Impact of Morphology Variations on Evolved Neural Controllers for Modular Robots

Eric Medvet[1(✉)] and Francesco Rusin[2]

[1] Department of Engineering and Architecture, University of Trieste, Trieste, Italy
emedvet@units.it
[2] Department of Mathematics and Geosciences, University of Trieste, Trieste, Italy

Abstract. Modular robots, in particular those in which the modules are physically interchangeable, are suitable to be evolved because they allow for many different designs. Moreover, they can constitute ecosystems where "old" robots are disassembled and the resulting modules are composed together, either within an external assembling facility or by self-assembly procedures, to form new robots. However, in practical settings, self-assembly may result in morphologies that are slightly different from the expected ones: this may cause a detrimental misalignment between controller and morphology. Here, we characterize experimentally the robustness of neural controllers for Voxel-based Soft Robots, a kind of modular robots, with respect to small variations in the morphology. We employ evolutionary computation for optimizing the controllers and assess the impact of morphology variations along two axes: kind of morphology and size of the robot. Moreover, we quantify the advantage of performing a re-optimization of the controller for the varied morphology. Our results show that small variations in the morphology are in general detrimental for the performance of the evolved neural controller. Yet, a short re-optimization is often sufficient for aligning back the performance of the modified robot to the original one.

Keywords: Embodied cognition · Soft robotics · Adaptation

1 Introduction and Related Works

Fully autonomous robotic systems require to be adaptable to environmental changes without an external intervention. One main path toward adaptation of an entire robotic ecosystem, instead of the single robot, consists in having a population of robots that are built, "live", and are disposed in such a way that their robotic material can be reused for building other robots [3]. Modular robots are particularly suitable to form such ecosystems, because their modularity eases the building and disposal phases [4,12]. Moreover, modular robots also favor adaptation, in particular through evolution [2], because they allow for great expressiveness for the morphology and the controller [1,9].

However, actually realizing the scenario of real (i.e., with hardware robots) evolution of a robotic ecosystem poses several challenges [2], ranging from the

well-known reality gap problem [14,20], to the much longer times for evaluating candidate solutions [13], to the need of employing automatic assembly and disassembly procedures for robots, in order to make the system scalable and really autonomous.

The recent progresses on manipulators [12,13], as well as on self-assembly [17], make the automatic assembly of modular robots feasible. However, it can happen that the morphology obtained by unassisted assembly differs from the designed one, in particular for soft modular robots [5]. In such a case, the controller that was designed to be associated with the expected morphology might not be equally effective with the slightly different one. The drop in effectiveness might be large, according to the embodied cognition paradigm that states that the intelligent behavior of an embodied agent depends on the combined work of both its body and brain, since the body-brain misalignment might be detrimental.

In this paper, we experimentally investigate the impact of small variations in the morphology on the effectiveness of evolved neural controllers, i.e., of closed-loop controllers based on neural networks whose parameters are optimized, for a given morphology and task, by means of an evolutionary algorithm (EA). We consider the case of Voxel-based Soft Robots (VSRs), a kind of modular robots whose modules are soft cubes that can expand or contract individually based on signals dictated by the controller.

We perform a number of experiments with six morphologies, consisting in three base morphologies in small and large versions, with controllers evolved for the task of locomotion, a classic task of evolutionary robotics. We apply small random variations to each morphology and measure the impact on the degree to which the resulting VSRs solve the task (i.e., their velocity, for locomotion). We find that the decrease in robot velocity, i.e., the controller effectiveness, is large even for small modifications (one voxel for small original morphologies) and we explain this finding in terms of the great potential for morphological computation that VSRs offer. This potential results in the body having a key role in determining VSR behavior and hence make the body-brain misalignment particularly detrimental.

We also experimentally verify whether a re-optimization for the slightly modified morphology can make the controller back on par with the effectiveness of the one evolved for the original morphology. We found that (a) modified morphologies are not intrinsically worse and (b) seeding the re-optimization with the original controller makes the re-optimization efficient, besides effective.

We believe that our results contribute to strengthen the understanding of how body and brain interact in modular robots and, more broadly, constitute a further step toward autonomous and evolvable robotic ecosystems.

