
category: Artificial Life

Coupling Morphology and Control in an Evolved Robot

Craig Mautner

Computer Science and Engineering Dept.
University of California, San Diego
La Jolla, CA 92109
cmautner@cs.ucsd.edu (619)453-4364

Richard K. Belew

Computer Science and Engineering Dept.
University of California, San Diego
La Jolla, CA 92109
rik@cs.ucsd.edu (619)534-2601

Abstract

The history of natural evolution displays an inseparable coupling between organic bodies and the nervous systems that control them. That is, animal nervous systems extend throughout the entire body and it is problematic to separate them for isolated study. In contrast to this almost all research in Evolutionary Robotics to date begins with a robot body whose features are fixed and proceeds to evolve a control structure for this body. Our research program is focused on exploring the coupled evolution of both the body and the control structure in real robots. In this paper we take early steps toward this goal by exploring the space of sensor and effector selection and positioning coupled with a neural network linking them. This space is explored using a single grammar for generating both the body and neural network. Results from several problem worlds are presented and analyzed.

1 INTRODUCTION

The most successful agents we know of are those found in real life. These agents are well adapted to their environment and can handle many small and large surprises to their world and themselves without failures. Because of this we look to biology for much of our inspiration for this work.

In the past it has not been feasible to explore the space of robot morphology. In contrast to this, the course of natural evolution shows a history of body, nervous system and environment all evolving simultaneously in cooperation with and in response to each other. The research described below investigates the interaction

between co-evolved body and control structures.

Evolutionary Robotics is a very new field dating back only to the early 1990's. The vast majority of work in this field has explored evolving control structures for fixed robotic platforms. Taking this approach is justified when the purpose of the research is to explore learning algorithms solely in which case the body can be considered to be part of the environment/task that is to be learned. This approach is also understandable when the availability of off the shelf robots with embedded controllers is considered.

The methods by which we explore morphology and control is to specify the agent's body and NNet using a grammar. Grammars offer modularity in terms of encompassing detailed structures at various levels of granularity. That is, a grammar can provide a compact representation for complicated and repeated structures. By using grammars we are able to build hierarchical solutions to problems based on the solutions found at lower levels. These grammars are then evolved using common Genetic Algorithm techniques based on the performance of the instantiated agent.

Many aspects of this research have been investigated in isolation by others. Examples of evolved robots that have implemented NNet controllers include six legged walking controllers ((D. Whitley and Pyeatt, 1995), (Kodjabachian and Meyer, 1997)), maze following ((Floreano and Mondada, 1994)), predator-prey behavior ((Floreano and Nolfi, 1997)), and food tracking ((Angeline et al., 1994)). An early experiment applying grammatical models to the construction of feed-forward NNets is due to (Kitano, 1990). Other researchers who have used grammars to develop NNets are (D. Whitley and Pyeatt, 1995) and (Lucas, 1995).

Very little prior work in evolving morphology exists. (Menczer and Belew, 1994) investigated the nature of sensor usage by providing their agents with an evolvable NNet connected to sensors and effectors. (Mark

et al., 1998) explored eye types and positioning. (Sims, 1994) demonstrated a simulation where the complete morphology of the individuals was involved. Sims created an artificial world in which each agent was grown from a genome that defined both the physical structure and the control structure. (Eggenberger, 1997) has developed an evolutionary system that simulates the growth of a body based on differential gene expression. (Lee et al., 1996) have also worked on evolving both control structures and body plans.

Our work most closely resembles (Lee et al., 1996) in its exploration of body and control space. However our work differs in two significant ways. First of all, Lee use independent evolutionary pathways for the body (GA) and the control structure (GP). Our approach uses a common mechanism to explore the body space and control space as a unified whole. Secondly their bodies and control structures are completely specified by their genetic representation. We feel that defining the body/control representation as a grammar provides a developmental model that gives our approach advantages described below.

