

Co-evolving Task-Dependent Visual Morphologies in Predator-Prey Experiments

Gunnar Buason and Tom Ziemke

Department of Computer Science, University of Skövde
Box 408, 541 28 Skövde, Sweden
{gunnar.buason,tom}@ida.his.se

Abstract. This article presents experiments that integrate competitive co-evolution of neural robot controllers with ‘co-evolution’ of robot morphologies and control systems. More specifically, the experiments investigate the influence of constraints on the evolved behavior of predator-prey robots, especially how task-dependent morphologies emerge as a result of competitive co-evolution. This is achieved by allowing the evolutionary process to evolve, in addition to the neural controllers, the view angle and range of the robot’s camera, and introducing dependencies between different parameters.

1 Introduction

The possibilities of evolving both behavior and structure of autonomous robots has been explored by a number of researchers [5, 7, 10, 15]. The artificial evolutionary approach is based upon the principles of natural evolution and the survival of the fittest. That means, robots are not pre-programmed to perform certain tasks, but instead they are able to ‘evolve’ their behavior. This, to a certain level, decreases human involvement in the design process as the task of designing the behavior of the robot is moved from the distal level of the human designer down to the more proximal level of the robot itself [13, 16]. As a result, the evolved robots are, at least in some cases, able to discover solutions that might not be obvious beforehand to human designers.

A further step in minimizing human involvement is adopting the principles of competitive co-evolution (CCE) from nature, where in many cases two or more species live, adapt and co-evolve together in a delicate balance. The adaptation of this approach in Evolutionary Robotics allows for simpler fitness function and that the evolved behavior of both robot species emerges in incremental stages [13]. The use of this approach has been extended, not only co-evolving the neural control system of two competing robotic species, but also ‘co-evolving’ the neural control system of a robot together with its morphology.

The experiments performed by Cliff and Miller [5, 6] can be mentioned as examples of demonstrations of CCE in evolutionary robotics, both concerning evolution of morphological parameters (such as ‘eye’ positions) and behavioral strategies between two robotic species. More recent experiments are the ones performed by Nolfi and Floreano [7, 8, 9, 12]. In a series of experiments they studied different aspects of CCE of neural robot controllers in a predator-prey scenario. In

one of their experiments [12] Nolfi and Floreano demonstrated that the robots' sensory-motor structure had a large impact on the evolution of behavioral (and learning) strategies, resulting in a more natural 'arms race' between the robotic species. Different authors have further pointed out in [14, 15] that an evolutionary process that allows the integrated evolution of morphology and control might lead to completely different solutions that are to a certain extent less biased by the human designer.

The aim of our overall work has been to further systematically investigate the tradeoffs and interdependencies between morphological parameters and behavioral strategies through a series of predator-prey experiments in which increasingly many aspects are subject to self-organization through CCE [1, 3]. In this article we only present experiments that extend the experiments of Nolfi and Floreano [12] considering two robots, both equipped with cameras, taking inspiration mostly from Cliff and Miller's [6] work on the evolution of "eye" positions. However, the focus will not be on evolving the positions of the sensors on the robot alone but instead on investigating the trade-offs the evolutionary process makes in the robot morphology as a result of different constraints and dependencies, both implicit and explicit. The latter is in line with the research of Lee et al. [10] and Lund et al. [11].

2 Experiments

The experiments described in this paper focus on evolving the weights of the neural network, i.e. the control system, and the view angle of the camera (0 to 360 degrees) as well as its range (5 to 500 mm) of two predator-prey robots. That means, only a limited number of morphological parameters were evolved. The size of the robot was kept constant, assuming a Khepera-like robot, using all the infrared sensors, for the sake of simplicity. In addition constraints and dependencies were introduced, e.g. by letting the view angle constrain the maximum speed, i.e. the larger the view angle, the lower the maximum speed the robot was allowed to accelerate to. This is in contrast to the experiments in [7, 8, 9, 12], where the predator's maximum speed was always set to half the prey's. All experiments were replicated three times.

2.1 Experimental Setup

For finding and testing the appropriate experimental settings a number of pilot experiments were performed [1]. The simulator used in this work is called YAKS [4], which is similar to the one used in [7, 8, 9, 12]. YAKS simulates the popular Khepera robot in a virtual environment defined by the experimenter (cf. Fig. 1). The simulation of the sensors is based on pre-recorded measurements of a real Khepera robot's infrared sensors and motor commands at different angles and distances [1].

