Beyond Body Shape and Brain: Evolving the Sensory Apparatus of Voxel-based Soft Robots

Andrea Ferigo[1][0000−0003−1795−011X], Giovanni Iacca[1][0000−0001−9723−1830], and Eric Medvet[2][0000−0001−5652−2113]

[1] Department of Information Engineering and Computer Science, University of Trento, Italy
[2] Department of Engineering and Architecture, University of Trieste, Italy

Abstract. Biological and artificial embodied agents behave by acquiring information through sensors, processing that information, and acting on the environment. The sensory apparatus, i.e., the location on the body of the sensors and the kind of information the sensors are able to capture, has a great impact on the agent ability of exhibiting complex behaviors. While in nature, the sensory apparatus is the result of a long-lasting evolution, in artificial agents (robots) it is usually the result of a design choice. However, when the agents are complex and the design space is large, making that choice can be hard. In this paper, we explore the possibility of evolving the sensory apparatus of voxel-based soft robots (VSRs), a kind of simulated robots composed of multiple deformable components. VSRs, due to their intrinsic modularity, allow for great freedom in how to shape the robot body, brain, and sensory apparatus. We consider a set of sensors that allow the agent to sense itself and the environment (using vision and touch) and we show, experimentally, that the effectiveness of the sensory apparatus depends on the shape of the body and on the actuation capability, i.e., the VSR strength. Then we show that evolutionary optimization is able to evolve an effective sensory apparatus, even when constraints on the availability of the sensors are posed. By extending the adaptation to the sensory apparatus, beyond the body shape and the brain, we believe that our study takes a step forward to the ambitious path towards self-building robots.

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

1 Introduction

Soft robots are deemed to be one of the key technologies for the future of mankind: compared to traditional hard robotics, they allow in fact better compliance with the environment and humans, leading to higher safety in mission-critical applications. Some of the most relevant examples of soft robots are the voxel-based soft robots (VSRs) [1], and the tensegrity soft modular robots (TSMRs) [2], although other paradigms also exist [3–5]. Common to most of these robotics structures, their main features are intrinsic softness and flexibility, which allow them to perform tasks that are otherwise incredibly difficult,

if not impossible, for hard robots: for instance, soft robots can perform smooth locomotion on rough terrain [6], or squeeze through tight spaces [7]. These possibilities make them ideal tools for complex robotic inspection applications, such as the exploration of hard-to-access environments [8]. In the medical domain, they have been proposed as a support for gait rehabilitation [9] and colonoscopies [10].

Another important feature of soft robots is that they are, usually, inherently prone to *modular design*. Often biologically inspired, this kind of robot can be assimilated from complex organs made of multiple components of the same kind, similar, e.g., to the myons that make up skeletal/muscle tissues. Furthermore, modularity facilitates manufacturing, redundancy, and repair [11, 12].

Despite these promises, designing soft robots' morphology and controller, (sometimes referred to as the *body* and the *brain*, respectively) is especially difficult. 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), often coupled with physics-based simulations, the main tool currently available for designing soft modular robots. The importance of EAs in designing soft modular robots has been highlighted, for instance, in [13], where the authors not only found that evolutionary optimized morphologies allow different kinds of interactions with the environment, but they have also shown that morphological development can, in turn, guide evolution to more robust designs. Other recent examples of evolutionary synthesis applied to soft robots are reported in [14, 15], where the authors used an EA to optimize the controller of VSRs. Overall, using EAs for soft modular robot design is clearly advantageous for a number of reasons. Artificial evolution has in fact the potential to uncover unconventional designs, difficult to anticipate for a human expert, that not only are optimized for their efficiency at the task at hand, but can also show improvements in terms of non-functional requirements, 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.

A currently overlooked aspect in the evolutionary design of soft modular robots is the sensory apparatus, i.e., the kinds of sensors available to the controller and their position in the body. Usually, sensors are considered as given, based on expert design choices and/or physical constraints, and controllers are optimized to use the available sensors. However, *what happens if the sensory apparatus can evolve?* Or, put it in other terms, 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 main research question. Our hypothesis is that evolving the sensory apparatus may either lead to use of fewer sensors (with respect to a manually designed sensory apparatus), or lead to more efficient use of the existing sensors, which can yield better robot performance and, potentially, better energy usage. To the best of our knowledge, the only work that addressed a similar research question are [16–18], where the sensory apparatus is evolved in different kinds of hard robots. Another work on this topic is [19], which has shown that sensor

placement can alter the landscape of the controller loss function, thus guiding evolution towards better controllers. However, the context of that work was unicycle non-holonomic mobile robots. On the other hand, no prior research so far addressed this question in the domain of soft modular robots and, in particular, of VSRs. Due to intrinsic and fine-grained modularity, VSRs offer indeed great freedom in the design of the sensory apparatus.

