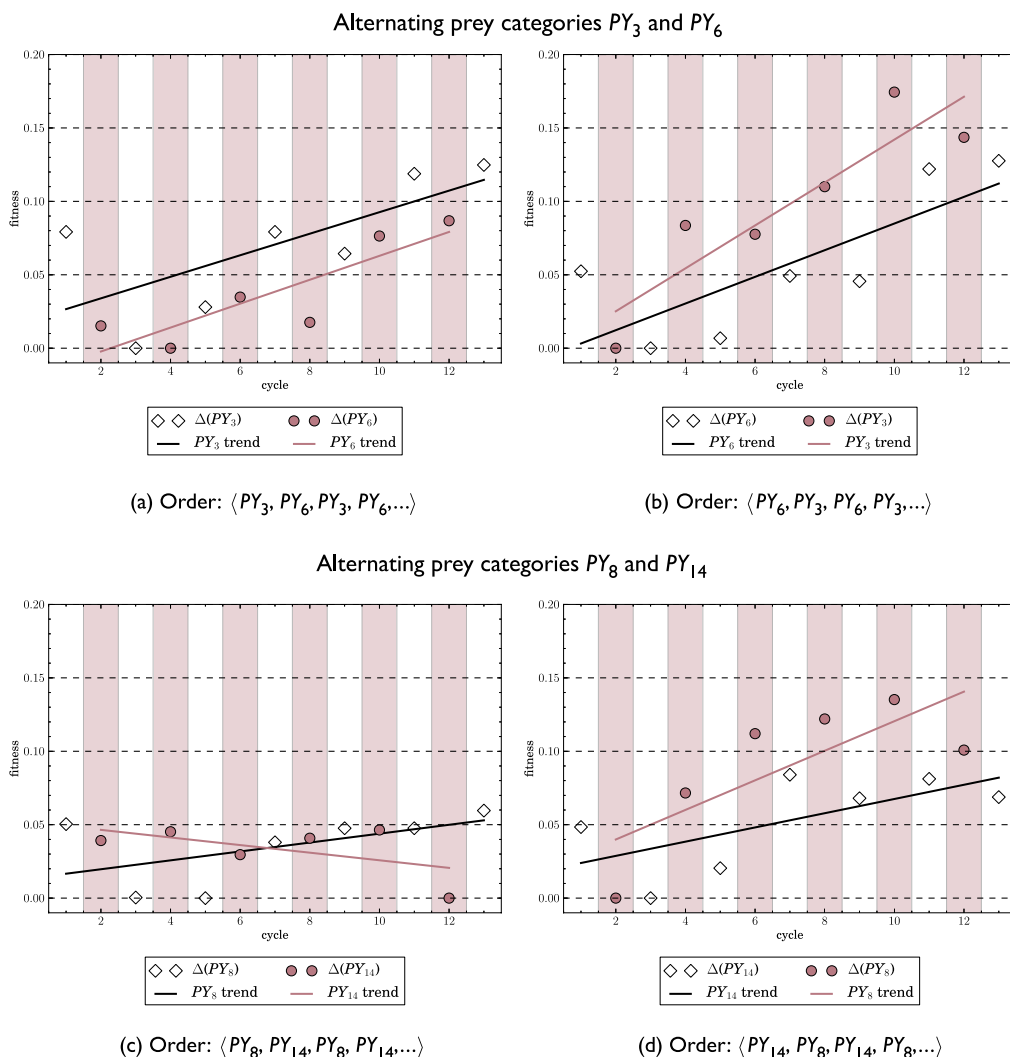


Figure 12. Relative master fitness increase of predator species against prey categories during the EMU. Shown is the within-cycle difference between cycles’ minima and maxima as observed in Figure 11 ( $\Delta$  in Figure 3). Vertically shaded areas mark the prey cycle interval. Each value represents an average over 500 trials.

Table 4. Overall predator master fitness increase (i.e., the fitness difference between the last and the first cycle of a particular prey category during a particular prey alternation sequence), in percentage points. Shown is the average fitness increase  $\Delta$  (where  $\Delta^+$  indicates that the first cycle is ignored) during the EMU over 500 trials (25 seeds, 20 trials per generation), and mean squared error (over 25 seeds) between parentheses.


|                        | $\langle PY_3, PY_6, \dots \rangle$ | $\langle PY_6, PY_3, \dots \rangle$ | $\langle PY_8, PY_{14}, \dots \rangle$ | $\langle PY_{14}, PY_8, \dots \rangle$ |
|------------------------|-------------------------------------|-------------------------------------|--|--|
| $\Delta$ ( $fam_1$ )   | 4.56 (4.88)                         | 7.52 (4.91)                         | 0.92 (4.89)                            | 2.04 (4.37)                            |
| $\Delta^+$ ( $fam_1$ ) | 12.48 (3.79)                        | 12.76 (4.06)                        | 5.92 (5.18)                            | 6.88 (5.49)                            |
| $\Delta$ ( $fam_2$ )   | 7.16 (3.91)                         | 14.36 (4.04)                        | -3.92 (4.28)                           | 10.08 (4.17)                           |

