Noname manuscript No.
(will be inserted by the editor)

Optimizing the Sensory Apparatus of Voxel-based Soft Robots through Evolution and Babbling

Andrea Ferigo · Eric Medvet · Giovanni Iacca*

Received: date / Accepted: date

Abstract The behavior of biological and artificial agents strongly depends, in general, on the data acquired through sensors while interacting with the environment. The sensory apparatus, namely the location and kind of sensors, has therefore a great impact on an agent's ability of exhibiting complex behaviors. Considering the case of robots, sensors are usually a design choice that is hard to take, due to the complexity of the robotic structure and a potentially large number of possible combinations. Here, we explore the possibility of using evolutionary algorithms to automatically design (and optimizing their use) the sensors of voxel-based soft robots (VSRs), a kind of robots composed of multiple deformable components. We chose these robots due to their intrinsic modularity, which allows to freely shape the robot body, brain, and sensory apparatus. We consider a set of sensors that allow agents to sense themselves and their environment and we show, experimentally, that the effectiveness of the sensory apparatus depends on the body shape and the actuation capability. Then we show that evolutionary optimization is able to evolve effective sensory apparatuses, even with constraints on the availability of sensors. We also consider how information from sensors can be exploited more efficiently by introducing the concept of "sensor babbling", which aims to enhance the robots' perception and, hence, their performances.

* Corresponding author

A. Ferigo
Department of Information Engineering and Computer Science, University of Trento, Italy
E-mail: andrea.ferigo@studenti.unitn.it

E. Medvet
Department of Engineering and Architecture, University of Trieste, Italy
E-mail: emedvet@units.it

G. Iacca
Department of Information Engineering and Computer Science, University of Trento, Italy
E-mail: giovanni.iacca@unitn.it

Keywords Adaptation · Morphological evolution · Embodied cognition · CMA-ES

Declarations

Funding

Not applicable.

Conflicts of interest

The authors declare that they have no conflict of interest.

Availability of data and material

Not applicable.

Code availability

The framework used in this work is divided into three parts:

- The evolver, available at: `https://github.com/ericmedvet/jgea`
- The simulator, available at: `https://github.com/ericmedvet/2dhmsr`
- A minimal framework used to combine the two previous components, available at: `https://github.com/ndr09/HSMRcoevo`

Author contributions

A. F.: investigation; software; data curation; visualization; writing - original draft. E. M.: conceptualization; methodology; software; visualization; writing - review & editing. G. I.: conceptualization; methodology; writing - review & editing.

Acknowledgements We thank Luca Zanella for the CMA-ES and vision sensor implementation. We gratefully acknowledge HPC-Cineca for making computing resources available.

1 Introduction

Soft robots are deemed to be one of the key technologies for the future of mankind [1]: compared to traditional hard (i.e., rigid) robotics, they allow in fact a better compliance with the environment and humans, which in turn

leads to higher safety in mission-critical applications. Two of the most relevant examples of soft robots are the voxel-based soft robots (VSRs) [2] and the tensegrity soft modular robots (TSMRs) [3], although other paradigms also have been proposed [4–6]. The main features of soft robots paradigms are their intrinsic softness and flexibility, which allows them to perform tasks that are considered incredibly difficult, if not impossible, for hard robots: for example, soft robots are able to perform smooth locomotion on uneven terrain [7] or squeeze through tight spaces [8]. These possibilities make them ideal tools for complex robotic inspection applications normally handled by either rigid robots or passive sensor agents [9]. In the medical domain, soft robots have been proposed as a support tool for gait rehabilitation [10] and colonoscopies [11].

Another important feature of soft robots is that they are, usually, inherently prone to *modular design*. Often biologically inspired, these kind of soft robots behave like organs made of multiple components of the same kind, similar, e.g., to the myons of which muscle tissues are made. Furthermore, modularity facilitates manufacturing, redundancy, and repair [12–14].

Despite these promises, modeling soft robots' *body* (or morphology) and *brain* (or controller) results in a very complex task. This is mainly due to the hard-to-model dynamics of soft materials being used in the body, as well as the non-linearity of the body-brain system. This, and the current lack of analytical design strategies, make Evolutionary Algorithms (EAs) coupled with physics-based simulations currently the main tool for designing soft modular robots. In a recent work, Kriegman et al. [15] underlined the usefulness of EAs for modeling soft modular robots: they not only found that morphologies optimized through EA allow different kinds of interactions with the environment, but they also suggested that morphological development can, in turn, guide evolution to more robust controller designs. Other recent examples of evolutionary synthesis applied to soft robots where the authors used an EA to optimize the body and/or brain of soft robots are reported in [16–18].

Overall, optimizing soft modular robots through EAs is clearly advantageous for a number of reasons. Firstly, EAs are able to discover unconventional designs, which are hard to design for a human expert. Moreover, these designs not only are better than the handmade ones for solving the desired task, but they can also improve non-functional aspects, such as reduced energy consumption (thus extended lifetime), higher robustness, etc. Furthermore, evolution is able to exploit synergistic effects between body and brain that, as discussed earlier, are often too hard to model analytically.

While most of the works applying EAs to soft robots design focused on the evolution of body and/or brain, we find that an important aspect that has been neglected so far in this field is the sensory apparatus, i.e., the type of sensors with which the robot is equipped and their position in the body. In fact, sensors are usually based on expert design choices and/or physical constraints that are considered *before* the evolutionary optimization, and kept fixed during the optimization process. However, *what happens if the sensory apparatus can evolve?* Is evolution able to optimize *which sensors to use, and*

where to position them in the body of a soft modular robot? This is our first research question. Our hypothesis is that evolving the sensory apparatus may lead to a greater effectiveness of the robot: fewer sensors (with respect to a manually designed sensory apparatus) might suffice for the robot for achieving its task, or better robot performance might be obtained.

Secondly, we focus our attention on the robot lifetime. In particular, we ask ourselves if evolutionary algorithms can be used not only to optimally design the sensory apparatus, but also to *optimize the use* of sensors during the lifetime. In this regard, our second research question is: *can robots be evolved to optimize the way they elaborate the information from sensors?* In this case, we hypothesize that increasing the proficiency in the use of sensors can lead to a better adaptability when facing new terrains.

To the best of our knowledge, the only works that addressed similar research questions are [19–21], although those works focused on evolving the sensory apparatus of hard robots. Another related work is [22], where it has been shown that the landscape of the controller loss function depends on the sensors placement, which in turn conducts evolution towards better controllers. However, the context of that work is unicycle non-holonomic mobile robots. On the other hand, no prior research so far studied the evolution of the sensory apparatus in the domain of soft modular robots and, in particular, of VSRs. Yet, VSRs offer a great freedom in the design of the sensory apparatus, due to their intrinsic and fine-grained modularity.

In order to answer our first research question, we use an EA to evolve the sensory apparatus on two different body shapes, a biped and a worm, which are tested on a locomotion task. In doing that, we compare the results of the evolutionary search with three baseline handcrafted sensory configurations, from "Low" to "High" sensor equipment. To answer our second question, we optimize the robot's controller capability to elaborate the sensory information during the robot lifetime, by including in the fitness function a measure of the sensor input dispersion observed during the first phase of the robot lifetime—that we call the "sensor babbling" phase. We compare then the results obtained in this case with two cases where the fitness does not include any measurement of "babbling", to evaluate any advantage possibly provided by this mechanism.

Summarizing our main results, we find that EAs are able to find configurations that reach the same performances of the ones manually designed based on our previous knowledge. This finding is confirmed also when evolution is constrained to use a smaller number of sensors than the ones available to manual designs. This additional constraint plays a key role from an engineering point of view: using fewer sensors decreases the complexity of the robot, thus reducing its possible points-of-failure and reducing energy use. Furthermore, we collect empirical evidence on the fact that the optimal sensory apparatus depends, in general, on the shape of the body and on the actuation capability. These results allow us to reason, qualitatively, on the link between actuation "strength" (i.e., how much the force exerted by actuators changes the body) and sensory apparatus, and the evolutionary importance of discovering the kinds of sensors that are actually beneficial to the task. Finally, concern-

ing the improving of the robot's capability to elaborate sensor readings, we observe that the optimization of the robots' perception can lead to better performances while also allowing an increase in the adaptability of the robots to new terrains.

This paper extends our previous work [23], where we first tackled the optimization of the sensory apparatus in VSRs, on two main aspects: firstly, we conduct a deeper analysis of the robot behavior (Section 5); secondly, and most importantly, here we introduce the "sensor babbling" concept (Section 4) and, in doing so, we extend the optimization scale from the evolutionary one to the robot lifetime (Section 5.2).

The rest of the paper is organized as follows. In the next section, we provide the background concepts on VSRs and the sensor/controller configuration considered in this work. In Section 3 and 4, respectively, we describe how we represent VSRs in a way that fits evolutionary optimization, and introduce the concept of "sensor babbling" as a means to improve the robots' capability to elaborate sensor readings. In Section 5, we present the experiments and discuss the results we obtained evolving the sensory apparatus and the ones achieved considering the concept of sensor babbling. Finally, in Section 6, we draw the conclusions of this work and discuss possible future research directions.

2 Background: voxel-based soft robots and environment

Voxel-based soft robots (VSRs) are robotic agents composed of several deformable blocks (*voxels*) that can actively vary their volume in response to a control signal [24]. In this study, we consider the 2-D version of the VSRs presented in [25] along with a simulation engine tailored to optimization. We remark that, as mentioned in the cited paper, working in 2-D (rather than in 3-D) eases optimization, since simulations are computationally cheaper and search spaces are smaller; on the other hand, while in principle the approach we propose in this study can be ported to the 3-D case, the generality of the experimental results that we find in the 2-D has to be verified.

