Received 29 November 2023, accepted 8 December 2023, date of publication 14 December 2023,
date of current version 21 December 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3342912

RESEARCH ARTICLE

Human Control of Simulated Modular Soft Robots May Predict the Performance of Optimized AI-Based Controllers

GIULIA MARCHIORI PIETROSANTI[1], GIORGIA NADIZAR[1], FEDERICO PIGOZZI[2], AND ERIC MEDVET[2], (Member, IEEE)
[1]Department of Mathematics and Geosciences, University of Trieste, 34127 Trieste, Italy
[2]Department of Engineering and Architecture, University of Trieste, 34127 Trieste, Italy

Corresponding author: Eric Medvet (emedvet@units.it)

ABSTRACT Robots with a modular body permit a wide range of physical configurations, which can be obtained by arranging the composing modules differently. While this freedom makes modular robots capable of performing different tasks, finding the optimal physical configuration for a given task is not trivial. In fact, practitioners attempt to jointly optimize the body and the controller of the robot for a given task, but the result is not always satisfactory. More broadly, it is not clear what factors make a physical configuration more or less successful. In this paper, we aim to fill this gap and verify if humans can be predictive with respect to the performance of an optimized controller for a given robot body. We consider the case of Voxel-based Soft Robots (VSRs), whose rich dynamic induced by the softness of the modules makes the body particularly relevant for the robot ability to perform a task. We instantiate a number of (simulated) VSR bodies, differing in shape and actuation mechanism, and let a panel of humans control them, by means of online interaction with the simulator, while performing the task of locomotion. We use the same bodies with controllers obtained with evolutionary optimization, for the same task. We compare the ranking of human- and optimized AI-based VSRs, finding them very similar. We believe that our results strengthen the hypothesis that intrinsic factors in the body of modular robots determine their success.

INDEX TERMS Control, voxel-based soft robots, evolutionary computation.

I. INTRODUCTION

Reconfigurability plays a key role in the road towards robots that are more adaptable in terms of task [1] and environment [2]. In this regard, modular robots [3] are a relevant family of robots. Modularity allows rearranging the composing modules to attain a wide range of physical configurations. With this freedom, modular robots have been capable of performing a wide array of tasks [4]. However, this freedom, albeit an advantage, makes the joint optimization of the morphology together with the controller a challenging problem [5]. As implied by the embodied cognition paradigm, which posits a deep entanglement between the brain, the body, and the environment of an agent [6], not every morphology is suitable for every controller; controllers must adapt not only to the task but also the morphology they are embodied within [7]. Additionally, the joint optimization of morphology and control entails a larger search space than when optimizing one of the two alone. Some works tackle this issue by embedding an inner optimization loop (e.g., for the controller) inside an outer loop (e.g., for the morphology) [8], [9], to the cost of higher computational requirements [4]. Shrinking the search space towards the most promising dimensions would simplify and speed up the optimization process.

Whereas the literature has delved into what factors make a physical configuration more or less successful [9], [10], [11], no work to date has considered how to augment the robot optimization with knowledge from external sources, despite it being common in other fields [12], like test [13] and content

The associate editor coordinating the review of this manuscript and approving it for publication was Wai-Keung Fung.

generation [14]. Just like most animal lifeforms, humans evolved powerful mental "world models" for predicting the future states of the world [15]. Here we consider humans as sources of external knowledge and ask if they can be predictive for the performance of a robot controller optimized for a given body and task.

Involving humans in prediction (a human-in-the-loop approach) may not be convenient for simple robots and environments, since it incurs a cost in terms of time and effort to implement human-machine interfaces. Yet, such a cost dissipates when considering the optimization of robots and environments that are burdensome to evaluate. Current robotics benchmarks require several days of computing on specialized equipment to be fully evaluated and solved, resources that are not accessible to the majority of academic labs [16]. Indeed, there is a widespread sentiment that current artificial intelligence algorithms require too high a cost in computing, data, and money [17], and the control of robots is no less [18].

As robots, we consider Voxel-based Soft Robots (VSRs) [19], [20]. They allow for a diverse set of morphologies [21], which are assemblies of homogeneous modules of soft material. Voxels contract or expand their volume following an actuation signal dictated by the controller. It is the concerted change of voxel volumes that allows for the emergence of the high-level behavior of the robot. The softness of the modules induces rich dynamics, making the VSR body particularly relevant for the robot ability to perform a task [22], [23]. Moreover, softness confers VSRs the desirable property of compliance, which could even enable their usage in direct human-robot interaction [24].

We instantiate simulated VSRs differing in shape and actuation mechanism, for a task of locomotion, and perform two different sets of experiments. In the first set, human participants control the VSRs by interacting online with the simulator. In the second set, we optimize AI-based controllers inspired by Talamini et al. [25] through Evolutionary Computation (EC) [26], for the same VSRs and task. Finally, we compare the rankings of human- and AI-based VSRs.

We find that, even though AI-based controllers perform better than humans at controlling the VSRs, the relative rankings are similar. This result suggests that humans can be predictive of the performance of a robot controller optimized for a given body and task. Looking forward, we believe our work to be a stepping stone in the road toward more adaptable modular robots: in the future, we envision surrogate models [27] built on top of human knowledge to predict what re-configurations are better for different tasks and environments.

II. RELATED WORKS

In this work, we explore human-robot interaction as a source of external knowledge to reduce the optimization space of the possible morphologies. The availability of reliable predictions on the performance of autonomous robot systems can help developers reduce the time needed to optimize the design of robots [28].

Previous papers already explored human-soft-robot interaction. Some of these works consider the exploitation of a joystick to control reconfigurable soft robots [29]. When possible, interactive interfaces similar to the robot shape are exploited. An example is given by [30], where a Rigid Link Robot Controller is exploited. Others exploit the resemblances with the human body, for example, in [31] a soft manipulator is controlled through the usage of a wearable glove.

More in general, human-robot interaction has been a subject of research since some time, albeit usually with different objectives than our study. Today autonomous systems might not reach the required performances or human supervision is required together with the possibility of effectively operating in emergency actions. Remote control continues to be relevant in applications such as operating in hazardous environments [32] and medical applications [33], [34]. Moreover, being humans good at performing several tasks, they can be also exploited to teach robots how to perform several tasks. This practice is studied in several works and takes the name of "robot learning from demonstration" [35], [36]. In this approach, robots leverage humans' expertise in specific tasks to acquire behaviors through the analysis of human demonstration data [37], [38] or by directly imitating the human motion [39]. In principle, a similar approach may be used for making robots learn from other robots [40].

However, few works exploit human interaction to make robot optimization more efficient and effective. One of them is by Matthews and Bongard [41], who proposed to use non-expert humans for "grading" robots being optimized for performing the actions described by humans themselves in natural language. Focusing on the same family of robots, Kafley [42] proposed a crowd-sourcing platform to allow the collective creation of soft robots, with the ultimate goal of scalability.

In other fields, such as text generation [12], computer vision [14], [43], natural language processing [44], and others [45], [46], [47], enhancing AI using external knowledge is quite common. Moreover, the usage of human interaction or human feedback is commonly included in the learning or optimization phase of many models in several AI and robotic fields [48]. Humans have also been successfully included in the training loop of AI system synthesis [49], [50], extensively exploited in Computation and Language AI systems such as [51] and [52], and their feedback has also been exploited in text-to-image systems [53]. A known example of how humans have been exploited specifically in the evolutionary optimization loop is given by Picbreeder [54], where humans and AI collaboratively evolve images. If humans turn out to be reliable predictors of the performance of autonomous robotic systems, systems similar to Picbreeder could be used to collaboratively evolve robot morphologies.

Within the Artificial Life community, the notion of *hybrid life* has recently been put forward, covering, among other aspects, hybrid architectures, which combine living and artificial systems, and/or the hybrid interactions among them [55]. Among the former ones, motor augmentation is one of the most relevant research areas, which is centered around the realization of cyborgs, integrating biological and machine parts. Although clearly not involving humans, this topic has various links to our study: an interesting example comes from the works of Kuwana et al. [56], [57], who were able to augment a mobile robot by implanting insect antennae on it as a pheromone sensor. Concerning hybrid interactions, instead, one of the central goals is the understanding of collective behaviors in biological systems by introducing virtual agents among them. Examples include virtual plankton [58], robot cockroaches [59], and many others. Interestingly, Shirado and Christakis [60] have even experimented with hybrid interactions with humans, including bots as players in a coordination game.