In order to answer our question, we conduct *in silico* experiments on two kinds of VSRs, a biped and a worm, and use an EA to optimize the sensory apparatus to perform 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. We also evaluate the generalizability of the solutions found by the EAs on different terrains characterized by various levels of roughness and the presence of different kinds of obstacles. In a nutshell, we find that EAs are able to find sensory configurations that perform at least as well as those manually designed based on our previous knowledge, even when the number of available sensors is constrained to be smaller than that available to manual designs. This additional constraint is important as 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 and sensory apparatus, and the evolutionary importance of discovering the kinds of sensors that are actually beneficial to the task.

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, we describe how we represent VSRs in a way that fits evolutionary optimization. In Section 4, we present the experiments and discuss the results. Finally, in Section 5, we draw the conclusions of this work and discuss possible future research directions.

2 Background: voxel-based soft robots

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 [20]. In this study, we consider the 2-D version of the VSRs presented in [21] along with a simulation engine tailored to optimization. We here briefly recap the main concepts of VSRs that are relevant to our study: we refer the reader to the aforecited paper and to [22] for further details.

A VSR is defined by its *morphology* and its *controller*. The morphology is itself defined by the number and placement of the voxels composing the VSR, that we call *shape*, and by the number, kind, and placement of sensors, that we call *sensory apparatus*. The controller is a law that determines, at each time step, the control signal to be applied to each voxel of the VSRs based on the readings of the sensors available in the sensory apparatus.

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). A voxel is a deformable square

that is modeled, in the simulation, as a compound of masses, spring-dampers systems, and distance constraints, see [21].

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 precise 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. 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 controlled with the same control signal.

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 outcome of the perception is consumed by the controller in order to determine the values of the control signal, which in turn affects the behavior of the VSR.

The sensory apparatus of a VSR consists of zero or more sensors for each voxel of the VSR shape. 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 sensor types. *Area* sensors perceive the ratio between the current area of the voxel and its rest area: the domain is hence $D =]0, +\infty[$. *Touch* sensors perceive whether the voxel is touching the ground ($\boldsymbol{s} = 1$) or not ($\boldsymbol{s} = 0$): the domain is $D = \{0, 1\}$. *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 $D = \mathbb{R}^2$. Finally, *vision* sensors perceive the distance towards close objects, as the terrain and the obstacles, within some field of view.

We designed the vision sensors for the purpose of this study and modeled them as 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 $D = [0, 1]^p$. Figure 1 shows two examples of VSRs equipped with vision sensors and highlights the rays cast by those sensors.

Controller. The controller determines the value of the control signals of the VSR voxels over time. Several forms may be employed for realizing the controller. Since we are interested in studying if and how the sensory apparatus of the VSR can be evolved, we use a controller that can exploit the sensor readings. In

particular, we followed the approach of [14], where the controller is a multi layer perceptron (MLP), a form of artificial neural network.

The MLP 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 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 [14]—we used a sinusoidal function with frequency of 1 Hz. The MLP has one output for each voxel: the value determines, at each time step, the control signal of the corresponding voxel. The input and output layers are connected through zero or more inner layers. In this work, we experimented with MLP without inner layer and we used tanh as activation function.

In this form, the controller is completely described by the weights $\boldsymbol{w}$ of the MLP. For a given VSR shape, the size $|\boldsymbol{w}|$ of the weight vector depends on the sensory apparatus of the VSR.

3 Evolution of the sensory apparatus

We are interested in verifying if the sensory apparatus of the VSR can be optimized by means of evolutionary computation, as the shape and the controller have been showed to be. However, the controller is intrinsically connected to the sensory apparatus: in the form that we consider in this paper, the controller size 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 according to which a numerical vector (the *genotype*) is mapped, given a shape, to a pair ⟨sensory apparatus, controller⟩ (the *phenotype*) suitable for that shape. We use a numerical vector as genotype since this enables us to use state-of-the-art evolutionary algorithms for the search. For this study, we rely on the Covariance Matrix Adaptation Evolution Strategies (CMA-ES) algorithm [23], since it has been shown to be effective for VSRs [24], as we confirmed with exploratory experimentation.

The two representations are direct, in the sense that an element of the numerical vector constituting the genotype directly determines the realization of a component of the controller and/or sensory apparatus. Both also exhibit some degree of redundancy, itself resulting in possible degeneracy of the representation [25]: in some portion of the genotype space, several different genotypes correspond to the same phenotype.

Furthermore, 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.

