

Co-evolving predator and prey robots: Do ‘arms races’ arise in artificial evolution?

Stefano Nolfi* Dario Floreano~

*Institute of Psychology, National Research Council
Viale Marx 15, Roma, Italy
stefano@kant.irmkant.rm.cnr.it

~LAMI - Laboratory of Microcomputing
Swiss Federal Institute of Technology
EPFL, Lausanne, Switzerland
Dario.Floreano@epfl.ch

Abstract

Co-evolution (i.e. the evolution of two or more competing populations with coupled fitness) has several features that may potentially enhance the power of adaptation of artificial evolution. In particular, as discussed by Dawkins and Krebs [3], competing populations may reciprocally drive one another to increasing levels of complexity by producing an evolutionary “arms race”. In this paper we will investigate the role of co-evolution in the context of evolutionary robotics. In particular, we will try to understand in what conditions co-evolution can lead to “arms races”. Moreover, we will show that in some cases artificial co-evolution has a higher adaptive power than simple evolution. Finally, by analyzing the dynamics of co-evolved populations, we will show that in some circumstances well adapted individuals would be better advised to adopt simple but easily modifiable strategies suited for the current competitor strategies rather than incorporate complex and general strategies that may be effective against a wide range of opposing counter-strategies.

1. Introduction

Co-evolution (i.e. the evolution of two or more competing populations with coupled fitness) has several features that may potentially enhance the adaptation power of artificial evolution¹.

First, the co-evolution of competing populations may produce increasingly complex evolving challenges. As discussed by Dawkins and Krebs [3] competing populations may reciprocally drive one another to increasing levels of complexity by producing an evolutionary “arms race”. Consider for example the well-studied case of two co-evolving populations of predators and prey [16]: the success of predators implies a failure of the prey and conversely, when prey evolve to overcome the predators they also create a new challenge for them. Similarly, when the predators overcome the new prey by adapting to them, they create a new challenge for the prey. Clearly the continuation of this process may produce an ever-greater level of complexity (although this does not necessarily happen, as we will see below). As Rosin and Belew [20] point out, it is like producing a *pedagogical* series of challenges that gradually increase the complexity of the corresponding solutions. For an example of how a progressive increase in the complexity of the training sample may allow a neural network to learn a complex task that cannot otherwise be learned see [4].

¹ By adaptive power we mean the ability to solve complex tasks. In the context of predator and prey, this means the ability to catch a very efficient prey or to escape a very efficient predator.

This nice property overcomes the problem that if we ask evolution to find a solution to a complex task we have a high probability of failure while if we ask evolution to find a solution first to a simple task and then for progressively more complex cases, we are more likely to succeed. Consider the predators and prey case again. At the beginning of the evolutionary process, the predator should be able to catch its prey which have a very simple behavior and are therefore easy to catch; likewise, prey should be able to escape simple predators. However, later on, both populations and their evolving challenges will become progressively more and more complex. Therefore, even if the selection criterion remains the same, the adaptation task may become progressively more complex.

Secondly, because the performance of the individual in a population depends also on the individual strategies of the other population which vary during the evolutionary process, the ability for which individuals are selected is more general² (i.e., it has to cope with a variety of different cases) than in the case of an evolutionary process in which co-evolution is not involved. The generality of the selection criterion is a very important property because the more general the criterion, the larger the number of ways of satisfying it (at least partially) and the greater the probability that better and better solutions will be found by the evolutionary process.

Let us again consider the predator and prey case. If we ask the evolutionary process to catch one individual prey we may easily fail. In fact, if the prey is very efficient, the probability that an individual with a randomly generated genotype may be able to catch it is very low. As a consequence, all individuals will be scored with the same null value and the selective process cannot operate. On the contrary, if we ask the evolutionary process to find a predator able to catch a variety of different prey, it is much more probable that it will find an individual in the initial generations able to catch at least one of them and then select better and better individuals until one predator able to catch the original individual prey is selected.

Finally, competing co-evolutionary systems are appealing because the ever-changing fitness landscape, due to changes in the co-evolving species is potentially useful in preventing stagnation in local minima. From this point of view, co-evolution may have consequences similar to evolving a single population in an ever-changing environment. Indeed the environment changes continuously given the fact that the co-evolving species is part of the environment of each evolving population.

Unfortunately a continuous increase in complexity is not guaranteed. In fact, co-evolving populations may cycle between alternative class of strategies that, although they do not produce advantages in the long run, may produce a temporary improvement over the co-evolving population. Imagine, for example, that in a particular moment population A adopts the strategy A_1 which gives population A an advantage over population B which adopts strategy B_1 . Imagine now that there is a strategy B_2 (genotypically similar to B_1) that gives population B an advantage over strategy A_1 . Population B will easily find and adopt strategy B_2 . Imagine now that there is a strategy A_2 (genotypically similar to A_1) that provides an adaptive advantage over strategy B_2 . Population A will easily find and adopt strategy A_2 . Finally imagine that previously discovered strategy B_1 provides an advantage over strategy A_2 . Population B will come back to strategy B_1 . At this point also population A will come back to strategy A_1

² We will use the term 'general strategy' or 'general solution' to indicate selected individuals able to cope with different tasks. In the context of predator and prey we will indicate with the term 'general' the strategy adopted by a predator which is able to catch a large number of prey adopting different, not necessarily complex, strategies.

(because, as explained above, it is effective against strategy B_1) and the cycle of the same strategies will be repeated over and over again (Figure 1).

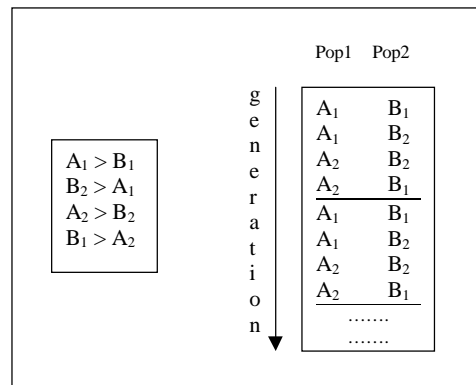


Figure 1. The same strategies (A_1 and A_2 in population A) and (B_1 and B_2 in population B) may be selected over and over again throughout generations as is shown in the right hand side of the figure if the interaction between them looks like what is represented on the left side of the Figure. In this case the repeated cycle corresponds to 4 different combinations of strategies.

Notice how the cycling may involve two or more different strategies for each population but also two or more different groups of strategies.

Parker [18] was the first to hypothesize that parent-offspring and intersexual ‘arms races’ may end up in cycles. Dawkins and Krebs [3] noted how this hypothesis can be applied to asymmetric arms races in general.

Of course this type of phenomenon may cancel out all the previously described advantages because, if co-evolution quickly falls into a cycling phase, the number of different solutions discovered might be quite limited. In fact, there is no need to discover progressively more complex strategies. It is sufficient to re-discover previously selected strategies that can be adopted with a limited number of changes. Moreover, it should be noted that cycling is not the only possible cause which may prevent the emergence of ‘arms races’.

In this paper we will investigate the role of co-evolution in the context of evolutionary robotics. In particular, we will try to understand in which conditions, if any, co-evolution can lead to “arms races” in which two populations reciprocally drive one another to increasing levels of complexity.

After introducing our experimental framework in section 2.1 and 2.2 we will describe the result of a first basic experiment in section 2.3. As we will see, the innovations produced in this first experiments may easily be lost because the evolutionary process quickly falls into a cycling phase in which the same type of solutions are adopted over and over by the two co-evolving populations. In section 2.4 we will show how the tendency to cycle between the same type of strategies may be reduced by preserving all previously discovered strategies and by using all of them to test the individual of the current population (we will refer to this technique as ‘Hall of Fame’ co-evolution). We will also point out that this technique, which is biologically implausible, has its own drawbacks. In section 2.5 in fact, we will see how ‘Hall of Fame’ co-evolution does not necessary produce better performance than simple co-evolution. On the contrary, in the case of the experiment described in this section, simple co-evolution tend to outperform ‘Hall of Fame’ co-evolution. In section 2.5 we will also see how ‘arms races’ can emerge and indeed produce better and better

solutions. In section 2.6 we will see how increasing the environmental richness may decrease the probability to fall in cycling phases. Finally, in section 2.7 we will see how co-evolution can solve problems that evolution alone cannot. In other words, we will show how in some circumstances co-evolution has an higher adaptive power than evolution of a single population.

2. Co-evolving predator and prey robots

Several researchers have investigated co-evolution in the context of predators and prey in simulation [11, 12, 1, 2]. More recently, we have tried to investigate this framework first by using realistic simulations based on the Khepera robot [7, 8] and subsequently the real robots [9]. Up to now, we replicated on the real robots the experiments which will be described in section 2.3. By comparing the results obtained in simulation with those obtained with the real robots in this case we did not observe any significant difference in term of performance and co-evolutionary dynamic. Although, not all the strategies observed in simulation were also observed in the experiments performed in the real environment. In this case, in fact, the presence of much larger noise filtered out brittle solutions [9].

In this section, we will first describe our experimental framework and the results obtained in a simple case. Then, we will describe other experimental conditions more suitable to the emergence of ‘arms races’ between the two competing populations.

2.1 The experimental framework

As often happens, predators and prey belong to different species with different sensory and motor characteristics. Thus, we employed two Khepera robots, one of which (the Predator) was equipped with a vision module while the other (the Prey) had a maximum available speed set to twice that of the predator. The prey has a black protuberance, which can be detected by the predator everywhere in the environment (see Figure 2). The two species could evolve in a square arena 47 x 47 cm in size with high white walls so that predator could always see the prey (within the visual angle) as a black spot on a white background.