2 Evolutionary Optimization of Voxel-based Soft Robots

Voxel-based Soft Robots (VSRs) are a kind of modular robots where each module, called voxel, is a deformable cube that is attached to adjacent cubes at the

vertexes. The volume of each voxel changes according to external forces acting on it, such as the gravity and those deriving from the contact with the ground, and to an internal force, itself modulated by a control value dictated by the controller of the robot. The way the controller varies the control value of each voxel over time determines, together with the robot-environment interaction, the behavior of the VSR.

Both the controller and the morphology of a VSR can be optimized in order to obtain a behavior that allows the robot to achieve a predefined task. For the purpose of this study, we are only concerned with the optimization of the controller; as an aside, previous works have shown that VSRs are particularly suitable for the concurrent optimization of controller and morphology [9,15].

In this study, we consider a 2-D variant of VSRs that can be simulated in discrete time and continuous space [8]. In the next sections, we describe in detail the morphology and the controller of our VSRs, as well as the way we optimize the latter by means of an EA.

2.1 VSR Morphology

A VSR *morphology* is unequivocally described by a 2-D grid in which each non-empty element describes a voxel. A *voxel* is modeled, in the simulation, as a compound of four rigid bodies (at the corners) and many spring-damper systems, for softness and elasticity, connecting the rigid bodies [8]. Adjacent voxels are glued together at the vertexes: that is, their rigid bodies at the corners are bound with joints that do not permit relative rotation, nor distance variation.

By varying the parameters of this mechanical model, the designer may vary the properties of the (simulated) material the voxel consists of, possibly impacting on the overall behavior of the VSR [11]. In this work, we assume that all the voxels consist of the same material, hence a morphology is unequivocally described by a Boolean grid (or matrix) $\boldsymbol{m} \in \{\mathrm{T}, \mathrm{F}\}^{w \times h}$, where the $m_{x,y}$ is set if there is a voxel at coordinates (x, y).

During the simulation, the area of voxels changes upon the combined effect of external forces and the voxel control value, dictated by the controller. For a voxel at (x, y), we denote the *control value* at time step k as $a_{x,y}^{(k)} \in [-1, 1]$, where -1 corresponding to maximum requested expansion, and 1 corresponding to maximum requested contraction. Contraction and expansion are modeled as linear variations of the rest-length of the spring-damper systems, proportional to value of $a_{x,y}^{(k)}$. We used the default values of the simulator 2D-VSR-Sim [8] for the minimum and maximum rest-length values.

2.2 VSR Controller

The controller of a VSR is in charge of determining the control value for each voxel of the VSR at each time step of the simulation.

In this study, we use a distributed neural controller [7] that determines the control values based on some sensory inputs acquired from the voxels, hence

realizing a closed-loop control of the robot. Our controller is distributed in the sense that, for each voxel, the control value is determined locally based on local sensory inputs and information acquired from adjacent voxels. Moreover, it is neural because the processing of that information is performed with a feed-forward artificial neural network (NN) embedded in the voxel.

In detail, the distributed neural controller works as follows. At each time step k and for each voxel at (x, y), we collect a *sensor reading* $\boldsymbol{s}_{x,y}^{(k)} \in [-1, 1]^4$ consisting of: (a) the ratio between the current area of the voxel and its rest area, (b) a binary value set to 1 if the voxel is in contact with the ground or -1 otherwise, and (c) the velocities of the center of mass of the voxel along the x- and y-axes. We normalize the values of the ratio and the velocities in order to ensure they are in the $[-1, 1]$ domain. Then, we use the NN located in the voxel to compute the control value $a_{x,y}^{(k)}$ and the information to be passed to adjacent voxels at the next time step:

$$\left[a_{x,y}^{(k)}\ \boldsymbol{i}_{x,y}^{\triangle(k)}\ \boldsymbol{i}_{x,y}^{\triangledown(k)}\ \boldsymbol{i}_{x,y}^{\triangleleft(k)}\ \boldsymbol{i}_{x,y}^{\triangleright(k)}\right] = \mathrm{NN}_{\boldsymbol{\theta}}\left(\left[\boldsymbol{s}_{x,y}^{(k)}\ \boldsymbol{i}_{x,y-1}^{\triangle(k-1)}\ \boldsymbol{i}_{x,y+1}^{\triangledown(k-1)}\ \boldsymbol{i}_{x+1,y}^{\triangleleft(k-1)}\ \boldsymbol{i}_{x+1,y}^{\triangleright(k-1)}\right]\right),$$

