

Co-evolving controller and sensing abilities in a simulated Mars Rover explorer

M. Peniak, D. Marocco, A. Cangelosi

Abstract—The paper presents an evolutionary robotics model of the Rover Mars robot. This work has the objective to investigate the possibility of using an alternative sensor system, based on infrared sensors, for future rovers capable of performing autonomous tasks in challenging planetary terrain environments. The simulation model of the robot and of Mars terrain is based on a physics engine. The robot control system consists of an artificial neural network trained using evolutionary computation techniques. An adaptive threshold on the infrared sensors has been evolved together with the neural control system to allow the robot to adapt itself to many different environmental conditions. The properties of the behavior obtained after the evolutionary process has been tested by measuring the generalization performance of the rover under various terrain conditions and especially under rough terrain conditions. In addition, the dynamics of the co-evolution between the controller and the threshold has been analyzed. Those analyses show that different pathways have been explored by the evolutionary process in order to adapt the sensing abilities and the control system.

I. INTRODUCTION

THE exploration of remote Solar System planets with human crews is currently an enterprise impossible to realise. Besides the technical difficulties, the main issue regards the huge distances involved and the long time required to reach such remote regions of the Solar System. It is expected that in the near future, robotics and autonomous robots in particular, will play an essential role in planetary exploratory mission. When communication delay between the robot and the Earth is hours, devising advanced autonomous capability for an exploring robot is the only route toward the expansion of our knowledge into deep space. Therefore, full autonomy of robots and their ability to rely on their own abilities to accomplish mission operations will be more and more crucial in the future.

Mars Pathfinder, launched in 1997, has been the first exploratory mission in which a semi-autonomous vehicle, called Sojourner, landed on the Martian surface. After Mars Pathfinder, more sophisticated robots, such as the rovers Spirit and Opportunity, were landed on Mars in 2004. The rovers were designed to withstand harsh Martian conditions for only 90 days, although after five years they are still exploring Mars and bringing new discoveries [1]. The future NASA's rover mission is called Mars Science Laboratory (MSL) and it is to be launched in 2011. This mission involves a rover carrying more sophisticated instruments that

will help answering questions about Mars history, climate, geology, possible life and it will also prepare for future human exploration. The rising interest on this type of missions is also strengthened by the fact that, alongside the NASA projects, several other projects are under development by the European Spatial Agency, as well as China and Japan.

The ability to navigate efficiently within an unknown environment, by autonomously avoiding obstacles, is a basic ability that every exploratory robot has to show in order to be effective. In addition, the obstacles can have different characteristics, such as rocks, holes in the terrain or a particularly rough surface that could be potentially dangerous for the robot. These differences require the robot to have the ability to distinguish between the different types of obstacles and actuate the appropriate avoidance maneuvers.

The above-mentioned rovers Sojourner, Spirit and Opportunity use stereo cameras for navigation and obstacle avoidance. The two more recent robots Spirit and Opportunity, in particular, are equipped with three sets of stereo camera pairs. One pair is looking forward, under the solar panel in front. Another pair is looking backward, under the solar panel in the back, and the last pair is placed on the mast. This camera is mainly used for navigation purposes. With the images taken by the cameras, a stereo algorithm calculate the 3D representation of the terrain in front of the robot and other algorithms are used to calculate a "traversability" map [2]. The information of this map is then used to calculate the next action of the robot. However, there are no other means for the rovers to sense the obstacles if these cameras failed. For this reason, it is worth to explore other possible solutions that allow the rovers to navigate and avoid obstacles, besides the use of stereo cameras. These alternative methods might represent useful complements in the sensory systems of robot which has to operate in difficult conditions into deep space, where any possible human intervention is prevented by the huge communication delays.

In this paper we will explore the feasibility of an alternative obstacle avoidance system based on a set of infrared sensors that provide the robots with information about the presence of obstacles within a given range in its proximity. The system presented is able to deal with different types of objects, such as rocks, holes and moderately rough surface.

To investigate this alternative methodology, a 3D physics rover as well as a terrain model was built using Open Dynamics Engine (ODE), which is an open source library for simulating rigid body dynamics (www.ode.org). The computer model of the rover is based on the approximate dimensions of the MSL rover and its control system consists

Manuscript received November 14, 2008. This work was supported in part by the EU euCognition grant.

M. Peniak, D. Marocco and A. Cangelosi are with the School of Computing, Communications and Electronics, University of Plymouth, Drake Circus, Plymouth PL4 8AA, United Kingdom (e-mail: [@plymouth.ac.uk](mailto:martin.peniak@davide.marocco/angelo.cangelosi)).

of an artificial neural network (ANN) whose synaptic weights were evolved using evolutionary computation techniques.

