Evolving Morphology and Control: A Distributed Approach

M. Mazzapioda, A. Cangelosi and S. Nolfi

Abstract—**In this paper we present a model which allows to co-evolve the morphology and the control system of realistically simulated robots (creatures). The method proposed is based on an artificial ontogenetic process in which the genotype does not specify directly the characteristics of the creatures but rather the growing rules that determine how an initial artificial embryo will develop on a fully formed individual. More specifically, the creatures are generated through a developmental process which occurs in time and space and which is realized through the progressive addition of both structural parts and regulatory substances which affect the successive course of the morphogenetic process. The creatures are provided with a distributed control system made up of several independent neural controllers embedded in the different body parts which only have access to local sensory information and which coordinate through the effects of physical actions mediated by the external environment through the emission/detection of signals which diffuse locally in space. The analysis of evolved creatures shows how they display effective morphology and control mechanisms which allow them to walk effectively and robustly both on regular and irregular terrains in all the replications of the experiment. Moreover, the obtained results show how the possibility to develop such skills can be improved by also selecting individuals on the basis of a task-independent component which reward them for the ability to coordinate the movements of their parts.**

I. Introduction

THE attempt to evolve complete artificial creatures (i.e. embodied and situated agents in which both morphological and control characteristics are adapted during the evolutionary process) has been and still represents a key long term goal for the Artificial Life and Evolutionary Robotics community.

After the pioneering work of Karl Sims [1-2], who demonstrated the feasibility of the idea, and the related model of Lipson and Pollack [3], who demonstrated how complete agents evolved in simulation could enter in the physical world through a semi-automatic manufacturing process, several other researchers attempted to developed models that could be more effective [4]-[8], [10]. The issue of whether these methods can be scaled up to tackle non–trivial problems and/or can be competitive with respect to alternative methods (for example with respect to methods in which parts of the characteristics of the robots are designed by the experimenter) still remains an open question.

The possibility to make progresses, in this respect, crucially depends from the possibility to identify a genotype-to-phenotype mapping with the following characteristics: evolvability, expressivity, and simplicity. By evolvability we mean that the probability that genetic variations lead to improvements of creatures' adaptive skills should not be too low. By expressivity we mean that variations at the level of the genotype should potentially lead to a large number of possible alternative solutions with respect to both the space of possible morphologies and the space of possible control systems. By simplicity, we mean that the rules that determine the relation between the genotype and the phenotype should be as simple as possible so to avoid the need, from the point of view of the experimenter, to deal with too many parameters to be chosen and optimized.

In this paper we describe a model which permits the co-evolution of the morphology and the control system of realistically simulated artificial creatures which are evolved for the ability to locomote on flat and irregular terrains. The method is characterized by the following features:

(i) An indirect encoding in which the genotype does not specify directly the characteristics of the phenotype but rather the way in which the embryo develop into a full formed individual in a realistic 3D space in which physical objects cannot overlap (as in most of the models referenced above).

(ii) The presence of regulatory processes which are realized through the synthesis in space of genetic products which later affect the morphogenetic process (as in [9]).

(iii) The adoption of a highly distributed approach in which robots are made up of several independent neural controllers embedded in the different body parts which only have access to local sensory information and which coordinate so to exhibit a coherent and effective behaviour. More precisely, the different body parts coordinate through the effects of physical actions mediated by the external environment through the emission/detection of signals which diffuse locally in space.

(iv) The adoption of a relatively simple formalism which

Manuscript received November 1, 2001. This work was supported by the ECAGENTS project funded by the Future and Emerging Technologies programme (IST-FET) of the European Community under EU R&D contract IST-1940.

M. Mazzapioda is with both the Institute of Cognitive Sciences and Technologies, National Research Council (ISTC-CNR), Rome, 00185 Italy (shared phone: +39 06 44595255; shared fax: +39 06 44595243; email: m.mazzapioda@istc.cnr.it) and School of Computing, Communications and Electronics, University of Plymouth (UK).

A. Cangelosi is with the School of Computing, Communications and Electronics, University of Plymouth, PL4 8AA Plymouth (UK) (e-mail: acangelosi@plymouth.ac.uk).

S. Nolfi is with the Institute of Cognitive Sciences and Technologies, National Research Council (ISTC-CNR), Rome, 00185 Italy (e-mail: stefano.nolfi@istc.cnr.it).

attempt to maximize the need to keep the model as simple as possible and the number of parameters as small as possible.

(v) The use of a selection criterion which also includes a task-independent component which favor the evolution of the required skill. More precisely the fitness function used include two components which reward the robots for the ability to accomplish the desired task (locomote at the highest possible speed, in the case of the experiments performed) and for the ability to coordinate the movements of their parts (for more details see below).

The model proposed is loosely inspired by evolutionary-ancient biological organisms which are not provided with a central nervous system and by simple organisms, such us stick insects [18] which display a high degree of decentralization. Our goal, however, is not that to model specific biological organisms, but rather to identify general mechanisms which can lead to the autonomous design of complete and evolvable artificial systems based on local distributed control. In particular, we aim to investigate whether creatures with high distributed organization formed by a collection of elements, which operate independently on the basis of local information, are evolvable, and if this evolvability results in creatures capable to perform requested task.