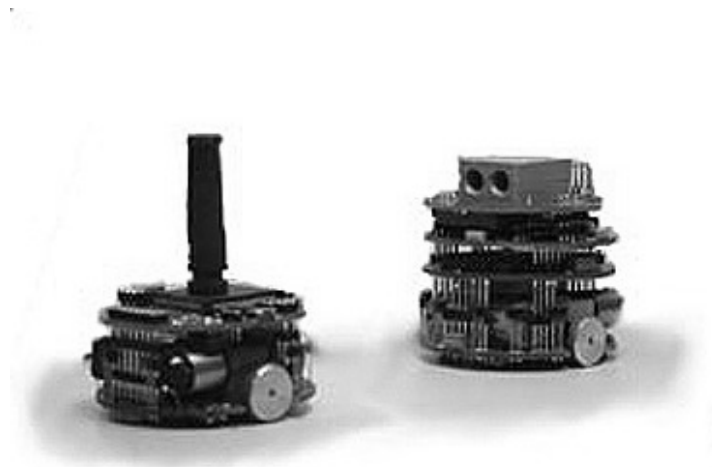


Figure 2. Prey and predator (left to right).

Both individuals were provided with eight infrared proximity sensors (six on the front side and two on the back) which had a maximum detection range of 3-4 cm in our environment. For the predator we considered the K213 module of Khepera which is an additional turret that can be plugged in directly on top of the basic platform. It consists of a 1D-array of 64 photoreceptors which provide a linear image composed of 64 pixels of 256 gray-levels each, subtending a view-angle of 36° . However the K213 module also allows detection of the position in the image corresponding to the pixel with minimal intensity. We used this facility by dividing the visual field into five sectors of about 7° each corresponding to five simulated photoreceptors (see Figure 3). If the pixel with minimal intensity lay inside the first sector, then the first simulated photoreceptor would become active; if the pixel lay inside the second sector, then the second photoreceptor would become active, etc. From the motor point of view, we set the maximum wheel speed in each direction to 80mm/s for the predator and 160mm/s for the prey.

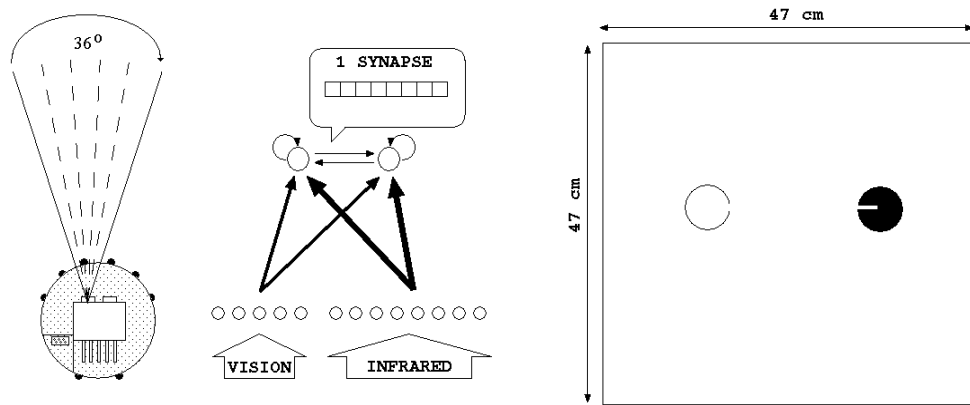


Figure 3. Left and center: details of simulation of vision, of neural network architecture, and of genetic encoding. The prey differs from the predator in that it does not have 5 input units for vision. Eight bits code each synapse in the network. Right: Initial starting position for prey (left, empty disk with small opening corresponding to frontal direction) and predator (right, back disk with line corresponding to frontal direction) in the arena. For each competition, the initial orientation is random.

In line with some of our previous work [6], the robot controller was a simple perceptron comprising two sigmoid units with recurrent connection at the output layer. The activation of each output unit was used to update the speed value of the corresponding wheel every 100ms. In the case of the predator, each output unit received connections from five photoreceptors and from eight infrared proximity sensors. In the case of the prey, each output unit received input only from 8 infrared proximity sensors, but its activation value was multiplied by 2 before setting the wheel speed.

In order to keep things as simple as possible and given the small size of the parameter set, we used direct genetic encoding [22]: each parameter (including recurrent connections and threshold values of output units) was encoded using 8 bits. For the same reason, the architecture was kept fixed, and only synaptic strengths and output units threshold values were evolved. Therefore, the genotype of the predator was $8 \times (30 \text{ synapses} + 2 \text{ thresholds})$ bits long while that of prey was $8 \times (20 \text{ synapses} + 2 \text{ thresholds})$ bits long. It should be noted that the type of architecture we selected may constraint the type of solutions which will be obtained during the evolutionary process.

In principle, it would be better to evolve both the architecture and the weights at the same time. However, how to encode the architecture of the network into the genotype is still an open and complex research issue in itself. Moreover, even more complex genotype-to-phenotype mappings (which would allow the evolution of the architecture too) might still constrain the evolutionary process in certain, albeit different ways.

Two populations of 100 individuals each were co-evolved for 100 generations. Each individual was tested against the best competitors of the previous generations (a similar procedure was used in [21, 2]). In order to improve co-evolutionary stability, each individual was tested against the best competitors of the ten previous generations (on this point see also below). At generation 0, competitors were randomly chosen within the same generation, whereas in the other 9 initial generations they were randomly chosen from the pool of available best individuals of previous generations.

For each competition, the prey and the predator were always positioned on a horizontal line in the middle of the environment at a distance corresponding to half the environment width, but always at a new random orientation. The competition ended either when the predator touched the prey or after 500 motor updates (corresponding to 50 seconds at maximum on the physical robot). The fitness function for each competition was simply 1 for the predator and 0 for the prey if the predator was able to catch the prey and, conversely 0 for the predator and 1 for the prey if the latter was able to escape the predator. Individuals were ranked after fitness performance in descending order and the best 20 were allowed to reproduce by generating 5 offspring each. Random mutation (bit substitution) was applied to each bit with a constant of probability $pm=0.02^3$.

For each set of experiments we ran 10 replications starting with different randomly assigned genotypes.

In this paper we will refer to data obtained in simulation. A simulator developed and extensively tested on Khepera by some of us was used [15].

2.2 Measuring adaptive progress in co-evolving populations

In competitive co-evolution the reproduction probability of an organism with certain traits can be modified by the competitors, that is, changes in one species affect the reproductive value of specific trait combinations in the other species. It might thus happen that progress achieved by one lineage is reduced or eliminated by the competing species. This phenomenon, which is referred to as the “Red Queen Effect” [19], makes it hard to monitor progress by taking measures of the fitness throughout generations. In fact, because fitnesses are defined relative to a co-evolving set of traits in the other individuals, the fitness landscapes for the co-evolving individuals vary. As a consequence, for instance, periods of stasis in the fitness value of the two populations may correspond to a period of tightly-coupled co-evolution.

In order to avoid this problem, different measure techniques have been proposed. Cliff and Miller [1] have devised a way of monitoring fitness performance by testing the

³ The parameters used in the simulations described in this paper are mostly the same as in the simulation described in [7]. However, in these experiments we used a simpler fitness formula (a binary value instead of a continuous value proportional to the time necessary for the predator to catch the prey). Moreover, to keep the number of parameters as small as possible, we did not use crossover. In the previous experiments, in fact, we did not notice any significant difference in experiments conducted with different crossover rates.

performance of the best individual in each generation against all the best competing ancestors which they call CIAO data (Current Individual vs. Ancestral Opponents).

A variant of this measure technique has been proposed by us and has been called Master Tournament [7]. It consists in testing the performance of the best individual of each generation against each best competitor of all generations. This latter technique may be used to select the best solutions from an optimization point of view (see [7]). Both techniques may be used to measure co-evolutionary progress (i.e. the discovery of more general and effective solutions).

2.3 Evolution of predator and prey robots: a simple case.

The results obtained by running a set of experiments with the parameter described in section 2.1 are shown below. Figure 4 represents the results of the Master Tournament, i.e the performance of the best individual of each generation tested against all best competitors from that replication. The top graph represents the average result of 10 simulations. The bottom graph represents the result of the best run.

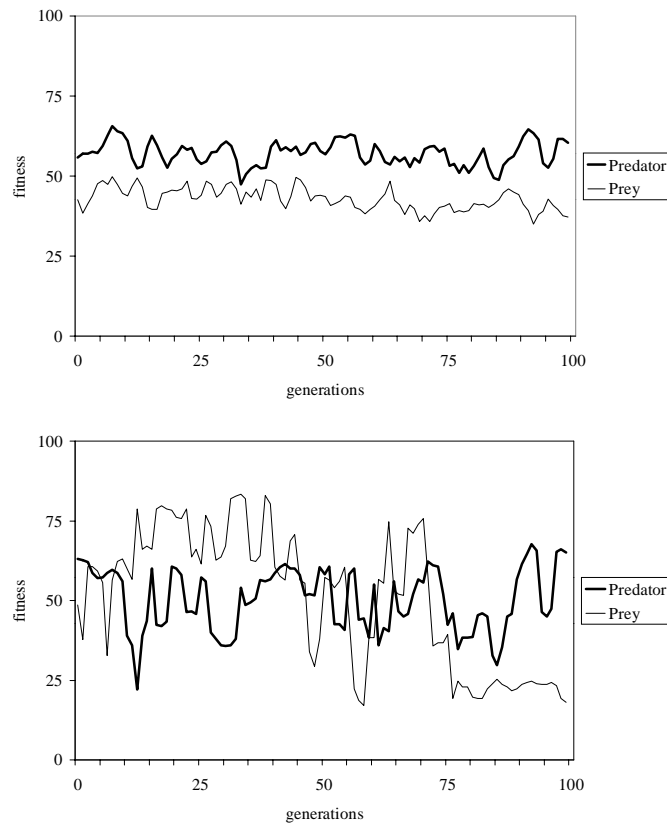


Figure 4. Performance of the best individuals of each generation tested against all the best opponents of each generation (Master Tournament). Performance may range from 0 to 100 because each individual is tested once against each best competitor of 100 generations. The top graph shows the average result of 10 different replications. The bottom graph shows the result in the best replication (i.e. the simulation in which predators and prey attain their best performance). Data were smoothed using rolling average over three data points.