The experimental framework that was implemented in the YAKS simulator was in many ways similar to the framework used in [7, 8, 9, 12]. What differed was that in our work we used a real-valued encoding to represent the genotype instead of direct




Fig. 1. Left: Neural network control architecture (adapted from [7]). Center: Environment and starting positions. The thicker circle represents the starting position of the predator while the thinner circle represents the starting position of the prey. The triangles indicate the starting orientation of the robots, which is random for each generation. Right: Khepera robot equipped with eight short-range infrared sensors and a vision module (a camera).

encoding, and the number of generations was extended from 100 to 250 generations to allow us to observe the morphological parameters over longer period of time. Beside that, most of the evolutionary parameters were ‘inherited’ such as the use of *elitism* as a selection method, choosing the 20 best individuals from a population of 100 for reproduction. In addition, a similar fitness function was used. Maximum fitness was one point while minimum fitness was zero points. The fitness was a simple time-to-contact measurement, giving the selection process finer granularity, where the prey achieved the highest fitness by avoiding the predator for as long as possible while the predator received the highest fitness by capturing the prey as soon as possible. The competition ended if the prey survived for 500 time steps or when the predator made contact with the prey before that.

For each generation the individuals were tested for ten epochs. During each epoch, the current individual was tested against one of the best competitors of the ten previous generations. At generation zero, competitors were randomly chosen within the same generation, whereas in the other nine initial generations they were randomly chosen from the pool of available best individuals of previous generations. This is in line with the work of [7, 8, 9, 12]. In addition, the same environment as in [7, 8, 9, 12] was used (cf. Fig. 1).

A simple recurrent neural network architecture was used, similar to the one used in [7, 8, 9, 12] (cf. Fig. 1). The experiments involved both robots using the camera so each control network had eight input neurons for receiving input from the infrared sensors and five input neurons for the camera. The neural network had one sigmoid output neuron for each motor of the robot. The vision module, which was only one-dimensional, was implemented with flexible view range and angle while the number of corresponding input neurons was kept constant.

For each experiment, the weights of the neural network were initially randomized and evolved using a Gaussian distribution with a standard deviation of 2.0. The starting values of angle and range were randomized using a uniform distribution function, and during evolution the values were mutated using Gaussian distribution with a standard deviation of 5.0. The view angle could evolve up to 360 degrees; if the random function generated a value of over 360 degrees then the view angle was

set to 360 degrees. The same was valid for the lower bounds of the view angle and also for the lower and upper bounds of the view range.

Constraints, such as those used in [7, 8, 9, 12], where the maximum speed of the predator was only half the prey's, were adapted here where speed was dependent on the view angle. For this, the view angle was divided into ten intervals covering 36 degrees each¹. The maximum speed of the robot was then reduced by 10% for each interval, e.g. if the view angle was between 0 and 36 degrees there were no constraints on the speed, and if it was a value between 36 and 72 degrees, the maximum speed of the robot was limited to 90% of its original maximum speed.

2.2 Results

The experiments were analyzed using fitness measurements, Master Tournament [7] and collection of CIAO data [5]. A Master Tournament shows the performance of the best individuals of each generation tested against all best competitors from that replication. CIAO data are fitness measurements collected by arranging a tournament where the current individual of each generation competes against all the best competing ancestors [5]. In addition some statistical calculations and behavioral observations were performed. Concerning analysis of the robots' behavior, trajectories from different tournaments will be presented together with qualitative descriptions. Here a summary of the most interesting results will be given (for further details see [1]).

Experiment A: Evolving the Vision Module

This experiment (cf. experiment 9 in [1]) extends Nolfi and Floreano's experiment in [12]. What differs is that here the view angle and range are evolved instead of being constant. In addition, the speed constraints were altered by setting the maximum speed to the same value for both robots, i.e. 1.0, and instead the maximum speed of the predator was constrained by its view angle.

Nolfi and Floreano [12] performed their experiments in order to investigate if more interesting arms races would emerge if the richness of the sensory mechanisms of the prey was increased by giving it a camera. The results showed that "by changing the initial conditions 'arms races' can continue to produce better and better solutions in both populations without falling into cycles" [12]. That is, the prey is able to refine its strategy to escape the predator instead of radically changing it. In our experiments the results varied between replications when considering this aspect, i.e. the prey was not always able to evolve a suitable evasion strategy.

