Advances in Knowledge-Based Systems, Data Science, and Cybersecurity Received 14-05-2026; Accepted 11-08-2026; Published 19-08-2026

Leveraging High Morphology Rates in Body-Brain Co-Evolution: A Case Study on 2D Robot Locomotion

Paolo Pagliuca paolo.pagliuca@istc.cnr.it
Institute of Cognitive Sciences and Technologies (ISTC)
National Research Council (CNR)
Via Giandomenico Romagnosi 18A, 00196, Roma, Italy

Corresponding Author: Paolo Pagliuca

Abstract

Body-brain co-evolution concerns the simultaneous optimization of body and brain to promote the discovery of optimized solutions. A key benefit of this approach lies in its adaptability, since it alleviates the burden of identifying suitable morphologies for a given problem. A widespread application of body-brain co-evolution regards robot locomotion, in which enabling morphological evolution provides notable advantages over using fixed (and potentially sub-optimal) morphologies. However, assessing the impact of morphological rate on the final performance is often overlooked. In this work, we delved into the analysis of how big the morphological rate should be in order to foster a significant performance enhancement. To this end, we performed an investigation on three 2D robot locomotion problems, BipedalWalker, Embryo and Halfcheetah2D, by considering both simple and challenging environmental conditions. To co-evolve body and brain, we employed the OpenAI Evolutionary Strategy (OpenAI-ES) and the Generational Genetic Algorithm (GGA). Our analysis indicates that, regardless of the considered algorithm, body-brain co-evolution is remarkably more effective than using fixed morphologies, and the advantage increases as the morphological rate becomes bigger.

Keywords: Body-Brain Co-Evolution, Robot Locomotion, BipedalWalker, Embryo, Halfcheetah2D, OpenAI-ES, GGA

1. INTRODUCTION

Body-brain co-evolution involves evolving simultaneously body and brain in order to find out effective solutions for a given problem. This idea comes from biology, in which living organisms undergo morphological and neural changes during their lives to cope with ecological and survival pressures [1, 2, 3]. A key advantage of body-brain co-evolution is to alleviate the burden of identifying suitable morphologies for a specific task.

Research in body-brain co-evolution saw a growing interest in the last three decades. A pioneering work was proposed in [4], in which virtual creatures underwent morphological and neural modifications during competition with other creatures. The author showed that effective morphologies

Citation: Paolo Pagliuca. Leveraging High Morphology Rates in Body-Brain Co-Evolution: A Case Study on 2D Robot Locomotion. **Advances in Knowledge-Based Systems, Data Science, and** ;Research. 2026;3(2).27.

Page no.-533-555.

and strategies emerged as a result of the competition. After this seminal work, other researchers strove to add knowledge to the field. For example, the author in [5] used Genetic Programming (GP) [6] and Genetic Algorithms (GAs) [7] to co-evolve, respectively, body and brain of a robot that must perform an obstacle-avoidance task. In [8], the author tackled body-brain co-evolution by using Novelty Search (NS) [9, 10]: results in two problems indicate that NS is a viable approach for body-brain co-evolution. In [11], the authors introduced the Neural Cellular Robot Substrate (NCRS) framework in which a robot undergoes morphological and neural development during its lifetime. The authors proved the efficacy of the approach on three different tasks. The authors in [12] demonstrated that body–brain co-evolution allows to optimize a two-wheeled differential drive robot that must solve maze navigation tasks. In [13, 14, 15], Lamarckian learning [16] has been employed to promote successful body-brain co-evolution in robot locomotion. In [17], the authors demonstrated that co-evolving body and brain in a real self-reconfiguring quadruped robot led to the development of high-performing and different solutions. A body of research showed the effectiveness of body-brain co-evolution in soft robots [18, 19, 20, 21, 22]. For a review of the effects of learning on the performance enhancement in body-brain co-evolution, the readers are referred to [23].

Among the existing works in literature, the investigation of body-brain co-evolution in robot locomotion tasks is worth mentioning. In [24], the authors analyzed the effects of co-evolving morphological and neural parameters in the Pybullet Halfcheetah and Walker2D tasks [25]. Their outcomes revealed that not only body-brain co-evolution is remarkably more effective than using a fixed morphology, but that the advantage is related to the mutual influence of the two optimization processes. In fact, the authors proved that evolving robot locomotion by using a pre-evolved body morphology under-performs compared to body-brain co-evolution. Similarly, the work in [26] demonstrated that body-brain co-evolution enables to better cope with varying environmental conditions [17, 27, 28]. However, a main limitation of these studies is the lack of any analyses about how big the morphological rate (i.e., the rate at which morphological parameters can be modified) should be in order to generate a notable performance enhancement.

In this research, we investigated the role of morphological rate in three 2D robot locomotion problems, BipedalWalker [29, 30], Embryo [31, 32] and Halfcheetah2D [26]. Furthermore, with regard to BipedalWalker and Halfcheetah2D, we considered two environmental scenarios — basic and hardcore — to analyze whether or not body-brain co-evolution provides significant advantages. To co-evolve body and brain, we employed the OpenAI Evolutionary Strategy (OpenAI-ES) [33], since it represents a widespread approach in robot locomotion [33, 34], and the Generational Genetic Algorithm [35], which is a popular method in Evolutionary Robotics (ER).

The results achieved in these three domains confirm that body-brain co-evolution outperforms the usage of fixed morphologies, thus aligning with previous works [24, 26]. Furthermore, we observed that the advantage becomes significant only when using morphological rates higher than 0.5, regardless of the algorithm used. Interestingly, with respect to OpenAI-ES, the benefit is higher for the BipedalWalker task than for the Halfcheetah problem, whereas the advantage is similar for GGA. Regarding Embryo, the impact of different morphological rates is not significant.

Our work contributes to shed light on identifying how body-brain co-evolution might foster the development of more adaptive and robust autonomous systems, with potential implications in real-world domains.

2. MATERIALS AND METHODS

In this section, we illustrate the 2D robot locomotion tasks (Subsection 2.1), the body-brain co-evolution strategy (Subsection 2.2), the OpenAI-ES algorithm (Subsection 2.3) and the experimental settings (Subsection 2.4) used in this work.

2.1 2D Robot Locomotion Problems

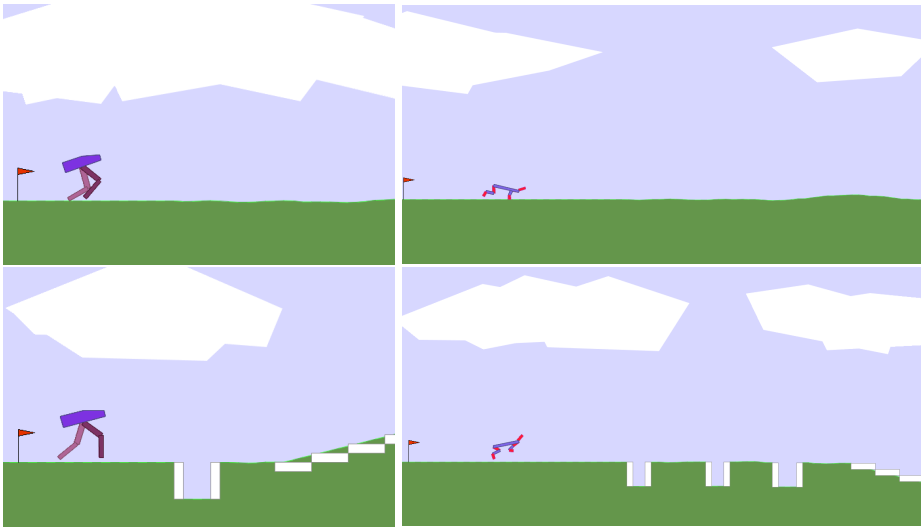

Figure 1: Illustration of the 2D robot locomotion problems: BipedalWalker (left side) and Halfcheetah2D (right side). Top figures refer to the basic environment, whereas bottom figures refer to the hardcore environment.

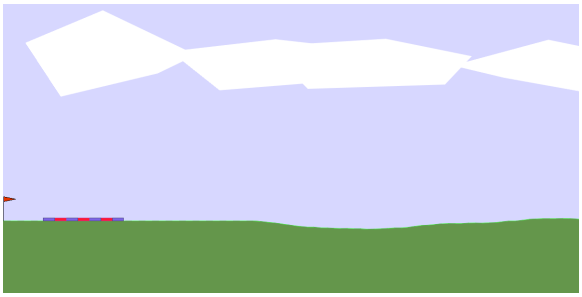

Figure 2: Illustration of the Embryo locomotion task.

As case studies, we considered the BipedalWalker [29, 30], Embryo [31, 32] and Halfcheetah2D [26] problems, which are illustrated in FIGURE 1 and FIGURE 2.

Specifically, the BipedalWalker task involves a two-legged robot with a hull that must locomote toward a target destination located at 300 meters away from the starting position (red flag in FIG-

URE 1). The task represents a benchmark in the Evolutionary Computation (EC) community [36, 37, 38].

