

# Co-evolving Predator and Prey Robots

**Stefano Nolfi**

Institute of Cognitive Sciences and Technologies  
Consiglio Nazionale delle Ricerche (CNR-ISTC)  
Via S. Martino della Battaglia, 44, 00185, Roma, Italy  
e-mail: stefano.nolfi@istc.cnr.it

## 1. Introduction

Today the study of animal and human behavior can rely also on a new approach --- the adaptive behavior approach --- that study the mechanisms underlying these phenomena and the way in which they originate and change by synthesizing adaptive artificial embodied agents (i.e. robots). Indeed, although in many cases this approach has been undertaken to solve engineering problems by taking inspiration from nature, it is clear that contributions can also be made in the opposite directions.

More precisely we have at our disposal a new family of approaches that vary with respect to (at least) three dimensions: (i) the level of embodiment and situatedness, (ii) the time scale of the phenomena under study, (iii) the level of bio-inspiration.

The first dimension concerns the extent to which the properties of the agents' body (dimension, mass, size, inertia, etc.) and the consequences of the fact that agents are situated in a physical and social world (the fact that agents have egocentric and partial information about the environment, the fact that agents' actions affect agents' sensory experiences, etc.) are taken into account. At one extreme of the first dimension we have weakly embodied models that take into account few selected aspects of the agent/environment interactions, such as for example the probability that an agent chooses the left or the right branch of a path on the basis of the current pheromone distribution along the two paths (Dussutour et al, 2004) without modeling in detail the agents' behavior (e.g. the sensory states experienced by the agents over time, the actual positions and orientations of the agents, the effects of the actions performed on the agents' sensors). On the other extreme of this dimension, we have robotic models or realistic robotic simulations, in which agents are constituted by physical agents that are situated in a physical environment, in which both the characteristics of the agents and of the environment are modeled in detail, and in which agent/environment interactions are subjected to the laws of physics (e.g. the experiments reviewed in this paper).

The second dimension concerns whether one attempts to model and study the evolutionary and/or developmental processes that give rise to a given agent with certain characteristics and skills or only an agent at a certain stage of its evolutionary-developmental process. In the latter case, the goal of the model is only to identify and study the key neural-physiological properties at the basis of the behavior (see for example Horchler et al, 2004). In the former case, instead, the objective also includes the study of the characteristics of the evolutionary-developmental process that can give rise to an agent with certain properties and skills and the way in which such agent is able to adapt to variation of the task/environment (see the experiment reviewed in this paper for an example).

The third and final aspect concerns the level of detail considered in the model. In some works the artificial agent represents a (simplified) model of an organism of a specific species and thus is designed by taking into consideration the specific characteristic of the body, and/or sensory-motor system, and/or nervous system of that specific species (see for example Ijspeert et al., 2007). In other works the artificial agent represents a model of a generic organism that shares general properties with many different natural organisms but does not

correspond in detail to any specific species (for an example see the experiments reviewed in this paper).

In general terms, variations along the three dimensions described above have advantages and drawbacks. For instance, the use of models that are simplified with respect to the embodiment/situated dimension and that do not take into account how agents' skills change phylogenetically and/or ontogenetically prevents analysis of the implications of these two important aspects but, on the other hand, might allow the use of analytic techniques that cannot be used in other type of models. Similarly, strong bio-inspired models enable the possibility to understand the implications of the detailed characteristics of the agents' body and agent/environment interactions but, on the other hand, typically prevent the possibility to study also the role of the evolutionary/developmental process.

The experiments reviewed in this paper study predator/prey behaviors by using robots or realistic robotic simulations that develop their skills through a competitive evolutionary process. This choice of using robots is motivated by the fact that, as we will see, the detailed characteristics of the robots' body and the characteristics and limitation of their sensory systems strongly affect the co-evolutionary dynamics and the behavior of the evolved agents. The choice of modeling the process through which the agent develop their skill by adapting to their physical and social task is motivated by the fact that competitive co-evolutionary dynamics are particularly interesting and by the fact that, as we will see, the behavior of predator and prey agents can hardly be understood without modeling the way in which the behavior of the two species co-adapted over time. The agents involved in the experiments, however, do not attempt to model two specific types of predator and prey organisms but rather two generic species that compete against each other. The agents thus are designed so to share general key properties with natural organisms (having a body, being provided with limited and local sensors, possessing a nervous system constituted by interconnected neural-like units). More detailed characteristics have been chosen arbitrarily and/or by taking into consideration criteria such as simplicity and technological feasibility. This choice has been motivated by an interest in the general aspects that characterized predator-prey co-evolution rather than by an interest in a specific case study and by the difficulty of combining a strong bio-inspired approach with an evolutionary/developmental approach.