III. VOXEL-BASED SOFT ROBOTS

In this study, we consider Voxel-based Soft Robots (VSRs), a kind of modular soft robots composed of several deformable cubes called *voxels*. Each voxel contracts or expands its volume, allowing for the emergence of the high-level behavior of the robot. Modular soft robots have been first introduced by [19].

In particular, we focus on a simulated (in discrete time, with a simulation frequency $f_{\text{sim}} = 60\,\text{Hz}$, and continuous space) 2D version of VSRs [20]. VSRs, due to their inherent nature, present various computational and control complexities. The softness of the modules induces rich dynamics, making the body particularly relevant for the robot ability to perform a task. Softness is also relevant for several applications where human safety is of fundamental importance such as in the medical field. Controlling these robots is difficult for humans, and therefore, utilizing such a 2D model may be advantageous for this study, as it provides an opportunity to conduct an initial investigation and, hopefully, pave the way for future studies exploring different objectives beyond those presented here. Indeed, working with fewer dimensions allows us to consider simpler interfaces for the interaction, both in terms of physical interface and video shown to the user. This opens the possibility of involving a diverse pool of individuals, including those not used to playing video games and interacting with systems such as simulators. Indeed, while one dimension makes these simulated VSRs less realistic, it simplifies the human-robot interaction.

In the following, we describe the characteristics of VSRs relevant to this study. A VSR is defined by its *body* and its *brain*, which we detail in the following subsections.

A. BODY

The body of a VSR consists of its structure and physical properties together with the sensors on each voxel. The VSR structure describes how the voxels are arranged in a 2D grid topology—the *shape*—of size $w \times h$ where w is the width and h is the height of the grid. We model each voxel as the assembly of spring-damper systems, masses, and distance constraints [20], and as being rigidly connected to its four adjacent voxels, when present.

Over time, voxels change their area, according to (a) external forces acting on the voxel (e.g., other bodies, including other voxels and the ground) and (b) internal forces computed by the controller. Contraction and expansion forces are modeled in the simulation as an instantaneous change in the resting length of the spring-damper systems of the voxel. The length change is linearly dependent on a control value residing in $[-1, 1]$, -1 being the greatest possible expansion and 1 being the greatest possible contraction.

Voxels can be equipped with zero or more sensors. Each sensor produces, at each time step, a sensor reading $\boldsymbol{s} \in \mathbb{R}^m$. In this work, we use four kinds of sensors. Touch sensors ($m = 1$) perceive whether the voxel is touching the ground ($\boldsymbol{s} = (1)$) or not ($\boldsymbol{s} = (0)$). Velocity sensors ($m = 2$) perceive the velocity of the center of mass of the voxel along the x- and y- axes of the voxel itself. Area sensors ($m = 1$) perceive the ratio between the current area of the voxel and its rest area. Sight sensors ($m = 1$) allow the voxel to perceive its distance from the ground, by computing the Euclidean distance between the center of the voxel and the ground. We put area sensors in each voxel, velocity sensors in the voxels of the top row of the body, and sight sensors in the voxels of the bottom row. To ensure all sensor readings lie in $[-1, 1]^m$, we apply the tanh function to every reading and then rescale it. After the normalization, to simulate real-world sensor noise, we perturb every sensor reading with additive Gaussian noise of mean $\mathbf{0}$ and deviation $\sigma_{\text{noise}} = 0.01$.

B. ACTUATION

VSRs behavior results from changing the area of the composing voxels over time: this depends directly on the control values. How controllers influence the control values depends on *actuation mechanisms*. Actuation mechanisms map a discrete action space A to a way to set the individual voxel control values after an action $a \in A$ is taken. Since controlling each voxel individually would be impractical for humans (as there could be tens of voxels in a VSR), actuation mechanisms are defined to not involve direct control of each individual voxel.

We consider four different actuation mechanisms, three of which are based on a partitioning of the shape voxels, and one, that we used only on the worm-shaped robot (see Section IV-A), induces a wave-like movement.

Partitioning-based actuation relies on partitions, i.e., disjoint sets of voxels that actuate altogether. There is one action a for each partition: whenever a is taken, all and only the voxels of the corresponding partition immediately contract for a short time (one time step). We use the following partitionings: $\leftrightarrow$ actuation, where the robot is divided into $|A| = 2$ partitions along the y-axes; $\updownarrow$ actuation, where the robot is divided into $|A| = 2$ partitions along the x-axes; $\boxtimes$ actuation, where the robot is divided into $|A| = 4$ partitions along both the x- and y-axes.

Wave-based actuation mechanism is used only with the worm shape and we will refer to this as the $\rightsquigarrow$ actuation; for $\rightsquigarrow$, $|A| = 2$. Intuitively, the voxels are actuated in the following way: whenever an action a is taken, a contraction wave starts at the left or right side of the robot (depending on the action), and goes on contracting subsequent columns while the previous expands again. This gives rise to a general contraction-expansion movement that resembles the movement of a wave. We made the wave last 1 s.

C. CONTROLLER

The controller itself together with the physics of the environment determines the whole movement of the VSRs. The controller generates the actions $a \in A$ that, depending on the actuation mechanism, make the different voxels contract or expand.

In this paper, we considered two types of controllers: AI-based controllers and human-based controllers.

1) AI-BASED CONTROLLER

The AI-based controller is inspired by one of [25] and consists of a Multi-Layer Perceptron (MLP)—other kinds of neural networks could be used, yet we opted for MLPs for their simplicity [61]. The MLP takes as inputs the concatenation of the sensor readings from all the voxels, $\mathbf{s} \in [-1, 1]^q$, with q the number of sensor readings, and outputs $|A|$ values: the taken action a is the one for each there is the largest MLP output. We use one hidden layer with $0.65\,q$ nodes and tanh as an activation function.

This controller may take one action at each time step, i.e., one every $\frac{1}{60}$ s, with a control frequency $f_c = 60\,\text{Hz}$. Human-based controllers would hardly be able to keep this pace. Hence, we experimented also with a slower version of the AI-based controller which takes an action every $\frac{1}{5}$ s, i.e., with $f_c = 5\,\text{Hz}$. We opted for such a frequency after conducting a few exploratory experiments. Indeed, during our testing, we found that $f_c = 5\,\text{Hz}$ was the lowest frequency that did not interfere with human control. We observed that smaller values were uncomfortable for humans, making it feel as though the simulator was lagging behind the given commands. As a result, we concluded that a frequency of $f_c = 5\,\text{Hz}$ closely approximates the rate at which a person operates on the simulator.

2) HUMAN-BASED CONTROLLER

The human-based controller consists of a human who interacts with the simulator using a keyboard. Taking inspiration from other works, we performed preliminary experiments with other devices than the keyboard, such as gamepads, however, we found no advantages.

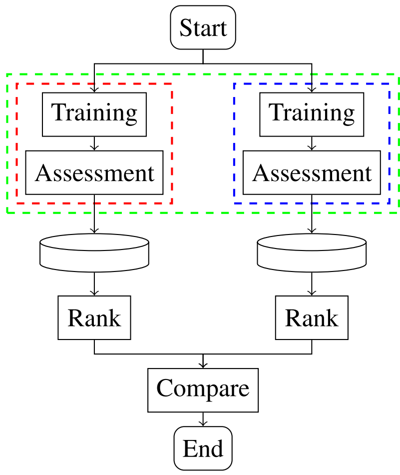

FIGURE 1. A schematic of the workflow of our experimental procedure. Rectangular boxes represent phases of the experiment. The dashed enclosing boxes represent loops: the green one - - iterates over shapes and actuation mechanisms, the red one - - over random seed for the ES AI-based optimizer, the blue one - - over human subjects.

During the tests, participants view a real-time video feed of the robot that they need to control. Using this video as a reference, individuals can press specific keyboard buttons to trigger the expansion of different voxels, thereby enabling the robot to perform the intended task. Each action $a \in A$, is associated with a particular keyboard key: whenever the human presses a key, the corresponding set of voxels is expanded, while, when he releases the key, the corresponding set of voxels return to their rest area. The program continuously monitors key pressure with a high frequency, ensuring that expansions occur immediately upon key presses. Furthermore, users have the option to press multiple keys simultaneously, allowing voxels from different partitions to expand together.

