

# Evolution of Gaits of a Legged Robot

Takashi Gomi and Koichi Ide  
Applied AI Systems, Inc.(AAI)  
Kanata, Ontario, Canada K2K 2E4  
E-Mail: gomi@Applied-AI.com; ide@Applied-AI.com

## Abstract

Gaits of a legged robot were generated from scratch using evolutionary robotic techniques. A physical robot was shared by 50 individuals (software invoked control processes that run the entire robot) in each generation. Each program competed for a better score for its gait and ability to lift the body. A genotype that describes leg motions of each robot was defined and used in behavior generation and reproduction of offsprings. An evaluation function that favors early generation of collaborative movements of legs for efficient forward motion of the robot was defined and successfully tested in the experiment.

## 1. Introduction: from learning to evolution

There have been previous attempts to generate the gait of a legged robot using automatic means. Since the creation of the first behavior-based legged robot, Genghis [1][2], legged robots have been programmed by humans until an attempt by P. Maes [3]. In this study, Maes used Genghis robot to learn its own gait from scratch using reinforcement learning. Maes attached "pain" sensors on the bottom surface of Genghis' body. She also attached a small wheel to the tail of the legged robot. The forward turning of the wheel would generate a sense of "pleasure" as it is pulled by the robot which is trying to acquire a forward movement. The artificial animal which learns to walk should try to move forward while avoiding hitting its belly. At the beginning, Maes' robot could only flutter its legs on the floor and was totally unable to stand up. After a few tens of seconds, it managed to partially stand up using a few of its legs, while collapsing and hitting its belly again soon afterwards. Reinforcement learning was applied to emphasize any deviations from the parameter settings defining the present motions whenever the new setting reduced "pain" and increased "pleasure". Small fluctuations in timing of the activation of step motors which drove the robot's six legs both vertically and

horizontally, as well as feedback from the legs as they hit the mostly even but still imperfect real floor, eventually allowed the robot to avoid hitting its belly and walk forward while coordinating the timing at which each leg was energized. One of the interesting achievements of Maes' study was that the bio-robot had actually learned to walk using what is known in animal ethology as "alternating tripod gait." In this gait, two sets of 3 legs {left front, right middle, and left back}, and {right front, left middle, and right back} took turns hitting the floor, and moving forward. The entire learning took place in an amazingly short time: from a totally random start to perfect leg coordination in about 90 seconds.

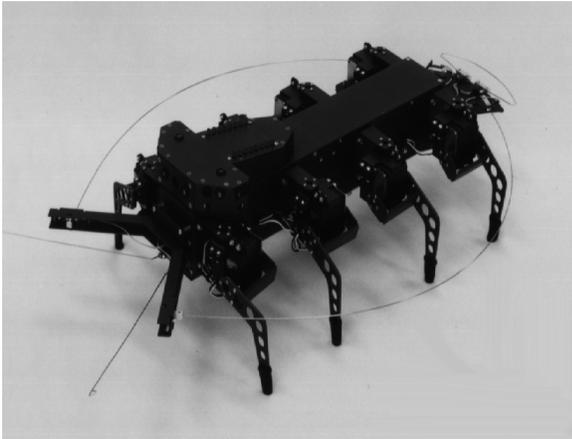
In 1991-92 Maes' experiment was re-implemented by us using Genghis-II robot. In our experiment, instead of using a trailing wheel to measure the progress, currents through the leg motors were used. This was so that more realistic conditions were realized for the experiment by making the morphology closer to the actual insect. In our experiment, the motors on which the current was monitored are those which generate alpha-motions or motions to swing a leg in a horizontal, forward-backward movement. This, combined with "pain" sensors on its belly, similarly implemented as in Maes' experiment, allowed Genghis-II to learn to walk in less than one minute. Again, the alternating tripod gait was observed when the robot began learning, the legs were moving almost in unison, every one of them attempting to stand up. This "egoistic" action on the part of each leg did not allow the robot to even stand up, let alone walk. As the robot learned to collaborate, its overall performance as an artificial animal improved drastically.

In 1995 we developed a revised version of a walking robot which was originally developed as an artificial lobster for a daily educational show at a robot aquarium [4]. The revised robot called OCT-Ib shown in Figure 1.1 was fitted, for the first time, to apply an evolutionary computational technique to create gaits. Its gaits were evolved using GA on genotypes that describe the governing parameters of each leg. Typically with population size of 50 and after 35 generations the artificial lobster began effective leg coordination, and at around 100 generations good gaits

emerged.