The study of competitive co-evolution through the synthesis of evolving agents started almost 30 years ago (Koza, 1991, 1992; Reynolds, 1994; Miller and Cliff, 1994; Cliff and Miller, 1995, 1996). In this article we review the first experiments carried out by using simulated and real robot that were performed by Dario Floreano, Francesco Mondada and myself a few years later (Floreano and Nolfi, 1997; Nolfi and Floreano, 1998; Floreano, Nolfi, and Mondada, 1998) as well as other related works (Buason and Ziemke, 2003; Nelson, Grant, and Henderson, 2004; Buason, Bergfeldt & Ziemke, 2005) and we briefly discuss the implications of this research for behavioral and evolutionary biology and for robotics. For related research on competitive co-evolutionary algorithms see also Rosin and Belew, 1997; Figici, 2004; Popovici, 2006; Bucci, 2007; Ebner, Watson and Alexander, 2010).

## **2. Co-evolving predator and prey robots**

In the experiments reported in (Floreano and Nolfi, 1997; Nolfi and Floreano, 1998; Floreano, Nolfi and Mondada, 1998) two populations of predator and prey robots were evolved for the ability to catch prey and avoid being caught by predators, respectively. Each population was formed by 100 Khepera robots (Mondada, Franzi, and Ienne, 1993) with identical morphological characteristics but different neural controllers (Figure 1).

Both the predator and prey robots were equipped with eight distance sensors (six on one side and two on the other side) able to detect obstacles up to a distance of about 4cm.

However, prey and predator robots differed in three ways. First, the maximum speed of the prey was twice that of the predator. Second, the predator had an additional vision system with a  $36^\circ$  field of view. Third, the prey had a black stick that could be visually perceived by the predator. These differences allowed predators to detect the prey at a distance of up to 100 cm, whereas prey could only infer the presence of nearby predator through its infrared sensors. On the other hand, the prey could outrun the predator.

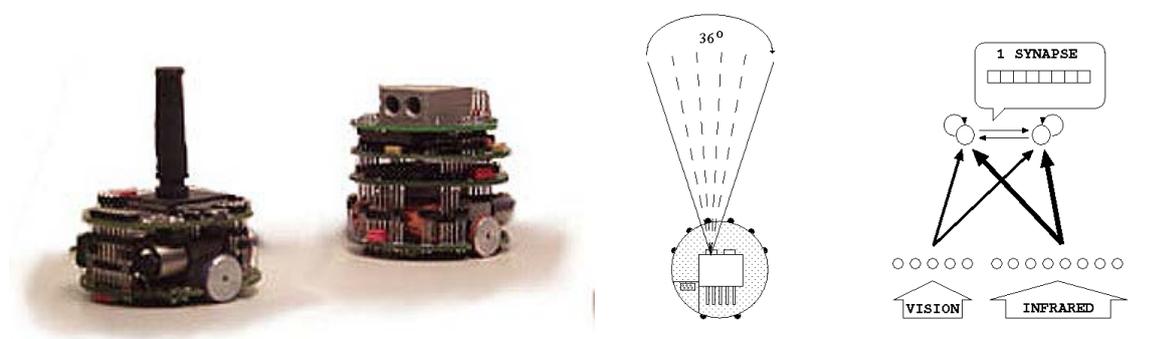


Figure 1. Left. Prey and predators robots (left and right, respectively). Center: positions of the sensors. Right. robots' neural controllers.