The analysis of the obtained results demonstrates how the method proposed allowed the synthesis of effective and robust behaviours after relatively short number of generations in all the replications of the experiments. Moreover, the analysis of the best evolved individuals demonstrates how the evolutionary process lead to the generation of a wide range of morphological structures and walking styles. The obtained results confirm the hypothesis that the use of an additional task-independent component, that reward the robots for coordinating the movement of their parts, facilitate the evolution of effective behaviours. Finally, the analysis of the characteristics of the robots throughout generations indicates that the major transitions during the evolutionary process typically involve significant variations of the robots morphology which enable the conditions for the synthesis and the retention of several adaptive variations during the immediate successive generations.

II. Model Description

In this section we describe the developmental process that determines how creatures develop from a single elementary unit into complete creatures provided with an articulated body, sensors, and actuators. In the following sub-sections we describe the genotype of creatures (initial embryo, growing rules and distributed control system) and the evolutionary process.

A. Genotype: Embryo

At the beginning of the developmental process, creatures consist of an initial 'embryo'.

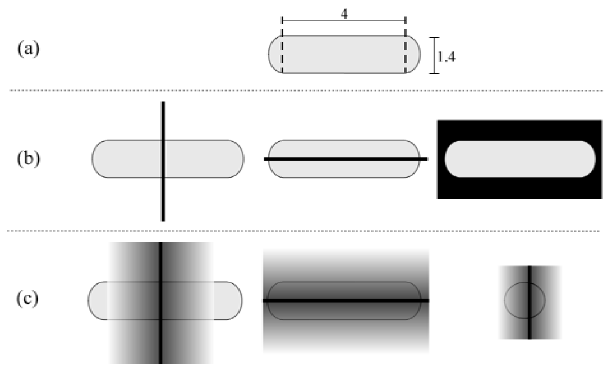

Figure 1: (a) an elementary unit; (b) the orthogonal planes representing the area of maximum concentration of the corresponding regulatory substances α, β, and γ; (c) concentration distribution of the three corresponding regulatory substances.

Each embryo (see Figure 1) is composed of an elementary unit (a cupped cylinder with a length of 5.4 cm and a diameter of 1.4 cm), and three regulatory substances, (α, β, and γ) distributed inside and outside the elementary unit.

The concentration of regulatory substance varies with respect to the distance from the plane of the maximum concentration. More precisely, the planes that determine the concentration of the three regulatory substances intersect the first elementary units and are oriented along the three main axis (Figure 1b) and the concentration of the three regulatory substances varies linearly along the left-right, dorsal-ventral and the rostral-caudal axis, respectively, with concentration of 1.0 on the corresponding plane which linearly decrease up to 0.0 for distances equal or greater than the maximum propagation range of the corresponding substance (Dα, Dβ, and Dγ).

The first elementary unit is identical for all individuals and is not subjected to the evolutionary process whereas all data regarding regulatory substances (plane position and rotation, and maximum propagation) are encoded inside the genotype.

B. Genotype: Growing Rules

Each genotype encodes 15 developmental rules and a vector of 15 indexes which indicates the order with which the developmental rules are executed.

All rules comply with the same pattern that is formed by a condition part and an action part which is executed in the location/s of the current embryo where the condition part holds. The condition part encodes the regulatory substances (within the following possibilities [α, β, γ, τ]) and the concentration of the substances (within the range [1.0, 0.0]) which leads to the execution of the action part. Since at each given stage of the developmental process, the condition part might be valid in any, one, o es: add-unit action and add-regulatory-substance action.

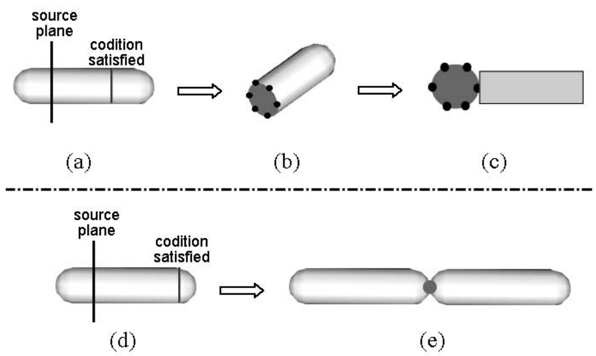

Figure 2: (a) An example of the plane in which the substance has the maximum distribution and of the circular surface in which the condition is satisfied. (b) An exemplification of the six points in which new elementary units can be attached. (c) An exemplification of a case in which the condition is satisfied at one of the ends of the elementary units. (e) An exemplification of the way in which the new element is attached in case (d).

The **add-unit action** adds from 0 to 6 elementary units in the location/s of the current embryo in which the condition part hold/s. A preliminary check has been done before adding new elements. In particular, if new element does not intersect any existing elements then it will be added, otherwise it will be discarded.

The parameters of this action include:

(1) Six binary values which encode whether or not the elementary units are grown along six possible perpendicular orientations (0°, 60°, 120°, 180°, 240°, 300°) with respect the longitudinal axis of the circular surface in which the condition holds (Figure 2a, b, c). When the condition is satisfied at one of the end points of the cylinder, the new elementary unit grows along the same axis of the elementary unit in which the condition is satisfied (Figure 2d, e).

