Neurocomputing 500 (2022) 279–294

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

A study of growth based morphological development in neural network controlled walkers

M. Naya-Varela [a,*], A. Faina [b], A. Mallo [a], R.J. Duro [a]

[a] Universidade da Coruña, Integrated Group for Engineering Research, CITIC (Centre for Information and Communications Technology Research), Spain
[b] IT University of Copenhagen, Robotics, Evolution and Art Lab (REAL), Computer Science Department, Denmark

ARTICLE INFO

Article history:
Received 25 March 2021
Revised 15 September 2021
Accepted 25 September 2021
Available online 25 May 2022

Keywords:
Morphological development
Robot learning
Fitness landscape
Legged locomotion

ABSTRACT

In nature, the physical development of the body that takes place in parallel to the cognitive development of the individual has been shown to facilitate learning. This opens up the question of whether the same principles could be applied to robots in order to accelerate the learning of controllers and, if so, how to apply them effectively. In this line, several authors have run experiments, usually quite complex and heterogeneous, with different levels of success. In some cases, morphological development seemed to provide an advantage and in others it was clearly irrelevant or even detrimental. Basically, morphological development seems to provide an advantage only under some specific conditions, which cannot be identified before running an experiment. This is due the fact that there is still no agreement on the underlying mechanisms that lead to success or on how to design morphological development processes for specific problems. In this paper, we address this issue through the execution of different experiments over a simple, replicable, and straightforward experimental setup that makes use of different neural network controlled walkers together with a morphological development strategy based on growth. The morphological development processes in these experiments are analyzed both in terms of the results obtained by the different walkers and in terms of how their fitness landscapes change as the morphologies develop. By comparing experiments where morphological development improves learning and where it does not, a series of initial insights have been extracted on how to design morphological development processes.

1. Introduction

The joint development of the morphology and cognitive system has been shown to facilitate learning in human beings [1,2], who must face complex and changing domains during their lives. Taking inspiration from this fact, several researchers have addressed the application of the developmental principles observed in nature to robots, with the aim of creating autonomous systems able to learn by themselves [3]. All of this work can be grouped under the umbrella of Developmental Robotics, which has been defined by Cangelosi as "the interdisciplinary approach to the autonomous design of behavioral and cognitive capabilities in artificial agents that takes direct inspiration from the developmental principles and mechanisms observed in the natural cognitive systems of children" [4].

Even though most published work on developmental robotics has concentrated on cognitive development over fixed morphologies, ignoring the effect of developing morphologies, some authors have carried out different experiments to try to ascertain what this effect could be. The main focus of these authors was on replicating the cognitive structures as well the morphological developments observed in human beings rather than trying to find the underlying principles of morphological development and how they could be applied to robots. Among these, we can find cases in which the application of morphological development principles seems to improve learning [5–8], others in which the application of these principles had no effect [9–11], and even some in which morphological development turned out to be detrimental [12,13]. In other words, reviewing the literature it is not possible to extract a clear notion of the effects of morphological development on the learning abilities of robots and much less how these techniques should be applied to be able to design successful morphological development processes.

Additionally, the different experiments carried out within the field present two important characteristics that hinder the study of morphological development in robots and finding the underlying principles that govern these processes: (1) As they try to mimic human development as faithfully as possible, most experiments

* Corresponding author.
E-mail addresses: martin.naya@udc.es (M. Naya-Varela), anfv@itu.dk (A. Faina), alma.mallo@udc.es (A. Mallo), richard.duro@udc.es (R.J. Duro).

https://doi.org/10.1016/j.neucom.2021.09.082
0925-2312/

employ very complex control systems and cognitive mechanisms [14,15], as well as often complex robotic devices. This makes it very hard to attribute causality to the effects observed, i.e., establish what is due to morphological development and what to the complex control systems and robotic structures. (2) The experiments themselves as well as the robots used are usually very specific and heterogeneous making it very difficult to compare and generalize results.

A representative example of this, considering the reaching task, is the study of Gomez et al. [16]. They studied the influence of development with an ad-hoc robot arm configuration using a complex neural network based control system whose topology changes along the robot's developmental process. Similar experimental characteristics are observed in Lee et al. [17] where they also tried to learn arm motions for reaching, but considering a different hardware structure and experimental configuration. Using these ad-hoc robotic arms, and their complex cognitive structure, hinders the task of identifying why morphological development in these cases has helped to improve the learning performance and makes it difficult to compare and generalize results.

This heterogeneity and lack of information can also be observed in grasping and walking examples. In the case of grasping using a three fingered hand controlled by a Continuous Time Recurrent Neural Network (CTRNN), Bongard [10] showed how the joint evolution of the morphology and the controller could help to improve the performance of the task. He hypothesized that morphological changes improved performance when the task was complex (grasping three different objects) but it seemed to be irrelevant in simple tasks (grasping one object). However, the complexity of the problem was given by the number of simultaneous objects to grasp, preventing the generalization of this definition to other non-grasping tasks.

In the case of walking, the heterogeneity becomes even more pronounced as different morphological development strategies such as growth or variation of the Range of Motion (ROM) have been used. In this line, Bongard [18] presented morphological development experiments based on these two strategies using quadruped and hexapod robots, also controlled by CTRNNs, whose task was to reach a source of light. He suggested that the learning improvement was motivated by the gradual growth of the legs, while abrupt changes in the morphology, related to the initial position of the legs led to a reduction in performance in ROM based development. Although he indicated that this fitness reduction was motivated by the abrupt change in the controller due to the sudden morphological changes that increase the ROM available, the paper did not dwell on the reasons underpinning the success of growth or the poor results of those abrupt changes in the morphology.

Furthermore, he indicates that "manually clamping some degrees of freedom in a robot with a fixed body plan, and then gradually releasing them during learning, can accelerate behavior acquisition". This statement is based on the hypothesis that in the early stages of development during growth, the movement of the legs has relatively little impact on the robot displacement, being the movement of the spine of the robot more relevant. Thus, initially, only the spine motors must be controlled, reducing the difficulty of coordinating the different robot motors.

The strategy of clamping or freezing and freeing Degrees of Freedom (DOF), a particular case of ROM, is based on the studies by Bernstein on human motor development [19]. He postulated that an initial reduction of the available DOF may help to learn motor skills. The success of this strategy is attributed to the initial reduction of the dimensionality of the search space due to the frozen DOF. This developmental strategy has been successfully applied by Lungarella and Berthouze over a bipedal robot [20]. In this case, starting with all the DOF reduced the likelihood of physical entrainment that takes place between morphology, neural, and environment dynamics, while an initial reduction of the DOF available helped to stabilize the system and to find more robust behaviors. However, they also found that this developmental strategy was not enough to produce a stable behavior when a perturbation was added to the system [21]. In this case, a single cycle of freezing and freeing DOF led to the collapse of the system.

Vujovic et al. [22] also postulated the advantage of a developmental approach while evolving and developing a morphology in a 3D printed robot instead of simply evolving the morphology. The authors stated that this performance improvement only happened "if the evolutionary and developmental terms are properly adjusted". However, the paper did not provide details on how this adjustment could be made. As they clearly mentioned "the interplay between development and evolution during this process is complex and not yet fully understood". Furthermore, they found that the evo-devo strategy could also be detrimental, suggesting that this decrease in performance could be attributed to the reduction of the search space, which could in some cases remove good choices.

Summarizing, although in the last twenty years several authors have carried out very relevant research on the application of morphological development in robots, the resulting body of research is very inconclusive in terms of when and what type of morphological development is useful. This is mainly due to three factors: (1) The authors have focused on replicating human-like development and not on extracting underlying principles that could facilitate the design of morphological development processes for robots. (2) Most of the experiments present very convoluted setups, with complex controllers and ad-hoc robots, making it very difficult to replicate the experiments or extract conclusions on what effects are actually due to morphological development. (3) The experiments were very heterogeneous, using different types of robots, different tasks, different morphologies, different controllers, thus making it very difficult to compare and generalize results.

Furthermore, very little is said with respect to how a morphological development process should be constructed in order to improve learning. In general, three main hypotheses have been proposed in the literature to explain why morphological development could be unsuccessful in improving learning [23]: (1) The development process is not well aligned with the nature of the task [12]; (2) The problem is too simple and, therefore, development does not provide any advantage [10,11]; (3) The development process also generates perturbations that negatively affect learning [18,21].

In this paper, we will analyze a developmental process and see if these hypotheses hold and, more importantly, we will try to provide some initial general insights on how morphological development processes could be constructed to improve their success. Thus, our aim is not to replicate human-like morphological development processes, but rather, to study morphological development in robots and to try to extract insights on how it can be applied to improve robot learning. Additionally, to alleviate complexity and improve replicability, we establish a very simple and controlled experimental framework based on standard tools and consider clear morphological development strategies and simple controllers. In fact, in this paper we only consider growth. This will allow us to extract what effects are due to morphological development and what effects are due to other factors in the experiments. Finally, to homogenize the setup, we will use the same neural network based controller and learning algorithm and the same family of morphologies over the same tasks so that results can be compared and generalized.

This paper is an extension of the paper we presented at the HAIS 2020 conference [24], where we began a process aimed at addressing these issues. Here, we take those experimental results, expand

them with new experiments and go one step further by proposing a formalization of morphological development. We also look into the evolution of the fitness landscapes of the problems addressed as development takes place in order to try to understand the underlying factors that determine the effects of morphological development.

To achieve these goals, the article is structured as follows: Section 2 provides a formal definition of morphology and morphological development as used in this paper. Given that the influence of morphological development seems to depend on problem difficulty, in Section 3 we discuss how this difficulty can be ascertained through fitness landscape representations and through their progression during development. In this sense, fitness landscape analysis is the main tool selected to achieve the goal of understanding how and why morphological development influences learning under a global and homogeneous perspective. Section 4 deals with the methodological aspects of the growth based morphological development experiments that have been carried out. The results of these experiments are presented in Section 5. Section 6 addresses a representation of the fitness landscapes that arise during the developmental process for each morphology and an analysis of their features. A discussion of the results obtained and how they are related to the features of the fitness landscapes is presented in Section 7. Finally, a series of conclusions of this work consisting in some insights on how to apply morphological development and some avenues for further research are presented in Section 8.

2. Morphological development

To provide a clear notion about what we refer to when talking about morphological development, we are going to propose a formalization of robot morphology and morphological development based on the one given by Naya-Varela et al. [23]. This formalization will be particularized to the growth developmental strategy, which is the developmental strategy used in the experiments of this article.