We are aware that within the field of evolutionary robotics, obstacle avoidance and navigation behaviors are well known topics that have been widely used in the past to demonstrate the feasibility of the evolutionary approach in the robotic domain. Alongside the new focus in this paper on the co-evolution of controller and sensor properties, such as the infrared sensors threshold (for previous works on sensor evolution see [22]), this paper is a preliminary attempt to extend the domain of the evolutionary techniques to interplanetary robotics to demonstrate the potential feasibility and robustness of such an approach when applied to the realm of planet exploration. This requires the consideration of the complexity of a hypothetical exploratory mission on a planetary surface, such as the issues of exploring, in a safe mode, an unknown environment by autonomously finding an effective route on a rough surface full of unpredictable obstacles and by taking into account the limited computational capability of the on-board hardware [5]. The accomplishment of such a task requires, on one hand, a control system that must be able to sense the different types of obstacles and to deal with a rough terrain that can often make it hard to navigate on it. The robot should be able to autonomously assess when a terrain is safe for navigation or when it is better to change direction. On the other hand, the limited on-board computing and electrical power forces us to reduce the complexity of the algorithms that provide the required navigation capabilities.

Navigation in rough terrain is a topic that has been addressed in different ways. A number of projects have employed behavior-based navigation methods [17][18]. In these studies, the whole behavior of the robot is the outcome of a complex interaction between simple sub-behaviors predefined by the researchers. Some researchers have used the so-called “arcs approach” [19][20]. In the arcs approach, an algorithm is devoted to generate several candidate arcs and, after a confrontation, one of the arc is chosen on the basis of some criteria (i.e. the arc with the largest clearance or, after calculating the costs along each arc, the one with the lowest cost is selected) The robot is then steered along the winning arc. In other works, the steering of the robot is calculated by creating a local terrain map surrounding the robot position into a grid-type traversability map [21].

In evolutionary robotics, the most recent studies that explicitly address the issue of the navigation in rough terrain, by avoiding obstacles and holes, are mainly based on coordinated motion behavior. This approach aims to solve the problem by the evolution of complex coordinated behaviors of simple interconnected mini-robots [6]. Another approach is based on the idea of reconfigurable robots, where robots can adopt different shapes in order to cope with different environmental conditions [7][8][9]. In contrast to the previous studies, our intention is to use a single robot, based on a model of the MSL rover, and investigate whether it would be possible to evolve a neural network controller

able to tackle obstacles like walls, different rocks, rough terrain as well as holes and cliffs.

In this paper we developed a simulation model of the MSL rover that is equipped with eighteen infrared sensors and a controller based on a Perceptron neural network.

Because it is necessary to evolve a robot that can deal with both rocks and holes, instead of providing the robot with different types of sensors, we provided it with an evolvable threshold on the sensors. This threshold adaptively modifies the activation range of the infrared sensors, in order to use front sensors for both rocks and holes detection. In this way we allow the evolutionary process to discover the best way of integrating the control system with the characteristic of the body, by adapting the sensing ability of the robot itself. As we will see in the next section, the threshold represents a modulation of the sensors that directly affects the way in which the robot “senses” the environment and, indirectly, its behavior. Therefore, the value of the threshold must be well tuned with the robot’s characteristics, as well as the characteristics of the environment.

The threshold, which is evolved together with the control system, can differentiate rocks and holes from the noise originating from rough terrain, and has been set by means of a co-evolutionary process between the rover’s behavior and the threshold itself. This suggests that both behavior and threshold are interdependent. The system was evolved in an environment that contained many different rocks, cliffs, holes, walls and areas of rough surface. Results from the experiments and testing showed that the system is very robust and it is able to adapt to different surface conditions.

In the following sections we describe our methodology, which involves a detailed description of the rover model, its neural network controller and the genetic algorithm (GA) used to evolve the connection weights of the neural network. We will present in detail the experimental setup used throughout all evolutionary runs and the obtained results. In order to show the reliability of the evolved system, we ran a series of tests that measured the robustness and adaptability to different environmental circumstances. Furthermore, we will present a series of tests and analysis that allow to clarify the role of the threshold and the different evolutionary pathways undertaken by the co-evolutionary process.

II. METHOD

As we have mentioned in the introduction, our approach is based on evolutionary robotics (ER) [3]. The ER approach emphasizes agent’s embodiment, which means that an emerging behavior is not only dependent on various properties of the actual robot such as its size, speed, degrees of freedom, sensors and actuators, but also on the environment with which a robot interacts [10]. ER is an excellent technique that allows us to create artificial control systems that autonomously develop their skills in close interaction with the environment and that exploit very simple, but extremely powerful sensory-motor coordination [11].

A. The Rover Model

The robot used in this experiment is a 3D simulation model of the MSL rover. The model cannot be considered as a trustful and detailed representation of the actual rover, but only an approximate copy. This is mainly because of the lack of information on the rover's real dimensions, weights and sizes of different parts, as well as of many other details. According to Centre National d'Etudes Spatiales [14], the dimensions of the real rover are 2900Lx2700Wx2200H mm and its weight is about 775 kg. The physical rover model was therefore built considering these details and several diagrams and pictures that were available. These limitations are in this study not crucial as at this stage we want to demonstrate that it is possible to use an ER approach and a simple sensory setup to develop a suitable controller able to handle complex obstacle avoidance tasks in unknown rough terrains.

The motor system of the rover model (see Fig. 1a) consists of six wheels where two front and two rear wheels are able to turn up to 90° to either side. The rover is capable of overcoming obstacles that are approximately of the size of its wheels. This is possible thanks to a rocker-bogie suspension system. This advanced suspension system is designed to be operating at low speed, and consists of two pivotal joints connecting two bogies with two rockers [15]. The rockers are connected together via a differential joint. This means the left and right part of the rocker-bogie system can move independently while keeping the main body leveled.

