Adaptive Behavior in Competing Co-Evolving Species

Dario Floreano
Center for Neuro-Mimetic Systems
Laboratory of Microcomputing
Swiss Federal Institute of Technology
EPFL, Lausanne, Switzerland
floreano@di.epfl.ch

Stefano Nolfi
Department of Neural Systems
and Artificial Life
National Research Council
Viale Marx 15, Roma, Italy
stefano@kant.irmkant.rm.cnr.it

Abstract

Co-evolution of competitive species provides an interesting testbed to study the role of adaptive behavior because it provides unpredictable and dynamic environments. In this paper we experimentally investigate some arguments for the co-evolution of different adaptive protean behaviors in competing species of predators and preys. Both species are implemented as simulated mobile robots (Kheperas) with infrared proximity sensors, but the predator has an additional vision module whereas the prey has a maximum speed set to twice that of the predator. Different types of variability during life for neurocontrollers with the same architecture and genetic length are compared. It is shown that simple forms of proteanism affect co-evolutionary dynamics and that preys rather exploit noisy controllers to generate random trajectories, whereas predators benefit from directional-change controllers to improve pursuit behavior.

1 Introduction

Adaptive behavior –as compared to innate and fixed behavior– might represent an advantage in unpredictable and dynamic environments. In this respect, co-evolution of competitive species provides an interesting testbed to study the role of adaptive behavior. In the simplest scenario of two competing species, such as a predator and a prey, the behavior of each individual is tightly related to the behavior of the competitor both on the evolutionary and on the ontogenetic time scale. On the evolutionary time scale, the coupled dynamics of the system give rise to the "Red Queen effect" whereby the fitness landscape of each population is continuously modified by the competing population [3]. Given the ubiquity of co-evolution in nature, the relative lack of bio-historical evidence for its role in adaptive progress, and the strong assumptions underlying simple mathematical models developed so far, Artificial Life techniques, such as computer simulations of artificial evolution, are a suitable method to study this phenomenon [1].

On the ontogenetic time-scale, it has been argued that pursuit-evasion contests might favor the emergence of "protean behaviors", that is behaviors which are adaptively unpredictable [4]. For example, preys could take advantage of unpredictable escape behaviors based on short sequences of stochastic motor actions. Similarly, predators could take advantage of enhanced perceptual characteristics and/or adaptive sensory-motor intelligence which could enable predictive tracking strategies. Miller and Cliff provided an excellent review of the biological significance of pursuit-evasion contests and several arguments for its relevance in the study of protean adaptive behavior [11]. Recently, they also described initial results from computer simulations of artificial co-evolution of competing agents controlled by continuous-time recurrent neural networks [2], and developed a set of techniques for analyzing and assessing adaptive progress of both populations [1]. Artificial co-evolution of competitive species has been studied also by other researchers using similar methods, such as Ray's "Tierra" system [13], Sim's creatures [15], and Reynolds' pursuer-evader systems [14]. In very recent work, which will be briefly summarized below, we have investigated the potentiality of the Red Queen effect for evolutionary robotics, and showed that, with a suitable combination of realistic simulations and measuring techniques, competitive co-evolution can develop a variety of efficient behaviors without effort in fitness design [7].

However, none of these experimental researches systematically explored the role of ontogenetic adaptive behavior in co-evolution of competing species. Although most of the evolved systems include some form of noise, it is difficult to say whether this plays an important role on the specific dynamics of co-evolving species or it is simply exploited for smoothing the fitness landscape. In general, all the results presented so far are based on single-run studies and do not include comparisons between different adaptation techniques.

The aim of this paper is that of presenting initial re-

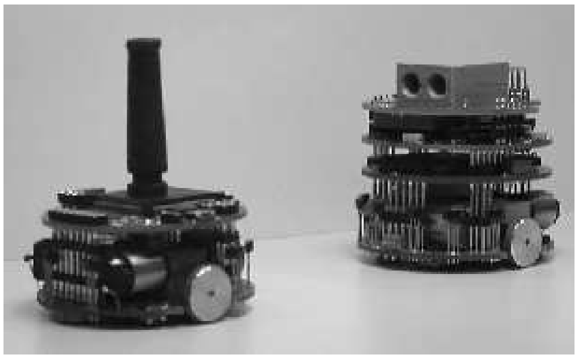

Figure 1: **Right**: The Predator is equipped with the vision module (1D-array of photo-receptors, visual angle of 36°). **Left**: The Prey has a black protuberance which can be detected by the predator everywhere in the environment, but its maximum speed is twice that of the predator. Both Predator and Prey are equipped with 8 Infrared proximity sensors (max detection range was 3 cm in our environment). Evolutionary runs have been conducted in simulation.