Fig. 2 presents the results of the Master Tournament. The graph presents the average results of ten runs, i.e. each best individual was tested for ten epochs against its opponent. Maximum fitness achievable was 250 points as there were 250 opponents. As Fig. 2 illustrates, both predator and prey make evolutionary progress initially, but in later generations only the prey exhibits steady improvement.

The text on the right in Fig. 2 summarizes the Master Tournament. The two upper columns describe in what generation it is possible to find the predator respectively the

¹ Alternatively, a linear relation between view angle and speed could be used.

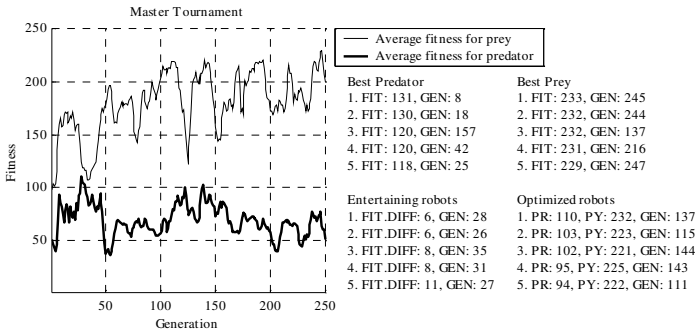


Fig. 2. Master Tournament (cf. Experiment 9 in [1, 2]). The data was smoothed using rolling average over three data points. The same is valid for all following Master Tournament graphs. Observe that the values in the text to the right have not been smoothed, and therefore do not necessarily fit the graph exactly.

best prey with the highest fitness score. The lower left column demonstrates where it is possible to find the most entertaining tournaments, i.e. robots that report similar fitness have a similar chance of winning. The lower right column demonstrates where in the graph the most optimized robots can be found, i.e. generations of robots where both robots have high fitness values.

The left graphs of Fig. 3 display the evolution of view angle and range for the predator and prey, i.e. the evolved values from the best individual from each generation. For the predator the average view range evolved was 344 mm and the average view angle evolved was 111°. It does not seem that the evolutionary process found a balance while evolving the view range as the standard deviation is 105 mm, but the view angle is more balanced with a standard deviation of 48°. The prey evolved an average view range of 247 mm (with a standard deviation of 125 mm) and an average view angle of 200° (with a standard deviation of 86°). These results indicate that the predator prefers a rather narrow view angle with a rather long view range (in the presence of explicit constraints), while the prey evolves a rather wide view angle with a rather short view range (in the absence of explicit constraints) (cf. Fig. 3).

Fig. 3, right graph, presents a histogram over the number of different angle intervals evolved by the predator. The number above each interval represents the maximum speed interval, e.g. in this case most of the predator individuals evolved a view angle between 108 and 144 degrees and therefore the speed were constrained to be within the interval of 0.0 to 0.7. The distribution seems to be rather normalized over the different view angle intervals (between 0 and 252 degrees) (cf. Fig. 3 right).

In other replications of this experiment, the evolutionary process found a different balance between view angle and speed, where a smaller view angle was evolved with high speed. Unlike the distribution in the right graph in Fig. 3 where a large number of predator individuals prefer to evolve a view angle between 108 and 144 degrees, in other replications the distribution was mostly between 0 and 72 degrees, implying




Fig. 3. Left: Morphological description of predator and prey (cf. Experiment 9 in [1, 2]). The graphs present the morphological description of view angle (left y-axis, thin line) and view range (right y-axis, thick line). The values in the upper left corner of the graphs are the mean and standard deviation for the view range over generations, calculated from the best individual from each generation. Corresponding values for the view angle are in the lower left corner. The data was smoothed using rolling average over ten data points. The same is valid for all following morphological description graphs. Right: Histogram over view angle of predator (cf. Experiment 9 in [1, 2]). The graph presents a histogram over view angle, i.e. the number of individuals that preferred a certain view angle. The values above each bin indicate the maximum speed interval.

small, focused view range and high speed. These results, however, depend on the behavior that the prey evolves. If the prey is not successful in evolving its evasion strategy, perhaps crashing into walls, then the predator could evolve a very focused view angle with a high speed. On the other hand, if the prey evolves a successful evasion strategy, moving fast in the environment, then the predator needs a larger view angle in order to be able to follow the prey.

In Fig. 4 a number of trajectories are presented. The first trajectory snapshot is taken from generation 43. This trajectory shows a predator with a view angle of 57° and a view range of 444 mm chasing a prey with a view angle of 136° and a view range of 226 mm. The snapshot is taken after 386 time steps. The prey starts by spinning in place until it notices the predator in its field of vision. Then it starts moving fast in the environment in an elliptical trajectory.