These results show that, at least in this case, phases in which both predators and prey produce increasingly better results are sometimes followed by sudden drops in performance (see the bottom graph of Figure 4). As a consequence, if we look at the average result of different replications in which increase and drop phases occur in different generations, we observe that performance does not increase at all throughout generations (see the top graph of Figure 4). In other words the efficacy and generality of the different selected strategies does not increase evolutionarily. In fact, individuals of later generations do not necessarily score well against competitors of much earlier generations (see Figure 5, right side). Similar cases have been described in [2, 21].

The ‘arms races’ hypothesis would be verified if, by measuring the performance of each best individual against each best competitor, a picture approximating that shown on the left side of Figure 5 could be obtained. In this ideal situation, the bottom-left part of the square, which corresponds to the cases in which predators belong to more recent generations than the prey, is black (i.e. the predator wins). Conversely, the top right part of the square, which corresponds to the cases in which the prey belong to more recent generations than the predators, is white (i.e. the prey wins). Unfortunately, what actually happens in a typical run is quite different (see right part of Figure 5). The distribution of black and white spots does not differ significantly in the two sub-parts of the square.

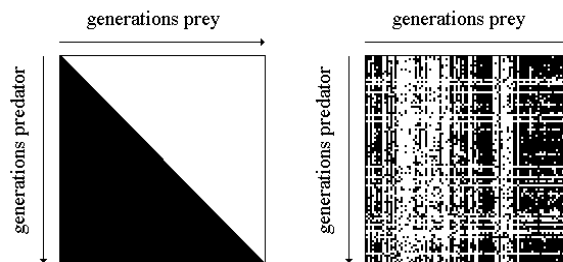


Figure 5. Performance of the best individuals of each generation tested against all the best opponents of each generation. The black dots represent individual tournaments in which the predators win while the white dots represent tournaments in which the prey wins. The picture on the left represents an ideal situation in which predators are able to catch all prey of previous generations and the prey are able to escape all predators of previous generations. The picture on the right represents the result for the best simulation (the same shown in Figure 4).

This does not imply that the co-evolutionary process is unable to find interesting solutions as we will show below (see also [7]). This merely means that effective strategies may be lost instead of being retained and refined. Such good strategies, in fact, are often replaced by other strategies that, although providing an advantage over the current opponents, are much less general and effective in the long run. In particular, this type of process may lead to the cycling process described in section 1 in which the same strategies are lost and re-discovered over and over again.

The cycling between the same class of strategies is actually what happens in these experiments. If we take a look at the qualitative aspects of the behavior of the best individuals of successive generations we see that in all replications, evolving predators discover and rediscover two different classes of strategies: (A_1) track the prey and try to catch it by approaching it; (A_2) track the prey while remaining more or less in the same area and attacking the prey only on very special occasions (when the prey is in a particular position relative to the predator). Similarly the prey cycles between two

classes of strategies: (B_1) stay still or hidden close to a wall waiting for the predator and eventually trying to escape when the IR sensors detect the predator (notice that predators usually keep themselves away from walls to avoid crashes); (B_2) move fast in the environment, avoiding both the predator and the walls.

Now, as in Figure 1, the strategy A_1 is generally effective against B_1 , in fact the predator will reach the prey if the prey does not move too much and has a good chance of succeeding given that the prey can only detect predators approaching from certain directions because of the uneven distributions of the infrared sensors around the body. Strategy B_2 is effective against strategy A_1 because the prey is faster than the predator and so, if the predator tries to approach a moving fast prey, it has little chance of catching it. Strategy A_2 is effective against strategy B_2 because, if the prey moves fast in the environment, the predator may be able to catch it easily by waiting for the prey itself to come close. Finally, strategy B_1 is very effective against strategy A_2 . In fact if the predator does not approach the prey and the prey stays still, the prey will never risk being caught. This type of relation between different strategies produces a cycling process similar to that described in Figure 1.

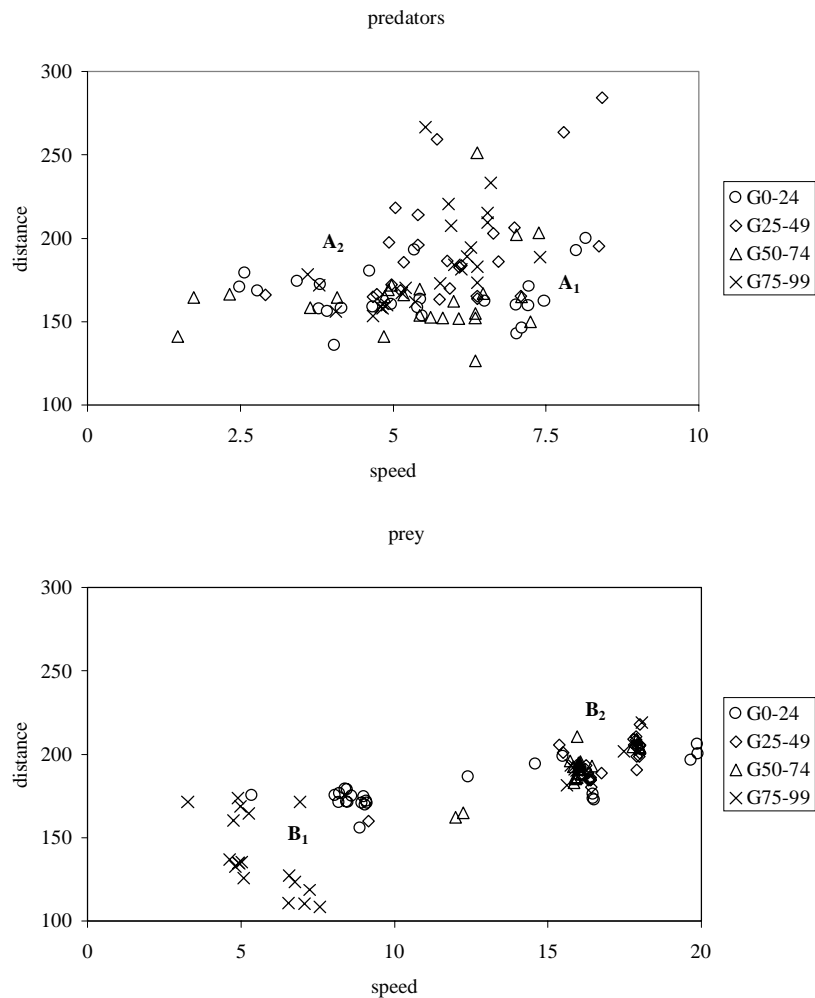


Figure 6. Position of the best predator and prey of successive generations in the phenotype space (top and bottom graph, respectively). The Y and X axes represent the average speed (i.e. computed as the absolute value of the algebraic sum of the two wheels) and the average distance (i.e. the distance in mm between

two competing individuals), respectively. Individuals of different generations are shown with different labels but, for graphic reasons, individuals of each 25 successive generations are shown with the same label. Average speed and distance have been computed by testing the best individual of each generation against the best competitor of each generation.

The cycling process is driven in general by prey which after adopting one of the two classes of strategies for several generations suddenly shift to the other strategy. This switch forces predators to shift their strategy accordingly. This is also shown in Figure 5 (right side) in which the reader can easily see that the main source of variation is on the X-axis which represents how performance vary for prey of different generations.

What actually happens in the experiments is not so simple as in the description we have just given because of several factors: (1) the strategies described are not single strategies but classes of similar strategies. So for example there are plenty of different ways for the predator to approach the prey and different ways may have different probabilities of being successful against the same opposing strategies; (2) the advantage or disadvantage of each strategy against another strategy varies quantitatively and is probabilistic (each strategy has a given probability of beating a competing strategy); (3) populations at a particular generation do not include only one strategy but a certain number of different strategies although they tend to converge toward a single one; (4) different strategies may be easier to discover or re-discover than others.

However the cycling process between the different classes of strategies described above can be clearly identified. By analyzing the behavior of the best individuals of the best simulation (the same as that described in Figures 3 and 4), for example, we can see that the strategy B_2 discovered and adopted by prey at generation 21 and then abandoned after 15 generations is rediscovered and re-adopted at generation 58 and then at generation 98. Similarly the strategy A_2 , first discovered and adopted by the predator at generation 10 and then abandoned after 28 generations for strategy A_1 , is then rediscovered at generation 57. Interestingly, however, prey also discover a variation of strategy B_1 that includes also some of the characteristics of strategy B_2 . In this case, prey move in circles waiting for the predator as in strategy B_1 . However, as soon as they detect the predator with their IR sensors, they start to move quickly exploring the environment as in strategy B_2 . This type of strategy may in principle be effective against both strategies A_1 and A_2 . However sometimes prey detect the predator too late, especially when the predator approaches the prey from its left or right rear side which is not provided with IR sensors.