The rover is equipped with a sensory apparatus that comprehends 18 infrared sensors in order to provide sufficient information from the surrounding environment. Two different sets of sensors were used to accommodate detection of various obstacles (see Fig. 1b). The first set consists of six lateral sensors which provide extra safety when the robot approaches obstacles from a side. These sensors have a range of three meters and are not able to detect holes. Lateral sensors cover an area of approximately 200° around the rover, leaving the front area deliberately uncovered. These sensors return either 0 (no obstacle) or 1 (obstacle present), when the sensor is activated by the presence of an obstacle within the activation range of the sensor.

The second set consists of 12 infrared sensors with the maximum reach of five and half meters. These infrared sensors, which we call ground sensors, are positioned on the rover's camera and pointing downward at 45° angle and reaching the ground approximately three meters in front of the rover. The twelve sensors are positioned and directed so that they are able to reach around 400 mm more than the level of the ground. Ground sensors constantly scan the distance from the surface and are able to detect both rocks and holes. Each of these sensors returns a floating point value from 0 (no feedback) to 1 (strongest feedback). Holes or cliffs can be detected by the rover when it loses sensory feedback from the ground (i.e. ground sensor returns a value

0). The same sensors allow the robot to detect dangerous rocks or excessively rough terrain. This is achieved thanks to a particular threshold. When the activation of a sensor reaches that threshold it means that the robot is facing an insurmountable rock or a potentially dangerous rough terrain. If a sensor's output goes over this threshold (a rock) or returns 0 (a hole) then its output value is changed from 0 (not active) to 1 (active). On the other hand, if the returned value stays within a certain boundary, which is given by the threshold, then a sensor returns 0. From this perspective a 0 activation can be seen as safe zone and 1 as an obstacle in the front. To model the lateral sensors and the ground sensors we aimed to simulate the existing infrared sensors Sharp 3A003 and Sharp 0A700, respectively.

In earlier experiments the threshold was manually set to 0.2, which was estimated from the terrain roughness. However, the robot did not evolve a good behaviour with this threshold value as it had problems distinguishing roughness from flat terrain. In order to provide the robot of more flexibility and allow the system to be completely free to adapt autonomously to the environment, the value of the threshold was not pre-set, but rather evolved throughout the evolutionary process. In this case the evolutionary process can find a threshold value which is more suitable to the physical characteristic of the rover and to a particular environment. Threshold can be in a range $[0,1]$.

In addition to the above sensors, the rover is provided with a couple of internal sensors measuring its speed and steering angle.

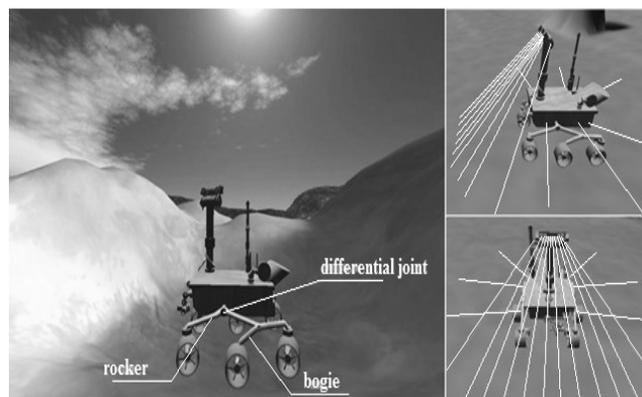


Fig. 1. 3D physics model of the rover highlighting different parts of the rocker-bogie suspension system (left). Side view (top right) and front view (bottom right) of the rover showing lateral and ground sensors and their positions.

B. System Architecture and Evolutionary Parameters

The control system is a fully-connected feedforward ANN with evolvable bias and discrete time (see Fig. 2). A set of 18 sensory neurons receive the activation from the 18 infrared sensors of the rover and an additional set of 2 proprioceptive neurons encode the value returned by the internal sensors, which provide information about the speed and the position of the wheels. The 20 sensory neurons are fully connected to 2 motor neurons that modulate the level

of the force which is applied to the actuators, which are directly responsible for rover's speed and steering, respectively. Motor neurons have sigmoid activation functions:

$$f(x) = \frac{1}{1+e^{-x}} \quad (1)$$

in the range [0, 1], where x is the weighted sum of the inputs minus the bias. Biases are implemented as a weight from an input neuron with an activation value set to -1. The ANN does not have a hidden layer as our previous experiments showed that it was redundant and did not help to achieve higher fitness. This simple architecture greatly reduces the computation demand of the control system, which is one of the most important requirements for designing a planetary rover.

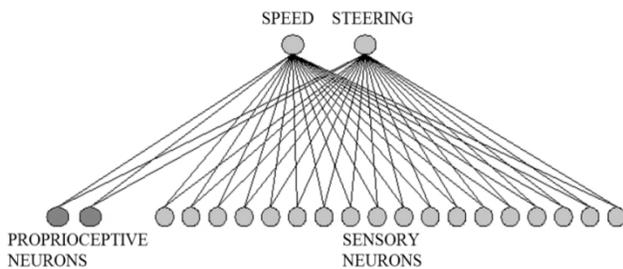


Fig. 2. Feed-forward neural network used as a control systems for the rover in the evolutionary experiments.