To determine the key configuration, i.e., which keys to associate with A actions, for each actuation mechanism we conducted multiple user tests. Our objective was to obtain the most intuitive configuration for the users. Since the actuation mechanisms require at most 4 inputs, the keys in the configurations are always selected from an original set of 4 keys. Therefore, the actuation mechanism with 4 possible actions uses the entire set, while the other configurations involve selecting keys to associate and excluding others. The entire set is composed of $\boxed{\text{W}}$, $\boxed{\text{A}}$, $\boxed{\text{S}}$, and $\boxed{\text{D}}$. This set of keys is common and widely used in modern PC video games. $\boxed{\text{A}}$ and $\boxed{\text{D}}$ correspond to the left and right partitions of $\leftrightarrow$ and $\boxtimes$ actuations. $\boxed{\text{W}}$ and $\boxed{\text{S}}$ correspond to the top and bottom partitions of $\updownarrow$ and $\boxtimes$ actuations. In $\rightsquigarrow$, $\boxed{\text{A}}$ and $\boxed{\text{D}}$ correspond to the wave starting from left and right. We summarize these key-bindings in Table 1.

IV. EXPERIMENTS AND DISCUSSION

As a goal of this work, we set out to evaluate whether humans could be predictive of the performance achieved by optimized

TABLE 1. Key-bindings for the four actuation mechanisms (↔, ⊠, ↕, ⇝). For the wave-based actuation, the binding shows in which direction the contraction wave flows. For the partition-based actuation mechanisms, the binding shows which voxels (in terms of colors, see Figure 2) contract.

Configuration	W	A	S	D
↔		■		■
⊠	■	■	■	■
↕	■		■	
↕	■		■	
⇝	■ → □		■ ← □	

controllers for VSRs. More in detail, we considered the following two research questions:

RQ1 Is human-based control predictive of optimized AI-based control for VSRs?

RQ2 Do humans need to try to actually control the VSRs to express a prediction? In other words, do people have some sort of innate knowledge with respect to which VSR can perform better upon optimization?

To address these questions we performed a thorough experimental evaluation, involving several VSRs of different shapes and actuation mechanisms. Namely, we first optimized the AI-based controller for all the VSRs and compared their performance with that attained by humans in the same setting. Then, we let another pool of people predict the performance of the very same agents by simple static observation and compared their forecast with the actual performance measured upon optimization. For both phases, we considered the task of locomotion as it is both widely explored within the robot optimization community and is intuitive enough for people to carry out without specific training required. Figure 1 summarizes the workflow of our experiments, which we describe in detail in the next sections.

A. RQ1: IS HUMAN CONTROL PREDICTIVE OF OPTIMIZED CONTROL PERFORMANCE?

To tackle the first research question, we performed a two-fold analysis: first, we relied on optimization, and then we let participants control the same VSRs.

We considered two VSR shapes, a biped of size 12×5 (by size we mean the size of its enclosing grid) and a worm of size 16×4, in combination with the aforementioned actuation mechanisms. We report the chosen VSRs in Figure 2, where we highlight with different colors the sections of the body that can independently be controlled with the given actuation mechanisms (with the exception of the ⇝ actuation, for which we use a gradient to give the impression of a "wave"); we summarize in Table 1 the key-bindings for the chosen VSRs, i.e., the actuation-key associations. We selected these VSR shapes as they have already been successfully employed quite often for the considered task [62], [63], although with a smaller scale. Here we opted for a larger scale to ease the control for humans, and to give the impression of faster movement, hence minimizing the frustration for the participants. As a side note, these two shapes resemble animals, making them more fun for people to control. However, this might induce bias as people might attempt to replicate the gait associated with these animals even if it might be not the most effective one for the two shapes.

To test and compare both control methods—humans and optimized controllers—we considered the task of *locomotion*, where the goal is to move the VSR as fast as possible along the positive x direction. Hence, we used the mean velocity of the VSR, v_x, to quantify the degree of accomplishment of said task, i.e., to measure the performance of a controller. To ease the task and arouse more engagement in the participants, we employed a smooth and slightly sloping terrain, tilted by 10 degrees.

1) AI-BASED CONTROL

We started our experimental evaluation by considering the AI-based controller for each of the VSRs of Figure 2. Namely, we used an MLP, as explained above, and we optimized its weights $\boldsymbol{\theta}$ by the means of EC for the task of locomotion. We experimented with two control frequencies, a *fast* one of 60 Hz, and a *slow* one of 5 Hz, which should be more akin to the timing of humans and should also prevent vibrating behaviors, as noted in [64].

For the optimization, we resorted to a simple form of non-overlapping Evolutionary Strategy (ES) for the optimization of the weights $\boldsymbol{\theta}$, as it has already been applied in similar settings [65], [66] and has also achieved remarkable results for continuous control tasks and game-playing [67]. The considered ES works by evolving a fixed-size population of the size of $n_{\text{pop}} = 36$ individuals for a total of $n_{\text{gen}} = 285$ generations (totaling 10 000 fitness evaluations). To initialize the population, we generate each individual by sampling a p-dimensional uniform distribution over $[-1, 1]^p$, where $p = |\boldsymbol{\theta}|$. Then, at each iteration, we obtain a new population by merging the offspring of size $n_{\text{pop}} - 1$ with the current fittest individual. To compute each element of the offspring, we take the element-wise mean $\boldsymbol{\mu}$ of the fittest quarter of the population, and we add a Gaussian noise $\mathcal{N}(0, \sigma)$, with $\sigma = 0.35$, to each of its components. Given the stochasticity of the ES execution, we performed 10 independent optimizations for each controller, each based on a different random seed. Clearly, given the locomotion task, we considered v_x as a fitness measure to drive the optimization of $\boldsymbol{\theta}$, which we measured in a simulation of 30 s.

We display the progression of the velocity during the optimization process for all the considered VSRs in Figure 3: for each VSR configuration, i.e., for each combination of shape, actuation mechanism, and control frequency f_c, we report the median and the inter-quartile range (across the 10 runs) of the velocity v_x achieved by the best individual in the population at each generation.

From the figure, we can make several observations. First, we see all lines consistently increasing until they plateau, which confirms that optimization is indeed occurring and also suggests that prolonging it would result in negligible

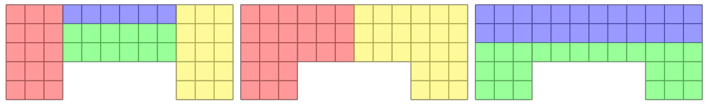

(a) Biped (B⊠, B↔, B↕)

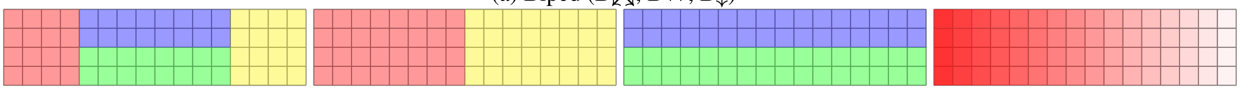

(b) Worm (W⊠, W↔, W↕, W⇝)

FIGURE 2. **Combinations of shape and actuation mechanisms.**

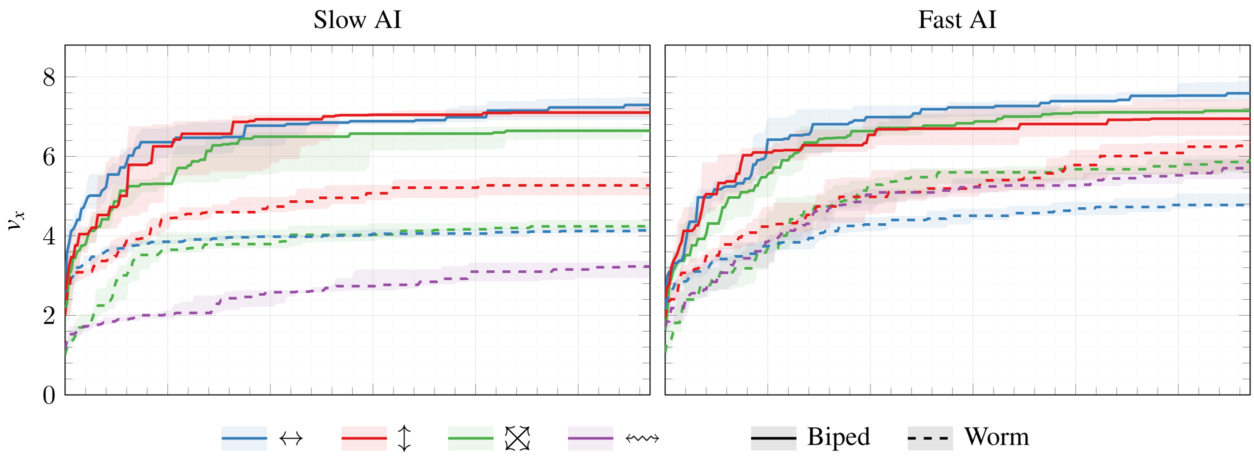

FIGURE 3. **Fitness v_x of the best controller during evolution. Optimization works for all shapes and actuation mechanisms, but some combinations achieve higher v_x.**

improvements in almost all cases. What is more interesting to note, is that the lines display some variability, not only in the result achieved at the end of evolution but also in how they reach it, i.e., in how fast they are optimized. This means that the considered configurations differ both in how "optimizable" they are and in how "easily optimizable" they are, and serves as a starting point for further reasoning also on how the participants will perform.