Thus, we consider that a robot morphology is made up of a set of l links $\boldsymbol{L} = \{l_1, l_2, \ldots, l_l\}$, a set of j joints $J = \{j_1, j_2, \ldots, j_j\}$, which can be actuated or not, and a set of s sensors $S = \{s_1, s_2, \ldots, s_s\}$. This morphology runs during time $t \in [0, T]$, where t = 0 is the beginning of its lifetime and t = T is the end. The robot has sets of properties for the links (L_p), joints (J_p) and sensors (S_P) with a cardinality of x, y and z, respectively. These property sets can be expressed as:

$$\boldsymbol{L_P} = \begin{bmatrix} lp_{11} & \cdots & lp_{1l} \\ \vdots & \ddots & \vdots \\ lp_{x1} & \cdots & lp_{xl} \end{bmatrix}$$

$$J_{\boldsymbol{P}} = \begin{bmatrix} jp_{11} & \cdots & jp_{1j} \\ \vdots & \ddots & \vdots \\ jp_{y1} & \cdots & jp_{yj} \end{bmatrix}$$

$$\boldsymbol{S_P} = \begin{bmatrix} sp_{11} & \cdots & sp_{1s} \\ \vdots & \ddots & \vdots \\ sp_{z1} & \cdots & sp_{zs} \end{bmatrix}$$

Therefore, a robot morphology can be defined as the set of links, joints and sensors that make up the robot and their properties$m = \{L, J, S, L_p, J_p, S_p\}$. This leads to the general definition of morphological development as a function $MD(t)$ that describes the evolution in time of the values for the properties in these property sets. Thus, morphological development can be formally defined as a non-stationary function $MD(t)$ that describes the values of these properties in time for the lifetime of the robot:

$$MD(t) = \{L_t, J_t, S_t\} \quad \forall t \in [0, T] \tag{1}$$

In this article, we are only considering morphological development based on growth, taken as changes in the lengths of the links (L_P). Thus, the morphology of the different robots presented in this article can be specified at the beginning of their lifetime as $\mathfrak{M}_{t=0} = \{L, J, S, {}^{L}P_{t=0}, {}^{J}P, {}^{s}P\}$ and by $\mathfrak{M}_{t=T} = \{L, J, S, {}^{L}P_{t=0}, {}^{J}P, {}^{s}P\}$ at the end. That is, the values of the properties represented by ${}^{J}P$ and ${}^{S}P$ remain constant throughout the lifetime of the robot and the only ones that change are those corresponding to L_p. Under these assumptions, morphological development as addressed in this article can be formally defined as a function $MD(t)$ that describes the values of these properties in time for the lifetime of the robot as:

$$MD(t) = \{ \quad {}^{L}P_t \} \ \forall t \in [0, T] \tag{2}$$

Where L_p specifies the length for each link of the robot:

$$L_p = [length_0, \ldots, length_l] \tag{3}$$

In the majority of natural systems, this $MD(t)$ function is a continuous function that describes the morphological development path an individual has followed from birth throughout its life, and which depends on intrinsic (genetic) and extrinsic (environment) conditions. As indicated later in the paper, for the experiments that are carried out here, this function will also be continuous and linearly growing in the length of the extremities, however, and for the sake of simplicity, this growth will not depend on environmental conditions.

3. Learning difficulty

One issue that always comes up when looking at different learning strategies and their performance, is that of problem or learning difficulty. This obviously impacts studies on morphological development when trying to discern the relationship of different morphological development strategies and results to the difficulty of the learning problem that is being considered. To address this question, it is possible to view learning as an optimization problem, where the objective consists in finding a controller, specified by a set of parameters, that allows the robot to optimally perform a task.

A consequence of this perspective on learning is that the difficulty of learning a task for a given parametrized controller structure acting over a particular morphology in a specific environment can be expressed in terms of the hardness of the optimization problem. In this line, there are quite a few studies that address this issue [25,26] by making use of the features of the fitness landscape to try to provide an indication of problem hardness.

A fitness value is a measure that expresses the quality or performance of each possible controller instance that may arise during learning [27]. This fitness value associates a numerical value to the combination of parameter values that define a specific instance of controller as a function of the performance of the task by the robot morphology in its environment using this controller instance. Thus, a fitness landscape or a landscape of possible solutions can be defined as the representation of the fitness in terms of the value for all the possible solutions available [28]. In this framework, different authors address features such as number and distribution of local optima with regards to global optima, and many others, in order to specify the relative hardness of a problem. In particular, a term that is frequently used to describe one characteristic of the landscape is the concept of attraction basin [29]. As Caa-

maño et al. mention "by definition an attraction basin of a local optimum is the set of points $x_1 \ldots x_k$ of the search space such that a local search algorithm starting at x_i ($1 \leq i \leq k$) ends at the local optimum." Using this concept of attraction basin, different landscape features may be considered based on the number of different attraction basins and the width of the base of these attraction basins [28].

Thus, based on the ideas stated above, there are a series of features of a fitness landscape that can be used to summarize its hardness:

- Deceptiveness. A deceptive fitness landscape means that the majority of the landscape would steer towards a local optimum (or more than one) that is not the global optimum. That is, the attraction basins of one or all the local optimums are much larger than that of the global one.
- Neutrality. A neutral fitness landscape is characterized by the lack of information about the situation of the global optimum. As a consequence, finding the global optimum becomes a random search process.
- Roughness. A rough fitness landscape is defined by a large number of local optima in the landscape with attraction basins that are comparable in width to the global one. This increases the chances of getting stuck in a local optimum for any optimization algorithm.
- Smoothness. A smooth fitness landscape means that there are few and shallow local optima and that the attraction basin of the global optimum dominates the fitness landscape.

A relationship between the fitness landscape and difficulty of the problem for a developmental stage was used by Vujovic et al. [22] and more in detail by Lungarella and Berthouze [20]. In this case, Lungarella and Berthouze studied the influence of changing some configuration parameters of the experiment to study how the fitness landscape in a developmental stage varied. However, what we want to address in this article is neither the characteristics of a fitness landscape for a specific developmental stage nor how that landscape is modified due to changes in the parameters of the morphology. Here we will study the evolution of the fitness landscapes throughout the different developmental stages for different successful and unsuccessful growth based morphological development experiments and try to determine how these changes in the fitness landscapes are related to the results about the relevance of growth for learning.

4. Experimental setup

To achieve the objective stated in the previous section, we have established an experimental framework based on a walking task on level ground using simple legged robots: a quadruped, a hexapod, and an octopod. The task and the different morphologies have been chosen for several reasons: (1) Due to the intrinsic stability of these morphologies, we consider that learning to walk will not be an overly difficult task. Thus, it will be possible to use simple controllers, which allow us to focus our efforts on studying the influence of morphological development. (2) The influence of the designer on learning how to walk is minimum compared to other tasks, such as grasping, because there are no decisions to be made about external factors. For example, grasping means "grasping something" and this object to grasp must be selected by the researcher (its shape, size, orientation, etc.), thus, biasing the process of learning. (3) Morphological development, especially body growth, has been successful in legged morphologies as shown by several articles [18,22,30]. Thus, we consider that legged morphologies on a walking task are a suitable starting point for understanding the effects of morphological development.

In order to perform the experiments, we employed the V-REP simulator (currently CoppeliaSim, [31]) with the Open Dynamics Engine [32] as physical engine. The selected morphologies are, as mentioned, a quadruped, a hexapod and an octopod (Fig. 1) and all were modelled in the simulator. All the morphologies are based on the same base morphology, made up of a central body and a series of limbs attached to it, as shown in Fig. 2 left. The limbs are the same for all morphologies. They are composed of an upper link and a lower link and the corresponding revolute joint (Fig. 2 right). The upper link measures 5x2.5x0.5 cm and has a mass of 250 g. The lower link is made up of two elements, with the same dimensions and mass as in the upper link, joined by a prismatic joint, which is the one that allows growth. The prismatic joint can apply a maximum force of 50 N and it is controlled by a proportional controller (P = 0.1). All the prismatic joints of the legs have a maximum stroke of 7.5 cm, which means that the length of the lower link may vary from 10 cm to 17.5 cm. In addition, there are two revolute joints in each limb that join the central body and the upper link, and the upper link and the lower link, respectively. All the revolute joints are actuated and have a maximum range of motion of [−90, 90] degrees. Their maximum torque is 2.5 Nm and they are also controlled through a proportional controller (P = 0.1).

Thus, the only differences between the morphologies are given by the number of limbs and the size and mass of the central body. The morphology of the quadruped is shown in Fig. 1 left. The size of its central body is 30x15x1 cm and its mass is 2 kg. The hexapod's central body dimensions are 60 × 15 × 1 cm and its mass is 4 kg (Fig. 1 middle). The dimensions of the central body of the octopod are 90 × 15 × 1 cm and its mass is 6 kg (Fig. 1 right). Of course, the number of limbs available also changes, and it is obviously 4 for the quadruped, 6 for the hexapod, and 8 for the octopod. Thus, the quadruped presents a total of 8 DOF, the hexapod 12 DOF and the octopod 16 DOF (we do not take the prismatic joints into account as they are part of the lower link, only used for growth, and they maintain the same position during each evaluation of the controllers).

Furthermore, we have created two additional types of morphologies for the quadruped: a quadruped with 16 DOF, and another one with 24 DOF. In the 16 DOF quadruped, the original upper link has been replaced by 3 elements of 1.3 × 2 × 3 cm and 0.39 kg of mass and 2 extra joints, see Fig. 3 left. Regarding the 24 DOF quadruped, the original upper link has been replaced by elements of 0.5 × 2 × 3 cm and 0.5 × 3 × 2 cm, according to their orientation, and four extra joints, see Fig. 3 right. The torque of all revolute joints is 2.5 Nm.

Many locomotion studies utilize neural oscillators [20] or CTRNNs [18,33] to generate rhythmic and periodic control signals. Here, and for the sake of simplicity we have chosen to use Neural Networks (NN) with sigmoidal activation functions that act over a sinusoidal input signal to generate the appropriate periodic signals. Sinusoidal inputs have been used by authors such as [34] to facilitate the generation of temporal patterns. Unlike them, in our case we have not made use of perceptual inputs. To avoid having to define the structure of the networks (number of hidden neurons, layers and connections), we have decided on the use of a neuroevolutionary approach for producing the appropriate network architecture and weight values. Furthermore, as we have planned to continue the research line with different morphologies and control signals, we have selected an ANN due to its flexible adaptation to different scenarios. In particular, a MultiNEAT [35] implementation of the NEAT [36] algorithm was chosen. The NEAT algorithm causes that the initial feedforward NN may end up including different numbers of hidden neurons with different configurations, and even recurrences depending how the architecture evolves (Fig. 4). Thus,

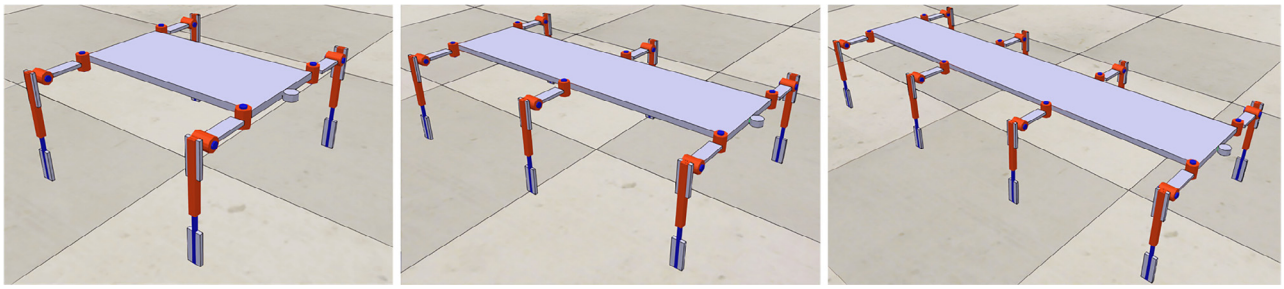

Fig. 1. Snapshots of each morphology considered in the experiments in their resting position. Left: quadruped. Middle: hexapod. Right: octopod. Each limb has three solid segments in a chain attached to the base by two actuated revolute joints (red cylinders) and a linear joint (red rectangular cuboid), which is used for the morphological development. The limbs in each side are equidistant from each other at a distance of 29 cm. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

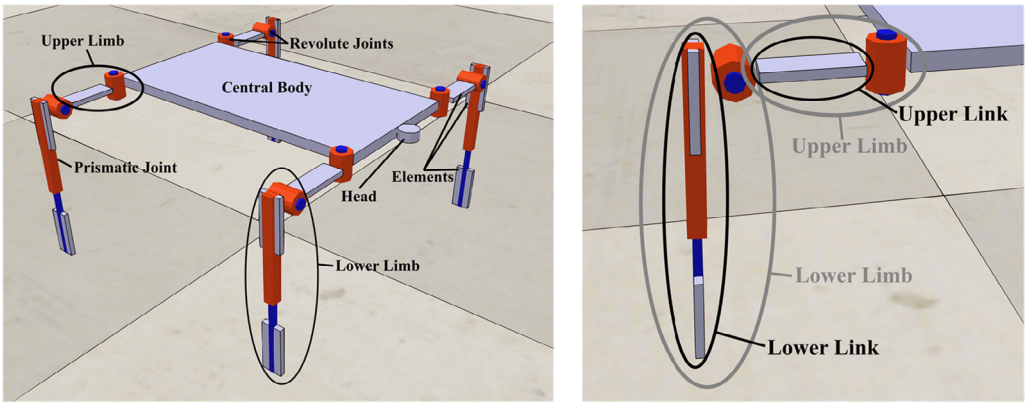

Fig. 2. Left: Configuration of a base morphology (quadruped) where the different parts are displayed. Right: Detail of the lower limb and upper limb, including the lower link and upper link respectively.