The performance of the solutions discovered for the BipedalWalker task is computed according to Eq. 1:

$$F = \frac{1}{N_{eps}} \frac{1}{N_{steps}} \sum_{e=1}^{N_{eps}} \sum_{i=1}^{N_{steps}} (\Delta p_i(e) - 5.0 \times |h_i(e)| - 0.028 \times clip(|a_i(e)|)) \quad (1)$$

where $\Delta p_i(e)$ is the robot's progress at the $i-th$ step of the $e-th$ episode, $h_i(e)$ denotes the hull position (height) at the $i-th$ step of the $e-th$ episode, $a_i(e)$ indicates the average size of actions performed at the $i-th$ step of the $e-th$ episode, $clip()$ is a function keeping its argument within the range $[0, 1]$, N_{steps} represents the number of episode steps, and N_{eps} refers to the number of episodes. The coefficients 5.0 and 0.028 are defined as in the Gymnasium engine [30].

The Embryo problem is the 2D variant of the task introduced in [24], in which a snake-like robot made of 7 segments should move toward a target location placed at 300 meters away. Differently from other tasks, given the initial configuration of the robot (see FIGURE 2), the evolution of neural parameters only does not enable the robot to locomote. Therefore, co-evolving body and brain is pivotal.

The performance of the discovered solutions is defined as in Eq. 2:

$$F = \frac{1}{N_{eps}} \frac{1}{N_{steps}} \sum_{e=1}^{N_{eps}} \sum_{i=1}^{N_{steps}} \Delta p_i(e) \quad (2)$$

Concerning the Halfcheetah2D problem, it represents the 2D variant of the Pybullet Halfcheetah task [25]: a cheetah robot has to locomote toward a target destination located at 300 meters. The task has been introduced in [26] and used to compare different Evolutionary Algorithms (EAs) in [39].

Solutions to this problem are evaluated by using the fitness function defined in Eq. 3.

$$F = \frac{1}{N_{eps}} \frac{1}{N_{steps}} \sum_{e=1}^{N_{eps}} \sum_{i=1}^{N_{steps}} (\Delta p_i(e) - 0.1 \times NJL_i(e)] \quad (3)$$

The symbol $NJL_i(e)$ in Eq. 3 indicates the number of joints at limit at the $i-th$ step of the $e-th$ episode. This penalty has been introduced to prevent unrealistic behaviors.

To evaluate the influence of body-brain co-evolution on the final performance, we considered two environmental scenarios for the BipedalWalker and Halfcheetah2D tasks: in the *basic* environment, robots are exposed to only limited variation of the terrain slope (see FIGURE 1, top figures). On the other hand, in the *hardcore* environment, the robot has to deal with challenges like holes, stairs

and obstacles (see FIGURE 1, bottom figures) that make the task remarkably more difficult. With respect to the Embryo problem, our analysis is limited to the basic environment (see FIGURE 2).

2.2 Body-Brain Co-Evolution Strategy

The morphological parameters are evolved by using the transformations provided in TABLE 1, most of which are derived from [26]. As can be seen, we adopted the same transformations for body size and thickness (see TABLE 1). Instead, concerning the angle range, we used two different transformations depending on the task. Specifically, for Embryo and Halfcheetah2D, we kept the same setting as in [26]. Conversely, we adopted a slightly different transformation for BipedalWalker, which guarantees that the robot configuration does not produce unrealistic behaviors. Furthermore, to prevent the identification and usage of unrealistic morphologies, we added a constraint check on body sizes and thicknesses, which cannot be inferior to 0.01.

Table 1: Transformations employed to modify body parameters in the BipedalWalker, Embryo and Halfcheetah2D problems. The symbol μ denotes the morphological rate, while p represents a generic evolutionary parameter.

Morphological Parameter	BipedalWalker	Embryo	Halfcheetah2D
Body size (s_b)	$s_b \times (1 + (\tanh(p) \times \mu))$	$s_b \times (1 + (\tanh(p) \times \mu))$	$s_b \times (1 + (\tanh(p) \times \mu))$
Body thickness (t_b)	$t_b \times (1 + (\tanh(p) \times \mu))$	$t_b \times (1 + (\tanh(p) \times \mu))$	$t_b \times (1 + (\tanh(p) \times \mu))$
Angle range (r_a)	$r_a \times \left(1.0 + \frac{\tanh(p) \times \mu}{2}\right)$	$r_a \times \mathrm{hardtanh}(p)$	$r_a \times \mathrm{hardtanh}(p)$
Hull size (s_h)	$s_h \times 2^{\tanh(p)}$	-	-

To assess how much the morphological rate μ affects the final performance, we considered the following set of values for μ: $\mu \in [0.0, 0.25, 0.5, 0.75, 1.0]$. This allowed us to analyze the evolution of brain parameters only (i.e., $\mu = 0.0$) and different body-brain co-evolution strategies with progressively increasing values of μ. Because Embryo requires evolving morphological parameters, for this task we considered $\mu \in [0.25, 0.5, 0.75, 1.0]$.

2.3 Evolutionary Algorithms

To align this investigation with the works in [24, 26], our solutions have been evolved by using the OpenAI Evolutionary Strategy (OpenAI-ES) [33]. This algorithm represents one of the most sophisticated techniques to evolve solutions in a broad spectrum of problems, with notable applications in domains like robot locomotion [24, 33, 34], collective aggregation and foraging [34, 40, 41, 42], pole balancing [34, 39], Atari games [33, 43], Deep Learning (DL) applications [44, 45] and Multi-Objective Optimization (MOO) scenarios [46, 47, 48].

The pseudocode of OpenAI-ES is shown in Alg. 1: the algorithm attempts to improve a set of parameters θ referred to as "centroid". At each iteration (line 2), the algorithm extracts λ samples from a Gaussian distribution (line 4) and evaluates pairs of symmetric solutions obtained by adding or subtracting each sample multiplied by a mutation rate σ (lines 5-8). This technique is known as "mirrored sampling" [49]. Once the set of solutions has been fully evaluated, a ranking procedure based on the obtained fitness is performed (line 10), and the gradient ∇_θ is computed

Algorithm 1 *OpenAI-ES*

1: Initialize: σ, λ, θ, F, N_{evals}, $opt()$
2: **for** $eval \leftarrow 0$ to N_{evals} **do**
3: **for** $i \leftarrow 0$ to λ **do**
4: draw sample: $s_i \sim \mathcal{N}(0, \mathbb{I})$
5: ${z_i}^+ \leftarrow \theta + \sigma \times s_i$
6: evaluate the fitness: $F({z_i}^+)$
7: ${z_i}^- \leftarrow \theta - \sigma \times s_i$
8: evaluate the fitness: $F({z_i}^-)$
9: **end for**
10: sort (s_i, z_i) w.r.t. $F(z_i)$ and compute normalized utilities u_i, $u_i \in [-0.5, 0.5]$
11: compute gradient ∇_θ:
$\nabla_\theta \leftarrow \frac{1}{\lambda} \sum_{i=1}^{\lambda} u_i \times s_i$
12: update θ:
$\theta \leftarrow \theta + opt(\nabla_\theta)$
13: **end for**

(line 11). Lastly, parameters θ are updated by the optimizer $opt()$ (line 12). In the canonical version of OpenAI-ES, Adam [50] is employed to compute the parameter update based on historical information about performance enhancements. To keep track of how fitness improved through past evolutionary steps, the algorithm stores two momentum vectors, which are updated within the Adam optimizer.

To broaden the analysis of the impact of body-brain co-evolution, we also considered the Generational Genetic Algorithm (GGA) [35], whose pseudocode is provided in Alg. 2. The algorithm evolves a population Θ of parameters. At each iteration, GGA evaluates the performance of each set of parameters θ (line 4). Once the set of solutions has been fully evaluated, a ranking procedure based on the obtained fitness is performed (line 6), which allows to determine the best N_{bests} solutions (line 7). Finally, the population is updated by generating N_{vars} offspring for each best solution through mutations (line 8). It is worth noting that the best N_{bests} solutions are retained within the population (line 8).

Algorithm 2 *GGA*

1: Initialize: σ, λ, $\Theta = \{\theta_i\}_{i=1,\ldots,\lambda}$, F, N_{evals}, N_{bests}, N_{vars}
2: **for** $eval \leftarrow 0$ to N_{evals} **do**
3: **for** $i \leftarrow 0$ to λ **do**
4: evaluate the fitness: $F(\theta_i)$
5: **end for**
6: sort θ_i w.r.t. $F(\theta_i)$
7: select N_{bests} individuals based on F
8: update Θ:
$\Theta \leftarrow \{N_{bests}\} \cup mutate(N_{bests}, N_{vars}, \sigma)$
9: **end for**

2.4 Experimental Settings

For the simulations reported here, we used the experimental settings provided in TABLE 2 and TABLE 3. In particular, we performed 20 replications of the experiments (random seeds: $1 \rightarrow 20$) to ensure replicability and statistical significance of our analyses, and comply with computational limitations. Regarding OpenAI-ES, parameters have been derived from similar works in literature [24, 26, 51]. Instead, for GGA we define the parameters so that they are as similar as possible to those used for OpenAI-ES. Clearly, the hyperparameter settings represent a crucial factor for determining the actual performance of an algorithm. However, our design choices enable fair comparisons with existing studies.

Table 2: List of parameters used in the simulations with OpenAI-ES. The values are inherited from previous studies [26, 51].