Delving more into detail about the observed differences, the most striking one is that between the two VSR shapes, i.e., between the two columns of the figure. Such difference is coherent across all actuation mechanisms and also for both the slow and fast control frequencies: the biped appears consistently faster and easier to optimize than the worm. In fact, from a brief examination of the videos of the VSRs performing locomotion, we see that the biped, thanks to its shape, can achieve a horse-like gait, which makes it more effective than the worm, which usually achieves motion through crawling.

Another element of difference among the VSRs is the actuation mechanism. Here we also expect to notice some diversity, as the mechanisms differ in expressivity, i.e., how wide of a behavior gamut they can generate, but also in search space size, i.e., the size of p. Clearly, on one side we have ⊠ with more freedom and larger p, whereas on the other side, we have all other actuation mechanisms. However, in the plots, we do not highlight much disparity related to these two aspects. Instead, we notice that the main role is played by the shape-actuation combination, although oftentimes some combinations become more or less effective with the different control frequencies employed. Indeed, B↔ and W↕ are always the best-performing combinations, with both fast and slow AI, while we notice some mismatch in the ordering of the others.

Hence, another point of necessary reflection lies in the control frequency employed. As already observed in other studies, when working at full control frequency, optimization might end up finding behaviors that are unintuitive and would not be practicable in the real-world—the so called *reality gap* problem [68], [69].

2) HUMAN-BASED CONTROL

Having observed these results upon optimization, that is, having confirmed that all the considered shape-actuation combinations can indeed perform locomotion, we moved on to the user study with humans as controllers for the VSRs. Namely, we involved 48 unpaid volunteer participants, 13 females and 35 males, of various ages, ranging from 18 to 60. Before starting the actual experiment, we surveyed the participants on their video-games playing habits, as these could have played some role in the final outcomes achieved. We observed different habits, both with respect to the frequency of playing and with respect to the types of games played, yet in the end we saw no particular dependency of the results from these factors.

After the survey, each participant started the actual experiment with a randomly chosen shape among the biped and the worm and then proceeded with the other one. We randomized the order of shapes to balance possible advantages of getting more accustomed to the environment or disadvantages of potential frustration and/or boredom. Concerning the actuation mechanism, all participants started with ⊠ and then proceeded with all the others. To help participants get used to the control of the VSR, we first let them interact with the simulator for 60 s prior to recording their performance. We allowed non-recorded interaction time in between every session, yet we decreased it to 30 s when we only changed the actuation and not the shape, to prevent annoyance in the participants. Regarding the actual experimental session, we let the participants control the VSR for 33 s, from which we discarded the initial 3 s to compute the mean velocity v_x. We chose to discard the initial transient to let participants properly focus and get settled before measuring their performance.

We report the outcomes of the user study, together with those obtained upon optimization in Figure 4. Namely, for each VSR, we display the distribution of the mean velocity v_x, achieved either during the user study or at the end of the evolutionary optimization. To order the x-axis in the plots, we follow the ranking induced by the median of the v_x achieved by humans for each VSR, in decreasing order.

To answer to RQ1, i.e., to assess whether human control is predictive of optimized control, we need to examine the three rows of plots in a comparative manner. Clearly, we do not aim at comparing the absolute values of v_x, as there is a myriad of factors making the two outcomes non-directly comparable. What we are indeed interested in analyzing, is whether the rankings among VSRs induced by human control are similar—or even equal—to those induced by optimization.

First, we notice that as in Figure 3, there is a rather clear predominance of the biped with respect to the worm, regardless of the actuation mechanisms employed. This hints that humans are strong predictors of the VSR shape. In fact, as in the optimized setting, humans would intuitively try to achieve a horse-like motion for the biped, which made it more effective than the worm.

Considering also the actuation mechanisms, we notice that the human ranking is the same as that attained with the optimized fast controller for the biped, whereas for the worm it matches that of the optimized slow controller. By performing the experiments, for different actuation mechanisms, we noticed that the frequency at which controllers are able to input the simulator might influence the final sequence found, and therefore the performance obtained and the final ranking. This result reminds us of the previously highlighted dependency of the performance on the control frequency, which derives from the deep entanglement of the controller

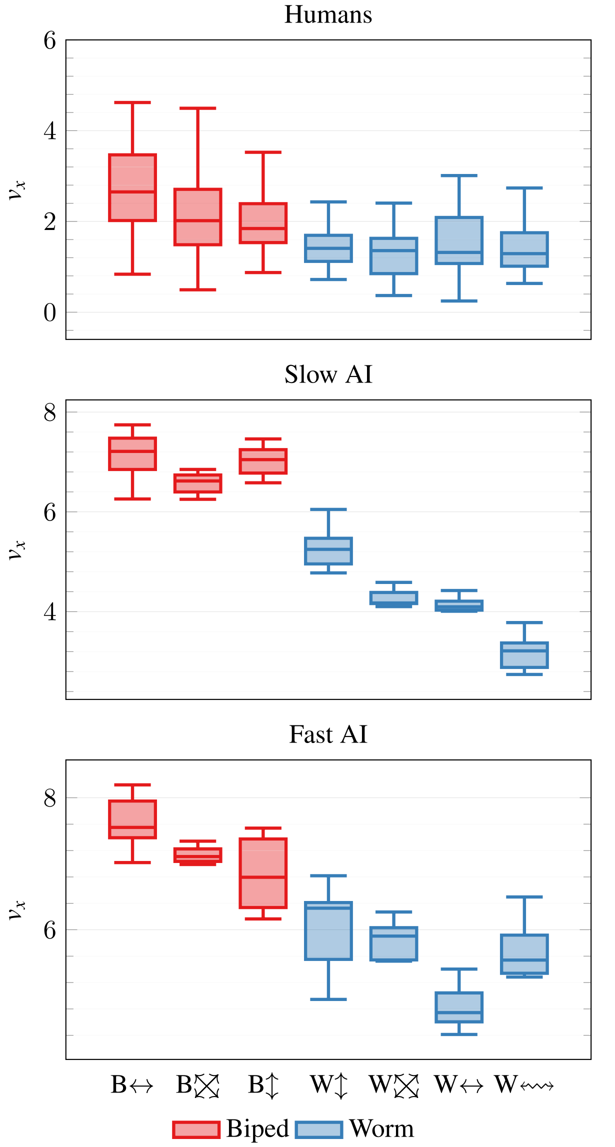

FIGURE 4. Distribution of the v_x achieved by humans, slow AI, and fast AI. The order in the x-axes is chosen to easily compare the resulting rankings. Shape and actuation mechanism combinations rank very similarly for humans and AI.

with the body dynamics. Hence, such intertwining makes it hard to achieve full human predictivity for every actuation mechanism considered.

Summarizing, from a high-level perspective, Figure 4 gives a positive answer to RQ1, especially for what concerns the VSR shape. Anyway, it is necessary to note that the distribution of velocities achieved by humans is fairly spread. This lack of consistency can descend from several aspects, such as the different starting points for the user study, the different individual capabilities, but also the different individual attitudes toward the simulator. Thus, our findings surely hold on average for an aggregation of the pool of participants considered, yet it is not clear whether any conclusion can be drawn at the single-participant level.

To further analyze this aspect, we considered the performance obtained by individual participants, rather than aggregation. Namely, we assessed the individual global rankings together with the individual shape rankings. For the former, we employed Spearman's Footrule [70], which is a measure of disarray for rankings, and computed the pair-wise comparisons among participants. Although we do not report the results here, we observed fairly large values, which indicate a general disagreement on the rankings obtained. Hence, we can draw few conclusions at the individual participant level.

