## An Artificial Visual Cortex Drives Behavioral Evolution in Co-evolved Predator and Prey Robots

Michael E. Palmer Stanford University 371 Serra Mall Stanford, CA 94305 415-867-3653

mepalmer@stanford.edu

## ABSTRACT

Using a rule-based system for growing artificial neural networks [1], we have evolved controllers for physically simulated 18degree-of-freedom robotic "spiders". These robots previously evolved "galloping" gaits and tracked a compass heading. Here, we introduce an artificial "visual cortex" that permits distinct predator and prey species to "see" one another. It has been claimed [2] that the "Cambrian explosion" of evolutionary diversity was due the emergence of animal eyes, which created a new dynamic in the interaction between predators and prey. The survival of both was for the first time determined by interactions between animals some distance apart. Here, we find that the introduction of vision does indeed generate complex interactions both between and within the two species, as they species evolve complex hunting or avoidance behavior, modulating their galloping gaits with dynamic input from their visual systems.

### **Categories and Subject Descriptors**

D.2.2 [Software Engineering]: Evolutionary Prototyping; I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search – *heuristic methods*; I.2.8 [Artificial Intelligence]: Robotics.

## **General Terms**

Algorithms, Experimentation, Theory.

#### Keywords

Open-ended evolution, complexity, predator-prey, artificial retina, generative, developmental systems.

## **1. INTRODUCTION**

One of the primary goals in the field of artificial developmental systems is evolve systems that show an open-ended increase in complexity over evolutionary time. "Generative" artificial developmental systems [3-6] are intended to provide the possibility of such open-ended increase, whereas, for example, a standard genetic algorithm with a fixed genome length, and fixed phenotypic meanings of all genetic loci, does not. Predator-prey interactions have long been thought to promote adaptive evolution in artificial systems [7, 8]. A "Red Queen's race" between predator and prey may create an arms race of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*GECCO'12 Companion*, July 7-11, 2012, Philadelphia, PA, USA. Copyright 2012 ACM 978-1-4503-1178-6/12/07 ...\$10.00.

Andrew K. Chou Stanford University 371 Serra Mall Stanford, CA 94305

## akchou@stanford.edu

adaptation between them. A series of work by Nolfi et al. [9, 10] and related work by Buason et al. [11, 12] explored co-evolution of robot predators and prey, using a simulated version of a hardware robot. Our work in this paper differs in that we use an open-ended method for "growing" our neural controllers.

## 2. METHODS

# **2.1 Robot Bodies and Actuation by Controllers**

We use a fixed hexapod "spider" robot body, three of which are visible in Figure 1. Each of the six legs has three degrees of freedom (DOF): from the center of the head, looking outward along one of the upper leg segments, the "thigh" segment can move left-right and up-down. The attached "foreleg" segment can move up-down only. Neither joint can twist. Thus, in total, each robot has thirteen rigid body parts, and twelve joints with 18 total DOF, all of which are actuated.



Figure 1: A focal spider sees three other objects in its environment along three of its six lines of sight (arrows).

Each joint axis has fixed limits to its range of motion. The neural network controlling each robot body has one Output neuron corresponding to each degree of freedom; when it assumes a value of +1, it is calling for its corresponding joint axis to be at its maximum range limit; a value of -1 calls for the minimum range limit. A simulated "spring" between the actual and requested positions generates a force on the joint axis.

We used similar robot bodies in [1], but here we have eliminated the body orientation and velocity sensor Inputs, and replaced them with an "artificial visual cortex".

## 2.2 Artificial Visual Cortex

An artificial "visual cortex" allows the spiders to "see" other objects of three types along six lines of sight. For example, in Figure 1, the green (which indicates the prey species) spider at lower right sees three objects situated around it: one spider of the same (green) species, one spider of the other (purple, indicating the predator) species, and one barrier object (cylinder). The lines of sight radiate from the spider at 60 degree angles; objects falling in a ~60 degree arc, centered on each line of sight, will register on the artificial visual cortex.



Figure 2: An "artificial visual cortex" registers the presence of three objects, indicated by the three large neurons.

The 18 Input neurons of the visual cortex are arranged in three rows of six, as shown in Figure 2. A particular neuron activates when an object of its type is in a particular 60-degree arc (centered on one of the six lines of sight): the rows encode the object type, and the columns encode the viewing direction. In Figure 2, three neurons in the visual cortex are activated (larger in size), indicating the presence of one spider of the same species to the front left ("same2"), one spider of the other species to the front right ("other3"), and one barrier object to the left ("barrier1"), as they were situated in Figure 1. The neurons activate more strongly for closer objects.