where $\boldsymbol{i}_{x,y}^{\triangle(k)}, \boldsymbol{i}_{x,y}^{\triangledown(k)}, \boldsymbol{i}_{x,y}^{\triangleleft(k)}, \boldsymbol{i}_{x,y}^{\triangleright(k)} \in [-1, 1]^{n_{\mathrm{comm}}}$ is the information output by the voxel at (x, y) for its neighbors at $(x, y + 1)$ ($\triangle$), $(x, y - 1)$ ($\triangledown$), $(x - 1, y)$ ($\triangleleft$), and $(x + 1, y)$ ($\triangleright$), respectively, and $\boldsymbol{\theta} \in \mathbb{R}^p$ is the vector of the parameters (or weights) of the NN. The NN hence acts as a function $\mathbb{R}^{4+4n_{\mathrm{comm}}} \to \mathbb{R}^{1+4n_{\mathrm{comm}}}$.

The information propagation to and from adjacent voxels allows the distributed neural controller to realize a form of collective intelligence that results, upon optimization, in an emergent behavior of the VSR, despite being composed of independents controller modules (the NNs) [16]. Moreover, the one-step delay in propagation resulting from the fact that, at time k, a NN uses the information produced by adjacent NNs at time $k-1$, makes the distributed neural controller capable of exhibiting complex dynamical behaviors, despite being composed of static modules. However, it has been observed that this richness may often lead to vibrating behaviors [9] that would hardly be effective in reality (hence resulting in the so-called reality gap problem [14,20]). For this reason, we actually change the control value at a lower frequency, namely 5 Hz, than the one of the simulation, being 60 Hz.

Concerning the dimension of the information passed to and from NNs and the architecture of each NN, we here use the same architecture, with one inner layer with 8 neurons, and parameters $\boldsymbol{\theta}$ for each NN and $n_{\mathrm{comm}} = 1$, as this setting has been proven to allow for effective control without increasing too much the search space [9,16]. We hence have $\boldsymbol{\theta} \in \mathbb{R}^p$, with $p = (4+4+1)\cdot 8+(8+1)\cdot(1+4) = 117$, where the trailing $+1$s are the bias. We use tanh as activation function.

A key property of the distributed neural controller that holds when, as in this case, all the NNs use the same parameters $\boldsymbol{\theta}$ is that a controller may be coupled with any morphology. This perfectly fits the scenario where voxels self-assemble to form a desired morphology on which a pre-trained controller is hence "installed". In our work, we exploit this this property to evolve a controller for one morphology and then couple it with slightly different morphologies to study its effectiveness in the new conditions.

2.3 Evolution of VSR Controller

In order to evolve the controller for a VSR, namely its parameters $\boldsymbol{\theta} \in \mathbb{R}^p$, we employ a standard $\mu + \lambda$ EA, as follows.

We iteratively evolve a population of n_{pop} solutions, i.e., numerical vectors, until we have performed n_{evals} fitness evaluations. At each iteration, we build the offspring by repeating n_{pop} times the following steps: (1) we randomly select the crossover (with probability p_{xover}) or the mutation (with probability $1 - p_{\text{xover}}$) genetic operator; (2) we select one or two parents (depending on the chosen operator) with a tournament selection of size n_{tour}; (3) we apply the operator to the parents obtaining a new individual for which we evaluate the fitness. Then, we merge the parents and the offspring, we retain only the n_{pop} best individuals, and proceed to the next generation. At the end of the evolution, we pick the solution in the population with the best fitness as the evolved controller.

For initializing the population, we build each initial solution $\boldsymbol{\theta}$ by sampling each θ_i from $U(-1, 1)$. Concerning the genetic operators, we use the extended geometric crossover, where the child $\boldsymbol{\theta} \in \mathbb{R}^p$ is obtained from the parents $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^p$ as $\boldsymbol{\theta} = \boldsymbol{\theta}_1 + \boldsymbol{\alpha}(\boldsymbol{\theta}_2 - \boldsymbol{\theta}_1) + \boldsymbol{\beta}$, where $\boldsymbol{\alpha} \in \mathbb{R}^p$ is sampled as $\alpha_i \sim U(-0.5, 1.5)$, and $\boldsymbol{\beta} \in \mathbb{R}^p$ is sampled as $\beta_i \sim N(0, \sigma_{\text{xover}})$. As mutation, we use the Gaussian mutation, where $\boldsymbol{\theta} = \boldsymbol{\theta}_1 + \boldsymbol{\beta}$, with betas sampled from $N(0, \sigma_{\text{mut}})$.