(2) An integer value which determines whether the new elementary units are connected through a fixed joint or a motorized joint with one degree of freedom (DOF) within the following possible cases (x, y, or z) and within the following limits [-30°: +30°]. The maximum force and the maximum velocity applied to the motorized joint correspond to 404 Nm and 9,6 rad/sec respectively.

(3) A vector of values which encodes the connection weights and biases of the artificial neural network (see below) which control the corresponding DOF (if any). Each parameter is encoded with 8 bits and normalized in a floating point value in the range [-15.0, +15.0]. This implies that the neural network controllers which are generated through the same rule have the same characteristics whereas those that are generated by different rules might differ.

The **add-regulatory-substance action** adds a new regulatory substance with a concentration which varies proportionally to the distance with respect to a new plane. This new plane is created in the point in which the condition is satisfied with the same orientation of the regulatory substance which triggered the execution of the action. (Figure 3).

The parameters of the add-regulatory-substance actions include:

(1) An integer value in the range [0, 3] which encodes the type of regulatory substance added.

(2) A floating point value in the range [3.0, 8.0] that encodes the size of the plane which determines the diffusion of the regulatory substance.

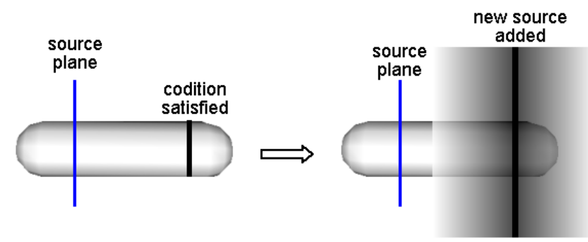

Figure 3: New source substance is added

C. *Genotype: Neural Controller*

Each elementary unit provided with a motorized joint includes a neural network controller with a fixed architecture consisting of five sensory neurons directly connected to five motor neurons (Figure 4).

The neural controller has access to the current angular position of the corresponding joint and regulates the frequency of oscillation of the joint. Neural modules are also allowed to communicate between themselves by producing up to four different signals and by detecting the signals produced by other neural controllers located within a maximum Euclidean distance.

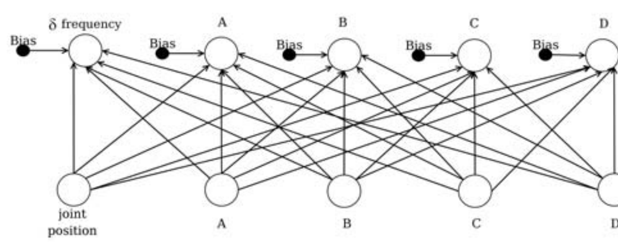

Figure 4: The architecture of neural controller

The first sensory neuron encodes the current angular position of the corresponding motorized joint (normalized in the range [-1.0, 1.0]). The other four sensory neurons encode how many signals produced by other neural modules are detected. Each neural module can produce four different signals (A, B, C, and D) that diffuse and can be detected up to a certain distance (D_A, D_B, D_C, and D_D in the case of signals A, B, C, and D respectively). Detection is a binary value, and the total sum of detected signals is normalized in the range [0.0, 1.0].

The desired position of each joint is determined by a sinusoidal oscillator with a frequency that is initially assigned randomly in the range [7.0, 14.0]Hz and that is later increased or decrease in each time step within the same range on the basis of the current output of the corresponding neural controller. More precisely the frequency increases in the range [0, 1.4]Hz for outputs in the range (0.5, 1.0], and drecreases in the range [0, -1.4]Hz, for outputs in the range (-0.5, -1.0]. The other four output neurons are threshold units which determine whether the signal A, B, C. and D are produced (value is 1) or not (value is 0). For more details on this type of neural controllers and for an analysis of how signals can lead to coordination in a distributed system used to control a robot with a fixed morphology see [13][14].

D. The Evolutionary Process

The initial population consists of 100 genotypes. Each genotype includes a set of genes which encode the initial substances in the embryo, the developmental rules and the vector which determine the order with which rules are executed. Parameters are randomly generated within the corresponding intervals.

The 20 best genotypes of each generation are allowed to reproduce by generating five copies each (elitism is applied in firsts five genotype only). During reproduction each gene is mutated (i.e. replaced with a new randomly selected value in the corresponding range) with a probability of 3%. Moreover, each element of the vector which encodes the indexes and the order with which developmental rules are expressed can be moved at the end of the vector with a probability of 2%. In order to increase the level of variability, the 10 worst-fitting individuals are replaced by new randomly generated ones every 50 generations To introduce more variability every 50 generations the 10 individuals with the worst fitness are replaced with new randomly generated ones (preliminary analysis indicate that this increase variability during the first generations only, since the probability that these randomly generated individuals will be selected becomes very low in successive generations).

The evolutionary process lasts for 500 generations (i.e. the process of testing, selecting and reproducing robots is iterated 500 times). Each individual genotype is allowed to develop in a free space into the corresponding phenotype and is tested in the environment for 5 trials (2 trials on a flat terrain and 3 trials on an irregular terrain).