2 Developmental Model

Our experiments consist of repeatedly evaluating populations of 200 agents in an environment. Each agent is assigned a fitness score based on its performance. Those agents with higher score are preferentially selected via a random process for inclusion in the next population. The agents are generated from a grammar which is represented by a string of characters. Prior to evaluation of the population the character strings encoding the grammars are crossed with one another with a probability of 0.8 and each character is randomly mutated with a probability of 0.001.

The grammars are used to generate agent bodies and an associated Neural Network (NNet) control systems. This process is described in the next sections.

2.1 PRODUCTION RULES

The grammar for an agent consists of a set of rules for rewriting non-terminal cells into terminal and non-terminal cells. Each agent starts as a single undifferentiated non-terminal cell and through repeated application of grammar rules is transformed into a body and NNet consisting only of several, tiled terminal cells.

A terminal cell is a set of directed, weighted edges from the sides of the cell to other sides. A typical terminal cell is shown in Figure 1.

A production rule of the grammar specifies how to re-

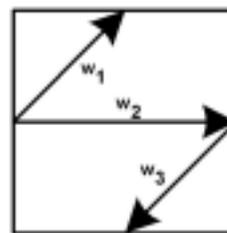


Figure 1: Terminal Cell Example

place a non-terminal cell with one or two terminal and non-terminal cells. The production rule also contains the orientation of the cell (i.e. whether it is to be horizontally or vertically flipped or left in its normal orientation) and it specifies whether the non-terminal produces one or two cells, and if two, the relative position (i.e. above, below, to the left of, or to the right of) of each. In general we indicate non-terminal cells by labeling them with the capital letters A-Z.

2.2 CELL DIVISION

As described in the previous section, each cell is either labeled with a non-terminal symbol or is a terminal cell. Initially each cell begins as a gamete which is labeled with the starting non-terminal symbol of the grammar. Cell differentiation proceeds by selecting and applying the rules of the grammar. For each cell labeled with a non-terminal, a rule is found whose left side matches the non-terminal. The cell is then replaced with the one or two cells specified by the right side of the rule. If there is no matching rule, the cell is replaced with a terminal cell with no weights or edges.

The process continues replacing non-terminals with terminals and non-terminals until there are only terminals left. The rules are applied only a limited number of times to keep rules of the form $A \rightarrow A$ from generating an infinite regress. In our experiments, the maximum number of cell divisions was set to 6. This permits a body to have at most 64 cells (one cell divided in half 6 times produces $2^6 = 64$ cells). A derivation that continues beyond the sixth rule application will replace the cell by a terminal cell with random weights and edges.

A derivation that takes four generations is shown in Figure 2. In this example the gamete is labeled with the starting symbol, A. The production rule $A \rightarrow B|t_1$ indicates that the non-terminal A is converted into two cells. The first cell, a non-terminal B, is to be placed to the left of the second which is the terminal cell, t_1 .

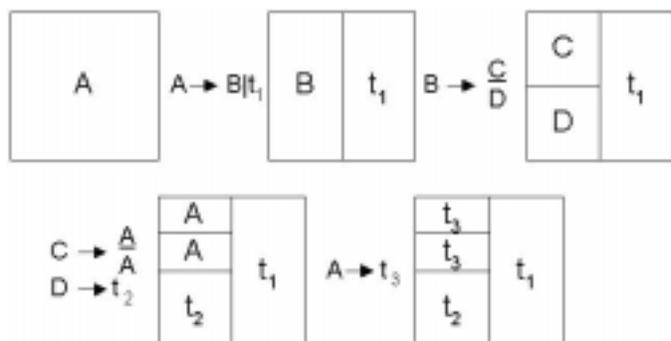


Figure 2: Development Process

2.3 NEURAL NETWORK INTERPRETATION

Once the cell division is complete, the body consists of a set of cells that have within them directed, weighted edges. The cells and edges are interpreted as sensors, effectors and the neural networks that connect them. Edges of a terminal cell are transformed into NNet edges through simple rules. These rules merge edges that point to the same terminal cell sides into a common NNet node that sums and squashes their output activation. When two cells abut one another, the nodes formed by the edges in one cell provide the input activation to the neighboring cells.

