

# AUTONOMOUS ROBOT EXPLORATION OF UNKNOWN TERRAIN: A PRELIMINARY MODEL OF MARS ROVER ROBOT

Martin Peniak<sup>(1)</sup>, Davide Marocco<sup>(1)</sup>, Angelo Cangelosi<sup>(1)</sup>

<sup>(1)</sup>*Adaptive Behaviour and Cognition Research Group - School of Computing, Communications & Electronics  
University of Plymouth, Drake Circus, Plymouth PL4 8AA, United Kingdom*

*martin.peniak@plymouth.ac.uk; davide.marocco@plymouth.ac.uk; angelo.cangelosi@plymouth.ac.uk*

## 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 behaviour obtained after the evolutionary process has been tested by measuring the performance of the rover under various terrain conditions. Simulations results show that the robot, at the end of the evolutionary process, is able to avoid rocks, holes and steep slopes based purely on the information provided by the infrared sensors.

## INTRODUCTION

The history of planetary exploration trace back to the 20<sup>th</sup> July 1969, when the first human footprint was impressed on the surface of the moon. However, exploring other planets with human crews is currently impossible to realize. 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. For that reason, robotics and autonomous robots in particular, will play an essential role in the future of planetary exploration. Autonomy is crucial as the more a robot is far from the earth, the more it should be able to rely on its own abilities to accomplish its 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.

Only recently, under the mission Mars Pathfinder, the first ever robotic exploration vehicle, called Sojourner, landed on the Martian surface in 1997. 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 four 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 2009. This mission involves a rover carrying more sophisticated instruments that will help answering the questions about Mars history, climate, geology, possible life and it will also prepare for future human exploration. Alongside the NASA projects, several other projects are under development by the European Spatial Agency, as well as China and Japan.

Among the several tasks that a robot devoted to explore a planet surface has to accomplish, the ability to move autonomously within an unknown environment is a basic one. In particular, such a robot must be capable of navigating in a new environment and avoiding obstacles that force the robot to deviate from its route. In addition, the obstacles can have different characteristics, such as big rocks or holes in the terrain. These differences require the robot to have the ability to distinguish between the different types of obstacles and actuate the appropriate avoidance manoeuvres. 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 and holes, and it is based on evolutionary robotics (ER) techniques. 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](http://www.ode.org)). The computer model of the rover is based on the approximate dimensions of the MSL rover and its “brain”, its control system, consists of an artificial neural network (ANN) which synaptic weights were evolved using evolutionary computation techniques. This approach is commonly known as evolutionary robotics [3]. Evolutionary robotics is inspired by the Darwinian principle of selective reproduction of the fittest and attempts to develop sensory-motor control systems for autonomous robots in an automated manner.

Within the field of evolutionary robotics, obstacle avoidance and navigation behaviours are well known topics that have been widely used in the past to demonstrate the feasibility of the evolutionary approach in the robotic domain. In particular, those behaviours have been the ideal test bed used by evolutionary robotics to show the inseparable interconnection between the control system, the body and the environment in which the robot is operating [4]. Alongside the scientific interests that often underpin the experiment in evolutionary robotics, the practical aim of this paper is to extend the domain of the evolutionary techniques to the realm of planet exploration. To do that, not only we will have to evolve a control system capable of avoiding obstacles, but we need to face all the complexity of an hypothetical exploratory mission on the planetary surface, i.e. exploring an unknown environment by autonomously finding an effective route on rough surface full of obstacles in a safe mode 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 often can make hard to navigate on it. That is, the robot should autonomously understand when a terrain is safe for navigation or when it is better to change direction. On the other hand, the limited on-board computing power forces us to reduce the complexity of the algorithms that provide the required navigation capabilities.

In evolutionary robotics, the most recent works that explicitly address the issue of the navigation in rough terrain, by avoiding obstacles and holes, are mainly based on coordinated motion behaviour. This approach aims to solve the problem by the evolution of complex coordinated behaviours 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 similar to 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 present the already mentioned rover model that is equipped with eighteen infrared sensors and a controller, which is based on a single layer neural network. Because it was necessary to evolve a robot that can deal with both rocks and holes, we provided the robot with an evolvable threshold. This threshold adaptively modifies the activation range of the infrared sensors, in order to use front sensors for both rocks and holes detection. 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 behaviour and the threshold itself, which suggests that both behaviour 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 measure the robustness and adaptability to different environmental circumstances. Finally, in the conclusion we will discuss the results of the experiments and their relevance for space exploration research in the future.