At the beginning of each trial the current phase and the current frequency of oscillation of each joint is set randomly within the corresponding range and the creature is placed at a height of approximately 5.5 cm from the ground (i.e. this implies that creatures should be able to stand on the posture/s in which they are able to locomote). Creatures are then allowed to move for 6000 time steps lasting 1.5ms each.

For each time step the state of the sensors and of the motors of each neural controller, the torque exerted by the motorized joints, and the dynamics of the creature/environmental interaction are updated. Creature and creature/environmental interactions are simulated by using the ODE dynamical simulation engine [19].

The fitness formula includes two components which score individuals for the ability to move as fast as possible and for the ability to produce coordinated movements. The first component is calculated by measuring the Euclidean distance travelled by an individual during its lifetime from time step 2000 on (i.e. the distance travelled during the first phase in which the creatures fall down and start coordinating does not effect their fitness). The second component consists of the average mutual information [20] calculated between the current frequency of oscillation of each couple of joints.

The total fitness is computed by normalizing, in the range [0.0, 1.5] the value of the two components with respect to an estimation of the maximum distance which can be travelled by a creature and with respect to the theoretical maximum of the mutual information, respectively.

III. Results

In the next three sections we describe; (A) the results obtained by running 10 replications of the experiment based on the model described above, (B) the results obtained in other control experiments in which we analyzed the effect obtained by varying some of the characteristics of the model proposed, (C) the results obtained by analyzing the course of the evolutionary process.

A. Results

From a qualitative and quantitative analysis of the performance of the best individual of the last generation, in each replication we observed that all individuals evolve an ability to locomote effectively (Table 1). The average distance covered by individuals varies for different replications and achieves appreciable results in the case of the best replications. Performance also differs with respect to type of terrain (Table 1). In some replications individuals display an ability to move effectively on both types of environments.

By analyzing the morphology of evolved creature (Figure 5) and the number of DOFs (Table 1) we can see how the evolutionary process leads to a large variety of morphological structures. This indicates that the model chosen has a good level of expressiveness. As can be seen from the figure, the evolved morphologies show a high degree of symmetry. This characteristic can be explained by considering that in the model proposed the developmental process occurs in the 3D space and is governed by regulatory substances which have a distribution which is symmetrical with respect to the plane of maximum concentration.

We can also observe how all evolved creatures are provided with morphologies which allow them to avoid falling on a side or tipping over, and which allow them to master obstacles of various sizes in irregular terrains. By visually inspecting the behavior of evolved creatures (see the movies available from http://laral.istc.cnr.it/esm/dros1-0) we can identify three phases: (a) an initial phase in which the creature falls on the terrain by assuming a certain posture, (b) an intermediate phase in which the movements of the joints, which are initially not-coordinated, become coordinated, (c) a final phase in which the creature is coordinated, move at its full speed, and try to compensate the perturbations which occur during motion (which tend to reduce the level of coordination).