Each robot was equipped with a neural controller with 8 sensory neurons that encode the state of the infrared sensors, 5 sensory neurons that encode the current state of the linear camera (in the case of predators only), and 2 motor neurons that encode the desired speed of the two motors controlling the two corresponding wheels. The motor neurons receive connections from the sensory neurons and from themselves (i.e. they have recurrent connections). The architecture of the neural controllers is fixed. The strength of the connection weights, instead, are encoded into the genome of the populations and evolved (Nolfi and Floreano, 2000).

The initial genome consists of two populations of 100 genotypes. Each genotype is formed by a randomly generated string of numbers that encode the connection weights of a neural controller that is embodied in a corresponding robot. Each robot is allowed to interact with the environment and with a competitor for 10 trials lasting 50s each. To improve co-evolutionary stability (see Nolfi and Floreano, 1998), each individual was tested against the best competitors of the ten previous generations. The fitness of the predator and of the prey is computed by calculating the percentage of trials in which they are able to catch or to escape their opponent, respectively (to catch a prey the predator should reach and touch it with its body, to escape a predator the prey should avoid being touched by the predator). After the fitness of all individuals has been calculated (i.e. after all individuals had the opportunity to interact with their opponent for 10 trials), the best 20 genotypes of each population are allowed to reproduce by generating 5 offspring each (i.e. 5 copies with 2% of their genes replaced with randomly selected values). This process of evaluation, selective reproduction, and variation is repeated for 100 generations. For more details, see Nolfi and Floreano (2008).

Ten independent replicates of 100 generations were carried out in physics-based computer simulations and three replicates of 25 generations were conducted with real robots (Floreano, Nolfi & Mondada, 1998). In another series of experiments, the robots were also able to adapt their connection weights while they interacted with the environment, as described below (Floreano and Nolfi, 1997b).

## 2.1 Co-evolutionary dynamic

As hypothesized by evolutionary biologists (Van Valen, 1973; Dawkins and Krebs, 1979) the analysis of the evolving robots indicates that competitive co-evolution produces a never-ending evolution of strategies and counterstrategies in the two competing populations. For example, by visually inspecting the behavior of the best individual of successive generations in a typical experiments, we can see that: (i) after a few generations, the prey developed fast motion in the environment whereas the predators visually tracked them so as to intercept their trajectories, (ii) some generations later, the predators refine their strategy and become very efficient in catching the prey, (iii) the prey then evolve a new strategy that consists of waiting for the predator and moving backward when the predator approaches them, (iv) the predators then change their strategy so to approach the prey from a side in order to exploit the low resolution of the prey sensory system on part of its body, (v) prey then resume on a fast moving strategy that this time is realized by coasting the walls, (vi) the predators then develop a “spider” strategy that consists in backing against one of the walls and waiting for the fast-moving prey whose sensors could not detect the predator sufficiently early to avoid it because its body reflected less infrared light than the white walls, (vii) the prey then display a novel variation of the wait-and-avoid strategy where they quickly rotated in place, which reduced the probability of being approached from the sides without sensors. The analysis ends here since the experiment is terminated after a certain number of generations, but the strategy displayed by the two population would keep changing if the evolutionary process is continued (see Floreano and Nolfi, 1997a, for more details). Overall, these results demonstrate that competitive coevolution can generate a large variety of sophisticated behavioral strategies.

Further experiments performed by using this experimental scenario demonstrated how the ever-changing challenge generated by the competing species enable the evolution of highly effective solutions that would not be discovered otherwise (Nolfi and Floreano, 1998). In particular, we observed that the probability to evolve predators able to catch a high performing prey was higher in the experiment in which predator and prey co-evolved than in control experiments in which the predators evolved against non-varying prey (even if the prey against which their performance is evaluated is the same that they faced during their evolutionary process). See Nolfi and Floreano (1998) for more details.