In this section we briefly describe the main concepts of VSRs that are important to our study: we refer the reader to the aforementioned paper for further details. A VSR is defined by its *morphology* and its *controller*. For our purposes, we consider the morphology as defined by two elements: the *shape* (of the body) and the *sensory apparatus*. The former is the number and placement of the voxels composing the VSR, while the latter is represented by the number, kind, and placement of sensors. The controller, instead, is a component that at each time step takes as input the readings of the sensors available to the body and determines the control signal applied to each voxel of the VSR.

2.1 Shape

The shape of a VSRs is a 2-D grid of voxels where adjacent voxels are rigidly connected at their vertices (see Figure 1 for an example of the two shapes considered in this study). A voxel is a deformable square that is modeled, in the simulation, as a compound of masses, spring-damper systems, and distance constraints, see [25] for details.

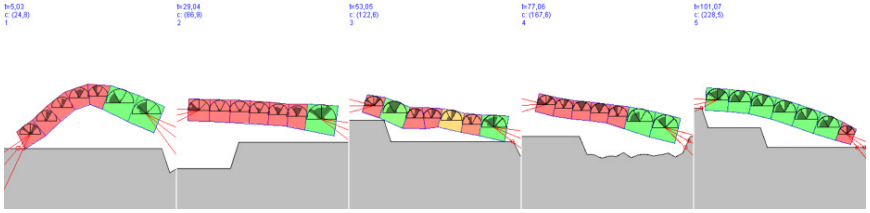

(a) Worm on the training terrain.

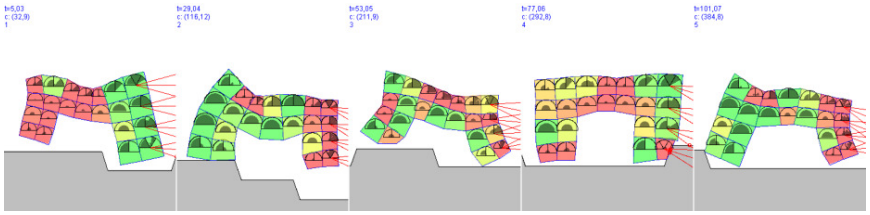

(b) Biped on the validation terrain.

Fig. 1 Frames of two VSRs (a worm and a biped, both with vision sensors) captured during two simulations. The color of each voxel encodes the ratio between its current area and its resting area: red indicates contraction, yellow no variation, and green expansion; the circular sector drawn at the center of each voxel indicates the current sensed values $\boldsymbol{s}(t)$ (see [25]). The rays of the vision sensors are shown in red.

A voxel changes its area depending on: (a) the control signal, and (b) the external forces applied by other voxels connected to it. The control signal is a value in $[-1, 1]$ representing the request of the controller to contract or expand the voxel: -1 corresponds to maximum expansion, 1 corresponds to maximum contraction. The actuation of the control signal is modeled in the simulation as an instantaneous change of the resting length of the spring-damper systems.

The exact amount of area change depends on the parameters of the voxel model, i.e., the properties of the voxel. In this work, we assume that all the voxels have the same properties. A parameter that is particularly relevant to this work is the *maximum area change* ρ_A, defined as follows. Let A be the area of a voxel not subjected to external forces and with a control signal $f = 0$, then the area of the voxel not subjected to external forces and a control signal f is $A(1 - \rho_A f)$. Intuitively, ρ_A represents the "strength" of a voxel: the larger

its value, the larger the area change when the voxel is controlled with the same control signal.

2.2 Sensory apparatus

The sensory apparatus of the VSR is a central point of this study and allows the VSR to perceive itself and the environment. The perception, represented by the vector of sensor readings, is processed by the controller in order to calculate the control signal, affecting in this way the behavior of the VSR.

The sensory apparatus of a VSR consists of the number, position and kind of available sensors; each voxel of the VSR shape can be equipped with zero or more sensors. Each sensor has a type, and for each voxel there can be at most one sensor of a given type. Each sensor produces, at each time step, a *sensor reading* $\boldsymbol{s} \in D \subseteq \mathbb{R}^p$, D being the domain of the sensor and p being the dimensionality of the sensor type.

In this work, we consider four types of sensors: two of them, *area* and *velocity*, perceive the internal state of voxels, while the other two, *touch* and *vision*, perceive the surrounding environment. More specifically, *area* sensors perceive the ratio between the current area of the voxel and its resting area: the domain is hence $D =]0, +\infty[$. *Velocity* sensors perceive the velocity of the center of mass of the voxel along the x- and y- axes integral with the voxel itself (i.e., the axes rotate with the voxel): the domain is $D = \mathbb{R}^2$. *Touch* sensors perceive whether the voxel is touching the ground ($\boldsymbol{s} = 1$) or not ($\boldsymbol{s} = 0$): the domain is $D = \{0, 1\}$. *Vision* sensors perceive the distance towards close objects, i.e., the terrain and the obstacles, within a given field of view. More specifically, vision sensors are modeled as lidars, i.e., p straight rays cast from the voxel center with angles $\alpha_1 \ \ldots \ \alpha_p$ with respect to the positive x-axis integral with the voxel. For each i-th ray, the corresponding value s_i of the sensor reading is $\min\left(\frac{d}{d_{\max}}, 1\right)$, where d is the distance between the voxel center and the point where the ray hits the closest object, and $d_{\max}$ is a parameter representing the maximum distance of sight. The domain of the vision sensors is hence $D = [0, 1]^p$. For illustration purposes, Figure 1 shows two examples of VSRs equipped with vision sensors and highlights the rays cast by those sensors.

Finally, we further wrap each sensor in two artifacts: the first one normalizes the values perceived in the range $[0, 1]$, in order to have the same domain for all sensors. The second one adds a random noise to the sensor, simulating the noise present in real robots. The wrapper modifies the sensor readings as follows:

$$\boldsymbol{s} := \boldsymbol{s} + \boldsymbol{\nu}, \text{ with } \boldsymbol{\nu} = [\nu_1 \ \ldots \ \nu_p] \text{ and } \nu_{i=1 \ldots p} \sim \gamma N(0, 1) \tag{1}$$

where γ is set to 0.01.

2.3 Controller

The role of the controller is to process the information perceived by sensors, determining the value of the control signals of the VSR voxels over time. Here we follow the control approach proposed in [26], i.e., the controller is a feed-forward neural network (FFNN), as depicted in Figure 2.

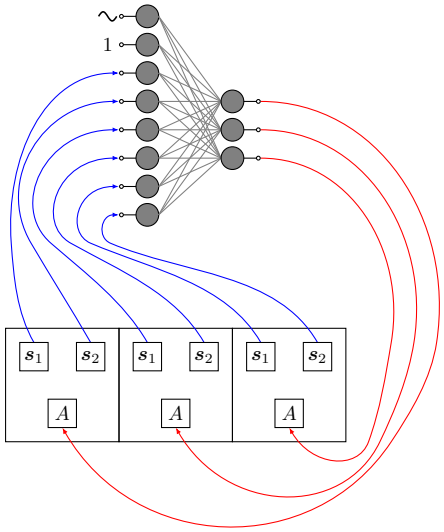

Fig. 2 Graphical representation of the controller for a robot composed of 3 voxels each one with 2 sensors. The 6 sensor readings provide 6 inputs to the FFNN (blue lines), the remaining two inputs (in the upper part) are the driving function and the bias respectively. The output layer returns the actuation value of each voxel (red lines).

The FFNN has one input for each of the values sensed by the sensors (i.e., one input for each touch sensor, two inputs for each velocity sensor, and so on), one input whose value varies over time according to an input driving function, and the bias. The role of the driving function is to facilitate the emergence of dynamics useful for the task to be accomplished by the VSR [26]: in particular, we use a sinusoidal function with frequency of 1 Hz. The FFNN has one output for each voxel. The value determines, at each time step k, the control signal $f_i^{(k)}$ of the corresponding i-th voxel:

$$\boldsymbol{f}^{(k)} = \text{FFNN}_{\boldsymbol{w}} \left(\left[\boldsymbol{s}_1^{(k)} \ \ldots \ \boldsymbol{s}_n^{(k)} \ \sin(2\pi k \Delta t) \right] \right) \tag{2}$$

where $\boldsymbol{s}_j^{(k)}$ is the j-th sensor reading at time step k, $\sin(2\pi k \Delta t)$ is the value of the driving function at time k, $\Delta t = 1/60\,\text{s}$ is the duration of a time step in the simulation, and $\boldsymbol{w}$ is the vector of the weights of the FFNN.

In our experiments, the input and output layers are connected without inner layers and each node in the FFNN is controlled by a tanh activation

function. We opt for this simple topology of the FFNN because it reduces the number of parameters, and hence the size of the search space when optimizing the controller. Indeed, previous works showed that the performance gap between VSRs with zero or more than zero hidden layers is not significant [16, 27].

In order to optimize the controller, we use a direct encoding of the weights $\boldsymbol{w}$ of the FFNN (more on this below). The total number $|\boldsymbol{w}|$ of weights depends, for a given shape, on the sensory apparatus of the VSR, as a change in the kind or number of available sensors corresponds to a change in the number of inputs.