#### 2.3 Growth of Neural Network

We use the L-Brain method [1] for "growing" neural networks according to inherited sets of growth rules. In the L-Brain method, a neural network unfolds in three dimensions according to cell division rules comprising: 1) a predicate type, 2) a conditional expression that indicates when and where the rule may be applied, and 3) two successor types. Beginning from a single protoneuron of a certain type, the rule set is repeatedly searched for applicable rules. If the predicate of a rule matches the type of a protoneuron, and the conditional expression evaluates to *true*, then the protoneuron divides into two protoneurons, each with one of the successor types. The conditional expressions are intended to control neural development in a space-, time-, and context-dependent way, analogous to natural gene regulation. When a full complement of protoneurons has been produced by a fixed number of cell divisions, the rules are applied again to convert some of the protoneurons into neurons. A fixed set of 18 Input neurons (green, the "visual cortex" neurons in Figure 2) and 18 Output neurons (red neurons in Figure 5, each of which will control one of the 18 DOF of the robot) are introduced Synaptic connections then grow according to the inherited rule set: briefly, the final neurons have a set of "preferred" types to which they would like form connections to, and from. The connections produced satisfy a combination of these preferences with a locality requirement. The L-Brain method itself is not the focus of this paper, but see [1] for much more detail. A video of the unfolding developmental process is available at the following URL: http://www.youtube.com/alifespider

## 2.4 Evolutionary Parameters

#### 2.4.1 Predator and prey interactions

In [1], we successfully used a single species to evolve a neural controller that would direct the 18 DOF of the robot to produce a "galloping" gait, and then track a compass heading to gallop to the North. In this paper, our goal is to study the interaction of two co-evolving species, one predator and one prey, selected for hunting and evasion behavior. The bodies and brains of the two species have identical physical capabilities and growth constraints, but the two species are scored differently. A predator individual receives credit for "eating" a prey individual, by physically colliding with it; the prey is penalized for being eaten, and rewarded for eating inanimate barrier objects.

## 2.4.2 Fitness evaluation in physically simulated local demes

We place N=25 individuals of each species into D separate demes (local populations), where D ranged from 16 to 320; thus the total metapopulation size is ND individuals, ranging from 400 to 8,000, of each species. Each individual has a distinct genotype, i.e., a distinct set of inherited rules. Both species are asexual. All the 2N robot bodies in a single deme are simulated together, along with N barrier objects; thus they may all physically interact. Fitness is relative among all individuals of each species, within one deme. A single evaluation lasts for 2,000 time steps of 1/30 second each, about 1 minute of simulated time. During this time, robots accumulate a score at each time step, according to the details of the physical simulation, including the velocities of the robots, and whether collision events occur. When a prey is "eaten", it receives a score penalty, but does not disappear from the simulation; rather it is "regenerated" (retaining its accumulated score) in a new random location and the simulation proceeds. When a barrier is eaten, it also moves to a new location. Individuals migrate to a new random deme at a rate of 0.01 per generation. Thus one evolutionary generation consists of: 1) fitness evaluation via 2,000 time steps of physical simulation; 2) reproduction according to relative fitnesses; 3) possible mutation of the "rules" making up each genotype (at a rate of 0.05 per rule per generation); and 4) migration among demes. We typically performed runs of 3,000 to 10,000 generations.

## 3. RESULTS

## 3.1 Tracking behavior

With the above scoring function, we had success evolving tracking behavior in both species, using N=25 and D=16. In Figure 3, an example of successful tracking behavior in both species is shown: a predator tracks a prey, which is itself

tracking a barrier object. The spiders leave colored "breadcrumbs" behind them to make their recent path visible; however, the breadcrumbs are not visible to the spiders. A video of tracking is available at: http://www.youtube.com/alifespider



Figure 3: A predator (purple) tracking a prey (green), which is, in turn, tracking a barrier object.

One alternative outcome to the evolution of hunting in both species is that the prey may become faster runners than the predators, such that even a predator that is successfully tracking a prey cannot catch up; this in turn reduces the selective advantage to the predators of good tracking, and they cease to improve it, or may even lose it. (Thus, interestingly, predators that are good at tracking get more "practice", and improve.)