Parameter	BipedalWalker	Embryo	Halfcheetah2D
N_{reps}	20	20	20
N_{evals}	5×10^7	5×10^7	5×10^7
λ	20	20	20
σ	0.02	0.02	0.02
N_{eps}	3	3	3
N_{steps}	500	500	500
μ	$[0.0, 0.25, 0.5, 0.75, 1.0]$	$[0.25, 0.5, 0.75, 1.0]$	$[0.0, 0.25, 0.5, 0.75, 1.0]$

Table 3: List of parameters used in the simulations with GGA. The values have been chosen to match the settings made for OpenAI-ES.

Parameter	BipedalWalker	Embryo	Halfcheetah2D
N_{reps}	20	20	20
N_{evals}	5×10^7	5×10^7	5×10^7
λ	40	40	40
σ	0.02	0.02	0.02
N_{bests}	4	4	4
N_{vars}	9	9	9
N_{eps}	3	3	3
N_{steps}	500	500	500
μ	$[0.0, 0.25, 0.5, 0.75, 1.0]$	$[0.25, 0.5, 0.75, 1.0]$	$[0.0, 0.25, 0.5, 0.75, 1.0]$

The robot controller is a Multi-Layer Perceptron (MLP) [52] consisting of a single internal layer with 50 neurons. The number of inputs/outputs of the MLP is 24/4, 23/6 and 26/6 for the BipedalWalker, Embryo and Halfcheetah2D, respectively. Regarding BipedalWalker, inputs to the MLP include the angle, linear velocity and angular velocity of the hull, the angle and speed of the joints, a flag activating if the leg(s) touches (touch) the ground and the Lidar sensory readings (see FIGURE 1, left side). Conversely, for the Embryo, inputs ot the MLP are the angle, linear velocity and angular velocity of the torso (i.e., the central body segment of the robot), the angle and speed of the joints, and a flag activating if the body segment(s) touches (touch) the ground. Lastly, for the Halfcheetah2D, inputs to the MLP encode the angle, linear velocity and angular velocity of the torso (i.e., the segment

connecting the legs and the head of the robot, see FIGURE 1, right side), the angle and speed of the joints, and a flag activating if the leg(s) touches (touch) the ground. For the three tasks, the outputs to the MLP are converted into joint torques determining the robot behavior.

We used evorobotpy3 [51], a simulation engine providing all the algorithms and tasks required for our investigation. This tool represents a state-of-the-art simulator for robotics [26, 53, 54] and is available at `https://github.com/PaoloP84/evorobotpy3`.

All the experiments have been performed on a Intel(R) Xeon(R) Gold 6252N CPU, with 1TB RAM and Ubuntu 20.04 Operating System. As regards the main libraries, we used Python version 3.8, Gymnasium version 0.29.0 and Box2D version 2.3.5.

3. RESULTS

In this section, we provide the results of our experimentation. To ease the readability of the section, we used the notations "BipedalWalkerHardcore" and "Halfcheetah2DHardcore" to represent the data obtained in the hardcore environmental scenario. For the statistical analysis, we used the Kruskal-Wallis H test [55] with the Dunn post-hoc test [56], and we considered significance at $p < 0.01$. Moreover, we reported the effect size for the Kruskal-Wallis H test, which is defined as in Eq. 4.

$$\eta_H^2 = \frac{H - k + 1}{n - k} \tag{4}$$

In Eq. 4, H indicates the Kruskal-Wallis test statistic, k denotes the number of groups (i.e., performance vectors) compared and n represents the size of each group.

With respect to correlation between variables, we adopted the Spearman correlation coefficient [57]. The information about the time required to run the simulations is illustrated in TABLE 4. As can be seen, GGA is slightly faster than OpenAI-ES.

Table 4: Execution time (in seconds) of the experiments. Data (shown as $avg \pm std$) are averaged over 20 replications.

Algorithm	μ	**BipedalWalker**	**BipedalWalkerHardcore**	**Embryo**	**Halfcheetah2D**	**Halfcheetah2DHardcore**
OpenAI-ES	0.0	$19,356.250 \pm 816.608$	$23,146.750 \pm 747.753$	-	$17,307.100 \pm 624.788$	$20,471.150 \pm 601.348$
	0.25	$19,881.900 \pm 943.850$	$23,128.400 \pm 767.604$	$13,707.100 \pm 511.964$	$17,017.050 \pm 573.911$	$19,867.200 \pm 688.472$
	0.5	$19,793.200 \pm 908.068$	$23,004.500 \pm 727.076$	$13,794.600 \pm 653.008$	$17,353.550 \pm 641.813$	$20,302.900 \pm 662.821$
	0.75	$19,491.650 \pm 831.364$	$22,701.450 \pm 747.894$	$13,917.450 \pm 604.007$	$17,596.050 \pm 822.575$	$20,530.450 \pm 694.080$
	1.0	$19,272.500 \pm 814.941$	$22,976.650 \pm 791.576$	$13,936.200 \pm 475.837$	$17,349.400 \pm 693.968$	$20,137.300 \pm 728.462$
GGA	0.0	$15,790.800 \pm 827.244$	$17,538.300 \pm 884.733$	-	$17,240.850 \pm 1,566.381$	$18,871.350 \pm 2,654.363$
	0.25	$15,930.750 \pm 1,128.944$	$17,577.300 \pm 1,003.880$	$11,157.850 \pm 410.402$	$15,636.450 \pm 1,078.950$	$16,502.350 \pm 748.480$
	0.5	$15,670.800 \pm 1,171.853$	$16,826.400 \pm 762.248$	$11,099.250 \pm 285.789$	$15,359.900 \pm 841.633$	$15,604.300 \pm 1,279.403$
	0.75	$15,231.900 \pm 804.460$	$16,627.750 \pm 762.340$	$11,189.700 \pm 364.799$	$17,596.050 \pm 822.575$	$16,631.950 \pm 896.750$
	1.0	$15,571.650 \pm 1,571.213$	$15,553.650 \pm 2,279.922$	$11,011.950 \pm 344.875$	$13,494.750 \pm 1,033.705$	$14,354.250 \pm 985.365$

TABLE 5 illustrates the final performance of the controllers evolved with OpenAI-ES by systematically varying the morphological rate μ in the BipedalWalker and Halfcheetah2D tasks. As can be seen, the condition $\mu = 1.0$ achieves the best fitness at the end of the evolutionary process, regardless

of both the problem (BipedalWalker or Halfcheetah2D) and the enviromental scenario (basic or hardcore). The performance distribution is also illustrated in FIGURE 3, whereas FIGURE 4 illustrates the medians and 95% bootstrapped confidence intervals of the performance. A noteworthy outcome concerns the low standard deviation achieved in the different tasks and scenarios, which indicates that OpenAI-ES generally converges toward a specific range of fitness values.

Table 5: Results of the experiments performed with OpenAI-ES in the BipedalWalker and Halfcheetah2D, for both basic and hardcore scenarios. Performance (shown as $avg \pm std$) is averaged over 20 replications of the experiments. Highest performance is denoted in bold.

μ	BipedalWalker	BipedalWalkerHardcore	Halfcheetah2D	Halfcheetah2DHardcore
0.0	348.703 ± 2.640	171.392 ± 28.833	316.645 ± 38.691	75.779 ± 13.567
0.25	348.632 ± 3.632	192.564 ± 41.560	333.576 ± 9.157	85.329 ± 18.708
0.5	350.633 ± 2.106	233.946 ± 40.926	336.938 ± 7.289	88.372 ± 18.351
0.75	353.647 ± 0.995	240.734 ± 40.134	335.470 ± 14.175	89.136 ± 17.705
1.0	**355.885 ± 1.530**	**278.373 ± 41.363**	**341.013 ± 9.695**	**92.791 ± 20.727**

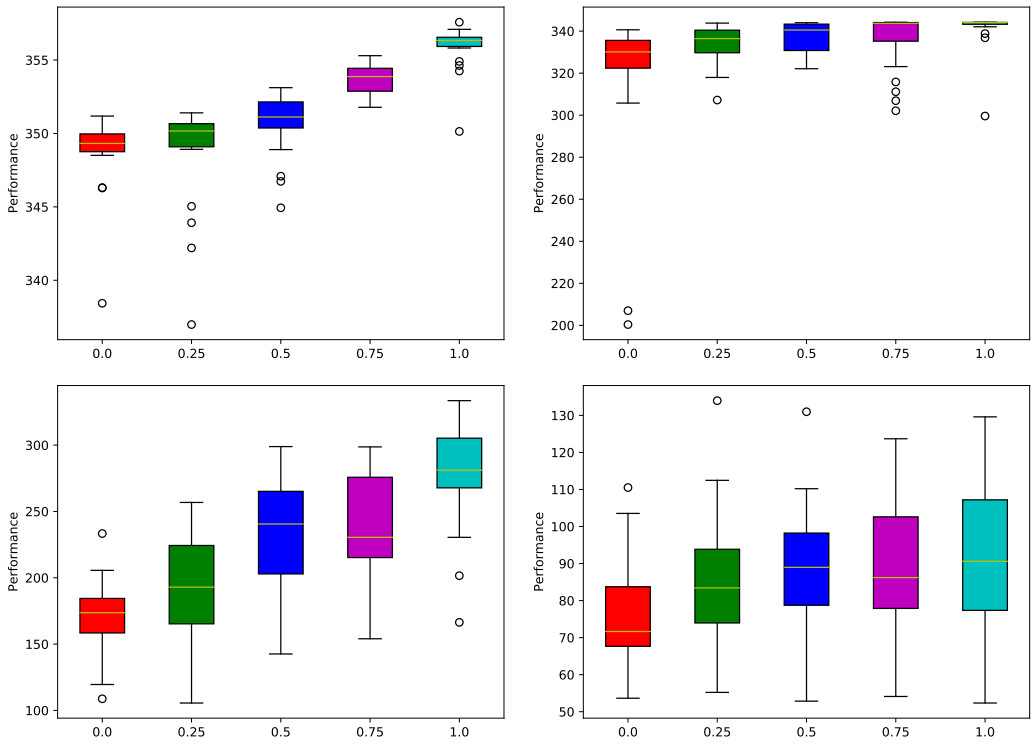

Figure 3: Performance achieved by OpenAI-ES by varying the morphological rate with regard to BipedalWalker (left side) and Halfcheetah2D (right side). Top figures refer to the basic environment, whereas bottom figures refer to the hardcore environment.