Moving this way the prey (cf. Fig. 4, left part) is able to escape the predator. This is an interesting behavior from the prey as it can only sense the walls with its infrared sensors while the predator needs only to follow the prey in its field of vision in a circular trajectory. However, after few generations the predator loses the ability to follow the prey and never really recovers in later generations. An example of this is the snapshot of a trajectory taken in generation 157 after 458 time steps (Fig 4., right). Here the predator has a 111° view angle and a 437 mm view range while the prey has an 86° view angle and a 251 mm view range. As previously, the prey starts by spinning until it notices the predator in the field of vision. Then it starts moving around in the environment, this time following walls. The predator does not demonstrate any good abilities in capturing the prey. Instead, it spins around in the center of the environment, trying to locate the prey.




Fig. 4. Trajectories from generation 43 (left) (predator: 57° , 444 mm; prey: 136° , 226 mm) and 157 (right) (predator: 111° , 437 mm; prey: 86° , 251 mm), after 386 and 458 time steps respectively (cf. Experiment 9 in [1]). The predator is marked with a thick black circle and the trajectory with a thick black line. The prey is marked with a thin black circle and the trajectory with a thin black line. Starting positions of both robots are marked with small circles. The view field of the predator is marked with two thick black lines. The angle between the lines represents the current view angle and the length of the lines represents the current view range.

Another interesting observation is that the prey mainly demonstrates the behavior described above, i.e. staying in the same place, spinning, until it sees the predator, and then starts its ‘moving around’ strategy.

Experiment B: Adding Constraints

This experiment (cf. experiment 10 in [1]) extends the previous experiment by adding a dependency between the view angle and the speed of the prey. As previously, the predator is implemented with this dependency. View angle and range of both species are then evolved. The result of this experiment was that the predator became the dominant species (cf. Fig. 5), despite the fact that the prey had certain advantages over the predator considering the starting distance and the fitness function being based on time-to-contact.

A Master Tournament (cf. Fig. 5) illustrates that evolutionary progress only occurs during the first generations and then the species come to a balance where minor changes in the strategy result in a valley in the fitness landscape. To investigate if the species cycle between behaviors, CIAO data was collected. Each competition was run ten times and the results were then averaged, i.e. zero in fitness score is the worst and one in fitness score is the best. The ‘Scottish tartan’ patterns in the graphs (Fig. 6) indicate periods of relative stasis interrupted by short and radical changes of behavior [7]. The CIAO data also shows that the predator is the dominating species. Stripes on the vertical axis in the graph for the prey indicate a good predator where the stripe is black respectively a bad predator where the stripe is white. This is more noticeable for the predator than for the prey, i.e. either the predator is overall good or overall bad while the prey is more balanced.

An interesting aspect is the evolution of the morphology (cf. Fig. 7). The predator, as in the previous experiment, evolves a rather small view angle with a rather long range. The prey also evolves a rather small view angle, in fact a smaller view angle than the predator, and a relative short view range with a relatively high standard deviation.

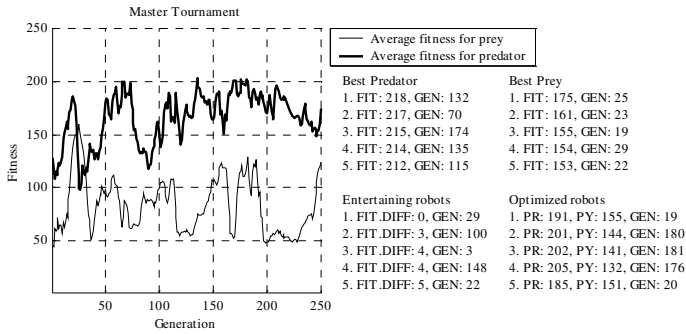


Fig. 5. Master Tournament.




Fig. 6. CIAO data (cf. Experiment 10 in [1, 2]). The colors in the graph represent fitness values of individuals from different tournaments. Higher fitness corresponds to darker colors.

When looking at the relation between view angle and view range in the morphological space then certain clusters can be observed (cf. Fig. 8). The predator descriptions form a cluster in the upper left corner of the area where the view angle is rather focused while the view range is rather long. The interesting part is that the prey also forms clusters with an even smaller view angle, i.e. it ‘chooses’ speed over vision. The clustering of the range varies from small range to very long range, indicating that for the prey the range is not so important.