For the shapes ranking, we computed, for each participant, the mean velocity for each shape averaged across actuation mechanisms, $\bar{v}_{x,\mathrm{B}}$ for the biped and $\bar{v}_{x,\mathrm{W}}$ for the worm. For comparing the results with those of optimized controllers, we repeated the same computation for the optimized controllers, where we considered each optimization seed as an "individual participant". We report the aforementioned values in the scatter plot of Figure 5 (left), where each point corresponds to a participant (or an optimization seed), and the axes are the $\bar{v}_{x,\mathrm{B}}$ and $\bar{v}_{x,\mathrm{W}}$, respectively. We also report the bisector to ease the comparison: the points lying on the lower right side of it rank the biped better than the worm. Observing all the AI-based controllers, both slow and fast, we note that they clearly lie on the lower right side, as expected from both Figures 3 and 4. For what concerns the human participants, we note that they are closer to the bisector and are more spread around, yet the vast majority (45 out of 48) achieved larger $\bar{v}_{x,\mathrm{B}}$ than $\bar{v}_{x,\mathrm{W}}$. This corroborates our previous finding, that humans have strong predictive capability with respect to the VSR shape.

3) IS WHAT WE FOUND IN RQ1 LIMITED TO THE TWO CONSIDERED SHAPES?

Given the results achieved in RQ1, an additional question arises on whether the findings are limited to the two considered shapes. In particular, we wonder if these results descend from the fact these shapes were chosen by humans—namely, us, as authors. To address this, and thus provide a more general answer to RQ1, we repeated the same experimental procedure for 4 additional shapes, reported in Figure 6. Namely, we let optimization "pick" the first two shapes, i.e., we considered two shapes optimized by EC for the locomotion task [71], [72], denoted by EC1 and EC2, while we selected by hand the other two shapes, a cross (X) and a table-like shape (T). As before, we experimented with different actuation mechanisms, $\boxtimes$, $\leftrightarrow$, and $\updownarrow$, in combination with each shape. We show all the combinations for this supplementary analysis in Figure 6.

To scale down the effort required, we reduced the participant sample to the 5 best performers, and we only experimented with fast AI, which proved more effective in the previous experiments. We report the distribution of velocities v_x in Figure 7, where again we order the x-axis according to the human ranking in decreasing order. We also summarize in Table 2 the mean v_x achieved by humans and fast AI with all the VSRs.

TABLE 2. Average velocity v_x achieved by humans and fast AI, for all combinations of shape and actuation mechanisms. Δr represents the difference between the ranking obtained in each combination between humans and fast AI. The rank r of most combinations is similar for both controllers: $\Delta r \leq 2$ for 13 on 19 combinations.

	Human (▼)		Fast AI (▲)			
Name	v_x	r	v_x	r	$\lvert\Delta r\rvert$	
B$\boxtimes$	2.02	8	7.15	6	2	▲
B$\leftrightarrow$	3.19	4	7.59	4	0	=
B$\updownarrow$	2.42	7	6.95	7	0	=
W$\boxtimes$	1.85	10	5.93	10	0	=
W$\leftrightarrow$	1.32	14	4.78	13	1	▲
W$\updownarrow$	1.71	13	6.34	9	4	▲
W$\leftrightsquigarrow$	1.76	12	5.70	11	1	▲
X$\boxtimes$	5.97	1	8.29	3	2	▼
X$\leftrightarrow$	5.64	2	9.16	2	0	=
X$\updownarrow$	5.39	3	9.32	1	2	▲
T$\boxtimes$	3.10	5	6.53	8	4	▼
T$\leftrightarrow$	1.77	11	3.71	16	5	▼
T$\updownarrow$	1.06	16	7.49	5	11	▲
EC1$\boxtimes$	1.09	15	3.82	15	0	=
EC1$\leftrightarrow$	1.99	9	5.62	12	3	▼
EC1$\updownarrow$	2.51	6	3.91	14	8	▼
EC2$\boxtimes$	0.13	18	0.80	19	1	▼
EC2$\leftrightarrow$	0.76	17	1.93	17	0	=
EC2$\updownarrow$	0.13	19	1.29	18	1	▲

Observing the two groups of box plots in a comparative manner we can note some similarities with respect to Figure 4. In particular, we observe some predictive power for the shapes—X always ranks highest, EC2 always ranks lowest—while the rankings for the actuation mechanisms are mostly different. Hence, we can confirm the already noted difficulty in predicting the behavior of different actuation mechanisms.

Regarding the shapes, we delve into a deeper analysis considering the mean velocity for each participant (or optimization seed) averaged across actuation mechanisms and evaluating the ranks induced by them. In this case, as suggested by Figure 7, there is indeed predictivity. More in detail, X is always the fastest, whereas EC2 is always the slowest, for both human control and optimized control. Concerning the two remaining shapes, T and EC1, we observe some mismatch in the relative rankings, worth a deeper investigation. Hence, to draw further conclusions, we rely on the scatter plot visualization of Figure 5 (right), where each point corresponds to an individual (or an optimization seed) and the axes measure $\bar{v}_{x,\mathrm{T}}$ and $\bar{v}_{x,\mathrm{EC1}}$, respectively. Namely, we note that, as before, the optimized results are extremely coherent, and lie all on the same side of the plane (i.e., $\bar{v}_{x,\mathrm{T}} > \bar{v}_{x,\mathrm{EC1}}$), whereas the human ones are more scattered, close but on both sides of the bisector. Therefore, in this case, the human predictive power proves weaker, also for the shapes only. We speculate this derives from the fact that the new shapes considered are less intuitive for humans to control, as there is less resemblance to real-world creatures.

Summarizing, for what concerns RQ1, we can conclude that (1) for "familiar" shapes individual humans are strong

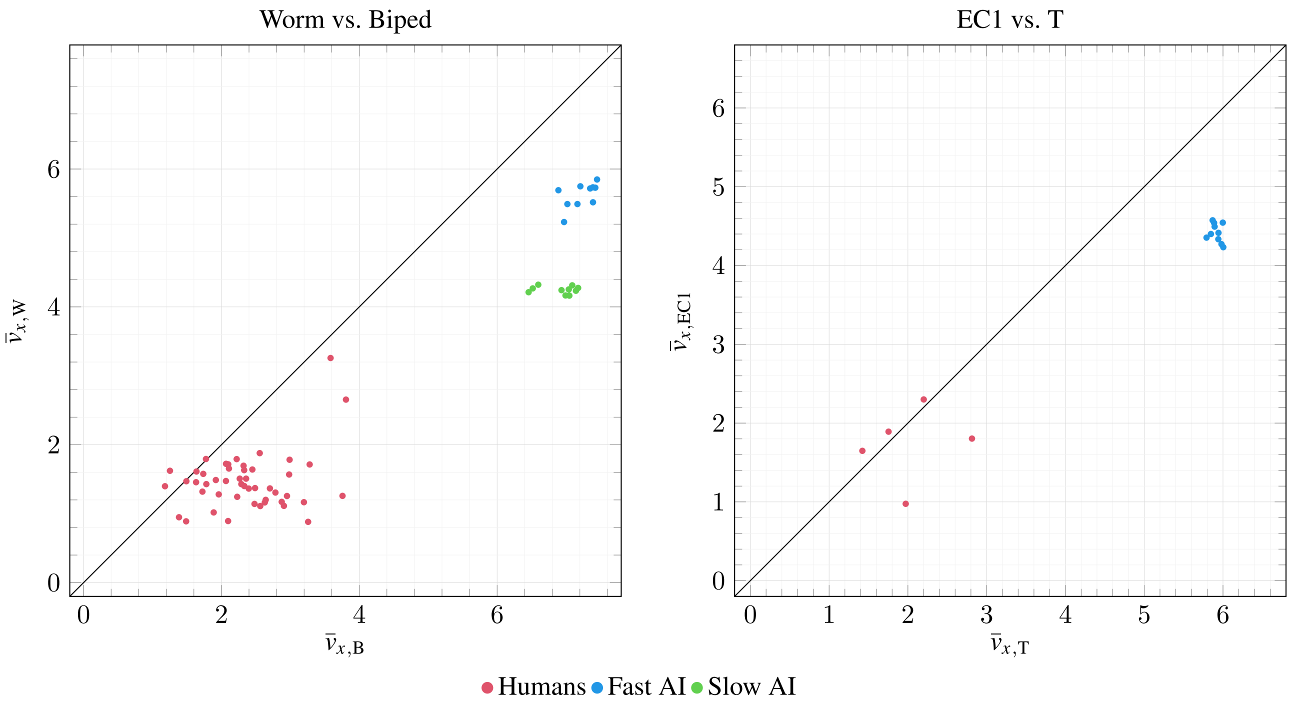

FIGURE 5. Scatterplot of the v_x achieved with the biped and the worm, on the left, and T and EC1, on the right, one point per participant or optimization seed. The biped is faster than the worm for both humans and the AI; with T and EC1, humans agree less.

