Evolving the Neural Controller for a Robotic Arm Able to Grasp Objects on the Basis of Tactile Sensors

Raffaele Bianco, Stefano Nolfi
Institute of Cognitive Science and Technologies
National Research Council (CNR)
Viale Marx, 15, 00137, Roma, Italy
E-mail: rbianco@ip.rm.cnr.it, nolfi@ip.rm.cnr.it

Short Title: Evolving a Robotic Arm-hand Controller

Corresponding author: Stefano Nolfi

Abstract

We describe the results of a set of evolutionary experiments in which a simulated robotic arm provided with a two-fingered hand has to reach and grasp objects with different shapes and orientations on the basis of simple tactile information. Obtained results are rather encouraging and demonstrate that the problem of grasping objects with characteristics that vary within a certain range can be solved by producing rather simple forms of behaviour. These forms of behaviour exploit emergent characteristics of the interaction between the body of the robot, its control system and the environment. In particular we show that evolved individuals do not try to keep the environment stable but on the contrary push and pull the objects thus producing a dynamic in the environment and exploit the interaction between the body of the robot and the dynamic environment to master different environmental conditions with similar control strategies.

Keywords: Evolutionary Robotics, Robotic Arm, Sensory-Motor Coordination.

1. Introduction

The problem of controlling a robotic arm is usually approached by assuming that the robot should have or should acquire through learning an internal model able to: (a) predict how the arm will move and the sensations that will arise, given a specific motor command (direct mapping), and (b) transform a desired sensory consequence into the motor command that would achieve it (inverse mapping), for a review see [7]. We do not deny that humans and other natural species rely on internal models of this form to control their motor behaviour. However, we do not believe that motor control and arm movements in particular are based on a detailed description of the sensory-motor effects of any given motor command and of a detailed specification of the desired sensory states.

Assuming that natural organisms act on the basis of a detailed direct and inverse mapping is implausible for at least two reasons. The first reason is that sensors provide only incomplete and noisy information about the external environment and, moreover, muscles have uncertain effects. The former aspect makes the task of producing a detailed direct mapping impossible, given that this would require a detailed description of the actual state of the environment. The latter aspect makes the task of producing an accurate inverse mapping impossible given that the sensory-motor effects of actions cannot be fully predicted. The second reason is that the environment might have its own dynamic and typically this dynamic can be predicted only to a certain extent. For these reasons, the role of the internal models is probably limited to the specification of macro-actions or simple behaviours, rather than to micro-actions that indicate the state of the actuators and the predicted sensory state in any given instant.

This leaves open the question of how macro-actions or simple behaviors might be turned into micro-actions. One possible solution to this problem is to imagine that macro-actions (i.e. basic motor behaviors such as grasping a certain class of objects in a certain class of environmental conditions) are produced through simple control mechanisms that exploit the emergent result of fine grained interactions between the control system of the organism, its body and the environment.

To investigate this issue we ran a set of experiments in which we evolved the control system of a simulated robotic arm provided with a two-fingered hand that has to reach and grasp objects with different shapes and orientations on the basis of simple tactile information. As we will see, evolving individuals develop an ability to grasp objects in different environmental conditions without relying on direct and inverse mappings.

In section 2 we describe the related work, in section 3 and 4 we present our experimental setup and the results obtained by evolving the robotic arm in different experimental conditions. In section 5 we discuss the obtained results and their implications.

2. Related Work

As far as we know, there have been only two attempts to apply evolutionary robotics techniques [1, 3] to the synthesis of robotic arms. The first attempt has been done by Moriarty and Mikkulainen [2] who evolved the neural controller of a robotic arm with three degrees of freedom in simulation. The arm is initially placed in a random position and is asked to reach a random target position by avoiding obstacles. At any time step the neural controller receives as input the relative distance between the hand and the target position with respect to the three geometrical axis (x,y,z) and the state of 6 directional proximity sensors located in the hand. The robotic arm and the environment are simulated in a rather simplified way (e.g. collisions between objects are not simulated -- the authors simply stop the arm when its end point moves into a position occupied by an obstacle). The second attempt has been done by Skopelitis [6] who evolved the control system of a robotic arm with three degrees of freedom that is asked to follow a moving target. At any time step the neural controller receives as input the (x,y,z) coordinate of the target, the coordinate of the hand, the Euclidean distance between the hand and the target, and the coordinates of the "elbow" joint. Also in this case, the robotic arm and the environment are simulated in a rather simplified way (e.g. the target is not a physical object but only an abstract point of the environment and the arm has no mass and is not subjected to physical forces or collisions).