2.4 Environment

We simulate robots on two different kinds of terrain: the first one, used for the evolution of the sensory apparatus, is an uneven terrain obtained by the alternation of three elements: a gap, a hill, and a rough section; the second one, used for the analysis of sensor babbling, is characterized by a sequence of hills with different slopes. It should be noted that both kinds of terrain are difficult enough to highlight the importance of the sensory apparatus. However, for the sensor babbling experiments we use the second terrain as the first one is computational heavier due to the numerous contacts between voxels and terrain, in particular in the rough sections.

3 Evolution of the sensory apparatus

As discussed earlier, previous works have applied EAs to evolve the morphology and controller of VSRs, while here we are interested in investigating if also the sensory apparatus can be explicitly evolved. However, the sensory apparatus is intrinsically connected to the controller, and vice versa: in particular, as we have seen the size of the FFNN is directly determined by the number and type of sensors available in the sensory apparatus. For this reason, we evolve the controller and the sensory apparatus together and propose, for this purpose, two representations (also called mappers)—that we call respectively "Unlimiting" and "Limiting"—which differ in the possibility of limiting or not the maximum number of sensors. According to these mappers, a numerical vector (the *genotype*) is mapped, given a certain shape, to a pair ⟨sensory apparatus, controller⟩ (the *phenotype*) suitable for that shape. As introduced before, we use the weights vector $\boldsymbol{w}$ as genome, since this vector completely describes the controller for a given shape and its numerical nature allows us to use well-known general-purpose optimizers to conduct the search. For our experiments, we rely on the Covariance Matrix Adaptation Evolution Strategies (CMA-ES) algorithm [28], which has been shown already to be effective for optimizing VSRs [18, 29].

In both representations, each element of the numerical vector (that constitutes the genotype) determines directly the realization of a component of the

sensory apparatus and its relative controller. Furthermore, the two representations exhibit some degree of redundancy: in some portion of the genotype space, several different genotypes correspond in fact to the same phenotype. This can result in degeneracy of the representation [30].

Finally, both representations are specific to a given body, i.e., to a pair consisting of a shape and a *maximal sensory apparatus*. The maximal sensory apparatus determines the most complex sensory apparatus (i.e., the one with the largest number of sensors) that can be represented for that body.

As for the mechanism used to limit (or not) the number of sensors in the two mappers, the implementation is the following. Let n_S be the overall dimensionality of the sensors of the maximal sensory apparatus and n_V the number of voxels of the shape. The genotype is defined in $\mathbb{R}^{(n_S+2)n_V}$ and directly encodes the weights $\boldsymbol{w}$ of a *maximal controller*, i.e., the FFNN that takes the inputs from the maximal sensory apparatus (plus the bias and the driving function) and applies its outputs to the n_V voxels. We denote with $\boldsymbol{w}_{(s)}$ the vector of weights corresponding to the inputs of the FFNN connected to the sensor s.

In the "Unlimiting" representation, given a threshold $\tau_w \in \mathbb{R}^+$, the sensory apparatus corresponding to a genotype $\boldsymbol{w}$ is composed of each sensor s of the maximal sensory apparatus for which at least one weight in $\boldsymbol{w}_{(s)}$ is greater, in absolute value, than τ_w, i.e., for which $w_{s,\max} \coloneqq \max_i |\boldsymbol{w}_{(s),i}| > \tau_w$. The sensors that satisfy this condition will be selected to compose the inputs of the FFNN while the weights will be directly obtained from the genotype, see Figure 3. It can be seen that there is no hard limit to the complexity of the sensory apparatus that can be represented: if enough weights are large enough, the sensory apparatus is the maximal one. On the other hand, this representation allows for a sensory apparatus consisting of no sensors at all.

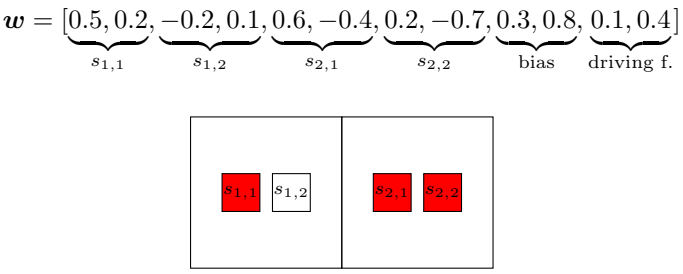

Fig. 3 Representation of the Unlimiting mapper: in the upper part the genotype is represented, while in the bottom the corresponding FFNN (the phenotype) is shown. Here the VSR is made of two voxels and its maximal sensory configuration is composed of 4 sensors $(s_{1,1}, s_{1,2}, s_{2,1}, s_{2,2})$ represented by the inner squares. The color of the square indicates if the sensor is enabled (red) or not (white). The genotype length, $|\boldsymbol{w}|$, is 12 as there are two weights for each sensor, plus the ones for the bias and the driving function. In this case we set the threshold τ_w to 0.3, hence the unlimiting mapper enables all the sensors except for $s_{1,2}$, whose weights are in absolute value smaller than τ.

In the "Limiting" representation, instead, given a maximum number of sensors n_{sensors}, the sensory apparatus corresponding to a genotype $\boldsymbol{w}$ is composed of the n_{sensors} sensors with the largest $w_{s,\max}$ (in absolute value), see Figure 4 for an example. It should be noted that in this case, regardless of the values in $\boldsymbol{w}$, the sensory apparatus has always the same complexity (i.e., it always consists of exactly n_{sensors} sensors). We remark, however, that weights that are very small correspond, in practice, to sensors that do not significantly impact the computation of the control signals, i.e., that are not actually used by the controller.

$$\boldsymbol{w} = [\underbrace{0.5, 0.2}_{s_{1,1}}, \underbrace{-0.2, 0.1}_{s_{1,2}}, \underbrace{0.6, -0.4}_{s_{2,1}}, \underbrace{0.2, -0.7}_{s_{2,2}}, \underbrace{0.3, 0.8}_{\text{bias}}, \underbrace{0.1, 0.4}_{\text{driving f.}}]$$

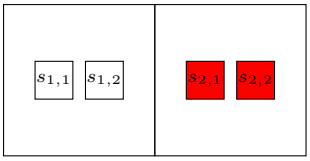

Fig. 4 Representation of the Limiting mapper: in the upper part the genotype is represented, while in the bottom the corresponding FFNN (the phenotype) is shown. The VSR is made of two voxels and its maximal sensory configuration is composed of 4 sensors $(s_{1,1}, s_{1,2}, s_{2,1}, s_{2,2})$ represented by the inner squares. The color of the square indicates if the sensor is enabled (red) or not (white). The genotype length, $|\boldsymbol{w}|$, is 12 as there are two weights for each sensor plus the ones for the bias and the driving function. We set the Limiting mapper to enable two sensors. The mapper ranks the sensors based on their greatest weight (in absolute value) and it enables the first two, which, in this case, are $s_{2,1}$ and $s_{2,2}$.

4 Sensor babbling

As discussed in Section 1, sensors allow the robot to perceive the environment and hence determine the robot's behavior. However, sensors are actually useful only if they produce readings that vary in a sensible way when the robot interacts with the environment; on the other hand, sensors that are never stimulated, regardless of the current situation the robot is in, are likely to be of poor utility. In order to capture, and possibly exploit for the optimization of the sensory apparatus, the usefulness of sensors, we get inspiration from the concept of *motor babbling* [31, 32], that has been introduced in robotics to imitate the process through which infants perform random movements in order to understand (i.e., build an internal model of) the relationship between their body and the environment. Since we deal with sensors, we call *sensor babbling* the phase of the robot's life when it explores its sensing ability. Our hypothesis is that, through this sensor babbling process, robots are exposed to a broad range of sensor inputs, which may lead to better performances on the task.

In order to be able to use the sensor babbling phase for the purpose of optimization (see Section 5.2), we need to measure how "intensively" the robot is doing sensor babbling. Once we can measure sensor babbling, we can promote robots that not only achieve a better performance on the task, by observing them in an "adult" phase of their life, but that also do an intense sensor babbling in an earlier "infant" phase of their life.

We measure quantitatively the "intensity" of the sensor babbling (to be maximized) by using, for each sensor kind, the average distance of the sensor readings from their centroid. In particular, first we collect the sensor readings for each sensor kind at each simulation tick; then, the vector of sensor inputs is filtered, keeping only the values read from the selected sensor kind; finally, we calculate the distance of each filtered vector from the corresponding centroid. In Algorithm 1 a the pseudo-code of the measurement algorithm used for sensor babbling is given.

Algorithm 1 Sensor babbling measurement.

```
1: procedure measureSensorBabbling(sensorKind)
2:     filteredInputs ← inputs.filtersBy(sensorKind)
3:     centroid ← calculateCentroid(filteredInputs)
4:     distance ← 0
5:     for input in filteredInputs do
6:         distance ← distance + euclideanDistance(input, centroid)
7:     end for
8:     return distance / filteredInputs.size()
```

5 Experimental analysis

5.1 Evolution of the sensory apparatus

We aim at investigating the possibility of evolving the sensory apparatus of VSRs from different points of view. More precisely, we address three facets of the problem, here represented in terms of research questions (RQ):

- RQ1 Are sensors beneficial to robot effectiveness? Is the benefit of sensing ability somehow diminished by greater strength? Is it dependent on the robot shape?
- RQ2 Can evolution discover an effective sensory apparatus?
- RQ3 When the complexity of the evolvable apparatus is limited, what sensors are "preferred" by evolution?