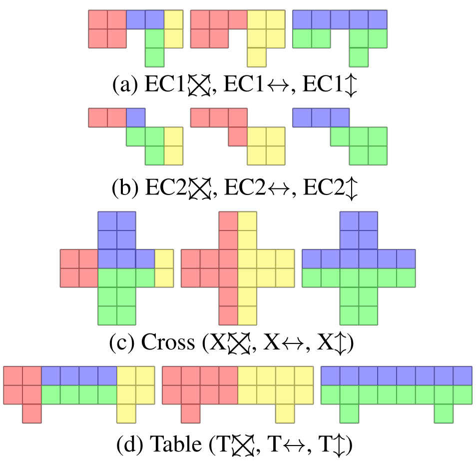

FIGURE 6. Combinations of shape and actuation mechanisms for the supplementary analysis.

predictors of the performance achieved upon optimization, and (2) such predictivity holds, though in a weaker way, even for less familiar shapes. Moreover, (3) we find some predictive power also for the actuation mechanisms when aggregating the results over a large pool of participants.

B. RQ2: DO HUMANS HAVE INNATE PREDICTIVE POWER?

To address the second research question, we relied on a survey where we asked people to rank some VSR shapes according to their expected velocities. In light of the previous results, this investigation becomes particularly relevant, as we noted

TABLE 3. Mean rank for the considered VSR shapes.

	B	T	X	EC1	EC2	W
Survey	2.84	3.29	3.35	3.65	3.45	4.42
Human control	2.00	4.00	1.00	4.00	6.00	4.00
Fast AI	2.00	3.00	1.00	5.00	6.00	4.00

some dependency of the predictivity on the familiarity level of the shapes.

Going more into detail about the evaluation performed, we involved a different pool of participants, disjoint from the previous experiment one, namely 31 university students of around 23 years of age on average. Before the actual survey, we provided a bit of context of VSRs to the participants, to enable them to make more informed choices. Namely, we briefly described VSRs, how they move, and the locomotion task and we also showed a video of a VSR of a different shape than those considered for the survey, performing locomotion in the same setting considered here. Then, we showed them the 6 shapes employed, i.e., those of Figures 2 and 6, and asked participants to rank them according to the expected velocity these VSRs would have when performing locomotion with a properly optimized controller. Differently from before, we did not consider different actuation mechanisms, as we observed low predictive power in that sense.

To assess whether the participants were coherent with the previous findings, we considered the mean ranking for each shape, reported in Table 3—we assigned score 1

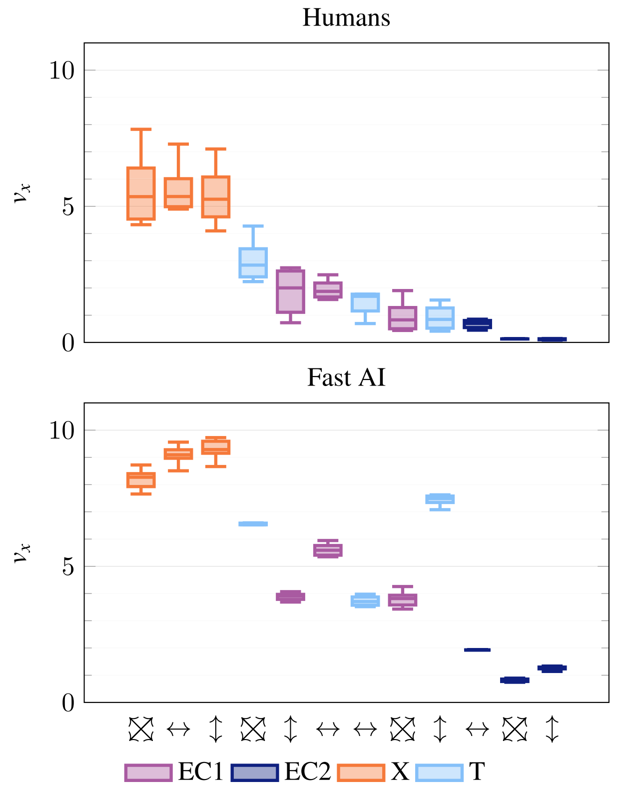

FIGURE 7. **Distribution of the v_x achieved by humans, slow AI, and fast AI. The order in the x-axes is chosen to easily compare the resulting rankings. Shapes rank similarly for humans and AI.**

to the fastest VSR and score 6 to the slowest in the ranking. To contextualize the results, we also report the mean ranking of human and optimized control considering the mean velocities for each participant (or optimization seed) for each shape (resulting from the evaluation done for RQ1). To compute these rankings we only considered the participants who experimented with all the shapes, and only the fast AI.

We start by examining the first row of the table, i.e., the expected performance. From that, we can notice that only two shapes stand out from the others, namely the biped and the worm, whereas for the others the mean ranking is approximately always the same, around 3.5. Conversely, the biped is on average expected to be the fastest, while the worm is the slowest. This relationship among the two shapes also holds at the single participant level: 75.8 % of participants expected the biped to perform better than the worm.

It is also interesting to look at the table in a comparative manner, evaluating the similarity between the rows. Namely, we notice that sometimes the human expectations are closer to the AI performance (which is, indeed, very consistent for all shapes) than the actual human performance (which is consistent only for three shapes). We believe this derives from the fact that people express their judgment by association. In particular, we see they rank the T right after the biped due to their similarity, which also makes them similarly effective upon optimization. Moreover, they express a comparable preference for EC2 and EC1, which are both small and not very familiar, and end up both performing poorly even with an optimized controller.

Another point worth mentioning is the difficulty of humans in predicting not immediately intuitive behaviors. That is the case for the X-shaped VSR, which was in general the fastest upon control, but not to people's expectations. In fact, such VSR achieved successful motion through rotation, yet without interacting with it people were not able to guess it.

To summarize and conclude, we cannot give a totally positive answer to RQ2. In fact, people's expectations, and hence their predictive power, appear to be strongly influenced by experience. In other words, we found people to be biased towards shapes they perceive as familiar: for those they can usually correctly predict the performance ranking, even without having to actually interact with the agents, whereas it becomes more difficult for more unusual shapes or for VSRs which achieve motion in unconventional ways. We believe this potential bias might limit the practical usefulness of humans' ability to predict AI performance, in particular when the shape of the robot is subjected to optimization, that is, when unfamiliar shapes may arise in the design process.

V. CONCLUSION AND FUTURE WORK

We investigating whether humans can be used as indicators of the performances achieved by AI-based controllers for different VSRs, a kind of modular soft robots. To this aim we performed several experiments: we first optimized and evaluated controllers in performing the locomotion task over VSRs differing for shapes and actuation mechanisms. Then, we let a panel of 48 humans control the very same VSRs to accomplish the very same task.

We found that VSR shapes rank very similarly, in terms of achievable velocity, when controlled by humans and AI; that is, humans are able, to some extent, to predict which shapes would be better than others to perform a task when controlled by AI-based controllers. This applies especially when considering shapes that resemble animals, such as the biped and the worm. On the other hand, when considering not only shapes but also actuation mechanisms, full predictivity is harder to achieve. Moreover, humans obtained results rather different from each other, which probably is a consequence of the different order in which the VSRs were proposed to humans.

We believe that our results strengthen the idea that there are intrinsic factors that make some shapes better than others to perform tasks and encourage further research in this direction [10].

While we experimented with a simple means of control, the keyboard, we think that new possibilities exist for future research where different physical interfaces are used for greater control granularity, such as the brain-human interface [73]. Moreover, our findings suggest that humans could also be employed for human-in-the-loop evolutionary robotics, similarly to other what is done in other domains [50], [51], [52], [74].

REFERENCES

[1] M. Nisser, L. Cheng, Y. Makaram, R. Suzuki, and S. Mueller, "ElectroVoxel: Electromagnetically actuated pivoting for scalable modular self-reconfigurable robots," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2022, pp. 4254–4260.

[2] D. Pathak, C. Lu, T. Darrell, P. Isola, and A. A. Efros, "Learning to control self-assembling morphologies: A study of generalization via modularity," 2019, *arXiv:1902.05546*.

[3] K. Gilpin and D. Rus, "Modular robot systems," *IEEE Robot. Autom. Mag.*, vol. 17, no. 3, pp. 38–55, Sep. 2010.

[4] J. Bhatia, H. Jackson, Y. Tian, J. Xu, and W. Matusik, "Evolution gym: A large-scale benchmark for evolving soft robots," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 2201–2214.

[5] N. Cheney, J. Bongard, V. Sunspiral, and H. Lipson, "On the difficulty of co-optimizing morphology and control in evolved virtual creatures," in *Proc. Artif. Life Conf.* Cambridge, MA, USA: MIT Press, 2016, pp. 226–233.