More generally speaking, however, the increase in evolvability (Table 4) must be attributed to either an increase in fitness maxima or a decrease in minima. Particularly in prey sequences  $\langle PY_6, PY_3, \dots \rangle$  (Figure 11b) and  $\langle PY_{14}, PY_8, \dots \rangle$  (Figure 11d) (where  $PY_{14} \subset PY_6$  and  $PY_3 \subset PY_8$ ), the effect is clearly visible, due to a large initial fitness difference between the first and second prey categories. More specifically, adaptability against the first category increases due to a decrease in cycle minima from  $\approx 0.6$  to  $\approx 0.5$ , while maintaining the maxima of  $\approx 0.8$ . Conversely, adaptability against the second category increases due to an increment in cycle maxima from  $\approx 0.4$  to  $\approx 0.6$ , while maintaining cycle minima of  $\approx 0.2$ . When prey sequences are reversed (Figure 11a and 11b), the effect is less demarcated and results from both a gradual decrease in minima and an increase in a maxima. Thus, in the latter case, the increase in adaptability seems to maintain the initial fitness equilibrium against both prey categories, while in the former the equilibrium is approached from the initially unbalanced distribution.

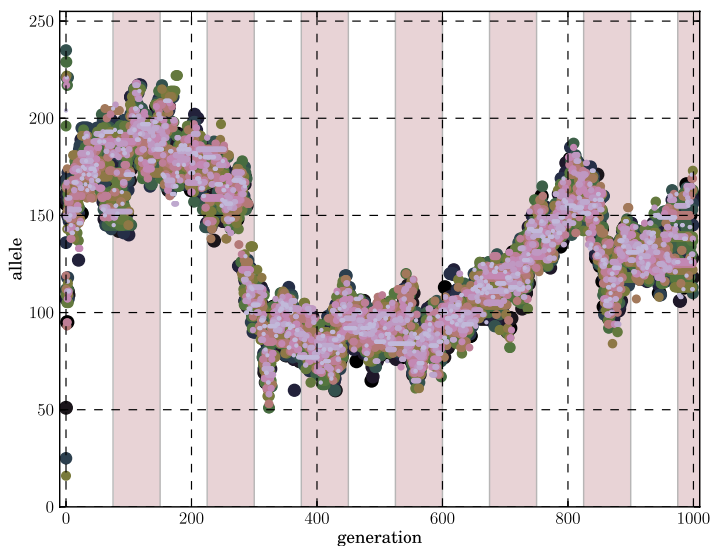
In the genetic domain, if we look at the allele distribution of switching gene 25 of seed 18 (Figure 13a), we see that the within-cycle prey response is characterized by linear genetic change. In the phenotypic domain, these within-cycle changes correspond to seemingly sigmoidal growth patterns that get interrupted by sudden drops caused by the opponent changing strategy (Figure 11d). Between these interruptions, however, the initial response is consistently acute. In a more long-term phenotypic time frame, however, we observe a between-cycle linear change (Figure 12), in the sense that the distance between the sigmoids' onset and upper asymptote generally increases when more alternations are encountered. The response's long-term linear profile shows us there is a notable delay between the actual encounter of a cycling opponent and the eventual development of an adaptive, pseudo-Baldwinian strategy. A long-term parallel is not immediately visible in the genetic domain, possibly, due to many genes being a factor in the phenotype (cf. complex traits).

Looking more closely at the allele distribution of gene 14 of seed 18, we can observe a more gradual genetic drift (Figure 13b). In contrast, the highly correlating allele fluctuations start to follow prey alternation intervals relatively abruptly (Figure 13b). Since we have already established that evolvability emerges gradually, while the allele alternation patterns show switching-genes appearing suddenly, we can safely reject the notion that the gradual emergence of adaptability is exclusively mediated by switching-genes. Again, this suggests the existence of complex traits. More specifically, it is likely that, if switching-genes are developed, the remaining, less correlating *context genes* co-provide the appropriate genetic environment for switching-genes to shift strategy efficiently. To hypothesize: Switching-genes control which opponent specialization gets expressed, but they depend on context genes to do so efficiently. Again, this is not unlike gene regulators such as *hox* genes or, for example, sex loci, although, again, our model does not allow for regulatory effects (Section 4.3.2). Overall, then, there are adaptive processes present at multiple scales, and they manifest differently in the phenotypic and the genetic domain. All the while, the correspondence between genotype and phenotype, even in our simplified model, is non-isomorphic. Multiple genes shape traits instead.

However, we need to report one anomaly when discussing the EMUs. First of all, geno-pheno PCC values obtained during the EMU (Figure 14a) and against the original prey categories in the



(a) Allele distribution of the highly correlated predator gene 25 (or  $g_{*,25}^{PD}$ )




(b) Allele distribution of the less correlated predator gene 14 (or  $g_{*,14}^{PD}$ )

Figure 13. Allele distributions of two predator genes during seed 18 evolved against prey sequence  $\langle PY_3, PY_6, PY_3, PY_6, \dots \rangle$  (Figure 14b). Allele shading indicates individual origin (20 shades, one shade for each of the 20 predator individuals). Allele values may be obscured by others. Vertical shaded bands correspond to prey alternation interval.