## **METHOD**

As we have mentioned in the introduction, our approach is based on evolutionary robotics (ER). The ER approach emphasizes agent’s embodiment, which means that an emerging behaviour 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. The behaviour is seen as an emergent result of the dynamical interaction between the control system, the body, and the external environment [10] and relies on the fact that, while moving, motor actions partially determine the sensory pattern that a robot receive from the environment. Thus, by coordinating sensory and motor processes it is possible to create control systems which are able to select favourable sensory patterns and, in turn, enhance their ability to achieve their goals [4]. ER is an excellent technique that allows us to create artificial control systems that autonomously develop their skill in close interaction with the environment and that exploit very simple, but extremely powerful sensory-motor coordination [11]. ER is mainly based upon two computational techniques: artificial neural networks and genetic algorithms. Artificial neural networks (ANNs) are very powerful brain-inspired computational models, which have been used in many different areas such as engineering, medicine, finance, and many others [12]. ANNs are constituted by a certain number of simple computational units, the neurons, massively interconnected through a series of connections, the synaptic weights. Synaptic weights can be associated to variable numerical values that can be modified in order to allow the ANN to show a specific behaviour. In ER the synaptic weights are usually modified through an automatic evolutionary process which is inspired to the Darwinian principles that govern the natural process of evolution. This process, called genetic algorithm [13], is based on a simple biological model of evolution where the survival of the fittest and a constant production of new offspring result in adaptation to changing environments and ability to respond to unexpected events. It usually works with a population of artificial chromosomes that are evaluated for their performance and best of these are selected for further reproduction. The optimal solution is obtained after a series of generations in which chromosomes are evaluated and selected on the basis of their adaptability (i.e. the fitness).

### **The Rover model**

The robot used in this experiment is a 3D physical 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 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 case minor as we want to demonstrate that it is possible to use ER approach and a simple sensory setup to develop a suitable controller abler to handle complex obstacle avoidance tasks.

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 join. This means that left and right part of the rocker-bogie system can move independently while keeping the main body levelled.

The rover is equipped with a sensory apparatus that comprehends eighteen infrared sensors in order to provide sufficient information from the surrounding environment. In order to accommodate detection of various obstacles, two different set of sensors were used (see Fig. 1b). The first set consists of six lateral sensors which provide extra safety when it 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 twelve infrared sensors with the maximum reach of five and half meters. These infrared sensors, that we call ground sensors, are positioned on the rover's camera and are pointing downwards in 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 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 the position of the wheels.



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

### Control 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 has no hidden layer as we have found out that same results can be achieved with simpler architecture, that greatly reduce the computation demand of the control systems.

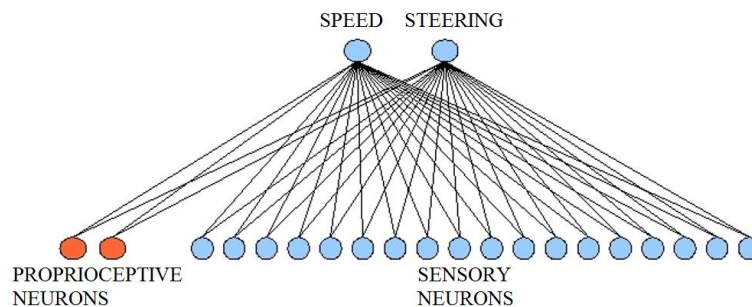


Fig. 2. Feed-forward neural network used as a control systems for the rover in the evolutionary experiments

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, i.e. genes, 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. The parameters are encoded as floating point values in the range  $[-1, 1]$  and the threshold in the range  $[0, 1]$ . In this way the ANN's weights can be simply represented as these genes and let GA to develop their strengths.

In our experiments we used a population size of 100 individuals, where the best 20 individuals were allowed to produce 5 offspring each. In practice, after a phase in which every 100 randomly generated ANNs were tested (i.e. were deployed in the rover and their performance were measured), a process of reproduction acted on the 20 best individuals where their genes were randomly mutated with a probability of 10% (a mutation occurs by adding to the original gene's value a quantity in the range  $[-1, 1]$ ). The reproduction and mutation processes were repeated 5 times for each of the best individuals, by generating 5 mutated copies of each of them. 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. On each generation, each control systems has been tested 10 times, by deploying it in the rover 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. To assure a good level of robustness of the evolved controllers, 15 evolutionary runs were conducted. 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 behaviour of the robot for effective and reliable exploration and obstacle avoidance behaviours:

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

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  $30^\circ$  from the centre. If for example the angle was  $15^\circ$  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  $30^\circ$  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.