In the experiments reported in this article we tried to evolve the neural controller for a much more complex robotic arm provided with a two-fingered hand that is asked to grasp objects. The arm and the environment are carefully simulated (see below) and the controller is only provided with simple touch and proprioceptive sensors (i.e. it does not have access to information that cannot be computed by local sensors such us the distance with respect to the target position).

3. The Robotic Arm and the Neural Controller

3.1 The robotic arm

The robot consists of an arm with six degrees of freedom (DOF) and a two-fingered hand provided with three DOF (Figure 1, left). The arm consists of three connected basic structures forming two segments and a wrist. Each basic structure consists of two bodies connected by two motorized joints (Figure 1, centre). More precisely, each basic structure (Figure 1, bottom-right) consists of a parallelepiped with a size of [x=25, y=15, z=25] cm and a weight of 4.6 Kg and a cylindrical object with a radius of 12.5 cm, a length of 50 cm, and a weight of 12.3 kg (a length of 10 cm and a weight of 2.5 Kg in the case of the last cylindrical object that forms the base of the hand). Parallelepipeds are connected to the previous segment or to a fixed point (in the case of the first segment) by means of a rotational joint (R_Joint) that provides one DOF on axis Y. Cylinders are connected with parallelepipeds by means of an elevation joint (E_Joint) that allows only one DOF on axis Z. In practice, the E_Joint controls the elevation/lowering of the next connected segments and the R_Joint controls the rotation in both directions of the next connected segments. The E_Joint is free to move only between 0 and π/2, just like an human arm that can bend the elbow solely in one direction. The range of R_Joint is [-π/2, +π/2] for the first two and is [0, π] for the last basic structure.

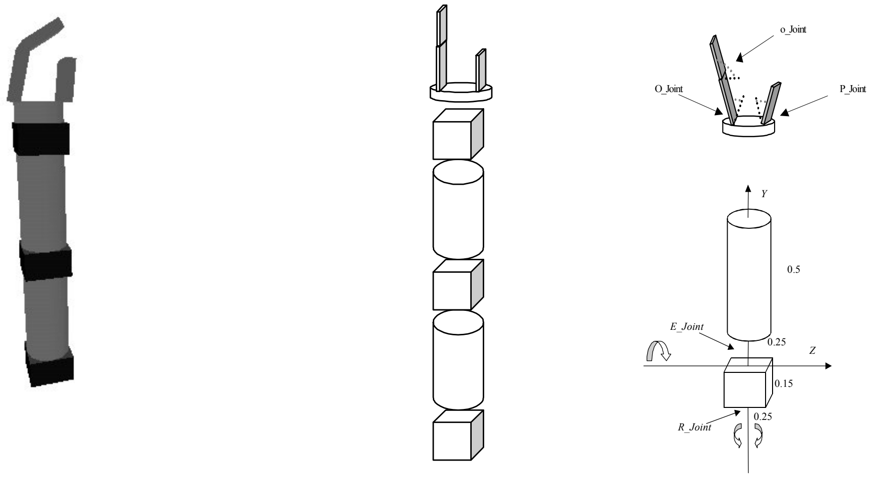

Figure 1. Left: The arm and the hand. **Centre**: A schematic description of the elements forming the arm and the hand. **Right**: A schematic description of the motorized joints that connect the different elements of the arm and of the hand. The exact orientation of the arm along the three axis is shown in Figure 3.