GENU (Table 2) are overall markedly similar. Notably, however, the high correlations of  $g_{S*,1}^{PD}$  and  $g_{*,27}^{PD}$  in the EMU were not observed in the GENU. This can, first of all, be explained by taking into account that the EMU shows PCC averages over 25 seeds instead of a single one. Furthermore, we designed the EMU to be more constrained than the GENU; it only presented a subset of prey opponents to the predator out of the GENU prey population (358 ( $= |PY_3| + |PY_6|$ ) out of 500 prey individuals). Altogether, we can expect small deviations in a species' genome when we compare the EMU with the initial GENU.



In conclusion, the gradual increase in predator adaptation against alternating prey categories demonstrates that species are not merely adapting to a particular situation in a strictly ad hoc fashion (Section 2.4). Instead, genotypes that allow for quick shifts in opponent specialization develop gradually as species are exposed to an increasing number of opponent cycles. Since individuals were controlled by simple perceptrons (Section A1.2), exploitation of phenotypic (e.g., neural) plasticity historically associated with the Baldwin effect (Section 2.4) is ruled out, and thus the increase in evolvability must be of a purely genetic nature. Furthermore, the demonstrated retention of historical



(a) PCC averages over 25 seeds



(b) Seed 18 PCC values

Figure 14. Geno-pheno correlation coefficients of the 28 predator genes evolved against the prey sequence ( $PY_3, PY_6, PY_3, PY_6, \dots$ ).

adaptations implies an instance of *genetic memory*, mechanistically not unlike that seen in cases of atavism (Section 4.3.2).

## 5 Conclusion

We investigated the evolutionary dynamics characterizing competitive coevolutionary scenarios through computer simulations involving two populations of predator and prey robots. By proposing and formalizing the concept of a *functionally extended phenotype* (Section 2.3) and the subsequent utilization of a *cluster analysis* technique (Section 3.3), we showed that the evolutionary process can converge toward limit cycling dynamics in which the two populations periodically rediscover previously discarded strategies. Thus, the evolutionary process does not lead to any arms races characterized by progressive complexification of agents' competence as is often hypothesized to occur in nature. In contrast to previous studies, our own study shows a high degree of formalization and automation in the analysis leading to these conclusions.

We hypothesized that the need to adapt to periodic variations in the environment could lead to the synthesis of genetic organizations characterized by a readiness to change toward specific directions (Section 1). We speculated that this form of evolvability could be realized through the utilization of *switching-genes*, that is, a limited number of genes enabling large shifts in the phenotype through only a few mutations. We first explored this idea by analyzing the temporal correlations between genetic and phenetic sequences (Section 3.4). If switching-genes exist and they control opponent specialization (i.e., they control large shifts in the phenotype), changes in those genes over time should correspond to changes in phenetic qualities (phenes). While we indeed showed that such correlations exist (Section 4.3), they did not appear to be distributed evenly against all opponent individuals. Here, we used the cluster analysis data (Section 4.2.1) to map correlations onto the category level. This showed that some genes show strong correlations only against certain opponent categories. Transplanting such highly correlating genes into individuals the genes did not originate from caused a shift in opponent specialization, providing further evidence for their role as switching-genes (Section 4.3). Conversely, transplanting less-correlating genes had negligible effect.

In conclusion, we proposed and demonstrated that switching-genes express specialization to different opponents (Section 4.3.2). This is made possible by retaining historical adaptations to previously encountered opponents. While our model abstracts away from reality and in particular does not include gene-regulatory mechanisms or any form of ontogenetics, this nevertheless parallels the idea that certain properties can remain dormant in an organism's genotype, possibly reexpressed by a single mutation such as is seen with *hox* genes. More conceptually, the fact that we observe the emergence of switching-genes only in some of the experiment replications (Section 4) demonstrates that they, in interaction with other genes, act as critical parameters in cyclic evolution. Like any dynamical system, the emergence of such heavily interaction-dependent organizations very much relies on initial conditions and on-line perturbations, and therefore does not develop in every case.

Finally, we investigated whether a disposition for evolvability, possibly mediated through switching-genes, would emerge abruptly or gradually. To this end, we used the cluster analysis to extract opponent categories with high geno-pheno correlations from a GENU conducted earlier (Section 3.5). We then presented these categories to the subject species in a controlled, alternating fashion, enabling fine-tune control over cycling dynamics. Here, it was shown that evolvability emerges gradually as evolution progresses (Section 4.4). Furthermore, it appeared that identical switching-genes that were observed during the GENU emerged now as well (the latter as a subset of the former). This provides yet another indication that these genes play a fundamental role in opponent specialization. Correspondingly, allele distributions showed switching-genes following an alternating pattern mirroring the opponent alternation interval but, somewhat unexpectedly, emerging quite abruptly. Less-correlating genes on the other hand showed a more gradual genetic drift (Section 4.4.2). This suggests that switching-genes control opponent specialization, provided that the remaining *context genes* supply the appropriate genetic background for switching-genes to shift strategy effectively.

In conclusion, it is clear that coevolving dynamics, while a plausible explanation for complexity in nature, is less self-evident than is often assumed, and involves highly interaction-driven interdependencies that require further advances in evolutionary thinking and subsequent analyses. Our study is an attempt to show how computer modeling can be of use in this effort.