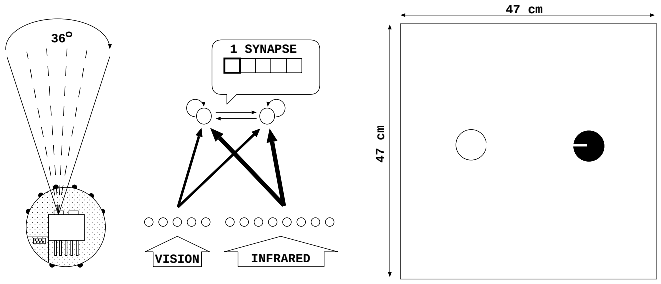

Figure 2: **Left and center**: Details of simulation of vision, of neural network architecture, and of genetic encoding. The prey differs from the predator in that it does not have 5 input units for vision. Each synapse in the network is coded by five bits, the first bit determining the sign of the synapse and the remaining four the other synaptic parameters. **Right**: Initial starting position for Prey (left, empty disk with small opening corresponding to frontal direction) and Predator (right, black disk with line corresponding to frontal direction) in the arena. For each competition, the initial orientation is random.

sults on the effect of adaptive protean behavior in co-evolving competing species. In particular, we want to address the following questions: Does protean behavior affect evolutionary dynamics? Do competing species exploit different types of protean strategies, and how does this affect the competitor's behavior? In the attempt to investigate these issues in very simple settings, we have compared co-evolution of competing species equipped with different types of simple adaptive controllers with results from previous experiments where the controllers were genetically determined [7].

2 Method

As often happens in nature, predators and preys 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 (Figure 1). Both individuals were also provided with eight infrared proximity sensors (six on the front side and two on the back) which had a maximum detection range of 3 cm in our environment. The two species evolved in a square arena of size 47 x 47 cm with high white walls so that the predator could always see the prey (if within the visual angle) as a black spot on a white background.

Since co-evolutionary experiments with real robots would require two separate cables for powering the units for several hours, *ad hoc* solutions must be devised to avoid that the cables twist on each other. Instead of building additional hardware, we have resorted to realistic computer simulations of the two Kheperas. It has been shown elsewhere that for geometrically-simple environments one can have small discrepancies between behaviors in simulation and on the real robot by sampling sensor activity at different distances and angles of the robot from the objects of the world (see [10] for details). We have thus employed this methodology and sampled infrared sensor activity of each robot in front of a wall and in front of another robot. These values were then separately stored away and accessed through a look-up table depending on the faced object.

Simulation of the visual input required different considerations. The vision module K213 of Khepera is an additional turret which can be plugged-in directly on top of the basic platform. It consists of a 1D-array of 64 photo-receptors which provide a linear image composed of 64 pixel of 256 gray-levels each, sub-tending a view-angle of 36°. The optics are designed to bring into focus objects situated at distances between 5cm and 50cm while an additional sensor of light intensity automatically adapts the scanning speed of the chip to keep the image stable and exploit at best the sensitivity of receptors under a large variety of illumination intensities. The K213 vision turret incorporates a private 68HC11 processor which is used for optional low-level processing of the scanned image before passing it to the robot controller. One of these options is position detection of the pixel with minimal activation in the image which in this case corresponds to the position of the prey in the visual field. Therefore, instead of simulating the response of the 1D-array of receptors resorting to complex and time-consuming ray-tracing techniques, we exploited the built-in facility for position detection of pixel with minimal intensity, and divided the visual angle in five sectors corresponding to five simulated photo-receptors (Figure 2, left). If the pixel with minimal intensity was within the first sector, then the first simulated photo-receptor would become active, if the pixel was within the second sector, then the second photo-receptor would become active, etc. We made sure in a set of preliminary measurements that this type of input reduction was largely sufficient to reliably capture and represent all the relevant visual information available to the predator.

Displacement of the robots was computed by passing to the simulator a vector of wheel velocities (positive and negative values standing for rotation in different directions) and calculating the new x, y position using a set of simple trigonometric equations which gave a maximum estimation error of $0.008mm$ for the predator and $0.016mm$ for the prey, absolutely negligible values with respect to the sensor characteristics (for more details, see [7]).

In line with some of our previous work (e.g., [5]), the robot controller was a simple perceptron of two sigmoid units with recurrent connections 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 photo-receptors and from eight infrared proximity sensors (Figure 2, center); 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. This structure, which is well-suited for evolution of Braitenberg-like obstacle avoidance, was chosen for being a minimally sufficient architecture to evolve something interesting while maintaining system complexity at a manageable level; for the same reason, the architecture was kept fixed, and only synaptic parameters were evolved.

In order to keep things as simple as possible and given the small size of the parameter set, we used direct genetic encoding [16]: each parameter (including recurrent connections and threshold values of output units) was encoded on five bits, the first bit determining the sign of the synapse and the remaining four bits different characteristics of the synapses, depending on the controller type under investigation. Therefore, the genotype of the predator was 5 x (30 synapses + 2 thresholds) bits long while that of the prey was 5 x (20 synapses + 2 thresholds) bits long. Two populations of 100 individuals each were co-evolved for 100 generations. Each individual was tested against the best competitors of the ten previous generations (a similar procedure was used in [15, 14, 1]) in order to improve co-evolutionary stability. For each competition, the prey and predator were always positioned on a horizontal line in the middle of the environment at a distance corresponding to half the environment width (Figure 2, right), 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 real robot).

	Bits for one synapse				
Condition	1	2	3	4	5
1	sign	strength			
2	sign	strength		noise	
3	sign	Hebb rule		rate	

Table 1: Genetic encoding of synaptic parameters for each co-evolutionary condition. 1: Genetically-determined controllers; 2: Adaptive-noise controllers; 3: Directional-change controllers.

