

# On the Impact of Body Material Properties on Neuroevolution for Embodied Agents: the Case of Voxel-based Soft Robots

## ABSTRACT

Artificial agents required to perform non-trivial tasks are commonly controlled with Artificial Neural Networks (ANNs), which need to be carefully fine-tuned. This is where ANN optimization comes into play, often in the form of Neuroevolution (NE). Among artificial agents, the *embodied* ones are characterized by a strong body-brain entanglement, i.e., a strong interdependence between the physical properties of the body and the controller. In this work, we aim at characterizing said interconnection, experimentally evaluating the impact body material properties have on NE for embodied agents. We consider the case of Voxel-based Soft Robots (VSRs), a class of simulated modular soft robots which achieve movement through the rhythmical contraction and expansion of their modules. We experiment varying several physical properties of VSRs and assess the effectiveness of the evolved controllers for the task of locomotion, together with their robustness and adaptability. Our results confirm the existence of a deep body-brain interrelationship for embodied agents, and highlight how NE fruitfully exploits the physical properties of the agents to give rise to a wide gamut of effective and adaptable behaviors.

## CCS CONCEPTS

• **Computing methodologies** → *Evolutionary robotics*; • **Theory of computation** → *Evolutionary algorithms*; • **Computer systems organization** → *Neural networks*.

## KEYWORDS

Evolutionary Robotics, Embodied Intelligence, Neuroevolution, Soft Robotics, Adaptability

## ACM Reference Format:

Eric Medvet, Giorgia Nadizar, and Federico Pigozzi. 2022. On the Impact of Body Material Properties on Neuroevolution for Embodied Agents: the Case of Voxel-based Soft Robots. In *Genetic and Evolutionary Computation Conference Companion (GECCO '22 Companion)*, July 9–13, 2022, Boston, MA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3520304.3533967>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

GECCO '22 Companion, July 9–13, 2022, Boston, MA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9268-6/22/07...\$15.00

<https://doi.org/10.1145/3520304.3533967>

## 1 INTRODUCTION

Embodied agents behave by performing actions on the environment based on the perception they acquire from the environment itself. A common assumption is that the *brain* of the agent takes decisions about the actions to be performed, whereas the *body* actuates actions and acquires perceptions. While this view might suggest that intelligence resides only in the brain, the embodied cognition paradigm states that intelligence is instead both in the brain and in the body of the agent [25]. What this paradigm does not say, though, is how much intelligence is in the brain and, hence, how much remains in the body—whatever intelligence might be; here, we mean the ability of the agent to exhibit complex behaviors that are useful for achieving some goal. In the extreme case, intelligence may reside entirely in the body. Examples exist both in the domain of artificial embodied agents, as bipedal passive walker robots [32], and in that of living, biological creatures, as the case of *Trichoplax* [3], a simple animal that does not have a nervous system, yet exhibits complex behaviors.

From another point of view, the embodied cognition paradigm suggests that both brain and body have the capacity for hosting intelligence, but it does not say in which proportion. Nevertheless, when facing the problem of searching the design space of artificial embodied agents, the focus is mainly on the brain, implicitly assuming that it has the largest capacity for hosting intelligence.

In many practical cases, the brain of an embodied agent is modeled as an Artificial Neural Network (ANN), because ANNs have good approximation capabilities and are, to some degree, biologically plausible. Thus, searching for a good brain means searching in the space of ANNs: Neuroevolution (NE), the sub-field of evolutionary computation that deals with ANNs, is an effective methodology for searching this space. Intuitively, the embodied cognition paradigm does impact on the ability of NE to find a good brain: the larger the proportion of intelligence that can be hosted in the brain, rather than in the body, the larger the search space and, hence, the longer the search and the more likely the existence of a good solution in the space. Unfortunately, this intuition is far too trivial to be useful in practice. The outcome of the search is more enigmatically determined by the structure of the search space and its relation with the goal of the search; understanding this interconnection involves analyzing the *fitness landscape* [29], that is induced mainly by the representation of the solutions and by the fitness function. In the case of NE on embodied agents, the fitness landscape is also implicitly impacted by the body: that is, for a given representation and a given task, the body shapes the fitness landscape.

In this paper, we attempt to experimentally characterize how the body shapes the fitness landscape of NE for a specific, but relevant, case of embodied agents: Voxel-based Soft Robots (VSRs) [11, 16].

VSRs are aggregates of soft modules that can change their volume in response to brain-determined control values: when equipped with sensors, VSRs employ ANNs as brains and, hence, can be optimized by means of NE [36]. VSRs constitute an optimal test-bed for this investigation for two reasons. First, since they are modular, there is a great freedom in their design: the brain, the body, and even the sensory apparatus [10] can be optimized. Second, since they are soft, their body has a rich dynamics; that is, the portion of intelligence that the body can host is likely significant. With VSRs, there is hence the potential for shaping the fitness landscape by varying the properties of the body.

We focus on the material constituting VSRs body modules, i.e., voxels. We consider three properties: the power of the material, i.e., to which degree it can expand and contract, its stiffness, and its friction with respect to other bodies as, e.g., the ground. By considering simple quantitative properties we facilitate the characterization of their impact on NE. We systematically study the impact of the three properties on the search *efficiency* and *effectiveness* of NE, i.e., on how long it takes to find a good brain and how good the found brain is. We perform experiments with three (simulated) VSR shapes using a simple evolutionary strategy as NE, that we use for optimizing the VSRs for the task of locomotion. We find that power and stiffness do impact on effectiveness: evolved VSRs are faster and NE is able to exploit the more favorable search space. For friction, we observe a less straightforward relation: NE effectiveness depends also on the shape of the body. Concerning efficiency, our results suggest that the impact of these body material properties is negligible. Beyond search efficiency and effectiveness, we analyze the behavior of the evolved VSRs and find that it significantly depends on material properties. Finally, we also analyze how the properties impact on the adaptability of evolved VSRs with respect to environmental conditions (the ground on which the robots move) and material they are built of, when modified after the evolutionary optimization.

The rest of this paper is organized as follows. In Section 2 we briefly survey previous works that are relevant to our study. In Section 3 we describe VSRs, their mechanical model for the body, and the way we use ANNs to control them. In Section 4 we describe the variant of NE we employed in the experiments, the shape of VSRs we optimized, and the material properties we varied. In Section 5 we present the experimental results and discuss them. Finally, in Section 6, we draw the conclusions.

## 2 RELATED WORKS

There exist a few previous works that tackled a research question similar, yet not identical, to the one addressed in our paper. Corucci et al. [6] explicitly studied the impact of material properties on the evolution of VSRs, by extending [4], but their VSRs employ open-loop controllers that are not based on ANNs: thus, they do not focus on NE. Talamini et al. [37] proposed to use the criticality—a measure of the complexity of a dynamical system—of a VSR body as a predictor of its the adaptability to different tasks: they did use ANN-based brains to validate this hypothesis and, hence, NE for optimizing them. However, they did not consider material properties and did not explicitly investigate the impact of criticality on NE efficiency and effectiveness.