## Acknowledgments

The authors would like to thank Dan Dediu of the Language and Genetics Department, Max Planck Institute for Psycholinguistics; Scott Moisik of the Division of Linguistics and Multilingual Studies, Nanyang Technological University; Daniel Tauritz of the Department of Computer Science, Missouri University of Science and Technology; and two anonymous reviewers for their comments and suggestions.

## References

- Allen, J. (1991). Chaos and coevolution: Evolutionary warfare in a chaotic predator-prey system. *Florida Entomologist*, 74(1), 50–59.
- Baldwin, J. (1896). A new factor in evolution. *American Naturalist*, 30(354), 441–451.
- Baronchelli, A., Chater, N., Christiansen, M., & Pastor-Satorras, R. (2013). Evolution in a changing environment. *PLoS One*, 8(1), e52742.
- Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72(1), 173–215.
- Carroll, S. B. (2005). *Endless forms most beautiful: The new science of evo devo and the making of the animal kingdom*. New York: W.W. Norton.
- Cartlidge, J. (2004). *Rules of engagement: Competitive coevolutionary dynamics in computational systems*. Ph.D. thesis, University of Leeds.
- Cartlidge, J., & Bullock, S. (2004). Unpicking tartan CIAO plots: Understanding irregular coevolutionary cycling. *Adaptive Behavior*, 12(2), 69–92.
- Cliff, D., & Miller, G. F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Morán, A. Moreno, J. J. Merelo, & P. Chacón (Eds.), *Advances in Artificial Life: Third European Conference on Artificial Life, Granada, Spain, 1995, Proceedings* (pp. 200–218). Berlin: Springer.
- Cliff, D., & Miller, G. F. (1996). Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 506–515). Cambridge: MIT Press.
- Dawkins, R. (1999). *The extended phenotype: The long reach of the gene*. Oxford, UK: Oxford University Press.
- Dawkins, R., & Krebs, J. (1979). Arms races between and within species. In *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 205, 489–511.
- Felsenstein, J., Archie, J., Day, W., Maddison, W., Meacham, C., Rohlf, F., & Swofford, D. (1986). The Newick tree format. Unpublished manuscript.
- Floreano, D., & Nolfi, S. (1997). Adaptive behavior in competing co-evolving species. In P. Husbands & I. Harvey (Eds.), *Proceedings of the Fourth European Conference on Artificial Life* (pp. 378–387). Cambridge, MA: MIT Press.
- Floreano, D., & Nolfi, S. (1997). God save the Red Queen! Competition in co-evolutionary robotics. In J. R. Koza, K. Deb, M. Dorigo, D. Foegel, B. Garzon, H. Iba, & R. L. Riolo (Eds.), *Genetic Programming 1997: Proceedings of the Second Annual Conference* (pp. 398–406). San Francisco, CA: Morgan Kaufmann.
- Floreano, D., Nolfi, S., & Mondada, F. (2001). Co-evolution and ontogenetic change in competing robots. In J. P. Mukesh, V. Honavar, & K. Balakrishnan (Eds.), *Advances in Evolutionary Synthesis of Neural Networks* (pp. 273–306). Cambridge, MA: MIT Press.
- Futuyma, D. J., & Slatkin, M. (Eds.) (1983). *Coevolution*. Sunderland, MA: Sinauer.
- Gould, S. (1983). *Hen's teeth and horse's toes*. New York: W.W. Norton.
- Hinton, G., & Nowlan, S. (1987). How learning can guide evolution. *Complex systems*, 1(1), 495–502.
- Huerta-Cepas, J., Dopazo, J., & Gabaldón, T. (2010). ETE: A Python environment for tree exploration. *BMC Bioinformatics*, 11(1), 24.
- Janssen, R. (2012). The royal family. Master's thesis, Radboud University Nijmegen, the Netherlands.