This cycling dynamics is shown also in Figure 6 which represents the position of the best predator and prey of successive generations in a two-dimensional phenotype space. To represent the phenotype space we considered two measures that are representative of the different strategies: the average speed and the average distance from the competitor (these two dimensions have been subjectively chosen to illustrate the qualitative features of the behaviors that we observed). In the case of the prey, two different classes of phenotype corresponding to the strategies B_1 and B_2 can be clearly identified. In the case of predators, on the other hand, a continuum of strategies can be observed between strategies that can be classified as typically A_1 or A_2 . In both cases, however, examples of each class of strategies can be found in the first and in successive generations, indicating that the same type of strategy is adopted over and over again.

2.4 Testing individuals against all discovered solutions

In a recent article, Rosin and Belew [20], in order to encourage the emergence of ‘arms races’ in a co-evolutionary framework, suggested saving and using as competitors all the best individuals of previous generations:

So, in competitive coevolution, we have two distinct reasons to save individuals. One reason is to contribute genetic material to future generations; this is important in any evolutionary algorithm. Selection serves this purpose. Elitism serves this purpose directly by making complete copies of top individuals.

The second reason to save individuals is for purposes of testing. To ensure progress, we may want to save individuals for an arbitrarily long time and continue testing against them. To this end, we introduce the ‘Hall of Fame’, which extends elitism in time for purposes of testing. The best individual from every generation, is retained for future testing.

From Rosin and Belew [20], pp. 8.

This type of solution is of course implausible from a biological point of view. Moreover, we may expect that, by adopting this technique, the effect of the co-evolutionary dynamic will be progressively reduced throughout generations with the increase in number of previous opponents. In fact, as the process goes on, there is less and less pressure to discover strategies that are effective against the opponent of the current generation and greater and greater pressure to develop solutions capable of improving performance against opponents of previous generations.

However, as the authors show, in some cases this method may be more effective than a ‘standard’ co-evolutionary framework in which individuals compete only with opponents of the same or of the previous generation. More specifically, we think, it may be a way to overcome the problem of the cycling of the same strategies. In this framework in fact, ad hoc solutions that compete successfully against the opponent of the current generation but do not generalize to opponents of previous generations cannot spread in evolving populations.

We applied the *Hall of Fame* selection regime to our predator and prey framework and measured the performance of each best individual against each best competitor (Master Tournament). Results are obtained by running a new set of 10 simulations in which each individual is tested against 10 opponents randomly selected from all previous generations (while in the previous experiments we selected 10 opponents from the immediately preceding generations). All the other parameters remain the same. As shown in Figure 7 and 8, in this case, we obtain a progressive increase in performance.

Figure 7 shows how in this case the average fitness of the best individuals tested against all best competitors progressively increases throughout generations, ultimately attaining near to optimal performances. Figure 8 shows how this is accomplished by being able to beat most of the opponents of previous generations. The results do not exactly match the ideal situation described in Figure 5 (left side) in which predators and prey are able to beat all individuals of previous generations. In the best simulation described in Figure 7 (bottom graph) and Figure 8, for example, there are two phases in

which prey are unable to beat most of the predators of few generations before. The general picture, however, approximates the ideal one.

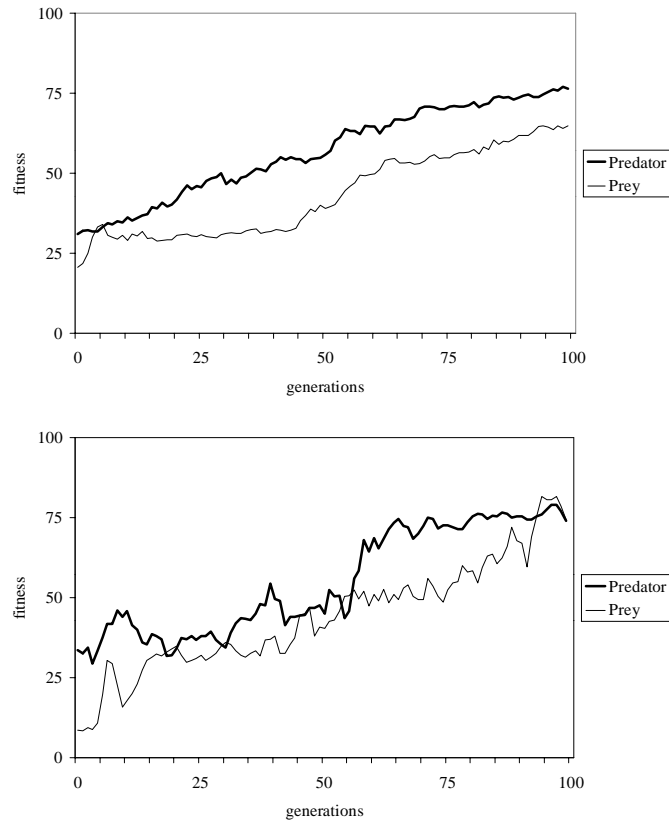


Figure 7. Performance of the best individuals of each generation tested against all the best opponents of each generation (Master Tournament). The top graph shows the average result of 10 different replications. The bottom graph shows the result in the best replication (i.e. the simulation in which predators and prey attain the best performance). Data were smoothed using a rolling average over three data points.

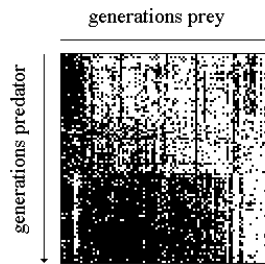


Figure 8. Performance of the best individuals of each generation tested against all the best opponents of each generation. Black dots represent individual tournaments in which the predators win while white dots represent tournaments in which the prey wins. Result for the best simulation (the same shown in Figure 7).

If we look at the individual strategies selected throughout generations in these experiments, we see that they are of the same class of those described in the previous section. However, in this case, the strategies are evolutionarily more stable (i.e. in general they are not suddenly replaced by another strategy of a different class). This enables the co-evolutionary process to progressively refine current strategies instead of

cycling between different classes of strategies, restarting each time from the same initial strategy.

The fact that individuals are tested against quite different strategies (i.e. competitors randomly selected from all previous generations) should enable the evolutionary process to find strategies that are more general (i.e. that are effective against a larger number of counter-strategies) than those obtained in the experiments described in the previous section. To verify this hypothesis we tested the best 10 predators and prey obtained with ‘standard’ co-evolution against the best 10 predators and prey obtained with ‘Hall of Fame’ co-evolution (i.e. the best predator and prey of each replication were selected). As can be seen, ‘standard’ individuals have a higher probability of defeating ‘standard’ individuals than ‘Hall of Fame’ individuals (Figure 9, left side). Similarly, ‘Hall of Fame’ individuals have a higher probability of defeating ‘standard’ individuals than ‘Hall of Fame’ individuals (Figure 9, right side). Although, variability in different replication is high, these results indicate that, in this case, ‘Hall of Fame’ co-evolution tends to produce more general solutions than ‘standard’ co-evolution. However, differences in performance are not as great as one could expect from the trends of the Master Tournaments in the two conditions, which are quite different (we will be return to this later on).

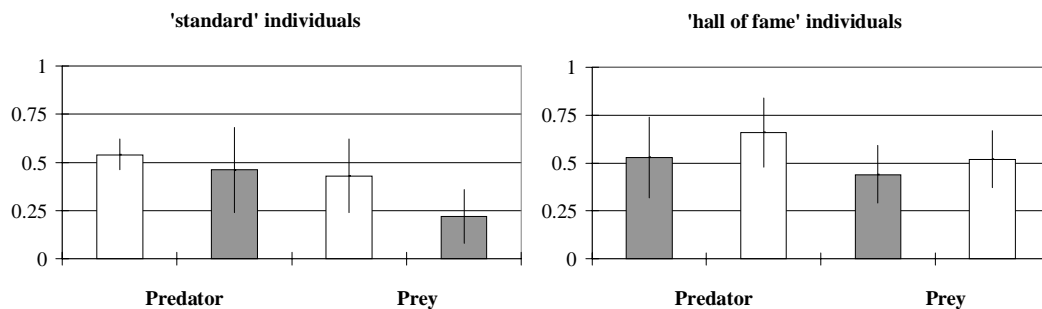


Figure 9. Graphs showing the average performance of the best individuals obtained with standard and with ‘Hall of Fame’ co-evolution (left and right side, respectively). Performance obtained by testing individuals against standard and ‘Hall of Fame’ competitors is shown using white and gray histograms, respectively. Vertical bars indicate standard deviation. Individuals are selected by picking the predator and the prey with the best score in the master tournament of each replication. Y-axis indicates the percentage of defeated competitors. Each column is the results of a separate test (individuals start with different randomly assigned orientations).

2.5 How the length of ‘arms races’ may vary in different conditions.

One of the simplification we adopted in our experiments is that the sensory-motor system of the two species is fixed. However, as we will show below, the structure of the sensory system can affect the course of the co-evolutionary process and the length of the ‘arms races’.

One thing to consider in our experiments is that the prey has a limited sensory system that enables it to perceive predators only at a very limited distance and not from all relative directions (there are no IR sensors able to detect predators approaching from the rear-left and rear-right side). Given this limitation, the prey cannot improve its strategy above a certain level. They can compete with co-evolving predators only by suddenly changing strategy as soon as predators select an effective strategy against them. However, if we increase the richness of the prey’s sensory system we may expect

that the prey will be able to overcome well adapted predators by refining their strategy instead of radically changing their behavior.