Figure 4: A predator (center, purple) engages in "orbit the barrier" behavior, waiting for prey (green), to approach.

## 3.2 "Orbit the barrier" baiting behavior

One alternative adaptive strategy taken by the predators, if a deme enters the slower predator / faster prey condition, is what we call predators' "orbit the barrier" behavior. In some runs, the predators would circle around a barrier object, apparently waiting for prey to track toward the barrier. When the prey finally approaches and "eats" the barrier, it is not difficult for the predator to move into the center of the barrier (which has just

disappeared, having been "eaten"), and capture the prey. In Figure 4, a predator circles a barrier as a prey approaches. A video of the "orbit the barrier" behavior is available at: http://www.youtube.com/alifespider

### 3.3 Larger metapopulations

With D=16, not every run would produce successful tracking in both species. Using a single node of our 20-node computing cluster (each node contains 2 E5520 4-core CPUs), we are able to conduct a D=16 run on a single cluster node at a rate of about 200 generations/hour, for N=25 (25 individuals of each species, and 25 barriers). In order to run larger metapopulations, we linked the 20 cluster nodes together by passing migrant individuals among them. That allowed us to run a large metapopulation of D=20\*16=320 demes (or ND=8,000 per species) on the cluster, at the same rate of 200 generations/hour.

Runs with larger metapopulations produced additional refinements to behavior. The "circle the barrier" behavior we previously described first arises by predators blindly bumping into a barrier, and having a gait that does not allow them to disengage from it. In larger runs, we commonly see this being refined by predators that can visually track barriers, close on them, and then circle them. This appears to occur when the prey are already accomplished trackers, which makes sense, since a "baiting" predator is relying on the prey's tracking ability.

We also observed prey that shy away from predators: if one of these prey individuals is tracking towards a barrier, and the experimenter manually places a predator in its path, it will detect the predator (the appropriate "other" Inputs are connected, and activate) and divert its course, in order not to collide with the predator. Interestingly, we have also observed prey that will shy away from other prey (and we can see that the appropriate "same" Inputs are activated during this behavior). The adaptive value of this may be that two spiders that collide usually end up with their legs tangled together, making them easy targets.

Interestingly, we have also, rarely, observed predators that track *toward* other predators, so that they collide with them; we are not sure whether this is adaptive or not. We have only observed it in the case of slow predators / fast prey, where the predators are also actively tracking the barriers; so it is possible that some predators benefit by tracking other predators when they are likely to already be circling barriers. We note that it is possible for blind predators to "orbit": when they run into a barrier, their leg motion may be such that they shuffle around the barrier, rather than detaching from it. It would thus be possible for one predator species to promote orbiting wherein a first predator blindly finds a barrier, and a second tracks to the first. This might perhaps be adaptive relative to being blind to barriers, but not relative to tracking the barriers directly. If this behavior is in fact selected, it would be a third-order interaction.

## 3.4 Brain structure for tracking behavior

A typical structure for a brain (here, a prey) that exhibits successful tracking behavior is shown in Figure 5. Only neurons that are "upstream" of some Output are shown in the figure, because only those can affect the gait. Two green Input neurons are so connected, "barrier0" and "barrier5". The typical behavior produced by such brains is to run in circles until a target object (a barrier in this case) appears near the spider. When one of the Inputs detects a target, the spider zigzags back and forth as it closes on the target, with the target alternately activating the two Input neurons as it passes into their line of sight. Each activation causes a "zig" or a "zag" that diverts the path of the spider back toward the target, until it eventually closes on the target.



Figure 5: A (prey) brain that exhibits successful tracking behavior by zig-zagging toward a barrier object.

A common way this evolved tracking algorithm may fail is when two or more target objects are nearby on either side of the spider; this can cause tracking anomalies, such that capture fails. In addition, with moving targets, the target may move across one of the lines of site, and outside the "tracking cone", also causing failure to capture. Commonly, a spider is under time pressure to quickly capture a target that it has sighted, lest a competing spider get to it first. Videos of spiders competing in this way are available at: http://www.youtube.com/alifespider

#### 4. CONCLUSIONS

One interesting effort that aims to make evolution more efficient, is called "novelty search" [13], which keeps track of the regions of phenotypic space that have been previously searched, and does not produce similar organisms again. When phenotypic space becomes very large and multidimensional, however, it may not be possible to efficiently characterize and record the previously searched volume of phenotype space; nor is it clear how to reduce the dimensionality of this representation for general problems. It may be the case that a long-term means to avoid evolutionary stagnation, and evolutionary cycling, is to make the dynamics so complex and multifaceted that the chance of exact cyclic behavior becomes remote. If we wish to design artifacts for which we cannot intelligently create stepping-stones, then at least we can increase the chance that evolution is not repeatedly searching the same domain.