## 2. The genotype

The genotype, shown diagrammatically in Figure 2.1, contains eight sets of genomes, each set describing characteristics of the motion of each leg. The characteristics



**Figure 1.1 OCT-Ib Robot**

are described by the amount of delay (up to 5 seconds) after which the leg begins to move, the current status of the leg (two bits indicating the direction of the leg's vertical and horizontal motions), end positions of both vertical and horizontal swings of the leg (-64 to +64, zero being center for both down/up and back/front maxima), and vertical and horizontal angular speed of the leg (-10 to 10, negative being upward and forward speed). The parameter sets are repeated for each robot. Each individual in the population is a process which enacts leg motions using a common software and in accordance with the genotype.

## 3. The fitness function

The evaluation function is set in favor of a robot which stands up, evolves coordination among its legs motions, and has a tendency to move forward. Positive feedback is given to an individual when a front-to-back stroke of a leg occurs while there is sufficient but not excessive vertical load on it. The electric current required to drive the leg is monitored at the corresponding servo motor. Current sensor outputs for both vertical and horizontal axis of the leg are used to identify the "proper loading" condition for each leg. Similarly, if a leg swings back to a starting position with no load, positive feedback is given. Other leg movements result in negative feedback. If any of the touch sensors

located on the belly of the robot is activated (meaning the robot is hitting the belly), a demerit point is added.

The first term of the fitness function represents how a forward motion would develop. When the leg lifts up and swings forward with no load, *lift\_swing* becomes 1. When the leg pushes down and swings backward with a proper load, the term also becomes 1. Otherwise *lift\_swing* is set to -1. The second term describes the energy efficiency and the smoothness of the leg's movement. If the current is over the threshold which is set separately for swing and lift, an over current flag is counted. The sum of over current readings for all motors during the lifespan of the individual is tallied. The third term indicates how well the motions are in general during the lifespan. It is calculated by summing up the possible range of motion for each axis over all legs. The last term is the number of belly hits detected by the two belly sensors.

At the end of each generation, the individual with the highest fitness score is allowed to continue in the following generation. The next five best fit individuals are selected as parents to produce 20 offsprings. A random selection of further parents is made from the rest of the population.

#### 4. The experiment

In upgrading the robot built for a museum exhibit (OCT-I) to a robot for continuous evolutionary runs (OCT-Ib), the engineering revisions were conducted in accordance with the empirical rule suggested by [5]. The OCT-Ib robot has a revised battery arrangement and a new motor type was chosen for durability. The weight of the robot was reduced by some 700 grams from the original 4500 grams. The weight reduction included a reduction in the number of connectors and cables by rearranging components and refining circuits themselves. An entirely new motherboard was developed. With the reduction of the connectors, much of the reliability problem was solved and the robot was capable of enduring evolution runs that often lasted a few days.

While the OCT-Ib robot has a relatively large number of sensors (active infrared sensors with the ability to receive modulated signals, contact sensors or whiskers, light sensors, and motor current sensors), only the motor current sensors and two belly contact sensors were used. In order to allow lengthy evolution runs, an external DC power source was used instead of on-board NiCd batteries. Cabling was arranged so that the robot which is trying to walk has a constant power source. As evolution progressed, this arrangement created a problem as the robot often successfully walked out beyond the reach of the power cable. Eventually, an observer had to be assigned to keep the robot from walking out too far from the field where the experiment was conducted.

$$Fitness = strides \times (1 - overcurrents) \times differences \times (1 - hits) \times 1000$$

where

$$strides = \frac{\sum_{i=1}^{lifespan} lift\_swing_i}{lifespan}$$

when the leg lifts up and swings back, lift\_swing = 1.  
 when the leg lifts down and swings forward, lift\_swing = -1.  
 otherwise lift\_swing = -1.