More broadly, our study is related to the possibility of relying on optimization for the design of both body and brains of robots [30]. Howard et al. [12] highlighted the potential benefits of evolutionary optimization for the automatic design of intelligent embodied agents at all levels, from materials to overall designs, possibly including morphological development plans [19, 39]. Indeed, experiments on concrete cases showed that co-evolving body and brain of robots can be particularly successful [17, 24]. The latter work focused on VSRs and also found that, depending on the representation and the Evolutionary Algorithm (EA), co-evolution of body and brain can result in a diverse set of good designs.

Finally, we remark that here we study simulated robots. Clearly, material properties play an even more important role when working with real robots. At least, differences between simulated materials used for optimizing the robot in simulation and real materials used for fabricating it might result in a reduced ability of the robot to perform its task. This problem is known as reality gap and affects many kinds of adaptation, including evolutionary optimization [8, 13] and reinforcement learning [34, 38]. However, bio-inspired optimization techniques still remain a viable option for facing hard optimization problems, such as those involving the automatic designs of soft robots [28]. In principle, our study could be extended to consider NE applied to real soft robots.

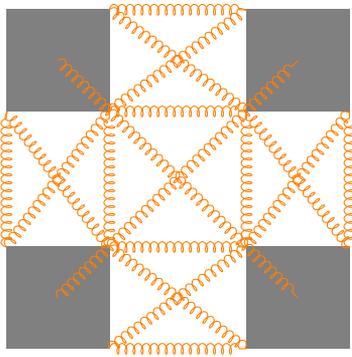
## 3 BACKGROUND ON VSRs

Voxel-based Soft Robots (VSRs), aggregations of elastic cubic blocks, are a kind of embodied agents. Actuation takes place by contracting or expanding the voxels volume; the overall symphony of volume contractions and expansions allows for the emergence of behavior at the agent level. Hiller and Lipson [11] first introduced VSRs together with a fabrication method. In this work, we employ a 2D simulated variant of VSRs [16], in discrete time and continuous space. While discarding one dimension makes the simulations less realistic, it eases modeling and reduces computational burden.

A VSR is composed of a *body* (i.e., the morphology) and a *brain* (i.e., the controller). The body, and its mechanical model, is the central piece of study of this work: we detail them thoroughly in the next section.

### 3.1 Mechanical model of VSRs body

While different approaches exist to model soft bodies [14, 31], we resort to spring-and-damper systems, which are simple yet effective and widely studied [7]. A VSR body consists in a number of equal modules organized in a 2D grid. Each cell of the grid may contain a *voxel*, the fundamental building block of a VSR; and the way voxels are assembled in the grid is the *shape* of the body. The mechanical model of the VSR body then rests on the mechanical model of the voxel. In the mechanical model of Medvet et al. [16], each voxel is a square of size  $l = 3$  m with four bodies of rigid material placed at the corners—adjacent bodies belonging to different adjacent voxels in the shape are welded together with weld joints not allowing relative displacement. Each body weighs 1 kg, is 1.05 m in side length, and has a friction  $\mu_k$ , describing how it resists the relative motion of other bodies (e.g., the ground) sliding on it. Within a voxel, we join pairs of bodies with spring-and-damper systems: each spring has a *frequency*  $f$ , in Hz), describing how many oscillations it executes



**Figure 1: The mechanical model of the voxel. Gray squares are rigid bodies and orange wiggly strings are springs.**

per second, and a damping ratio  $d$ , dimensionless, describing how rapidly oscillations decay over time. Since springs have no mass, the corner bodies are what endows the voxel with a mass and a sense of an “embodiment”, including the ability to react to forces and collide with other objects. Springs, by oscillating, allow the voxel to alter its area and thus endow it with softness. We present a schematic view of the voxel model in Figure 1.

A VSR agent acts by changing the area of the composing voxels over time. In the mechanical model of [16], actuation happens by varying the resting length of springs. Given a control signal  $a \in [-1, 1]$  for a voxel of side  $l$ , the resting length of the voxel springs changes instantaneously such that the voxel side becomes  $l' = \sqrt{l^2(1 - a\rho_a)}$ , with  $\rho_a > 0$  being the *active range*.  $\rho_a$  describes the maximum rate of increase or decrease of the voxel area:  $\rho_a = 0$  corresponds to no change, regardless of  $a$ ,  $\rho_a = 0.3$  corresponds to a  $\pm 30\%$  change of area with  $a = \mp 1$ . We remark that the control signal  $a$  causes the rest length of the springs, not the actual length, to change instantaneously: the actual length changes gradually depending on  $f$  and  $d$ .

In this work, we investigate how the four body properties of active range  $\rho_a$ , spring  $f$  and  $d$ , and friction  $\mu_k$  impact on NE. In fact, by varying these four, we impact on the properties of the material constituting the soft body of the VSR. Active range  $\rho_a$  determines the maximum amount of kinetic energy that actuation can produce, spring  $f$  and  $d$  impact on the softness of the voxel, while friction  $\mu_k$  determines how smooth is the body surface to the contact with other objects. More broadly,  $\rho_a$  models the power of the voxel,  $f$  and  $d$  model its material softness (or stiffness), and  $\mu_k$  models the relative friction of the voxel and the ground (when the VSR interacts with the ground, as in the locomotion problem we deal with in this paper).

For NE to take place, there must be a ANN-based controller, or neural controller [36]. Neural controllers are closed-loop controllers and require an input, usually in the form of sensor readings. To this end, we equip voxels with *sensors*. We use touch, area, and velocity sensors. The touch sensor perceives whether the voxel is touching other objects (e.g., the ground) or not and outputs 1 if yes, 0 otherwise. The area sensor perceives the ratio of the current area to the rest area. Velocity sensors perceive the  $x$ - and  $y$ -velocity of the voxel center of mass. We apply soft normalization to all sensor

readings to ensure they lie in  $[0, 1]$ . Finally, we apply Gaussian noise of mean 0 and standard deviation 0.01 to every sensor reading to simulate real-world sensor noise.

## 3.2 Neural network-based controller for VSRs

Since we focus on the impact of body properties on NE, we adopt ANNs for the controller. In particular, for a given VSR of  $n$  voxels, there is one centralized ANN controller that, at every time step, receives as input the concatenation  $s \in [0, 1]^q$  of  $q$  sensor readings from all voxels and outputs  $a \in [-1, 1]^n$  as the actuation values for all voxels. As a result, there are as many input neurons as the number of sensor readings of all voxels and one output neuron for every voxel.