The fitness function Φ_c for each tournament did not require complex and/or global measures; it was simply *TimetoContact* normalized by the maximum number of motor updates (500) TtC for the predator pr, and $1 - TtC$ for the prey py, further averaged over the number of tournaments (10). Therefore the fitness values were always between 0 and 1, where 0 means worst. Individuals were ranked after fitness performance in descending order and the best 20 were allowed to reproduce. One-point crossover was applied on all randomly paired strings with constant probability $pc = 0.1$, and random mutation (bit switching) was applied to each bit with constant probability $pm = 0.05$.[1].

2.1 Protean controllers

Being the definition of protean behavior only qualitative in the literature, we decided to provide the organisms with "protean controllers", that is neural networks which could potentially display forms of adaptive unpredictability or directional change and compare them with neural networks whose behavior was fixed and genetically evolved. For sake of comparison, all the neural networks had the same architecture, the same genotype length (5 bits per synapse), and used a comparable encoding technique. Here we describe three evolutionary conditions, each one corresponding to a different controller type (Table 1). In all conditions, the first bit of each synapse coded its sign (whether excitatory or inhibitory).

In the first condition, the remaining four bits coded the synaptic strength as a value in the range $[0, 1]$: since no changes take place during the life of the individuals,

[1] In preliminary experiments, we compared various evolutionary runs with different pc, and a less severe selection scheme (50 parents out of 100 individuals), but none of these changes affected evolutionary dynamics and outcomes.

let us call them *genetically-determined controllers*.

In the second condition, only two bits coded the synaptic strength (again, in the range $[0, 1]$), and the remaining two bits coded the level of random noise applied to the synaptic value. Each level corresponded to the lower and upper bounds of a uniform noise distribution: 0.0 (no noise), ± 0.337, ± 0.667, and ± 1.0. At every network activation, each synapse had its own newly-computed noise value added to its strength (with a final check to level out sums below 0.0 or above 1.0). We shall call this condition *adaptive-noise controllers* because each species can evolve the most appropriate noise level for each synapse.

In the third condition, two bits coded four hebbian rules and the remaining two bits the learning rate (0.0, 0.337, 0.667, and 1.0). Four variations of the Hebb rule were used: "pure Hebb" whereby the synaptic strength can only increase when both presynaptic and postsynaptic units are active, "presynaptic" whereby the synapse changes only when the presynaptic unit is active (strengthened when the postsynaptic unit is active, and weakened when the postsynaptic unit is inactive), "postsynaptic" whereby the synapse changes only when the postsynaptic unit is active (strengthened when the presynaptic unit is active, and weakened when the presynaptic unit is inactive), and "covariance" whereby the synapse is strengthened if the difference between pre- and post-synaptic activations is smaller than a threshold (half the activation level, that is 0.5) and is weakened if the difference is larger than such threshold. After decoding a genotype into the corresponding controller, each synapse was randomly initialised to a value in the range $[0, 1]$ and modified at each time step according to the corresponding hebbian rule and learning rate. In a previous paper, we have shown that this evolutionary scheme in a single-agent static environment can develop stable controllers which quickly develop navigation strategies starting from small random synaptic strengths [6]; interested readers will find more details in that paper. Flotzinger has recently replicated those results and studied in more detail the synaptic dynamics, showing that the continuously changing synaptic values reflect (to a certain approximation) input and output values of the network [8]. Therefore, let us call this condition *directional-change controllers*, simply indicating that synaptic changes depend on sensory activation and motor actions.

3 Results

For each condition, 6 different evolutionary runs were performed, each starting with a different seed for initializing the computer random functions. A set of pairwise two-tail t-tests of the average fitness and best fitness along generations among all the six runs, performed to check whether different seeds significantly affected the experimental outcomes, gave negative results at significance level 0.05. Therefore, for each condition below, we shall plot only data referring to seed 1 (arbitrarily chosen), but the statistical tests reported will be based on all the runs. Each run was carried out separately on a Sun SparcStation 20, lasting approximately 8 to 10 hours (time varied depending on the different controllers and on the performance of the competitors). For the first condition, we briefly summarize the basic results already described in a recent paper where we also presented additional analyses and considerations for robotic applications [7].

As we shall see below, the changing fitness landscape due to the "Red Queen effect" demands novel measuring techniques to monitor the dynamics of co-evolutionary systems. For example, a stationary fitness value over several generations could hide to the observer a set of coupled and rapid changes in both populations. Therefore, for all conditions we show the average fitness of the two competing species along with two types of analysis developed by Cliff and Miller to measure progress in co-evolutionary competitive systems [1]. The first, which they dubbed "CIAO data" (Current Individual vs. Ancestral Opponents), shows the performance of the best individuals of each generation against the best competing ancestors; in applying this technique to our populations after co-evolutionary training, we test each individual ten times (that is ten different individual tournaments) against each best competing ancestor, and plot the average fitness as darker squares for higher values. The second, which they dubbed "Ancestral Hamming Maps", shows the normalized Hamming distance between the genotypes of best individuals along generations (darker squares for higher distance). Both these measures reveal some of the underlying dynamics, such as continuous or instantaneous progress, genetic change, cycling through strategies, etc.