The evolution of the view angle is further illustrated in Fig. 9. While the predator seems to prefer to evolve a view angle between 36 and 72 degrees, the prey prefers to evolve a view angle between 0 and 36 degrees. This indicates that, in this case the prey prefers speed to vision. The reason behind this lies in the morphology of the robots. The robots have eight infrared sensors, two of them on the rear side of the robots and six of them on the front side. The camera on the robots is placed in a frontal direction, i.e. in the same direction as the six infrared sensors. The robots then mainly use the front infrared sensors for obstacle avoidance. Therefore when the prey evolves a strategy to move fast in the environment because the predator follows it, it

has more use of moving fast than being able to see. Therefore, it more or less ‘ignores’ the camera and evolves the ability to move fast, relying on its infrared sensors.

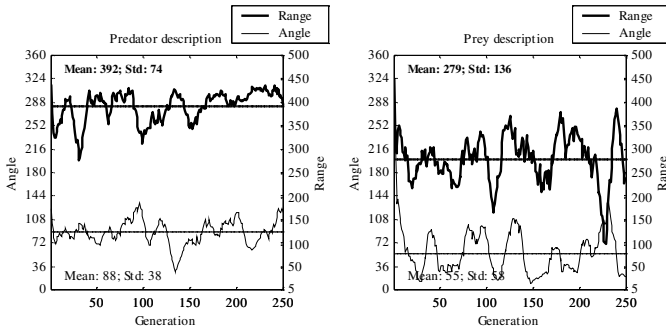


Fig. 7. Morphological descriptions (cf. Experiment 10 in [1, 2]).

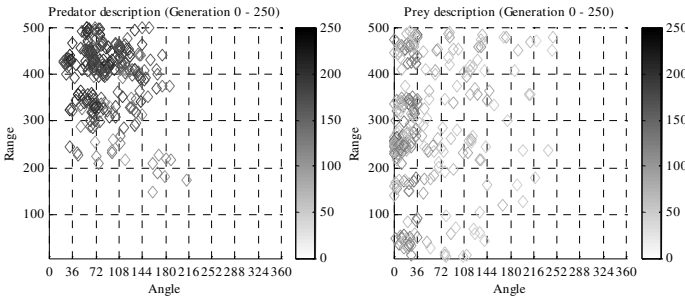


Fig. 8. Morphological space (cf. Experiment 10 in [1, 2]). The graphs present relations between view angle and view range in the morphological space. Each diamond represents an individual from a certain generation. The gray level of the diamond indicates the fitness achieved during a Master Tournament, wit. darker diamonds indicating higher fitness.

A number of trajectories in Fig. 10 display the basic behavior observed during the tournaments. On the left is a trajectory snapshot taken in generation 23 after 377 time steps. The predator has evolved a 99° view angle and a 261 mm view range, while the prey has evolved a 35° view angle and a 484 mm view range. The prey tries to avoid the predator by moving fast in the environment following the walls. The predator tries to chase the prey but the prey is faster than the predator so no capture occurs. In this tournament, the predator also has the strategy of waiting for the prey until it appears in its view field, and then attack (which in this case fails). Although this strategy was successful in a number of tournaments, this strategy was rarely seen in the overall evolutionary process.




Fig. 9. Histogram over view angle of predator and prey (cf. Experiment 10 in [1, 2]).




Fig. 10. Trajectories from generations 23 (predator: 99°, 261 mm; prey: 35°, 484 mm), 134 (predator: 34°, 432 mm; prey: 11°, 331 mm) and 166 (predator: 80°, 412 mm; prey: 79°, 190 mm), after 377, 54 and 64 time steps respectively (cf. Experiment 10 in [1]).

In the middle snapshot (cf. Fig. 10), both predator and prey have evolved narrow view angle (less than 36°), which implies maximum speed. As soon as the predator localizes the prey, it moves straight ahead trying to capture it.

The snapshot on the right demonstrates that for a few generations (snapshot taken in generation 166) the prey tried to change strategy by starting to spin in the same place and as soon as it had seen the predator in its field of vision it started moving around. The prey has a view angle of 80° and a view range of 190 mm. This, however, implies constraints on the speed and therefore the predator soon captures the prey. The strategy was only observed for a few generations.

3 Summary and Conclusions

The experiments described in this article involved evolving camera angle and range of both predator and prey robots. Different constraints were added into the behaviors of both robots, manipulating the maximum speed of the robots.