$$strides = \frac{\sum_{i=1}^{lifespan} \sum_{j=1}^{legs} (lift\_current_{ij} + swing\_current_{ij})}{lifespan}$$

difference =

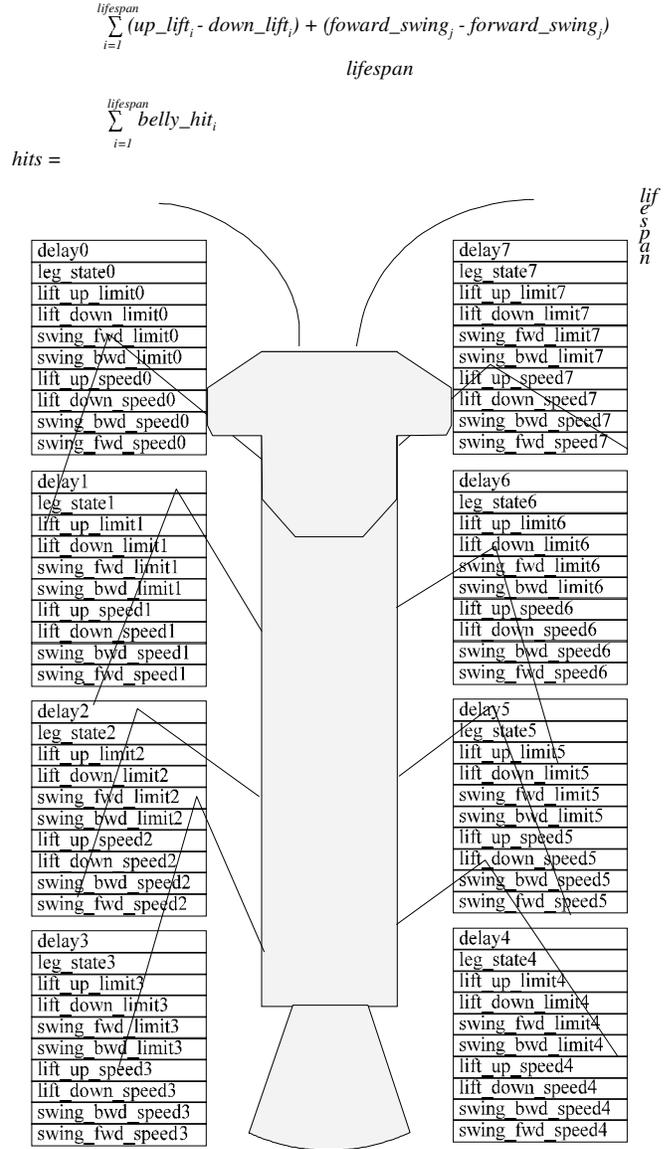


Figure 2.1 Genotype

Figure 3.1 Fitness Function

The manufacturer of the motors guarantees a minimum 100 hours of continuous operation before failure. An individual in this experiment is given 40 seconds of lifespan. Thus the

tethered robot can sustain at least 180 generations of experiments run at population size of 50.

In order to avoid costly loss of time by having to restart an evolution halfway through the experiment, an evolution backup system was implemented. It stores crucial data from each generation (i.e., individuals' genotype and parameters for the generation itself). This allowed us to go back to any point in the history of an evolution run to resume the process from there, often with a modified evaluation scheme.

## 5. Results

### 5.1 Evolution of gaits as cooperation among multiple agents

After generation 10, most individuals, except for a sporadic few, stop hitting the belly, succeed in achieving stand-and-walk status with a faint gait. After another several to a few dozen generations in a typical run, the robot begins to show noticeable forward movements. The results vary from run to run, and some specialization across an entire species often emerges. Some runs had to be "reset," as they obviously went astray, typically leaving a creature not capable of completely standing up, or one walking in circles in the test field. Obviously, in these cases an early convergence took place on an undesirable local peak in the fitness landscape.

Most dominant among the emerged gaits, however, was a mixture between alternating tetrapod and wave gaits. Two sets of legs alternate in the former, while sequential wave-like propagation of leg motions backwards signify the latter. In some of the runs, generation 21 or 22 became a watershed. Steady improvements in walking patterns were observed by then, but some deteriorated afterwards. Another common trend found was the way the robot swings its legs. It typically swings the leg slowly to evade over current detection and at the same time avoiding slipping or loss of friction with the walking surface.

In other runs, generation 35 became a peak, and gaits diverged afterwards. In some of the runs, an odd individual learned to move backward. Another oddity that happened once was a crawler which used crawling to move forward like a sea turtle during an egg-laying expedition on the beach. In such an individual, the last four legs gave the impression of being paralysed, while the four front legs pushed backward in unison to give the body a forward move. In one generation, about 20 per cent of the population was dominated by crawlers.

One generation takes a minimum 33.3 minutes at a life span of 40 seconds per individual, if run continuously. In actual experiments, it typically takes something like one and one-half hours since the run is halted frequently for

observations, recording, and modification of parameters. The longest evolutionary run so far is 110 generations.