3.1 Genetically-determined controllers

When both controllers were genetically determined, a set of oscillations in fitness values emerged after an initial short period (figure 3, left center), as in [15, p. 36]. The onset and amplitude of these oscillations varied across different seeds, but the general pattern was always the same and it kept repeating for several hundred generations, as we could observe in a continuation of this run up to 500 generations (data not shown). However, continuing a run for more than one hundred generations does not reveal anything new with respect to the analyses reported below. We never observed dominance of one species over the other in any of the evolutionary runs, although the preys tended to display higher peaks due to the initial position advantage.

A *relational measure* of performance gives us additional information on the coupled dynamics of such a co-evolved system: for example, one can derive an index of *relative performance* r_i^c by counting how often one

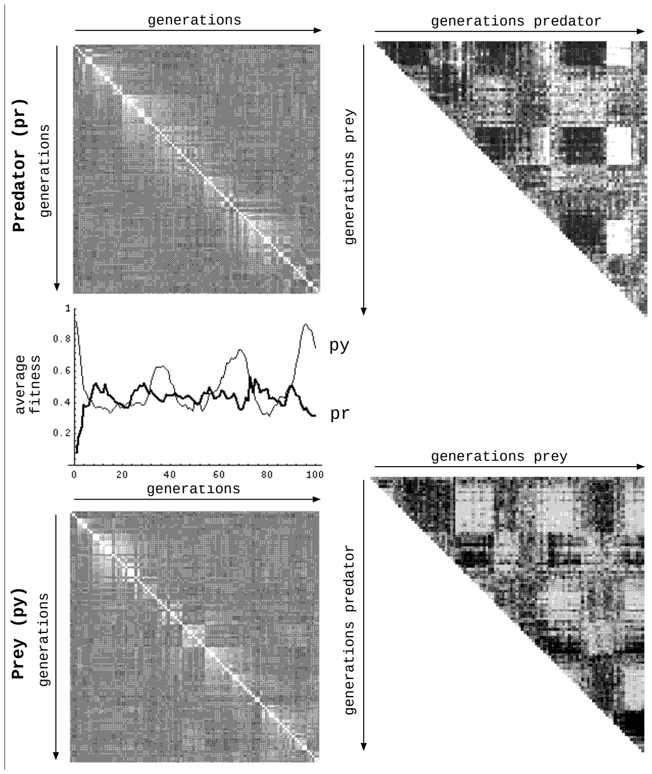

Figure 3: *Genetically-determined controllers.* **Left, center**: Average fitness across generations for predator (pr) and prey (py). **Left, top and bottom**: Ancestral Hamming Maps for predator and prey, respectively: see section 3 for explanation. **Right, top and bottom**: CIAO graphs for predator and prey. respectively: see section 3 for explanation.

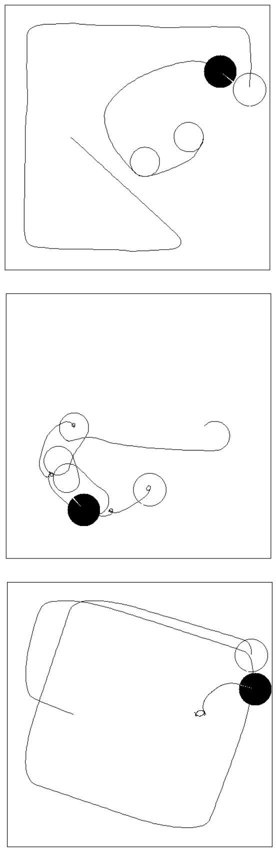

Figure 4: Behaviors of *genetically-determined controllers.* Black disk is predator, white is prey. **Top**. Generation 20: The predator has developed good visual tracking strategies. The prey is a wall-follower with obstacle avoidance. **Center**. Generation 70: The prey turns on place quickly backing whenever the predator attacks. **Bottom**. Generation 90. The predator develops a "spider-strategy" slowly backing against a wall and waiting for the fast-approaching prey.

species reports higher fitness than the competing species at each generation for each separate run i in a specific condition c. In our co-evolutionary runs which lasted 100 generations, such index will be in the range $[-100, 100]$, where -100 means that the preys always outperformed the predators, 0 means that both species were equally better or worse than the competitors, and 100 means that the predators always outperformed the preys. In this condition ($c = 1$), the average value over different runs is $\overline{R^1} = 16.67$ with standard deviation of the sample mean $\sigma = 38$, indicating that both species reported similar performances. The development of a better strategies by one of the species corresponded to a decrement in performance of the competing species.

Major changes in behavioral strategies are reflected by the genotype of the best individuals selected for reproduction. The Ancestral Hamming Maps reported on top and bottom left of figure 3 show how each individual differs from the other individuals of its own population. The white diagonal line (Hamming distance zero) is the identity comparison. Small white areas, which indicate almost identical genotypes, correspond to periods of similar fitness in the performance graph. Since these areas appear in the neighborhood of the matrix diagonal, only best individuals which are close in evolutionary time have similar genotypes. It also means that individuals that report similar fitness, but are distant in time, do not have the same genotype. Small dark lines between white zones indicate an abrupt change of behavioral strategy, whereas a gradual fading of white into gray indicates gradual genetic change (which is quite rare here).