21. Johannsen, W. (1911). The genotype conception of heredity. *The American Naturalist*, 45(531), 129–159.
22. John, B., & Lewis, K. (1966). Chromosome variability and geographic distribution in insects. *Science*, 152(3723), 711–721.
23. Jonker, C., Snoep, J., Treur, J., Westerhoff, H., & Wijngaards, W. (2002). Putting intentions into cell biochemistry: An artificial intelligence perspective. *Journal of Theoretical Biology*, 214(1), 105–134.
24. Kerr, B., Riley, M., Feldman, M., & Bohannan, B. (2002). Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors. *Nature*, 418(6894), 171–174.
25. Kirkup, B., & Riley, M. (2004). Antibiotic-mediated antagonism leads to a bacterial game of rock–paper–scissors in vivo. *Nature*, 428(6981), 412–414.
26. Lotka, A. (1910). Contribution to the theory of periodic reactions. *The Journal of Physical Chemistry*, 14(3), 271–274.
27. Miller, G., & Cliff, D. (1994). *Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations*. Technical Report CSRP311. University of Sussex, School of Cognitive and Computing Sciences.
28. Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptoz, A., Magnenat, S., Zufferey, J., Floreano, D., & Martinoli, A. (2009). The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions* (pp. 59–65).
29. Nolfi, S. (2012). Co-evolving predator and prey robots. *Adaptive Behavior*, 20(1), 10–15.
30. Nolfi, S., & Floreano, D. (1998). Coevolving predator and prey robots: Do “arms races” arise in artificial evolution? *Artificial Life*, 4(4), 311–335.
31. Nolfi, S., & Gigliotta, O. (2010). Evorobot\*: A tool for running experiments on the evolution of communication. In S. Nolfi & M. Mirolli (Eds.), *Evolution of Communication and Language in Embodied Agents* (pp. 297–301). Berlin: Springer.
32. O’Brien, G., & Yule, W. (1995). *Behavioural phenotypes*. Cambridge, UK: Cambridge University Press.
33. Palsson, B. (2006). *Systems biology*. Cambridge, UK: Cambridge University Press.
34. Parisi, D. (1997). Artificial life and higher level cognition. *Brain and Cognition*, 34(1), 160–184.
35. Parker, G. (1979). Sexual selection and sexual conflict. In M. Blum (Ed.), *Sexual selection and reproductive competition in insects* (pp. 123–166). New York: Academic Press.
36. Parker, G., & Macnair, M. (1979). Models of parent-offspring conflict. IV. Suppression: Evolutionary retaliation by the parent. *Animal Behaviour*, 27, 1210–1235.
37. Ridley, M. (1994). *The Red Queen: Sex and the evolution of human nature*. Harmondsworth, UK: Penguin.
38. Rodgers, J., & Nicewander, W. (1988). Thirteen ways to look at the correlation coefficient. *The American Statistician*, 42(1), 59–66.
39. Roughgarden, J. (1983). Coevolution between competitors. In D. J. Futuyama & M. Slatkin (Eds.), *Coevolution* (pp. 383–403). Sunderland, MA: Sinauer.
40. Sharkey, N., & Heemskerk, J. (1997). The neural mind and the robot. In A. J. Browne (Ed.), *Neural network perspectives on cognition and adaptive robotics* (pp. 169–194). London: IOP Press.
41. Sinervo, B., & Lively, C. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380(6571), 240–243.
42. Sinervo, B., Miles, D., Frankino, W., Klukowski, M., & DeNardo, D. (2000). Testosterone, endurance, and Darwinian fitness: Natural and sexual selection on the physiological bases of alternative male behaviors in side-blotched lizards. *Hormones and Behavior*, 38(4), 222–233.
43. Sokal, R., & Michener, C. (1958). A statistical method for evaluating systematic relationships. *University of Kansas Science Bulletin*, 38, 1409–1438.
44. Turney, P. (1996). Myths and legends of the Baldwin effect. In *Proceedings of the Workshop on Evolutionary Computing and Machine Learning at the 13th International Conference on Machine Learning* (pp. 135–142).
45. Van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1(1), 1–30.
46. Waddington, C. (1953). Genetic assimilation of an acquired character. *Evolution*, 7(2), 118–126.
47. Xu, R., & Wunsch, D. (2005). Survey of clustering algorithms. *IEEE Transactions on Neural Networks*, 16(3), 645–678.

## Appendix I: Simulation Implementation

### AI.1 Robots and Environment

We situated two 75-mm-diameter e-puck robots [28] in a  $60 \times 60$ -cm environment surrounded by walls. We equipped both the predator and prey robots with eight infrared sensors (spaced  $45^\circ$  apart) that enabled them to detect nearby objects, a camera that enabled them to detect the relative position of the other robot, and two actuators controlling the speed of two corresponding wheels (Figure 15). The relative position of an object (i.e., wall or opponent robot) to an infrared sensor was matched to a sample table that contained activation measurements from real infrared sensors of the e-puck. The state of each sensor was encoded in a corresponding sensory neuron. The state of the neural controller, of the sensors, and of the actuators is updated every 100 ms.

The predator and prey differed in the maximum speed at which the robots could move, and in the field of view of the camera. We set the robots' maximum speed to  $v_{max} = 80$  mm/s for the predator and to  $v_{max} = 100$  mm/s for the prey; we set the cameras' view fields to  $45^\circ$  and  $360^\circ$ . We divided the cameras' angles of view into five sectors of  $9^\circ$  and  $72^\circ$ , respectively. The state of each of the five sectors was fed into five sensory neurons that encoded the average gray level of the corresponding  $1^\circ$  photoreceptors.

The offset in speed between the predator and the prey was set so to balance the relative efficacy of respectively catching or escaping the opponent. The difference in the field of view was set to take




Figure 15. Layout of sensors and actuators (predator configuration shown). The arrows show the robot's forward direction. (a) Infrared sensors, linear camera, and the two wheels are represented by dotted lines, dashed lines, and the two dotted rectangles respectively. (b) Layout of the linear camera. Solid lines and dashed lines represent sectors and photoreceptors, respectively.

into account that evasive behavior generally requires a larger field of view than pursuit behavior. This mirrors binocular vision in predators and prey in nature, as also demonstrated in the context of morphological evolutionary robotics [8].