In statistical terms, we found that the conditions $\mu = 0.75$ and $\mu = 1.0$ significantly outperform the conditions $\mu = 0.0$ and $\mu = 0.25$ in the BipedalWalker problem ($p < 0.01$, see FIGURE 5). Moreover, the condition $\mu = 1.0$ is also significantly better than $\mu = 0.5$ in this task ($p < 0.01$).

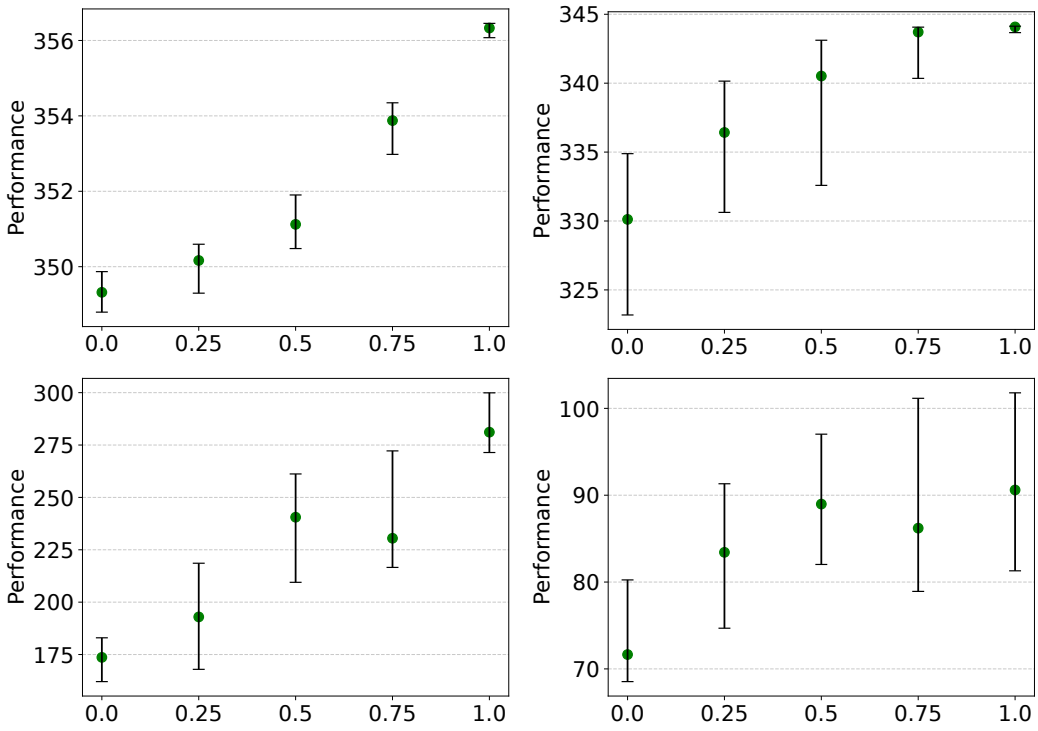

Figure 4: Medians (green circles) and 95% bootstrapped confidence intervals of the performance achieved by OpenAI-ES by varying the morphological rate. Data are averaged over 20 replications. Left side figures refer to BipedalWalker, while right side plots concern Halfcheetah2D. Top figures refer to the basic environment, whereas bottom figures refer to the hardcore environment.

Similarly, the conditions $\mu = 0.5$, $\mu = 0.75$ and $\mu = 1.0$ are remarkably superior to the condition $\mu = 0.0$ in the BipedalWalkerHardcore task ($p < 0.01$, see FIGURE 5), and the condition $\mu = 1.0$ is better than $\mu = 0.25$ ($p < 0.01$). On the other hand, the conditions $\mu = 0.0$ and $\mu = 0.25$ are statistically equivalent ($p > 0.01$). Overall, this implies that body-brain co-evolution enhances performance when $\mu \geq 0.5$, while a low morphological rate (i.e., $\mu = 0.25$) does not provide any advantages over using a fixed morphology.

With regard to Halfcheetah2D, our results indicated that the conditions $\mu = 0.75$ and $\mu = 1.0$ significantly outperform the condition $\mu = 0.0$ in the basic scenario ($p < 0.01$, see FIGURE 5). Furthermore, the condition $\mu = 1.0$ is also superior to the condition $\mu = 0.25$ ($p < 0.01$). Conversely, there is no difference between the other conditions ($p > 0.01$ for each pair of comparisons). If we consider the Halfcheetah2DHardcore task, we did not find any statistical differences between the different values of μ ($p > 0.01$, see FIGURE 5). For both tasks, the effect sizes are provided in TABLE 6.

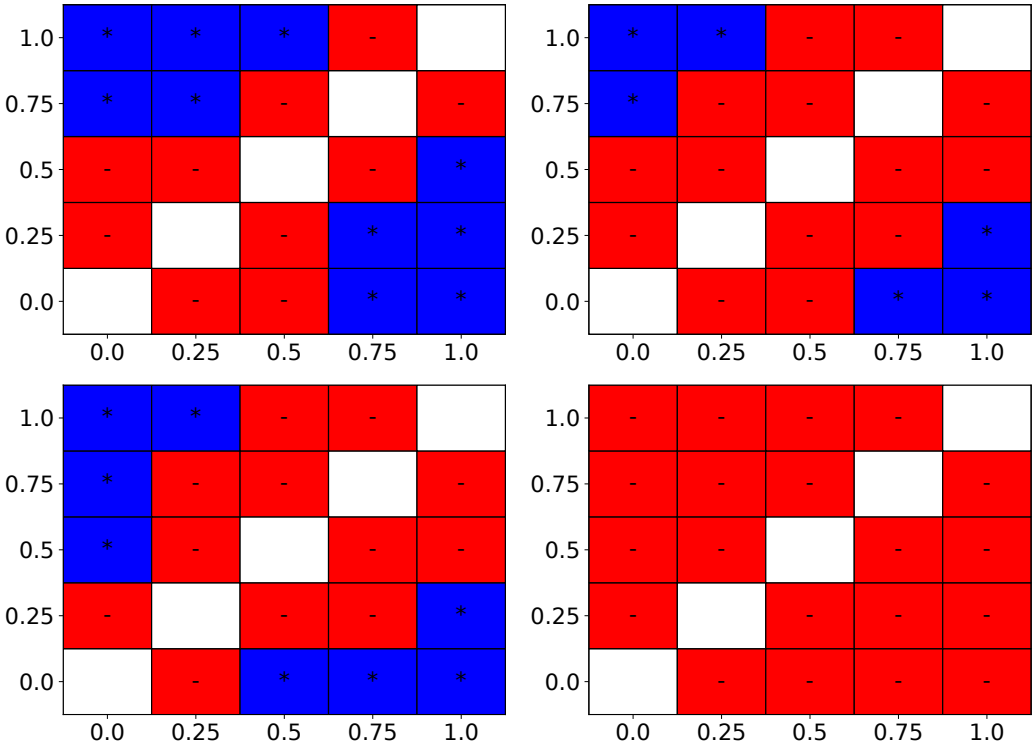

Figure 5: Statistical comparison of the performance achieved by OpenAI-ES by varying the morphological rate. Symbols are defined as follows: "-": not significant; "*": significant at $p < 0.01$. Data are averaged over 20 replications. Left side figures refer to BipedalWalker, while right side plots concern Halfcheetah2D. Top figures refer to the basic environment, whereas bottom figures refer to the hardcore environment.

With regard to GGA, TABLE 7 illustrates the final performance of the controllers evolved in the BipedalWalker and Halfcheetah2D tasks. Also in this case, the condition $\mu = 1.0$ achieves the best fitness at the end of the evolutionary process, regardless of both the problem (BipedalWalker or

Table 6: Effect size η^2_H of the Kruskal-Wallis H test for the BipedalWalker and Halfcheetah2D problems. Data refer to OpenAI-ES.

	BipedalWalker	**BipedalWalkerHardcore**	**Halfcheetah2D**	**Halfcheetah2DHardcore**
η^2_H	0.741	0.454	0.334	0.069

Halfcheetah2D) and the enviromental scenario (basic or hardcore). The performance distribution is also illustrated in FIGURE 6, while FIGURE 10 shows the medians and 95% bootstrapped confidence intervals of the performance. Differently from OpenAI-ES, the standard deviation is high, particularly for the basic scenarios (FIGURE 6, top figures).

Table 7: Results of the experiments performed with GGA in the BipedalWalker and Halfcheetah2D, for both basic and hardcore scenarios. Performance (shown as $avg \pm std$) is averaged over 20 replications of the experiments. Highest performance is denoted in bold.