Whether genetic change implies real progress (that is increasingly more complex and efficient pursuit and evasion strategies) or not, is revealed by the CIAO graphs displayed on top and bottom right of figure 3. These graphs show that individuals in later generations do not necessarily score well against competitors of much earlier generations. For example, best predators around generation 90 can hardly catch best preys of generations 0-10, 35-50, and 70-80, despite the fact that the average population fitness is relatively high. This indicates that around generation 90 the predators developed a behavioral strategy tuned to their preceding ten best competitors (during co-evolutionary training, each individual is tested against the best competitors of the preceding ten generations). By carefully comparing the fitness graphs with the CIAO graphs, one can find several indications for the development of behavioral strategies specifically tuned to competitors' behaviors. Another case is shown by the performance of the best preys in the final generations which is quite high in the instantaneous fitness graph, but does not compare so well in the CIAO graphs when the competitors are taken from much earlier generations. Finally, the the Scottish tartan pattern of the CIAO graphs, together with the Ancestral Hamming Maps, indicate periods of relative stasis and fast evolutionary changes in both species (a method for picking out best individuals of each species for efficiency and/or fun purposes has been indicated in a previous paper [7]).

Figure 4 shows some behaviors recorded at interesting points of evolution. At generation 20 we already have challenging tournaments for both species. The prey has developed a good wall-following and obstacle-avoidance behavior, and the predator displays good pursuit strategies. Later, at generation 70, the prey turns in place until it perceives the approaching predator; it then quickly backs (faster than the predator) and starts again rotating. Anytime the prey escapes, the predator performs a half turn on one wheel and re-attacks. The prey is caught when, by chance, the predator attacks it on the side with the motor (where there are no infrared sensors). Finally, at generation 90 the prey has resumed a very fast wall-following strategy. Since high speed combined with a short-range sensor resolution for the prey[2] is such that it cannot avoid an incoming predator, the predator adopts a spider-strategy: it slowly backs towards a wall and there it waits for the fast-approaching prey. However, this predator strategy does not pay off for all the other prey strategies described before, as it can also be seen by the CIAO graph of figure 3.

3.2 Adaptive-noise controllers

The condition with evolutionary adaptive noise ($c = 2$) displayed an average relative performance $\overline{R^2} = 11.66$ with standard deviation of the sample mean $\sigma = 32.5$ which was not statistically different from that of the condition of genetically-determined controllers (probability value was 0.83 for t-test of the difference of the means between the two conditions, i.e. much bigger than significance level 0.05 typically used for rejecting the equality hypothesis). However, in all runs predator and prey performance did not display the high oscillations observed in condition 1 (except for one run with a small counter-phase oscillation in both populations around generation 80); rather, both species displayed similar fitness values, or either the prey or the predator was slightly better than the other, as shown in the center left of figure 5. Furthermore, in all cases the fitness value of the two species required roughly twice the number of initial generations –as compared to condition 1– to reach the intersection point (20 against 10) which marks the establishment of challenging pursuit-evasion tournaments.

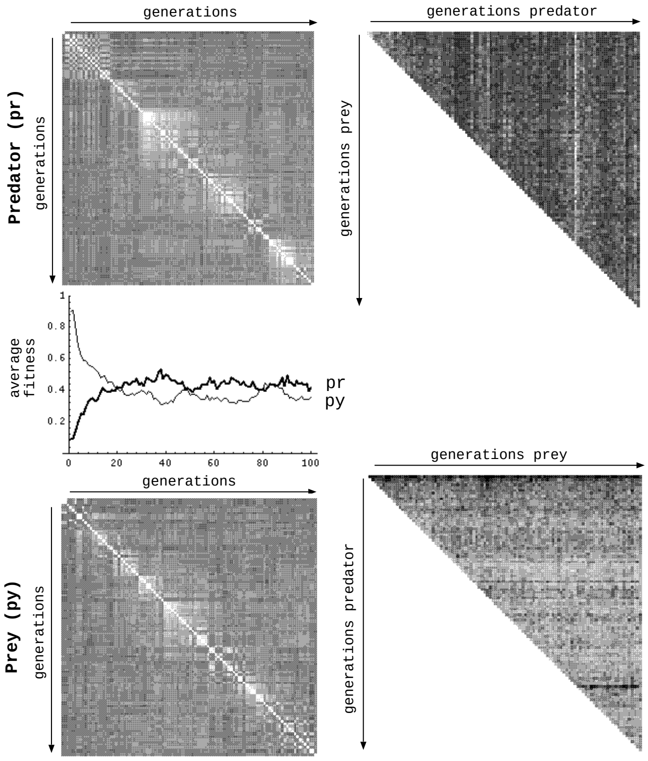

Figure 5: *Adaptive-noise controllers.* **Left, center**: Average fitness across generations for predator (pr) and prey (py). **Left, top and bottom**: Ancestral Hamming Maps for predator and prey, respectively: see section 3 for explanation. **Right, top and bottom**: CIAO graphs for predator and prey, respectively: see section 3 for explanation.