In experiment A the prey ‘prefers’ a camera with a wide view angle and a short view range. This can be considered as a result of coping with the lack of depth perception, i.e. not being able to know how far away the predator is. In the presence of constraints in experiment B, the prey made a trade-off between speed and vision,

preferring the former. The predator, on the other hand, in both experiments preferred a rather narrow view angle with a relative long view range. Unlike the prey it did not make the same trade-off between speed and vision, i.e. although speed was needed to chase the prey, vision was also needed for that task. Therefore, the predator evolved a balance between view angle and speed.

In sum, this paper has demonstrated the possibilities of allowing the evolutionary process to evolve appropriate morphologies suited for the robots' specific tasks. It has also demonstrated how different constraints can affect both the morphology and the behavior of the robots, and how the evolutionary process was able to make trade-offs, finding appropriate balance. Although these experiments definitely have limitations, e.g. concerning the possibilities of transfer to real robots, and only reflect on certain parts of evolving robot morphology, we still consider this work as a further step towards removing the human designer from the loop, suggesting a mixture of CCE and 'co-evolution' of brain and body.

References

1. Buason, G. (2002a). *Competitive co-evolution of sensory-motor systems*. Masters Dissertation HS-IDA-MD-02-004. Department of Computer Science, University of Skövde, Sweden.
2. Buason, G. (2002b). *Competitive co-evolution of sensory-motor systems - Appendix*. Technical Report HS-IDA-TR-02-004. Department of Computer Science, University of Skövde, Sweden.
3. Buason, G. & Ziemke, T. (in press). *Competitive Co-Evolution of Predator and Prey Sensory-Motor Systems*. In: *Second European Workshop on Evolutionary Robotics*. Springer Verlag, to appear.
4. Carlsson, J. & Ziemke, T. (2001). YAKS - Yet Another Khepera Simulator. In: Rückert, Sitte & Witkowski (eds.), *Autonomous minirobots for research and entertainment - Proceedings of the fifth international Heinz Nixdorf Symposium* (pp. 235–241). Paderborn, Germany: HNI-Verlagsschriftenreihe.
5. Cliff, D. & Miller, G. F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In: F. Moran, A. Moreano, J. J. Merelo, & P. Chacon, (eds.), *Advances in Artificial Life: Proceedings of the third european conference on Artificial Life*. Berlin: Springer-Verlag.
6. 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 & S. W. Wilson (eds.), *From animals to animats IV: Proceedings of the fourth international conference on simulation of adaptive behavior (SAB96)* (pp. 506-515). Cambridge, MA: MIT Press.
7. Floreano, D. & Nolfi, S. (1997a). 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*. San Francisco, CA: Morgan Kaufmann.
8. Floreano, D. & Nolfi, S. (1997b). Adaptive behavior in competing co-evolving species. In P. Husbands, & 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, J.-A. Meyer, & S. W. Wilson (eds.), *From animals to animats V: Proceedings of the fifth international conference on simulation of adaptive behavior*. Cambridge, MA: MIT Press.

10. Lee, W-P, Hallam, J. & Lund, H.H. (1996). A hybrid GP/GA Approach for co-evolving controllers and robot bodies to achieve fitness-specified tasks. In: *Proceedings of IEEE third international conference on evolutionary computation* (pp. 384–389). New York: IEEE Press.
11. Lund, H., Hallam, J. & Lee, W. (1997). Evolving robot morphology. In: IEEE International Conference on Evolutionary Computation (ed.), *Proceedings of IEEE fourth international conference on evolutionary computation* (pp. 197–202). New York: IEEE Press.
12. Nolfi, S. & Floreano, D. (1998). Co-evolving predator and prey robots: Do ‘arms races’ arise in artificial evolution? *Artificial Life*, 4, 311–335.
13. Nolfi, S. & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
14. Nolfi, S. & Floreano, D. (2002). Synthesis of autonomous robots through artificial evolution. *Trends in Cognitive Sciences*, 6, 31–37.
15. Pollack, J. B., Lipson, H., Hornby, G. & Funes, P. (2001). Three generations of automatically designed robots. *Artificial Life*, 7, 215–223.
16. Sharkey, N. E. & Heemskerk, J. N. H. (1997). The neural mind and the robot. In: Browne, A. (ed.), *Neural network perspectives on cognition and adaptive robotics* (pp. 169–194). Institute of Physics Publishing, Bristol, UK.