3 Experiments and Discussion

Our broad aim is to study the impact of small variations in the morphology on the effectiveness of evolved neural controllers. More precisely, we aim at answering the following research questions:

RQ1 To which degree do small morphology variations affect the controller effectiveness? Do overall robot size and initial morphology have an impact on how variations affect effectiveness?

RQ2 Is it possible to mitigate the decrease in effectiveness of an evolved controller by re-optimizing it on the varied morphology?

In order to answer these questions, we performed a large number of experiments in which we (i) evolved the controllers for a few pre-defined morphologies, taking note of their effectiveness upon evolution on a pre-defined task, (ii) modified the morphologies to different extents, and (iii) measured the effectiveness on the task of the evolved controller applied to the modified morphology. Moreover, for the purpose of addressing RQ2, we performed a further re-optimization of the evolved controller on the modified morphology, taking note of its effectiveness.

We experimented with the task of locomotion, in which the robot is placed on a flat surface and has to move the farthest possible towards the right (i.e., along the positive direction of the x-axis). For this task, we measure the effectiveness as the average velocity v_x along the x-axis computed by considering the x-position of the center of mass of the VSR at the beginning and at the end of a simulation lasting $t_{\text{sim}} = 30$ s. We use v_x as the fitness of the robot during the evolution: clearly, the larger, the better.

Concerning the pre-defined morphologies, i.e., the ones on which we evolved the controllers, we experimented with 6 different cases, that we show in Fig. 1 (in darker green). They are a biped, a comb, and a worm, each one in two sizes (small and large): the biped consists of an horizontal body and two legs, the comb extends this concept to four legs, resembling indeed a comb, and the worm is just a full rectangle. They are composed by 10, 11, and 10 voxels (biped, comb, worm), for small morphologies, and 40, 44, and 40 voxels, for large ones.

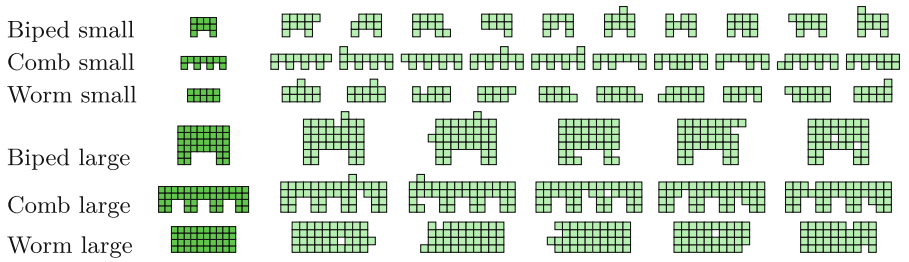

Fig. 1. The 6 morphologies of our experiments (in darker green) with 10 (small, with $\delta = 1$) or 5 (large, with $\delta = 2$) example variations (in lighter green). (Color figure online)

For producing the variations in the morphology, that are a key component of this study, we proceeded as follows. Given a morphology $\boldsymbol{m}$ and a target extent δ of the variation, in number of voxels, we iteratively and randomly removed or added voxels, i.e., we randomly changed the Boolean values of $\boldsymbol{m}$, until the Hamming distance between the varied morphology $\boldsymbol{m}'$ and the original one was $d_{\text{hamming}}(\boldsymbol{m}, \boldsymbol{m}') = \delta$. At each random addition or removal of a voxel from a morphology, we ensured that it remained connected, i.e., a single body; that is, we always dealt with proper polyominoes. Figure 1 shows, for each original morphology and in lighter green, some example variations obtained with $\delta = 1$ (small) or $\delta = 2$ (large).

We performed our experiments with 2D-VSR-Sim [8], for simulating the VSRs, and with JGEA [10], for the evolutionary optimization. Concerning the former, we set all the parameters to their default values. For the latter, we set $n_{\text{pop}} = 100$, $n_{\text{evals}} = 10000$, $p_{\text{xover}} = 75\%$, $n_{\text{tour}} = 5$, and $\sigma_{\text{mut}} = \sigma_{\text{xover}} = 0.35$. We used $\delta \in \{1, 2, 3\}$ for small morphologies and $\delta \in \{2, 4, 6\}$ for large ones.