To answer these questions, we perform a number of experiments on the uneven terrain described in Section 2.4. We evolve two shapes (shown previously in Figure 1): the first, that we call worm, is a 7×1 (x and y size of the voxel

grid constituting the shape) rectangle; while the second one resembles a biped structured as a 7×4 rectangle with 3×2 missing voxels at the bottom-center.

Concerning the task, we consider *locomotion*, i.e., a limited time span, the *episode*, in which the robot has to travel as far as possible along the positive x direction. We measure the effectiveness of a VSR in performing locomotion as its velocity with respect to the x-axis during the episode: $\overline{v}_x = \frac{x(t_f)-x(0)}{t_f}$ (in m/s), where t_f is the episode duration and $x(t)$ is the position of the VSR center of mass at time t. Of note, locomotion is a classic task in evolutionary robotics and usually consists in making the robot run along a flat terrain. However, robots that perceive the environment, in particular the ones that perceive obstacles, should be favored with respect to robots that base their locomotion capability on a regular gait that does not depend on sensors [26]. Hence, in this work, we use an uneven terrain which we consider better suited for investigating the robots' perception capabilities. Moreover, for a subset of the experiments, we measure the locomotion velocity $\overline{v}_x$ over two different terrains, one used for evolution, one after: the motivation is to verify if the robot actually exploits the sensory apparatus for running faster on an "unknown" terrain.

For all the experiments, we use CMA-ES for the optimization process, with default parameter settings (as indicated in [33]—the main ones being the initial step size $\sigma = 0.5$ and the population size $\lambda = 4 + \lfloor 3 \log |\boldsymbol{w}| \rfloor$)—and with the initial vector of means set by sampling uniformly the interval $[-1, 1]$ for each vector element. We simulate episodes lasting $t_f = 180\,\mathrm{s}$ (simulated time). For the simulation, we use the default parameters of the 2D-VSR-SIM software [25], unless otherwise specified. The code used for the experiments, based on JGEA (`https://github.com/ericmedvet/jgea`) for the evolutionary optimization part, is publicly available at: `https://github.com/ndr09/HSMRcoevo`.

5.1.1 RQ1: Sensor potential benefit

In order to answer this question, first we manually design, based on domain knowledge, three sensory apparatuses for each of the two shapes. For each resulting body, we then evolve the controller with different values of ρ_A, i.e., different values of strength (see Section 2.1).

Figure 5 shows the six bodies, i.e., shapes and sensory apparatuses. For each shape, the three apparatuses are derived from a single maximal sensory apparatus that consists of a number of sensors of different types placed in the shape in order to favor the perception of a robot that moves towards the right, i.e., in the positive x direction. In particular, in the maximal sensory apparatus we place:

- a vision sensor with three rays (with $\alpha_1 = 0°$, $\alpha_2 = -15°$, and $\alpha_3 = -30°$) on each rightmost voxel of the shape; for the worm, we also place one vision sensor on the leftmost voxel (with $\alpha_1 = 180°$, $\alpha_2 = 195°$, and $\alpha_3 = 210°$) to compensate the limited extension of the forward front of the shape;
- a touch sensor on each voxel of the bottom row of the grid (i.e., all the voxels in the worm and the voxels corresponding to the "feet" of the biped);

– area and velocity sensors spread over the body, as shown in Figure 5.

We then differentiate the three apparatuses based on the total number of available sensors with respect to the corresponding maximal sensory apparatus. In the "Low" perception apparatus, robots are equipped with few area sensors. In the "Medium" perception apparatus, we remove all the vision sensors. In the "High" perception apparatus, we do not remove any sensor, i.e., this apparatus is the maximal sensory apparatus and is the only one that allows the robot to "see" its surroundings. As a consequence, the genome size, i.e., the number of weights $|\boldsymbol{w}|$, greatly varies for the three controllers and the two shapes: from 35 variables for the worm with the Low apparatus to 1188 variables for the biped with maximal sensory configuration. The genome size of all the combinations of shape and sensory apparatus considered in the experimentation is shown in Table 1 (where, for completeness, we also show the genome size for the Limiting and Unlimiting representations, which for both mappers coincides with that of the maximal sensory apparatus, i.e., the High one).

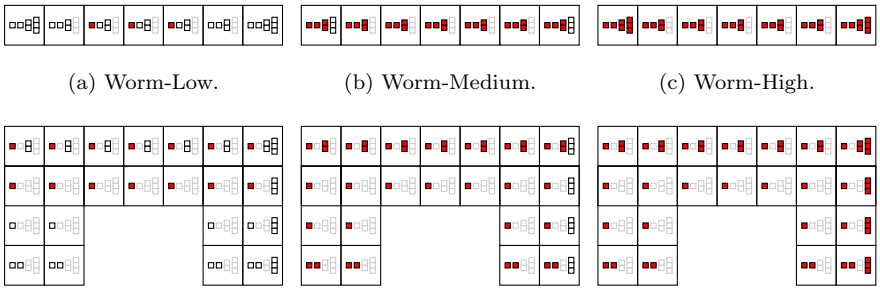

(a) Worm-Low. (b) Worm-Medium. (c) Worm-High.

(d) Biped-Low. (e) Biped-Medium. (f) Biped-High.

Fig. 5 The six bodies (shape and sensory apparatus) used in the experiments. In each voxel, sensors are represented as stacks of p squares, p being the dimensionality of the sensor (see Section 2.2). Area sensor ($p = 1$) is the first (leftmost) column; touch sensor ($p = 1$) is the second column; velocity sensor is the third one ($p = 2$); vision sensor is the last one ($p = 3$). The color of the sensor represents its presence in the apparatus: gray border means not present; black border means not present, but present in the maximal sensory apparatus; red fill means present.

Table 1 Genome size for all the combinations of shape and sensory apparatus considered in the experimentation.

Shape	Low	Medium	High	Limiting	Unlimiting
Worm	35	210	252	252	252
Biped	352	924	1188	1188	1188

We consider 8 values for ρ_A, evenly distributed in $[0.1, 0.8]$. For each one of the $2 \times 3 \times 8$ combinations of shape (worm, biped), sensory apparatus (Low, Medium, High), and ρ_A, we perform 10 independent evolutionary runs of the controller with CMA-ES. We use the velocity $\overline{v}_x$ as fitness and we stop the evolutionary process after 2000 fitness evaluations.

Figure 6 summarizes the results of this experiment with the two shapes, the three fixed sensory apparatuses, and the eight values of ρ_A. The solid lines indicate the average performances, across the 10 runs, in the locomotion task, namely the average velocity $\overline{v}_x$, of the best VSR found by CMA-ES on the *training* terrain. The figure also shows, with dashed lines, the average performances that the best VSRs score on the *validation* terrain, i.e., an uneven terrain different from the one using during evolution, which is created with a different seed. The shaded areas indicate the std. dev. across the 10 runs.

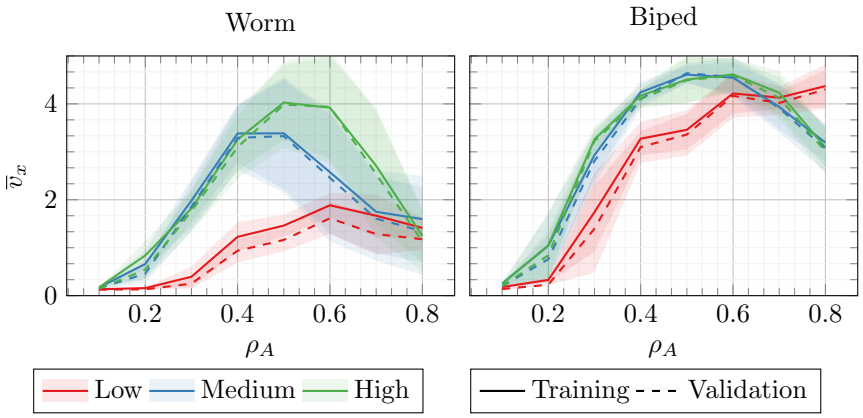

Fig. 6 Velocity $\overline{v}_x$ of the best evolved VSRs on different terrains (line style) and sensory apparatuses (line color) vs. the strength ρ_A.

We can highlight three points observing Figure 6. First of all, as expected, the locomotion capability of the robots depends on their strength: the stronger the robot, the faster it will be in locomotion; however, this holds only for values of ρ_A lower than 0.5. For $\rho_A > 0.5$, the results show a lowering in the performance of the VSRs: we further analyze the behavior of the evolved robots in this case and we discover that increasing the value of ρ_A in practice leads to the limit of the physics model of the simulator; this results in robots that are difficult to control and hence ineffective, due to stronger contractions and expansions of their voxels.

Secondly, the differences in velocity among sensory apparatuses depend on the shape. For the worm, the more complex the sensory apparatus, the faster is the robot: moreover, the three apparatuses peak at different values of strength and the Low apparatus is greatly outperformed by the Medium and High configurations. For the biped, the differences are in general smaller: the Medium and the High apparatuses result in roughly the same peak $\overline{v}_x \approx 4.3$

for the same value of ρ_A. The gap of the Low apparatus with respect to the other two is not negligible only for values of ρ_A in the range $[0.2, 0.5]$.

Finally, when the VSR bases its locomotion also on the perception of environment, it is more capable of coping with unseen environmental conditions. We can draw this conclusion observing the difference in fitness on the training and validation terrains: it appears to be larger for the Low sensory apparatus—this is more apparent for the worm shape.

We believe that these findings are important, in particular the second one, since they suggest that there is not a one-fits-all solution for the sensory apparatus. On the contrary, the best apparatus depends (at least) on the shape and strength of the VSR. Thus, the idea of optimizing a sensory apparatus for a specific VSR shape appears sound.