### 5.2 Simulation vs physical evolution

The results from the evolution indicate that a wide variety of behaviors may be generated through evolution despite a specific behavior implied in the evaluation function. Physical characteristics of the sensor reception circuits which coordinate sensor signals, a processor that interprets sensor signals, circuits that send out drive currents, and motors that drive the legs all give rise to fluctuations and deviations to the assumed physical conditions. This is contrary to the assumption of steady physical parameters assumed by the designer of simulated evolution in a virtual environment. More impacting on the outcome of the artificial evolution process would be the interactions that must take place between physical entities involved in the process, each of which is driven by or influenced by these fluctuations. For these reasons, the authors believe that it is important to actually carry out embodied evolution as opposed to achieving evolution entirely or mainly through simulation. It is also important to notice the value in conducting evolution experiments without depending on simulation at all even in the early part of the evolution when one attempts to find out if interesting behaviors, such as the ones we saw, would come out at all.

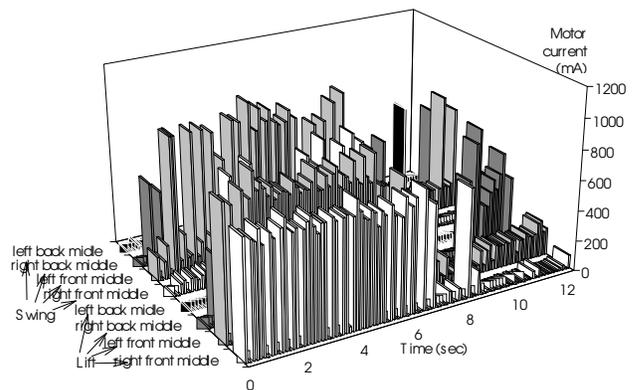
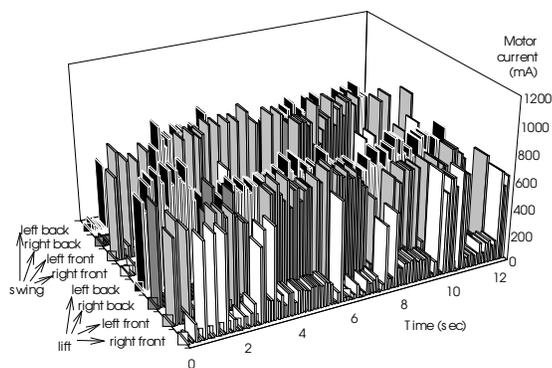


Figure 5.1b Motor current to four middle legs of an individual at generation 0

Increasingly more attempts are being made to evolve robot behaviors using solely simulated process as typified in [6]. While Jacobi succeeded in obtaining satisfactory leg coordination or a gait to move around and avoid obstacles [7], his robot also acquired a set of rough leg motions which would have been eliminated very early in the evolution had

embodied physical evolution been used instead of simulation. Excessive currents through the motors caused by jagged and excessive motions of the legs would have been damaging as happened in replacement of three motors of Jacobi's robot in a short period. In physical evolution, one of the factors taken into consideration when designing an evaluation functions is the energy requirement. In our experiment a term in the evaluation function discouraged the use of energy in general, and penalized rough motions which drain excessive current. In simulated evolution, it is impossible to measure the actual current sunk into each of the motors. It is also not easily calculable as a large number of non-linear elements are involved in the process: current-torque relationships, uneven friction of rubber tipped legs on the floor surface and unpredictable interference from other legs. The last item would be

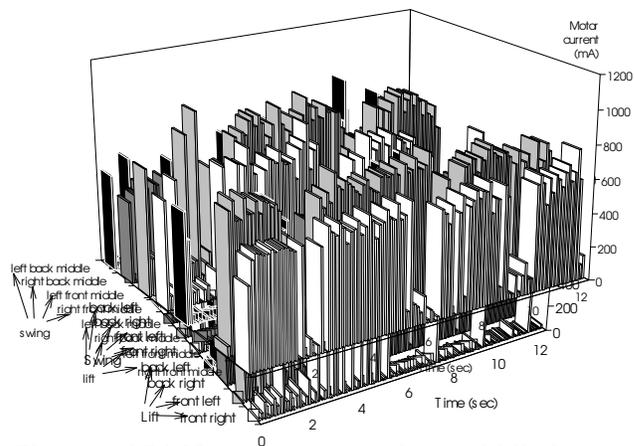


**Figure 5.2(a) Motor current to four front and back legs of an individual at generation 99**

particularly important. The purpose of applying evolutionary computation to robots is partly to study the emergent phenomena associated with the evolution as opposed to simply engineer a set of behaviors on a robot.