The two representations differ in the possibility of further limiting the maximum number of sensors: we hence call them *Limiting* and *Unlimiting*. They work as follows. 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 MLP 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 by $\boldsymbol{w}_{(s)}$ the vector of weights corresponding to the inputs of the MLP connected to the sensor s.

Unlimiting. In this 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 controller is the MLP obtained by considering only the maximal controller inputs connected to the sensory apparatus obtained from the genotype. 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, the representation allows for a sensory apparatus consisting of no sensors at all.

Limiting. In this representation, given a maximum number $n_{\text{sensors}} \in \mathbb{N}$ of sensors, the sensory apparatus corresponding to a genotype $\boldsymbol{w}$ is composed of the n_{sensors} sensors with the largest $w_{s,\max}$. The controller is set as in the Unlimiting case. It can be seen 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.

4 Experimental analysis

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

- 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 the evolution discover an effective sensory apparatus?
- RQ3 When the complexity of the evolvable apparatus is limited, what sensors are "preferred" by the evolution?

To answer these questions, we performed a number of experiments considering two VSR shapes and the task of locomotion on an uneven terrain. Concerning the shapes, we experimented with a 7×1 (x and y size of the voxel grid constituting the shape) rectangle, that we call *worm*, and a 7×4 rectangle with 3×2 missing voxels at bottom-center, that we call *biped*. Figure 1 shows the two shapes during one simulation.

Concerning the task, we considered *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 measured the effectiveness of a VSR in performing locomotion as its average speed $\overline{v}_x = \frac{x(t_f)-x(0)}{t_f}$ during the episode, t_f being the episode

duration and $x(t)$ being 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 surface. On the other hand, here we used an uneven surface, since we believe that it is better suited for investigating sensing ability: robots that are able to perceive obstacles along the way should be favored with respect to robots that base their locomotion effectiveness on a regular gait not impacted by current perception [14]. Moreover, for a subset of the experiments, we measured the locomotion effectiveness $\overline{v}_x$ on two different surfaces, one used for the evolution, one after: the motivation was again to verify if the robot actually exploited the sensory apparatus for running faster on an "unknown" surface.

For all the experiments, we used CMA-ES for the optimization with default parameter settings (as indicated in [26], 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 simulated episodes lasting $t_f = 180\,\text{s}$ (simulated time). For the simulation, we used the default parameters of the 2D-VSR-SIM software [21], unless otherwise specified. We made the code used for the experiments, based on JGEA (`https://github.com/ericmedvet/jgea`) for the evolutionary optimization part, publicly available on `https://github.com/ndr09/HSMRcoevo`.

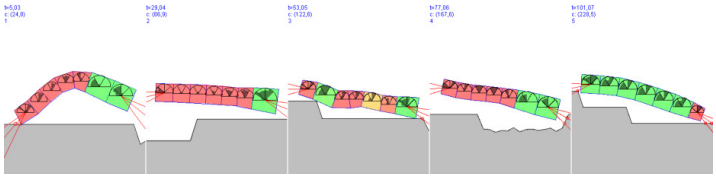

(a) Worm on the training surface.

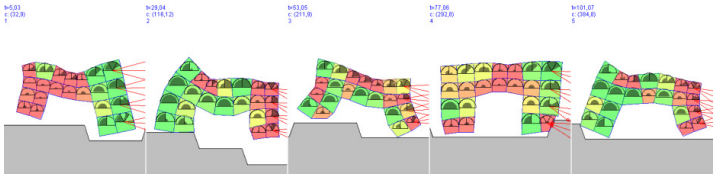

(b) Biped on the validation surface.

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 rest 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 [21]). The rays of the vision sensors are shown in red.

4.1 RQ1: Sensor potential benefit

In order to answer this question, we manually designed, based on domain knowledge, three sensory apparatuses for each of the two shapes. For each resulting

body, we evolved the controller with different values of ρ_A, i.e., different strengths of the bodies (see Section 2).

Figure 2 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, we placed in the maximal sensory apparatus:

- 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 placed 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 2.

We then differentiated the three apparatuses based on the number of sensors that we removed from the corresponding maximal sensory apparatus. In the *Low* perception apparatus, we left just a few area sensors. In the *Medium* perception apparatus, we removed all the vision sensors. In the *High* perception apparatus, we did 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 number $|\boldsymbol{w}|$ of weights for the three controllers, i.e., the size of the search space from the point of view of the evolutionary optimization, was 35, 210, and 252, for the worm, and 352, 924, and 1188, for the biped, respectively for the Low, Medium, and High apparatuses.