5.1.2 RQ2: Effectiveness of the evolved sensory apparatus

In order to answer this second question, we perform a longer optimization process evolving the sensory apparatus of the two shapes, the worm and the biped, using the two mappers described in Section 3. We then compare the results of this evolutionary optimization against the three fixed apparatuses previously described (Low, Medium, High). In this case, when we evolve the apparatus we use the High configuration as the maximal sensory apparatus: the number of weights $|\boldsymbol{w}|$ for the evolvable apparatuses is hence 252 and 1188, respectively for the worm and the biped.

We set the parameters of the two representations as follows. For the Limiting representation, we set $n_{\text{sensors}} = 20$ for the worm and $n_{\text{sensors}} = 30$ for the biped, i.e., approximately the number of sensors in the Medium apparatus. For the Unlimiting representation, we determine the value of the weight threshold τ_w by examining the values of the weights of the $10 + 10$ (worm and biped) best evolved VSRs with the High apparatus found in the previous experiment (see Figure 7). Then we choose, for each shape, the value corresponding to the tenth percentile of $w_{s,\max}$ (see Section 3), that is $\tau_w = 1.96$ for the worm and $\tau_w = 2.87$ for the biped. We recall that the Unlimiting representation, differently from the Limiting one, can potentially allow all the sensors of the High configuration, which are 23 and 37 respectively for the worm and the biped.

In this case, we fix the value of $\rho_A = 0.3$. We choose this value (that is also the default value of the simulator) because, based on the findings of the previous experiment (see Figure 6), with $\rho_A = 0.3$ the differences among sensory apparatuses are large enough to allow investigation. As in the previous experiment, we perform 10 independent evolutionary runs, in this case ended after 25 000 fitness evaluations. We consider a total of 10 different combinations of shapes and sensory apparatuses, namely 2 shapes (worm, biped), each one tested with the three statically defined sensor configurations (Low, Medium High), and two evolvable ones (Limiting, Unlimiting). Finally, we modify CMA-ES in order to have an initial population starting with all the sensors disabled.

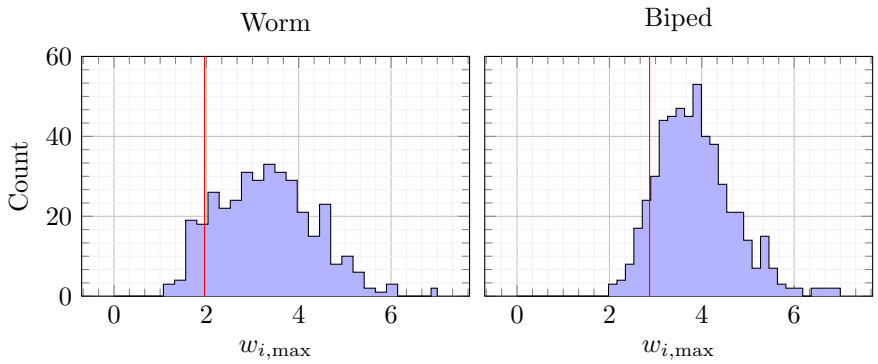

Fig. 7 Distribution of the $w_{s,\max}$ values for the sensors evolved with the High sensory apparatus in the experiment of Section 5.1.1. The red line corresponds to the 10-th percentile.

Figure 8 summarizes the results of this experiment: it shows the trend of the velocity $\overline{v}_x$ (mean $\pm$ std. dev. across the 10 runs) of the best VSR at different stages of the evolutionary process (i.e., number of births). The foremost finding that can be noted from is that the evolved sensory apparatuses are not, in general, worse than the hand-designed ones. By looking in detail at the values of $\overline{v}_x$ at the end of the runs, which are shown in Figure 9 in the form of boxplots (where the whiskers represent the 1.5 IQR), it can be seen that for the worm there are not significant differences between the Medium, the High, and the two evolved apparatuses, while for the biped both the evolved apparatuses and the High one outperform the others.

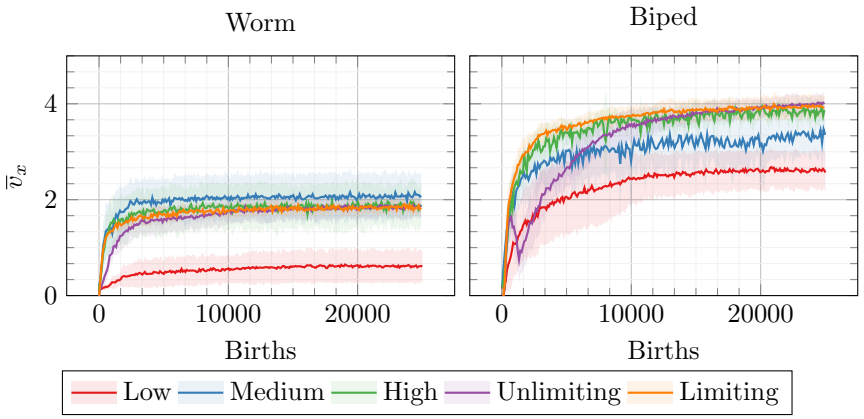

Fig. 8 Velocity $\overline{v}_x$ of the best evolved VSRs found during the evolutionary process in the experiment described in Section 5.1.2.

We further analyze the results for the worm with the Medium and High configurations to understand why they reach similar performances. Inspecting

the behavior of the robots in the Medium configuration, we observe that the majority of them vibrate at a high frequency, causing a behavior that produces a faster, but less realistic locomotion with respect to the one shown by the other robots. In Figure 10, we show the velocity (mean ± std. dev. across the 10 runs) of the robots classified by visual inspection based on their behavior (i.e., vibrating or not vibrating). As it can be seen, the performance of the robots that do not vibrate is similar to the one of the High configuration. Moreover, we observe that the vibration behavior is not favorable to avoid difficult obstacles, such as deeper gaps.

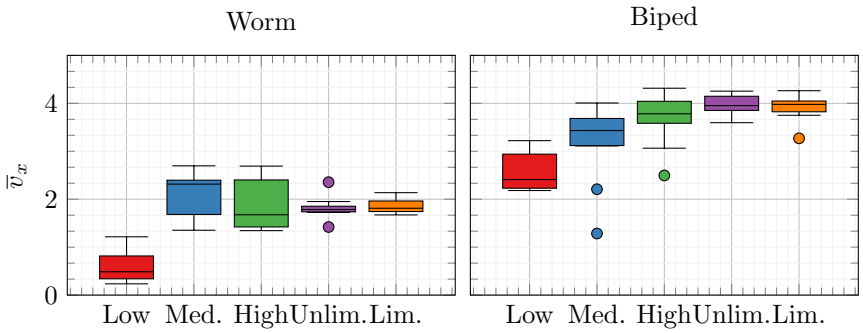

Fig. 9 Distribution of the velocity $\overline{v}_x$ of the best evolved VSRs found at the end of the evolutionary process in the experiment described in Section 5.1.2.

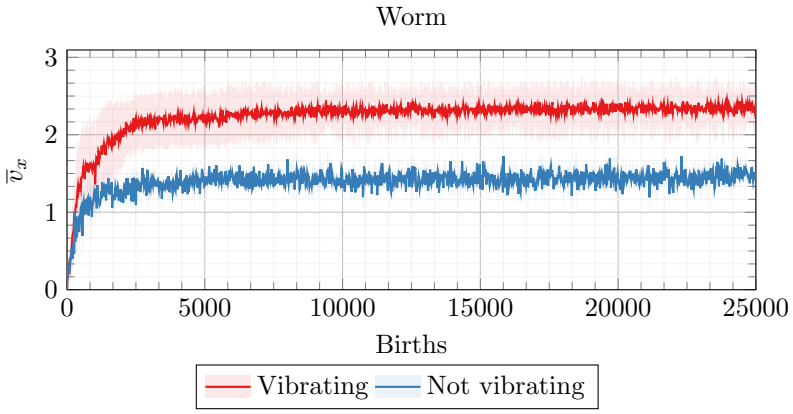

Fig. 10 Velocity $\overline{v}_x$ of the best evolved VSRs found during the evolutionary process with the worm and the Medium configuration, classified by behavior. The velocity of the VSRs that vibrate is much bigger than that of the ones that do not vibrate.

As a further confirmation of the results shown in Figure 8, we perform the Mann-Whitney U test (after having verified the relevant hypotheses) between

pairs of samples from different apparatuses for each shape. Table 2 shows the resulting p-values. By comparing the two representations for the evolvable apparatuses it can be seen that there are no significant differences. In particular, this means that when evolution uses fewer sensors it is still capable to find a pair of sensors apparatus and controller that reaches the same performances of the controller evolved with a fixed configuration.

Table 2 p-values of the pairwise Mann-Whitney U test performed on the $\overline{v}_x$ resulting from the experiment of Section 5.1.2 (values shown in Figure 9). Values corresponding to a significant difference (Bonferroni correction, $\alpha = 0.05$) are highlighted in boldface.