The hand consists of two fingers made of two parallelepipeds with a size of [x=5, y=20, z=20] cm and a weight of 1 Kg connected by two motorized joints (*O_Joint* and *P_Joint*) to the last cylindrical object forming the arm. These two joints, that allow the fingers to open and close, can move only in a range of [-π/10, π/6] and [-π/10, π/4] respectively (Figure 1, right up). The first finger has an additional phalange consisting of a parallelepiped with a size of [x=5, y=20, z=25] cm and a weight of 1.25 Kg connected by a motorized joint (*o_Joint*) to the previous part of the finger. This additional joint that allows the finger to close its upper part can move in the range of [0, π/2].

Each actuator is provided with a corresponding motor that can apply a maximum torque of 10 Nm. Friction coefficient is set to 0.7 and the acceleration of gravity is –0.098 m/ds^2.

This means that to reach and grasp an object the robot has to appropriately control 9 joints and to deal with the constraints due to gravity and collisions.

The sensory system consists of six contact sensors (three placed on the three cylindrical objects forming the arm and the wrist and three placed on the three parallelepipeds forming the two fingers) that detect, in a binary fashion, whether these bodies collide with other bodies. Moreover, robots have nine proprioceptive sensors that encode the current angular position of the nine corresponding motor joints controlling the arm and the fingers.

The environment consists of a planar surface (at height 0) and an object (e.g. a ball, a cube or a bar) placed on the surface (see Figure 3). The first element of the arm is anchored to a fixed point [x=0, y=115, z=0] cm and is oriented along the vector [0,-0.86,-0.5].

To reduce the time necessary to test individual behaviours and to model the real physical dynamics as accurately as possible, we used the rigid body dynamics simulation SDK of Vortex^{TM} (see http://www.cm-labs.com/products/vortex/). This software allowed us to build a rather realistic simulation and to speed up the evolutionary process by allowing simulated robots to move faster than real physical robots.

3.2 The neural controller

Each individual is controlled by a fully connected neural network with 15 sensory neurons and 9 motor neurons. Neurons are updated with the logistic function.

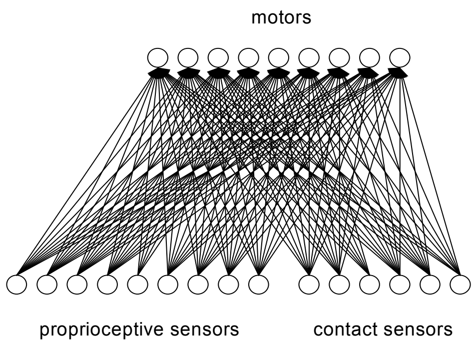

Figure 2. The architecture of the neural controller.

The sensory neurons encode the angular position (normalized between 0.0 and 1.0) of the 9 DOF of the joints and the state of the six contact sensors located in the arm and in the fingers. The motor neurons control the actuators of the 9 corresponding joints. The output of the neurons is normalized within the range of movement of the corresponding joint and is used to encode the desired position of the corresponding joint. More precisely, motors are activated so as to reach a speed proportional to the difference between the current and the desired position of the joint (maximum motor speed is 500 deg/sec).

The genotype of evolving individuals encodes the connections' weights and the biases of the neural controller. Each parameter is encoded with 8 bits. Weights and biases are normalized between –10.0 and 10.0. Population size is 100.

The 20 best individuals of each generation were allowed to reproduce by generating 5 copies of their genotype which were mutated by replacing 2% of randomly selected bits with a new randomly chosen value. Each experiment was replicated 10 times.

Each individual of the population was tested for a given number of trials, with each trial consisting of a given number of steps (each step lasts 5 ms of real time). At the beginning of each trial the arm and the hand is set in the initial position and an object is placed on top of a planar surface. The starting angles of the three E_Joints are set to 0, 70 and 30 degrees respectively, those of R_Joints are set to 0, -30 and 20 degrees, those of the O_Joints, o_Joint and P_Joint are set to 0 degrees (see Figure 3). During each time step: (1) the state of the sensory neurons is updated, (2) the activation state of internal (if present) and motor neurons is determined, (3) the desired speed of the motors controlling the joints is set according to the actual state of the motor neurons, and (4) the dynamic of the arm and of the environment is updated on the basis of the length of the time step and of the current forces and collisions.