μ	**BipedalWalker**	**BipedalWalkerHardcore**	**Halfcheetah2D**	**Halfcheetah2DHardcore**
0.0	140.168 ± 91.736	28.453 ± 14.109	45.208 ± 16.365	36.853 ± 11.805
0.25	206.011 ± 73.135	32.117 ± 8.855	68.068 ± 25.567	52.400 ± 8.748
0.5	239.749 ± 61.085	59.578 ± 25.414	85.004 ± 39.561	57.166 ± 10.484
0.75	257.135 ± 55.697	103.908 ± 49.217	128.173 ± 99.413	82.384 ± 19.016
1.0	**275.607 ± 50.605**	**162.917 ± 74.376**	**208.819 ± 88.451**	**138.827 ± 61.476**

From a statistical point of view, we found that the conditions $\mu = 0.5$, $\mu = 0.75$ and $\mu = 1.0$ significantly outperform the conditions $\mu = 0.0$ in the BipedalWalker, BipedalWalkerHardcore and Halfcheetah2D ($p < 0.01$, see FIGURE 8). There exist notable differences also between conditions $\mu = 1.0$ and $\mu = 0.25$, regardless of the problem and scenario. In addition, condition $\mu = 1.0$ is also superior to condition $\mu = 0.5$ in the BipedalWalkerHardcore, Halfcheetah2D and Halfceetah2DHardcore ($p < 0.01$). Overall, despite GGA is remarkably inferior to OpenAI-ES ($p < 0.01$ for each pair of comparisons), the usage of high mutation rates $\mu \geq 0.75$ allows GGA to achieve remarkable benefits in all cases over the evolution of neural parameters only. The effect sizes are provided in TABLE 8.

Table 8: Effect size η^2_H of the Kruskal-Wallis H test for the BipedalWalker and Halfcheetah2D problems. Data refer to GGA.

	BipedalWalker	**BipedalWalkerHardcore**	**Halfcheetah2D**	**Halfcheetah2DHardcore**
η^2_H	0.326	0.621	0.432	0.617

Regarding the Embryo problem, TABLE 9 shows the performance achieved by OpenAI-ES and GGA. A first interesting observation comes from the fact that there exists little variance among the different values of μ, regardless of the considered algorithm. This is also confirmed by the performance distribution shown in FIGURE 9. Overall, we did not find any statistical difference in this case ($p > 0.01$, see FIGURE 11). Another worthwhile outcome is the fact that GGA is competitive with OpenAI-ES in this domain ($p > 0.01$ for each pair of comparisons). The effect sizes are provided in TABLE 10.

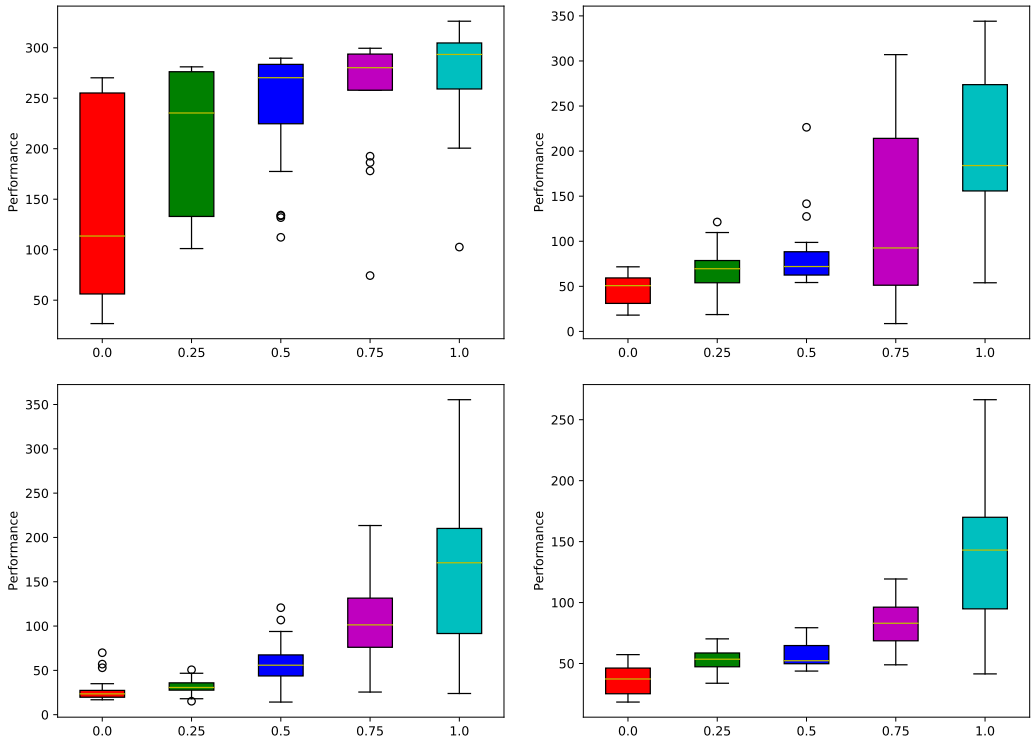

Figure 6: Performance achieved by GGA by varying the morphological rate with regard to BipedalWalker (left side) and Halfcheetah2D (right side). Top figures refer to the basic environment, whereas bottom figures refer to the hardcore environment.

Table 9: Results of the experiments performed with OpenAI-ES and GGA in the Embryo problem. Performance (shown as $avg \pm std$) is averaged over 20 replications of the experiments. Highest performance is denoted in bold.

μ	OpenAI-ES	GGA
0.25	**326.266 ± 43.380**	318.516 ± 25.340
0.5	309.294 ± 82.157	320.966 ± 30.587
0.75	323.700 ± 44.287	**326.960 ± 22.104**
1.0	316.887 ± 60.230	304.023 ± 34.061

Table 10: Effect size η^2_H of the Kruskal-Wallis H test for the Embryo task.

	OpenAI-ES	GGA
η^2_H	0.0	0.050

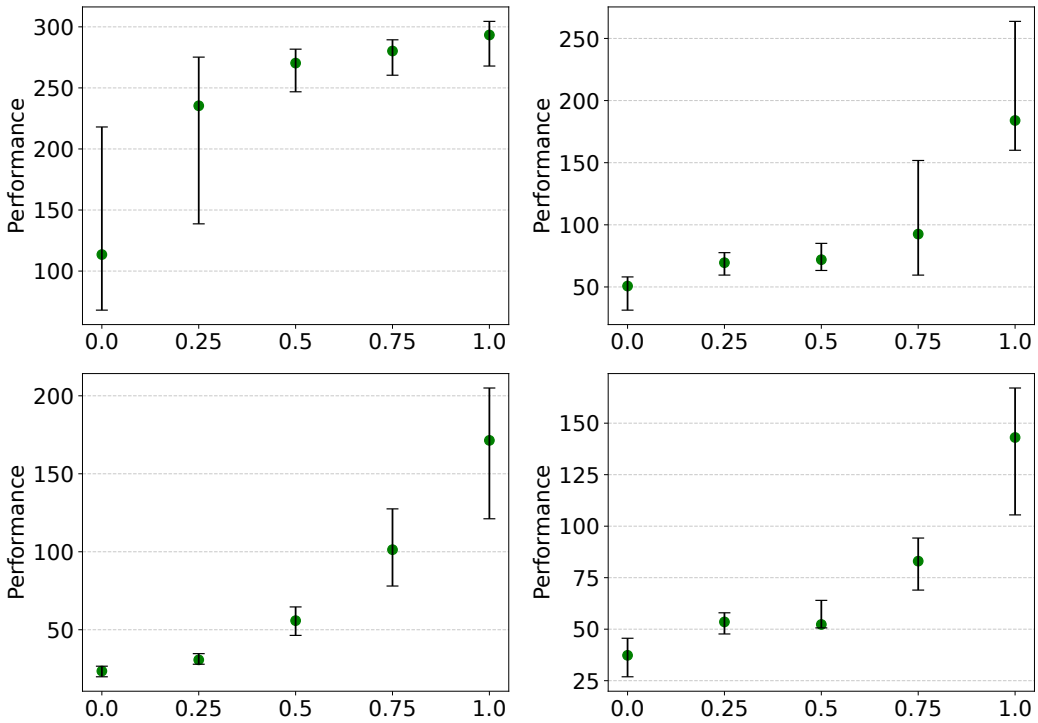

Figure 7: Medians (green circles) and 95% bootstrapped confidence intervals of the performance achieved by GGA by varying the morphological rate. Data are averaged over 20 replications. Left side figures refer to BipedalWalker, while right side plots concern Halfcheetah2D. Top figures refer to the basic environment, whereas bottom figures refer to the hardcore environment.