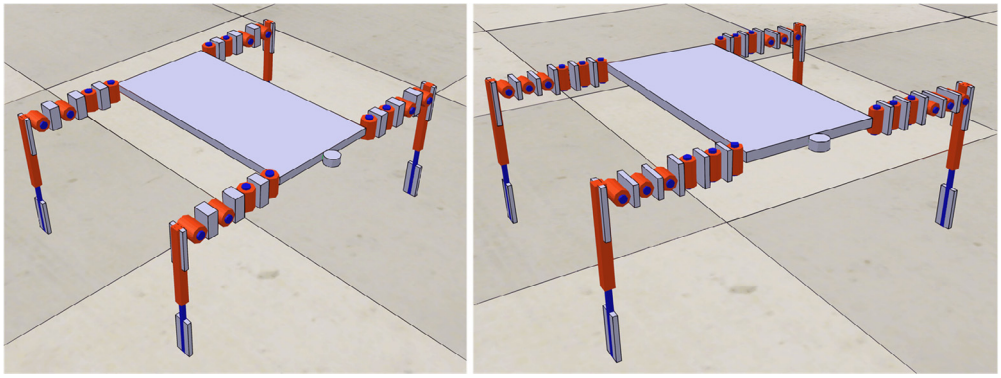

Fig. 3. Snapshots of the quadruped morphology with 16 DOF (left) and 24 DOF (right).

outputs with small phases can be obtained thanks to the recurrent connections, or larger phase differences (phase and anti-phase) are obtained due to different signs in the outputs. The sinusoidal input has an amplitude of 2 rad and an angular velocity of 10 rad/s and it is input without any normalization to one of the two inputs to the networks. The other input is a bias with a value of 1. The outputs of the NN correspond to the control signals for each one of the joints through a denormalization from the NN sigmoidal output interval [0,1] and the range of motion of each joint [−90, 90]. Consequently, the number of outputs depends on the number of DOF of the morphology. The NEAT processes start from an initial population of networks consisting of just the input layer and the output layer and random weights (the weights are always in the [-10, 10] interval). Initially, each neuron of the input layer is connected to all of the output neurons, without any recurrence. Although the NEAT algorithm has a large number of configuration parameters, the most basic and relevant ones are presented on Table 1 (and all parameters are displayed in the GitHub repository of the project related to this article[1]).

Every independent run of NEAT evolves a population of 50 individuals for 300 generations. To gather statistical data, a total of 40 independent runs have been carried out for each experiment. Due to the large number of experiments that were carried out, we have used of a high performance cluster: CESGA [37]. The fitness of the individual is calculated directly as the distance traveled in a

[1] [1]https://github.com/GII/morphological_development/blob/main/publications/2021_neurocomputing/source_code/config/evolution_config.py.

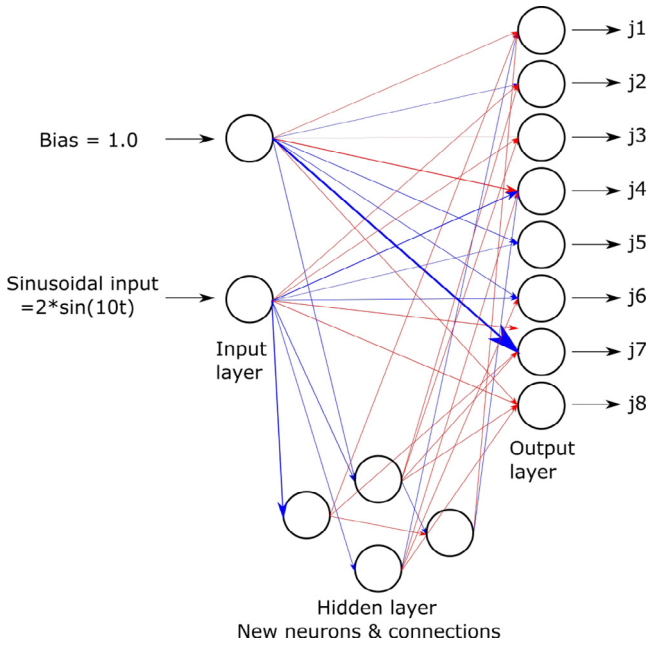

Fig. 4. Example of an NN topology obtained with the NEAT algorithm for the 8 DOF quadruped. The input layer consists of a sinusoidal input and the bias, while each neuron of the output layer controls each joint. Initially, the input layer is fully connected with the output layer without any hidden layer. However, as the evolution advances, NEAT adds new neurons and connections to the ANN, creating a heterogeneous hidden layer. Blue lines represent weights with positive values and red lines weights with negative ones and the thickness of the lines represent the value of the weights. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
Basic configuration parameters of the MultiNEAT algorithm.

MultiNEAT parameter	Value
Population Size	50
Minimum Species Size	2
Maximum Species Size	7
Survival Rate	0.3
Crossover Rate	0.6
Mutate Weights Probability	0.5
Elite Fraction	0.02
Mutate Add Neuron Probability	0.5
Mutate Add Link Probability	0.5
Recurrent Probability	0.1

straight line by the head of the robot. Each individual of the population is tested for 3 s with a simulation time step of 50 ms and a physics engine time step of 5 ms. The source code of each experiment can be found in the repository[2].

To study how morphological development influences learning in different morphologies, we have carried out two kinds of experiments for each morphological configuration:

1. *No development experiment*: This experiment is run with a fixed morphology (the same as the final morphology for the rest of the experiments) from the beginning to the end. The robot starts at generation 0 with the maximum length of the lower links and the neuro-evolutionary algorithm seeks a neural network-based controller to achieve maximum displacement.
2. *Developmental experiment*: The robot morphology starts with the shortest version of the links at generation 0 (the prismatic joints are fully contracted, with an extension of 0 cm, and therefore, the length of the links is 10 cm). The link length is grown linearly and simultaneously for all the legs for a number of generations until it reaches the maximum length of 17.5 cm. Please, note that each growth step is applied before each evaluation starts and the robot does not grow during the evaluations. Additionally, growth takes place in a set number of generations for each experiment. That is, the final morphology is reached at generation 20, 40, 60, 80, 100 or 120 depending on the experiment. This permits studying the relevance of the growth rate with regards to performance.

More formally, as defined in the previous section, for $t \in [0, T]$, where t = 0 is the beginning of the robot's lifespan and t = T the end (300 generations in our experiments), the quadruped undergoes variations in its ^{L}P parameters (length of each link), in this case linearly from no extension to full extension of the links, until the end of development ($t = \tau$), where τ is set to 20, 40, 60, 80, 100 or 120, depending on the experiment. Thus, the morphological development function $MD(t)$ is defined as:

$$MD(t) = \begin{cases} L_P(t), & t < \tau \\ L_P(\tau), & t \geqslant \tau \end{cases} \tag{4}$$

With:

$$l_{LLl}(t) = L_{init} + \frac{\Delta L}{\tau} * t \quad \forall l \in [1,\ 4] \tag{5}$$

$$l_{UL1}, l_{UL2}, l_{UL3}, l_{UL4} = 5\ cm$$

$$L_{init} = 10\ cm$$

$$\Delta L = 7.5\ cm$$

where *ULx* indicates upper link *x* and *LLy* indicates lower link *y*. Also, as we are considering only one modifiable property for the links (their length) the link properties matrix of the previous chapter is presented here as a vector for simplicity.

The results obtained from each of the experiments have allowed us to compare the differences that exist between the methods applied to each morphology, and the differences in the results for different designs.

5. Results

To present the results, we have split the experiments into two different groups. The first one presents the results for the quadruped with 8 DOF, the hexapod and the octopod. The second one corresponds to the results of the three different quadrupeds (8, 16 and 24 DOF) in order to evaluate whether the number of degrees of freedom has any bearance. An example of the best gait obtained for each morphology and type of experiment can be viewed in the GitHub repository.[3]

The results of the learning process and the statistical results for each morphology of the first group are displayed in Fig. 5, left. Specifically, it shows the results obtained after the learning process through neuro-evolution for the no-development case and just one of the growth experiments (growth up to generation 60). The solid and dashed lines display the median of the best fitness obtained in 40 independent runs for the growth experiment and the no-development one, respectively. The shaded areas represent the

[2] [2]https://github.com/GII/morphological_development/tree/main/publications/2021_neurocomputing/source_code.

[3] [3] https://github.com/GII/morphological_development/tree/main/publications/2021_neurocomputing/videos.