To verify in a simple way if the object has been correctly grasped or not, in every trial we removed the planar surface after a certain number of time steps and after another time interval we checked whether the object fell down (i.e. if the arm did not grasp it) or not (i.e. if the arm did grasp it). Please notice that we cannot remove the planar surface from the beginning because otherwise the object would fall down before the arm can reach it.

Evolving individuals were scored on the basis of the number of objects they were able to grasp during a fixed number of trials. In addition, to facilitate the emergence of an ability to grasp objects, we also rewarded individuals for their ability to touch the object with their fingers (i.e. we used a form of incremental evolution [3]). More precisely the fitness of an individual was computed according to the following equation:

$$\text{fitness} = (GP * 10000) + NC$$

where *GP* is the number of objects that have been successfully grasped (i.e. that have an y coordinate higher or equal to 25 cm and that collided with at least one of the finger at the end of the corresponding trial), and *NC* is the number of collision between the objects and the fingers of the hand during the whole lifetime of an individual.

4. Experimental results

4.1 Grasping Cubic or Spherical Objects with Different Size, Weight and Orientation

In a first set of experiments we asked evolving individuals to grasp cubic or spherical objects with different size, weight and orientation. Each individual was tested for 20 trials, each trial consisting of 800 steps (the plane is removed after 700 steps). During its life each individual experienced 10 spherical and 10 cubic objects that were located in the following position [x=100, z=0] cm (the y coordinate was set proportionally to the size so as to assure that the object lays on the plane). The size, weight and orientation of the objects was randomly chosen in each trial from a given range. The side of cubic objects varied between 15 and 20 cm and the radius of spherical objects varied between 10 and 15 cm. The density of the objects varied between 100 and 500 kg/m^3. The orientation varied between 0 and 90 degrees along the Y axis.

By running 10 evolutionary experiments for 50 generations we observed (see Figure 4) that evolving individuals display rather good performance (up to 100% of successful trials in the case of the best replication and up to 83.6% of successful trials on average).

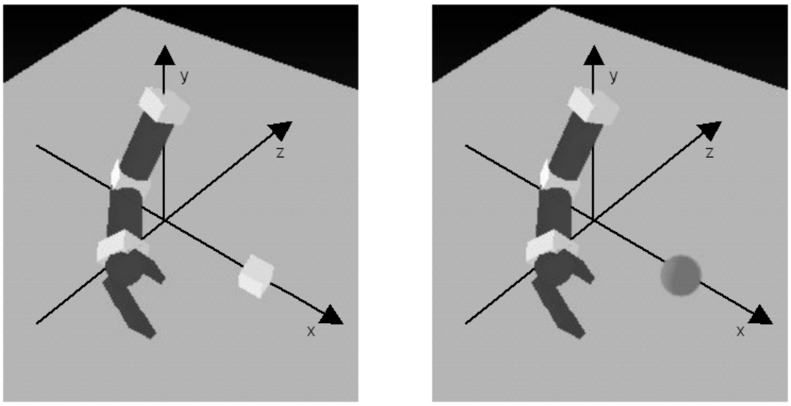

Figure 3. The arm and the hand in their initial position and the environment. The environment consists of a cubic or a spherical object laying on a planar surface.

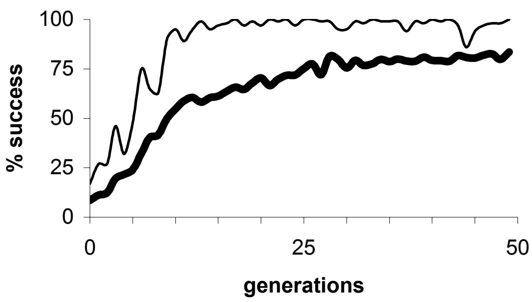

Figure 4. Fitness obtained by testing the best individuals of each generation for 100 trials. **Thin line**: performance in the case of the best replication. **Thick line**: average performance of 10 replications of the experiment.