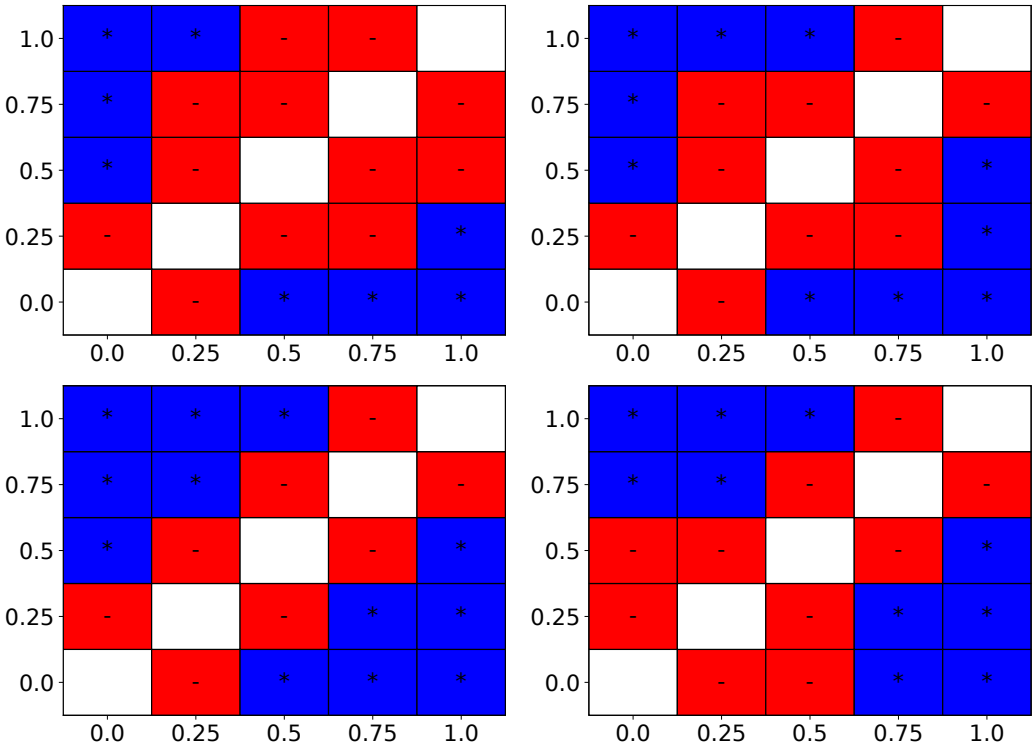

Figure 8: Statistical comparison of the performance achieved by GGA by varying the morphological rate. Symbols are defined as follows: "-": not significant; "*": significant at $p < 0.01$. Data are averaged over 20 replications. Left side figures refer to BipedalWalker, while right side plots concern Halfcheetah2D. Top figures refer to the basic environment, whereas bottom figures refer to the hardcore environment.

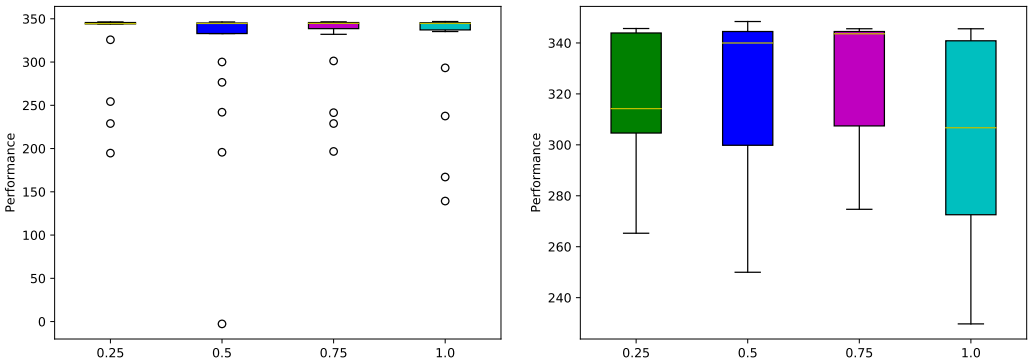

Figure 9: Performance achieved by OpenAI-ES (left) and GGA (right) by varying the morphological rate with regard to Embryo.

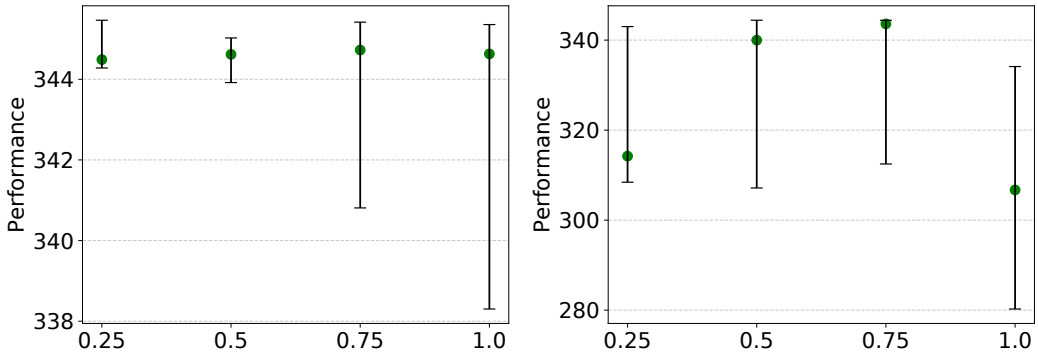

Figure 10: Medians (green circles) and 95% bootstrapped confidence intervals of the performance achieved by OpenAI-ES and GGA in the Embryo problem by varying the morphological rate. Data are averaged over 20 replications.

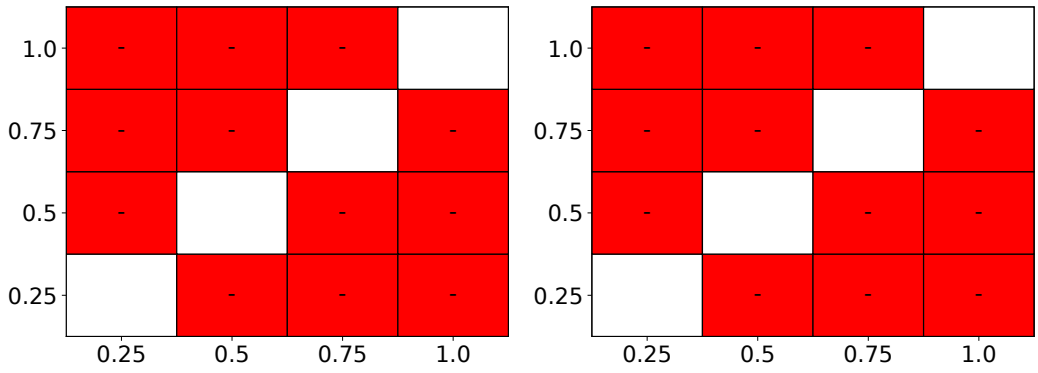

Figure 11: Statistical comparison of the performance achieved by OpenAI-ES (left) and GGA (right) by varying the morphological rate. Symbols are defined as follows: "-": not significant; "*": significant at $p < 0.01$. Data are averaged over 20 replications.

To quantify how big the morphological rate μ must be in order to promote the discovery of enhanced solutions, we analyzed the correlation between μ and the performance F. Regarding OpenAI-ES, we found a strong positive correlation in the BipedalWalker ($\rho = 0.851$, significant at $p < 0.01$), BipedalWalkerHardcore ($\rho = 0.679$, significant at $p < 0.01$), Halfcheetah2D ($\rho = 0.598$, significant at $p < 0.01$) and Halfcheetah2DHardcore ($\rho = 0.295$, significant at $p < 0.01$). Therefore, the higher the morphological rate μ, the higher the performance of the discovered solutions. A similar trend has been observed for GGA: a strong positive correlation exists in the BipedalWalker ($\rho = 0.587$, significant at $p < 0.01$), BipedalWalkerHardcore ($\rho = 0.792$, significant at $p < 0.01$), Halfcheetah2D ($\rho = 0.654$, significant at $p < 0.01$) and Halfcheetah2DHardcore ($\rho = 0.783$, significant at $p < 0.01$). Instead, no correlation exists in the Embryo, irrespective of the considered algorithm.

4. DISCUSSION

The results reported in this work confirm the positive influence of body-brain co-evolution on performance. Specifically, the analysis provided in Section 3 indicate that using morphological rates μ higher than 0.5 is beneficial in the BipedalWalker (both basic and hardcore scenarios) and Halfcheetah2D (basic scenario) problems. In fact, enabling the evolution of robot morphology leads to solutions that better cope with the environmental conditions. For example, in the BipedalWalker task, evolved solutions are generally characterized by longer and thinner legs (see FIGURE 12, top figures). Morever, the density of legs and hull is lower than in the fixed morpholohy case. Both features are particularly useful to make robot locomote easily and maintain a better balance on the terrain. Moreover, they turn out to be crucial for tackling challenges like holes and obstacles in the hardcore scenario. Concerning the Embryo problem, the evolved solutions are mostly characterized by modified angle ranges, while the body sizes and thicknesses do not vary significantly (see FIGURE 12, middle figures). This is due to the need to discover morphologies enabling the robot to move on the terrain. However, the different evolved morphologies lead to different locomotion strategies. With respect to the Halfcheetah2D task, the evolved solutions typically display thicker feet (see FIGURE 12, bottom figures), which allow the robot to maintain stability on the terrain. Furthermore, robot bodies are characterized by a low density, which allows the robots to move faster. In the hardcore scenario, the evolved solutions make the robots produce jumping behaviors, which are useful to avoid holes and obstacles.