Rover's actions depend on the value of the synaptic weights of the ANN. So that, each weight must be set to an appropriate value to produce a desired output and, as we mentioned before, a genetic algorithm was used to evolve them. The free parameters that constitute the genotype of the control system and that are subject to evolution consist of: 42 synaptic weights (the 40 synaptic weights that connect the 20 sensory neurons to the 2 motors neurons, plus the 2 biases) and a single gene which encodes the threshold applied to the ground sensors. Weights and biases are encoded as floating point values in the range [-1, 1] and the threshold in the range [0, 1].

In our experiments we used a population size of 100 individuals, where the best 20 individuals were allowed to produce 5 offspring each with a mutation probability of 10% (a mutation occurs by adding to the original gene's value a quantity in the range [-1, 1]). The only exception was the first offspring of the best individual, which was copied to the next generation without mutation. This is often known as *elitism* where the best solution is always preserved by not allowing mutations to change its genes. In this way we produced a new generation of 100 individuals that inherit their genes from the best individuals of the previous generation. The whole evolutionary process lasted 100 generations. In each generation, each control system was tested 10 times by deploying it in the rover (randomly positioned and rotated) and allowing it to act in the environment for up to 3000 sensory-motor cycles, that is, 3000 activations of the ANN. However, this was not always

the case, as the evaluation of a particular genotype was terminated when a rover fell into a hole or crashes into an obstacle. 10 evolutionary runs were conducted where each of these was initialized with a different randomly generated population.

The performance of every single control system was evaluated according to the fitness function (2) that was carefully designed to shape the behavior of the robot for effective and reliable exploration and obstacle avoidance behaviors:

$$F = \frac{1}{S \cdot T} (Sp \cdot St) \quad (2)$$

where the fitness F is a function of the measured speed Sp and steering angle St , where Sp and St are in the range [0,1]. Speed Sp is 1 when the rover goes at the maximum speed and 0 when it does not move or goes backward. Steering angle St is 1 when wheels are straight and 0 when they are turned over an angle of 30° from the centre. If for example the angle was 15° then St would be 0.5. T is the number of trials (10 in these experiments) and S is the number of sensory-motor cycles per trial (3000 in these experiments). Equation (2) shows how the fitness is calculated at every sensory-motor cycle. Thus, the GA has to maximize the fitness by increasing the value of Sp and St , which implies that a rover has to move at a maximum possible speed while steering only when necessary. In fact, if a rover goes forward at the maximum speed but keeping the steering angle over 30° then its final fitness would be 0. Similarly, if a rover goes backwards or does not move at all, its fitness would also be 0 regardless the steering angle. The maximum fitness contribution at each time step is therefore $1/(S \cdot T)$. The final fitness of each individual is in a range [0, 1] and it is the sum of all contributions from all time steps of all trials.

In order to evolve a good controller, it was necessary to create a suitable environment to allow the robot to experience different surface conditions (see Fig. 3.). The environment that was modeled for this purpose is an arena of 60x60m and contains inclined and declined surface, three high and three small rocks, holes and rough areas. 111m² of the terrain is covered by obstacles and hence not traversable.

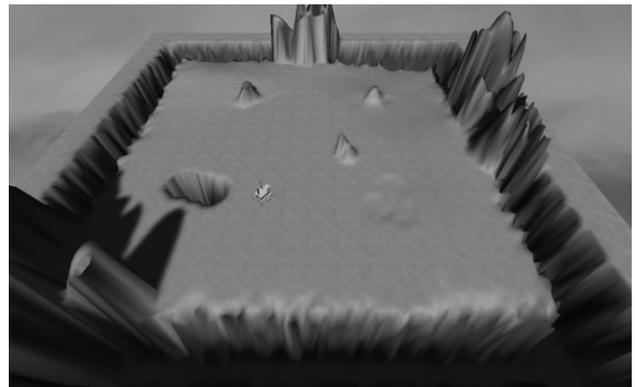


Fig. 3. Environment that was used during all evolutionary runs.

III. RESULTS

The results obtained from all the ten evolutionary experiments show that an effective behavior emerged in all evolutionary runs. In particular, thanks to the general behavior optimized by the fitness function and the evolutionary threshold, we obtained robots that can navigate the environment with a certain degree of efficacy and are able to avoid obstacles of different types by dealing with a rough terrain.

The chart in Fig.4 shows the average results of the ten evolutionary runs. The graph was created by averaging values from all the ten runs. The dark grey line shows the maximum fitness obtained by the best individuals, the light grey line the average fitness of all the populations and the dashed line shows the threshold value across the generations. By looking at the graph it can be noticed that while the maximum and the average fitness are increasing the threshold is decreasing and reaching the optimum value of about 0.3 by 50th generation. With this optimized threshold the rover can detect all the rocks present in the terrain while not being confused by its roughness.