Sensor nodes detect signals of the environment. They provide the input that is propagated through the NNet of the body. Their activation is proportional to $\frac{1}{d^2}$ where d is the distance to the signal source. Any directed edge that originates from a cell side on the perimeter of the body becomes a sensor or input node.

Similarly effector nodes are the outputs of the NNet. They provide propulsion to the agent's body. Any directed edge that terminates on a cell side on the perimeter of the body becomes an effector or output node. The force of this propulsion is proportional to their output activation.

The direction of propulsion of the agent is a result of summing all of these forces based on their position on the body. The vector sum of all effector outputs is broken into two forces: The first is a pressure, which acts through the center of the body and translates the body through the environment. The second force is torque which acts perpendicular to a line through the center of the body and causes the body to rotate. The net effect is demonstrated in Figure 3.

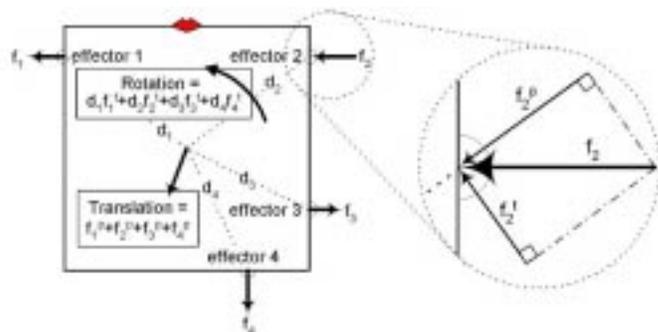


Figure 3: Conversion of Effector Outputs to Rotation and Translation

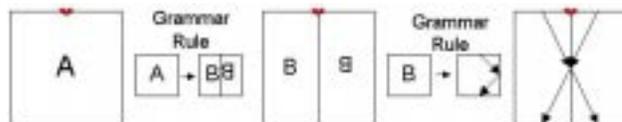


Figure 4: Grammar for Generating Braitenberg Vehicle

3 THE BRAITENBERG VEHICLE

A classic design in the field of robotic control is the Braitenberg Vehicle 2b described in (Braitenberg, 1984). This agent has two sensors on the front and two effectors on the back. The agent's body is bi-laterally symmetric with each sensor connected via a positive weight to the effector on the opposite side. The effect of this connectivity is to steer the agent to the side with the stronger sensor. The effectiveness of this design has been demonstrated in a number of robots.

A simple grammar that generates a complete Braitenberg body is shown in Figure 4. The grammar consists of two rules. The first rule rewrites the undifferentiated cell body (start symbol A) into two nonterminal B cells one of which is horizontally flipped relative to the other. The second rule converts a B nonterminal cells into a terminal cell with two edges defined. Figure 5 shows Braitenberg's Vehicle 2b alongside the one generated by this hand-designed grammar. The parsimonious nature of the grammar that generates the Braitenberg Vehicle under our system shows the representational adequacy of the grammar system.

4 A SIMPLE CENTER-SOURCE ENVIRONMENT

Our first experiments evolved agents that would approach a source of reward placed in the center of the world. The world consists of a 500x500 unit, toroidal space that has at its center a source of both signal and

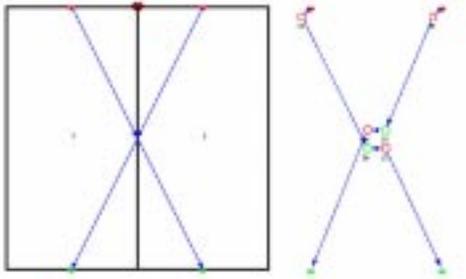


Figure 5: Braitenberg Vehicle with Cell and NNet Representations

reward each of which falls off as the inverse-square of the distance. Agents’ bodies are 20 units on each edge. The body has a mouth placed at its top center.