The fact that the strategies of the co-evolving species keep changing without reaching a stable state, however, does not necessarily imply that the efficacy of the solutions keep increasing throughout evolution. In fact, the co-evolutionary process might enter in a limit cycle dynamic in which the same type of strategies are abandoned and rediscovered over and over again (Nolfi and Floreano, 1998). Suppose for example that at a certain evolutionary stage population A adopts the strategy  $A_1$  that is effective against the strategy  $B_1$  currently adopted by population B (Figure 2). Imagine now that there is a strategy  $B_2$  (similar to  $B_1$ ) that is effective against the strategy  $A_1$ . This will create the adaptive condition for retaining the genetic variations that lead to strategy  $B_2$ . Imagine now that there is a strategy  $A_2$  (similar to  $A_1$ ) that is effective against the strategy  $B_2$ . Population B will sooner or later abandon strategy  $B_1$  and will discover strategy  $B_2$ . Finally imagine that the previously discovered strategy  $B_1$  is effective against strategy  $A_2$ . Population B will come back to strategy  $B_1$ . At this point also population A will come back to strategy  $A_1$  (because, as explained above, it is effective against strategy  $B_1$ ). Overall this implies that, after a certain number of generations, the two populations might re-discover the same strategies they were displaying before and that the evolutionary dynamics might enter into a limit cycle in which the same types of strategies are abandoned and rediscovered over and over again. The analysis of the behavior exhibited by the robot during these evolutionary experiments confirms that this is indeed what happens. Prey tend to rediscover, refine, and then abandon over and over again strategies belonging to the following two families:  $B_1$ : moving fast by avoiding obstacles (that is effective against strategy  $A_2$  but not against strategy  $A_1$ ), and  $B_2$ : wait for the predator and

avoid it with sharp movements when it comes nearby (that is effective against strategy  $A_1$  but not  $A_2$ ). Predators, on the other hand, tend to discover, refine, abandon, and then rediscover over and over again the strategies belonging to the following two families:  $SA_1$ : moving toward the prey by trying to anticipate it (that is effective advantage against strategy  $B_2$  but not  $B_1$ ), and  $A_2$ : stay still by waiting the right moment to move toward the prey by trying to anticipate its trajectory (that provides an advantage against strategy  $B_1$  but not against strategy  $B_2$ ).

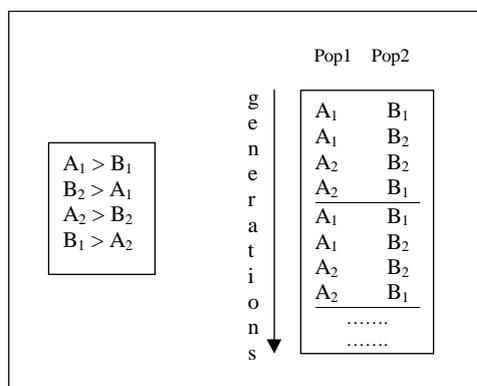


Figure 2. The same strategies ( $A_1$  and  $A_2$  in population A) and ( $B_1$  and  $B_2$  in population B) may be selected over and over again throughout generations as is shown in the right hand side of the figure if the interaction between them looks like what is represented on the left side of the Figure. In this example the repeated cycle includes 4 different combinations of strategies.

## 2.2 Change, innovations, and open-ended evolution

In competitive co-evolution, progresses in one species often creates challenges for the other species and vice versa. The fact that the adaptive task faced by evolving individuals is initially simple, when the competing populations have limited capabilities, and progressively increases in complexity, when the capabilities of the individuals expands, leads to a form of incremental evolutionary process (see also Rosin and Belew 1997) that might facilitate the development of complex skills (thanks to the possibility to re-use previously acquired skills). This potential advantage is indeed confirmed by the experimental data reviewed in the previous section that demonstrate that co-evolved predator robots outperform predators evolved against fixed prey.

In principle, as hypothesized by Dawkins and Krebs (1979), the continuation of this process might lead an ever increasing level of skills/abilities analogous to that observed in an “arms race”. If this hypothesis is true, competitive co-evolution might represent an important drive for change and innovation in evolution (Futuyama & Slatkin, 1983). As we have demonstrated in the previous section, however, the data collected on evolving robots indicates that co-evolutionary phases leading to real innovations do not last forever since the evolutionary dynamic at a certain point enters into a limit cycle in which approximately the same strategies and skills are re-discovered, refined, and abandoned over and over again.

One interesting question, in this respect, is which are the factors that might impact on the length of these innovation phases and that might reduce the tendency to fall into limit cycle dynamics? Also in this case, the data collected by evolving robots, provide useful indications.

By varying the experimental conditions we observed that certain factors might affect the length of the innovation phases and the complexity of the evolved solutions.