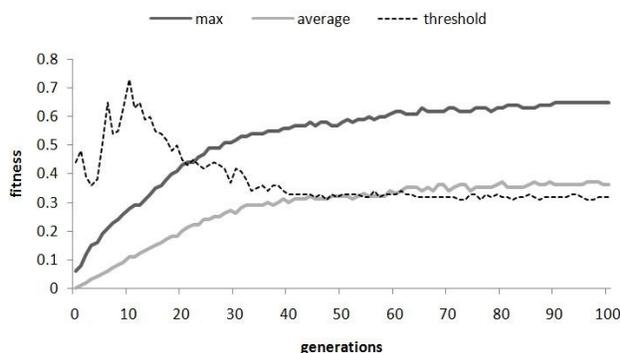


Fig. 4. Fitness graph showing maximum and average fitness as well as the threshold. Note that the fitness can never reach 1.0 as the rover needs to turn and decrease its speed to avoid obstacles

A number of results from different evolutionary runs showed dramatic changes in the fitness after a suitable threshold value was found. This suggests that a good behavior can only emerge if a suitable threshold value is found. This point will be discussed in the next section. Another interesting finding was that even the few evolutionary runs that did not end up with high fitness were capable to evolve good obstacle avoidance. In order to understand the changes in fitness, as well as the differences between different experiments, several tests were conducted. In particular, tests were designed to evaluate the system's robustness in terms of performance, reliability and adaptability to new conditions. These properties of the evolved controllers were examined using two different tests where the time for genotype evaluation was lengthened to 10,000 sensory-motor cycles to make sure the system is robust. The first test measured the fitness of the best ten controllers. For this purpose, the best controller from the last

generation of each run was evaluated. Each of these controllers was tested 100 times from random initial positions/rotations and average fitness was recorded. This process was repeated on two other terrains (same width and length). One terrain had the same obstacles but extra roughness, which means that the whole surface was rough. The other terrain had extra rocks and holes. The left graph in Fig. 5. shows the average fitness of all evolutionary runs for the basic terrain. Average fitness value of controllers tested on original or rough terrain is around 0.5. This number drops dramatically on the terrain with more obstacles and reaches the value of 0.38. However, this is not surprising as the fitness is affected by the rover steering. In this terrain, the rover had to turn much more than in the original terrain, which reflected in the lower fitness. The second test measured the exploration ability of the best controllers. The main purpose of this test was to have a more reliable measure of the system performance. It is not surprising that fitness decreases when the rover is tested in an environment in which it is required to steer much more. Therefore, we conducted an additional test, which should reveal whether our system is robust or not. For this purpose, the three terrains were therefore divided into 400 square blocks (20x20), each being 3x3meters long. In this test, we recorded the number of squares that a particular controller was able to visit. Same as in the previous test, each controller was tested 100 times from random positions/rotations. The average of these trials was taken and used for the statistics where we show the percentage of the terrain that the rover was able to explore within a given time (10,000 sensory-motor cycles). Note that this percentage considers only those squares that the rover can visit. Hence, squares covering areas with holes and rocks were not considered as it can be seen from (3), where E is the percentage of the explored terrain, $S_{visited}$ is the number of visited squares, S_{total} is the total number of squares and finally $S_{obstacles}$ is the number of squares covering obstacles (37 for the first two terrains and 91 for the terrain with more obstacles).

$$E = \frac{S_{visited}}{S_{total} - S_{obstacles}} \quad (3)$$

This approach help us to understand the extent to which the evolved system is robust as this test is not so much affected by the number of obstacles in the terrain. As it can be seen from the right graph in Fig. 5 there is only a slight difference in exploration success on the three terrains. The average exploration was 41.8% on the original terrain, 42.4% on the rough terrain and 38.3% on the terrain with more obstacles. The results obtained from the terrain with more obstacles deviate more (3.5%) from the original terrain than the results from the rough terrain (0.6%). However, this small difference is negligible and it seems to be caused by the fact that the rover tends to explore more often same areas of the

terrain. It is more likely for the rover to explore less of the environment if there are many obstacles, which cause the rover to visit the same places more than once, rather than moving over new areas. In other words, the presence of many obstacles makes it less likely that all parts of the terrain are explored within 10,000 sensory-motor cycles.

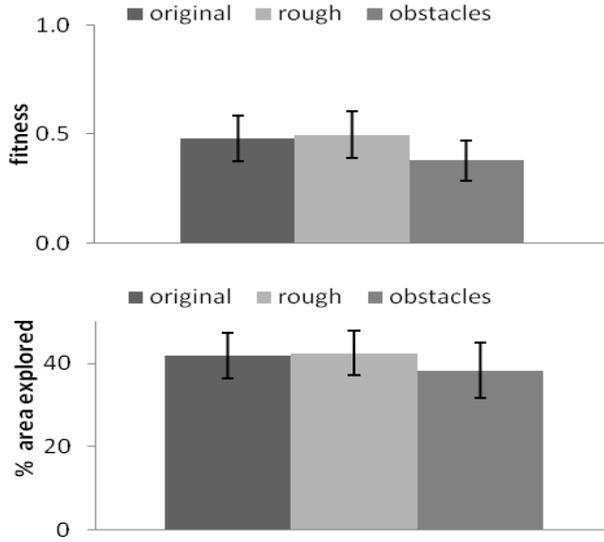


Fig. 5. Graphs showing average fitness (left) and exploration (right) of all evolutionary runs.