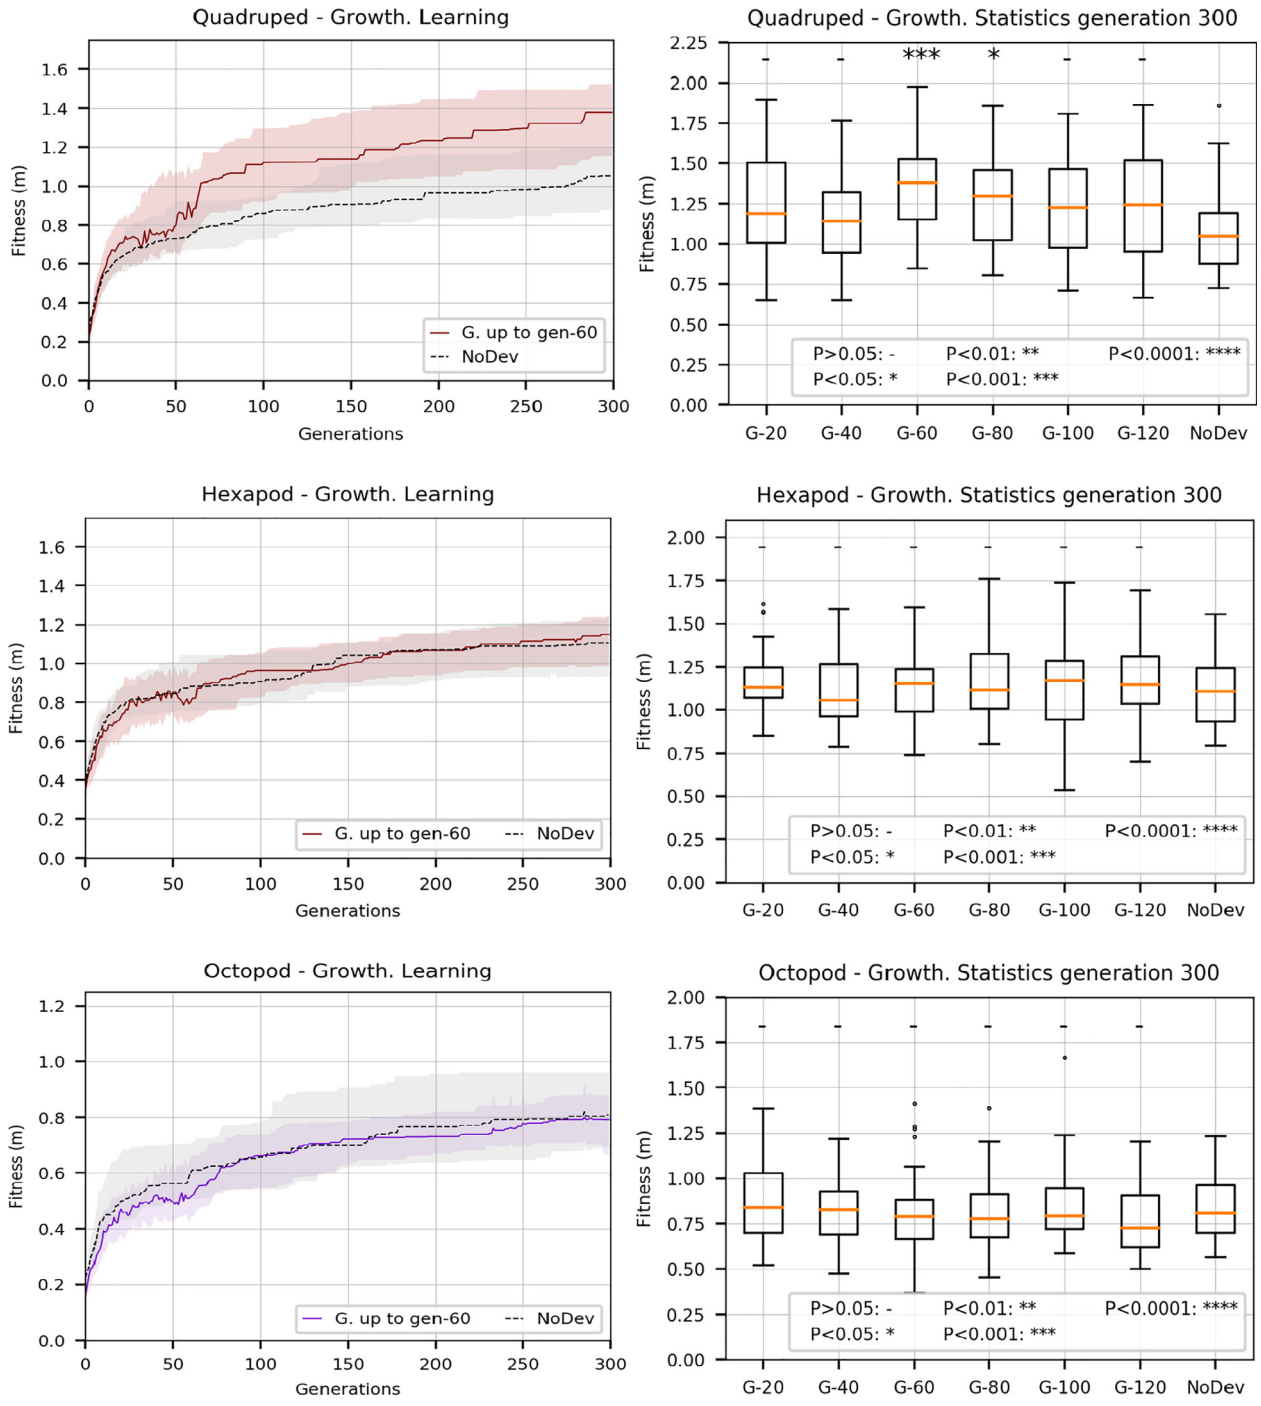

Fig. 5. Results of the learning process (left) and the statistical analysis at the end of learning (right) for the quadruped (top), hexapod (middle) and octopod (bottom). All the developmental experiments are compared to the no-development experiment and the p-values of the Mann-Whitney test, corrected by Bonferroni correction, have been replaced by asterisks. The greater the number of asterisks, the greater the statistical difference as indicted in the legend. G. up to means "Growth up to".

areas between percentiles 75 and 25 for each experiment. Although we have run the experiments with different growth rates, as mentioned in section 4, for the sake of clarity in this graph, we only display the results of the growth up to generation 60 experiment as it provides the most relevant results. A statistical analysis based on the Mann-Whitney *U* test [38] was carried out to test for the statistical significance of the growth experiments. Specifically, all the growth experiments were compared against the no-development one. The statistical results are represented by a series of boxplots (Fig. 5 right). Each boxplot represents the median and the 75 and 25 quartiles in the last generation for 40 independent runs of each of the different types of experiments. The whiskers are extended to 1.5 of the interquartile range (IQR). Single points represent values that are out of the IQR. A p-value of 0.05 is taken as the significance value for accepting or rejecting the null hypothesis. All the p-values have been adjusted using the Bonferroni correction [39]. Fig. 5 displays the statistical results for the three morphologies once this correction has been applied.

These results show that the developmental strategy based on growth only offers better results than no development in the case of the quadruped and only at specific growth rates. Furthermore, growth was irrelevant for learning, but not detrimental, in the case

of the hexapod and octopod. Analyzing the results obtained in more detail, it can be observed that:

1. During learning, as shown in the graphs in the left column of Fig. 5, the median fitness value of the no-development experiments increases progressively and continuously from the beginning to the end of learning. However, the median of the growth experiments displays a noisy behavior during the growth phase ($t < \tau$). These oscillations in the fitness value are motivated by the continuous adaptation of the controller to the morphology at each developmental stage. This obviously ends once the growth phase ends.
2. Different growth ratios give rise to different results. The top-right graph of Fig. 5 shows how significant statistical results are only obtained for growth up to generations 60 (p-value 0.0002) and 80 (p-value 0.0158). Furthermore, performing a statistical analysis at different stages of the learning process, we can observe how growth is also relevant in the early stages of learning and this relevance increases with the number of generations, as displayed in Fig. 6. Thus, there exists an optimal growth rate interval around which the best results are obtained during learning. This fact seems to indicate that the success or failure of a given morphological development strategy also depends on the capacity of the learning algorithm to adapt to the changes in morphology that take place between stages.
3. The absolute fitness value obtained by the hexapod is lower than the fitness obtained by the quadruped, and the fitness of the octopod is even lower. This can be motivated by the configuration of the control system. With a single sinusoidal input, it is very hard to obtain an ANN that results in any other phase behaviors than the most basic phase or anti-phase ones (with slight variations if there is any recurrence) with respect to the inputs. Thus, optimum phase adjustments for each joint of a robot may not be readily available, hindering the task of finding optimum gaits. This effect may increase with the number of DOF of the robot and affects both the no-development and growth cases.

Given these results, a pertinent question would be whether the higher success of growth based morphological development in the quadruped as compared to the hexapod and octopod is simply related to the lower number of degrees of freedom of the morphology (making the search space smaller) or, whether the intrinsic characteristics of the morphology somehow play a role in this. For example, although the joints have the same torque and the leg design is the same for the three morphologies, the quadruped presents both a smaller size and weight as well as a reduced number of limbs compared to the hexapod and octopod. These characteristics may imply that some walking solutions for the quadruped, such as flexible dynamic gaits, may not be available to the hexapod and octopod.

To elucidate the relevance of the number of degrees of freedom and the characteristics of the morphology, a series of experiments over quadrupeds with different DOF were carried out. In these experiments, we have analyzed and compared the results obtained after applying growth up to generation 60 for a quadruped with 8 DOF (the one presented above), 16 DOF (as in the case of the octopod), and 24 DOF. The results of these experiments are displayed in Fig. 7. They show that growth based morphological development improves learning compared to the no-development experiment in the quadruped with 16 DOF (p-value of 0.0234), although to a lesser extent than in the 8 DOF quadruped. In the case of the 24 DOF quadruped, no statistically significant improvement can be appreciated (p-value of 0.1673). These results are especially relevant in the quadruped with 16 DOF because, even though it shares the same DOF as the octopod, unlike in the case of the octopod, growth based morphological development actually improves learning.

The results obtained from the experiments seem to indicate that there are two opposing factors that influence the relevance of growth as a morphological development strategy in these cases. On the one hand, the intrinsic characteristics of the morphology may provide more avenues of improvement and thus favor learning using morphological development. That is, they may make the problem easier. This is the case of the quadruped with the same DOF as the octopod, for which morphological development favors learning. However, at the same time, increasing the number of DOFs (parameters to be controlled) makes the search space larger and thus works against learning. In fact, it can be seen that as the number of degrees of freedom increases, the absolute fitness value obtained decreases until a point is reached where the difficulty induced by the large number of DOFs cancels the advantages of the morphology and makes morphological development irrelevant, as in the case of the quadruped with 24 DOF.

6. Fitness landscape representation

As stated in Section 3, to try to ascertain the learning difficulty induced over a task by different morphologies as well as by different morphological development processes in a morphology independent representation, we will resort to studying the fitness landscapes resulting from the different experiments. More specifically, we want to compare the fitness landscape of the quadruped (as it represents a successful application of growth) to the fitness landscape of the octopod (as it is the case for which the performance difference between growth and no-development has been the smallest, independently of the growth speed). However, this type of analysis is not straightforward for two main reasons:

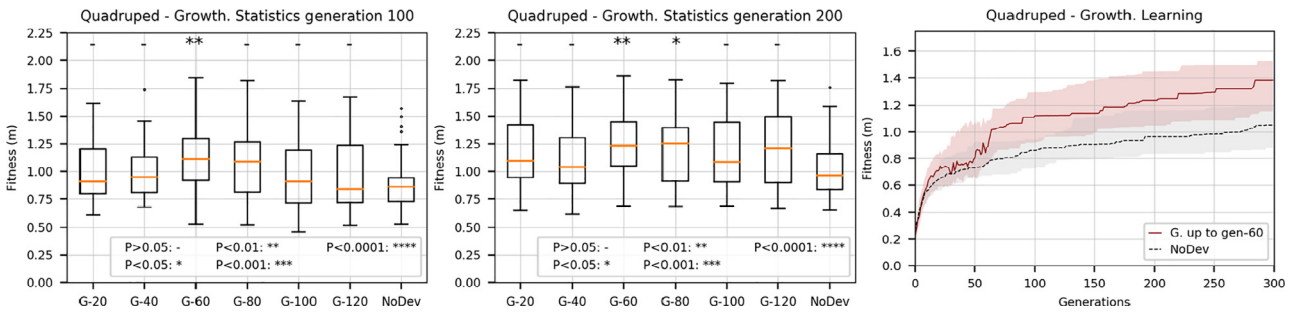

Fig. 6. Statistical analysis for different growth rates at generation 100 (left), generation 200 (middle), and generation 300 (right). The asterisks represent the p-value of the Mann-Whitney test after applying the Bonferroni correction.