We note that predator-prey interactions do force a sort of temporally-local "novelty search" due to the frequency-dependent fitness of predator-prey interactions: when the predator adopts strategy A, and the prey adopts strategy B to counter it, then *at least for a short time*, the predator is forced to find a novel, non-A solution. No manual dimensionality-reduction of the phenotype space is required: the discouraged strategy (A) is encoded, in a sense, in the genome of the prey. When the predator changes to another strategy, and the prey follows, then this "memory" of the previously covered region of phenotype space is lost – or is it? It is possible for second-order selective effects related to evolvability to shape the genome: even through the prey is no longer currently expressing the B strategy, its genome may now be more easily

able to re-evolve the B strategy. This produces a "memory" on a longer timescale: if the predator adopts A again, it may be more quickly countered with B. Predators that evolve a novel, non-A strategy, can thus be rewarded on this longer timescale. Only when A has been avoided – and novelty has been enforced – for a relatively long time may this "memory" eventually fade.

This Extended Abstract is a summary of a longer paper submitted to the ALife XIII conference, 2012.

#### 5. ACKNOWLEDGMENTS

Many thanks to Marc Feldman and the members of the Feldman lab at Stanford University. This work was supported in part by NIH grant GM28016.

#### 6. REFERENCES

- Palmer, M. E. Evolved Neurogenesis and Synaptogenesis for Robotic Control: The L-brain Model. in *Genetic and Evolutionary Computation Conference (GECCO 2011)* (Dublin, Ireland, July 12-16, 2011).
- [2] Parker, A. *In the blink of an eye*. Perseus Pub., Cambridge, MA, 2003.
- [3] Boers, E. J. W. and Kuiper, H. Biological metaphors and the design of modular artificial neural networks. Master's Thesis, Leiden University, the Netherlands, 1992.
- [4] Gruau, F. Neural Network Synthesis Using Cellular Encoding and the Genetic Algorithm. Ph.D. thesis, l'Ecole Normal Supérieure de Lyon, 1994.
- [5] Hornby, G. S. and Pollack, J. B. Body-Brain Co-evolution Using L-systems as a Generative Encoding. in *Genetic and Evolutionary Computation Conference (GECCO 2001)* (San Francisco, CA, 2001), 868-875.
- [6] Jacob, C. and Rehder, J. Evolution of Neural Net Architectures by a Hierarchical Grammar-based Genetic System. in *Artificial Neural Nets and Genetic Algorithms (ANNGA)* '93 (1993), 72-79.
- [7] Koza, J. R. Evolution and Co-Evolution of Computer Programs to Control Independently-Acting Agents. MIT Press, City, 1991.
- [8] Sims, K. Evolving 3D Morphology and Behavior by Competition. MIT Press, City, 1994.
- [9] Floreano, D. and Nolfi, S. Adaptive behavior in competing coevolving species. in *Proceedings of the Fourth European Conference on Artificial Life* (1997). MIT Press, 378-387.
- [10] Nolfi, S. and Floreano, D. Co-evolving predator and prey robots: Do "arm races" arise in artificial evolution? *Artificial Life*, 4, 4 (1998), 311-335.
- [11] Buason, G., Bergfeldt, N. and Ziemke, T. Brains, Bodies, and Beyond: Competitive Co-Evolution of Robot Controllers, Morphologies and Environments. *Genetic Programming and Evolvable Machines*, 6 (2005), 25-51.
- [12] Buason, G. and Ziemke, T. Competitive Co-evolution of Predator and Prey Sensory-Motor Systems. in *EvoWorkshops* 2003 (2003). Springer-Verlag, 605-615.
- [13] Lehman, J. and Stanley, K. O. Efficiently Evolving Programs through the Search for Novelty. in *Proceedings of the Genetic* and Evolutionary Computation Conference (GECCO 2010). (Portland, OR, 2010).