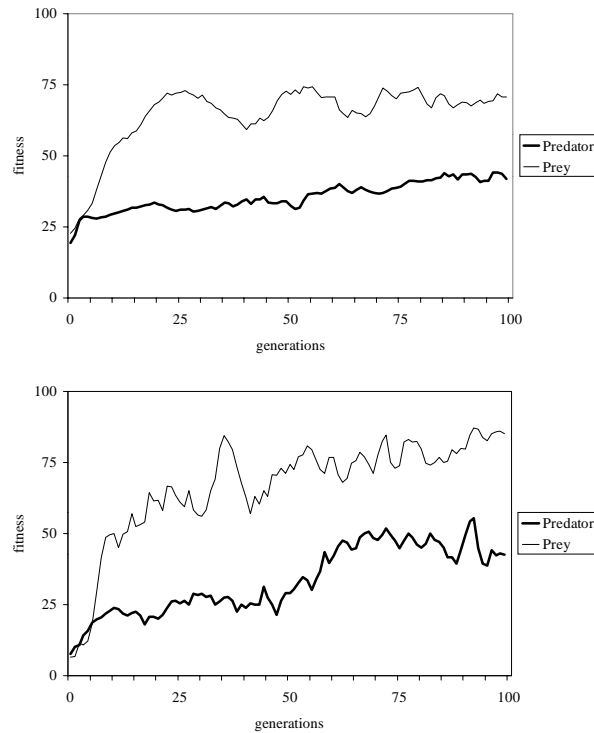


Figure 10. Experiments with standard co-evolution (i.e., not Hall of Fame) and prey with camera too. Performance of the best individuals of each generation tested against all the best opponents of each generation (Master Tournament). The top graph shows the average result of 10 different replications. The bottom graph shows the result in the best replication (i.e. the simulation in which predators and prey attain the best performance). Data were smoothed using a rolling average over three data points.

To investigate this hypothesis we ran a new set of simulations in which the prey also was provided with a camera able to detect the predators' relative position. For the prey we considered another turret under development at LAMI, which consists of an 1D-array of 150 photoreceptors which provide a linear image composed of 150 pixels of 256 gray levels each subtending a view-angle of 240° [14]. We chose this wider camera because the prey, by escaping the predators, will only occasionally perceive opponents in their frontal direction. As, in the case of predators, the visual field was divided into five sectors of 48° corresponding to five simulated photoreceptors. As a consequence, in this experiment, both predator and prey are controlled by a neural network with 13 sensory neurons. Moreover, in this case, both predator and prey could see their competitors as a black spot against a white background. 'Standard' co-evolution was used (i.e. individuals were tested against the best competitors of the 10 previous generations and not against competitors selected from all previous generations as in the experiments described in section 2.4). All the other parameters remained the same.

If we measure the average performance of the best predators and prey of each generation tested against all the best opponents of each generation (Master Tournament)

we see that, although the prey in general overcomes predators⁴, a significant increase in performance throughout generations is observed in both populations (Figure 10). Figure 11 shows the performance against each competitor for the best replication also shown in Figure 10 (bottom graph).

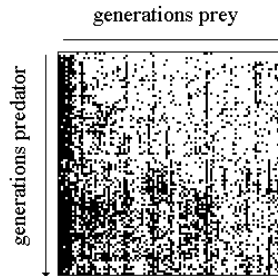


Figure 11. Performance of the best individuals of each generation tested against all the best opponents of each generation. Black dots represent individual tournaments in which the predators win while white dots represent tournaments in which the prey wins. Result for the best simulation (the same as that shown in Figure 10).

These results show how by changing the initial conditions (in this case by changing the sensory system of one population) ‘arms races’ can continue to produce better and better solutions in both populations for several generations without falling into cycles.

Interestingly, in their simulations in which also the sensory system of the two co-evolving populations was under evolution, Cliff and Miller observed that “.. pursuers usually evolved eyes on the front of their bodies (like cheetahs), while evaders usually evolved eyes pointing sideways or even backwards (like gazelles).” [2, pp. 506]⁵.

To investigate whether also in this case ‘Hall of Fame’ co-evolution outperforms standard co-evolution we ran another set of experiments identical to those described in this section but using the ‘Hall of Fame’ selection regime. Figure 12 shows the Master Tournament measures obtained on average in this second set of experiments. As expected, performance measured using Master Tournament increased even more in this second set of simulations (in particular, a larger increase in performance throughout generations can be observed in the case of prey). However, if we test individuals obtained with standard co-evolution against individuals obtained with ‘Hall of Fame’ co-evolution we find that the latter do not outperform the standard individuals (see

⁴ This may be due to the fact that in this and in the experiments which will be presented in the next sections the sensory system of the prey have been enriched with respect to the experiments described in section 2.3 and 2.4.

⁵ The authors did not provide enough data in their paper to understand whether their simulations fell into solution cycles. However, even though both the nervous system and the sensory system were under co-evolution in their case, it seems that Cliff and Miller did not observe any co-evolutionary progress toward increasingly general solutions. In fact, they report that ‘co-evolution works to produce good pursuers and good evaders through a pure bootstrapping process, but both types are rather specially adapted to their opponents’ current counter-strategies.’ [2, pp. 506]. However, it should be noted that there are several differences between Cliff and Miller experiments and ours. The fitness function used in their experiments, in fact, is more complex and includes additional constraints that try to force evolution in a certain direction (e.g. predators are scored for their ability to approach the prey and not only for their ability to catch it). Moreover, the genotype-to-phenotype mapping is much more complex in their cases and includes several additional parameters that may effect the results obtained.

Figure 13). On the contrary, individuals obtained with standard co-evolution tend to outperform individuals obtained with ‘Hall of Fame’ co-evolution.

As can be seen, individuals obtained by means of ‘standard’ co-evolution have a higher probability of defeating ‘Hall of Fame’ than ‘standard’ competitors (Figure 13, left side). Similarly, ‘Hall of Fame’ prey has a higher probability of defeating ‘Hall of Fame’ than ‘standard’ predators (Figure 13, right side). Notice however that also in this case, there is a high variability between different replications. Thus ‘standard’ individuals tend to be more effective than individuals obtained by ‘Hall of Fame’ co-evolution. However, ‘Hall of Fame’ predators are more likely to defeat ‘standard’ than ‘Hall of Fame’ prey.

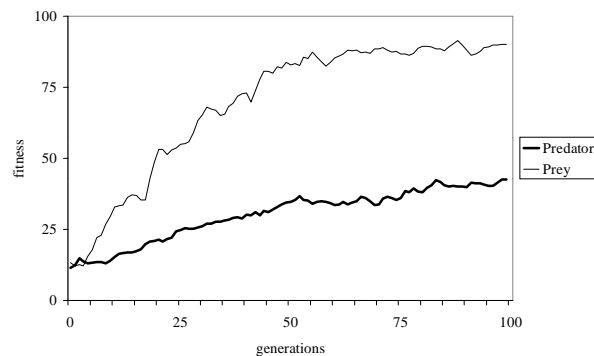


Figure 12. Experiments with Hall of Fame; the prey is equipped with vision too. Performance of the best individuals of each generation tested against all the best opponents of each generation (Master Tournament). Average result of 10 different replications. Data were smoothed using a rolling average over three data points.

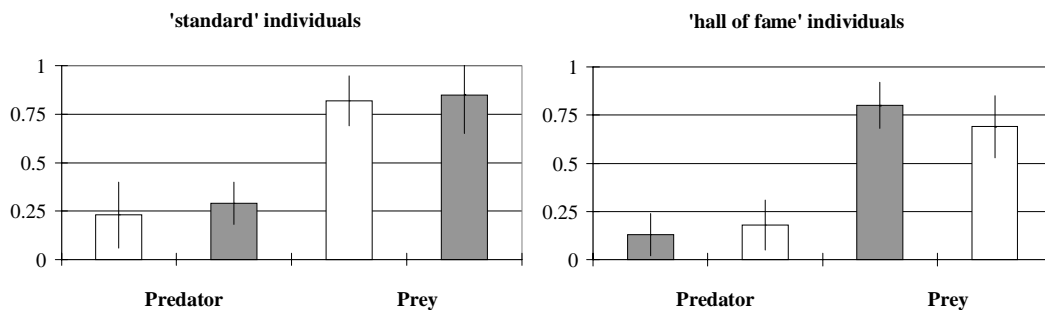


Figure 13. Graphs showing the average performance of the best individuals obtained with ‘standard’ and with ‘Hall of Fame’ co-evolution (left and right side, respectively). Performance obtained by testing individuals against ‘standard’ and ‘Hall of Fame’ competitors is shown using white and gray histograms, respectively. Vertical bars indicate standard deviation. Individuals are selected by picking the predator and the prey with the best score in the master tournament for each replication. Y-axis indicates the percentage of defeated competitors. Each column is the results of a separate test (individuals start with different randomly assigned orientations).

From these results it may be concluded that although the ‘Hall of Fame’ selection regime always tends to reduce the probability of falling into limit cycles (see Figure 12 which shows how progressively more general solutions are selected), it does not necessarily produce better solutions than ‘standard’ co-evolution (see Figure 13). When, as in the case described in this section, ‘standard’ co-evolution can produce arms races of significant length, it may outperform ‘Hall of Fame’ co-evolution. Furthermore, by continuing the evolutionary process, the ‘Hall of Fame’ selection regime might be even

less effective than ‘standard’ co-evolution given that, as mentioned earlier, the co-evolutionary dynamics tends to become progressively less effective throughout the generations with an increasing probability of opponents from previous generations being selected.