We performed statistical significance tests with the Mann-Whitney U rank test with the null hypothesis of equality of the median, after having verified all other relevant hypotheses, and with $\alpha = 0.05$.

3.1 Results and Discussion for RQ1: Impact of Variations

Figure 2 shows the results concerning the impact on controller effectiveness v_x (on the y-axis) of morphology variations of extent δ (on the x-axis) applied after the evolution of the neural controller (red line) and without any re-optimization, for the six considered morphologies—the blue line is discussed in the next section

as it pertains to RQ2. The value of v_x for $\delta = 0$ corresponds to the fitness at the end of the evolution on the original morphology, i.e., without any variation, and is computed as the median across 10 independent evolutionary runs. The value of v_x for $\delta > 0$ is the median of the velocities of the 10 controllers obtained with the original morphology applied each one to 10 variations, for each δ, of the original morphology; i.e., for each $\delta > 0$, there are 100 values.

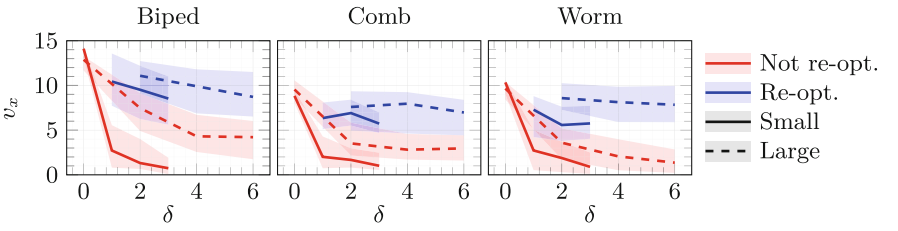

Fig. 2. Controller effectiveness v_x (median and interquartile range) vs. the extent δ of the variation for different original morphologies, for the original controller (red) or with re-optimization (blue). (Color figure online)

The main finding we infer from Fig. 2 is that there is a clear decrease in the effectiveness of the neural controller once the morphology it has been evolved on changes. The decrease is apparent for all the original morphologies in both sizes and it is statistically significant for any $\delta > 0$. The difference between the v_x on the original morphology and the one on the modified morphology ranges from ≈ 10 ($\approx 70\%$) for the small biped to ≈ 5 ($\approx 55\%$) for the large comb.

Interestingly, the extent δ of the morphology variation appears to play a minor role in the decrease of effectiveness, in particular for the small original morphologies—the differences for pairs of $\delta > 0$ are statistically significant only in a few cases. The drop in effectiveness looks smoother, with respect to δ, only for the large biped.

For understanding the reason for such a drastic decrease in the effectiveness, even when the variation in the morphology consists in just one voxel, we analyzed many simulations visually, i.e., we observed the behaviors in locomotion of the corresponding VSRs. In general, the degree to which the behavior appears successful in the task of locomotion depends on the original morphology, the small biped being the best one. For this case, it is sharply clear that the gait is very negatively affected by addition of removal of just one voxel: we made a video for the gait of the VSR with the original morphology and with $\delta = 1$ publicly available at https://youtu.be/bB1u3Yj6FTo. The gait for the original small biped is effective because the controller evolved to master the dynamics of the body: the latter has a peculiar periodicity that can be exploited to obtain an effective gait. When just one block is removed or added, the periodicity changes and the controller is no more able to exploit it. More broadly, the VSR is a dynamical system whose attractor in the space of poses is cyclic and highly functional to the locomotion: when some properties of the morphology change, the attractor becomes much less effective.

We believe that these experiments further corroborate the validity of the embodied cognition paradigm: the ability of the VSR to perform a task is hosted jointly in the VSR brain and body, not just, nor mostly in the former. This is particularly true for VSRs because their body, being an aggregation of many soft components, each corresponding to a simple dynamical system, offers a great potential for performing morphological computation [18,22].

3.2 Results and Discussion for RQ2: Re-Optimization

With the experiments discussed in the previous section, we found that the small variations in the morphology are greatly detrimental for the effectiveness of the evolved neural controllers. We also hypothesized that the decrease in effectiveness is rooted in the misalignment between body and brain that arises from the morphology variation. Here we wonder if another, not necessary alternative, motivation is in the fact that the modified morphologies are intrinsically less suitable for performing the task. For investigating this scenario, we considered, for each original morphology and each value of $\delta > 0$, the $10 \cdot 10$ varied morphologies and performed one evolutionary optimization of a neural controller for each of them. We used the same EA, with the same parameters, we used for the previous experiments.