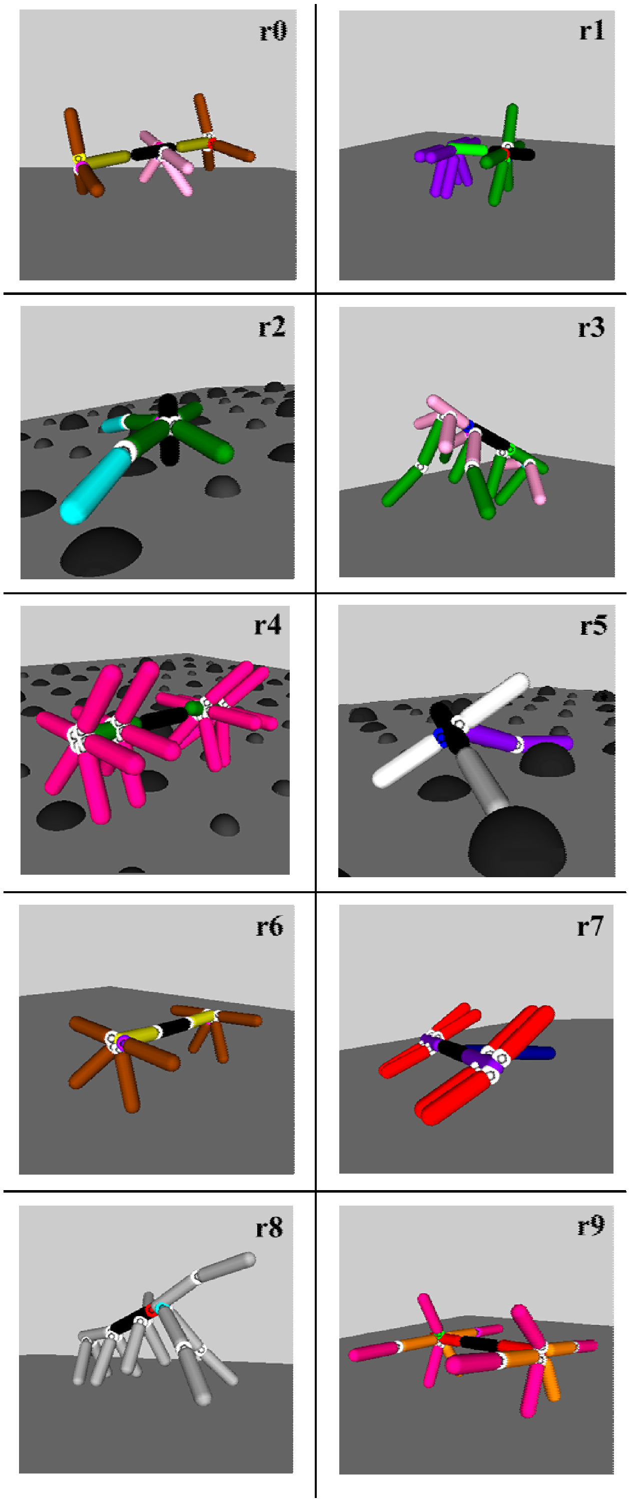

Figure 5: Morphology of the best evolved creatures of each replication of the experiment. Some creatures are shown on the flat and some creatures are shown on the irregular terrain. The grey scale colours of the segments and of the joints correspond to the id of the developmental rules which generated them. Joints of the same colour are provided with identical neural controllers. A true colour version of this picture is provided in the following web page: http://laral.istc.cnr.it/esm/dros1-0.

During the first phase, all the creatures (with the exception of replications r2 and r4, which are illustrated below) show an ability to assume a precise posture, after falling down on the ground, from which they are able to coordinate and locomote. This is accomplished thanks to the fact that the evolved morphologies are suitable both to stand on a preferred posture and to walk from that posture. Indeed, some of the evolved morphologies display segments which do not play a major role (or any role) for locomotion but which reduce the risk of tipping over (see for example r9). In some cases, creatures are able to assume more than a single posture from which they can locomote effectively. This is the case of creatures r2 and r4 which can assume two and four legged postures respectively, from which they can walk equally well.

During the second phase, all creatures show an ability to coordinate on an effective gait (with the exception of creature r0 and r3 which succeeds in only part of the trials). Some of the creatures (r7 and r9), do not use signals and coordinate by only exploiting the effects of the forces applied to each joint, mediated by the collisions with the external environment, on the posture assumed by the entire creature and the angular position of each of its joints. These physical interactions, in fact, codetermine the current state of each joint which, in turn, affects the propriosensors of each corresponding neural controller and thus the output of the neural controller itself which determines the desired speed of oscillation of the corresponding joint. It is important to note that this ability to coordinate depends on two aspects: (a) an ability to modify the frequency of oscillation of each joint on the basis of the actual position of the same joint which is determined by its previous position, by the previous postures of the entire creature, and by the previous actions executed by each joint, (b) the morphology of the creature.

The rest of the creatures (r0, r1, r2, r3, r4, r5, r6, r8) mainly exploit signals to coordinate. The analysis conducted indicate that coordination is achieved by producing signals which encode information about the current position of the joint emitting the signals and which are used by the neural controllers of the joints receiving the signal to regulate their speed of oscillation so to improve the coordination between the different part of the creature. For a detailed analysis of how a similar type of coordination process is achieved in robots with a fixed hexapod morphology provided with a similar distributed control system, see in [13] [14].

During the third phase, a qualitative observation of creatures' behavior indicate that they move at their full speed in a coordinated manner by compensating for relative misalignments between the joints which originate during motion especially in irregular terrains. Evolved creatures differ significantly in the way in which they locomote (see movies available from http://laral.istc.cnr.it/esm/dros1-0). In particular, we can identify three families of locomotion styles: (1) a quadruped gait style (displayed by r0, r4, r6 and r9); (2) a jumping style (displayed by r1, r3, r7, r9); (3) a dragging style in which the effect of the friction between un-actuated elements and the ground is minimized or maximized when the pulling legs are moving in the direction of motion or in the opposite direction (displayed by r2 and r5). As can be observed by inspecting the morphology and the behavior of evolved creature, the control mechanisms and the developmental rules are tightly co-adapted.

	Covered Distance (cm)		**Mutual Information**		**DOF**	**Signals**
	Flat Terrain	Irregular Terrain	Flat Terrain	Irregular Terrain		
r0	84.4890	54.3952	0.0137	0.0054	10	4
r1	73.3929	40.2574	1.5745	1.4689	5	4
r2	49.8966	25.4535	1.0629	1.1363	6	4
r3	79.1521	45.7687	0.0107	0.0090	15	4
r4	71.6282	59.7889	0.7690	0.6650	20	4
r5	62.1776	41.8614	0.2337	0.1038	4	4
r6	145.091	84.8877	0.0234	0.0230	10	2
r7	81.5839	71.2868	0.0029	0.0018	8	3
r8	134.195	87.2367	0.0244	0.0194	14	3
r9	46.3102	40.7141	0.0681	0.0718	8	3

Table 1: Performance and characteristics of the best evolved individuals of each replication. Covered distances values indicate the average distance in cm covered by creatures during 4000 steps (6 sec) after 2000 step from the beginning of each trial. During the test creatures have been situated for 50 trials on a flat terrain and 50 trials on a irregular terrain. Mutual information (DOF) values indicate the average mutual information calculated on each couple of joint. DOF indicate the number of DOF possessed by the individual. Signals indicate whether the possibility of the neural controller of each joint to produce signals and to regulate the current speed on the basis of detected signal is used or not by the evolved creature.