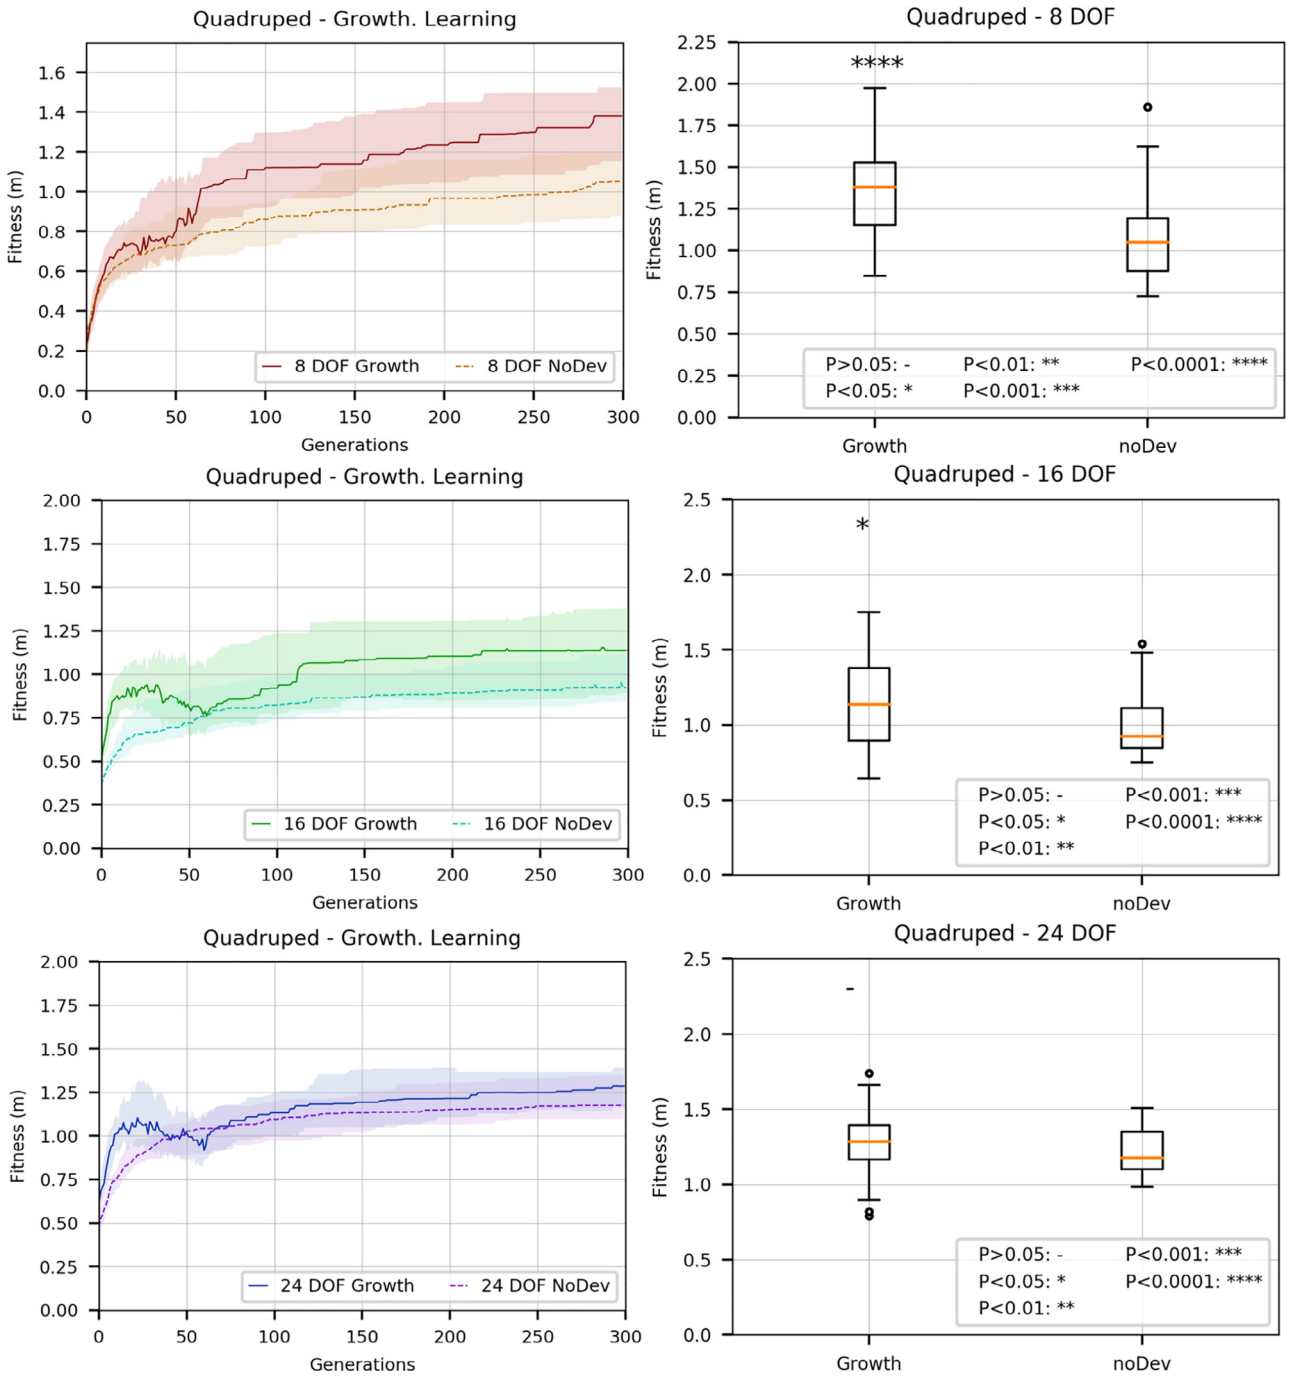

Fig. 7. Results of the learning process (left) and the statistical analysis at the end of learning (right) for the quadruped with 8 DOF (top), the quadruped with 16 DOF (middle) and the quadruped with 24 DOF (bottom). The numerical values of the Mann-Whitney test have been replaced by asterisks to simplify the figure. The larger the number of asterisks, the greater the statistical difference. Due to the comparison of only two samples for each morphology, the Bonferroni correction has not been applied here.

1. The NEAT neuro-evolutionary algorithm optimizes both network weights and the topology of the network. Thus, once the learning process has finished, we can find individuals with different neural network topologies and dimensions, that is, the number of weights in the neural network may vary from one individual to another. This implies that we would need to compare fitness landscapes of different dimensionalities.
2. The high dimensionality of the fitness landscapes obtained for neural networks, which usually involve tens, hundreds or even larger numbers of parameters, and thus, dimensions, makes them very difficult analyze and understand directly.

To overcome these problems and work with manageable fitness landscapes, we have projected the N-dimensional parametric spaces of the resulting neural networks into a two-dimensional space. That is, into a space that is defined by two parameters, in this case P1 and P2, similarly to the technique used by Koos et al. in the transferability approach to minimize the reality gap [40]. A description of the projection process used to perform this transformation is presented in the Appendix (see Figure A1 and A2). It is

[4] [4]https://github.com/GII/morphological_development/tree/main/publications/2021_neurocomputing/videos/fitness_landscapes.

obviously the case that any projection from a high dimensional space onto a lower dimensional space will hide details on the real configuration of the original space, especially when the differences in dimensionality are very large. However, as we are not interested in the precise details of these spaces, but on general features that may be relevant to the evolution of the difficulty of the problem in the different developmental stages, we have found these projections very informative.

To try to extract relevant information on the difficulty and variations of the fitness landscapes along the development process, the controller for each morphology was tested at 11 developmental stages, from the initial morphology to the final one. However, due to space constraints, in the figures for the quadruped, hexapod, and octopod we will only show the fitness landscape of the initial morphology, an intermediate developmental stage and the final one (Fig. 8).

In this figure, we can observe and analyze the differences in the fitness landscapes for the successful example, the quadruped, and the unsuccessful ones, those of the octopod and hexapod in three different developmental stages and how they develop. Thus, comparing the fitness landscape of the quadruped and octopod it can be observed that:

- For the initial morphology, the fitness landscape of the quadruped (Fig. 8 top-left) is smooth and non-deceptive. In this sense, it is possible to see how the fitness value increases gradually and continuously towards the optimum from the borders. However, in the case of the octopod, its initial morphology (Fig. 8 top-right) already presents a rough and deceptive fitness landscape, similar to that of the intermediate morphology and also the final one.

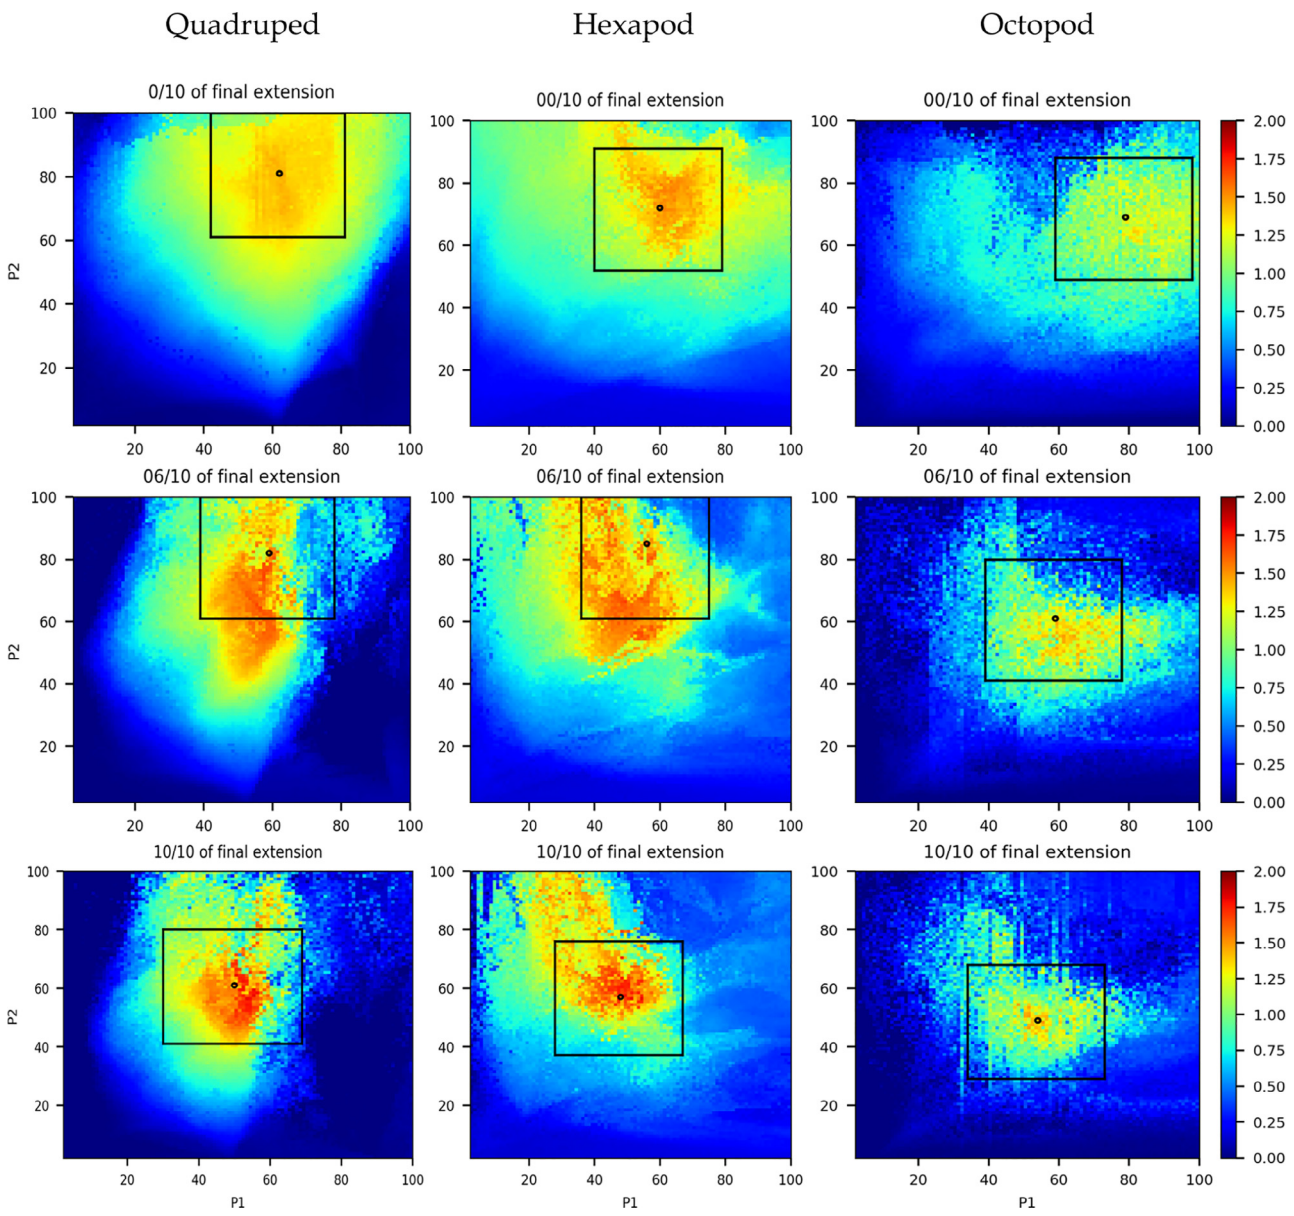

Fig. 8. Fitness landscape of the 8DOF quadruped (left), the hexapod (middle), and the octopod (right) for the initial morphology (top), an intermediate growth stage (middle) and the final morphology (bottom). The color scale indicates the fitness (m). The global optimum for each landscape is marked as a small circle and the squares indicate the area around the optimum used to calculate the roughness.