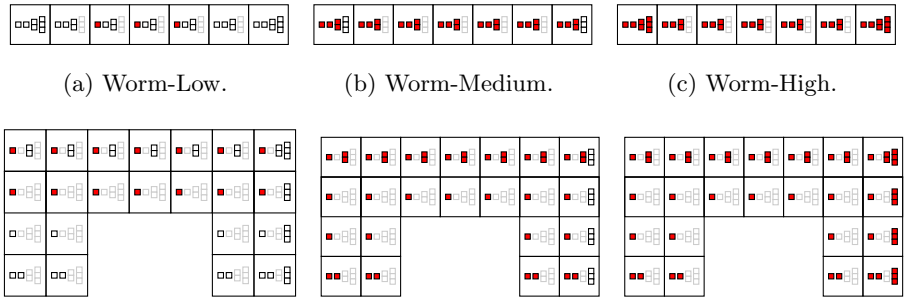

(a) Worm-Low. (b) Worm-Medium. (c) Worm-High.

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

Fig. 2. 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). Area sensor ($p = 1$) is the first (leftmost) column; touch sensor ($p = 1$) is the second column; velocity sensor is the third ($p = 2$); vision sensor is the last ($p = 3$). 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.

We considered 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, sensory apparatus, and ρ_A, we performed 10 independent evolutionary optimizations of the controller with CMA-ES. We used the average speed $\overline{v}_x$ as fitness and we stopped the evolution after 2000 fitness evaluations.

Figure 3 summarizes the results of this experiment. It shows the mean (across the 10 runs) effectiveness in locomotion (i.e., $\overline{v}_x$) of the best VSR at the last iteration of CMA-ES with the three sensory apparatuses on the two shapes for each value of ρ_A. The figure also shows the $\overline{v}_x$ that the best VSR scores on a second *validation* surface, i.e., an uneven surface different from the one used for computing the fitness (*training* surface).

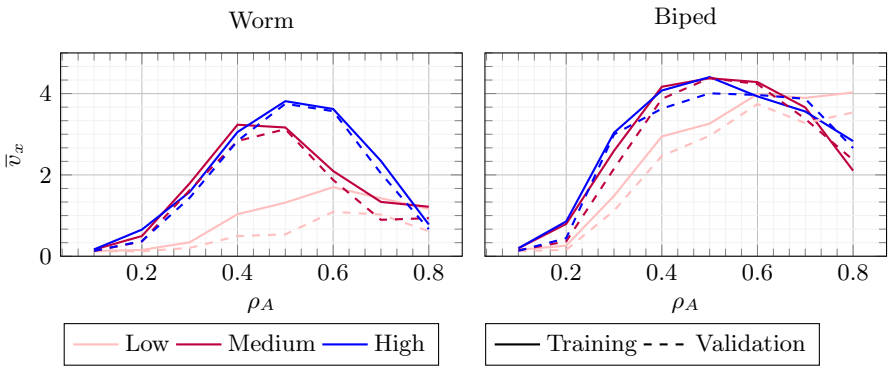

Fig. 3. Locomotion effectiveness $\overline{v}_x$ of evolved VSRs with different shapes (plot) and sensory apparatuses (line color) vs. the strength ρ_A.

Three observations can be made based on Figure 3. First, the locomotion effectiveness depends on the robot strength. As expected, the stronger the robot, the faster in locomotion; but this holds only for values of ρ_A lower than 0.5. Instead, for $\rho_A > 0.5$, the VSRs of both shapes become less effective with greater strength: we examined the behavior of the evolved VSRs in details and found that with large ρ_A the physics model of the simulator is brought to the limit and results in voxels that contract too much, eventually making the robot difficult to control and hence ineffective.

Second, the differences in effectiveness among sensory apparatuses depend on the shape. For the worm, the more complex the sensory apparatus, the faster the robot: moreover, the three apparatuses peak at different values of strength and the Low apparatus is greatly outperformed by the Medium and High. 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 $[0.2, 0.5]$.

Third, the difference between the effectiveness on the training and validation surfaces appears to be larger for the Low sensory apparatus—this is more apparent for the worm shape. This means that when the VSR bases its locomotion

effectiveness also on the perception of the environment, it is more capable of coping with unseen environmental conditions.

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. 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.

4.2 RQ2: Effectiveness of evolved sensory apparatus

For answering this question, we performed a second experiment where we compared the results of the evolutionary optimization of the sensory apparatus of the two shapes, using the two representations described in Section 3, against the three apparatuses defined in the previous section (Low, Medium, and High). Moreover, we used the High apparatus as the maximal sensory apparatus when evolving the apparatus: the number $|\boldsymbol{w}|$ of weights for the evolvable apparatuses was 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 determined the value of the weight threshold τ_w by examining the values of the weights of the $10 + 10$ (worm and biped) evolved VSRs with the High apparatus in the previous experiment (see Figure 4): we chose, for each shape, the value corresponding to the tenth percentile of $w_{s,\max}$ (see Section 3), that was $\tau_w = 1.96$ for the worm and $\tau_w = 2.87$ for the biped.