IV. CO-EVOLUTION OF THRESHOLD AND CONTROL SYSTEM

The co-evolution between the threshold value and the neural network weights resulted in several different evolutionary scenarios. All of these are different, however, they all point out that this co-evolution is important for achieving a good behavior of the robot.

We conducted a series of tests to confirm our assumptions that the neural network and the threshold co-evolved differently in each of the cases in Fig. 6.

In some cases the evolution started to exploit behavior only when a suitable threshold value predominated a particular population. This can be noticed from the Fig. 6 1.a), where at around the 20th generation the threshold value started to oscillate between 0.29 and 0.49, which seems to be very close to the upper and lower boundary for a usable threshold. Only when this threshold settled and spread throughout the whole population the evolution started to work on exploiting the behavior. It seems that there is a certain degree of threshold stability that is necessary in order to enable the evolution to discover and exploit a good behavior. This is apparent from the Fig. 6.1a), where the fitness began to increase at around 50th generation when the threshold value got stabilized. Tests did not show any difference in fitness when the best evolved individual was tested with the threshold value of 0.369 instead of its original threshold of 0.346 that was characteristic for the evolutionary stage where the obstacle avoidance behavior was not present (see Fig. 6 1.b)).

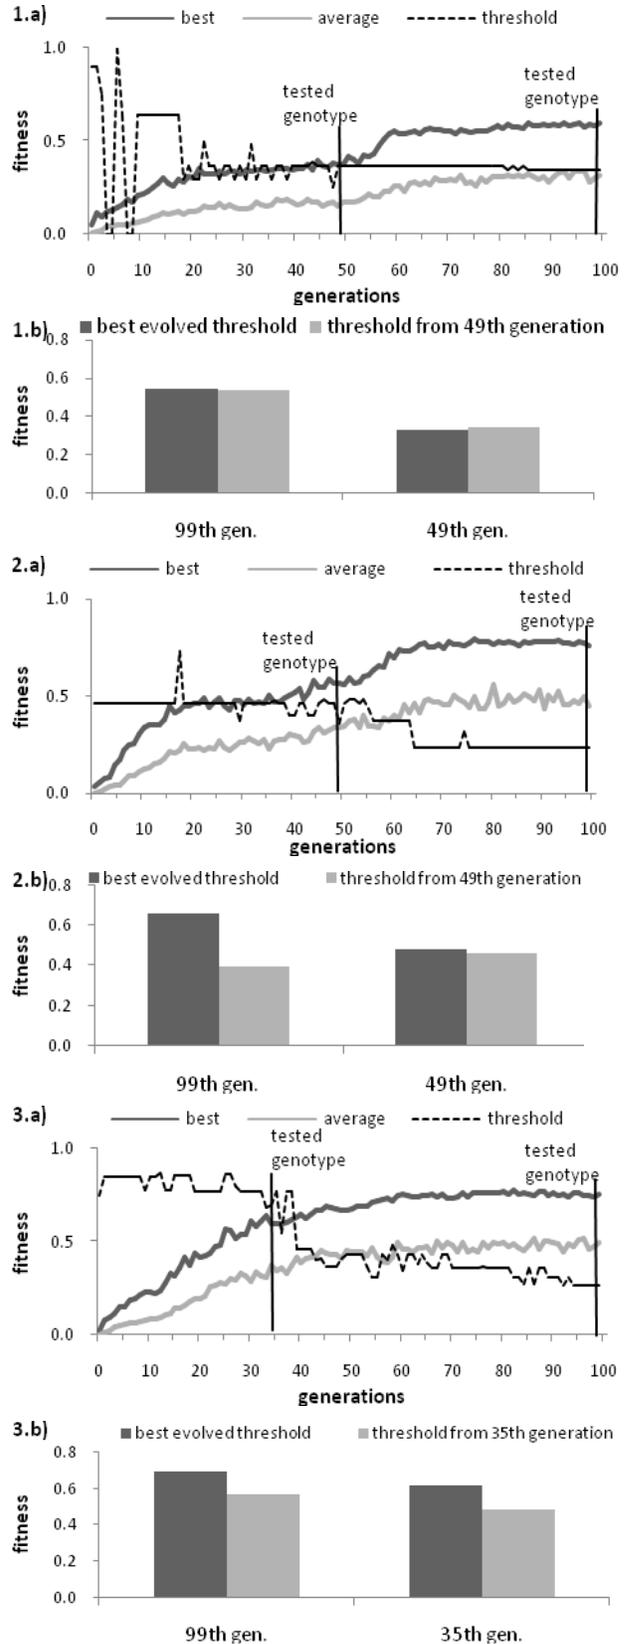


Fig. 6. Fitness graphs together with changing threshold value from selected evolutionary runs. Column graphs show the results from testing of this threshold on best evolved individuals and also on those previous evolutionary stages (see text for details).