Two motor neurons were used to encode the desired speed and turning angle of the robot, respectively. The wheel speeds  $s_{\text{left}}(\tau_n)$  and  $s_{\text{right}}(\tau_n)$ , where  $-s_{\text{max}} \leq s_{\text{left}}, s_{\text{right}} \leq s_{\text{max}}$ , was translated from the motor neuron activations  $o_v(\tau_{n-1})$  and  $o_w(\tau_{n-1})$  (both in the range  $[0, 1]$ ) at time step  $\tau_{n-1}$  (Section A1.2). Neuron  $o_v$  encoded the baseline robot speed (Equations 9 and 10). Neuron  $o_w$  encoded the robot turning-rate: When  $o_w$ 's activation started to deviate from  $\approx 0.5$ , one of the wheels decreased or reversed speed, causing the robot to turn (Equation 11 and Figure 16). Turning neuron  $o_w$ 's falloff rate  $\psi = 1$  determined how abruptly a robot's velocity changed when turning. Calculated wheel speeds were matched with an empirical sample table.

$$s_{\text{left}}(\tau_n) = \begin{cases} s_{\text{max}} o_v(\tau_{n-1}) \lambda(o_w(\tau_{n-1})) & \text{if } o_w(\tau_{n-1}) < 0.5 \\ s_{\text{max}} o_v(\tau_{n-1}) & \text{otherwise} \end{cases} \tag{9}$$

$$s_{\text{right}}(\tau_n) = \begin{cases} s_{\text{max}} o_v(\tau_{n-1}) \lambda(o_w(\tau_{n-1})) & \text{if } o_w(\tau_{n-1}) > 0.5 \\ s_{\text{max}} o_v(\tau_{n-1}) & \text{otherwise} \end{cases} \tag{10}$$

$$\lambda(x) = -2^{2\psi+1} (x - 0.5)^{2\psi} + 1 \tag{11}$$

### A1.2 Robot Neural Network Controller

We provided both predator and prey robots with a perceptron neural network controller with an input layer  $I$  of 13 neurons  $x_i$  ( $i = 1, \dots, 13$ ) and an output layer with two neurons  $y_j$  ( $j = 1, 2$ ) (Figure 17).




Figure 16. Schematization of the relation between the activation state of the turning motor neuron ( $o_w(\tau_{n-1})$ ) and the desired speed of the left ( $s_{\text{left}}(\tau_n)$ ) and right ( $s_{\text{right}}(\tau_n)$ ) wheels, when  $s_{\text{max}} = 1$ ,  $o_v = 1$ , and  $\psi = 1$ . The dashed and solid lines show the left and right wheel speeds, respectively.




Figure 17. The NN architecture. Arrows indicate layers are fully connected. Graded neurons denote an evolvable responsiveness parameter.

The input and output layers were fully connected by 26 connections with weights  $w_{ij} \in \mathcal{W}$ . We parameterized two motor neurons each with a responsiveness parameter  $\beta_j \in B$ . Connection weights and responsiveness parameters varied in the range  $[-5.0, 5.0]$ . The state of the sensor neurons was scaled in the range  $[0.0, 1.0]$ . Sensors and motor neurons were updated at each time step  $\tau_n$ . The activation of the output neurons was computed according to

$$y_j(\tau_n) = \sigma \left( \beta_j \left( \sum_i^{|I|} w_{ij} x_i(\tau_n) \right) \right) \tag{12}$$

using the sigmoid function  $\sigma(\alpha) = 1/(1 + e^{-\alpha})$ , and where  $\beta_j \in B$  denotes  $y_j$ 's responsiveness parameter, which determines the steepness of the sigmoid. The output nodes correspond to the motor neurons ( $y_1 = \omega_v$  and  $y_2 = \omega_\omega$ ) and determine the rotation speed of the wheels (Section A1.1). The 28 NN parameters (encoding 26 connection weights and two responsiveness parameters) were genetically encoded and evolved (Section A1.3).

### A1.3 Evolutionary Algorithm

We initialized (random genotypes) and evolved two populations (predator and prey), each of size  $\mu = 20$  in the GENU (Table 5 and Algorithm 2). Subsequently, we evolved only the predator in a second series against a fixed, alternating opponent (prey) during the EMU (Section 3.5). In the EMU, we extracted the prey individuals from the GENU using the cluster analysis (Section 3.3). The NN's connection weights  $\mathcal{W}$  and output neurons' responsiveness parameters  $B$  (Section A1.2) were encoded in a genotype  $\mathbf{g} = \langle \gamma_0, \gamma_1, \dots, \gamma_{28} \rangle$  of 28 ( $= |\mathcal{W}| + |B|$ ) integer values 0, 1, ..., 255 (see also Definition 1), and subsequently evolved by the EA.

First, we assigned pairs of predator and prey fitness scores by evaluating them exhaustively ( $\mu^2$ ) in the simulated environment (Section A1.1); each individual was evaluated against each opponent of the current generation during  $\mu$  trials. At the beginning of each trial, the position and the orientation of the robots were randomly initialized. The state of the robots was updated in discrete, 100-ms time steps  $0 \leq \tau_n \leq \tau_{\max}$ . Each trial ended at  $\tau_{\max} = 500$  or when the prey got caught. In order to instantiate the robot controllers, each allele was linearly transformed into the NN's range of  $[-5, 5]$ .