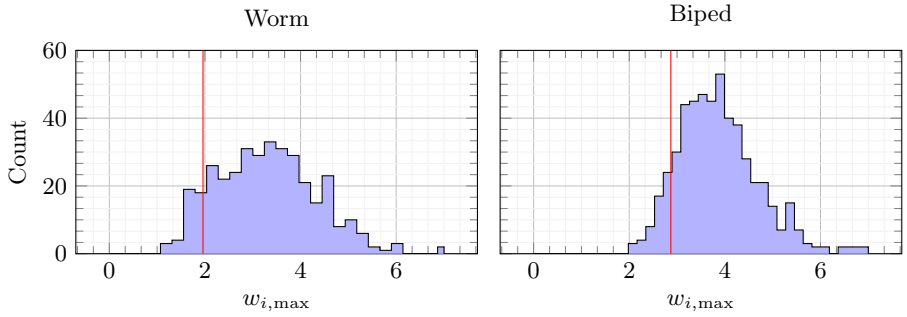

Fig. 4. Histograms of the value of $w_{s,\max}$ for sensors of the VSRs evolved with the High sensory apparatus in the experiment of Section 4.1. The red line corresponds to the 10-th percentile.

We set $\rho_A = 0.3$ and we performed 10 independent evolutionary optimizations for each one of the $2 \times (3 + 2)$ configurations (two shapes and three statically defined apparatuses plus two evolvable ones). We stopped the evolution after 25 000 fitness evaluations. For the evolvable apparatuses, we initialized the vector of means in CMA-ES to $\mathbf{0}$ instead of sampling each element $[-1, 1]$: in this way, the sensors started disabled and was up to the evolution to enable them when convenient.

Figure 5 summarizes the results of this experiment: it shows the mean fitness $\overline{v}_x$ (across the 10 runs) of the best VSR at different stages of the evolution (i.e., number of births).

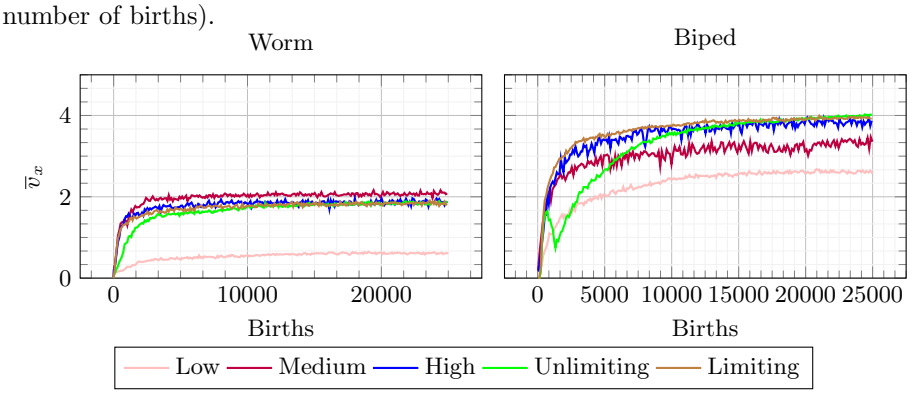

Fig. 5. Mean fitness $\overline{v}_x$ (across 10 runs) during the evolution in the second experiment.

The foremost finding that can be seen in Figure 5 is that the evolved sensory apparatuses are not, in general, worse than the hand-designed ones. By looking in details at the value of $\overline{v}_x$ at the end of the optimizations, which is shown in Figure 6 in the form of box plots, 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. As a further confirmation, we performed the Mann-Whitney U test (after having verified the relevant hypotheses) between pairs of samples from different apparatuses for each shape: Table 1 shows the resulting p-values.

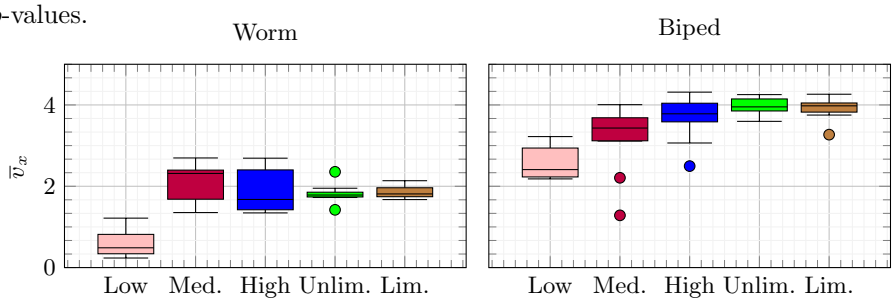

Fig. 6. Box plots of the fitness $\overline{v}_x$ at the end of the evolution in the second experiment.