This effect is mostly due to the similarity of both threshold values. We can conclude that in this case the evolution was not able to discover a good behavior when there was a lot of instability in the threshold. As we have already mentioned, only when the threshold is stabilized the good behavior can be found. In such a case, therefore, the co-evolution resulted in the development of good controllers upon stabilizing the threshold value.

In a different run the evolution was not able to exploit a certain behavior with the threshold of around 0.48. Only when the threshold started to decrease after 55th generation the fitness started to rise. This can be noticed from the graph in Fig. 6. 2.a), where the best fitness line looks like a mirror image of the threshold line. This evolutionary process of lowering the threshold value from 0.48 down to 0.23 led to increased fitness from 0.55 up to 0.79. This suggests that the evolutionary process could not shape the weights in the neural network when the threshold was high. To confirm this and to understand better the interdependence between a particular behavior and a threshold we ran the same test as in the previous case. In contrast to the first co-evolutionary type, where the evolution was waiting for the threshold stability, this co-evolutionary case displays different properties. The test that was conducted (see Fig.6 2.b) shows that fitness of the best evolved individual is significantly affected when its original threshold value (0.23) is changed for the threshold value present in the 49th generation (0.36). This suggests that the behavior was very dependent on the original threshold, which also highlights the importance of the co-evolution between the two. In contrast, when the best individual from 49th generation was tested with the best threshold the fitness did not significantly change, which implies that at that evolutionary stage the obstacle avoidance behavior as well as its dependence on the threshold was not present.

The last two co-evolutionary examples showed that fitness increased either by stabilizing a threshold value (the first case) or by minimizing it to a suitable number (the second case). However, in both of these runs a reasonably good obstacle avoidance behavior was not present. It only began to emerge after either of these changes took place. In contrast to this, the last case (see Fig. 6 3.a) is very different. Here the fitness increased to 0.59 with a surprisingly high threshold of 0.77. With such a high threshold the rover is not able to detect the obstacles, however, it is still able to detect holes and walls. It was apparent that even at this evolutionary stage there was already a reasonably good behavior that allowed the rover to avoid most of the obstacles apart from the rocks and the rough surface. It seems that during this evolutionary stage only those individuals were selected that did not come across many rocks. It is worth to note that these individuals had already a good behavior, however, the high threshold value did not allow them to detect rocks in most of the cases. To demonstrate that this was the case we ran the same test as in the previous two cases. This time we tested the individual from 35th generation, as this was the early evolutionary stage after which the threshold value started to decrease. When the best individual was tested with the original

threshold (0.26) the fitness reached the value of 0.62. On the other hand, when the same individual was tested with the threshold value typical for 35th generation (0.77), the fitness dropped to 0.57. In this case, the high threshold simply prevented the rover from detecting these rocks, which resulted in lower fitness. More interesting was the test in which the best individual from 35th generation was tested with the best threshold from 99th generation (0.26). The results proved that our assumption that the behavior was already present at these early evolutionary stages was right because the fitness increased dramatically when the rover could use the best threshold value. As it can be seen from the Fig. 6. 3.b) the fitness increased from 0.49 to 0.62, which is very close to the fitness of 0.69 achieved by the best evolved rover. This test proves that the good obstacle avoidance behavior was already present during early evolutionary stages. However, higher fitness was not achieved due to the fact that the system could not fully optimize as a result of high threshold.

The present analysis on the co-evolutionary process show that, in order to obtain a suitable behavior a fine adaptation of the control system and the threshold is required and the evolutionary strategy exploited by the system can be different. This fact indicates that the co-adaptation between the sensing abilities and the controller is a complex process that takes into account all the factors involved in the development of the behavior: i) the body and its possibility to adapt to the environmental characteristics, through the evolution of the threshold; ii) the environment, that acts in shaping the value of the threshold and, finally, iii) the control system that has to move the robot on the basis of the information provided by sensors. None of these factors is predominant in the evolution of the adaptive behavior and the analysis clearly show that the evolution can work indifferently by exploiting any possible interaction between parameters: by stabilizing the threshold and then adapting the control systems, by stabilizing the control system and then working on the threshold or by working at the same time on both factors.

V. CONCLUSIONS

We have shown that the rover model equipped with the evolved neural network controller is able to deal with different types of obstacles by distinguishing between terrain roughness noise and dangerous obstacles thanks to the evolvable threshold. The threshold allows the rover to adapt its "sense" to the characteristic of the environment and it is particularly relevant for dealing with terrains of different roughness.

Our tests indicate that the system is very robust and able to maintain the obstacle avoidance behavior under different circumstances and in different environments.

It is worth noting that the exploration and the obstacle avoidance behaviors are not obtained through a pre-designed pattern of interaction between the rover and the environment. Rather, they are the emergent product of a fitness function that works at the level of the whole behavior of the robot. Those behaviors are actually discovered autonomously by the evolutionary process and are functional

to the optimization of the global fitness used for the evolution. We are convinced that this property of evolutionary robotics can be very useful to design a robust and computationally light controller, capable to deal with some of the peculiar problems which will be facing the future planetary robotics missions. As we have shown in this work, the evolved neural network controllers can be extremely simple, require only a minimum processing power and yet be very robust and reliable.