B. Variations of the model and their effect

To investigate how some of characteristics of the model described in the previous section influence the obtained results we have conducted a series of experiments in which we tested systematic variations of the model and in which variations has been retained or discarded on the basis of their effects in the observed results.

In most of test experiments, the variations considered did not produced significant effects. Below we restrict our analysis to the variations that produced the most significant effects on the agent’s morphology/behaviour and which highlights the crucial role of some of the model parameters or features. One of these variations concerns the introduction of the mutual information component in the fitness function. By comparing the results obtained in the experiment described above and in a control experiment in which evolving creatures were only rewarded for the ability to locomote, we observed that the creature evolved in the experiment in which the fitness include the Mutual Information component display better performance on the average but lower performance with respect to the best replication (Table 2).

The comparison between the performance exhibited by the best individuals of all replications of the experiment with and without mutual information indicates that, while in the former case all individuals display reasonably good performance on both flat and irregular terrains, in the latter case many individuals display good performance on flat but rather poor performance on irregular terrains. These differences are also reflected in the morphology which, in the latter case, is generally simpler and in the behaviour which, in the latter case, involves jumping or rolling strategies but never walking strategies.

		Flat Terrain		**Uneven Terrain**
Experiment with M.I.	Ave	82.791 cm	Ave.	55.165 cm
	Best	145.091 cm	Best.	84.887 cm
Test 1 without M.I.	Ave.	86.3055 cm	Ave.	46.0613 cm
	Best	149.434 cm	Best.	121.839 cm

Table 2: The values indicate the average distance in cm covered during 4000 steps (6 sec) after 2000 step from the beginning of each trial by the best individuals of 10 replications (Ave) and by the best individual of the best replication (Best) on flat and uneven terrains. All data refer to performance measured with respect to the distance traveled only.

Overall these data indicates that the addition of mutual information channel the evolutionary process toward types of solution which are more complex with respect to the morphology, more robust on the average with respect to different environmental conditions, less sensitive to the characteristics of the initial population, but not necessarily better in terms of performance (at least in the case of this task/environment).

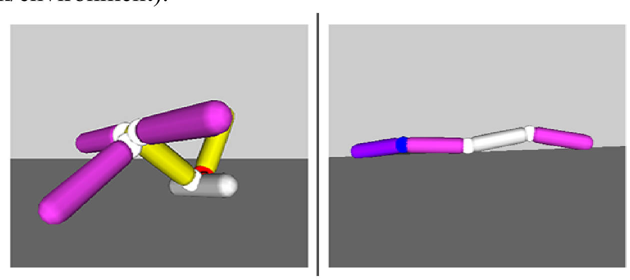

Figure 6: Examples of morphology of the best evolved creatures in test experiments in which creatures were evolved solely on flat terrain.

A second important variation consisted in the type of environment (flat and/or irregular) in which the creatures are evolved. The analysis of the results observed in a control experiment in which the creatures were evolved solely on flat terrain indicates that in this case the creatures tend to develop simpler morphologies which are less effective and which only work properly on flat terrains (see Figure 6). This data indicate that the complexity of the environment represents a crucial prerequisite for the development of effective creatures.

A third important variation consisted in increasing the number of orientations in which new segments can grow from four (0°, 90°,180°, 270°) to six (0°, 60°, 120°, 180°, 240°, 300°) orientations. The analysis of the results obtained in a control experiment in which there were only four possible orientations indicates that evolved creatures are much less stable, on average, with respect to the baseline experiment. Moreover, the evolved morphologies are much simple and less effective with respect to the basic experiment. This confirms the hypothesis that the expressiveness of the model might play a crucial role on the obtained results.

C. *Analysis of the evolutionary process*

To analyze the course of the evolutionary process we investigate, for each replication of the experiment, how the performance, the morphology, and the behaviour exhibited by the best creatures vary throughout generations. Rather than analysing the characteristics of the best individuals of each generation, for each replication, we reconstructed the lineage (i.e. the 499 ancestors) of the best individual of the last generation which includes only the genetic variations which have been retained in successive generations.