By comparing the two representations for the evolvable apparatuses it can be seen that there are no significant differences. This means that with the Limiting representation the evolution is still able to find a sensory apparatus (and the corresponding controller) that is effective in locomotion, even if the VSR can use fewer sensors.

Finally, by looking at the $\overline{v}_x$ during the evolution with Unlimiting representation, i.e., the the green line in Figure 5, it can be seen that it stays behind the Limiting representation in the early stage of the evolution. For the biped in particular, the peculiar shape of the line shows that after an initial, rapid improvement of the fitness, it then decreases up to ≈ 150 births and then starts

Table 1. p-value of the pairwise Mann-Whitney U test performed on the $\overline{v}_x$ resulting from the second experiment (values shown in Figure 6). Values corresponding to a significant difference (with significance level set to 0.05) are highlighted in boldface.

	Worm				Biped			
	Med.	High	Unlim.	Lim.	Med.	High	Unlim.	Lim.
Low	**0.0002**	**0.0002**	**0.0002**	**0.0002**	**0.0284**	**0.0012**	**0.0002**	**0.0002**
Med.		0.5967	**0.0413**	**0.0156**		0.0821	0.0588	0.1988
High			0.4057	0.2899			0.7624	0.4497
Unlim.				0.3258				0.4057

to improve again—recall that CMA-ES does not guarantee the monotonicity of the fitness. We looked at the raw results and found that this is because of the progressive enabling of the sensors: at the beginning, the 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 the evolution to "adjust" the controller, but they later become beneficial allowing the fitness to increase again. This interpretation is supported by Figure 7, that shows how the fraction of VSRs in the population that do not have sensors (that is initially 100 % because of the initialization of CMA-ES means to $\mathbf{0}$) quickly goes to 0 % in the early stage of the evolution. This phenomenon is less visible in the fitness of the Unlimiting representation for the worm since it has much fewer sensors to be activated. Moreover, it is not present at all for the Limiting representation that, by design, results in a sensory apparatus that has always n_{sensors} sensors.

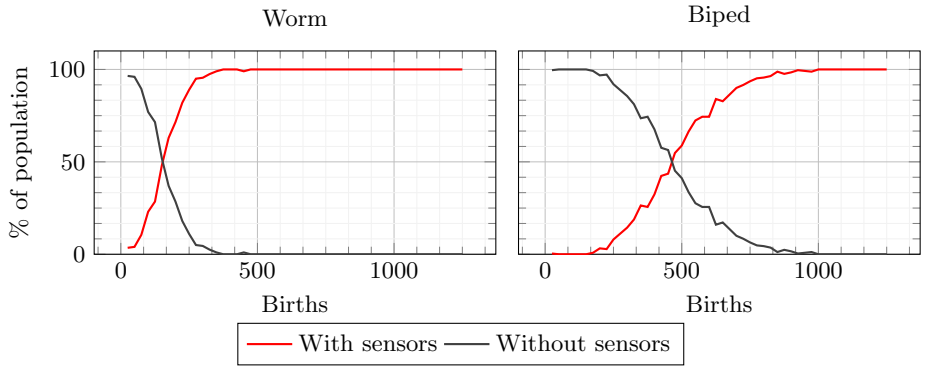

Fig. 7. Percentage of the individual of the population not using (gray) or using (red) sensors in the early stage of the evolution (up to 1250 births) with the Unlimiting representation.

4.3 RQ3: Sensors preferred by the evolution

We attempted to answer this question in two ways. First, quantitatively, by analyzing which sensors remain not enabled (with respect to the maximal sensory apparatus) in the apparatuses evolved with the Limiting representation—we recall that, by design, the Limiting representation does not enable all the sensors.

Second, qualitatively, by examining the enabling order of the sensors in the evolutionary optimizations performed with the Unlimiting—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 repeated the experiments of the previous section after having slightly modified the two representations. In this experiment, we allowed single dimensions of the sensors to be enabled independently: e.g., a worm could have just the front vision sensor with $\alpha_3 = -30°$ enabled, and the other dimensions ($\alpha_1 = 0°$ and $\alpha_2 = -15°$) disabled.

Table 2 shows the results corresponding to the first, quantitative analysis. It shows the number and percentage of sensors, for each type and shape, that are enabled at the end of the evolution with the Limiting representation. We considered a sensor as enabled if *at least one dimension* of the sensor was enabled.