In the future we plan to use this system together with an active vision pan/tilt camera that would provide the rover with navigation capabilities. Active computer vision systems are inspired by information gathering of mammals and insects. Such systems can greatly simplify the computational complexity as they only use information from an environment that is necessary to solve a certain task while the rest is ignored. Past research in this field demonstrated that it is possible to combine an active vision system together with feature selection to acquire and integrate information from an environment in order to solve a specific task [16]. Hence, our future goal is to use both the active vision system and the current system to achieve complex, robust and reliable, yet computationally cheap behaviors. We are aware that future planetary robotics missions will have to face many challenges and we are convinced that evolutionary robotics is worth to be considered as a possible approach that could address several problems that are hard to overcome using conventional methods.

REFERENCES

- [1] Jet Propulsion Laboratory, "Summary," *Mars Exploration Rover Mission*, 12 July 2007, <http://marsrovers.jpl.nasa.gov/overview/> (accessed October 20, 2008).
- [2] S. Goldberg, M. Maimone and L. Matthies, "Stereo Vision and Rover Navigation Software for Planetary Exploration," *IEEE Aerospace Conference*, Big Sky, Montana, March 2002.
- [3] S. Nolfi and D. Floreano, *Evolutionary Robotics*, Cambridge, Massachusetts: MIT Press, 2001.
- [4] S. Nolfi and D. Marocco, "Active perception: A sensorimotor account of object categorization," in *From Animals to Animats 7, Proceedings of the 7th Intl. Conf. on Simulation of Adaptive Behavior*, B. Hallam, D. Floreano, J. Hallam, G. Hayes and J. A. Meyer, Eds. Cambridge, MA: MIT Press, 2002.
- [5] M. Maimone, A. Johnson, Y. Cheng, R. Willson and L. Matthies, "Autonomous Navigation Results from the Mars Exploration Rover (MER) Mission," *Proceedings of the 9th International Symposium on Experimental Robotics (ISER)*, Singapore, June 2004.
- [6] V. Trianni, S. Nolfi and M. Dorigo, "Hole avoidance: Experiments in coordinated motion on rough terrain," *Proceedings of IAS-8: Intelligent Autonomous Systems Conference*, pp. 29-36, 2004.
- [7] A. Castano, W. Shen and P. Will, "CONRO: Towards deployable robots with inter-robot metamorphic capabilities," *Autonomous Robots*, vol. 8, pp. 309-324, 2000.
- [8] K. Sty, W. M. Shen and P. Will, "Global locomotion from local interaction in self-reconfigurable robots," in *Proceedings of the 7th International Conference on Intelligent Autonomous Systems (IAS-7)*, W. M. Shen, C. Torras and H. Yuasa, Eds. Amsterdam, The Netherlands: IOS Press, 2002.
- [9] M. Yim, D. G. Du and K. D. Roufas, "PolyBot: A modular reconfigurable robot," *Proceedings of the 2000 IEEE/RAS International Conference on Robotics and Automation*, vol. 1, pp. 514-520, 2000.
- [10] R. Beer, "A dynamical systems perspective on agent environment interaction," *Artificial Intelligence*, vol. 72, pp. 173-215, 1995.
- [11] S. Nolfi, "Power and Limits of Reactive Agents," *Neurocomputing*, vol. 42, pp.119-145, 2002.
- [12] H. Haykin, *Neural Networks: A comprehensive Foundation*, 2nd ed., Singapore: Pearson Education, 2005.
- [13] J. Holland, *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: University of Michigan Press, 1975.
- [14] Centre National d'Etudes Spatiales, *MSL 09 at a glance*, <http://www.cnes.fr/web/5719-msl-09-at-a-glance.php> (accessed 20 September 2008).
- [15] D. P. Miller and T. L. Lee, "High-Speed Traversal of Rough Terrain Using a Rocker-Bogie Mobility System," *Proceedings of Robotics 2002: The 5th International Conference and Exposition on Robotics for Challenging Situations and Environments*, 2002.
- [16] D. Floreano, T. Kato, D. Marocco and E. Sauser, "Co-evolution of active vision and feature selection," *Biological Cybernetics*, vol. 90, pp. 218-228, 2004.
- [17] D. Langer, et al., A Behavior-based system for off-road navigation, *IEEE Transactions on Robotics and Automation*, vol. 10, no. 6 pp. 776-783, 1994.
- [18] H. Seraji and A. Howard, Behaviour-based robot navigation on challenging terrain: a fuzzy logic approach, *IEEE Transactions on Robotics and Automation*, vol. 18, no. 3, 2002.
- [19] S. Singh et al., Recent progress in local and global traversability for planetary rovers, *Proc. IEEE International Conference on Robotics and Automation*, pp. 1194-1200, 2000.
- [20] S. Lacroix et al., Autonomous Rover Navigation on Unknown Terrain: Functions and Integration, *International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 917-942, 2002.
- [21] C. Ye and J. Borestein, A Method for Mobile Robot Navigation on Rough Terrain, *Proc. IEEE International Conference on Robotics and Automation*, pp. 3863-3869, 2004.
- [22] K. Dautenhahn, D. Polani and T. Uthman, Guest Editors' Introduction: Special Issue on Sensor Evolution, *Artificial Life Spring 2001*, vol. 7, no. 2, 2001.