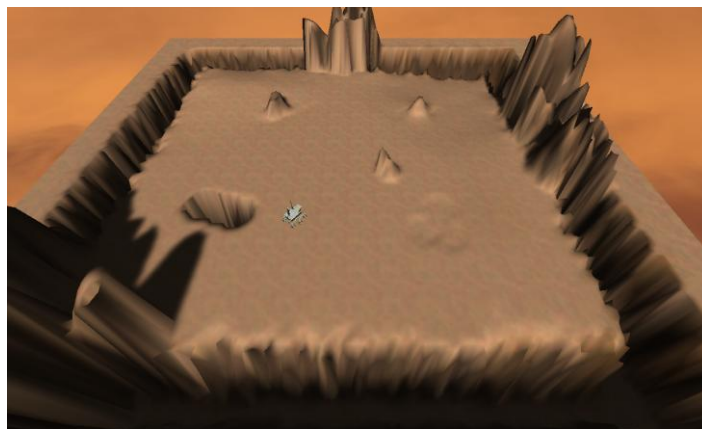


Fig. 3. Environment that was used during all evolutionary runs

In order to evolve a good controller, it was necessary to create a suitable environment (see Fig. 3.) and to allow the rover to interact with it. The environment that we modelled for this purpose is an arena of 60x60 m surrounded by holes and walls and containing obstacles and holes.

## RESULTS

The results obtained from all the fifteen evolutionary experiments show that an effective behaviour emerged in all evolutionary runs. In particular, thanks to the general behaviour optimised 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 results from all evolutionary runs. The graph was created by averaging values from all the fifteen runs. The blue line shows the maximum fitness obtained by the best individuals, the red line the average fitness of all the populations and the green 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 50<sup>th</sup> generation. With this optimized threshold the rover can detect all the rocks present in the terrain while not being confused by its roughness.

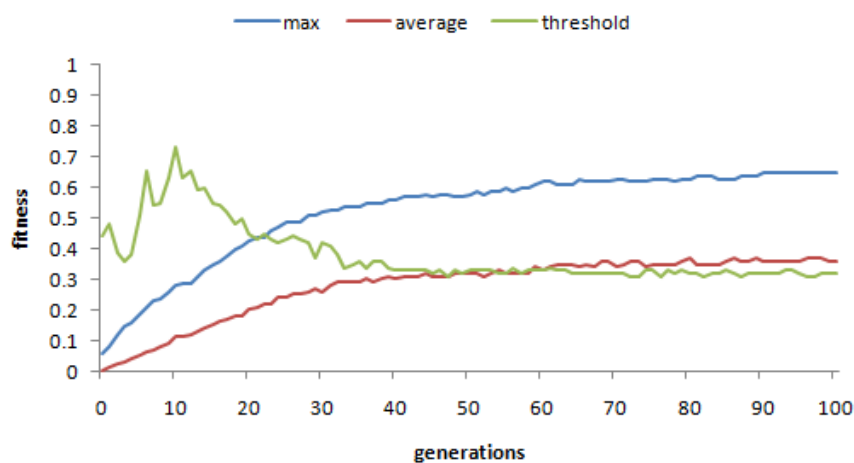


Fig. 4. Fitness graph showing maximum and average fitness as well as the threshold. Note that the fitness can never reach 1.0 as the rover needs to turn and decrease its speed to avoid obstacles

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 behaviour can only emerge if a suitable threshold value is found. Another interesting finding was that even a 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 certain experiments, several tests were conducted. In particular, tests were designed to evaluate the system 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 fifteen 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, and 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 was clear that the fitness will decrease if the rover is tested in such environment where 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 was explored within a given time. 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). This approach helps 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 make it less likely that all parts of the terrain are explored within 10,000 sensory-motor cycles.

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

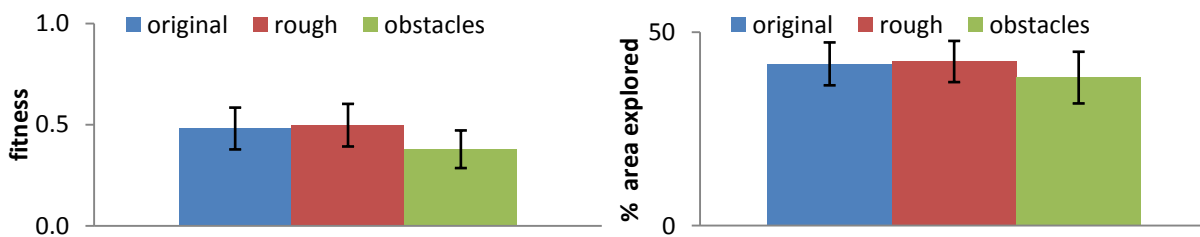


Fig. 5. Graphs showing average fitness (left) and exploration(right) of all evolutionary runs

## 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 an evolvable threshold. Our tests indicate that the system is very robust and able to maintain the obstacle avoidance behaviour under different circumstances and in different environments.

It is worth to note that the exploration and the obstacle avoidance behaviours 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 behaviour of the robot. Those behaviours 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 behaviours. 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 Animals 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*, 2<sup>nd</sup> 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.