Each agent is evaluated for 300 time steps. The accumulated reward after this time is used to determine fitness for the purpose of reproduction in the next generation. The agents are placed at random locations in the world and allowed to wander freely (or to just sit in the case of many agents) for up to 30 time steps after which they are moved to a new location. All agents in a given generation are started from the same set of locations. The set of locations are changed at each generation but are never within 28 units of the center of the world.

The fitness, f_i , at time step i is $\frac{1}{d^2}$ for $d \geq 28$ and 1 for $d < 28$ where d is the distance from the agent’s mouth to the center of the world. If an agent’s mouth gets within 28 units of the center of the world then it is given a fitness of 1 for that time step and immediately moved to a new location. If, after 30 time steps, the mouth has not moved within 28 units of the center the agent is automatically moved to another random location. A perfect agent would turn it’s mouth toward the center of the world and approach the center in as few time steps as possible.

4.1 EVOLUTION FROM BRAITENBERG AGENTS

Ten runs of populations seeded with an initial population of Braitenberg agents were performed. Examining the best individuals of the final generation of each run we discovered the following:

- In all runs the original bi-lateral symmetry was retained.
- Four of the ten runs divided horizontally into two cells following the initial vertical division.
- Two of the ten runs lost the crossed pathways

that characterized the original Braitenberg, however none of the remaining crossed pathways resembled the original pathways of the Braitenberg.

Figure 6 shows the body of the best agent in the last generation of the best run. Next to the body is shown the NNet and a trace of its behavior starting from five random positions. Compared with the original Braitenberg ancestors, we see that effectors were added to the front and sides of the body and sensors were added to the rear. The crossed pathways from the front sensors to the opposite side rear effectors were replaced with direct pathways from front to rear on the same sides. The new front effectors are controlled by negative connections from the front and rear sensors on the opposite sides and by a positive weight from the rear sensor on the same side. The side effectors are controlled by a positive weight from the front sensor and a negative weight on the rear sensor of the same sides. Bias weights on the rear effectors provide constant propulsion which are negated by the front effectors as the agent approaches the center of the world.

Behavioral traces of this agent starting from five random positions and orientations are shown in the last figure. (The two long jumps shown are artifacts of the toroidal world). Note the accuracy of the final jumps into the center of the world. The trace beginning midway up on the left side actually falls short on its second jump but then takes a very small step to fall within the center ring.

4.2 EVOLUTION FROM RANDOM GENOMES

We next initialized a population with completely random genomes as opposed to seeding the population with Braitenberg agents above. The results achieved in ten runs are much less consistent than those achieved with the Braitenberg progenitors. Each run produced a solution however these solutions varied wildly in their implementations and quality of solution.

- Only two of the ten runs produced bi-laterally symmetric agents.
- Of the ten runs performed, nine of them produced agents that were inferior to those produced by the previous experiment.
- The tenth run produced an agent that outperformed all of the Braitenberg descendants. This agent is shown in Figure 7.

This successful agent combined negative bias weights on the front with positive bias weights on the rear