[6] R. Pfeifer and J. Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA, USA: MIT Press, 2006.

[7] A. E. Eiben and E. Hart, "If it evolves it needs to learn," in *Proc. Genetic Evol. Comput. Conf. Companion*, Jul. 2020, pp. 1383–1384.

[8] K. Miras, M. D. Carlo, S. Akhatou, and A. Eiben, "Evolving-controllers versus learning-controllers for morphologically evolvable robots," in *Proc. Int. Conf. Appl. Evol. Comput.* Cham, Switzerland: Springer, 2020, pp. 86–99.

[9] A. Gupta, S. Savarese, S. Ganguli, and L. Fei-Fei, "Embodied intelligence via learning and evolution," 2021, *arXiv:2102.02202*.

[10] J. Talamini, E. Medvet, and S. Nichele, "Criticality-driven evolution of adaptable morphologies of voxel-based soft-robots," *Frontiers Robot. AI*, vol. 8, 2021, Art. no. 673156.

[11] N. Milan and S. Nolfi, "Phenotypic complexity and evolvability in evolving robots," *Frontiers Robot. AI*, vol. 9, Oct. 2022.

[12] W. Yu, C. Zhu, Z. Li, Z. Hu, Q. Wang, H. Ji, and M. Jiang, "A survey of knowledge-enhanced text generation," *ACM Comput. Surv.*, vol. 54, no. 11s, pp. 1–38, Jan. 2022.

[13] A. J. Offutt, Z. Jin, and J. Pan, "The dynamic domain reduction procedure for test data generation," *Softw., Pract. Exper.*, vol. 29, no. 2, pp. 167–193, Feb. 1999.

[14] B. Wu, C. Xu, X. Dai, A. Wan, P. Zhang, Z. Yan, M. Tomizuka, J. Gonzalez, K. Keutzer, and P. Vajda, "Visual transformers: Token-based image representation and processing for computer vision," 2020, *arXiv:2006.03677*.

[15] Y. LeCun, "A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27," *Open Rev.*, vol. 62, no. 1, 2022.

[16] J. S. O. Ceron and P. S. Castro, "Revisiting rainbow: Promoting more insightful and inclusive deep reinforcement learning research," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 1373–1383.

[17] N. Thompson, K. Greenewald, K. Lee, and G. F. Manso, "The computational limits of deep learning," in *Proc. 9th Comput. Limits*, 2023.

[18] P. Thodoroff, W. Li, and N. D. Lawrence, "Benchmarking real-time reinforcement learning," in *Proc. NeurIPS Workshop Pre-Registration Mach. Learn.*, 2022, pp. 26–41.

[19] J. Hiller and H. Lipson, "Automatic design and manufacture of soft robots," *IEEE Trans. Robot.*, vol. 28, no. 2, pp. 457–466, Apr. 2012.

[20] E. Medvet, A. Bartoli, A. De Lorenzo, and S. Seriani, "2D-VSR-Sim: A simulation tool for the optimization of 2-D voxel-based soft robots," *SoftwareX*, vol. 12, Jul. 2020, Art. no. 100573.

[21] F. Pigozzi, E. Medvet, A. Bartoli, and M. Rochelli, "Factors impacting diversity and effectiveness of evolved modular robots," *ACM Trans. Evol. Learn. Optim.*, vol. 2023, pp. 1–13, Mar. 2023, doi: 10.1145/3587101.

[22] F. Corucci, N. Cheney, H. Lipson, C. Laschi, and J. Bongard, "Material properties affect evolutions ability to exploit morphological computation in growing soft-bodied creatures," in *Proc. Artif. Life Conf.* Cambridge, MA, USA: MIT Press, 2016, pp. 234–241.

[23] E. Medvet, G. Nadizar, and F. Pigozzi, "On the impact of body material properties on neuroevolution for embodied agents: The case of voxel-based soft robots," in *Proc. Genetic Evol. Comput. Conf. Companion*, Jul. 2022, pp. 2122–2130.

[24] M. Mazare, S. Tolu, and M. Taghizadeh, "Adaptive variable impedance control for a modular soft robot manipulator in configuration space," *Meccanica*, pp. 1–15, 2022.

[25] J. Talamini, E. Medvet, A. Bartoli, and A. De Lorenzo, "Evolutionary synthesis of sensing controllers for voxel-based soft robots," in *Proc. Conf. Artif. Life*. Cambridge, MA, USA: MIT Press, 2019, pp. 574–581.

[26] K. A. De Jong, *Evolutionary Computation: A Unified Approach*. Cambridge, MA, USA: MIT Press, 2006.

[27] S. Koos, J.-B. Mouret, and S. Doncieux, "The transferability approach: Crossing the reality gap in evolutionary robotics," *IEEE Trans. Evol. Comput.*, vol. 17, no. 1, pp. 122–145, Feb. 2013.

[28] M. Luperto, V. Castelli, and F. Amigoni, "Predicting performance of SLAM algorithms," 2021, *arXiv:2109.02329*.

[29] Y. Zhang, T. Wang, W. He, and S. Zhu, "Human-powered master controllers for reconfigurable fluidic soft robots," *Soft Robot.*, 2023.

[30] C. G. Frazelle, A. D. Kapadia, K. E. Fry, and Ian. D. Walker, "Teleoperation mappings from rigid link robots to their extensible continuum counterparts," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2016, pp. 4093–4100.

[31] B. T. Phillips, K. P. Becker, S. Kurumaya, K. C. Galloway, G. Whittredge, D. M. Vogt, C. B. Teeple, M. H. Rosen, V. A. Pieribone, and D. F. Gruber, "A dexterous, glove-based teleoperable low-power soft robotic arm for delicate deep-sea biological exploration," *Sci. Rep.*, vol. 8, no. 1, 2018, Art. no. 14779.

[32] H. Shin, S. H. Jung, Y. R. Choi, and C. Kim, "Development of a shared remote control robot for aerial work in nuclear power plants," *Nucl. Eng. Technol.*, vol. 50, no. 4, pp. 613–618, May 2018.

[33] H. Shen, C. Wang, L. Xie, S. Zhou, L. Gu, and H. Xie, "A novel remote-controlled robotic system for cerebrovascular intervention," *Int. J. Med. Robot. Comput. Assist. Surg.*, vol. 14, no. 6, p. e1943, Dec. 2018.

[34] S. Nertinger, R. J. Kirschner, A. Naceri, and S. Haddadin, "Acceptance of remote assistive robots with and without human-in-the-loop for healthcare applications," *Int. J. Social Robot.*, vol. 2022, pp. 1–20, Oct. 2022.

[35] S. Schaal, "Learning from demonstration," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 9, 1996.

[36] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robot. Auto. Syst.*, vol. 57, no. 5, pp. 469–483, May 2009.

[37] J. F.-S. Lin, P. Carreno-Medrano, M. Parsapour, M. Sakr, and D. Kulić, "Objective learning from human demonstrations," *Annu. Rev. Control*, vol. 51, pp. 111–129, Jan. 2021.

[38] Z. Zhu and H. Hu, "Robot learning from demonstration in robotic assembly: A survey," *Robotics*, vol. 7, no. 2, p. 17, Apr. 2018.

[39] N. Maruyama, A. Masumori, and T. Ikegami, "Studying embodied cognition in the Android alter," in *Proc. Conf. Artif. Life*. Cambridge, MA, USA: MIT Press, 2021.

[40] E. Medvet and G. Nadizar, "GP for continuous control: Teacher or learner? The case of simulated modular soft robots," in *Genetic Programming Theory and Practice XX*. Cham, Switzerland: Springer, 2023, pp. 1–23.

[41] D. Matthews and J. Bongard, "Crowd grounding: Finding semantic and behavioral alignment through human robot interaction," in *Proc. Conf. Artif. Life*. Cambridge, MA, USA: MIT Press, 2020, pp. 148–156.

[42] S. Kafley, "Twitch plays simulated soft robotics," Tech. Rep., 2022.

[43] G. Li, H. Su, and W. Zhu, "Incorporating external knowledge to answer open-domain visual questions with dynamic memory networks," 2017, *arXiv:1712.00733*.

[44] X. Wang, P. Kapanipathi, R. Musa, M. Yu, K. Talamadupula, I. Abdelaziz, M. Chang, A. Fokoue, B. Makni, and N. Mattei, "Improving natural language inference using external knowledge in the science questions domain," in *Proc. AAAI Conf. Artif. Intell.*, 2019, vol. 33, no. 1, pp. 7208–7215.