The CIAO graphs revealed a smooth grey texture indicating almost equal performance for all tournaments, except for the initial generations. In some runs the predator was slightly better (darker grey patterns), in others the prey was. In any case, these data showed that both

[2] IR-sensor activity is smaller when the robot faces a small dark object like another Khepera than when it faces a large bright surface like the walls

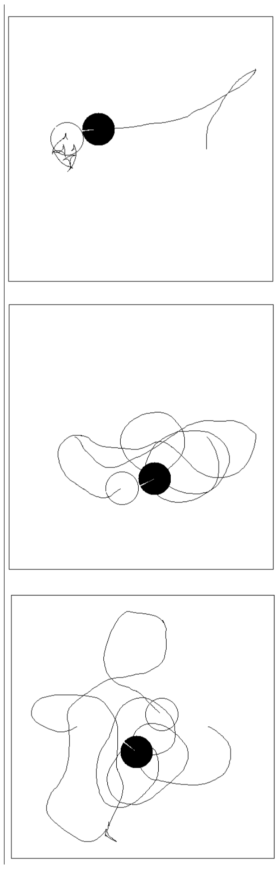

Figure 6: Behaviors of *adaptive-noise controllers*. Black disk is predator, white is prey. **Top**. Generation 20. **Center**. Generation 50. **Bottom**. Generation 80.

species employed behavioral strategies that were comparably challenging for the competitor. However, the fact that small patterns of change in the CIAO graphs were always perpendicular to the predator axis (both for the predator and the prey CIAO data) implied that performance changes were due to a behavioral change in the predator only. This could be also seen from the Ancestral Hamming Maps of figure 5. Here, the slow variation in fitness values during the first 20 generations is accompanied by a corresponding pattern of change in the genotype of the predator only. On the other hand, the genotype of the prey displays roughly the same amount of change along generations.

The hypothesis that the prey exploited noise to develop unpredictable controllers (that is, not improving much on initial random controllers) while the predator tried to develop more stable pursuit strategies was confirmed by the analysis of noise levels in the two species across generations. Similarly to what was done for the the fitness values, we compared at each generation population noise values and noise values of the best individual between the two species. In all runs, the preys reported higher noise values than the predators, except for one run where the noise levels were roughly equal (the same run where the oscillation was observed). Two separate t-tests for checking differences in average noise level and in noise level of the best individuals for all evolutionary runs both displayed a significant difference ($p \ll 0.1$).

A qualitative analysis of behavioral patterns showed unpredictable manoeuvres for the prey and a not-so-smooth target-oriented navigation for the predator (figure 6). There were no detectable changes after generation 20. In all cases, the prey's trajectory was often changing while retaining sufficient obstacle-avoidance abilities (it sometimes stopped near a wall for a few instants and later moved away). The predator's behaviours were more predictable. In general, it was quite good at keeping the prey within the visual field, but its actions were not precise.

3.3 Directional-change controllers

Relative performance of the two species in this condition significantly differed from condition 1 (and also from condition 2), $\overline{R^3} = 72$ with standard deviation of the sample mean $\sigma = 15.39$, $p < 0.01$ for a two-tailed t-test of the difference of the means. In all the six evolutionary runs predators reported higher average and best fitness values than preys, except for short temporary oscillations (figure 7). Furthermore, in all runs, the average fitness of the predator population was more stable than that of the preys.

Although CIAO graphs revealed the presence of behavioral strategies specifically tuned to the behavior of the competitor, such as at generation 90 for the predator in the run plotted here in figure 7, this pattern was less marked than in condition 1. All CIAO graphs for all runs consistently displayed higher performance of the predator (darker grey levels). Ancestral Hamming Maps were not qualitatively different from previous conditions, except for a higher genetic variation between temporally adjacent generations.

More information can be gained by observing behavioral patterns of the two competitors during individual tournaments (figure 8). There is not much variation in the behavior of the predator. It always displays a very good tracking ability across generations: once the prey has been locked in its visual field, it quickly accelerates to maximum speed until contact. As a matter of fact, for the predator it is sufficient to get the sign of the synapses right. Then, independently of their initial random values, the synapses from active sensors will be increased causing an acceleration in the right direction. As compared to condition 1, where the predator tended to efficiently track in only one direction, here it can turn in both directions at equal speed. In condition 1 proper tracking in both directions would have required accu-

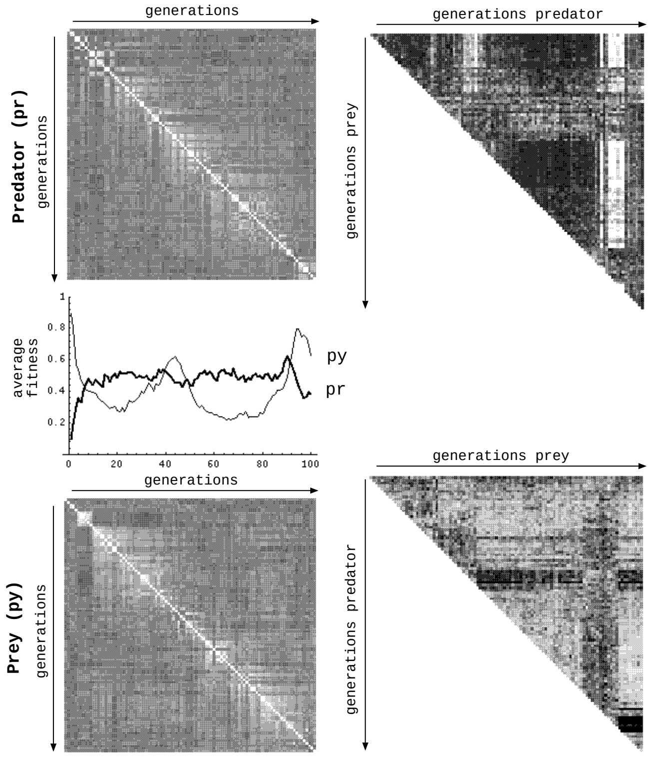

Figure 7: *Directional-change controllers.* **Left, center**: Average fitness across generations for predator (pr) and prey (py). **Left, top and bottom**: Ancestral Hamming Maps for predator and prey, respectively: see section 3 for explanation. **Right, top and bottom**: CIAO graphs for predator and prey. respectively: see section 3 for explanation.