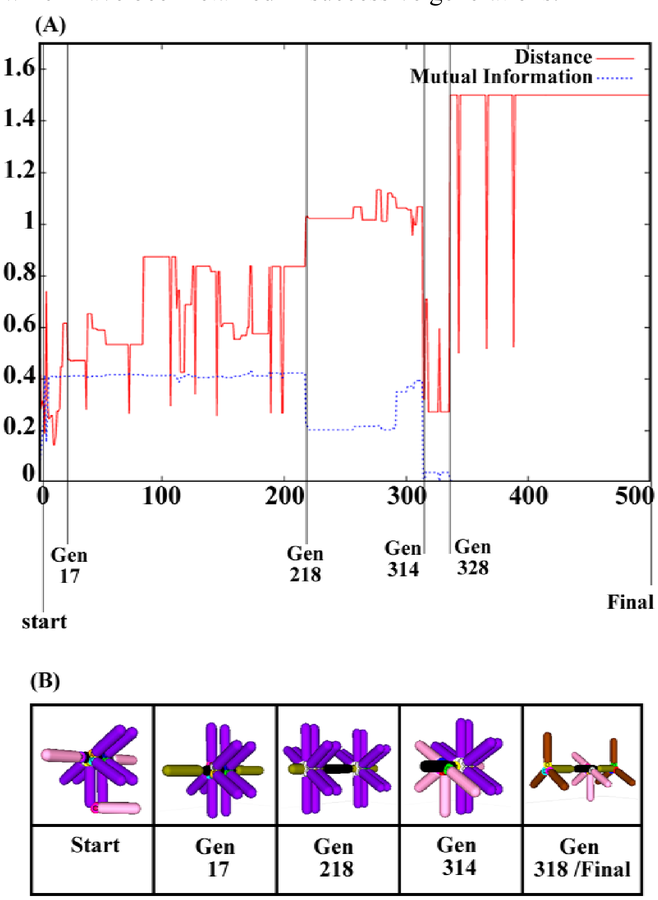

Figure 7 Average distance traveled and average mutual information of the ancestors of the best individual of the last generation of replication r0 throughout 500 generations..The full line indicates the distance traveled in cm. The dotted line indicates the mutual information (normalized in the range [0.0, 1.5]). The four vertical lines at generation 17, 218, 313, and 328 indicate the four major transitions with respect to morphological changes. The pictures at the bottom indicate the morphology of the individuals of the corresponding generations.

The analysis of these data indicates that, as for the case of replication r0 (see Figure 7):

(i) After the first generations, the morphology of the creatures undergoes a few crucial adaptive changes (at generation 17, 218, 314, and 318 in the case of r0), which represents crucial transition in the evolutionary process while remains rather stable between the phases that precede and follow transitions.

(ii) Once the performance of creatures increases, the probability of observing variations of the morphology which are adaptive become progressively lower. For this reason, after the first part of the evolutionary process, the occurrence of a major transitions which produce a significant modification of the morphology often arise as a result of variations which are retained despite they are maladaptive but which later become adaptive thanks to the verification and retention of further changes (as in the case of the variations occurring between generation 314 and generation 318 in the case of r0).

(iii) Mutual information tends to increase and then decrease during the initial and final phase of the evolutionary process, respectively (as in the case of r0). This can be explained by considering that the optimization of the distance travelled tend to interfere with the need to vary the desired speed of the joint in a coordinated manner. This data confirm the observation reported above which indicate that the attempt to maximize mutual information channel the evolutionary process toward the development of creatures which are able to coordinate the movements of their part, which in turn tend to lead to effective solutions. On the other hand the attempt to optimize mutual information interferes with the need to optimize the locomotion behaviour.

IV. DISCUSSION AND FUTURE WORK

We presented a model which allows to co-evolve the morphology and the control system of realistically simulated creatures. The method proposed is based on an artificial ontogenetic process in which the genotype does not specify directly the characteristics of the creatures but rather the growing rules that determine how an initial artificial embryo will develop on a fully formed individual. More specifically, the creatures are generated through a developmental process which occurs in time and space and which is realized through the progressive addition of both structural parts and regulatory substances which affect the successive course of the morphogenetic process. The creatures are provided with a distributed control system made up of several independent neural controllers embedded in the different body parts which only have access to local sensory information and which coordinate through the effects of physical actions mediated by the external environment through the emission/detection of signals which diffuse locally in space.

The analysis of the obtained results demonstrates how the method proposed allow the synthesis of effective and robust solutions which allow the evolved creature to effectively locomote both on flat and irregular terrains after relatively short number of generations in all the replications of the experiment. This general capability involves a capacity of assuming a preferred posture (by avoiding to fall on a side or to tip over), a capacity of coordinating the movement of several body parts so to produce a coherent behavior which allow the creatures to locomote, and a capacity to keep moving in a coordinated manner while compensating the perturbations which arise during motion and during the interaction with obstacles of variable size in the uneven terrain.

The fact that evolving creatures display good performance in all replication of the environment indicates that the model proposed guarantees a good level of evolvability. The fact that the evolved creatures display a variety of different morphologies and a variety of behavior strategies indicates that the model proposed display a good level of expressiveness. Finally, the relative simplicity of the model allowed us to test several variations of the model itself in order to identify the features and the parameters of the model which play an important role.

Interestingly, the analysis of the results obtained indicate than increasing the complexity of the task/environment (i.e. by asking the creature to locomote both on a flat and an irregular terrain rather than only on a flat terrain) might improve the performance of the evolved creatures also with respect to their ability to walk on a flat terrain.

The obtained results also demonstrates how the use of a task-independent measures (in the case of the experiment reported in this paper, the possibility to reward evolving individuals on the basis of the average mutual information computed between the output of each couple of neural controllers) combined with task-dependent measures (in this case the distance traveled by creatures within a limited time interval) can facilitate the maximization of the task dependent criteria (on this point see also [21]). More precisely, the inclusion of the task-independent component in the fitness function lead to solution which are more robust (with respect to environmental variations) and less sensitive to the initial conditions which, however, are not necessarily better in terms of absolute performance.

Finally, the analysis of the evolutionary process indicate how, after the very first generations, the morphology of evolved creatures tend to remain rather stable beside few significant changes which represent major transitions in the evolutionary process.