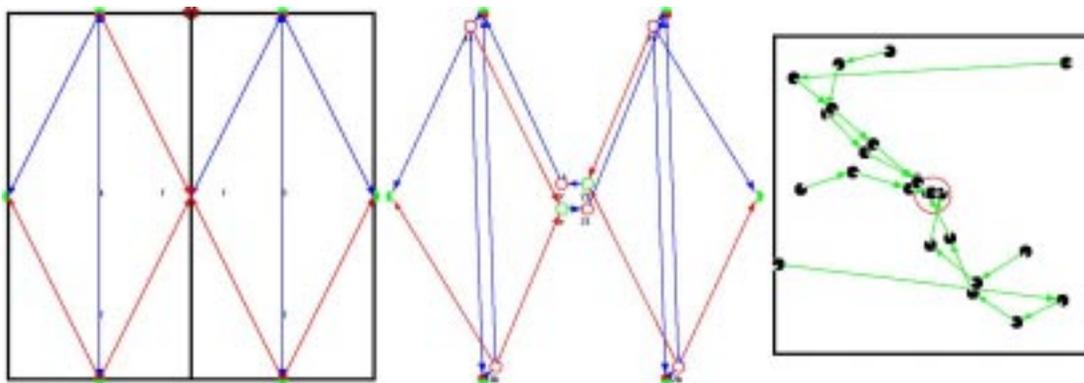


Figure 6: Center-Source Environment: Best Agent Evolved From Braitenberg

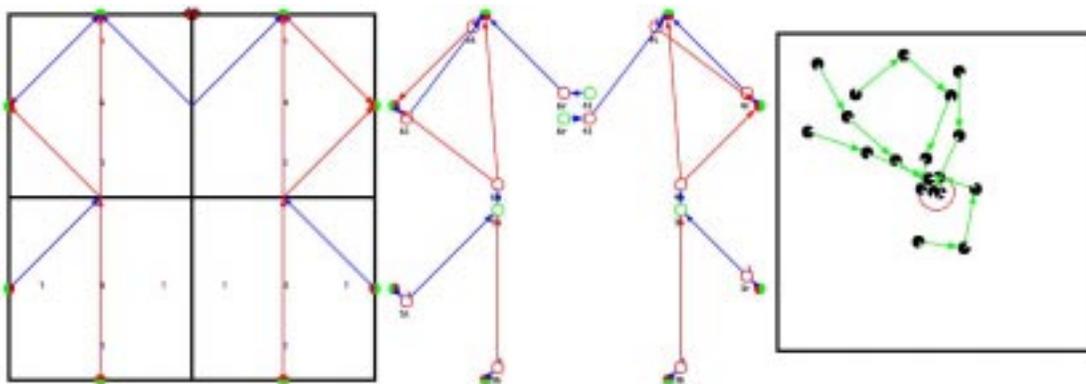


Figure 7: Simple Environment: Best Agent Evolved From Random

effectors to produce a strong jump. In addition the two effectors on the sides towards the front of the agent provided an accurate steering mechanism for pointing the front of the agent at the center of the world.

5 CENTER-SOURCE ENVIRONMENT WITH A VARIATION

The same fitness environment described above was modified slightly to speed up processing. In making this change a small opportunity for receiving high fitness was allowed. This could only happen if an agent stopped at the edge of the reward ring ($d = 28$) with its mouth outside of the ring and then rotated its body without moving it so that its mouth fell within the reward ring. Ten runs of the Braitenberg seeded population were run in this environment. Seven out of the ten runs were able to discover the strategy for receiving high reward. This was surprising to us because we were not aware of the opportunity for receiving this high reward until the agents fitness began to skyrocket.

The best agent of the final generation of one of the

best runs is shown in Figure 8. Note the behavioral trace in this Figure. Note the two paths in the bottom left corner of the behavioral trace. Despite the fact that the second to the last step is nearly twice as far away in one trace than in the other the final jump takes the agent right to the edge of the ring in both cases. This demonstrates the extreme control that has evolved in the agent in order to exploit the opportunity for high reward when landing at the edge. Compare the location of all of the final jumps to the previous environments where the agents ended up very close to the center of the reward ring and far from its edges.

When the population was seeded with random agents rather than Braitenberg only one experimental run was able to discover and exploit the ring effect. The agent that displayed the behavior was similar to the one shown in Figure 8 in that it had positive bias weights at the bottom that serve to push it forward and positive weights from the bottom to the top that inhibit the bias weights when the agent approaches the center of the world. There were also weights that went from sensors on one side to effectors on the other but they were not bi-laterally symmetric like those found above.