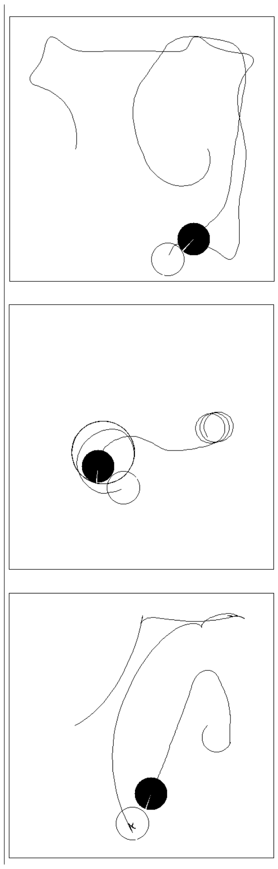

Figure 8: Behaviors of *directional-change controllers.* Black disk is predator, white is prey. **Top**. Generation 20. **Center**. Generation 70. **Bottom**. Generation 95.

rate settings of all synaptic strengths from visual inputs. Here, instead, since synapses are temporarily increased depending on active units [6, 8], individual adjustments of synapses take place when and where required depending on current sensory input. The trajectory in the center image of figure 8 shows another example of synaptic adjustment. Here, while the prey rotates always around the same circle, the predator performs three turns during which synaptic values from the visual units are gradually increased; at the fourth turn, the synaptic values will be sufficiently high to cause a straight pursuit (eventually, the prey will try to avoid the predator without success). Finally, the temporary drop in performance of the predator after generation 90 is due a more precise tracking combined with a slower motion (bottom image of figure 8). Such behavior was probably developed because the preys were also slower and more careful in avoiding obstacles (including the predator).

Although activity-dependent synaptic change are exploited by the far-sighted predator, not the same happens for the prey. Preys are faster than in condition 1 and 2, especially when turning near walls (where IR sensors become active and synapses temporarily strengthen), but they cannot increase their behavioral repertoire with respect to condition 1. Not even can they improve it because volatile changes of the synaptic values imply that most of the time they must re-develop on-the-fly appropriate strengths; although this can be well-suited for avoidance of static obstacles, it does not represent an advantage when facing another moving object such as the predator.

4 Discussion

Introducing protean controllers in co-evolutionary competition significantly affected various aspects of the system dynamics, both on the evolutionary and on the ontogenetic time-scale. On the evolutionary time-scale, noisy controllers ($c = 2$) caused a relaxation of the tightly coupled dynamics observed in the benchmark condition ($c = 1$). High behavioral variation during life of the competitors, especially in the case of preys, was such that only a sufficiently general behavioral strategy could pay off, both for the predator and for the prey. Co-evolutionary search here had higher probability of selecting individuals located in better zones of the fitness landscape, a well-known phenomenon in single-agent evolutionary systems with local variability [9, 12] Instantaneous fitness values reflected more closely behavioral progress, as in traditional single-agent static environments. On the other hand. directional-change controllers ($c = 3$) clearly favored dominance of one species – the predator– whose sensory-motor system profited most of non-random changes of synaptic values.

On the ontogenetic time scale, that is at the level of individual tournaments, the two species differentially exploited the two types of protean controllers. In condition 2, both species reported similar performances, but they differently exploited adaptive noise: preys employed higher noise levels to generate unpredictable and hard-to-track trajectories, whereas predators reduced noise level to maintain sufficient pursuit strategies. In condition 3, predators seemed to benefit from directional synaptic change to improve their pursuit abilities with respect to condition 1.

In order to check whether predators' superior performances in condition 3 were due to a real advantage of the predator rather than to some difficulties of the preys to cope with directional-change controllers, we performed two Master Tournaments. In a Master Tournament each individual is tested against each best competitor of *all* generations (see [7] for more details) and the resulting average fitness (in this case, over 100 tournaments) of each individual across generations is called Master fitness. The graph on the left of figure 9 shows the Master fitness for predators and preys co-evolved in condition 1 (run 1 displayed in figure 3), giving a relative performance $\overline{R} = -12$ (relative performance for the average fitness data of this run was $\overline{r_1^1} = -4$). The graph on the right of figure 9 instead shows the Master fitness for predators evolved in condition 3 (run 1 displayed in figure 7) against preys evolved in condition 1, giving a relative performance $\overline{R} = 42$ (relative performance for the average fitness data of this run was $\overline{r_1^3} = 50$). Had the advantage reported in section 3.3 been caused by underdeveloped preys rather than better predators, the Master Tournament between species evolved in different conditions should have generated opposite results.