[45] J. Gu, H. Zhao, Z. Lin, S. Li, J. Cai, and M. Ling, "Scene graph generation with external knowledge and image reconstruction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 1969–1978.

[46] X. Pan, K. Sun, D. Yu, J. Chen, H. Ji, C. Cardie, and D. Yu, "Improving question answering with external knowledge," 2019, *arXiv:1902.00993*.

[47] H. Sharma and A. S. Jalal, "Incorporating external knowledge for image captioning using CNN and LSTM," *Mod. Phys. Lett. B*, vol. 34, no. 28, Oct. 2020, Art. no. 2050315.

[48] J. Powers, S. Pell, and J. Bongard, "A framework for search and application agnostic interactive optimization," in *Proc. Conf. Artif. Life*, 2020.

[49] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.

[50] M. Virgolin, A. De Lorenzo, F. Randone, E. Medvet, and M. Wahde, "Model learning with personalized interpretability estimation (ML-PIE)," in *Proc. Genetic Evol. Comput. Conf. Companion*, Jul. 2021, pp. 1355–1364.

[51] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe, ''Training language models to follow instructions with human feedback,'' 2022, *arXiv:2203.02155*.
[52] Y. Bai et al., ''Training a helpful and harmless assistant with reinforcement learning from human feedback,'' 2022, *arXiv:2204.05862*.
[53] K. Lee, H. Liu, M. Ryu, O. Watkins, Y. Du, C. Boutilier, P. Abbeel, M. Ghavamzadeh, and S. S. Gu, ''Aligning text-to-image models using human feedback,'' 2023, *arXiv:2302.12192*.
[54] J. Secretan, N. Beato, D. B. D. Ambrosio, A. Rodriguez, A. Campbell, and K. O. Stanley, ''Picbreeder: Evolving pictures collaboratively online,'' in *Proc. SIGCHI Conf. Hum. Factors Comput. Syst.*, Apr. 2008, pp. 1759–1768.
[55] M. Baltieri, H. Iizuka, O. Witkowski, L. Sinapayen, and K. Suzuki, ''Hybrid life: Integrating biological, artificial, and cognitive systems,'' 2022, *arXiv:2212.00285*.
[56] Y. Kuwana, I. Shimoyama, and H. Miura, ''Steering control of a mobile robot using insect antennae,'' in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst., Hum. Robot Interact. Cooperat. Robots*, 1995, pp. 530–535.
[57] Y. Kuwana, S. Nagasawa, I. Shimoyama, and R. Kanzaki, ''Synthesis of the pheromone-oriented behaviour of silkworm moths by a mobile robot with moth antennae as pheromone sensors,'' *Biosensors Bioelectron.*, vol. 14, no. 2, pp. 195–202, 1999.
[58] W. Matsunaga and E. Watanabe, ''Visual motion with pink noise induces predation behaviour,'' *Sci. Rep.*, vol. 2, no. 1, p. 219, Jan. 2012.
[59] J. Halloy, G. Sempo, G. Caprari, C. Rivault, M. Asadpour, V. Durier, S. Canonge, C. Detrain, N. Correll, A. Martinoli, F. Mondada, R. Siegwart, and J. L. Deneubourg, ''Social integration of robots into groups of cockroaches to control self-organized choices,'' *Science*, vol. 318, no. 5853, pp. 1155–1158, Nov. 2007.
[60] H. Shirado and N. A. Christakis, ''Locally noisy autonomous agents improve global human coordination in network experiments,'' *Nature*, vol. 545, no. 7654, pp. 370–374, May 2017.
[61] G. Nadizar, E. Medvet, S. Nichele, and S. Pontes-Filho, ''An experimental comparison of evolved neural network models for controlling simulated modular soft robots,'' *Appl. Soft Comput.*, vol. 145, Sep. 2023, Art. no. 110610.
[62] G. Nadizar, E. Medvet, H. Huse Ramstad, S. Nichele, F. A. Pellegrino, and M. Zullich, ''Merging pruning and neuroevolution: Towards robust and efficient controllers for modular soft robots,'' *Knowl. Eng. Rev.*, vol. 37, Nov. 2022.
[63] F. Pigozzi, Y. Tang, E. Medvet, and D. Ha, ''Evolving modular soft robots without explicit inter-module communication using local self-attention,'' in *Proc. Genetic Evol. Comput. Conf.*, Jul. 2022.
[64] E. Medvet and F. Rusin, ''Impact of morphology variations on evolved neural controllers for modular robots,'' in *Proc. 14th Int. Workshop Artif. Life Evol. Comput. (WIVACE)*, 2022.
[65] A. Ferigo, G. Iacca, E. Medvet, and F. Pigozzi, ''Evolving Hebbian learning rules in voxel-based soft robots,'' *IEEE Trans. Cognit. Develop. Syst.*, 2022.
[66] G. Nadizar, E. Medvet, F. A. Pellegrino, M. Zullich, and S. Nichele, ''On the effects of pruning on evolved neural controllers for soft robots,'' in *Proc. Genetic Evol. Comput. Conf. Companion*, Jul. 2021, pp. 1744–1752.
[67] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, ''Evolution strategies as a scalable alternative to reinforcement learning,'' 2017, *arXiv:1703.03864*.
[68] F. van Diggelen, E. Ferrante, N. Harrak, J. Luo, D. Zeeuwe, and A. E. Eiben, ''The influence of robot traits and evolutionary dynamics on the reality gap,'' *IEEE Trans. Cognit. Develop. Syst.*, 2021.
[69] E. Salvato, G. Fenu, E. Medvet, and F. A. Pellegrino, ''Crossing the reality gap: A survey on sim-to-real transferability of robot controllers in reinforcement learning,'' *IEEE Access*, vol. 9, pp. 153171–153187, 2021.
[70] P. Diaconis and R. L. Graham, ''Spearman's footrule as a measure of disarray,'' *J. Roy. Stat. Soc., B, Methodol.*, vol. 39, no. 2, pp. 262–268, Jan. 1977.
[71] E. Medvet, A. Bartoli, F. Pigozzi, and M. Rochelli, ''Biodiversity in evolved voxel-based soft robots,'' in *Proc. Genetic Evol. Comput. Conf.*, Jun. 2021, pp. 129–137.
[72] A. Ferigo, L. B. Soros, E. Medvet, and G. Iacca, ''On the entanglement between evolvability and fitness: An experimental study on voxel-based soft robots,'' in *Proc. Conf. Artif. Life*. Cambridge, MA, USA: MIT Press, 2022.
[73] M. Alimardani and K. Hiraki, ''Passive brain–computer interfaces for enhanced human–robot interaction,'' *Frontiers Robot. AI*, vol. 7, p. 125, Oct. 2020.
[74] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley, ''Picbreeder: A case study in collaborative evolutionary exploration of design space,'' *Evol. Comput.*, vol. 19, no. 3, pp. 373–403, Sep. 2011.

GIULIA MARCHIORI PIETROSANTI received the bachelor's degree (cum laude) in computer science from the University of Rome ''La Sapienza,'' Italy, in 2019, and the master's degree (cum laude) in data science and scientific computing from the University of Trieste, Italy, in 2023. While working on the M.Sc. thesis with the Evolutionary Robotics and Artificial Life Laboratory, she discovered her interests in evolutionary computation and artificial life. She is also interested in reinforcement learning.

GIORGIA NADIZAR received the bachelor's and master's degrees (cum laude) in electronics and computer engineering from the University of Trieste, Italy, in 2019 and 2021, respectively, where she is currently pursuing the Ph.D. degree with the Evolutionary Robotics and Artificial Life Laboratory. Her research interests include the realm of embodied AI and explainable AI and the intersection thereof.

FEDERICO PIGOZZI received the bachelor's degree (cum laude) in economics and the master's degree (cum laude) in data science and scientific computing from the University of Trieste, Italy, where he is currently pursuing the Ph.D. degree in computer engineering. He has been a Visiting Research Fellow (twice) with the University of Vermont, USA. His research interests include artificial life, evolutionary computation, and artificial intelligence.

ERIC MEDVET (Member, IEEE) received the degree (cum laude) in electronic engineering and the Ph.D. degree in computer engineering from the University of Trieste, Italy, in 2004 and 2008, respectively. He is currently an Associate Professor of computer engineering with the University of Trieste, the Director of the Evolutionary Robotics and Artificial Life Laboratory, and the Co-Director of the Machine Learning Laboratory. His research interests include evolutionary robotics, artificial life, evolutionary computation, and applications of machine learning.

Open Access funding provided by 'Università degli Studi di Trieste' within the CRUI CARE Agreement