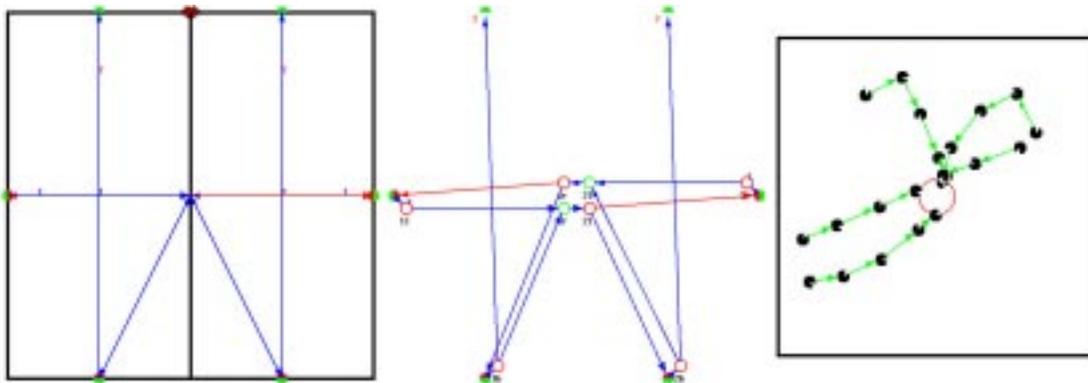


Figure 8: Environment with Edge Bonus: Best Agent that Evolved from Braitenberg

6 DUAL-SOURCE ENVIRONMENTS

We next developed two more sophisticated environments. The purpose of this environment is twofold: first, to learn how our techniques scaled up with increased complexity; and secondly to explore evolution when multiple sensor types are available.

The new world is characterized by two sources of reward and two sensor types. One source of reward produces positive fitness while the other produces negative. Each source generates a fitness that is inversely proportional to the square of the distance. The total fitness is then the sum of the two component fitnesses. If the sum of all the fitnesses awarded over this period is less than 0, then the fitness for the agent is set to 0.

The sources are positioned in the center of opposite quadrants of the world as shown in Figure 9. Similarly to the earlier environment, if the agent moves itself within the rings shown by dashed lines, it receives a fitness of +1.0 for the upper left ring and -1.0 for the lower right ring. The agent is then immediately repositioned to a new starting point. Each agent is evaluated for 300 time steps spending no more than 30 time steps starting from a location before being given a new starting position.

In contrast to the single source world this world is not toroidal. Agents who attempt to wander off the edge are stopped at the boundary.

The two types of sensors available to the grammar are each capable of detecting exactly one of the sources of reward. The level of detection is proportional to the inverse square of the distance from the center of the dashed rings.

Our expectations were that an ideal agent would follow a path to the positive source while avoiding the neg-

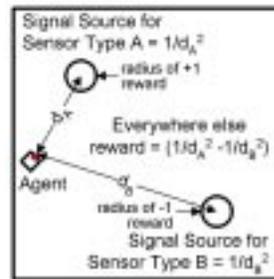


Figure 9: Dual-Source Environment

ative source. For this to occur the ideal agent must make use of both types of sensors. Agents that have only sensors that detect the positive source of reward would regularly achieve positive fitness scores but will occasionally receive negative fitness because they cannot detect the negative source and will occasionally stumble into the negative ring. Agents that detect only the negative source and avoid it will also score higher than an agent with random behavior. Such agents will occasionally accidentally pass through the positive ring although it is just as likely that they will be driven into a wall and receive very little reward for a given trial. As in prior experiments our agents surprised us.

All experimental runs were seeded with random genomes. Of the ten runs in this environment only one produced bi-lateral symmetry (Figure 10) and the bi-laterally symmetric agent again performed the best of all runs. As can be seen from the behavioral trace this agent occasionally wandered into the negative ring. Only one run produced agents with sensors for both the positive and negative sources and that run performed the worst of all runs. All other runs produced agents that used positive reward sensors only.