		Low	Med.	High	Unlim.	Lim.
Worm	Low		**0.00078**	**0.00078**	**0.00078**	**0.00078**
	Med.	**0.00078**		1.0	0.2062	0.0778
	High	**0.00078**	1.0		1.0	1.0
	Unlim.	**0.00078**	0.2062	1.0		1.0
	Lim.	**0.00078**	0.0778	1.0	1.0	
Biped	Low		0.1418	**0.0057**	**0.0007**	**0.0007**
	Med.	0.1418		0.4105	0.2939	0.9938
	High	**0.0057**	0.4105		1.0	1.0
	Unlim.	**0.0007**	0.2939	1.0		1.0
	Lim.	**0.0007**	0.9938	1.0	1.0	

Finally, by looking at the values of $\overline{v}_x$ during the evolutionary process with the Unlimiting representation, it can be seen that they stay below the values obtained with the Limiting representation in the early stage of evolution. For the biped in particular, the peculiar behavior of the fitness shows that after an initial, rapid improvement the performance decreases up to ≈ 150 births and then starts to improve again—recall that CMA-ES does not guarantee the monotonicity of the fitness. Analyzing the raw results, we find that this behavior is due to the progressive discovery of the sensors: at the beginning, evolution pushes towards a controller that is able to run without perceiving the environment (thanks to the dynamics generated by the driving function). As sensors are discovered, they initially require evolution to "adjust" the controller, and in this phase the fitness decreases. Later, sensors become beneficial, thus allowing the fitness to increase again. This interpretation is supported by Figure 11, that shows the distribution of VSRs that have at least one sensor enabled and the ones that do not have any sensor enabled. Initially all VSRs belong to the second population, because of the initialization of CMA-ES means to $\mathbf{0}$. Then, the number of VSRs that do not use sensors quickly goes to 0 % in the early stage of the evolutionary process. This phenomenon is less visible in the fitness trend of the Unlimiting representation for the worm, since this configuration has much fewer sensors to be activated and, hence, the changing phase happens faster. Moreover, the same phenomenon can not

be present using the Limiting representation, because this configuration maintains a controller with n_{sensors} sensors enabled by design.

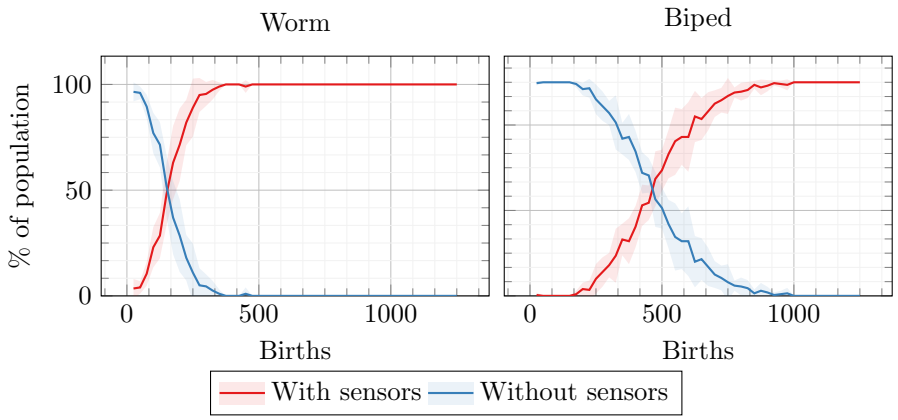

Fig. 11 Percentage (mean ± std. dev. across 10 runs) of the individuals in the population not using (blue) or using (red) sensors in the early stage of the evolutionary process (up to 1250 births) with the biped and the Unlimiting representation, see Figure 8 for the full evolutionary trend.

5.1.3 RQ3: Sensors preferred by evolution

We attempt to answer this question in two ways. First, quantitatively, looking at the robots evolved with the Limiting representation and analyzing which sensors of the maximal sensory configuration remain deactivated—we recall that, by design, the Limiting representation does not enable all the sensors. Second, qualitatively, through the examination of the activation order, during the evolutionary time, of the sensors in the evolutionary runs performed with the Unlimiting representation—we recall that, by design, the Unlimiting representation starts with all the sensors disabled and can end with all the sensors enabled. The rationale is that the sensors that appear to be more beneficial for locomotion should be enabled more often, in the first case, and earlier, in the second case.

In order to better investigate sensor importance, we repeat the experiment of the previous section after having slightly modified the two representations. In this experiment, we allow the possibility to enable only specific domains of a sensor, e.g., a biped can have just one front vision sensor with $\alpha_3 = -30°$ enabled, and the other dimensions ($\alpha_1 = 0°$ and $\alpha_2 = -15°$) disabled.

Table 3 shows the results corresponding to the first, quantitative analysis. It shows the number of sensors (mean and std. dev. across 10 runs), for each type and shape, that are enabled at the end of the evolutionary process with the Limiting representation. We consider a sensor as enabled if *at least one dimension* of the sensor is enabled.

Table 3 Mean and std. dev. (across the 10 runs) for the number of sensors, for each type and shape, that are enabled at the end of the evolutionary process with the Limiting representation.

Shape	Sensory app.	Area	Touch	Velocity	Vision
Worm	Maximal	7.0	7.0	7.0	2.0
	Evolved	5.7±1.0	5.8±0.6	6.7±0.6	2.0±0.0
Biped	Maximal	22.0	4.0	7.0	4.0
	Evolved	16.4±0.9	3.4±0.7	6.8±0.3	4.0±0.0

From the table, it can be seen that for both shapes the vision sensors are the only ones that result all enabled (in all runs). While for the area sensors, that are the least used ones with both shapes, there is a great difference between the worm and biped case, velocity and touch sensors share, more or less, the same activation percentages for both the shapes. With the biped, ≈25 % of the area sensors remain not used on average: interestingly, this is also the most numerous sensor in the maximal sensory apparatus of this shape. A possible interpretation of the latter finding is that in the manually designed maximal sensory apparatus of the biped the area sensors are redundant: under the constraint posed by the Limiting representation to use fewer sensors, evolution still finds an effective solution for locomotion (see Section 5.1.2) and hence it "is aware" of this redundancy.

Concerning the qualitative analysis, we report here the detailed outcome for one of the runs with the Unlimiting representation–several other runs exhibit a similar behavior. Figure 12 shows the sensory apparatus of the best biped (represented with the graphical notation of Figure 5) in four salient points of the evolutionary process, that are referred to the corresponding values of the fitness $\overline{v}_x$ and the percentage of sensor dimensions enabled with respect to the maximal sensory apparatus. Besides the beginning of the evolutionary process, we select three other salient moments: (i) the discovery of the first sensor, (ii) the phase of maximum improvement of the fitness after the initial drop, and (iii) the first time all sensors are enabled.

Two qualitative observations can be made by looking at Figure 12. Firstly, in the moment of maximum improvement after the initial drop in the fitness, the number of area, velocity and touch sensors enabled is much lower with respect to the vision ones, which are almost all enabled. This clue seems to confirm the findings discussed above (based on Table 3) about the relevance of the vision sensors for locomotion. Secondly, the first peak value of the fitness and the discovery of the sensors occurs at the same evolutionary time, namely just after the beginning of the evolutionary process. After that moment, as previously observed in Section 5.1.2 (in particular in Figure 8 and 11), evolution needs to "learn" how to effectively use the discovered sensors for a faster locomotion. For this reason, initially the fitness rapidly decreases for some hundreds of births and then it starts to increase again.

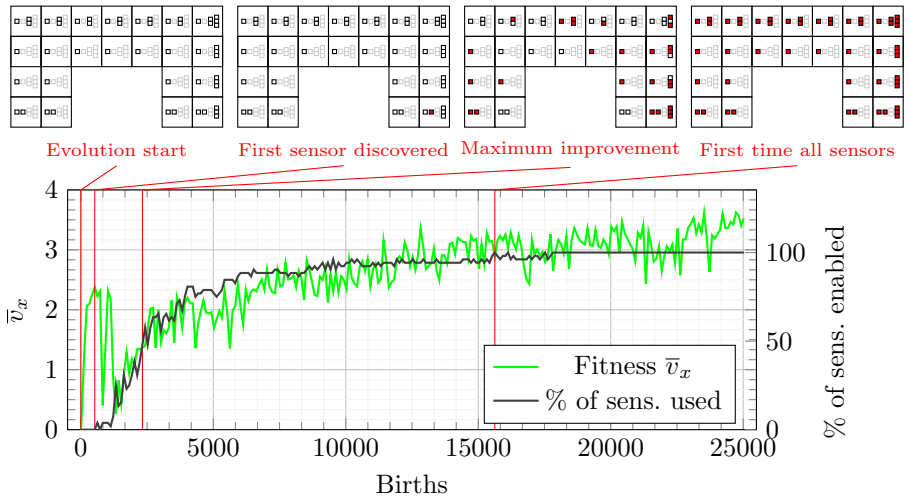

Fig. 12 Sensory apparatus of the biped with the Unlimiting representation in four salient moments of the evolutionary process for one of the evolutionary runs. The plots of the VSRs are referred to the corresponding values of fitness $\overline{v}_x$ (green) and percentage of enabled sensor dimensions (black) in the bottom plot.

5.2 Sensor babbling

The previous experiments show that robots with an evolved sensory apparatus perform not worse than those with a sensory apparatus designed by hand. Here, we experimentally investigate whether, given a fixed sensory apparatus, explicitly favoring sensor babbling during the optimization can lead to robots that are more effective in the locomotion task. More precisely, we aim at answering the following two research questions:

RQ1 Is sensor babbling useful for driving evolution of faster robots?
RQ2 Is the adaptability of the VSRs affected by sensor babbling?

In order to answer these questions, we consider the worm and the biped with the High sensory apparatus and we perform the evolutionary optimization of their controllers by using a fitness function that explicitly takes into account the measure of sensor babbling defined in Section 4. In particular, given a simulation of 60 s, we use the first 20 s for measuring the "intensity" of the sensor babbling (separately for each sensor kind), and the remaining 40 s for measuring the ability in locomotion, i.e., the velocity $\overline{v}_x$. Considering that the two metrics are not in contrast, we evaluate the overall fitness of the robots by summing these two values.