Taking inspiration from previous works [17, 21, 26, 36], we set one intermediate layer, with as many neurons as the number of input neurons and, to ensure activations lie in  $[-1, 1]$ , tanh as activation function.

The weights and biases of the ANN controller form the vector  $\theta$  of parameters that we subject to evolutionary optimization; we denote by  $p = |\theta|$  the number of parameters of the ANN.

## 4 METHODS

### 4.1 Shapes

In order to understand how body properties affect NE across different morphological conditions, we experiment with three different shapes.

In particular, we consider a  $4 \times 3$  (size of the voxels grid enclosing the VSR) rectangle with a  $2 \times 1$  rectangle of missing voxels at the bottom-center, that we call *biped* , a  $7 \times 2$  rectangle with empty voxels at the odd  $x$  positions in the bottom row, that we call *comb* , a  $5 \times 2$  rectangle, that we call *worm* . Thus, we experiment with fundamental shapes like the worm, that has no limbs, as well as more sophisticated ones like the comb, that has many limbs, with the biped being an intermediate step between the two.

Concerning the sensory apparatus of the VSRs, we put area sensors in each voxel, velocity sensors in the voxels of the top row, and touch sensors in the voxels of the bottom row.

The ANN input size  $|s|$  is thus 20, 29, and 25 for the biped, the comb, and the worm, respectively. As a result, the number  $p$  of parameters is 630, 1 200, and 910, respectively.

### 4.2 Evolutionary algorithm

We cast the problem of NE as a numerical optimization problem in the space of parameters  $\theta$  of the ANN controller; indeed, NE is capable of optimizing neural controllers for soft robots [27]. In particular, we adopt a simple variant of Evolution Strategies (ES) [2]. Indeed, ES have achieved state-of-the-art results for continuous control tasks requiring NE [33] and have already succeeded in optimizing neural controllers for VSRs [9, 20, 22]. We describe our EA in Algorithm 1.

The EA evolves a fixed-size population of  $n_{\text{pop}}$  individuals, i.e., numerical vectors  $\theta$  of dimension  $p$ . At first,  $n_{\text{pop}}$  individuals are uniformly initialized by independently sampling  $[-1, +1]$  for each vector element. Then, at every iteration (i.e., generation), the fittest quarter of the individuals is chosen as parents.  $n_{\text{pop}} - 1$  children

```

1 function evolve():
2    $P \leftarrow \emptyset$ 
3   foreach  $i \in \{1, \dots, n_{\text{pop}}\}$  do
4      $P \leftarrow P \cup \{0 + U(-1, 1)^p\}$ 
5   end
6   foreach  $g \in \{1, \dots, n_{\text{gen}}\}$  do
7      $P_{\text{parents}} \leftarrow \text{bestIndividuals}(P, \lfloor \frac{|P|}{4} \rfloor)$ 
8      $\mu \leftarrow \text{mean}(P_{\text{parents}})$ 
9      $P' \leftarrow \{\text{bestIndividuals}(P, 1)\}$ 
10    while  $|P'| < n_{\text{pop}}$  do
11       $P' \leftarrow P' \cup \{\mu + N(0, \sigma)^p\}$ 
12    end
13     $P \leftarrow P'$ 
14  end
15  return bestIndividuals(P, 1)
16 end

```

**Algorithm 1:** The simple ES used in our experiments.

are born from the parents, each one obtained by adding a Gaussian noise sampled from  $N(0, \sigma)$  to each of the  $p$  elements of the element-wise mean  $\mu$  of all parents. Finally, the generated offspring is merged with the fittest parent to form the population for the next generation. The algorithm iterates this process for  $n_{\text{gen}}$  generations and the fittest individual is returned at the end.

After preliminary experiments, we set  $n_{\text{pop}} = 36$ ,  $n_{\text{gen}} = 285$  (corresponding to 10 000 fitness evaluations), and  $\sigma = 0.35$ . We verified that, for the chosen value of  $n_{\text{pop}}$  and  $n_{\text{gen}}$ , evolution was in general capable of converging to a solution, i.e., longer evolutions would have resulted in negligible fitness improvements.

### 4.3 Body properties values

We focus on the impact that the four properties, namely  $\rho_a$ ,  $f$ ,  $d$ , and  $\mu_k$ , have on NE. To this end, we sample several values for each of the properties to experiment with. The sample space for each property is orthogonal with respect to the other properties; in other words, we test the values of a property while fixing the other properties to default values. The reason for this is two-fold: (a) simplifying the investigation by excluding any interaction among body properties; (b) easing the computational requirements by reducing the number of combinations.

Relying on our previous knowledge, we set the default values to be 0.2 for  $\rho_a$ , 8 Hz for  $f$ , 0.3 for  $d$ , and 10 for  $\mu_k$ . After preliminary experiments, we set the sample ranges to be  $[0.1, 0.3]$  for  $\rho_a$ ,  $[3, 10]$  for  $f$ ,  $[0.1, 0.99]$  for  $d$ , and  $[0.05, 25]$  for  $\mu_k$ . For each of them, we sample 8 values at regular intervals for  $\rho_a$ ,  $f$ ,  $d$ , and 8 values at regular intervals on a logarithmic scale for  $\mu_k$ .

## 5 EXPERIMENTS AND DISCUSSION

We performed an experimental campaign with the aim of verifying what impact the key body properties have on: (a) effectiveness; (b) search efficiency; (c) robustness with respect to terrain and parameter changes.

To this end, we experimented with the three shapes of Section 4.1 and the EA of Section 4.2 on a locomotion task, where the goal is

to travel as fast as possible along the positive  $x$  direction over an amount of simulated time  $t_{\text{final}}$ , on a flat surface. The fitness of an individual of Algorithm 1 is then the average velocity  $v_x$  computed by considering the position of the center of mass of the VSR at the beginning and at the end of the simulation. We set  $t_{\text{final}} = 30$  s. Locomotion is not only a widespread task in evolutionary robotics [23], but is also fundamental enough to dispense with the impact that factors other than body properties, including the environment and other agents, might have on NE.

For the implementation, we built a software tool on top of 2D-VSR-Sim [16], for the VSR simulation, and JGEA [18]<sup>1</sup>, for the evolutionary optimization. For the simulation, we set frequency to 60 Hz and all other parameters to default values. We made the code to replicate the experiments publicly available at <https://github.com/giorgia-nadizar/BodyParamsInfluenceOnNE>. We performed 10 independent evolutionary runs for each combination of  $\rho_a$ ,  $f$ ,  $d$ ,  $\mu_k$  values, varying the random seed for the EA. We remark that simulations are deterministic.