In order to assess the impact of each type of morphological parameter (i.e., body size, body thickness, angle range), we performed an ablation study by isolating the effect of each parameter. As case studies, we considered the BipedalWalker and Halfcheetah2D tasks in the basic scenario. Specifically, we defined a post-evaluation test in which the controllers evolved with OpenAI-ES have been compared on 10 fixed episodes with variants obtained by keeping only one type of morphological parameter, whereas the remaining parameters have been set with default values. TABLE 11 reports the outcomes of this test. Unsurprisingly, evolved controllers outperform the variants in all scenarios, particularly for $\mu \geq 0.5$. This clearly indicates that there is no single type of morphological parameter contributing to the success of body-brain co-evolution. Instead, the possibility to simultaneously optimize multiple types leads to the discovery of morphologies allowing enhanced performance.

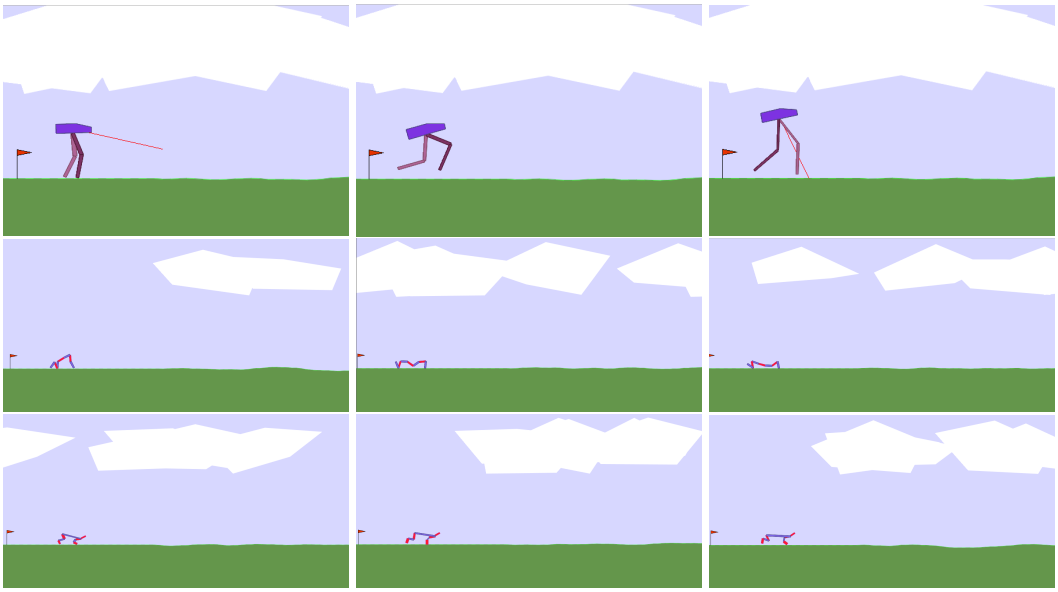

Figure 12: Examples of morphologies evolved with OpenAI-ES in the BipedalWalker (top), Embryo (middle) and Halfcheetah2D (bottom).

Table 11: Performance (shown as $avg \pm std$) of controllers evolved with OpenAI-ES and their variants during the ablation test.

	μ	Evolved	Only body size	Only body thickness	Only joint angle	Only hull
BipedalWalker	0.25	333.385 ± 26.685	307.098 ± 83.280	239.322 ± 120.619	242.475 ± 126.592	232.708 ± 101.938
	0.5	342.232 ± 13.635	281.227 ± 92.767	135.178 ± 117.399	117.410 ± 104.956	92.683 ± 98.083
	0.75	344.372 ± 17.143	169.053 ± 137.451	26.994 ± 56.151	8.153 ± 31.135	8.652 ± 29.651
	1.0	349.912 ± 12.331	157.528 ± 140.981	0.633 ± 16.684	−1.245 ± 12.914	−2.660 ± 15.716
Halfcheetah2D	0.25	274.954 ± 38.106	205.703 ± 68.262	206.072 ± 77.279	201.034 ± 70.307	-
	0.5	274.948 ± 47.611	169.374 ± 87.272	154.624 ± 87.481	148.983 ± 78.465	-
	0.75	293.984 ± 29.333	142.177 ± 78.014	88.927 ± 77.781	79.917 ± 64.168	-
	1.0	307.669 ± 34.194	125.204 ± 119.750	73.972 ± 78.005	62.474 ± 69.140	-

In terms of performance, OpenAI-ES discovers better solutions for the BipedalWalker than for the Halfcheetah2D, regardless of the scenario (basic or hardcore) ($p < 0.01$ for each pair of comparisons). Furthermore, the performance drop in the Halfcheetah2D is higher than for the BipedalWalker (see TABLE 5 and TABLE 12). Interestingly, the fitness ratio between the basic and hardcore scenarios reduces as the morphological rate μ increases for both BipedalWalker and Halfcheetah2D ($\rho = -1.0$, significant at $p < 0.01$). This implies that the possibility of co-evolving body and brain enables to better cope with the experienced environmental conditions and develop more effective strategies. These findings align with previous research works [24, 26].

Concerning GGA, it generally finds more effective solutions for the BipedalWalker than for the Halfcheetah2D in both scenarios. Differently from OpenAI-ES, the performance drop is more marked in the BipedalWalker than in the Halfcheetah2D (see TABLE 12). Moreover, the fitness ratio between the basic and hardcore scenarios reduces as the morphological rate μ increases for the BipedalWalker, while the opposite is true for the Halfcheetah2D.

Table 12: Performance drop exhibited in the BipedalWalker and Halfcheetah2D problems when moving from basic to hardcore scenario.

	OpenAI-ES		**GGA**	
μ	**BipedalWalker**	**Halfcheetah2D**	**BipedalWalker**	**Halfcheetah2D**
0.0	2.035	4.179	4.926	1.227
0.25	1.810	3.909	6.433	1.299
0.5	1.499	3.813	4.024	1.485
0.75	1.469	3.764	2.475	1.556
1.0	1.278	3.675	1.692	1.504

5. CONCLUSIONS

Research in body-brain co-evolution draws inspiration from morphological and neural maturation in living organisms adapting to ecological pressures. Although it was originally intended to replicate these biological mechanisms in artificial life and robotics, a lot of studies focus also on several techniques to make this co-evolution as much effective as possible. In addition, the possibility to simultaneously co-evolve body and brain frees the experimenter(s) from the burden of designing suitable morphologies for a given problem, which requires task-specific expertise and might lead to sub-optimality. A classic benchmark in the community involves robot locomotion, in which evolving the morphology could be beneficial to tackle the environmental constraints and adapt to different situations. For example, body-brain co-evolution in Pybullet robot locomotion tasks proved to be more effective than using a fixed morphology. Nevertheless, assessing the impact of the morphological rate (i.e., the rate at which morphological parameters vary) on the efficacy of body-brain co-evolution has not been previously analyzed.

This work presents an investigation on the impact of morphological rate μ in three 2D locomotion problems , i.e. BipedalWalker, Embryo and Halfcheetah2D. As regards BipedalWalker and Halfcheetah2D, the analysis has been conducted in two environmental scenarios of increasing complexity. We used the OpenAI-ES and GGA algorithms to co-evolve morphological and neural pa-

rameters, and we considered different values of μ. Our results indicate that the superiority of body-brain co-evolution over using a fixed morphology arises when $\mu \geq 0.5$. In general, we observe a positive correlation between the morphological rate μ and the final performance in the BipedalWalker and Halfcheetah2D, irrespective of the considered scenario. Moreover, our analysis underscores that body-brain co-evolution produces better performance in the BipedalWalker problem than in the Halfcheetah2D task. With respect to the Embryo, we do not observe significant differences based on the morphological rate μ. Lastly, we run an ablation study to assess the impact of each type of morphological parameter (body size, body thickness, angle range) on the effectiveness of body-brain co-evolution, and we prove that no single type of morphological parameter is responsible for this success. Instead, the possibility to simultaneously optimize all the different types plays a major role in the efficacy of co-optimizing body and brain.

As future research directions, we are interested in extending our investigation to other problems, including soft robots, maze navigation, other locomotion tasks and real-world problems, in order to provide a general framework. Furthermore, we plan to delve into the role of the algorithm used for body-brain co-evolution: our results have been obtained by using OpenAI-ES and GGA, but we do not know if similar outcomes hold with different Evolutionary Algorithms (EAs). Therefore, extending our analysis to other methods will be object of future works. Finally, we will investigate the impact of hyperparameter settings in the efficacy of body-brain co-evolution through automated tools [58].

References

[1] Bongard J. Morphological Change in Machines Accelerates the Evolution of Robust Behavior. Proc Natl Acad Sci. 2011;108:1234-1239.

[2] Carroll SB. Chance and Necessity: The Evolution of Morphological Complexity and Diversity. Nature. 2001;409:1102-1109.

[3] Koehl MAR. When Does Morphology Matter? Annu Rev Ecol Syst. 1996;27:501-452.

[4] Sims K. Evolving 3D Morphology and Behavior by Competition. Artif Life. 1994;1:353-372.

[5] Lee WP. Evolving Robot Brains and Bodies Together: An Experimental Investigation. J Chin Inst Eng. 2003;26:125-132.

[6] Koza JR. On the Programming of Computers by Means of Natural Selection. Genet Program. 1992.

[7] Holland JH. Adaptation in Natural and Artificial Systems. Ann Arbor: University of Michigan Press; 1976;18:529-530.

[8] Krčah P. Solving Deceptive Tasks in Robot Body-Brain Co-Evolution by Searching for Behavioral Novelty. In: 10th International Conference on Intelligent Systems Design and Applications. 2010:284-289.