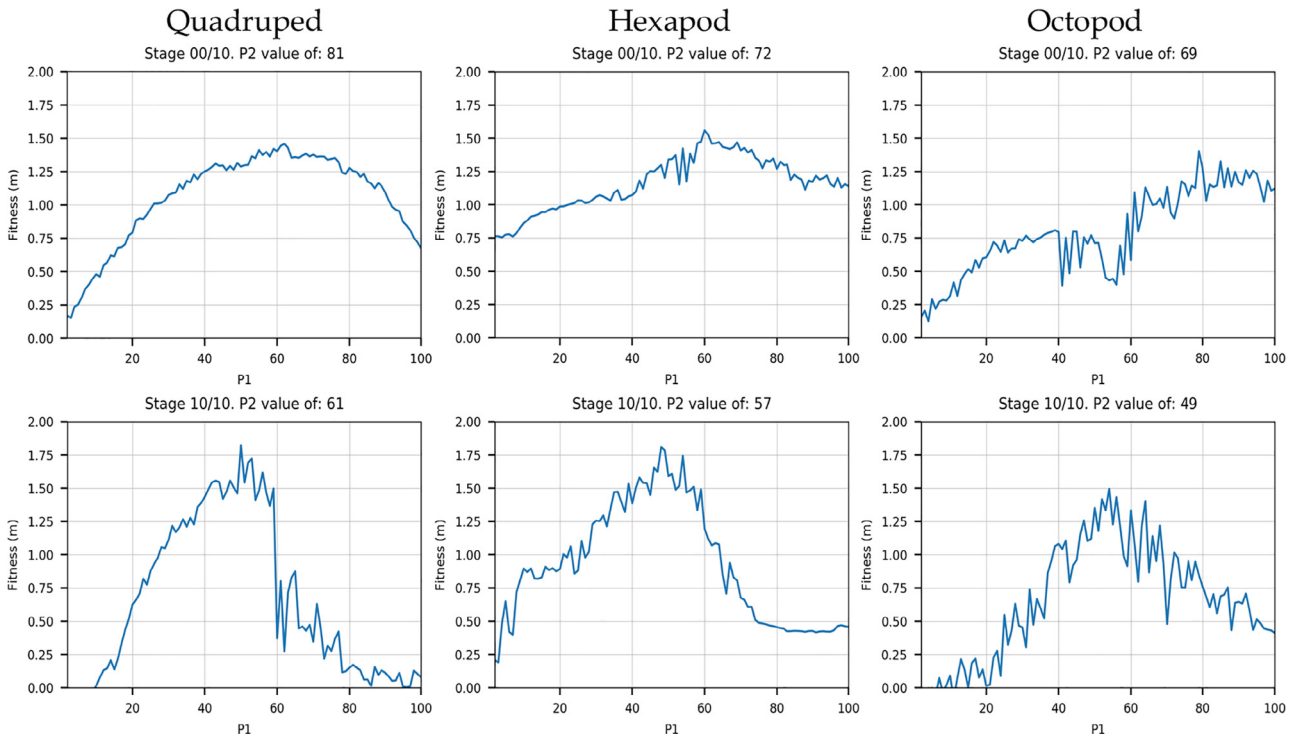

Fig. 9. Representation of the attraction basin area of the quadruped, hexapod, and octopod landscape. These figures represent the 2D representation of the initial (top) and final morphology (bottom) for a fixed P2 value: Left column top: P2 value of 81. Left column bottom: P2 value of 61. Middle column top: P2 value of 72. Middle column bottom: P2 value of 57. Right column top: P2 value 69. Right column bottom: P2 value of 49.

- These characteristics of the fitness landscape change as the morphologies grow. In the case of the quadruped, for the intermediate morphology (Fig. 8 middle-left) the fitness landscape becomes rougher and more deceptive compared to the initial one: there are areas with large color changes between points that are close to one another, indicating the existence of large fitness gradients. These fitness gradients suggest the presence of areas around the optimum that are suboptimal solutions. These areas can confuse the learning algorithm in its path towards the optimum, characteristics that are more relevant in the case of the adult morphology (Fig. 8 bottom-left). This is probably due to the fact that longer legs make the robots more prone to falls and leg interference during walking. For the octopod, although the fitness landscapes are different for each developmental stage, we cannot say that their general characteristics change with the growth of the morphology. That is, the roughness and deceptiveness of the landscapes is present from the initial morphology to the final one (transition from Fig. 8 top-right to Fig. 8 bottom-right). In other words, in terms of general difficulty in finding a solution, there seems to be very little difference between the initial and the final morphology. We suspect that these similarities in the fitness landscape are motivated by the intrinsic stability of the octopod. The effects that growth has over the quadruped (more stability during the early stages of development and increasing dynamicity as growth takes place) are not present in this case due to the large number of legs, making morphological development irrelevant.
- Finally, we can observe a clear difference in the maximum fitness achieved for each morphology at each developmental stage, which is relatively logical considering that the size of the legs increases with development and therefore it is possible to cover a greater distance with the same movement. For the quadruped, the maximum median fitness reached in the initial morphology is 1.45 m, while the maximum median fitness achieved in the final morphology is 1.82 m (a difference of 0.37 m, which implies an increase of 25.5% with respect to the initial fitness achieved). However, we cannot observe such a clear difference in the absolute fitness value achieved in the initial morphology and the final one in the octopod. The area of optimal solutions in the final morphology is indeed larger than in the initial one, but the difference is not as relevant as it is in the quadruped. In this case, the maximum fitness achieved in the initial and final morphology are 1.40 m and 1.49 m respectively, which implies a difference of only of 0.09 m (which is just an increase of 6.4% with respect to the initial fitness achieved).

Regarding the fitness landscape characteristics of the initial, intermediate, and final morphologies in the case of the hexapod, they display characteristics that are between those of the octopod and the quadruped. It is interesting to point out the difference in the fitness value achieved between the 1.56 m of the initial morphology and the 1.80 m achieved by the final one. The difference is 0.24 m which is an intermediate value between the difference obtained in the quadruped (0.37 m) and octopod (0.09 m) and, although it represents a notable increase of 15.4%, the characteristics of the evolution of the fitness landscapes were not enough to provide an advantage when learning during morphological development.

These observations are confirmed producing a slice (constant P2) of the fitness landscape through the optimum solution for the initial morphology (Fig. 9 top) and for the final one (Fig. 9 bottom) for the three morphologies. In this representation, the differences in terms of roughness between the fitness landscapes of the initial and final morphologies can be easily appreciated. In the case of the quadruped, the unidimensional fitness landscape for the initial morphology is characterized by a relatively smooth line, whilst the landscape for the adult morphology is characterized by many local maxima and minima in the fitness value. Especially relevant is the large decrease in the fitness observed near the optimum in

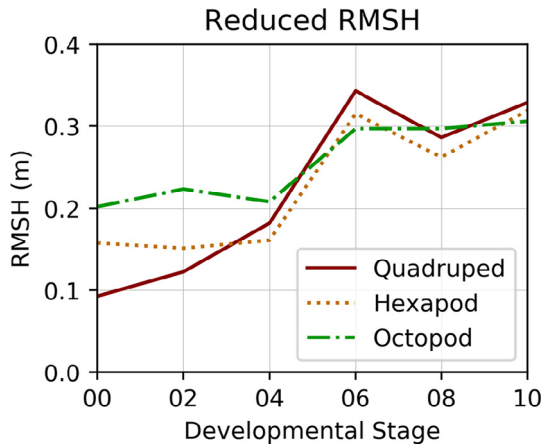

Fig. 10. Representation of the RMSH value for the reduced landscape of Fig. 7 for several developmental stages of the quadruped, hexapod, and octopod.

the final morphology (an example of the different gaits for P1 values of 59 and 60 can be viewed in the repository[4]). Furthermore, as it was shown in Fig. 8, this roughness of the landscape arises in the area close to the optimal solution, not so much in its periphery. This contrasts with the characteristics of the fitness landscape of the octopod. The landscape profile already displays a high level of roughness in the initial morphology with many suboptimal solutions where the learning algorithm can become stuck (Fig. 9, right column). Furthermore, notice the evolution of the position of the optimum throughout development. The position of the optimum in the final morphology corresponds to a minimum in the initial one. In fact, the initial fitness landscape leads the optimization process away from this point, representing a deceptive path for the developmental process in its objective of reaching the optimum of the final morphology. Finally, it can be observed that the hexapod presents graphs that are in between those of the octopod and the quadruped. They display a rougher fitness landscape than the quadruped, but they are not as deceptive as those of the octopod.

These previous statements are supported by a numerical characterization of the landscapes. The different landscapes of Fig. 8 have been characterized by the Root-Mean-Square Height (RMSH) considering a step size of 2 [41,42]. As there are large areas of almost cero fitness in most of the landscapes which increase in size during development, the RMSH is practically constant throughout development (around 0.5, 0.4 and 0.3 for the quadruped, hexapod and octopod respectively). To discard their effect, we have considered an area around the optimum of each landscape (areas marked by a black box in Fig. 8). Then, it can be seen that, although the final RMSH (roughness) is similar for the three final morphologies, the initial RMSH for the hexapod and the octopod are significantly larger than for the quadruped (Fig. 10). Consequently, the RMSH difference between the initial morphology and the final one decreases with the number of legs of the morphology (0.236 for the quadruped, 0.156 for the hexapod, which means a reduction of 33.8% with respect to the quadruped, and 0.104 for the octopod, which means a reduction of 55.9% with respect to the quadruped).

In addition, another factor that must be taken into account in the case of successful morphological development processes, is that during the development of the fitness landscapes their overall structure is preserved. The attraction basin of the optimum in the initial morphology is large. As the morphology increases in size, the attraction basin is reduced becoming rougher and sometimes fragmented. However, it preserves its initial structure. That is, there is initially an area of better solutions and this area does not change throughout the growth process.

7. Discussion

We can observe a relationship between the influence of growth for the three morphologies and how this influence is related to the characteristics of the fitness landscapes along development. These observations allow us to infer a relationship between problem difficulty and the relevance of growth, based on the characteristics of the problem. First of all, we want to clarify that the adult morphology is given by the problem. Thus, a designer of a morphological development process can only control the initial morphologies and how they may lead to the final one through the appropriate configuration of the *MD(t)* function. And, usually, this final morphology will show a rough and deceptive fitness landscape, otherwise it would be easy to learn the controller directly over the adult morphology and there would be no need for morphological development. This is the case of the fitness landscapes of the three adult morphologies we have considered: their final fitness landscapes are rough. Consequently, this is a necessary, but not sufficient condition for morphological development to make sense.

Secondly, from the results obtained, it seems that in successful cases, the initial morphologies display smooth and non-deceptive fitness landscapes with a wide attraction basing towards the optimum that guide the learning algorithm easily to it. Furthermore, the areas of good solutions increase their fitness value as development advances. All of this helps to focus the learning process towards the most promising areas of the solution space in the adult morphology, avoiding deceptive paths during development. This is the case of the fitness landscape of the quadruped initial morphology, but not that of the hexapod nor octopod. Thus, we can say that growth can improve learning by starting on a series of smoother and less deceptive fitness landscapes than the ones of the final morphologies. Now, depending on the morphology one considers and the task to be carried out, smoothing the initial fitness landscape may imply different morphological changes.

Thirdly, in the case of successful experiments, the optimum of the initial fitness landscapes was in the same area of the optimum in the final fitness landscapes or a path of overlapping optimal areas was provided from this initial optimum to that of the final morphology, thus guiding the process towards the final optimum through the fitness landscapes of successive stages of development. This guidance of the learning algorithm happens neither for the hexapod nor for the octopod. On the one hand, the initial fitness landscapes in these cases present higher roughness and deceptiveness characteristics, closer to those of the final morphologies, which means that the initial morphologies that were chosen for the hexapod and the octopod did not provide as much of an advantage in the optimization process. On the other hand, in these cases, as the morphologies developed, this roughness and deceptiveness was present throughout, thus hiding any possible path towards the final optimum. In fact, in the case of the octopod, these fitness landscapes would even guide the learning process away from the optimum.