It was clear from these results that there was not

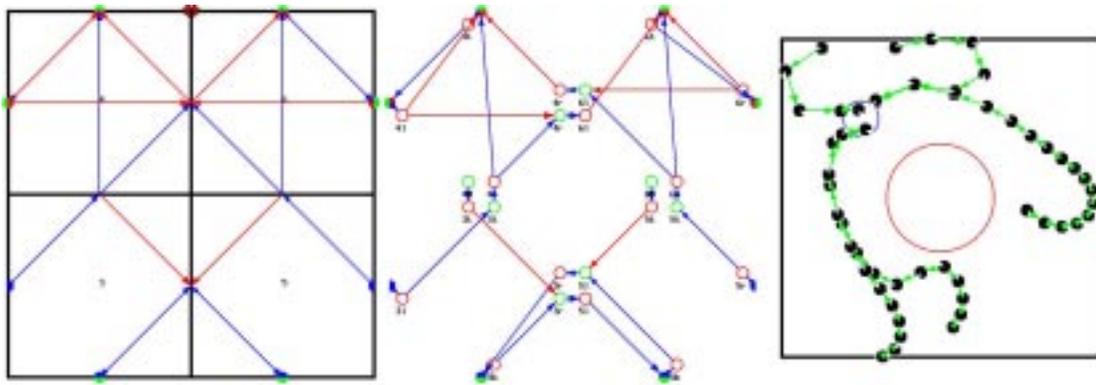


Figure 11: More Challenging Dual-Source Environment: Best Agent

evolving LEGO robots where both the physical body and control structure will be formed by our grammars.

References

- Angeline, P., Saunders, G., and Pollack, J. (1994). An evolutionary algorithm that constructs recurrent networks. *IEEE Trans. on Neural Networks*, 5:54–65.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge.
- D. Whitley, F. G. and Pyeatt, L. (1995). Cellular encoding applied to neurocontrol. In Eshelman, L., editor, *International Conference on Genetic Algorithms*. Morgan Kaufmann.
- Eggenberger, P. (1997). Evolving morphologies of simulated 3d organisms based on differential gene expression. In Husbands, P. and Harvey, I., editors, *Fourth European Conference on Artificial Life*, pages 205–213. MIT Press.
- Floreano, D. and Mondada, F. (1994). Automatic creation of an autonomous agent: Genetic evolution of a neural network driven robot. In D. Cliff, P. Husbands, J.-A. M. and Wilson, S., editors, *From Animals to Animals III*, Cambridge, MA. MIT Press.
- Floreano, D. and Nolfi, S. (1997). Adaptive behavior in competing co-evolving species. In Husbands, P. and Harvey, I., editors, *Fourth European Conference on Artificial Life*. MIT Press.
- Kitano, H. (1990). Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4(4).
- Kodjabachian, J. and Meyer, J. A. (1997). Evolution and development of neural networks controlling locomotion, gradient following, and obstacle avoidance in artificial insects. <http://www.biologie.ens.fr/fr/-animatlab/perso/kodjaba/kodjaba.html>.
- Lee, W., Hallam, J., and Lund, H. (1996). Hybrid gp/ga approach for co-evolving controllers and robot bodies to achieve fitness-specified tasks. In *Proceeding of IEEE 3rd International Conference on Evolutionary Programming*, pages 384–389, New York. IEEE Press.
- Lucas, S. M. (1995). Growing adaptive neural networks with graph grammars. In *Proceedings of European Symposium on Artificial Neural Networks (ESANN '95)*, pages 235–240. <http://esewww.essex.ac.uk/sml/papers.html>.
- Mark, A., Polani, D., and Uthmann, T. (1998). A framework for sensor evolution in a population of braitenberg vehicle-like agents. In Adami, C., Belew, R. K., Kitano, H., and Taylor, C. E., editors, *Artificial Life VI*, pages 428–432, Cambridge, Mass. MIT Press.
- Menczer, F. and Belew, R. K. (1994). Evolving sensors in environments of controlled complexity. In Brooks, R., editor, *Proc. Fourth Conf. on Artificial Life*.
- Sims, K. (1994). Evolving 3d morphology and behavior by competition. In Brooks, R. and Maes, P., editors, *Proceedings of the International Conference Artificial Life IV*, Cambridge MA. MIT Press.