Figure 2 presents the results of this experiment. It shows, through the blue lines, the controller effectiveness v_x (on the y-axis) obtained at the end of the re-optimization vs. the morphology variations of extent δ (on the x-axis), for the six considered morphologies.

The foremost finding is that the controller effectiveness v_x for the modified morphologies is, upon re-optimization, almost on par with the effectiveness of the original controller on the original morphology—the small differences are always statistically significant, with the exception of the large comb and worm. From another point of view, if a new neural controller is evolved for the modified morphology, it is clearly better than the non re-optimized controller that was obtained for the original morphology.

A second observation concerns the variability of v_x for the re-optimized controllers—we recall that both red and blue lines of Fig. 4 are, for $\delta > 0$, computed on 100 values. The figure suggests that, for half of the morphologies, the interquartile range for re-optimized controllers is larger than the one for the original controllers. We looked at a few behaviors of the VSRs with re-optimized controllers and we found that, indeed, some of the morphologies appeared more suitable for locomotion than others. While it has already been showed that "regular", hand-designed morphologies are not, in general, better than evolved ones [9,21], we were not able to identify a single criterion of improvement. In particular, it was not the size of the VSR, i.e., bigger (and hence stronger) robots were not in general faster—we discuss this analysis more deeply in Sect. 3.2.

Summarizing, this experiment showed that the modified morphologies are not, in general clearly worse than the original ones and that a evolving a new controller for them can make them work.

Smarter Re-Optimization. Re-optimizing a new neural controller for a modified morphology proved to be effective, according to the experiment discussed above. In a practical realization of the scenario considered in this study, however, this finding would hardly be exploitable. The requirement of fully evolving from scratch a neural controller for a morphology resulting from self-assembly that just slightly differ from the expected one would make self-assembly a fragile component of the adaptation of the robotic ecosystem. On the other hand, every skill acquired by optimizing the controller for the original morphology would be dropped entirely, hence wasting previous optimizations effort.

To address this point, we explored the possibility of starting the re-optimization from the neural controller evolved for the original morphology instead of starting from scratch. In detail, for each one of the 10 controllers evolved for each original morphology, we took each of the 10 corresponding variations and performed an evolutionary optimization in which the initial population was not entirely random, but partly built based on the original controller $\boldsymbol{\theta}^\star$. In particular, we built $\frac{1}{2}n_{\text{pop}}$ individuals randomly, i.e., by sampling $U(-1,1)$, $\frac{1}{2}n_{\text{pop}}-1$ individuals by applying the Gaussian mutation to $\boldsymbol{\theta}^\star$, and finally included $\boldsymbol{\theta}^\star$ itself—that is, we *seeded* the initial population with $\boldsymbol{\theta}^\star$ as done in [7,19]. For reducing the computational effort, we considered only the small morphologies with $\delta = 1$.

Figure 3 shows the results of this experiments. It shows how v_x on the modified morphologies changes during the evolution when starting from a random initial population (blue line) or from a population seeded with $\boldsymbol{\theta}^\star$ (green line). For comparison, the figure also shows the evolution of v_x on the original morphology (red line). This is computed on 10 values, differently than the other two lines that are computed on 100 values, and results hence less smooth.

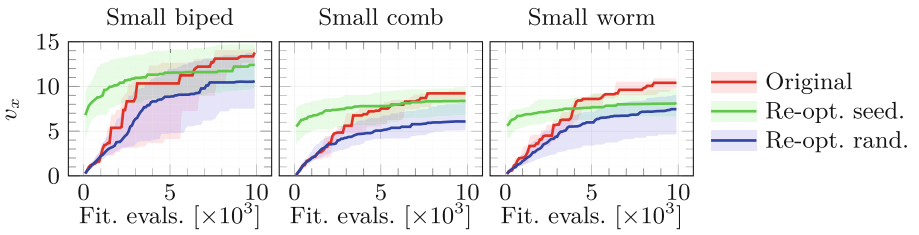

Fig. 3. Controller effectiveness v_x (median and interquartile range) during the evolution with the original morphology (red), a modified morphologies with $\delta = 1$ with re-optimized controller with random (blue) or seeded (green) population. (Color figure online)

The main finding arising from Fig. 3 is that re-using the controller evolved for the original morphology is very effective. For the biped and worm morphologies, v_x for the best individual at the first iteration of the EA with the seeded initial population is very close to the random initial population case after 5000 fitness evaluations. For the comb, the initial seeded best is on par with the final best

of the random case. In other words, with the seeded population, one can obtain an effective neural controller even with a very short re-optimization, i.e., much more efficiently.