The fitness function scored for the predator was inversely related to the time required to catch the prey (i.e., 0 when the predator failed to catch the prey), and for the prey was proportional to the

Table 5. The EA summarized.

|                       |   |
|-----------------------|---|
| Representation        | Integer-valued (0, 1, . . . , 255) vector (i.e., 8-bit numbers) |
| Population size       | $\mu = 20$  |
| Recombination         | None  |
| Mutation              | Random bitwise, Gaussian  |
| Parent selection      | Exhaustive  |
| Survivor selection    | $\mu$ times $\mu + 1$   |
| Termination condition | 500 or 1000 generations   |

time required to be caught (i.e., 1 when the prey was not caught). The predator and prey fitnesses therefore complementarily summed to 1 for a single trial:

$$fit_{\text{trial}}(p) = \begin{cases} 1 - \frac{\tau_p}{\tau_{\max}} & \text{if } p \in PD \\ \frac{\tau_p}{\tau_{\max}} & \text{if } p \in PY \end{cases} \quad (13)$$

The total fitness of an individual was computed by averaging the fitnesses obtained during all trials (possibly spanning multiple generations).

After we evaluated all parents in a generation, each one produced an offspring by means of mutation (Figure 18b, pointer 1). This offspring was then evaluated against the entire opponent population (Figure 18b, pointer 2; the opponent's fitness was not updated in this phase, to avoid introducing fitness inflation), resulting in  $2\mu^2$  trials ( $\mu$  per offspring, two species). For the first generation, no offspring were produced, in order to establish a more solid fitness baseline.

---

**Algorithm 2.** The evolutionary algorithm.

---

```

for generation ← 0 to 500 do // Evaluate parents
  foreach Predator do
    foreach Prey do
      pdFitness, pyFitness ← trial(Predator, Prey);
      totalPdFitness ← totalPdFitness + pdFitness;
      totalPyFitness ← totalPyFitness + pyFitness;
    if generation > 0 then // Create and evaluate offspring
      foreach Predator, Prey do // First predator population, then prey
        foreach Individual do
          Child ← generateChild(Individual);
          foreach Opponent do
            childFitness ← trial(Child, Opponent);
            totalChildFitness ← totalChildFitness + childFitness;
          // Only consider parents from current species
          WorstParent ← getWorstParent();
          if avgChildFitness ≥ WorstParent.avgFitness then
            replace(WorstParent, Child);

```

---



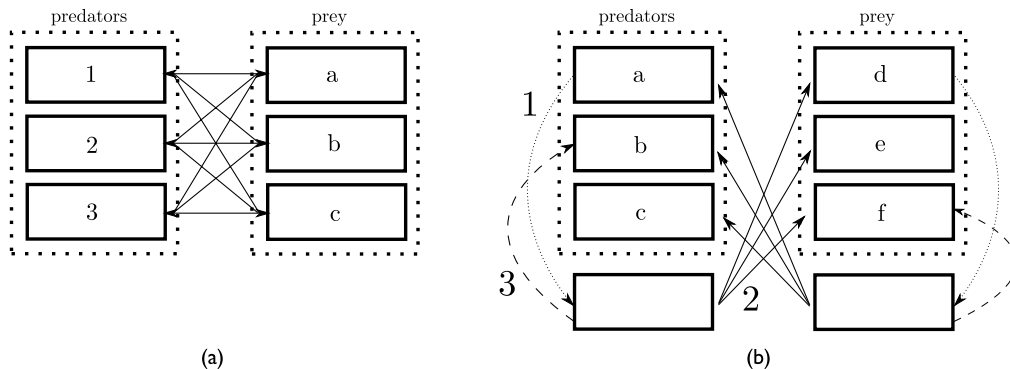


Figure 18. The EA in a schematic visualization. Shown are the steps in progressing from one generation to the next. The dotted boxes represent the predator and prey populations, while the solid ones represent individuals. The solid arrows mark the pairs of opponent species playing in the trials. (a) The first step in producing a new generation is to establish the fitness of both parent populations. Alphanumeric indicate the sequence of evaluation (1 versus a, 1 versus b, 1 versus c, 2 versus a, etc.). Solid arrows indicate evaluations. (b) The second step involves evaluating the offspring. The dotted arrows (1) show the creation of offspring (only the first individual is shown here), in the order of alphabetic lettering. The solid arrows (2) show how that offspring is evaluated against all opponent parents. The dashed arrows (3) show how the newly generated offspring might replace a conspecific parent individual.

During the GENU we used a random bitwise mutator. This allowed for relatively unrestrained yet unguided exploration of the genospace while simultaneously not deviating too much from random genetic mutations in nature. Each gene in the genotype got base-2 converted from an integer to a vector of eight bits. Every bit had a 0.02 chance of being flipped, averaging to  $\approx 4.2$  mutations in an entire genotype. Depending on the representational position  $i$  of a mutated bit, the change in the gene’s integer value it represented would increase or decrease with a value of  $2^i$  when that bit got mutated, where  $0 \leq i \leq 7$ . In the EMU we used a Gaussian mutator to provide more control in that the genospace would now be explored in a more structured way while also allowing experimental manipulation of the speed with which the subject species could change opponent specialization. Contrary to the bitwise mutator, the Gaussian one ( $\mu = 0, \sigma = 4$ ) was applied to every gene and then rounded to the nearest integer value 0, 1, ..., 255.

Offspring replaced the worse parent if the latter was outperformed (Figure 18b, pointer 3). The other offspring were discarded.<sup>7</sup> We repeated the selection and reproduction process for 500 generations in the GENU, and for 1000 generations in the EMU.