At this stage, we want to point out an important difference between morphological and cognitive development in the way they handle the learning process. In cognitive development, the first units of knowledge to learn are basic features, usually about motor coordination and sensor processing, that serve as basic knowledge structures to build more complex structures in later stages of the learning process [43]. For example, scaffolding learning [44] is used to learn complex behaviors, like avoiding obstacles, which are based on simpler ones, such as walking in a straight line or steering. However, in each developmental stage of morphological development, the cognitive structures that were created in the previous one constitute the starting point for those that will be applied to the new morphology. These cognitive structures are

modified and adapted to be able to work with the physical changes that happened during the morphological development process. The modifications of the cognitive structures during morphological development produce new ones, which are better adapted to the new morphologies and improve the performance of the interaction with the environment. This leads us to another important observation in the experiments presented in this paper: the growth rate is relevant. That is, in the case of successful morphological development experiments, not all growth rates led to the same results. In fact, we have seen that there are growth rate intervals that produce much better results than others. This makes us think that any successful morphological development strategy must involve matching the capabilities of the learning algorithm for adapting to changing morphologies with the speed at which these morphologies change. As shown in the results, developmental speeds that were too slow or, more importantly, too fast, did not lead to successful results over the quadruped. In other words, morphological development and its influence also seems to be closely related to the capabilities of the learning algorithm used.

8. Conclusions and future work

The objective of this work was to study the influence of morphological development during learning in simple controlled experiments over the task of walking. To achieve this objective, first, a formal definition of morphology as well as of morphological development strategies has been provided. This definition has been established to bring together in one formalism the variety of morphological development strategies found in the literature. Under the umbrella of such definition, a series of experiments based on a growth based morphological development strategy have been carried out. These experiments correspond to a particular implementation of the *MD(t)* function.

We have observed that growth seems to facilitate learning in quadrupeds, producing better results than in the no-development case, something not observed either in the case of hexapods or octopods. However, growth does not seem to be detrimental to performance in any of these morphologies. In fact, the poor results in the case of hexapods and octopods cannot be attributed only to the fact that they present more DOFs, leading to larger search spaces, as in the tests carried out using a quadruped morphology with the same degrees of freedom as the octopod, the growth *MD(t)* improved over the no-development case. Consequently, the morphology itself as related to the task also has an effect.

Furthermore, through the experiments carried out and the study of the succession of fitness landscapes corresponding to the morphological development processes (*MD(t)*) constructed for the different morphologies, we have also observed a series of characteristics that define the influence of growth during learning and that could potentially be used to better design growth related morphological development processes.

As a morphological development process is usually stated in terms of designing a *MD(t)* function that leads to a final morphology (given) that must accomplish a task, we have observed that in successful morphological development processes the *MD(t)* function usually starts with an initial morphology that conforms a smooth and non-deceptive fitness landscape. This landscape changes progressively, as the morphology develops, towards the fitness landscape generated by the final morphology. However, this sequence of fitness landscapes must lead the learning process

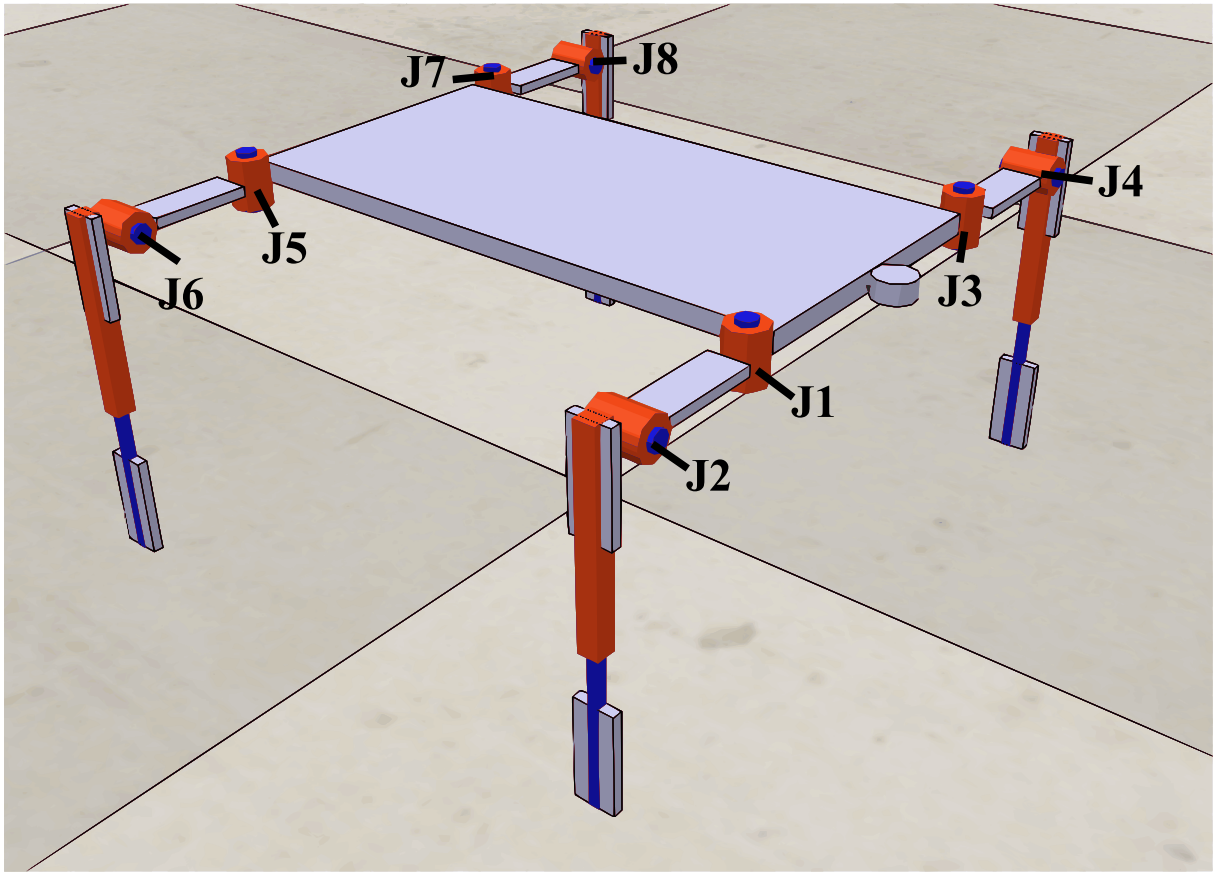

Fig. A1. Labeling of the different joints for the quadruped.

(seen as an optimization problem) towards the optimum of the final morphology, whose fitness landscape is usually rough and deceptive. This implies that the successive fitness landscapes must provide a path of overlapping optimum areas towards that of the final morphology. Additionally, different growth rates lead to different results, some of them successful and others unsuccessful. This hints at the need for a suitable synergy between the growth rate and the adaptation capabilities of the learning algorithm.

These are initial results and, obviously, more work is needed in this area. Thus, we consider that it would be interesting to check whether growth-based morphological development can also be successfully applied over legged robots with morphological configurations different from the ones presented here, as well as to real robots, such as the one described by Nygaard et al. [45]. In addition, we are carrying out experiments using other morphological development strategies such as variation in the range of motion or freeing and freezing degrees of freedom and over different tasks, as well as testing different controllers and learning algorithms, in order to extend and generalize the results presented here.

CRediT authorship contribution statement

M. Naya-Varela: Methodology, Software, Writing – original draft, Data curation, Investigation, Resources, Writing – review & editing. **A. Faina:** Writing – original draft, Formal analysis, Investigation, Supervision, Writing – review & editing. **A. Mallo:** Software, Validation, Data curation, Visualization. **R.J. Duro:** Writing – original draft, Funding acquisition, Investigation, Supervision, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work has been partially funded by the Ministerio de Ciencia, Innovación y Universidades of Spain/FEDER (grant RTI2018-101114-B-I00), Xunta de Galicia (EDC431C-2021/39) and the Centro de Investigación de Galicia "CITIC", funded by Xunta de Galicia and the European Union (European Regional Development Fund-Galicia 2014-2020 Program), by grant ED431G 2019/01. Funding for open access charge: Universidade da Coruña/CISUG. We also want to thank CESGA (Centro de Supercomputación de Galicia. https://www.cesga.es/) for the possibility of using its resources.

Appendix

This appendix describes how the projection of the n-dimensional parameter space of the neural networks used as con-

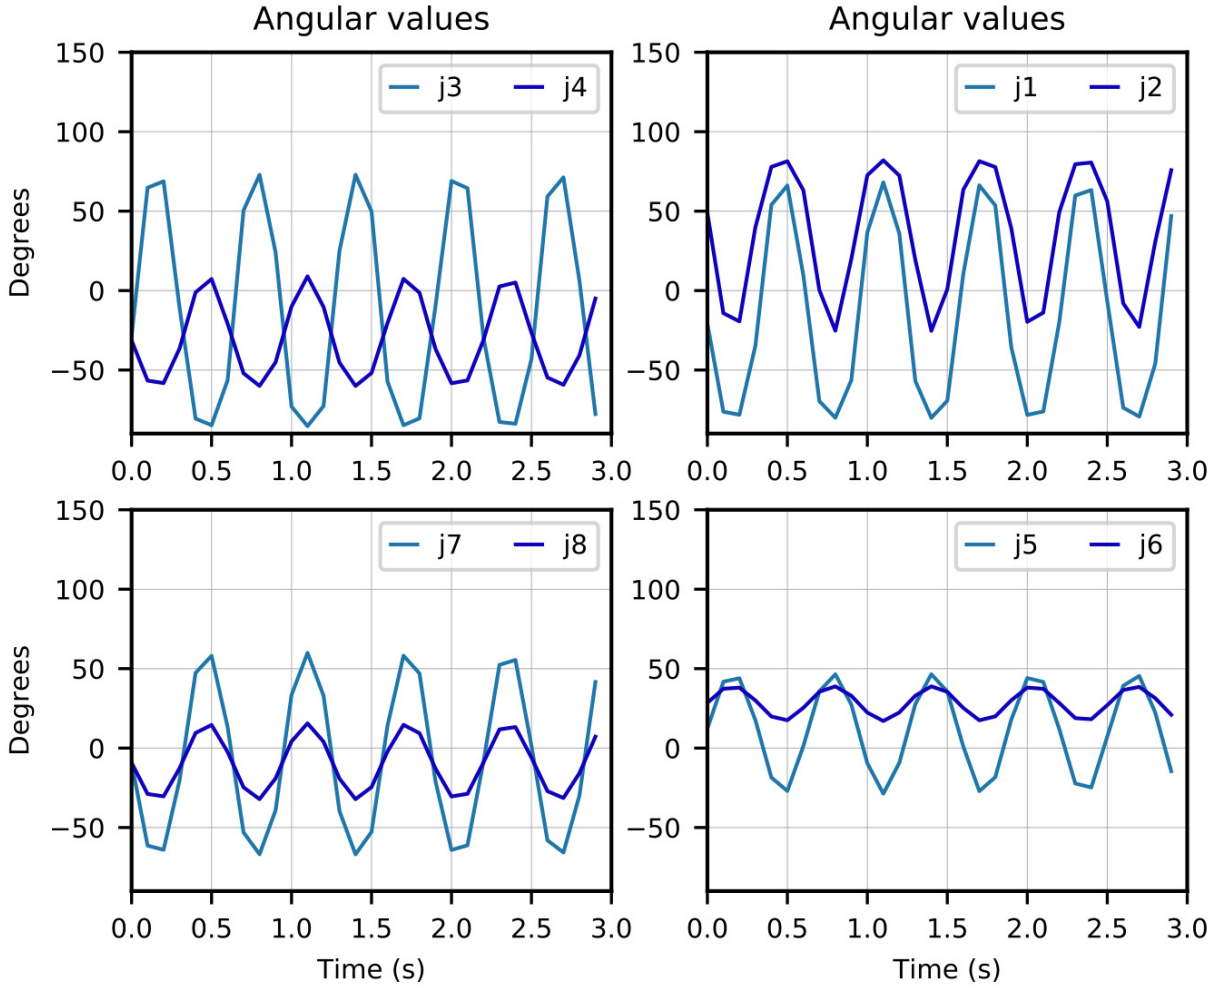

Fig. A2. Waveforms of the outputs of the neural network (target angles) described by each joint during locomotion for the quadruped.

trollers to a two-dimensional parameter space was carried out. This allowed representing the fitness values (f) for the whole fitness space in terms of two parameters (P1 and P2).