A secondary, yet important finding, is that not only the seeded population positively impacts on efficiency of the optimization, but it also allows to obtain more effective controllers. For the biped and comb morphologies, the re-optimized controllers on the modified morphologies are not statistically significantly worse than the original controller on the original morphology.

Overall, this experiments show that a short re-optimization of the controller is enough to re-align it to the modified body, hence making the VSR ecosystem truly capable of adaptation and robustness.

Discovery of More Effective Morphologies. While analyzing the raw results for the discussion of Sect. 3.2, we discovered that for some modified morphologies, the neural controllers resulting from re-optimization (from random initial population) achieved sharply better v_x than those obtained with the original morphology.

In Fig. 4 we present the v_x obtained for all the variants of the small biped morphology, including the original, unmodified one. We chose this case because it was the one resulting, in general, in the most effective gaits. The figure shows one point for each of morphology, with coordinates given by its v_x (upon re-optimization, on the x-axis) and its size (number of voxels, on the y-axis) and color based on the morphology: green for the original one, gray for other cases, and blue and red for the Talos and horse variants.

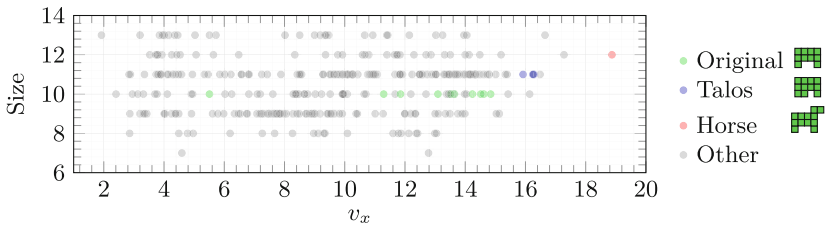

Fig. 4. Controller effectiveness v_x vs. size (num. of voxels) for all the 100 small biped morphologies, each with a re-optimized controller.

The latter two morphologies exhibited larger v_x values that were easily explainable by visually inspecting their behavior. The Talos variant, which we named after the main character of the book "Lo scudo di Talos" [6], had a thicker rear leg that conferred it greater strength for hopping faster. The horse variant used its "head", moved in counter-phase with other parts of its body, to balance the overall movement and hence advance faster.

In order to validate this casual observation, we performed 10 evolutionary runs for each of the two variants and found that their suitability for locomotion was systematic. The evolved controllers obtained a better v_x than the one for

the original biped, with statistically significant differences. We deem this finding particularly promising, in perspective. If new, more effective morphologies can be discovered "by chance" by re-optimizing existing controllers from morphologies that present small, erratic variations, then there is an opportunity for a further adaptation that propagates these small, yet positive advancements, hence making the robotic ecosystem, as a whole, more adaptive. We plan to investigate this possibility more deeply in future works.

Finally, Fig. 4 also shows that there is no correlation between the size of the VSR and v_x ($R^2 \approx 0.005$). In other words, for the biped variants, it does not hold that bigger morphologies result in faster robots, despite having, in principle, greater available strength.

4 Concluding Remarks

We considered the case of Voxel-based Soft Robots (VSRs) and experimentally characterized the impact of small morphology variations on the effectiveness of evolved neural controllers. We found that even small variations are highly detrimental for controller effectiveness, i.e., robots with a controller evolved for a morphology are much smaller if applied on morphologies differing in one or more voxels. By analyzing robot behaviors, we motivated this finding in terms of body-brain misalignment: since VSRs are soft, their bodies exhibit a rich dynamics which plays a key role in determining the behavior.

We also experimented with the re-optimization of controllers for the modified morphologies and found that it is both efficient and effective if it is seeded with the controller evolved for the original morphology. As an aside, we discovered that random morphology variants can give rise to morphologies that are more effective than the original ones. We believe this constitute an opportunity that can potentially be exploited for making the full process (automated assembly, possibly with errors, and robot life with evaluation being the key steps) more adaptable. We plan to investigate this possibility in future works.