### Appendix 2: Generality of Seeds

There is some variation in cycling dynamics between seeds (Figure 19), but they all show similar patterns to that of seed 6 that was the focus of our study. Compared to Figure 5a, Figure 19a and c show more high-frequency cycling, and Figure 19c also shows a transition into a phase where the prey is slightly dominant (darker coloring in the later generations). Figure 19c and d show more chaotic dynamics; in Figure 19b the earlier predator generations seem notably susceptible to variations in prey strategies (indicated by the black and white blocks at the top of the plot). All cases, however, show rectilinear banding associated with cycling dynamics and so should be perfectly amenable to the methods we propose in this study.

<sup>7</sup> Note that the parent population was updated continuously, on an offspring-by-offspring basis. Thus, no distinct offspring population exists, and no survivors are actively being selected from distinct parent and offspring populations. Thus, it might happen that an offspring got promoted to the parent population, and was immediately replaced by another offspring from the same generation (if the latter outperformed the former). Most importantly, however, both parents and offspring have a fair chance to survive into the next generation, maintaining genetic diversity as well as emphasizing fitness-based selection criteria.

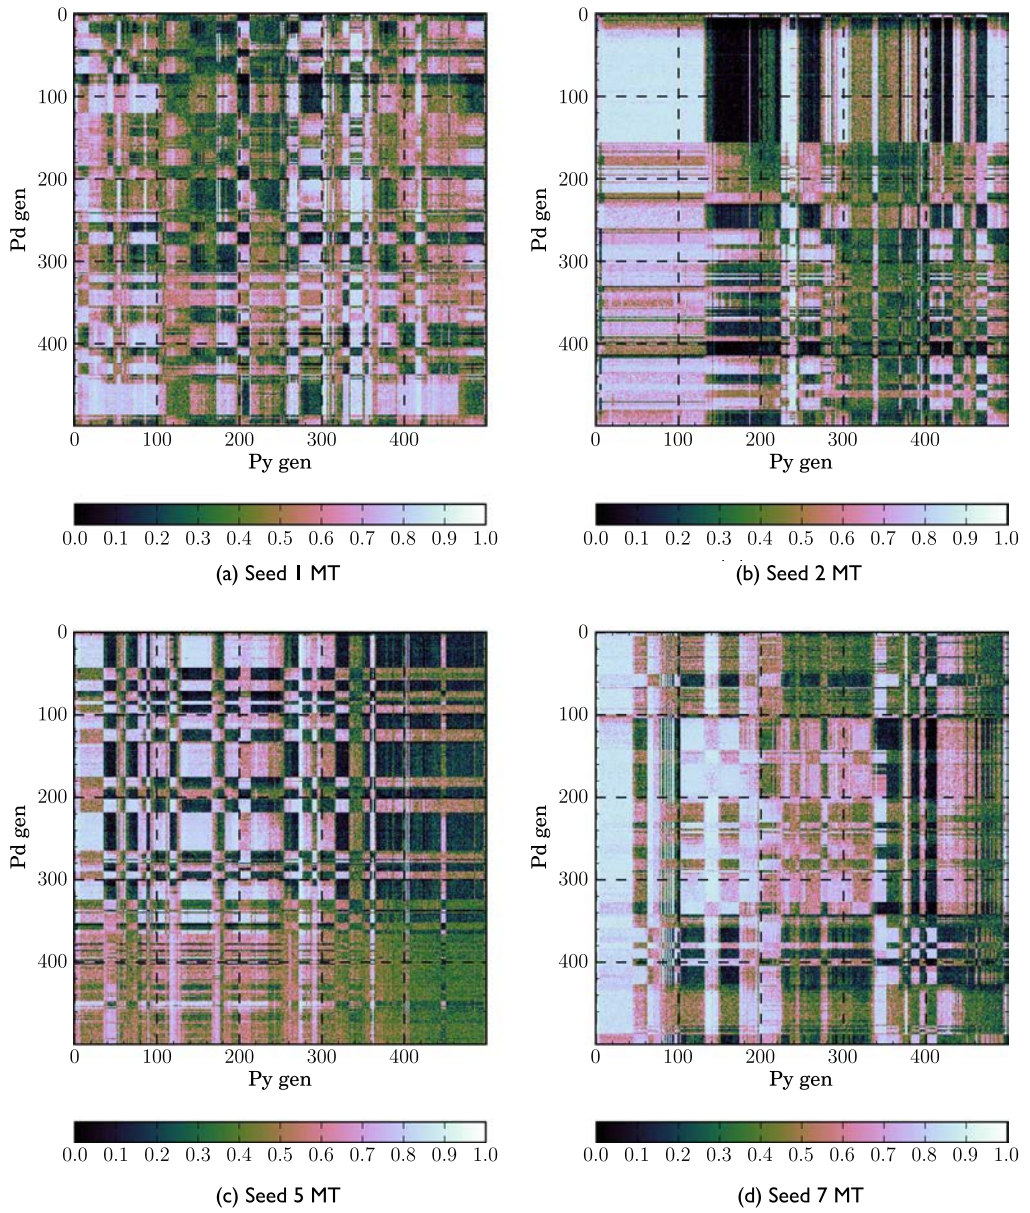


Figure 19. MTs of four seeds other than seed 6. Seeds shown were selected on the basis of a fair representation of the diversity within all ten seeds.