The analysis of the behaviour displayed by evolved individuals shows that evolved robots are able to reach and grasp different types of objects with different sizes and weights by mastering the dynamical interaction between the objects and the hand. Indeed, as shown in Figure 5, the robot approaches the object from its top-left side with its hand open and starts to close the hand as soon as it detects the object with the contact sensors of the fingers. However, while grasping the object, the robot also rotates the hand on the right side. This movement from left to right allows the robot to block the movement of the object (caused by the collision between the hand and the object) with the hand itself and allows the robot to exploit the properties emerging from the dynamical interaction between the moving object and the hand. Indeed, as a result of the initial collision between the hand and the object and of the successive rotation of the hand, objects tend to move toward the inner part of the hand and spontaneously adjust small misplacements resulting from the fact that objects have different sizes and orientations.

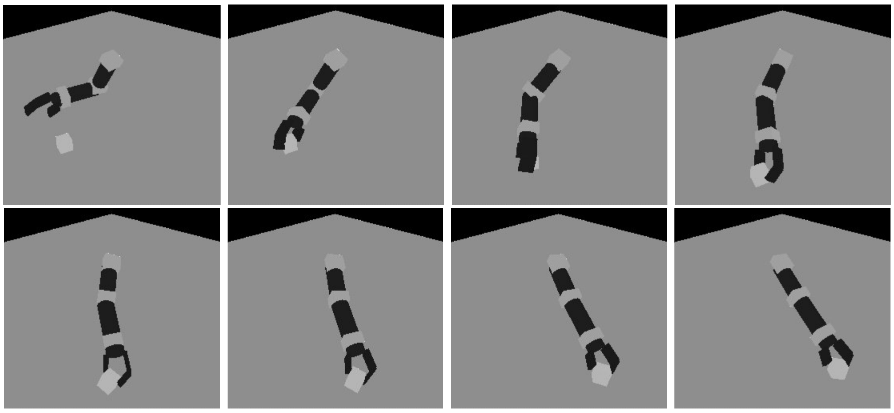

Figure 5. A typical behaviour displayed by an evolved individual. The eight pictures (from left to right and from top to bottom) show eight snapshots of a trial in which the robot approaches and grasps a cubic object.

4.2 Grasping Bars with Different Weights and Orientations

In a second set of experiments we asked evolving individuals to grasp bars with different weights and orientations (Figure 6). Given that bars can only be grasped by placing the hand in the right relative orientation with respect to the object, we might expect that to solve this problem robots should first detect the orientation of the bar and then approach the object appropriately. As in the case of the previous experiment, however, by exploiting the interaction between the hand and the object evolving individuals develop a simpler solution that consists in modifying the orientation of the bar.

The size of bars is [x=15, y=15, z=50] cm. Each individual was tested for 30 trials, with each trial consisting of 1000 steps (the plane is removed after 900 steps). In each trial the barycentre of the bar is initially placed in the position [x=100, y=7.5, z=0], the orientation is randomly chosen between 0 and 180 degrees along the Y axis, and the weight is randomly chosen between 100 and 200 kg/m^3 (see Figure 6). The evolutionary process was continued for 100 generations. All other parameters are identical to those of the experiment described in section 4.1.

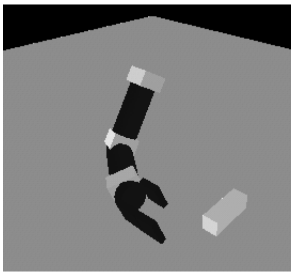

Figure 6. The arm and the bar at the beginning of a trial. The orientation and the weight of the bar is randomly chosen in every trial within a given range (see text).

Also in this experiment evolving individuals display a rather good performance (up to 97% of successful trials in the case of the best replication and up to 76.4% of successful trials on the average) (see Figure 7).