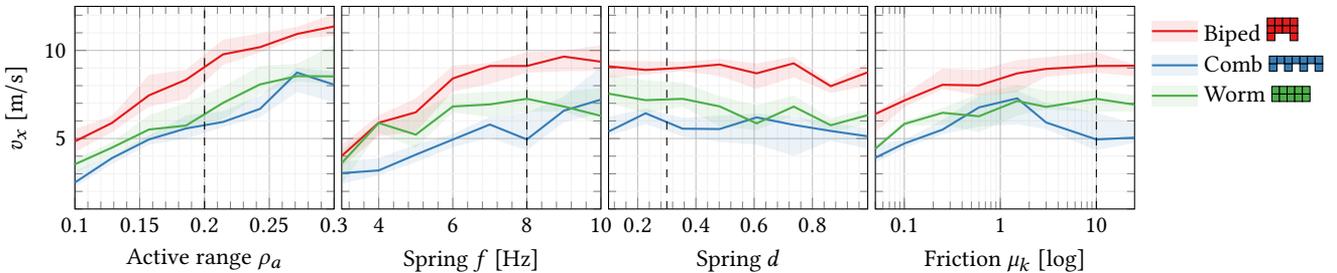
### 5.1 Impact on effectiveness

At first, we are interested in studying the impact of body properties on effectiveness of NE. We measure effectiveness as the fitness of the best individual found in each evolutionary run. Then, the performance index is  $v_x$ . We summarize the results in Figure 2, which shows  $v_x$  in terms of median and interquartile range, across the value for the four body properties. For each body property, we mark its default value with a dashed vertical line.

From the figure, we can gain two interesting insights: (a) some properties do affect NE effectiveness, while others do not; (b) the impact of body properties does not qualitatively vary with different shapes.

As far as the former insight is concerned,  $\rho_a$  and  $f$  greatly affect effectiveness measured in  $v_x$ ; the same is arguably not true for  $d$ , while it is mildly true for  $\mu_k$ . We computed the Pearson correlation coefficient and found it to be  $\approx 0.3$  for both  $\rho_a$  and  $f$ , meaning a positive linear correlation between  $v_x$  for both  $\rho_a$  and  $f$ . For  $d$  and  $\mu_k$ , on the other side, it is  $\approx 0.0$ , meaning no correlation. The higher the active range  $\rho_a$ , the more voxels can increase or decrease their area; as a result, area contraction and expansion can produce a greater amount of kinetic energy, that, intuitively, helps in generating effective locomotion gaits. At the same time, the higher the spring  $f$ , the stiffer the body. It means that, in our simulation, stiffness is key to generating effective gaits. Intuitively the impact of stiffness is two-fold: (a) stiff voxels are more resilient to external forces, namely, forces caused by other adjacent voxels; (b) stiff voxels react more promptly to area changes triggered by the control signal, since springs elongate or shorten at higher speed—we recall that the actuation signal impacts instantly on the rest length of the spring, but the actual length varies depending on the spring properties  $f$  and  $d$ . While we cannot tell apart the two effects and their quantitative impact on  $v_x$ , we believe that the latter might be stronger. Indeed, not much differently from  $\rho_a$ , this promptness to area change determined by  $f$  can be associated with the ability of the material to efficiently convert energy into action. We believe the finding to be relevant, as it ties VSRs (and soft bodies in general) to

<sup>1</sup><https://github.com/ericmedvet/jgea>



**Figure 2: Median and interquartile range of the velocity  $v_x$  of the best individuals found in each of the 10 evolutionary runs, for 8 values of four body properties while fixing the other properties to default values (dashed lines). Active range  $\rho_a$  and spring  $f$  do affect effectiveness.**

“embodied energy” [1], according to which energy in autonomous agents could be best exploited by embodying it inside the materials rather than, e.g., through battery packs.

Spring  $d$ , on the other side, does not seem to have an impact or, if it has, NE “equalizes” it (which does not happen for the other properties). We hypothesize the reason to be that  $d$  does not affect the amount of energy of the body (modulated by active range  $\rho_a$  and spring  $f$ ). Finally, for friction  $\mu_k$ , it might be the case that the higher  $\mu_k$ , the more effective the VSR at generating locomotion gaits; nevertheless, the improvements plateau off after a while. We believe the reason to be that, initially, increasing friction (thus making the rigid bodies of the voxels less “slippery”) is beneficial as walking with rough voxels is easier than walking with slippery voxels. Snakes, for example, exploit the friction of their scales against contact points to ease locomotion [15]. If friction increases too much, though, it becomes more of a handicap as it requires the VSR to generate more kinetic energy to overcome it.

Regarding insight (b), although it is true that overall effectiveness differs by shape (in particular, bipeds seem to be more effective than combs and worms), the trend for the body properties does not generally change for bipeds and worms. For combs, friction has a more peculiar impact: since combs have many little limbs that must synchronize for locomotion, too much friction makes them stick to the ground to the point of disrupting their synchronization. This fact points to the generality of our results in the context of VSRs.

We remark that, despite the differences in  $v_x$ , all values reported in Figure 2 correspond to effective locomotion behaviors. Indeed, we visually inspected the best evolved individuals and found them to be generally adapted to locomotion. At the same time, their gaits allowed us to appreciate the impact of body properties on NE. For example, bipeds with low  $\rho_a$  hop on their feet, while bipeds with high  $\rho_a$  use their rear leg to jump forward as rabbits do. Individuals with low spring  $f$  are “flabby” and trudge forward as if they were weary. Individuals with low  $\mu_k$  struggle to move forward as they slide on the ground. We show sample time-lapses for VSRs evolved with different values of  $\rho_a$  and  $f$  in Figure 3, while videos can be found at <https://giorgia-nadizar.github.io/BodyParamsInfluenceOnNE/>.

As a final comment, we find that NE can always optimize ANNs effective for the task of locomotion, regardless of the value of the body properties; this result is not trivial, as body properties have a clear effect on the amount of energy available to the VSR. NE

thus optimizes brains that are adapted to the body enough to be effective.

## 5.2 Impact on search efficiency

Second, we are interested in studying the impact of body properties on NE search efficiency, as it provides a rough estimate of how the fitness landscape changes. We consider two aspects of NE search efficiency: smoothness and multimodality of the fitness landscape.

We measure smoothness in terms of how long it takes to achieve a “reasonably good” performance: as performance index, we define  $g_{80}$  as the earliest generation at which the best individual has a velocity of at least 80% of that of the best at the end of evolution. We report the results in Figure 4, in terms of the median and interquartile range, across the value for the four body properties.

From the figure, we can gain the following insights: (a) there is no clear impact of body properties on NE search efficiency; (b) the impact of body properties does not vary with the shape. Insight (a) tells us that, contrary to NE effectiveness (see Section 5.1), it is hard to tell which body properties do affect NE search efficiency in terms of smoothness and which others do not. In particular,  $\rho_a$ ,  $d$ , and  $\mu_k$  have no linear correlation with  $g_{80}$ . Spring  $f$ , on the other side, might be negatively correlated with  $g_{80}$ , but that happens only for the biped shape; with this exception, all body properties impact on smoothness comparatively across the three shapes, as remarked by insight (b). These facts altogether reveal that the fitness landscape does not significantly change varying the body properties.