The fact that the structure of the sensory-motor system of the two species can significantly affect the course of the evolutionary process demonstrates the importance of using real robots instead of simulated agents. Real robots and natural agents, in fact, have sensory-motor apparatus which rely on measures of physical entities (light, speed, etc.), which have limited precision, which are affected by noise etc. Simulated agents instead, often adopt sensors and motors which have idealized characteristics (e.g. sensors which have infinite precision or which measure abstract entities such as distances between objects). Moreover, in the case of simulated agents, the experimenter may unintentionally introduce constraints which channel the experiment in a certain direction.

2.6 The role of environmental richness

In the previous section we showed how the length of ‘arms races’ (i.e. the number of generations in which co-evolving populations produce strategies able to defeat a progressively larger number of counter-strategies) may vary in different conditions. If both co-evolving populations can produce better and better strategies given their initial organization, ‘arms races’ may last several generations. Conversely, if one or both populations fail to improve their current strategy sufficiently, it is likely that the co-evolutionary dynamics will quickly lead to a limit cycle in which similar strategies are rediscovered over and over again.

Another factor that may prevent the cycling of the same strategies is the richness of the environment. In the case of co-evolution, competing individuals are part of the environment. This means that part, but not all, of the environment is undergoing co-evolution. We may hypothesize that the probability that a sudden shift in behavior will produce viable individuals is inversely proportional to the richness of the environment that is not undergoing co-evolution. Imagine, for example, that an ability acquired under co-evolution, such as the ability to avoid inanimate obstacles, involves a characteristic of the environment which is not undergoing co-evolution. In this case it is less likely that a sudden shift in strategy involving the loss of this ability will be retained. In fact, the acquired character will always have an adaptive value independently of the current strategies adopted by the co-evolving population. The same argument applies to any case in which one population is co-evolving against more than one other population. The probability of retaining changes involving a sudden shift in behavior will decrease because, in order to be retained, such changes would have to provide an advantage over both co-evolving populations.

To verify this hypothesis we ran a new set of experiments in which individuals experienced 5 different environments (i.e. they were tested for 2 epochs in each of the 5 environments instead of 10 epochs in the same environment). All the other parameters were the same as those described in section 2.1. In particular ‘standard’ co-evolution and prey without camera were used. Figure 14 shows the five environments which varied in shape and in the number and type of obstacles within the arena.

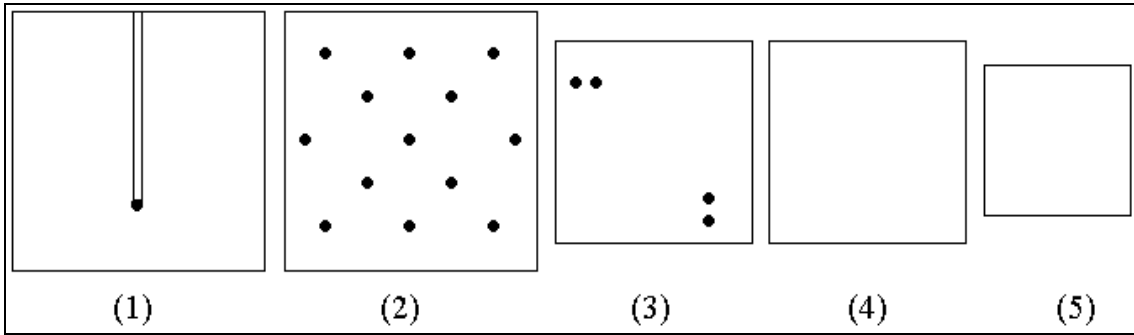


Figure 14. Five different environments. Environments 1 and 2 consisted of an arena measuring 60x60 cm, 3 and 4 by an arena of 47x47 cm, and 5 by an arena of 35x35 cm. Environments 1, 2, and 3 also contained cylindrical obstacles with a diameter of 2.3 cm. Finally, environment 1 included an inside wall 2 cm thick dividing the arena into two rooms connected by an open door about 14 cm wide.

If we measure the average performance of the best predators and prey of each generation tested against all the best opponents of each generation (Master Tournament) a significant increase in performance throughout generations is observed in some replications (see for example the bottom graph in Figure 15 showing the results of the best replication). The average results, however, show a slight increase only in the first 20 generations (see top graph in Figure 15).

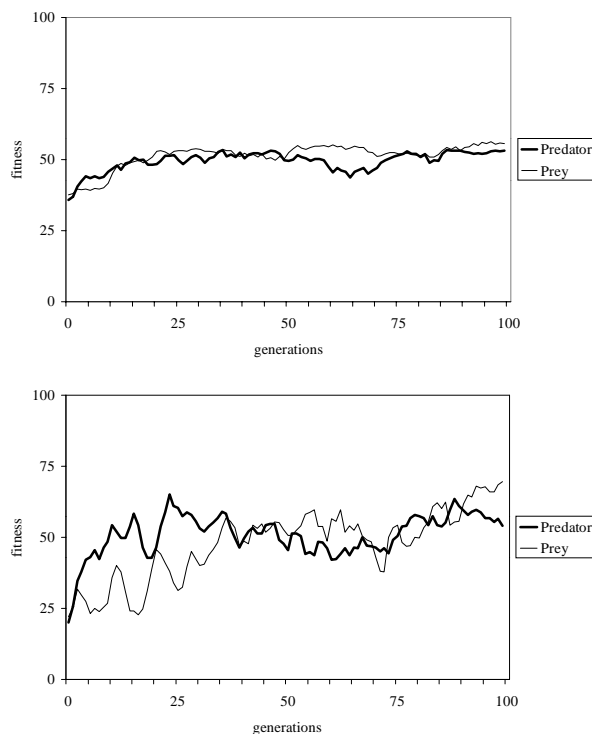


Figure 15. Performance of the best individuals of each generation tested against all the best opponents of each generation (Master Tournament). The top graph shows the average result of 10 different replications. The bottom graph shows the result in the best replication (i.e. the simulation in which predators and prey attain the best performance). Data were smoothed using a rolling average over three data points.

From these results it may be concluded that, as hypothesized at the beginning of this section, the richness of the environment may delay the convergence of the co-

evolutionary process towards a limit cycle. On the other hand we should consider that larger the number and importance of fixed constraints is, the lower the importance of the co-evolutionary dynamic may be. A rich continuum of possibilities between an extreme in which the environment is constituted only by the competitor and another extreme in which the competitor is one within several different sources of constraints should be considered. It may be that interesting co-evolutionary dynamics only arise in a given interval between these two extremes.

2.7 How co-evolution can enhance the adaptive power of artificial evolution

In the previous sections we showed how ‘arms races’ between co-evolving populations can arise. At this point we should try to verify whether co-evolution can really enhance the adaptive power of artificial evolution. In other words, can artificial co-evolution solve tasks that cannot be solved using a simple evolutionary process?

There are two reasons for hypothesizing that co-evolution can have an higher adaptive power than evolution. The first reason is that individuals evolving in co-evolutionary frameworks experience a larger number of different environmental events. The second and more important reason is related to the emergence of ‘arms races’.

To verify if co-evolution can produce solutions to problems that evolution alone is unable to solve we tried using simple evolution (i.e. an evolutionary process in which only a single population was evolved through selective reproduction and mutation). More specifically we tried to evolve predators able to catch the best prey obtained using artificial co-evolution. Likewise, we tried to evolve prey able to escape the best predators obtained by co-evolution. If evolution fails, at least in some cases, we may conclude that co-evolution is able to select better individuals than simple evolution. In other words, we may conclude that co-evolution is able to produce solutions to problems that evolution is unable to solve.

We ran several sets of simulations in which we tried to evolve individuals able to catch the best co-evolved prey and to escape the best co-evolved predator obtained from all the experiments described in the previous sections. The parameters used in the simulations were the same as those described in section 2.1 although only one population was subjected to the evolutionary process (the predator to escape or the prey to be caught remained identical over the entire evolutionary process). As a consequence individuals were tested for 10 epochs and 100 generations against exactly the same opponent. In all cases simple evolution was able to produce better and better individuals until optimal or close to optimal performance was obtained.

In order to produce a challenge that was too complex for simple evolution it was necessary to change the sensory system of the predator and prey and to use a more complex environment than the simple arena involved in most of the experiments described above.

We ran a new set of co-evolutionary experiments in which predator and prey were not equipped with cameras but were allowed to use the 8 ambient light sensors added in the basic Khepera module. Moreover, we included a 1 watt lamp on the top of both predator and prey so that each could obtain an indirect measure of the angle and distance of the other. The genotype of both predator and prey was 8 x (36 synapses + 2 thresholds) bits long. As environment, we used an arena measuring 60x60cm with 13 cylindrical obstacles (see environment 2 in Figure 14). The ‘standard’ selection regime was used. For all other parameters the same values described in section 2.1 were used.