In future work we plan to analyze in more details the characteristics of the evolutionary process by also identifying quantitative measure for analyzing how evolvability varies throughout generations. Moreover, we plan to analyze in more details how the different parts of the creature coordinate, which are the principles that regulate how the morphology and the control system are co-adapted, and to what extent evolved creatures rely on form of morphological computations [22].

References

[1] K. Sims, “Evolving virtual Creatures”, *Computer Graphics*, vol. 28, pp. 15-34, 1994.

[2] K. Sims, “Evolving 3D morphology and behavior by competition. In *Proc. Artificial Life IV*, MIT Press, Cambridge, MA, 1994, pp 28-39.

[3] H. Lipson and J.B. Pollack, “Automatic design and manufacture of artificial life forms”, *Nature*, vol. 406, pp. 974-978, 2000.

[4] G.S. Hornby and J.B. Pollack, “Body-Brain Co-evolution using L-systems as a generative encoding”. In *Proc. of the Genetic and Evolutionary Computation*, Morgan Kauffman, San Francisco, CA, 2001, pp. 868-875.

[5] G.S. Hornby and J.B. Pollack, “Creating High-Level Components with a Generative Representation for Body-Brain Evolution”, *Artificial Life*, vol. 8, no. 3, pp.223-246, 2002.

[6] J. Ventrella, “Explorations in the emergence of morphology and locomotion behaviour in animated characters”, In *Proc. Artificial Life IV*, MIT Press, Cambridge, MA, 1994, pp 436-441.

[7] M. Komosinski and S. Ulatowski, “Framsticks: Towards a Simulation of a Nature-Like World, Creatures and Evolution”. In *Advances in Artificial Life*, LNAI vol. 1674, J.D. Nicoud, D. Floreano and F. Mondada, Ed. Berlin:Springer-Verlag, 1998, pp. 261-265.

[8] P. Eggenberger, “Evolving morphologies of simulated 3D organisms based on differential gene expression”, In *Proc. of the Fourth European Conference on Artificial Life*, MIT Press, Cambridge, MA, 1997, pp 205-213.

[9] J.C. Bongard and R. Pfeifer, “Evolving complete agents using artificial ontogeny”. In *Morpho-functional Machines: The new species (Designing Embodied Intelligence)*. F. Hara and R. Pfeifer Ed. Berlin: Springer-Verlag, 2003, pp. 237-258.

[10] K.O. Stanley and R. Miikkulainen, “A Taxonomy for artificial embryogeny”, *Artificial Life*, vol. 9, no. 2, pp. 93-130, 2003.

[11] S. Kumar and P.J. Bentley *On Growth, Form and Computers*. Amsterdam: Elsevier Academy Press, 2003.

[12] S. Wischmann, M. Hulse and F. Pasemann, “(Co)evolution of (De)centralized neural control for a gravitationally driven machine”. In *Advances in Artificial Life*, LNAI vol. 3630, M. Capcarrere, A.A. Freitas, P.J. Bentley, C.G. Johnson and J. Timmis Ed. Berlin:Springer-Verlag, 2005, pp. 179-188.

[13] M. Mazzapioda and S. Nolfi, “Synchronization within Homogeneous Neural Modules Controlling a Simulated Hexapod Robot”, In *Proc. Artificial Life X*, MIT Press, Cambridge, MA, 2006, pp. 199-205.

[14] M. Mazzapioda and S. Nolfi, “Synchronization and Gait Adaptation in Evolving Hexapod Robots”. In *Proc. of the Nineth International Conference on Simulation of Adaptive Behavior - From Animals to Animats 9*, Springer-Verlag, Germany, 2006, pp 113-125.

[15] M.R. Elphick, G. Kemenes, K. Staras and M. O’Shea, “Behavioural role for nitric oxide in chemosensory activation of feeding in a mollusc”, *Journal of Neuroscience*, vol. 15, no. 11, pp. 7653-7664, 1995.

[16] M.R. Elphick, L. Williams and M. O’Shea, “New features of the locust optic lobe: evidence of a role for nitric oxid in insect vision”, *Journal of Experimental Biology*, vol. 199, pp. 2395-2407, 1996.

[17] P. Husbands, A. Philippides, T.M.C. Smith and M. O’Shea, “Volume Signalling in Real and Robot Nervous Systems”, *Theory in Biosciences*, vol. 120, no. 3-4, pp. 253-269, 2001.

[18] H. Cruse, T. Kindermann, M. Schumm, J. Dean and J. Schmitz, “Walknet: a biologically inspired network to control six-legged walking”, *Neural Network, Special Issue*, vol. 11, no. 7-8, pp. 1432-1447, 1998.

[19] Open Dynamic Engine. Available: www.ode.org

[20] C. E. Shannon, “A Mathematical Theory of Communication”. *Bell System Technical Journal*, vol. 27, pp. 379-423, 623-656, 1948.

[21] M. Prokopenko, V. Gerasimov and I. Tanev, “Evolving Spatiotemporal Coordination in a Modular Robotic System”. In *Proc. of the Nineth International Conference on Simulation of Adaptive Behavior - From Animals to Animats 9*, Springer-Verlag, Germany, 2006, pp 558-569.

[22] R. Pfeifer, F. Iida and G. Gómez, “Morphological computation for adaptive behaviour and cognition”. In *International Congress Series*, vol. 1291, pp. 22-29, 2006.