As far as multimodality is concerned, we say the fitness landscape is multimodal if the solutions found in different evolutionary runs are diverse, as that is arguably a sign of different local minima. To this end, we report in Figure 5 the pairwise Euclidean distances  $d_\theta$  among the genotypes  $\theta$  of the best individuals of the 10 evolutionary runs, in terms of median and interquartile range.

Considering that the dimension of the genotype is roughly the same for the three shapes, the results are surprising. Distances are significantly different among the three shapes; in particular, solutions for the biped shape are, in general, less diverse. However, for the same shape, there are no significant differences by body property, with the only exception of spring  $f$  for the biped: for low values (i.e., flabby bipeds), genotypes might be more diverse.

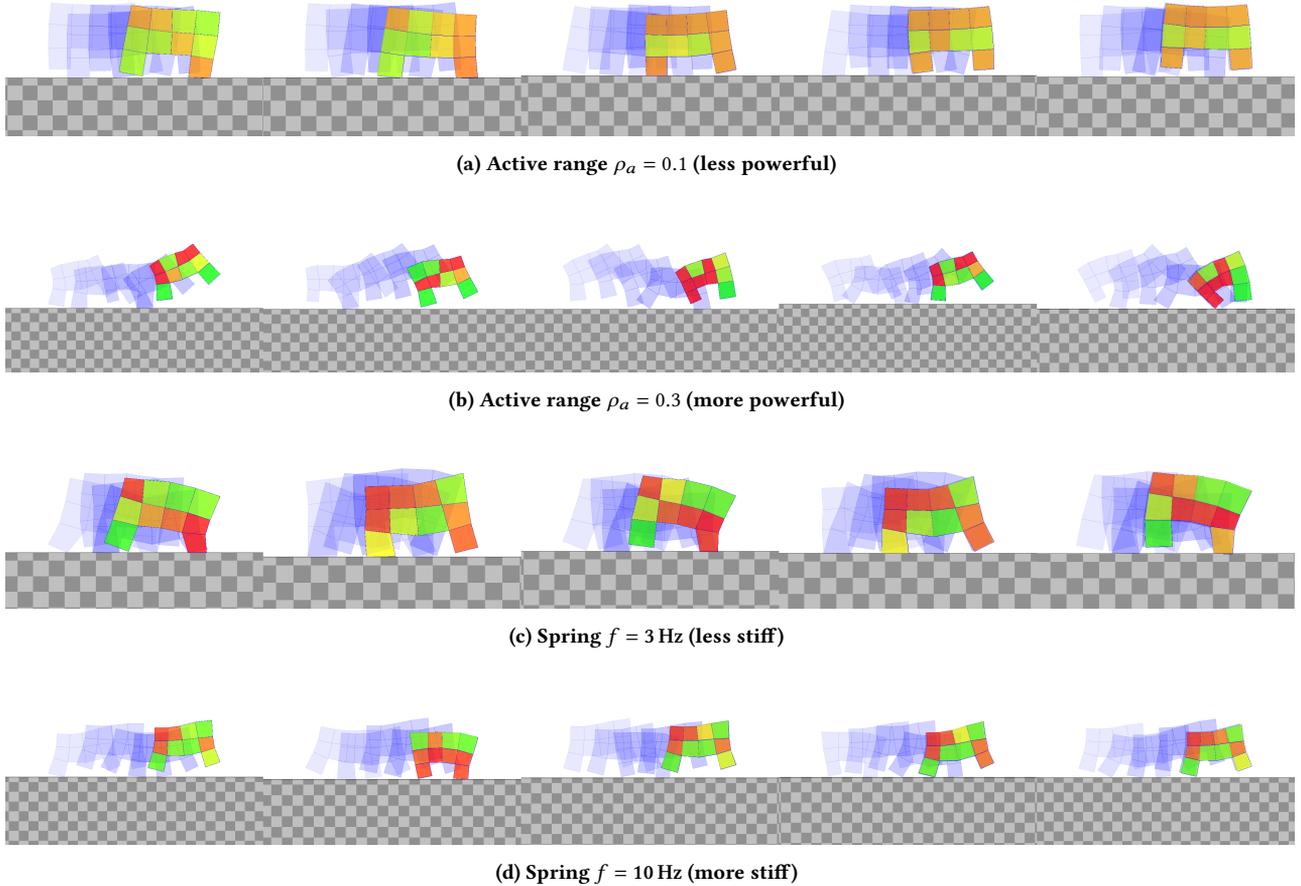


Figure 3: Time-lapses showing locomotion for two bipeds evolved with different values of active range  $\rho_a$  (in 3a and 3b) and of spring  $f$  (in 3c and 3d). Frames are taken at 1 s intervals; in each frame, the blue shades show the position of the robots at earlier 0.5 s intervals. The color of each voxel represents the current area ratio with respect to the rest area (red for contracted, i.e.,  $< 1$ , green for expanded, i.e.,  $> 1$ , yellow otherwise, i.e.,  $\approx 1$ ). For  $\rho_a$ , the higher the value, the more kinetic energy the VSR can produce: the “strongest” VSR (in 3b) jumps higher, touches the ground with a single foot, and, as a consequence, keeps a different pose. For  $f$ , the lower the value, the less stiff the body, to the point that voxels react less promptly to area changes and struggle in jumping high (in 3c).

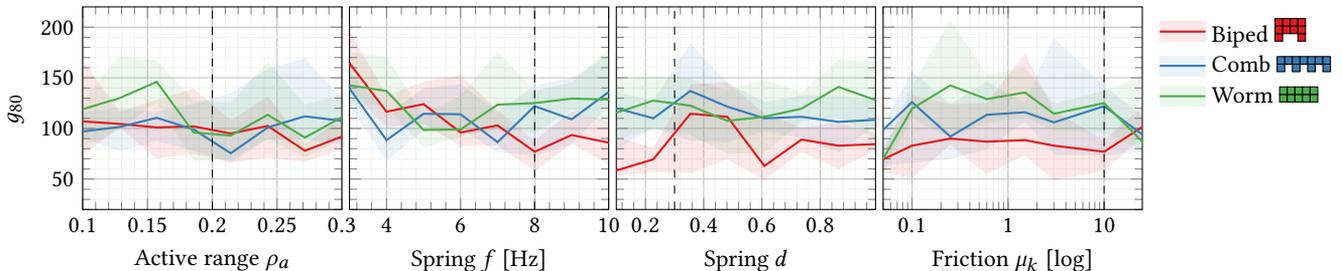
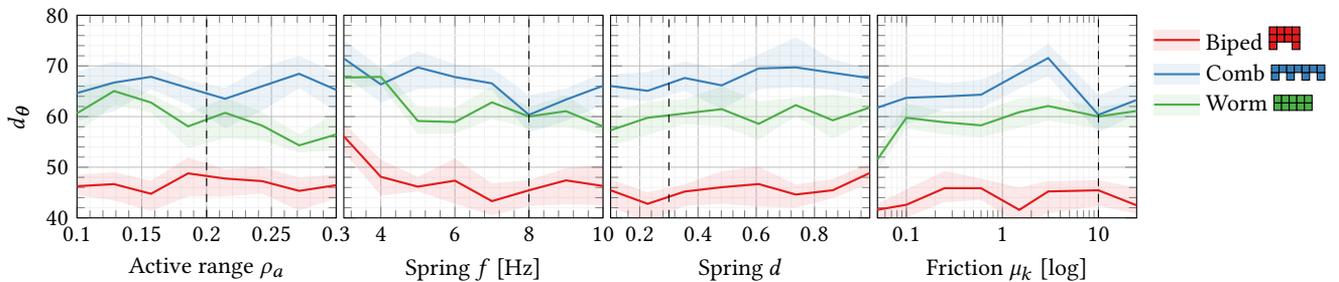