We perform 10 independent runs on the two shapes, each run stopping after 10 000 fitness evaluations. After each evolutionary run, we take the best robot (in terms of fitness), and we measure its velocity over a flat terrain and

over 10 uneven terrains different from the one used for fitness evaluation—we use this as a proxy measure for the adaptability of the robot.

We repeat this entire procedure twice: in the first case (that we call "Fixed" terrain policy), each VSR faces the same uneven terrain during the evolutionary process; in the second one (that we call "Random" terrain policy), at each fitness evaluation the VSR faces a different terrain with random hills, generated with a distinct seed.

5.2.1 RQ1: Sensor babbling effectiveness

Analyzing the setting with sensor babbling, we find that the component of the fitness related to babbling remains constant during the entire evolution, while the velocity component is the only part of the fitness that increases. Therefore, in order to evaluate the actual effectiveness of babbling, we compare the results of the setting with sensor babbling with the results obtained from two other settings where the fitness does not take into account any babbling measure. For a fair comparison, we consider two comparative cases: in the first one, that we refer to as "only velocity", the entire locomotion tasks lasts 40 s and the fitness is the velocity over this period; in the second one, that we refer to as "last velocity", the duration of the task is set to 60 s, but velocity is measured only in the last 40 s, because the first 20 s are not considered. We perform this comparison in order to verify if, considering the different fitness function definitions, the sensor babbling component assumes an important role in enhancing the locomotion capability of the VSRs and that only leaving robots free to move at the beginning of the locomotion task (without measuring babbling) is not sufficient to increase the performance.

The results, shown in Figure 13 (mean ± std. dev. across the 10 runs), indicate that, considering the sensor babbling as the first phase of the robot lifetime, only in some configurations this allows to produce better performances than the ones achieved in the settings where no babbling fitness component is measured. More in detail, with the worm the results obtained with babbling are in general comparable to, or slightly worse than, those obtained in the two settings without babbling. In the case of the biped, instead, the sensor babbling setting shows comparable or better performances (with both the terrain policies) for all the sensors except for the vision one. The reason for this is quite straightforward: to read different values from vision sensors, the biped has to present a behavior which is not favorable to a fast locomotion, i.e., in order to allow vision sensors on the top to read different values it has to bend the front "leg".

5.2.2 RQ2: Sensor babbling adaptability

For this analysis, we test over a flat terrain and over 10 uneven (randomly generated) ones the best 10 robots (one for each run) found in each configuration from the previous experiment. The results, which are shown in Figure 14, indicate that over unseen terrains bipeds evolved with sensor babbling are at

least as capable as the ones without babbling: in particular, if the "Random" terrain policy is adopted, the adaptability of the VSRs slighly increases. The results for the worm show another interesting aspect: while with the "Random" terrain policy the results are consistent with the previous ones, with the "Fixed" terrain policy the touch and vision settings are capable of reaching similar performances with respect to the ones evolved considering the "only velocity" setting, which instead outperforms them in the previous experiment. However, we should note that we cannot draw any sharp conclusion on this regard since the differences in the results shown in the figure do not have statistical significance according to the Mann-Whitney U test with Bonferroni correction ($\alpha = 0.05$), except for a few occasional cases.

6 Concluding remarks and future works

In this paper we have investigated two aspects related to the sensory apparatus of a 2-D variant of simulated Voxel-based Soft Robots.

In the first part of the work, we evaluated if and how evolutionary optimization can be used to evolve the sensory apparatus. In order to do that, we considered a locomotion task with two predefined shapes, a worm and a biped, that remain fixed during the evolutionary process. We compared the results of the evolution of the sensory apparatus under two genotype representations, referred to as "Limiting" and "Unlimiting", against three manually designed sensory configurations with different levels of available sensory information. In all cases, the results of the evolutionary search were at least comparable to those of the handcrafted sensory configurations. Furthermore, we analyzed the results in terms of performance difference between robots that use sensors and robots that do not use them: although different kinds of sensors have different effects on the fitness, we demonstrated that VSRs equipped with sensors have a clear evolutionary advantage over the others. The resulting evolutionary trend indeed presents phases that indicate that evolution effectively "learns" to use the sensors that are most beneficial to the task. From a practical point of view, it is also important to note that the use of the Limiting representation can lead to minimal yet efficient sensor configurations and thus simplify the manufacturing of the soft robots, while also reducing their points-of-failure and energy consumption.

In the second part of the work, we focused on the lifetime of the VSRs and, in particular, we investigated if and how the information perceived from sensors can be elaborated more efficiently by introducing a robot life phase that we called "sensor babbling". The latter occurs in the first part of the robot lifetime and consists in leaving the robot "free to move". Our hypothesis was that thanks to this procedure, evolution can optimize the way robots elaborate the sensor readings and thus improve their performance on the task at hand (in this case, locomotion). The results have shown that, at least for the biped shape, this approach produces comparable or better performance with respect to those obtained not taking into account babbling as part of the

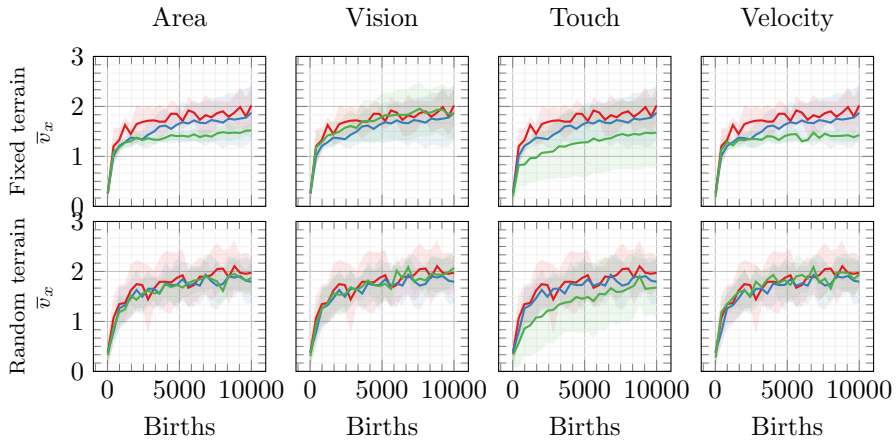

(a) Worm shape.

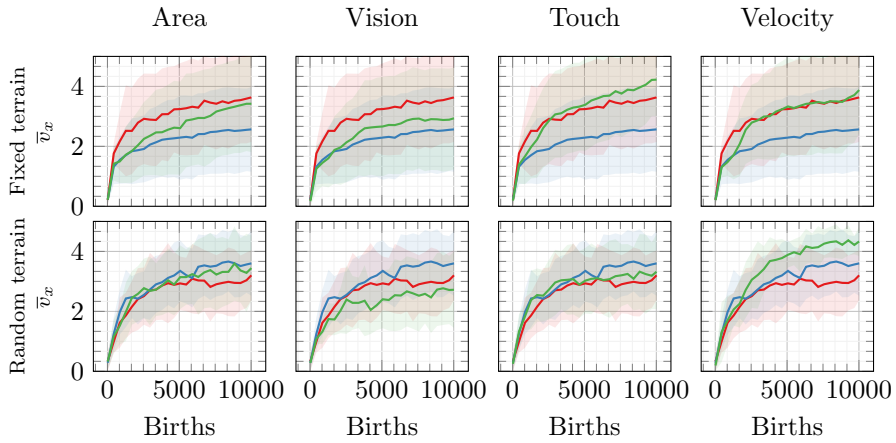

(b) Biped shape.

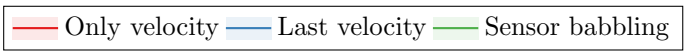

Fig. 13 Velocity $\overline{v}_x$ of the best evolved VSRs with sensor babbling and in two cases where babbling is not taken into account in the fitness (separate runs for each sensor kind): "only velocity" indicates the case where the locomotion task lasts for 40 s; "last velocity" indicates the case where the first 20 s are discarded.

fitness. Concerning the adaptability, our results indicate that in most cases the robots evolved with sensor babbling are at least as capable as the ones without babbling ad dealing with over unseen terrains.

While we cannot state that our findings apply to more complex robots (e.g., 3-D simulated VSRs, simulated rigid-body robots, real robots) or robot with different controllers (e.g., neural networks with a more complex topology than ours), we believe that they indicate that the sensory apparatus of robots

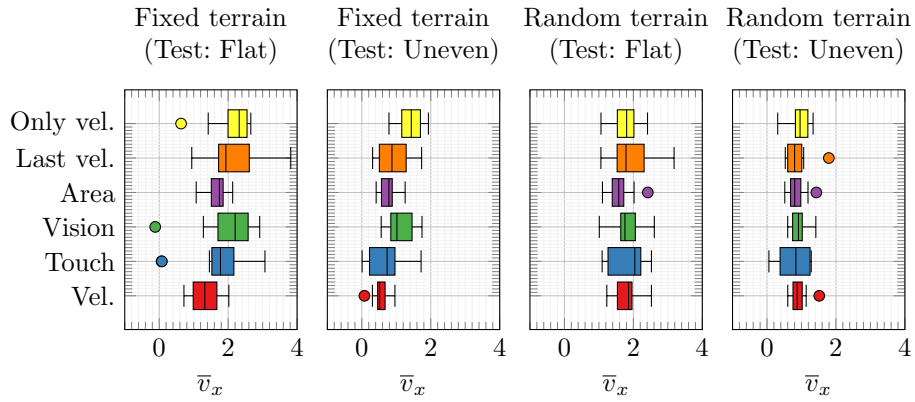

(a) Worm shape.