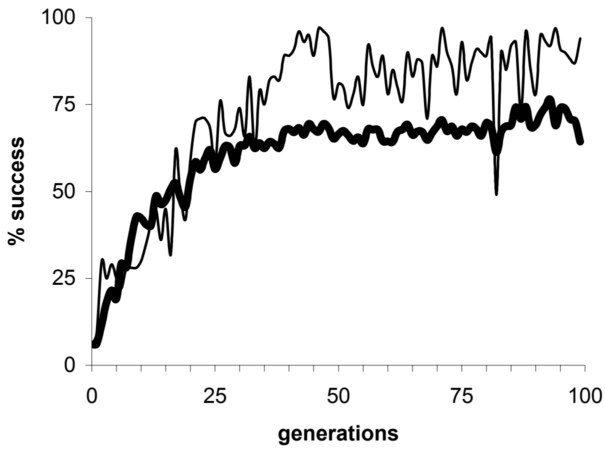

Figure 7. Fitness obtained by testing the best individuals of each generation for 100 trials. **Thin line**: performance in the case of the best replication. **Thick line**: average performance of 10 replications.

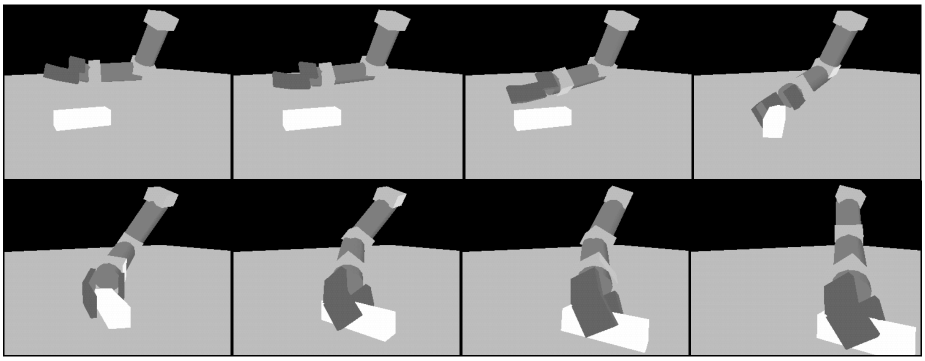

Figure 8. A typical behaviour displayed by an evolved individual. The eight pictures (from left to right and from top to bottom) show eight snapshots of a trial in which the robot approaches and grasps a bar with a randomly selected orientation.

The analysis of the behaviour displayed by evolved individuals shows that most of the times, evolved robots are able to reach and grasp bars independently from their relative orientation and weight. As shown in Figure 8, the robot approaches the object from the left side and, while grasping the object, also rotates the bar toward the preferred orientation (in the case of this individual, the orientation that the bar has in the bottom-right picture). In other words, evolved robots do not need to detect the current orientation of the bar and then approach the object from different orientations according to the actual position of the bar. They solve the problem with a rather simple behaviour by exploiting the dynamical interaction between the hand and the environment. In particular, the bar-rotation behaviour emerges from the simple approaching behaviour produced by the robot, the different length of the two fingers, and the effect of the collisions between the hand and the bar, produced by the movements of the arm and of the fingers.

4.3 Grasping Cubic Objects with Different Weights, Orientations and Positions

In the third set of experiments we asked evolving individuals to grasp cubes with varying weights, orientations and positions. The goal of this third set of experiments was to verify the possibility to evolve robots that apart from being able to grasp objects in varying conditions, are also able to find objects by exploring the environment.

Each individual was tested for 20 trials, with each trial consisting of 800 steps (the plane is removed after 700 steps). During its lifetime each individual dealt with cubic objects with a size of 20 cm that were randomly located in a rectangular area with an upper left corner in [x=125, y=10, z=-50] and a lower right corner in [x=75, y=10, z=50]. The weights and orientations of the objects were randomly chosen, within a given range, in each trial. The density of the objects varied between 100 and 500 kg/m^3. The orientations varied between 0 and 90 degrees along the Y axis.