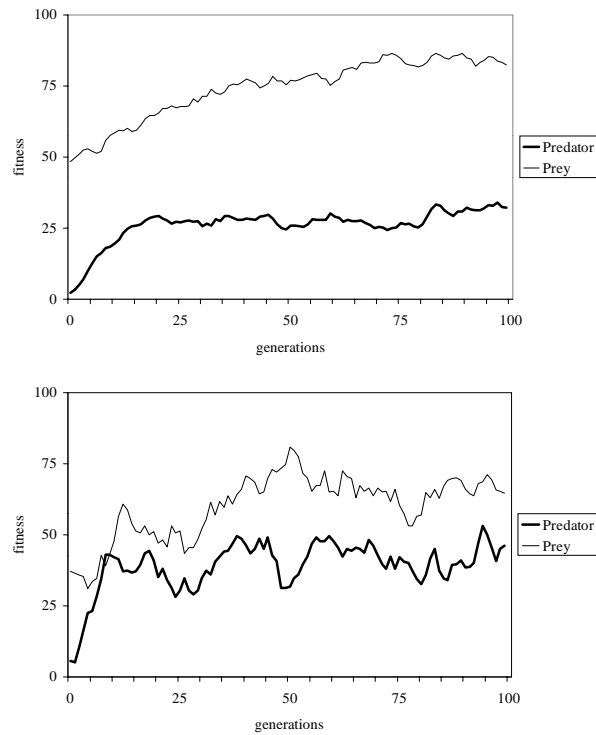


Figure 16. Experiments with the complex environment and predators and prey equipped with ambient light sensors ('standard' co-evolution). Performance of the best individuals of each generation tested against all the best opponents of each generation (Master Tournament). The top graph shows the average result of 10 different replications. The bottom graph shows the result in the best replication (i.e. the simulation in which predators and prey attain the best performance). Data were smoothed using a rolling average over three data points.

If we look at the Master Tournament performance, we can see how a significant increase may be observed both on average and in the case of the best replication. Moreover, we can see how, unlike the experiments described in the previous section, in this case predators of the very first generations have close to null performance. This implies that they are unable to catch most of the prey of succeeding generations.

We then ran a new set of experiments in which simple evolution (i.e. evolution of a single population against a fixed opponent) was used to select predators able to catch the best co-evolved prey obtained in the experiments just described. Similarly we used simple evolution to select prey able to escape the best co-evolved predators.

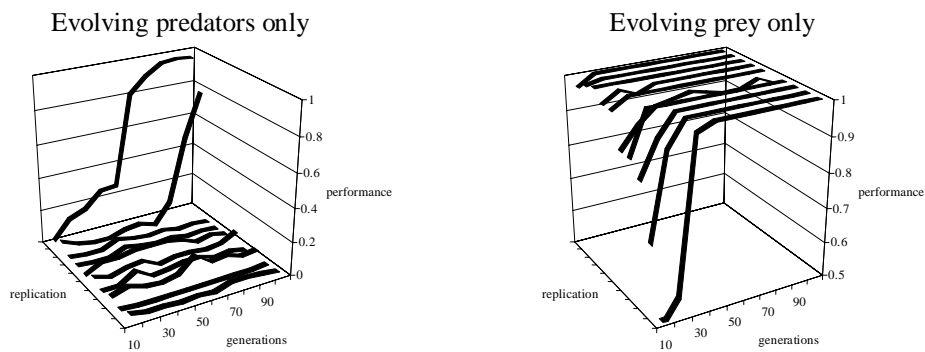


Figure 17. Performance of evolving predators and prey against fixed co-evolved opponents (left to right, respectively). Each graph shows the results of 10 experiments in which the best prey and predator obtained in the 10 co-evolutionary experiments described above were used as fixed opponents. Average

results of 10 succeeding generations are shown. Results have been sorted on the replication axis to enhance readability.

As can be seen in Figure 17, which shows the performance of the best evolving predators and prey, in 8 cases out of 10 simple evolution failed to select predators able to catch the co-evolved prey (the best individuals of 8 simulations are able to catch the prey less than 15% of the time). Conversely, the best co-evolved predators were able to catch the best co-evolved prey at least 25% of the time in 9 out of 10 simulations. The reason why simple evolution was not always successful was that predators of the first generations were scored with a null value (predators were able to catch the fixed prey once only occasionally, while their offspring usually failed to catch the same prey). As a consequence, the selection mechanism could not operate.

The fact that in this case co-evolution was able to produce more complex challenges than in the other experiments described in the previous section seems to be due to the ability of the prey to use the information coming from the ambient light sensors. Most of the co-evolved prey waited for the predator until it reached a distance of about 100 mm and only then did they start to escape. This allowed the prey to force the predator to follow them with little chance of catching them given the difference in speed. Moreover, it eliminated the risk of facing the predator head on by moving fast even when the predator was far away.

It should be noticed, however, that evolution of a single population can create very effective prey against the best of co-evolved predators (see Figure 17, right side). This implies that, in this case, it is always possible to find a simple strategy able to defeat each single individual predator. As we mentioned above, this is what happens for both predator and prey in all experiments described in the previous sections.

3. Discussion

Evolutionary Robotics is a promising new approach to the development of mobile robots able to act quickly and robustly in real environments. One of the most interesting features of this approach is that it is a completely automatic process in which the intervention of the experimenter is practically limited to the specification of a criterion for evaluating the extent to which evolving individuals accomplish the desired task. However, it is still not clear how far this approach can scale up.

From this point of view, one difficult problem is given by the fact that the probability that one individual within the initial generations is able to accomplish the desired task, at least in part, is inversely proportional to the complexity of the task itself. For complex tasks, it is very likely that all individuals of the initial generations are scored with the same zero value and, as consequence, the selection mechanism might result in a mere random process. We will refer to this problem as the *bootstrap problem*.

This problem arises from the fact that in artificial evolution people usually start from scratch (i.e. from individuals obtained with randomly generated genotypes). In fact, one possible solution to this problem is the use of ‘incremental evolution’. In this case, we start with a simplified version of the task and, after we get individuals able to solve such a simple case, we progressively move to more and more complex cases [5, 10, 13]. This type of approach can overcome the bootstrap problem, although it also has the negative consequence of increasing the amount of supervision required and the risk of introducing inappropriate constraints. In the case of incremental evolution in fact, the experimenter should determine not only an evaluation criterion but also a ‘pedagogical’

list of simplified criteria. In addition the experimenter should decide when to change the selection criterion during the evolutionary process. Some of these problems may arise also when the selection criterion includes rewards for sub-components of the desired behavior (although, in this case, the selection criterion is left unchanged throughout the evolutionary process) [17].

Another possible solution of the bootstrap problem is the use of co-evolution. Co-evolution of competing populations, in fact, may produce increasingly complex evolving challenges spontaneously without any additional human intervention. Unfortunately, no continuous increase in complexity is guaranteed. For example, the co-evolutionary process may fall into a limit cycle in which the same solutions are adopted by both populations over and over again (we will refer to this problem as the *cycling problem*). What happens is that at a certain point one population, in order to overcome the other population, finds it more useful to suddenly change its strategy instead of continuing to refine it. This is usually followed by a similar rapid change of strategy in the other population. The overall result of this process is that most of the characters previously acquired are not appropriate in the new context and therefore are lost. However, later on, a similar sudden change may bring the two populations back to the original type of strategy so that the lost characters are likely to be rediscovered again.

The effect of the cycling problem may be reduced by preserving all the solutions previously discovered for testing the individuals of the current generations [20]. However, this method has drawbacks that may affect some of the advantages of co-evolution. In fact, as the process goes on there is less and less pressure to discover strategies that are effective against the opponent of the current generation and increasing pressure to develop solutions able to improve performance against opponents of previous generations which are no longer under co-evolution. While in some cases testing individuals against a sample of all previously selected competitors may produce better performance (as shown in section 2.4), in other cases this might not be true. Indeed it may even result in less effective individuals (see section 2.5).

We believe that the cycling problem, like the local minima problem in gradient-descent methods (i.e. the risk of getting trapped in a sub-optimal solution when all neighboring solutions produce a decrease in performance), is an intrinsic problem of co-evolution that cannot be eliminated completely. However, as we have shown in sections 2.5 and 2.7 the cycling problem does not always affect the co-evolutionary dynamics so strongly as to prevent the emergence of ‘arms races’. When both co-evolving populations can produce better and better strategies, ‘arms races’ may last several generations and produce progressively more complex and general solutions. On the other hand, if one or both populations cannot improve their current strategy sufficiently, the co-evolutionary dynamics will probably quickly lead to a limiting cycle in which similar strategies are rediscovered over and over again.

Despite the cycling problem, it can be shown that in some cases co-evolution may succeed in producing individuals able to cope with very effective competitors (by selecting the competitors at the same time) while simple evolution is unable to do so (see section 2.7). The reason for this is that co-evolution, by selecting also the competitors that determine the complexity of the task, is not affected by the ‘bootstrap problem’. On the other hand, when simple evolution is faced with fixed co-evolved competitors, it may happen that the genetic operators are unable to generate any individual able to defeat the competitor, even in a few cases. As a consequence the selection process does not work.

Moreover, it should be noted that some factors may limit the cycling problem. One of these factors is, as we have shown in section 2.6, the richness of the environment. In the case of co-evolution, competing individuals are part of the environment. This means that part, but not all of the environment, is undergoing co-evolution. The probability that a sudden shift in behavior will produce viable individuals is inversely proportional to the richness of the environment that is not undergoing co-evolution. In fact, if an acquired ability involves a characteristic of the environment which is not undergoing co-evolution it is less likely that a sudden shift in strategy involving the loss of such an ability will be retained. Indeed the acquired character will always have an adaptive value independently of the strategies adopted by the co-evolving population. This effect may be particularly significant in the case of natural evolution in which, in general, the environment is much richer than in the case of the experiments performed in artificial evolution.

Another factor that may limit the effect of the cycling problem is ontogenetic plasticity. Plastic individuals, in fact, may be able to cope with different classes of strategies adopted by the second population by adapting to the current opponent's strategy during their lifetime, thus reducing the adaptive advantage of a sudden shift in behavior which causes the cycling problem. The experiments described in this paper did not support this issue (none of the experiments described in this paper involved ontogenetic plasticity). On the effects of some forms of ontogenetic plasticity within a co-evolutionary framework see [8].