Figures 5.1 (a) and (b) show the profile of current drawn on sixteen motors that are responsible for both the horizontal ('swing') and vertical ('lift') motions of the eight legs during an early stage of our evolution experiment (an individual in generation 0). Readings were made through a small serial resistor inserted in each of the 16 actual motor drive circuits. Because the reading of the current value was driven by an instruction embedded in the leg drive command, not all transient phenomena could be captured. Nevertheless, the record depicts most of what happened to motors of the OCT-1b robot in an early stage of evolution. Some of the legs were stuck to the floor while there was an obvious effort to swing them either forward or backward (e.g., 'front left' and 'right back' legs). Others (e.g., 'front right' and 'right back middle' legs) were succeeding in some form of swinging action while having

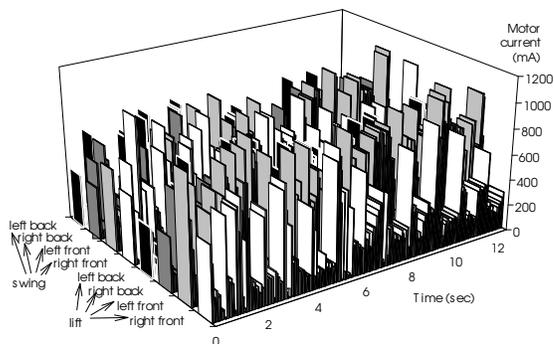
the leg touching the floor (occasional 'swing' current while having positive 'lift' current). It is clear that most of the motions were erratic and arbitrary, and could not be said to have composed a gait at this point.



**Figure 5.2(b) Motor current to four middle legs of an individual at generation 99**  
**Figure 5.1(a) Motor currents to four front and back legs of an individual at generation 0**

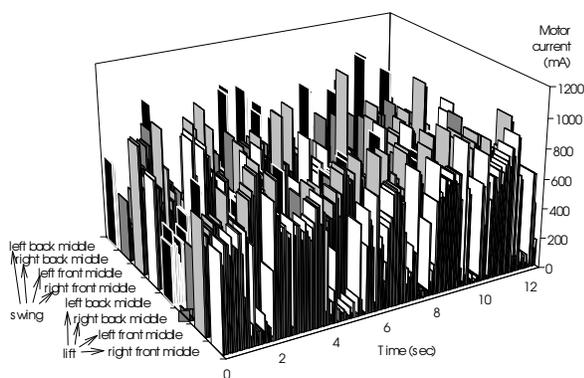
Figures 5.2 (a) and (b) are profiles of motor current for an individual caught in our experiment at generation 99. In general, there is a periodicity in current profiles, indicating some regularity in leg motions. The complementary nature of drive current between the 'right front' and the 'left front' legs, for example, indicate alternating sharing of carrying the weight at the front of the robot. A good synchronization between vertical and horizontal motions is seen in the current for the 'right front middle' and the 'left front middle' legs. The peak current to the motors was well stabilized at about 700 mA with a few exceptions which reach 800 mA.

We requested Nick Jacobi of Sussex University to send us a copy of the program he had evolved for 400 generations using a simulator and then tested on their OCT-1b. Steps to



**Figure 5.3(a) Motor current to four front and back legs of Jacobi's robot at generation 400**

allow reading of the motor current were inserted into his original source code and then compiled and down-loaded to our identical OCT-Ib. When run, it exhibited similar rapid but jagged movements as seen on the Sussex robot. Motor current was measured using inserted steps and the resistors put into the motor drive circuit. The result of a measurement is shown in Figure 5.3 (a) and (b). In contrast to Figure 5.2, Figures 5.3 (a) and (b) show far less regularity in the current pattern despite the fact it was taken after 400 generations of evolution using a simulator. The average high current is about 900 mA. The higher current and the rapid rate of current changes reflect a faster and galloping pace with which the robot walks about and obviously are the cause for the frequent destruction of the motors. The servo motor on OCT-Ib (Futaba S9303) is designed mainly for PID control where position changes are not too frequent, such as controlling the aileron of a model airplane. Although the exact specification of the motor is not given, it is empirically known that rapid and prolonged changing of current above several hundred



**Figure 5.3(b) Motor current to four middle legs of Jacobi's robot at generation 400**

milliamperes is damaging.