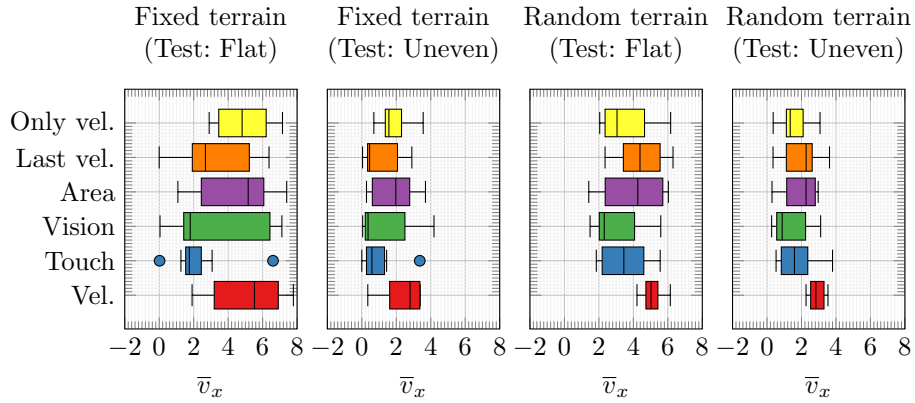

(b) Biped shape.

Fig. 14 Distribution of the velocity $\overline{v}_x$ of the best evolved VSRs found in the experiment of Section 5.2.1, facing a terrain unseen during the evolutionary process. The labels "Fixed terrain" and "Random terrain" indicate respectively if the VSRs have been evolved using the "Fixed" terrain policy or the "Random" one. The labels "Test: Flat" and "Test: Uneven" indicate respectively if the evolved VSRs have been tested on a flat or an uneven terrain.

is a key component whose design may benefit from evolutionary optimization. More broadly, our results confirm the importance of EAs in soft robot design, extending their use to the evolution of the sensory apparatus, beyond the conventional optimization of the body shape and controller. Furthermore, the present work encourages research directions that aim to consider alternative evolutionary paradigms, such as MAP-Elites [17] and quality diversity algorithms [34], similar to the recent work on modular rigid robots [35]. These diversity-driven algorithms might indeed be able to explore the search space even more effectively than traditional fitness-driven EAs. Other interesting opportunities would be to include explicit energy constraints in the sensory apparatus evolution, with the aim to reduce the overall robot complexity and

reduce the reality gap. Finally, another interesting research direction would be the co-evolution of body (including both shape *and* sensory apparatus) and controller in VSRs.

References

1. Rossiter, J.: Soft robotics: the route to true robotic organisms. Artificial Life and Robotics (2021) 1–6
2. Cheney, N., MacCurdy, R., Clune, J., Lipson, H.: Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In: Proceedings of the Genetic and Evolutionary Computation Conference. (2013) 167–174
3. Zappetti, D., Mintchev, S., Shintake, J., Floreano, D.: Bio-inspired tensegrity soft modular robots. In: Proceedings of the Conference on Biomimetic and Biohybrid Systems, Springer (2017) 497–508
4. Lee, C., Kim, M., Kim, Y.J., Hong, N., Ryu, S., Kim, H.J., Kim, S.: Soft robot review. International Journal of Control, Automation and Systems **15**(1) (2017) 3–15
5. Shah, D., Yang, B., Kriegman, S., Levin, M., Bongard, J., Kramer-Bottiglio, R.: Shape Changing Robots: Bioinspiration, Simulation, and Physical Realization. Advanced Materials (2020) 2002882
6. Howison, T., Hauser, S., Hughes, J., Iida, F.: Reality-assisted evolution of soft robots through large-scale physical experimentation: a review. arXiv preprint arXiv:2009.13960 (2020)
7. Mintchev, S., Zappetti, D., Willemin, J., Floreano, D.: A soft robot for random exploration of terrestrial environments. In: Proceedings of the International Conference on Robotics and Automation, IEEE (2018) 7492–7497
8. Cheney, N., Bongard, J., Lipson, H.: Evolving soft robots in tight spaces. In: Proceedings of the Genetic and Evolutionary Computation Conference. (2015) 935–942
9. Hallawa, A., Iacca, G., Sariman, C., Rahman, T., Cochez, M., Ascheid, G.: Morphological evolution for pipe inspection using robot operating system (ROS). Materials and Manufacturing Processes **35**(6) (2020) 714–724
10. Song, Y.S., Sun, Y., Van Den Brand, R., Von Zitzewitz, J., Micera, S., Courtine, G., Paik, J.: Soft robot for gait rehabilitation of spinalized rodents. In: Proceedings of the International Conference on Intelligent Robots and Systems, IEEE (2013) 971–976
11. Zhang, B., Fan, Y., Yang, P., Cao, T., Liao, H.: Worm-like soft robot for complicated tubular environments. Soft robotics **6**(3) (2019) 399–413
12. Hiller, J., Lipson, H.: Automatic design and manufacture of soft robots. IEEE Transactions on Robotics **28**(2) (2011) 457–466
13. Lee, H., Jang, Y., Choe, J.K., Lee, S., Song, H., Lee, J.P., Lone, N., Kim, J.: 3D-printed programmable tensegrity for soft robotics. Science Robotics **5**(45) (2020)
14. Eiben, A.E., Hart, E., Timmis, J., Tyrrell, A.M., Winfield, A.F.: Towards autonomous robot evolution. In: Software Engineering for Robotics. Springer (2021) 29–51
15. Kriegman, S., Cheney, N., Bongard, J.: How morphological development can guide evolution. Scientific reports **8**(1) (2018) 1–10
16. Medvet, E., Bartoli, A., De Lorenzo, A., Fidel, G.: Evolution of distributed neural controllers for voxel-based soft robots. In: Proceedings of the Genetic and Evolutionary Computation Conference. (2020) 112–120
17. Zardini, E., Zappetti, D., Zambrano, D., Iacca, G., Floreano, D.: Seeking quality diversity in evolutionary co-design of morphology and control of soft tensegrity modular robots. In: Proceedings of the Genetic and Evolutionary Computation Conference. (2021) 189–197
18. Medvet, E., Bartoli, A., Pigozzi, F., Rochelli, M.: Biodiversity in evolved voxel-based soft robots. In: Proceedings of the Genetic and Evolutionary Computation Conference. (2021) 129–137
19. Sims, K.: Evolving virtual creatures. In: Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques. (1994) 15–22

20. Balakrishnan, K., Honavar, V.: On sensor evolution in robotics. In: Proceedings of the First International Conference on Genetic Programming, Citeseer (1996) 455–460
21. Mautner, C., Belew, R.K.: Evolving robot morphology and control. Artificial Life and Robotics **4**(3) (2000) 130–136
22. Powers, J., Grindle, R., Kriegman, S., Frati, L., Cheney, N., Bongard, J.: Morphology dictates learnability in neural controllers. In: Proceedings of the Artificial Life Conference, MIT Press (2020) 52–59
23. Ferigo, A., Iacca, G., Medvet, E.: Beyond body shape and brain: Evolving the sensory apparatus of voxel-based soft robots. In: Applications of Evolutionary Computation. Volume 12694., Springer, Cham (2021) 210–226
24. Hiller, J., Lipson, H.: Dynamic Simulation of Soft Multimaterial 3D-printed Objects. Soft Robotics **1**(1) (2014) 88–101
25. 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** (2020) 100573
26. Talamini, J., Medvet, E., Bartoli, A., De Lorenzo, A.: Evolutionary Synthesis of Sensing Controllers for Voxel-based Soft Robots. In: Proceedings of the Artificial Life Conference, MIT Press (2019) 574–581
27. Nadizar, G., Medvet, E., Pellegrino, F.A., Zullich, M., Nichele, S.: On the effects of pruning on evolved neural controllers for soft robots. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion. (2021) 1744–1752
28. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. Evolutionary Computation **9**(2) (2001) 159–195
29. Medvet, E., Bartoli, A.: Evolutionary optimization of graphs with graphea. In: Proceedings of International Conference of the Italian Association for Artificial Intelligence, Springer (2020) 83–98
30. Rothlauf, F., Goldberg, D.E.: Redundant representations in evolutionary computation. Evolutionary Computation **11**(4) (2003) 381–415
31. Demiris, Y., Dearden, A.: From motor babbling to hierarchical learning by imitation: a robot developmental pathway. In: Proceedings of the Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems. Volume 123., Lund University Cognitive Studies (2005) 31–37
32. Saegusa, R., Metta, G., Sandini, G., Sakka, S.: Active motor babbling for sensorimotor learning. In: Proceedings of the International Conference on Robotics and Biomimetics, IEEE (2009) 794–799
33. Hansen, N.: The CMA evolution strategy: a comparing review. In: Towards a new evolutionary computation. Springer (2006) 75–102
34. Auerbach, J.E., Iacca, G., Floreano, D.: Gaining insight into quality diversity. In: Proceedings of the Genetic and Evolutionary Computation Conference – Companion. (2016) 1061–1064
35. Nordmoen, J., Veenstra, F., Ellefsen, K.O., Glette, K.: Quality and Diversity in Evolutionary Modular Robotics. arXiv preprint arXiv:2008.02116 (2020)