3.1 A dynamical view of adaptation

We have thus been able to show that, at least in one case, co-evolution can produce a strategy that is too complex for simple evolution to cope with (section 2.7). However, in the other 3 cases examined (see section 2.3, 2.5, and 2.6) evolution was quickly able to select individuals that proved very effective against such complex strategies. In particular this happened also with the strategies obtained in the experiments described in section 2.5 and 2.6 in which Master Tournament measures clearly indicated a progress throughout generations. This means that, even though more and more general strategies were selected in these experiments through co-evolution, it was always easy to select individuals able to defeat these strategies by starting from scratch. Further proof of this is that if we look at the performance of the best individuals of the last generations we see that, even though they score increasingly better against individuals of previous generations on average, they may sometimes be defeated by individuals of many generations before (see for example Figures 8 and 11).

These results point to the conclusion that in certain tasks it is always possible to find a simple strategy that is able to defeat another single, albeit complex and general, strategy (although such simple strategy is a specialized strategy, i.e. it is able to defeat only that individual complex and general strategy and, of course, other similar strategies). If this is really true, in other words, if completely general solutions do not exist in some cases, we should re-consider the 'cycling problem'. From this point of view, the fact of a co-evolutionary dynamics leading to a limit cycle in which the same type of solutions are adopted over and over again should not be considered as a failure but as an optimal solution. We cannot complain that co-evolution does not find a more general strategy able to cope with all the strategies adopted by the co-evolving population during a cycle if such general strategies do not exist. The best that can be

done is to select the appropriate strategy for the current counter-strategy, which is actually what happens when the co-evolutionary dynamics ends in a limit cycle.

More generally we can predict that co-evolution will lead to a progressive increase in complexity when complete general solutions (i.e. solutions which are successful against all the strategies adopted by previous opponents) exist and can be selected by modifying the current solutions. Conversely, if complete general solutions do not exist or the probability of generating them is too low, co-evolution may lead to a cycling dynamics in which solutions appropriate to the strategy of the co-evolving population but which can also easily be transformed so to match other strategies will be selected. In other words, when general solutions cannot be found, it becomes important for each evolving population to be able to dynamically change its own strategy into one of a set of appropriate strategies. From the individuals' point of view, we may say that individuals with a predisposition to change in certain directions will be selected.

Interestingly, one can argue that these dynamics may be an ideal situation for the emergence of ontogenetic adaptation. The ability to adapt during one's lifetime to the opponent's strategy would in fact produce a significant increase in the adaptation level of a single individual because ontogenetic adaptations are much faster than phylogenetic ones. Therefore, we may hypothesize that, when a co-evolving dynamics leads to a limiting cycle, there will be a high selective pressure in the direction of ontogenetic adaptation. At the same time, the cycling dynamics will create the conditions in which ontogenetic adaptation may more easily arise because, as we have seen, individuals with a predisposition to change in certain directions will be selected. It is plausible to argue that, for such individuals, a limited number of changes during ontogeny will be able to produce the required behavioral shift. In other words, we can argue that it will be easier for co-evolving individuals to change their behavior during their lifetime in order to adopt strategies already selected by their close ancestors thanks to the cycles occurring in previous generations.

Notice that although an individual which has a single strategy able to defeat a set of counter-strategies and an individual which possesses a set of different strategies able to defeat the same set of counter-strategies are equivalent at a certain level of description there are some important differences (to distinguish the two cases let us call the former 'full-general' and the latter 'plastic-general'). The plastic-general individual should be able to select the right strategy given the current competitor. In other words, should be able to adapt through ontogenetic adaptation. From this point of view the full-general individual will be more effective because it does not require such adaptation process and may provide immediately the correct answer to the current competitor. On the other hand, as we said above, it may be that in certain conditions a fully-general individual cannot be selected because a fully-general strategy does not exist or because it is too improbable that the evolutionary process is able to find it. In this case the only option left is that of plastic-general solutions. However also a plastic-general individual is difficult to obtain because it implies that such individual should be able to display a variety of different strategies and because it should also be able to select the right strategy at the right moment (depending on the behavior of the current competitor). What is interesting is that, when a fully-general strategy cannot be found, co-evolution will fall into a cycling dynamics in which a set of 'specialistic' strategy will be discovered over and over again. Now, because during this phase the best thing individuals can do to improve the chances of survival of their offspring is to produce offspring which can evolutionary change their strategy as fast as possible (in other words individuals which

have a predisposition to change in certain directions) we may expect that the length of the cycles will be progressively shortened throughout successive generations. At this point, we might speculate that co-evolution might favor the emergence of individuals with the ability to modify their behavior during lifetime in the most appropriate directions (the same directions for which a predisposition to change have been genetically acquired), if the genotype could allow for some type of phenotypic modification.

Of course, this is only a hypothesis. The only results we have from the experiments we described in this paper is that in most of our experiments simple ‘specialist’ solutions can be found while fully-general solutions cannot. It remain to be ascertained if plastic-general solutions (i.e. solutions which consists of a set of simple ‘specialist’ solutions and a mechanism for selecting the right one during lifetime) can be selected. Preliminary evidences have been described in [8].

Acknowledgments

We thank the anonymous referees for their valuable suggestions.

References

1. Cliff, D., & Miller, G. F. (1995). Tracking the red queen: Measurement of adaptive progress in co-evolutionary simulations. In F. Moran, A. Moreno, J. J. Merelo & P. Chacon (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, Berlin: Springer Verlag.
2. Cliff, D., & Miller, G. F. (1996). Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes, M. Mataric, J-A Meyer, J. Pollack, H. Roitblat & S. Wilson (Eds.), *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cambridge, MA: MIT Press-Bradford Books.
3. Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London B*, 205, 489-511.
4. Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunket, K. (1996). *Rethinking Innateness. A Connectionist Perspective on Development*. Cambridge, MA: MIT Press.
5. Floreano, D. (1993). Emergence of home-based foraging strategies in ecosystems of neural networks. In J. Meyer, H.L. Roitblat & S.W. Wilson (Eds.), *From Animals to Animats II: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, Cambridge, MA: MIT Press-Bradford Books.
6. Floreano, D., & Mondada, F. (1994). Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In D. Cliff, F. Husband, J-A. Meyer & S. W. Wilson (Eds.), *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, Cambridge, MA: MIT Press-Bradford Books.
7. Floreano, D., & Nolfi, S. (1997). God save the red queen! Competition in co-evolutionary robotics. In J. R. Koza, D. Kalyanmoy, M. Dorigo, D. B. Fogel, M. Garzon, H. Iba & R. L. Riolo (Eds.), *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pp. 398-406, San Francisco, CA: Morgan Kaufmann.
8. Floreano, D., & Nolfi, S. (1997). Adaptive behavior in competing co-evolving species. In P. Husband & I. Harvey (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, Cambridge, MA: MIT Press.
9. Floreano, D., Nolfi, S., & Mondada, F. (1998). Competitive co-evolutionary robotics: From theory to practice. In R. Pfeifer, B. Blumberg & H. Kobayashi (Eds.), *Proceedings of the Fifth International Conference of the Society for Adaptive Behavior (SAB98)*, Cambridge, MA: MIT Press
10. Harvey, I., Husbands, P., & Cliff, D. (1994). Seeing the light: artificial evolution, real vision. In D. Cliff, P. Husbands, J-A. Meyer & S.W. Wilson (Eds.), *From Animals to Animats III: Proceedings of*

the Third International Conference on Simulation of Adaptive Behavior, Cambridge, MA: MIT Press-Bradford Books.

11. Koza, J. R. (1991). Evolution and co-evolution of computer programs to control independently-acting agents. In J.A. Meyer & S. Wilson (Eds.), *From Animals to Animats: Proceeding of the First International Conference on Simulation of Adaptive Behavior*, Cambridge, MA: MIT Press.
12. Koza, J. R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, Cambridge, MA: MIT Press.
13. Gomez, F., & Miikkulainen, R. (1997). Incremental evolution of complex general behavior, *Adaptive Behavior*, 5, 317-342.
14. Landolt, O. (1996). Description et mise en oeuvre du chip ED084V2A, *Technical Report 16-11-95*, Centre Suisse d'Electronique et Microtechnique, Switzerland.
15. Miglino, O., Lund, H. H., & Nolfi, S. (1995). Evolving mobile robots in simulated and real environments. *Artificial Life*, 4 (2), 417-434.
16. Miller, G. F., & Cliff, D. (1994). Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations, *Technical Report CSRP311*, School of Cognitive and Computing Sciences, University of Sussex.
17. Nolfi, S. (1997). Evolving non-trivial behaviors on real robots: a garbage collecting robot, *Robotics and Autonomous System*, 22, 187-198.
18. Parker, G.A. (1979). Sexual selection and sexual conflict. In M. A. Blum, & N. A. Blum (Eds.), *Reproductive competition and sexual selection*. New York: Academic Press.
19. Ridley, M. (1993). *The Red Queen: Sex and the evolution of human nature*. London: Viking Press.
20. Rosin, C. D. & Belew, R. D. (1997). New methods for competitive coevolution, *Evolutionary Computation*, 5 (1), 1-29.
21. Sims, K. (1995). Evolving 3D morphology and behavior by competition. *Artificial Life*, 1, 353-372.
22. Yao, X. (1993). A review of evolutionary artificial neural networks, *International Journal of Intelligent Systems*, 4, 203-222.