The transformation involved the following steps:

1. The best controller for each morphology is selected.
2. This controller is run on the morphology with the objective of extracting the angular positions described by the robot's links during locomotion. These values are obtained from the output of the neural network of the controller. Figure A-I displays the nomenclature of the different joints of the quadruped and Figure A-II displays the representation of these values for the quadruped.
3. A sinusoidal function is fitted to the sequence of values corresponding to each joint. This function has the form of:

$$Angular\ value = A_0 + A * \sin(w * t\ +\ \varphi) \tag{A.1}$$

4. Finally, the joint's equations are modified to make them dependent on two parameters, P1 and P2, each of them has a range of [0,100]. These parameters modify the value of A_0 and the value of A respectively. These functions have the form:

$$Angular\ value = P1 * \left(-\frac{A_0}{50}\right) + P2 * \left(\frac{A}{50}\right) * \sin(\omega * t + \varphi) \tag{A.2}$$

Thus, it is necessary to mention that:

- w is a constant with the same value for all joints.
- A, A_0 and φ are also constants but different for each joint.

As a result, the controller of each morphology depends exclusively on two parameters (P1, P2). At the end of the simulation time, the distance travelled by the quadruped, which is considered the fitness of the controller (f), is recorded. Thus, the fitness landscape can be constructed representing the fitness value F as a function of P1 and P2. Formally:

$$F = f(P1,\ P2) \tag{A.3}$$

This reverse engineering sequence is repeated for each of the morphologies to obtain the corresponding controller and then produce the low dimensional fitness landscape.

References

[1] E. Thelen, Motor development as foundation and future of developmental psychology, Int. J. Behav. Dev. 24 (4) (2000) 385–397.
[2] J. Piaget, M. Cook, The origins of intelligence in children, vol. 8, no. 5. International Universities Press New York, 1952.
[3] J. Law, M. Lee, M. Hülse, A. Tomassetti, The infant development timeline and its application to robot shaping, Adapt. Behav. 19 (5) (2011) 335–358.
[4] A. Cangelosi, M. Schlesinger, Developmental robotics: From babies to robots, MIT Press, 2015.
[5] A. Baranes, P.-Y. Oudeyer, Maturationally-constrained competence-based intrinsically motivated learning, in: *IEEE 9th International Conference on Development and Learning*, 2010, pp. 197–203.
[6] L. Natale, G. Metta, G. Sandini, A developmental approach to grasping, in *Developmental robotics AAAI spring symposium*, 2005, vol. 44.
[7] S. Kriegman, N. Cheney, F. Corucci, J. Bongard, A minimal developmental model can increase evolvability in soft robots, in: *Proceedings of the Genetic and Evolutionary Computation Conference*, 2017, pp. 131–138.
[8] F. Chao, M.H. Lee, J.J. Lee, A developmental algorithm for ocular-motor coordination, Rob. Auton. Syst. 58 (3) (2010) 239–248.
[9] P. Savastano, S. Nolfi, Incremental learning in a 14 DOF simulated iCub robot: Modeling infant reach/grasp development, Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics) 7375 LNAI (2012) 250–261.
[10] J. Bongard, The utility of evolving simulated robot morphology increases with task complexity for object manipulation, Artif. Life 16 (3) (2010) 201–223.
[11] D. Buckingham, J. Bongard, Physical scaffolding accelerates the evolution of robot behavior, Artif. Life 23 (3) (2017) 351–373.
[12] V. Ivanchenko, R.A. Jacobs, A developmental approach aids motor learning, Neural Comput. 15 (9) (2003) 2051–2065.
[13] M. Lungarella, L. Berthouze, Adaptivity via alternate freeing and freezing of degrees of freedom, ICONIP 2002 - Proc. 9th Int. Conf. Neural Inf. Process. Comput. Intell. E-Age, vol. 1, pp. 482–487, 2002.
[14] M.H. Lee, Q. Meng, F. Chao, Developmental learning for autonomous robots, Rob. Auton. Syst. 55 (9) (2007) 750–759.
[15] M. Schlesinger, D. Parisi, J. Langer, Learning to reach by constraining the movement search space, Dev. Sci. 3 (1) (2000) 67–80.
[16] G. Gomez, M. Lungarella, P. Eggenberger-Hotz, Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity, in: *Proceedings of the 4th International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 2004, pp. 119–122.
[17] M.H. Lee, Q. Meng, F. Chao, Staged competence learning in developmental robotics, Adapt. Behav. 15 (3) (2007) 241–255.
[18] J. Bongard, Morphological change in machines accelerates the evolution of robust behavior, Proc. Natl. Acad. Sci. 108 (4) (2011) 1234–1239.
[19] N. Bernstein, The Coordination and Regulation of Movements, Pergamon Press, London: London, 1967.
[20] M. Lungarella, L. Berthouze, On the Interplay Between Morphological, Neural, and Environmental Dynamics: A Robotic Case Study, Adapt. Behav. 10 (3) (2002) 223–241.
[21] L. Berthouze, M. Lungarella, Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom, Adapt. Behav. 12 (1) (2004) 47–64.
[22] V. Vujovic, A. Rosendo, L. Brodbeck, F. Iida, Evolutionary developmental robotics: Improving morphology and control of physical robots, Artif. Life 23 (2) (2017) 169–185.
[23] M. Naya-Varela, A. Faina, R.J. Duro, Morphological Development in robotic learning: A survey, IEEE Trans. Cogn. Dev. Syst. 13 (4) (2021) 750–768.
[24] M. Naya-Varela, A. Faina, and R. J. Duro, "Some Experiments on the influence of Problem Hardness in Morphological Development based Learning of Neural Controllers," in Hybrid Artificial Intelligent Systems. HAIS 2020. Lecture Notes in Computer Science, 2020, pp. 362–373.
[25] G. Lu, J. Li, and X. Yao, "Fitness Landscapes and Problem Difficulty in Evolutionary Algorithms: From Theory to Applications BT - Recent Advances in the Theory and Application of Fitness Landscapes," H. Richter and A. Engelbrecht, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 133–152.
[26] B. Naudts, L. Kallel, A comparison of predictive measures of problem difficulty in evolutionary algorithms, IEEE Trans. Evol. Comput. 4 (1) (2000) 1–15.
[27] P.F. Stadler, "Fitness landscapes", in *Biological evolution and statistical physics*, in: M. Lässig, A. Valleriani (Eds.), Lecture Notes in PhysicsBiological Evolution and Statistical Physics, Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 183–204.
[28] E. Pitzer, M. Affenzeller, A comprehensive survey on fitness landscape analysis, in: Recent advances in intelligent engineering systems, Springer, 2012, pp. 161–191.
[29] P. Caamaño, F. Bellas, J.A. Becerra, R.J. Duro, Evolutionary algorithm characterization in real parameter optimization problems, Appl. Soft Comput. 13 (4) (2013) 1902–1921.
[30] S. Kriegman, N. Cheney, J. Bongard, How morphological development can guide evolution, Sci. Rep. 8 (1) (2018) 1–17.
[31] E. Rohmer, S.P.N. Singh, M. Freese, V-REP: A versatile and scalable robot simulation framework, IEEE/RSJ International Conference on Intelligent Robots and Systems 2013 (2013) 1321–1326.
[32] R.L. Smith, "Open Dynamics Engine." [Online]. Available: https://www.ode.org/.
[33] J. Santos, Á. Campo, Biped locomotion control with evolved adaptive center-crossing continuous time recurrent neural networks, Neurocomputing 86 (2012) 86–96.
[34] A. Ferigo, G. Iacca, E. Medvet, Beyond Body Shape and Brain: Evolving the Sensory Apparatus of Voxel-Based Soft Robots, Appl. Evol. Comput. (2021) 210–226.
[35] P. Chervenski, S. Ryan, "MultiNEAT, project website," URL http//www.multineat.com/, 2012.
[36] K.O. Stanley, R. Miikkulainen, Evolving neural networks through augmenting topologies, Evol. Comput. 10 (2) (2002) 99–127.
[37] CESGA, CESGA. Centro de Supecomputacion de Galicia," 2020. [Online]. Available: http://www.cesga.es/.
[38] P.E. McKnight, J. Najab, Mann-Whitney U Test, Corsini Encycl. Psychol., p. 1, 2010.
[39] H. Abdi, Holm's sequential Bonferroni procedure, Encycl. Res. Des. 1 (8) (2010) 1–8.
[40] S. Koos, J.-B. Mouret, S. Doncieux, The transferability approach: Crossing the reality gap in evolutionary robotics, IEEE Trans. Evol. Comput. 17 (1) (2013) 122–145.
[41] J. Wu, Q. Yang, Y. Li, Partitioning of Terrain Features Based on Roughness, Remote Sens. 10 (12) (2018) 1985.
[42] M.K. Shepard, B.A. Campbell, M.H. Bulmer, T.G. Farr, L.R. Gaddis, J.J. Plaut, The roughness of natural terrain: A planetary and remote sensing perspective, J. Geophys. Res. Planets 106 (E12) (2001) 32777–32795.

[43] F. Bellas, R.J. Duro, A. Faiña, D. Souto, Multilevel darwinist brain (MDB): Artificial evolution in a cognitive architecture for real robots, IEEE Trans. Auton. Ment. Dev. 2 (4) (2010) 340–354.
[44] J. Saunders, C.L. Nehaniv, K. Dautenhahn, Teaching robots by moulding behavior and scaffolding the environment, in: *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, 2006, pp. 118–125.
[45] T.F. Nygaard, C.P. Martin, J. Torresen, K. Glette, Self-Modifying Morphology Experiments with DyRET: Dynamic Robot for Embodied Testing, International Conference on Robotics and Automation (ICRA) 2019 (2019) 9446–9452.

Martin Naya Varela received a B.S. in industrial engineering from the University of A Coruña in 2013 and a Ph.D. in 2021 from the same university. He is currently working as a postdoc at the Integrated Group for Engineering Research, at the University of A Coruña, Spain. His research interests include the development of educational robotics and AI fields, such as developmental robotics, embodied cognition and evolutionary algorithms.

Andrés Faíña received a M.S. degree in Industrial Engineering in 2006 and a Ph.D. in 2011 from the University of A Coruña, Spain. He is currently working as an Assistant Professor in the Department of Computer Science and is a member of the Robotics, Evolution and Art Lab (REAL) at the IT University of Copenhagen, Denmark. His research interests include modular and self-reconfigurable robots, evolutionary robotics and electronic and mechanical design.

Alma Mallo Casdelo received a B.S. in Computer Engineering from the University of A Coruña in 2002, a M.S. in Free Software from the Open University of Catalonia in 2014 and a M.S. in Computer Engineering from the same university in 2021. She is currently working as an associate lecturer at the University of A Coruña and as a software engineer at MINT. Her research interests include robotics, machine learning and computer vision.

Richard J. Duro received a M.S. degree in Physics from the University of Santiago de Compostela, Spain, in 1989, and a PhD in Physics from the same University in 1992. He is currently a *Full Professor* in the Department of Computer Science and head of the Integrated Group for Engineering Research at the University of A Coruña. His research interests include higher order neural network structures, signal processing and autonomous and evolutionary robotics.