[9] Lehman J, Stanley KO, et al. Exploiting Open-Endedness to Solve Problems Through the Search for Novelty. In: Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI). MIT Press, 2008. 2008;65:329-336.

[10] Lehman J, Stanley KO. Abandoning Objectives: Evolution Through the Search for Novelty Alone. Evol Comput. 2011;19:189-223.

[11] Pontes-Filho S, Walker K, Najarro E, Nichele S, Risi S. A Single Neural Cellular Automaton for Body-Brain Co-Evolution. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion. 2022:148-151.

[12] Medvet E, Lippolis S, Nadizar G. Quality-Diversity in Problems With Composite Solutions: A Case Study on Body–Brain Robot Optimization. Genet Program Evolvable Mach. 2025;26:22.

[13] Jelisavcic M, Glette K, Haasdijk E, Eiben AE. Lamarckian Evolution of Simulated Modular Robots. Front Robot AI. 2019;6:9.

[14] Luo J, Miras K, Tomczak J, Eiben AE. Enhancing Robot Evolution Through Lamarckian Principles. Sci Rep. 2023;13:21109.

[15] Luo J, Miras K, Longhi C, Weissl O, Eiben AE. Lamarckian Inheritance Improves Robot Evolution in Dynamic Environments. IEEE Trans Evol Comput. 2026;30:1495-1504.

[16] Available at: https://www.biodiversitylibrary.org/bibliography/26297

[17] Nygaard TF, Martin CP, Howard D, Torresen J, Glette K. Environmental Adaptation of Robot Morphology and Control Through Real-World Evolution. Evol Comput. 2021;29:441-461.

[18] de Bruin E, Glette K, Ellefsen KO, Nadizar G, Medvet E. Social Learning Strategies for Evolved Virtual Soft Robots. In: Proceedings of the Genetic and Evolutionary Computation Conference. 2026:96-105.

[19] Ferigo A, Iacca G, Medvet E, Nadizar G. Totipotent Neural Controllers for Modular Soft Robots: Achieving Specialization in Body–Brain Co-Evolution Through Hebbian Learning. Neurocomputing. 2025;614:128811.

[20] Kriegman S, Cheney N, Bongard J. How Morphological Development Can Guide Evolution. Sci Rep. 2018;8:13934.

[21] Mertan A, Cheney N. Modular Controllers Facilitate the Co-Optimization of Morphology and Control in Soft Robots. In: Proceedings of the Genetic and Evolutionary Computation Conference. New York, USA: ACM; 2023:174-183.

[22] Mertan A, Cheney N. Investigating Premature Convergence in Co-Optimization of Morphology and Control in Evolved Virtual Soft Robots. In: European Conference on Genetic Programming (Part of EvoStar). 2024:38-55.

[23] Hart E, Le Goff LK. Artificial Evolution of Robot Bodies and Control: On the Interaction Between Evolution, Learning and Culture. Philos Trans R Soc Lond B Biol Sci. 2021;377:20210117.

[24] Pagliuca P, Nolfi S. The Dynamic of Body and Brain Co-Evolution. Adapt Behav. 2022;30:245-255.

[25] Coumans E, Bai Y. Pybullet, a Python Module for Physics Simulation for Games, Robotics and Machine Learning. 2016.

[26] Pagliuca P, Vitanza A. Learning Locomotion by Co-Evolution of Morphological and Neural Parameters. In: Proceedings of the 2025 IEEE International Conference on Development and Learning (ICDL). 2025:1-6.

[27] Carvalho JT, Nolfi S. The Role of Morphological Variation in Evolutionary Robotics: Maximizing Performance and Robustness. Evol Comput. 2024;32:125-142.

[28] Pagliuca P, Nolfi S. Robust Optimization Through Neuroevolution. PLOS One. 2019;14:e0213193.

[29] Brockman G, Cheung V, Pettersson L, Schneider J, Schulman J, et al. OpenAI Gym. Arxiv preprint: https://arxiv.org/pdf/1606.01540.

[30] Towers M, Kwiatkowski A, Balis J, De Cola G, Deleu T, et al. Gymnasium: A Standard Interface for Reinforcement Learning Environments. Advances in Neural Information Processing Systems. 2026:38.

[31] Nolfi S. Behavioral and Cognitive Robotics: An Adaptive Perspective. 2021.

[32] Pagliuca P, Vitanza A. The Interplay Between Functionality and Complexity in Neural Network Efficiency. In: Proceedings of the 24th International Conference of the Italian Association for Artificial Intelligence (AIxIA 2026); 2026.

[33] Salimans T, Ho J, Chen X, Sidor S, Sutskever I. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. Arxiv preprint: https://arxiv.org/pdf/1703.03864

[34] Pagliuca P, Milano N, Nolfi S. Efficacy of Modern Neuro-Evolutionary Strategies for Continuous Control Optimization. Front Robot AI. 2020;7:98.

[35] Grefenstette, John J. Genetic Algorithms for Changing Environments. Parallel Problem Solving from Nature. 1992.

[36] Aydogmus O, Yilmaz M. Comparative Analysis of Reinforcement Learning Algorithms for Bipedal Robot Locomotion. IEEE Access. 2024;12:7490-7499.

[37] Ha D. Reinforcement Learning for Improving Agent Design. Artif Life. 2019;25:352-365.

[38] Song DR, Yang C, McGreavy C, Li Z. Recurrent Deterministic Policy Gradient Method for Bipedal Locomotion on Rough Terrain Challenge. In: 15th International Conference on Control, Automation, Robotics and Vision (ICARCV). 2018:311-318.

[39] Pagliuca P. Whether, How and When Modern Evolutionary Strategies Can Be Improved. J Artif Intell Auton Intell. 2025;2:425-449.

[40] Pagliuca P, Vitanza A. Self-Organized Aggregation in Group of Robots With OpenAI-ES. In: International Conference on Soft Computing and Pattern Recognition. 2022:770-780.

[41] Pagliuca P, Vitanza A. A Comparative Study of Evolutionary Strategies for Aggregation Tasks in Robot Swarms: Macro- and Micro-Level Behavioral Analysis. IEEE Access. 2025;13:72721-72735.

[42] Pagliuca P, Vitanza A. What Evolutionary Algorithm(s) Should We Use? Towards Preliminary Guidelines From an Exploratory Analysis of Foraging Problems. In: Proceedings of the 22nd International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT); 2026:22-24.

[43] Chrabaszcz P, Loshchilov I, Hutter F. Back to Basics: Benchmarking Canonical Evolution Strategies for Playing Atari. Arxiv preprint: https://arxiv.org/pdf/1802.08842.

[44] Lange R, Tang Y, Tian Y. NeuroEvoBench: Benchmarking Evolutionary Optimizers for Deep Learning Applications. Adv Neural Inf Process Syst. 2023;36:32160-72.

[45] Zhang X, Clune J, Stanley KO. On the Relationship Between the OpenAI Evolution Strategy and Stochastic Gradient Descent. Arxiv preprint: https://arxiv.org/pdf/1712.06564

[46] Pagliuca P, Vitanza A. Enhancing Aggregation in Locomotor Multi-Agent Systems: A Theoretical Framework. Proceedings of the 25th Edition of the Workshop From Object to Agents (WOA24). 2024;3735:42-57.

[47] Pagliuca P. Discovering a Single Neural Network Controller for Multiple Tasks With Evolutionary Algorithms. J Artif Intell Auton Intell. 2025;2:322-348.

[48] Pagliuca P, Trivisano G, Vitanza A. How to Evolve Aggregation in Robotic Multi-Agent Systems. In: Proceedings of the 26th Edition of the Workshop From Object to Agents (WOA25). 2025:140-156.

[49] Brockhoff D, Auger A, Hansen N, Arnold DV, Hohm T. Mirrored Sampling and Sequential Selection for Evolution Strategies. In: International Conference on Parallel Problem Solving From Nature. 2010:11-21.

[50] Kingma DP, Ba J. Adam: A Method for Stochastic Optimization. ArXiv preprint arXiv:1412.6980, 2014.

[51] Pagliuca P, Nolfi S, Vitanza A. EvorobotPy3: A Flexible and Easy-to-Use Simulation Tool for Evolutionary Robotics. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2025 Companion). 2025:155-158.

[52] Rosenblatt F. Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms. Spartan Books. 1962;55.

[53] Pagliuca P. Emerging Trends in Multi-Agent Systems From an Evolutionary Perspective. In: Proceedings of the 27th Edition of the Workshop From Objects to Agents (WOA26). 2026.

[54] Pagliuca P. Enhancing Performance of Evolutionary Strategies With Symmetric Sampling. Algorithms. 2026;19:504.

[55] Kruskal WH, Wallis WA. Use of Ranks in One-Criterion Variance Analysis. J Am Stat Assoc. 1952;47:583-621.

[56] Dunn OJ. Multiple Comparisons Using Rank Sums. Technometrics. 1964;6:241-252.

[57] Spearman C. The Proof and Measurement of Association Between Two Things. Am J Psychol. 1904;15:72-101.

[58] López-Ibáñez M, Dubois-Lacoste J, Pérez Cáceres L, Birattari M, Stützle T. The Irace Package: Iterated Racing for Automatic Algorithm Configuration. Oper Res Perspect. 2016;3:43-58.