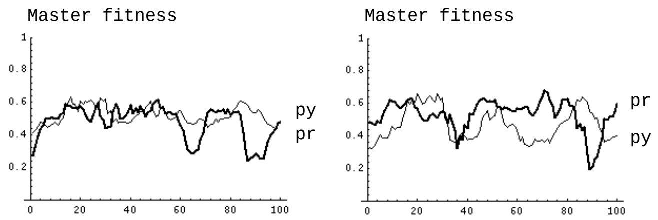

Figure 9: Master Tournament between species evolved in condition 1 (left) and Master Tournament between predator evolved in condition 3 and prey evolved in condition 1 (right).

5 Conclusion

The results reported in this paper provide initial experimental support to the arguments given in section 1 for the differential exploitation of protean behaviors in different co-evolving and competing species [11]. However, one should be careful before generalizing too much these results. Firstly, differential exploitation of sources of controller variability here heavily depended on the sensory-motor systems of our species which were assumed to be different and unmodifiable. Secondly, the controller architecture could have been a limiting factor, being a simple perceptron with discrete-time dynamics: this might be a reason why in these experiments we had no evidence for the development of predictive behavioral strategies during life of each individual. Thirdly, the type of controller variability that we employed in these experiments was quite arbitrary, although specifically chosen to address the issue of random strategies against goal-directed strategies. It would be necessary to carry out further tests with different types of controller adaptation. Additionally, while still keeping the types of controllers employed in condition 2 and 3, it would be interesting to let each species evolve the most suitable controller among the two; for example, one could introduce an extra gene for each individual which specifies how the genotype should be decoded and observe whether there is a significant choice of type 2 for preys and type 3 for predators.

Finally, we think that competing co-evolutionary systems are an interesting testbed to study also other aspects of artificial evolution. One of these is genetic encoding, which here was deliberately kept as simple as possible. However, for co-evolution between different species to be a powerful engine of evolutionary progress, it would be advisable to employ a genetic encoding which takes better care of achieved progress, for example as in the work by Cliff and Miller [1].

Acknowledgments

D.F. thanks Francesco Mondada for providing the Khepera robots with the vision module and Inman Harvey for insightful comments on earlier parts of this work.

References

[1] D. Cliff and G. F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Morán, A. Moreno,

J. J. Merelo, and P. Chacón, editors, *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, pages 200–218. Springer Verlag, Berlin, 1995.

[2] D. Cliff and G. F. Miller. Co-evolution of Pursuit and Evasion II: Simulation Methods and Results. In P. Maes, M. Mataric, J-A. Meyer, J. Pollack, H. Roitblat, and S. Wilson, editors, *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. MIT Press-Bradford Books, Cambridge, MA, 1996.

[3] R. Dawkins. *The Blind Watchmaker*. Longman, Essex, 1986.

[4] P. Driver and N. Humphries. *Protean behavior: The biology of unpredictability*. Oxford University Press, Oxford, 1988.

[5] D. Floreano and F. Mondada. Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics-Part B*, 26:396–407, 1996.

[6] D. Floreano and F. Mondada. Evolution of plastic neurocontrollers for situated agents. In P. Maes, M. Mataric, J-A. Meyer, J. Pollack, H. Roitblat, and S. Wilson, editors, *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. MIT Press-Bradford Books, Cambridge, MA, 1996.

[7] D. Floreano and S. Nolfi. God save the red queen! competition in co-evolutionary robotics. In J. Koza, K. Deb, M. Dorigo, D. Fogel, M. Garzon, H. Iba, and R. L. Riolo, editors, *Proceedings of the 2nd International Conference on Genetic Programming*, Stanford University, 1997.

[8] D. Flotzinger. Evolving plastic neural network controllers for autonomous robots. Msc dissertation 9580131, COGS, University of Sussex at Brighton, 1996.

[9] G. E. Hinton and S. J. Nowlan. How learning can guide evolution. *Complex Systems*, 1:495–502, 1987.

[10] O. Miglino, H. H. Lund, and S. Nolfi. Evolving Mobile Robots in Simulated and Real Environments. *Artificial Life*, 2:417–434, 1996.

[11] G. F. Miller and D. Cliff. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In D. Cliff, P. Husbands, J. Meyer, and S. W. Wilson, editors, *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press-Bradford Books, Cambridge, MA, 1994.

[12] S. Nolfi, J. L. Elman, and D. Parisi. Learning and evolution in neural networks. *Adaptive Behavior*, 3:5–28, 1994.

[13] T. S. Ray. An approach to the synthesis of life. In C.G. Langton, J.D. Farmer, S. Rasmussen, and C. Taylor, editors, *Artificial Life II: Proceedings Volume of Santa Fe Conference*, volume XI. Addison Wesley: series of the Santa Fe Institute Studies in the Sciences of Complexities, Redwood City, CA, 1992.

[14] C. W. Reynolds. Competition, Coevolution and the Game of Tag. In R. Brooks and P. Maes, editors, *Proceedings of the Fourth Workshop on Artificial Life*, pages 59–69, Boston, MA, 1994. MIT Press.

[15] K. Sims. Evolving 3D Morphology and Behavior by Competition. In R. Brooks and P. Maes, editors, *Proceedings of the Fourth Workshop on Artificial Life*, pages 28–39, Boston, MA, 1994. MIT Press.

[16] X. Yao. A review of evolutionary artificial neural networks. *International Journal of Intelligent Systems*, 4:203–222, 1993.