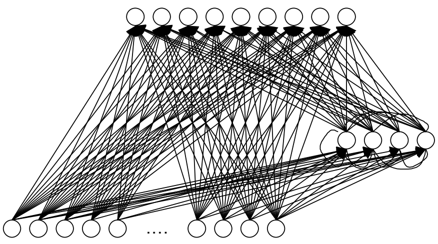

Figure 9. The architecture of the neural controller.

By running several set of experiments in which we varied the neural architecture and the sensory structure of the hand (results not shown), we observed that to solve this problem evolving robots should have: (1) more resolution on contact sensors (i.e. the possibility to detect more precisely the location of the hand that gets in contact with external bodies or other parts of the robot itself), and (2) additional internal neurons (see Figure 9). In particular close to optimal performances (i.e. up to 89% of objects grasped correctly, see Figure 10) were achieved: (1) by providing each segment of the fingers with two different contact sensors placed on the interior and the exterior part of each segment and by providing the two fingers with two additional contact sensors placed on the tip parts of the fingers, (2) by providing the neural controllers with four internal neurons receiving connections from the sensory neurons and from themselves and projecting connections to the motor neurons. Therefore, in these experiments the neural controllers have 20 sensory neurons (9 encoding the angular position of the joints and 11 encoding the state of the contact sensors), 9 motor neurons controlling the 9 corresponding motorized joints, and 4 internal neurons receiving connections from sensory neurons and from themselves and sending connections to motor neurons (see Figure 9).

The activation state of internal neurons was updated accordingly to the following equation (see [4] for more details):

$$A_j = t_j + \sum w_{ij} O_i \qquad O_j = \tau_j O_j^{(t-1)} + \left(1 - \tau_j\right)\left(1 + e^{-A_j}\right)^{-1} \qquad 0 \le \tau_j \le 1 \quad \textbf{(1)}$$

With *Aj* being the netinput of the *j*th neuron, *tj* the bias of the *j*th neuron, *Wij* the weight from the *i*th to the *j*th neuron, *Oi* the output of the *i*th neuron. *Oj* is the output of the *j*th neuron, *τj* the time constant of the *j*th neuron. The time constants of neurons were genetically encoded.

The evolutionary process was continued for 150 generations. All other parameters are as those of the experiments reported in the previous sections.

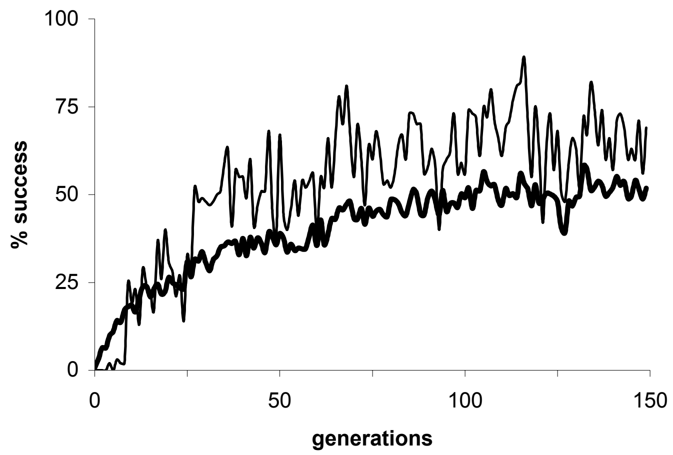

Figure 10. Fitness obtained by testing the best individuals of the best replication of each generation for 100 trials. **Thin line**: performance in the case of the best replication. **Thick line**: average performance of 10 replications.