### 5.3 Other observations

The emergence of specialization of gaits throughout an evolution run possibly implies less than optimal selection of evolutionary parameters (e.g., the method and rate of crossover, the rate of mutation, the percentage of individuals selected by Darwinian selection, etc.). More analysis is needed in this area.

The quick manner in which the belly-hitting is eliminated seems positive. However, in the eyes of an observer, this also appears as the emergence of cautiousness or even timidity very early in evolution, which might have set the course for the rest of evolution. Most larger swings of the legs were eliminated as the creature tries to avoid

'falls' or accidental squatting that activates one of two belly contact sensors.

An experiment which took about 35 generations created a population of creatures most members of which abandoned the use of the mid-legs. We speculated that this was because of the rigid body structure which is unlike most natural creatures: it was not necessary to activate more than the four legs which supported the robot in the front and back. In other runs, the middle four legs often bore the majority of the burden, possibly for a similar reason.

All motors lasted considerably longer than the guaranteed 100 hours, allowing sufficient experimentation. Application of the drive current is cut if it exceeded a threshold set at about 750mA. The sudden stopping of a leg often resulted in achieving a synchronization between that leg and others. Other legs unloaded the load carried by the stopped leg which resumed its contribution only after the load is sufficiently reduced.

Despite having them in the genotype, the angular velocity did not play a major role in evolving gaits. The speed with which the legs are swung more or less converged to a relatively low average. Along with the timidity observed in swing ranges, we concluded that this was a cautious manner acquired by the robot through evolution. Similar precautionary slowing down in speed of an evolving robot was observed in [8].

In some of the runs, evolutionary parameters were changed during the course of the experiment, as often as every few generations. The backup system was effective in allowing frequent and impromptu halting, modification, and resumption of the run. In one such modified experiment, a system of feeding back subjective evaluation of robots' performance was implemented. This is to directly inject the subjective judgement of the observer of the experiment to the selection process. A keyboard is used to simply tag an individual if its performance is good. This results in a substantial hike in the value of its evaluation function, in most cases enough to have it selected for parenthood. In another run, the evolution was halted to introduce a copy function which copied high scoring leg parameters on one side of the body to legs on the other side across the center line of the body. Desirable results were obtained in both cases.

At this stage of the project we are not systematically chasing the evolutionary parameters. However, once a reasonable rough tuning of the parameters of the experiment is done, we would conduct a more systematic investigation of the runs.

## 6. Conclusions

The gait of an eight-legged walking robot was successfully evolved from scratch without depending on simulation. A few interesting behaviors that defy the intentions of the

designer of the experiment were observed. Unlike simulated evolution, the evolution honored preset hardware limits to protect itself from the destruction of electro-mechanical components of the robot used in the experiment.

## Acknowledgement

The authors are grateful for the cooperation they received from Nick Jacobi of Sussex University for allowing them to study his evolved code on their robot.

## References

- [1] C. Angle, "Genghis, A Six Legged Autonomous Walking Robot", BSc Thesis in Electrical and Computer Science, MIT, May, 1989.
- [2] Brooks, R.A., "A Robot that Walks: Emergent Behavior from a Carefully Evolved Network", *Neural Computation* 1:2, pp. 253-363, 1989.
- [3] P. Maes, and R.A. Brooks, "Learning to Coordinate Behaviors", Proceedings AAAI90, Boston.
- [4] G. Nishi, S. Shinbori, T. Gomi, K. Ide, and P. Maheral., "Mechanimals for Mechquarium". AAI Technical Report 96-03, 1996.
- [5] R.A. Brooks, Lecture note at NATO Advanced Study Institute, Trento, Italy, 1993.
- [6] N. Jacobi, "Half-baked, Ad-hoc and Noisy: Minimal Simulations for Evolutionary Robotics", *Fourth European Conference on Artificial Life (ECAL'97)*, 1997, pp 348-357.
- [7] N. Jacobi, Personal communication, 1997.
- [8] D. Floreano, and F. Mondada,, "Automatic Creation of an Autonomous Agent: Genetic Evolution of a Neural-Network Driven Robot", In *From Animals to Animats 3: Proceedings of Third Conference on Simulation of Adaptive Behavior*, edited by D. Cliff, P. Husbands, J. Meyer, S.W. Wilson. MIT Press, 1993.