Figure 4: Median and interquartile range of number of generations  $g_{80}$  to reach 80% of fitness of the best individuals found in each of the 10 evolutionary runs, for 8 values of four body properties while fixing the other properties to default values (dashed lines). There is no clear impact of body properties on NE search efficiency in terms  $g_{80}$ .



**Figure 5: Median and interquartile range for the pairwise Euclidean distances  $d_\theta$  between the genotypes  $\theta$  of the best individuals of the 10 evolutionary runs. Best individuals are, generally, diverse; in terms of NE search efficiency, it means the fitness landscape is likely multimodal.**

We visually inspected the best evolved individuals and indeed found them to behave generally diversely, especially for the worm and comb shapes, confirming the results of Figure 5.

### 5.3 Impact on generalization

Since we deal with NE of embodied agents that interact with an environment, we are naturally interested in how body properties affect their ability to generalize to novel environments and morphological conditions. In particular, we consider generalization to two settings: terrain changes and body properties changes. For both, we take the best VSR of each evolutionary run and re-assess it in different conditions.

**5.3.1 Terrain changes.** For this setting, we re-assess the best VSR of each run on 16 new terrains, unseen during evolution: hilly with 6 combinations of heights and distances between the hills, steppy with 6 combinations of steps heights and widths, downhill with 2 different inclinations, and uphill with 2 different inclinations. As performance index, we use the ratio  $\rho_{v_x} = \frac{\bar{v}_x}{v_x}$  between the mean  $\bar{v}_x$  velocity of the best VSR over the unseen terrains and its fitness  $v_x$ , i.e., its velocity on the flat terrain. The lower  $\rho_{v_x}$ , the lower the adaptability. We show the results in Figure 6, in terms of the median and interquartile range.

Not surprisingly, all VSRs are slower, on average, on the 16 terrains used for re-assessment than on the flat terrain on which they have been evolved:  $\rho_{v_x}$  is approx 0.4 in almost all conditions. Concerning the impact of the properties, the plots of Figure 6 suggest that the only property that has an apparent effect on adaptability is the friction  $\mu_k$ . VSRs evolved with low values for the friction appear to lose less speed when facing unseen terrains, i.e., they generalize better. Interestingly, this observation holds for all the three shapes.

**5.3.2 Body properties changes.** For this setting, we re-assess the best VSR of each run changing its body properties. Specifically, we carry on two different experiments: (a) change the body property under investigation; (b) change the body properties not under investigation.

In experiment (a), we change the body property under investigation and keep the others to their default values. As performance index, we use  $v_x$  obtained by re-assessing the individual on flat terrain. To ease comparison between evolution and re-assessment,

Figure 7 reports, in the form of heatmaps,  $v_x$  for the different combinations of body property values set during evolution and during re-assessment. We do not report values explicitly, but ordered in the same way as Figures 2 and 4 to 6 (i.e., the closer the cells the closer the values).

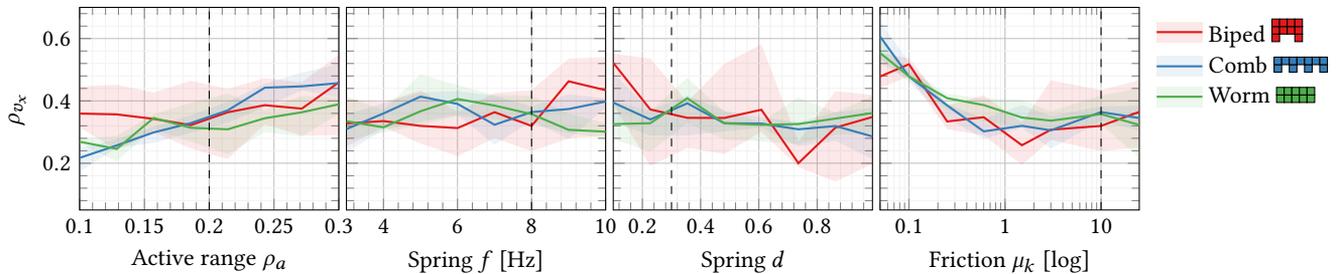
To summarize, Figure 7 corroborates our previous findings: active range  $\rho_a$  and spring  $f$  affect generalization the most, while spring  $d$  and friction  $\mu_k$  do not. In fact, the heatmaps for  $d$  and  $\mu_k$  are dark and “blurred”, meaning that there are no drastic drops in  $v_x$ . Contrarily, the heatmaps for  $\rho_a$  and  $f$  are lighter (away from the diagonal), with drastic drops in  $v_x$  moving away from the diagonal. In particular, the lower-right corner is always the lightest for a given body property and shape: it means that re-assessing with a body property (be it  $\rho_a$  or  $f$ ) value lower than during evolution results in worse  $v_x$  than re-assessing with a higher value.

Finally, as a sanity check, we notice that diagonals are darker than off-diagonal elements, meaning that re-assessing an individual with the same body conditions as evolution results in good performance.