One first factor is constituted by the characteristics of the robots’ body and sensory-motor system. Indeed, by running a new set of experiments in which the sensory system of the prey was extended (i.e. in which prey were provided with a camera with a view angle of  $240^\circ$ ) we observed that the evolutionary process leads to much longer innovation phases and to more

complex behaviors (Nolfi and Floreano, 1988). Interestingly, by running a series of experiments with simulated predator and prey agents in which also the characteristics of the agents' sensory system was encoded in the genotype and subjected to variation, Cliff and Miller (1996) observed that predators usually evolve eyes on the front of their bodies (like cheetahs) while prey usually evolved eyes pointing sideways or even backward (like gazelles). Similarly, in a series of experiments performed by using the same experimental scenario described in section 2.1 but in which the view angle and range of the robots' sensory system was subjected to variations, Buason and Ziemke (2003) observed that predators evolved a sensory system with a relatively narrow view angle and long view range while prey evolved a sensory system with wide view angle and shorter view range (see also Buason, Bergfeldt, and Ziemke, 2005). Overall these results indicate that the possibility to subject also the characteristics of the robots' body and sensory-motor system to variations represents a crucial prerequisite for allowing the synthesis of more effective and general solutions.

A second factor is constituted by the ability of the agents to adapt ontogenetically to their task/environment. Indeed, by running a new set of experiments, in which the robots are allowed to vary their connection weights online (on the basis of hebbian learning rules) while they interact with their competitors we observed the emergence of predators able to deal with a larger variety of prey by adapting online to their current competitor (Floreano and Nolfi, 1997b). This is realized, for example, by initially displaying a strategy that consists in "waiting for the right moment to anticipate the prey" which is then transformed into a "move toward prey" strategy when the prey tend to exhibit a "wait for the predator and then escape" behavior. This ability to display multiple strategies and to select the right strategy on the basis of the behaviors of the competitor allows evolving robots to cope with evolutionary variations of the competitor behavior thus reducing the tendency of the competitor population to try to gain an advantage by keep changing strategy. This in turn leads to an incremental evolutionary process in which the new skills are added to the previously developed skills that tend to be preserved. This results indicates that the possibility to adapt at a shorter time rate with respect to genetic evolution might represent a second import pre-requisite for allowing the synthesis of more effective and general solutions.

Finally, another important aspect is constituted by the richness of the task/environment that can be increased, for example, by including in the environment different types of objects (Nolfi and Floreano, 1988; Nelson, Grant, and Henderson, 2004), by studying scenarios with multiple predators or prey (Nelson, Grant, and Henderson, 2004), and by introducing the need to face additional adaptive needs such as for example foraging and saving energy. Indeed, the richness of the task/environment, together with the other two factors mentioned above, represent a crucial prerequisite for establishing a truly open-ended process in which evolution keeps producing changes and innovations by never entering in a stable state or in a limit cycle dynamics.

### **3. Implications for Robotics**

Pursuit and evasion behavior arising from predator prey interactions does not only represent one of the most common and challenging problems for natural organisms but also constitutes a interesting and challenging setup for robotics (Miller and Cliff, 1994). Indeed the need to face highly dynamic, largely unpredictable, and hostile environments requires the development of fast, robust, and reliable solutions. Moreover mastering predator and prey competition, even in the relatively simple experimental scenarios described above, requires the development and display of a wide variety of behavioral and cognitive capabilities such as avoiding fixed and moving obstacles, exploring the environment, exhibiting goal directed navigation, integrating different type of sensory information, integrating sensory-motor

information over time, displaying sequential behaviors, coping with the temporary unavailability of crucial sensory information, arbitrating between different behaviors, anticipating events, adapting online to environment variations, as well as an ability to integrate all these capacities in a single system. For all these reasons predator-prey experimental scenarios represents an ideal test bed for robotics research.