Table 2. Number and percentage (mean across 10 runs) of sensors, for each type and shape, that are enabled at the end of the evolution with the Limiting representation, based on the number of sensors of the corresponding type that are present in the maximal sensory apparatus of the corresponding shape (first row of each shape).

		Area	Touch	Velocity	Vision
Worm	Number in maximal	7	7	7	2
	Number in evolved	5.7	5.8	6.7	2.0
	% in evolved	81.4	82.9	95.7	100.0
Biped	Number in maximal	22	4	7	4
	Number in evolved	16.4	3.4	6.8	4.0
	% in evolved	74.5	85.0	97.1	100.0

It can be seen from Table 2 that for both shapes the vision sensors result all enabled (in all runs). Velocity and touch sensors are enabled less frequently and with similar percentages in the two shapes. Finally, the area sensors appear to be the least used in both shapes. In the biped, $\approx$25 % of them 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, the evolution still finds an effective solution for locomotion (see Section 4.2) and hence "is aware" of this redundancy.

Concerning the qualitative analysis, we report here the detailed outcome for one of the runs with the Unlimiting representation. Figure 8 shows the sensory apparatus of the best biped (represented with the graphical notation of Figure 2) in four salient points of the evolution, 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 evolution, we chose as 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.

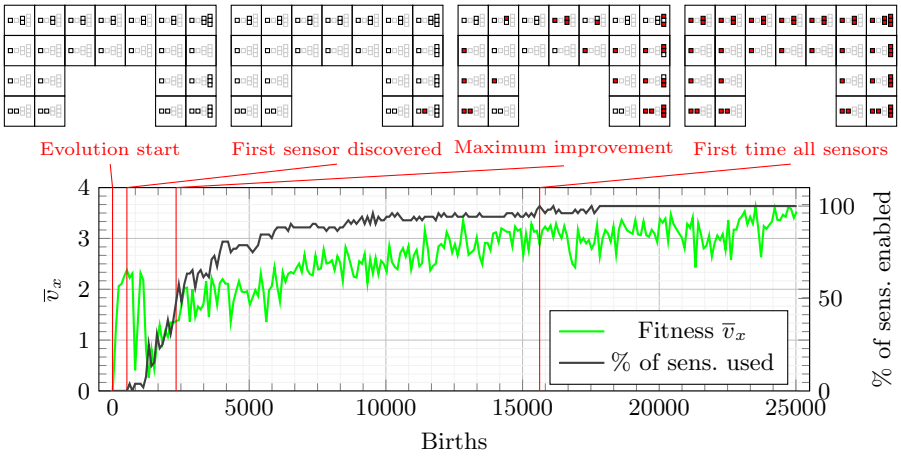

Fig. 8. The sensory apparatus of the biped in four salient moments of the evolution (see text) for one of the evolutionary runs. The plots of the VSRs are referred to the corresponding values of fitness $\overline{v}_x$ (green) and percentage of enabled sensor dimensions (gray) in the bottom plot.

Two qualitative observations can be made by looking at Figure 8. First, in the moment of maximum improvement after the initial drop in the fitness, almost all the vision sensors are enabled, whereas the percentage of enabled sensors for the other sensor types is much lower. This clue appears to support the findings discussed above (based on Table 2) about the importance of the vision sensors.

Second, the first peak value of the fitness just after the beginning of the evolution occurs at the same (evolutionary) time of the discovery of the first sensor. After that moment, the fitness rapidly decreases for some hundreds of births while other sensors are progressively enabled, until the evolution manages to "understand" how to effectively use the discovered sensors for a faster locomotion. This observation seems to be in line with what we showed in Section 4.2 (in particular in Figure 7).

5 Concluding remarks and future work

We have investigated how evolutionary search can be used to optimize the sensory apparatus of Voxel-based Soft Robots. To do that, we considered a locomotion task performed with two predefined shapes, a worm and a biped, and compared the results of the evolution of sensory apparatus under two genotype representations, Limiting and Unlimiting, against three manually designed sensory configurations with different levels of available sensory information. Our results confirmed the importance of Evolutionary Algorithms in soft robot design, broadening their use beyond the conventional optimization of body shape and controller. In all cases, the results of evolution were at least comparable to those of the handcrafted sensory configurations. Furthermore, we found that robots using sensors have a clear evolutionary advantage over those ones that do

not use them, although different kinds of sensors have different effects on the fitness. The resulting evolutionary trends present indeed phases that indicate that evolution effectively "learns" to use the sensors that are most beneficial to the task. In practical applications, this characteristic can be used to find minimal yet efficient sensory apparatuses and thus simplify the manufacturing of the soft robots, while also reducing their points-of-failure and energy consumption.