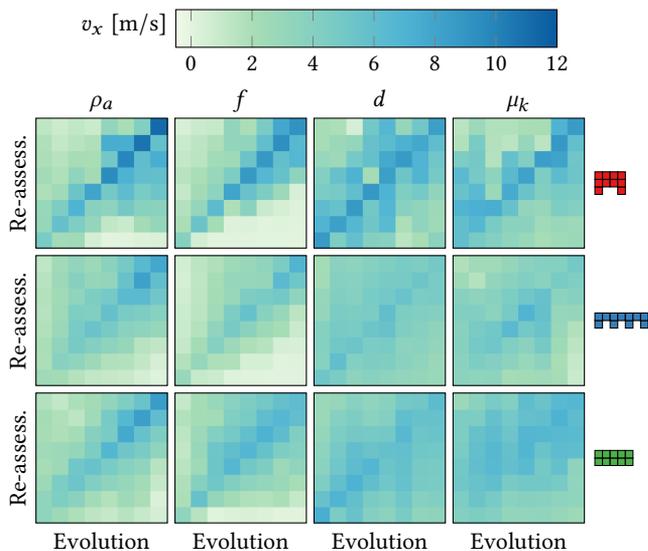
In experiment (b), we change the body properties not under investigation. As performance index, we use the ratio  $\rho_{v_x} = \frac{\bar{v}_x}{v_x}$  between the mean  $\bar{v}_x$  velocity of the best VSR over all the combinations of body property values in which the property under investigation is untouched and all the other ones have values different than default and its fitness  $v_x$ . The lower  $\rho_{v_x}$ , the lower the adaptability. Figure 8 reports the results in terms of median and interquartile range.

According to Figure 8, we discover that changing body properties not under investigation results in large drops in  $v_x$ . The values for  $\rho_{v_x}$  are  $\approx 0.4$ , that means that by changing the body properties the speed decreases on average by 60%. The figure is similar to case of terrain change, but differently from that case, the differences among shapes are more apparent. In particular, the bipeds appear to be affected more by body property changes, i.e., for that shape, NE evolves brains that are less adaptable to these changes. Intuitively, the biped shape tends to “overfit” more, as it is especially suited for locomotion, having two limbs that can be used for jumping; on the other hand, the degree to which it can actually exploit them depends more on the body dynamics.

The results of Figure 8 are relevant for this work: by virtue of the embodied cognition paradigm, changing body properties deeply alters how the brain interacts with the environment. Thus, we



**Figure 6: Median and interquartile range of the ratio  $\rho_{v_x}$  between the mean velocity  $\bar{v}_x$  over 16 unseen terrains (hilly, steepy, uphill, downhill) and the velocity  $v_x$  on the flat terrain. Friction  $\mu_k$  does affect generalization to terrain changes.**



**Figure 7: Heatmap of velocity  $v_x$  obtained by re-assessing an individual—evolved with the body property value on the  $x$ -axis—with the body property value on the  $y$ -axis. Active range  $\rho_a$  and spring  $f$  affect generalization the most.**

can conclude that generalization to unseen body properties is a difficult task for NE of embodied agents. In other words, this finding corroborates the hypothesis that the ability of these robots to do locomotion resides, in a significant portion, in the body, despite the apparent larger complexity of the brain.

## 6 CONCLUDING REMARKS

We investigated the impact of body material properties of embodied agents on Neuroevolution (NE) with the aim of experimentally characterizing the body-brain interconnection. To this end, we considered the case study of Voxel-based Soft Robots (VSRs), a form of modular soft robots, and optimized their controller for the task of locomotion by means of NE.

More in details, we experimented with three VSRs shapes, varying some material properties of the VSR constituting modules, such as the available power, the stiffness, and the friction. We considered different evaluation metrics, namely, effectiveness, search efficiency,

and generalization ability, and found some properties to be more impacting than others on the final outcomes of the NE, unveiling a non-equal body-brain entanglement across them. Although our results cannot be deemed completely general, we believe our study gives significant evidence of the relevance of the embodied cognition paradigm. In addition, having observed that the generalization ability of VSRs does not only reside in their body, we believe our work should foster the inclusion of adaptation mechanisms in the controllers of the embodied agents [9].

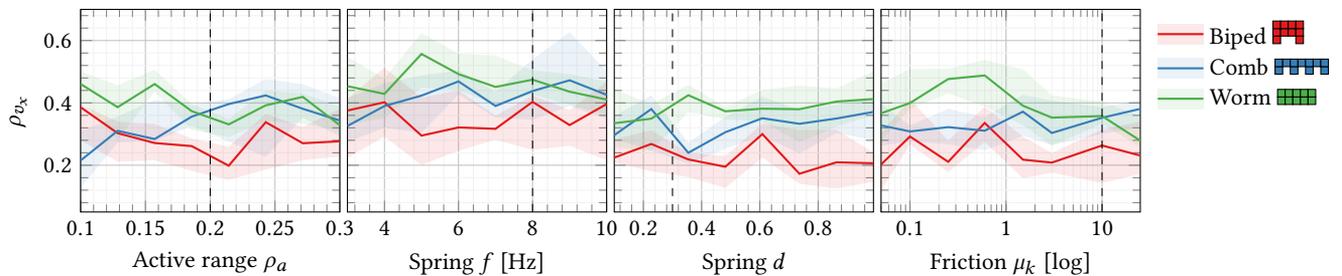
As an extension of this work, it might be noteworthy to rely on more sophisticated analysis tools, as, e.g., [35], to deepen the study of NE dynamics, or to compare the results achieved through NE with “classical” gradient based methods [5].

## ACKNOWLEDGMENTS

The experimental evaluations of this work was done on the CINECA HPC cluster under the CINECA-University of Trieste agreement.

## REFERENCES

- [1] Cameron A Aubin, Benjamin Gorissen, Edoardo Milana, Philip R Buskohl, Nathan Lazarus, Geoffrey A Slipper, Christoph Keplinger, Josh Bongard, Fumiya Iida, Jennifer A Lewis, et al. 2022. Towards enduring autonomous robots via embodied energy. *Nature* 602, 7897 (2022), 393–402.
- [2] Hans-Georg Beyer and Hans-Paul Schwefel. 2002. Evolution strategies—a comprehensive introduction. *Natural computing* 1, 1 (2002), 3–52.
- [3] Matthew S Bull, Laurel A Kroo, and Manu Prakash. 2021. Excitable mechanics embodied in a walking cilium. *arXiv preprint arXiv:2107.02930* (2021).
- [4] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. 2014. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. *ACM SIGEVOLUTION* 7, 1 (2014), 11–23.
- [5] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. 2015. The loss surfaces of multilayer networks. In *Artificial intelligence and statistics*. PMLR, 192–204.
- [6] Francesco Corucci, Nick Cheney, Francesco Giorgio-Serchi, Josh Bongard, and Cecilia Laschi. 2018. Evolving soft locomotion in aquatic and terrestrial environments: effects of material properties and environmental transitions. *Soft robotics* 5, 4 (2018), 475–495.
- [7] Cosimo Della Santina, Robert K Katzschmann, Antonio Biechi, and Daniela Rus. 2018. Dynamic control of soft robots interacting with the environment. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 46–53.
- [8] Stephane Doncieux, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E Eiben. 2015. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI* 2 (2015), 4.
- [9] Andrea Ferigo, Giovanni Iacca, Eric Medvet, and Federico Pigozzi. 2021. Evolving Hebbian Learning Rules in Voxel-based Soft Robots. (2021).
- [10] Andrea Ferigo, Eric Medvet, and Giovanni Iacca. 2022. Optimizing the Sensory Apparatus of Voxel-Based Soft Robots Through Evolution and Babbling. *SN Computer Science* 3, 2 (2022), 1–17.
- [11] Jonathan Hiller and Hod Lipson. 2011. Automatic design and manufacture of soft robots. *IEEE Transactions on Robotics* 28, 2 (2011), 457–466.