## References

- Bucci A. (1997). Emergent Geometric Organization and Informative Dimensions in Coevolutionary Algorithms. PhD Thesis. Brandeis University.
- Buason G. and Ziemke T. (2003). Competitive co-evolution of predator and prey sensory-motor systems. In: Günther Raidl et al. (Eds.), Applications of Evolutionary Computing – EvoWorkshops. Berlin: Springer Verlag.
- Buason, Bergfeldt & Ziemke (2005). Brains, bodies, and beyond: Competitive co-evolution of robot controllers, morphologies, and environments. Genetic Programming and Evolvable Machines, 6 (1), 25-51.
- 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.
- 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.
- Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. Proceedings of the Royal Society of London B, 205, 489-511.
- Dussutour A, Fourcassie V, Helbing D and Deneubourg JL. (2004). Optimal traffic organization in ants under crowded conditions. Nature. 428 (6978):70-3.
- Ebner M., Watson R. and J Alexander (2010). Co-evolutionary dynamics on a deformable landscape. In C. Di Chio et al. (Eds), Proceedings of the European Conference on the Applications of Evolutionary Computation. Berlin: Springer Verlag.
- Figici S.G. (2004). Solutions Concepts in Coevolutionary Algorithms. Phd Thesis. Brandeis University.
- Futuyama D.J. and Slatkin M. (1983). Coevolution. Sunderland, MA, USA: Sinauer Press.
- Horchler A., Reeve R., Webb B. and Quinn R. (2004). Robot phonotaxis in the wild: a biologically inspired approach to outdoor sound localization. Advanced Robotics (2004), 18(8):801-816
- Ijspeert A.J., Crespi A., Ryczko D., Cabelguen J.M. (2007). From swimming to walking with a salamander robot driven by a spinal cord model. Science 315 (5817): 1416–1420.
- 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.
- Koza, J. R. (1992). Genetic Programming: On the Programming of Computers by Means of Natural Selection, Cambridge, MA: MIT Press.
- Nolfi S., Floreano D. (1998). Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution? Artificial Life, 4(4):311-335.
- Floreano D., Nolfi S. (1997a). God Save the red Queen! Competition in co-evolutionary robotics. In J.R.Koza, K.Deb, M.Dorigo, D.Foegel, B.Garzon, H.Iba & R.L.Riolo (Eds.)

- Genetic Programming 1997: Proceedings of the Second Annual Conference. San Francisco, CA: Morgan Kaufmann.
- Floreano D., Nolfi S. (1997b). Adaptive behavior in competing co-evolving species. In P. Husband & I. Harvey (Eds), Proceedings of the Fourth Conference on Artificial Life, MIT Press, Cambridge, MA.
- Floreano D, Nolfi S, Mondada F (1998) Competitive co-evolutionary robotics: from theory to practice. In: Pfeifer R, ed. Proceedings of the fifth International Conference on Simulation of Adaptive Behavior: From Animals to Animats 5. Cambridge, MA: MIT Press.
- Miller, G.F. and Cliff D. (1994). Protean behaviour in dynamic games: Arguments for the co-evolution of pursuit and evasion tactics. In D. Cliff, P. Husbands, J-A Meyer, and 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.
- Mondada F., Franzi E., and Ienne P. (1993). Mobile Robot miniaturisation: A tool for investigation in control algorithms. In: Proceedings of the third International Symposium on Experimental Robotics, Kyoto, Japan.
- Nolfi S. & Floreano D. (2000). Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. Cambridge, MA: MIT Press/Bradford Books
- Nelson A.L., Grant E., and Henderson T.C. (2004). Evolution of neural controllers for competitive game playing with teams of mobile robots. *Robotics and Autonomous Systems*, 46:135-150.
- Popovici E. (2006). An Analysis of Multi-population Co-evolution. PhD Thesis. George Mason University.
- Reynolds C.W. (1994). Competition, coevolution and the game of tag. In R.A. Brooks and P. Maes (Eds.) Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems. Cambridge, MA: MIT Press.
- Miconi T. (2009). Why coevolution doesn't "work": superiority and progress in coevolution. In L. Vanneschi et al. (Eds). Proceedings of the European Conference on Evolutionary Computation, Lecture Notes in Computer Science, 2009, Volume 5481/2009:49-60. Berlin: Springer Verlag.
- Rosin, C. D. & Belew, R. D. (1997). New methods for competitive coevolution. *Evolutionary Computation*, 5 (1), 1-29.
- Van Valen L (1973). A new evolutionary law. *Evolutionary Theory* 1: 1–30.