The present work encourages a number of research directions as to what concerns, for instance, the use of alternative evolutionary paradigms, such as MAP-Elites and quality diversity algorithms [27], similar to the recent work on modular rigid robots [28]. 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, and study the co-evolution of body (including both shape *and* sensor apparatus) and controller in VSRs.

Acknowledgments and author contributions

We thank Luca Zanella for the CMA-ES and Lidar sensor implementation. We gratefully acknowledge HPC-Cineca for making computing resources available. A. F.: Investigation; Software; Data curation; Visualization; Writing - original draft. G. I.: Conceptualization; Methodology; Writing - review & editing. E. M.: Conceptualization; Methodology; Software; Visualization; Writing - review & editing.

References

1. Cheney, N., MacCurdy, R., Clune, J., Lipson, H.: Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In: Genetic and Evolutionary Computation Conference. (2013) 167–174
2. Zappetti, D., Mintchev, S., Shintake, J., Floreano, D.: Bio-inspired tensegrity soft modular robots. In: Conference on Biomimetic and Biohybrid Systems, Springer (2017) 497–508
3. 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
4. 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
5. 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)
6. Mintchev, S., Zappetti, D., Willemin, J., Floreano, D.: A soft robot for random exploration of terrestrial environments. In: International Conference on Robotics and Automation, IEEE (2018) 7492–7497
7. Cheney, N., Bongard, J., Lipson, H.: Evolving soft robots in tight spaces. In: Genetic and Evolutionary Computation Conference. (2015) 935–942
8. 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

9. 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: International Conference on Intelligent Robots and Systems, IEEE (2013) 971–976
10. 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
11. Hiller, J., Lipson, H.: Automatic design and manufacture of soft robots. IEEE Transactions on Robotics **28**(2) (2011) 457–466
12. 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)
13. Kriegman, S., Cheney, N., Bongard, J.: How morphological development can guide evolution. Scientific reports **8**(1) (2018) 1–10
14. Talamini, J., Medvet, E., Bartoli, A., De Lorenzo, A.: Evolutionary Synthesis of Sensing Controllers for Voxel-based Soft Robots. In: Artificial Life Conference, MIT Press (2019) 574–581
15. Medvet, E., Bartoli, A., De Lorenzo, A., Fidel, G.: Evolution of distributed neural controllers for voxel-based soft robots. In: Genetic and Evolutionary Computation Conference. (2020) 112–120
16. Sims, K.: Evolving virtual creatures. In: Proceedings of the 21st annual conference on Computer graphics and interactive techniques. (1994) 15–22
17. Balakrishnan, K., Honavar, V.: On sensor evolution in robotics. In: Proceedings of the First International Conference on Genetic Programming, Citeseer (1996) 455–460
18. Mautner, C., Belew, R.K.: Evolving robot morphology and control. Artificial Life and Robotics **4**(3) (2000) 130–136
19. Powers, J., Grindle, R., Kriegman, S., Frati, L., Cheney, N., Bongard, J.: Morphology dictates learnability in neural controllers. In: Artificial Life Conference, MIT Press (2020) 52–59
20. Hiller, J., Lipson, H.: Dynamic Simulation of Soft Multimaterial 3D-printed Objects. Soft Robotics **1**(1) (2014) 88–101
21. 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
22. Medvet, E., Bartoli, A., De Lorenzo, A., Seriani, S.: Design, Validation, and Case Studies of 2D-VSR-Sim, an Optimization-friendly Simulator of 2-D Voxel-based Soft Robots. arXiv (2020) arXiv–2001
23. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. Evolutionary Computation **9**(2) (2001) 159–195
24. Medvet, E., Bartoli, A.: GraphEA: a Versatile Representation and Evolutionary Algorithm for Graphs. In: Workshop on Evolutionary and Population-based Optimization (WEPO@AIxIA). (2020)
25. Rothlauf, F., Goldberg, D.E.: Redundant representations in evolutionary computation. Evolutionary Computation **11**(4) (2003) 381–415
26. Hansen, N.: The CMA evolution strategy: a comparing review. In: Towards a new evolutionary computation. Springer (2006) 75–102
27. Auerbach, J.E., Iacca, G., Floreano, D.: Gaining insight into quality diversity. In: Genetic and Evolutionary Computation Conference – Companion. (2016) 1061–1064
28. Nordmoen, J., Veenstra, F., Ellefsen, K.O., Glette, K.: Quality and Diversity in Evolutionary Modular Robotics. arXiv preprint arXiv:2008.02116 (2020)