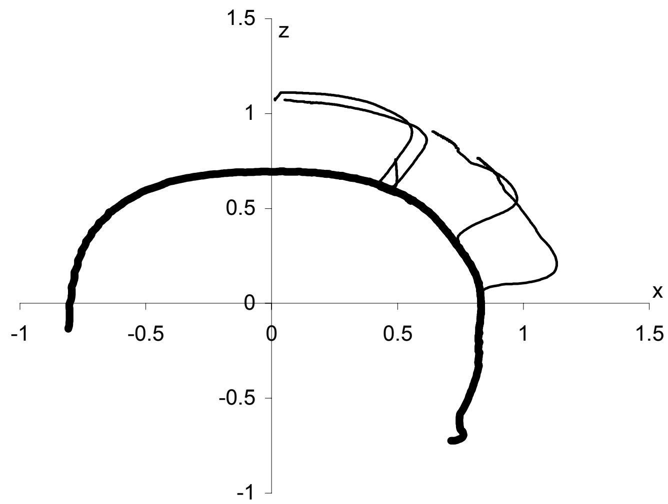

Figure 11. The trajectory on the x and z plane of the hand during different trials. **Thick line**: trajectory in an environment without objects. **Thin lines**: trajectory with the cubic object placed in randomly varying positions and orientations.

The need of a higher resolution in contact sensors can be explained by considering that the variation in the position of the object causes a significant variation in the relative position between the hand and the object when they first come into contact. This, in turn, causes the need to tune the grasping behaviour on the basis of the relative position of the hand and of the object. This position can be properly detected only by having more contact sensors (i.e. more detailed information on the positions in which contacts occur). Indeed, as shown in Figure 11, evolved individuals display an ability to tune their approaching and grasping behaviour for different positions of the object. As shown in the figure, the trajectory of the hand significantly varies in different trials as soon as the hand first detects the object with its touch sensors according to the actual position and orientation of the object itself.

The need of internal neurons with recurrent connections can be explained by considering that in this experiment, the robot should significantly modify its behaviour according to the current position of the object to be grasped (see Figure 11) by taking into account not only the current but also the previous sensory states.

5. Discussion

In this paper we present a set of evolutionary robotics experiments in which simulated robotic arms provided with a two-fingered hand develop the ability to reach and grasp objects with different locations, shapes and orientations on the basis of simple tactile information. Obtained results are rather encouraging and demonstrate that the problem of grasping objects with characteristics that vary within a certain range can be solved by producing rather simple behaviours. These behaviours exploit emergent characteristics of the interaction between the body of the robot, its control system, and the environment. In particular we showed that in all cases, evolved individuals do not try to keep the environment stable but, on the contrary, push and pull the objects thus producing a dynamic in the environment. Moreover, they exploit the interaction between their body and the dynamical environment to master rather different environmental conditions with rather similar control strategies.

The results of these experiments demonstrate that the evolutionary robotic approach can scale up to the development of robots with many degrees of freedom that are able to operate robustly in varying and dynamical environmental conditions. Indeed, the ability to exploit properties that emerge from the dynamical interaction between the control system of the robot, its body, and the external environment allows the evolutionary process to find solutions that are simple and robust (see also [4-5]).

In future work we plan to: (a) apply this method to the development of a control system for a real robotic arm, (b) develop controllers provided with different neural modules capable of displaying different classes of behaviours and capable of arbitrating between the different modules.

Acknowledgments

This research has been partially supported by MIUR (Italian Ministry of Education, University and Research) under project RoboCare (A Multi-Agent System with Intelligent Fixed and Mobile Robotic Components).

References

1. Cliff, D., Harvey, I., & Husbands, P. (1993) Explorations in evolutionary robotics. *Adaptive Behavior*, 2, 73-110
2. Moriarty, D.E., & Mikkulainen, R. (1996). Evolving obstacle avoidance behavior in a robot arm. In Maes P., Mataric M., Meyer J.-A., Pollack J., and Wilson S.W (eds.), *Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviors*. Cambridge MA: MIT Press.
3. Nolfi, S., & Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books
4. Nolfi, S., & Marocco, D. (2001). Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences,* 120, 287-310.
5. Nolfi, S. (2002). Power and Limits of Reactive Agents. *Neurocomputing*, 42, 119-145.
6. Skopelitis C. (2002). Control System for a Robotic Arm. *Master Thesis*. School of Cognitive and Computing Sciences (COGS), University of Sussex, U.K.
7. Torras C. (2002). Robot arm control. In M.A. Arbib (Ed.) *The Handbook of Brain Theory and Neural Networks*, Second edition. Cambridge, MA: The MIT Press.