References

1. Corucci, F., Cheney, N., Giorgio-Serchi, F., Bongard, J., Laschi, C.: Evolving soft locomotion in aquatic and terrestrial environments: effects of material properties and environmental transitions. Soft Rob. **5**(4), 475–495 (2018)
2. Faiña, A.: Evolving modular robots: challenges and opportunities. In: ALIFE 2021: The 2021 Conference on Artificial Life. MIT Press (2021)
3. Hale, M., et al.: The are robot fabricator: how to (re) produce robots that can evolve in the real world. In: International Society for Artificial Life: ALIFE2019, pp. 95–102. York (2019)
4. Li, S., et al.: Scaling up soft robotics: a meter-scale, modular, and reconfigurable soft robotic system. Soft Rob. **9**(2), 324–336 (2022)
5. Malley, M., Haghighat, B., Houe, L., Nagpal, R.: Eciton robotica: design and algorithms for an adaptive self-assembling soft robot collective. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), pp. 4565–4571. IEEE (2020)

6. Manfredi, V.M.: Lo Scudo di Talos. Edizioni Mondadori, Milan (2013)
7. Medvet, E., Bartoli, A., De Lorenzo, A., Fidel, G.: Evolution of distributed neural controllers for voxel-based soft robots. In: Proceedings of the 2020 Genetic and Evolutionary Computation Conference, pp. 112–120 (2020)
8. Medvet, E., Bartoli, A., De Lorenzo, A., Seriani, S.: 2D-VSR-Sim: a simulation tool for the optimization of 2-D voxel-based soft robots. SoftwareX **12**, 100573 (2020)
9. Medvet, E., Bartoli, A., Pigozzi, F., Rochelli, M.: Biodiversity in evolved voxel-based soft robots. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 129–137 (2021)
10. Medvet, E., Nadizar, G., Manzoni, L.: JGEA: a modular java framework for experimenting with evolutionary computation. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion (2022)
11. Medvet, E., Nadizar, G., Pigozzi, F.: On the impact of body material properties on neuroevolution for embodied agents: the case of voxel-based soft robots. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion (2022)
12. Moreno, R., Faiña, A.: EMERGE modular robot: a tool for fast deployment of evolved robots. Front. Robot. AI **8**, 198 (2021)
13. Moreno, R., Faiña, A.: Out of time: on the constrains that evolution in hardware faces when evolving modular robots. In: Jiménez Laredo, J.L., Hidalgo, J.I., Babaagba, K.O. (eds.) EvoApplications 2022. LNCS, vol. 13224, pp. 667–682. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-02462-7_42
14. Mouret, J.B., Chatzilygeroudis, K.: 20 years of reality gap: a few thoughts about simulators in evolutionary robotics. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion, pp. 1121–1124 (2017)
15. Nadizar, G., Medvet, E., Miras, K.: On the schedule for morphological development of evolved modular soft robots. In: Medvet, E., Pappa, G., Xue, B. (eds.) Genetic Programming. EuroGP 2022. Lecture Notes in Computer Science (Part of EvoStar), vol. 13223, pp. 146–161. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-02056-8_10
16. Nadizar, G., Medvet, E., Nichele, S., Pontes-Filho, S.: Collective control of modular soft robots via embodied Spiking Neural Cellular Automata. arXiv preprint arXiv:2204.02099 (2022)
17. Peck, R.H., Timmis, J., Tyrrell, A.M.: Self-assembly and self-repair during motion with modular robots. Electronics **11**(10), 1595 (2022)
18. Pfeifer, R., Gómez, G.: Morphological computation-connecting brain, body, and environment. In: Creating Brain-Like Intelligence, pp. 66–83. Springer, Cham (2009)
19. Pigozzi, F., Tang, Y., Medvet, E., Ha, D.: Evolving modular soft robots without explicit inter-module communication using local self-attention. In: Proceedings of the Genetic and Evolutionary Computation Conference (2022)
20. Salvato, E., Fenu, G., Medvet, E., Pellegrino, F.A.: Crossing the reality gap: a survey on sim-to-real transferability of robot controllers in reinforcement learning. IEEE Access **9**, 153171–153187 (2021)
21. Talamini, J., Medvet, E., Nichele, S.: Criticality-driven evolution of adaptable morphologies of voxel-based soft-robots. Front. Robot. AI **8**, 673156 (2021)
22. Zahedi, K., Ay, N.: Quantifying morphological computation. Entropy **15**(5), 1887–1915 (2013)