**Figure 8: Median and interquartile range of the ratio  $\rho_{v_x}$  between the mean velocity  $\bar{v}_x$  with properties (except the one under investigations) changed and the velocity  $v_x$  with properties as in evolution. Generalization to unseen body properties depends more on the shape than on body properties.**

- [12] David Howard, Agoston E Eiben, Danielle Frances Kennedy, Jean-Baptiste Mouret, Philip Valencia, and Dave Winkler. 2019. Evolving embodied intelligence from materials to machines. *Nature Machine Intelligence* 1, 1 (2019), 12–19.
- [13] Sylvain Koos, Jean-Baptiste Mouret, and Stéphane Doncieux. 2012. The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation* 17, 1 (2012), 122–145.
- [14] Cecilia Laschi, Barbara Mazzolai, and Matteo Cianchetti. 2016. Soft robotics: Technologies and systems pushing the boundaries of robot abilities. *Science Robotics* 1, 1 (2016), eaah3690.
- [15] Hamidreza Marvi and David L Hu. 2012. Friction enhancement in concertina locomotion of snakes. *Journal of the Royal Society Interface* 9, 76 (2012), 3067–3080.
- [16] Eric Medvet, Alberto Bartoli, Andrea De Lorenzo, and Stefano Seriani. 2020. 2D-VSR-Sim: A simulation tool for the optimization of 2-D voxel-based soft robots. *SoftwareX* 12 (2020), 100573.
- [17] Eric Medvet, Alberto Bartoli, Federico Pigozzi, and Marco Rochelli. 2021. Biodiversity in evolved voxel-based soft robots. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 129–137.
- [18] Eric Medvet, Giorgia Nadizar, and Luca Manzoni. 2022. JGEA: a Modular Java Framework for Experimenting with Evolutionary Computation. In *Proceedings of the genetic and evolutionary computation conference companion*.
- [19] Giorgia Nadizar, Eric Medvet, and Miras Karine. 2022. On the Schedule for Morphological Development of Evolved Modular Soft Robots. In *European Conference on Genetic Programming*. Springer.
- [20] Giorgia Nadizar, Eric Medvet, Stefano Nichele, and Sidney Pontes-Filho. 2022. Collective control of modular soft robots via embodied Spiking Neural Cellular Automata. *arXiv preprint arXiv:2204.02099* (2022).
- [21] Giorgia Nadizar, Eric Medvet, Felice Andrea Pellegrino, Marco Zulich, and Stefano Nichele. 2021. On the effects of pruning on evolved neural controllers for soft robots. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 1744–1752.
- [22] Giorgia Nadizar, Eric Medvet, Hola Huse Ramstad, Stefano Nichele, Felice Andrea Pellegrino, and Marco Zulich. 2022. Merging pruning and neuroevolution: towards robust and efficient controllers for modular soft robots. *The Knowledge Engineering Review* 37 (2022).
- [23] Stefano Nolfi and Dario Floreano. 2000. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press.
- [24] Paolo Pagliuca and Stefano Nolfi. 2020. The dynamic of body and brain co-evolution. *Adaptive Behavior* (2020), 1059712321994685.
- [25] Rolf Pfeifer and Josh Bongard. 2006. *How the body shapes the way we think: a new view of intelligence*. MIT press.
- [26] Federico Pigozzi and Eric Medvet. 2022. Evolving Modularity in Soft Robots through an Embodied and Self-Organizing Neural Controller. *Artificial life* (2022).
- [27] Federico Pigozzi, Yujin Tang, Eric Medvet, and David Ha. 2022. Evolving Modular Soft Robots without Explicit Inter-Module Communication using Local Self-Attention. *arXiv preprint arXiv:2204.06481* (2022).
- [28] Joshua Pinskiar and David Howard. 2022. From bioinspiration to computer generation: Developments in autonomous soft robot design. *Advanced Intelligent Systems* 4, 1 (2022), 2100086.
- [29] Erik Pitzer and Michael Affenzeller. 2012. A comprehensive survey on fitness landscape analysis. *Recent advances in intelligent engineering systems* (2012), 161–191.
- [30] Andre Rosendo, Marco Von Atzigen, and Fumiya Iida. 2017. The trade-off between morphology and control in the co-optimized design of robots. *PloS one* 12, 10 (2017), e0186107.
- [31] Daniela Rus and Michael T Tolley. 2015. Design, fabrication and control of soft robots. *Nature* 521, 7553 (2015), 467–475.
- [32] Kazi Rushdi, Derek Koop, and Christine Q Wu. 2014. Experimental studies on passive dynamic bipedal walking. *Robotics and Autonomous Systems* 62, 4 (2014), 446–455.
- [33] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864* (2017).
- [34] Erica Salvato, Gianfranco Fenu, Eric Medvet, and Felice Andrea Pellegrino. 2021. Crossing the Reality Gap: a Survey on Sim-to-Real Transferability of Robot Controllers in Reinforcement Learning. *IEEE Access* (2021), 1–19. <https://doi.org/10.1109/ACCESS.2021.3126658>
- [35] Stefano Sarti, Jason Adair, and Gabriela Ochoa. 2022. Recombination and Novelty in Neuroevolution: A Visual Analysis. *SN Computer Science* 3, 3 (2022), 1–15.
- [36] Jacopo Talamini, Eric Medvet, Alberto Bartoli, and Andrea De Lorenzo. 2019. Evolutionary synthesis of sensing controllers for voxel-based soft robots. In *Artificial Life Conference Proceedings*. MIT Press, 574–581.
- [37] Jacopo Talamini, Eric Medvet, and Stefano Nichele. 2021. Criticality-Driven Evolution of Adaptable Morphologies of Voxel-Based Soft-Robots. *Frontiers in Robotics and AI* 8 (2021), 172.
- [38] Fuda van Diggelen, Eliseo Ferrante, Nihed Harrak, Jie Luo, Daan Zeeuw, and AE Eiben. 2021. The Influence of Robot Traits and Evolutionary Dynamics on the Reality Gap. *IEEE Transactions on Cognitive and Developmental Systems* (2021).
- [39] Kathryn Walker, Helmut Hauser, and Sebastian Risi. 2021. Growing simulated robots with environmental feedback: an eco-evo-devo approach. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 113–114.