

Material properties affect evolution’s ability to exploit morphological computation in growing soft-bodied creatures

Francesco Corucci^{*1,3}, Nick Cheney^{2,3}, Hod Lipson⁴, Cecilia Laschi¹ and Josh C. Bongard³

¹The BioRobotics Institute, Scuola Superiore Sant’Anna, Pisa, Italy

²Department of Biological Statistics and Computational Biology, Cornell University, Ithaca, NY, USA

³Morphology, Evolution & Cognition Lab, University of Vermont, Burlington, VT, USA

⁴Creative Machines Lab, Columbia University, NY, USA

*f.corucci@sssup.it

Abstract

The concept of morphological computation holds that the body of an agent can, under certain circumstances, exploit the interaction with the environment to achieve useful behavior, potentially reducing the computational burden of the brain/controller. The conditions under which such phenomenon arises are, however, unclear. We hypothesize that morphological computation will be facilitated by body plans with appropriate geometric, material, and growth properties, while it will be hindered by other body plans in which one or more of these three properties is not well suited to the task. We test this by evolving the geometries and growth processes of soft robots, with either manually-set softer or stiffer material properties. Results support our hypothesis: we find that for the task investigated, evolved softer robots achieve better performances with simpler growth processes than evolved stiffer ones. We hold that the softer robots succeed because they are better able to exploit morphological computation. This four-way interaction among geometry, growth, material properties and morphological computation is but one example phenomenon that can be investigated using the system here introduced, that could enable future studies on the evolution and development of generic soft-bodied creatures.

Introduction

Evolving complete and intelligent artificial creatures is one of the long-term goals of artificial life and evolutionary robotics researchers. More than two decades after the first pioneering attempts (Sims, 1994), we are still far from matching the complexity exhibited by even the simplest organisms. Nevertheless, many insights have been gained to date, and many limitations overcome.

Hand in hand with similar developments in robotics (Kim et al., 2013; Rus and Tolley, 2015), substantial steps forwards in the complexity and interestingness of evolved virtual creatures have been recently made, by allowing evolution to make use of soft materials (Hiller and Lipson, 2010, 2012; Joachimczak and Wróbel, 2012; Joachimczak et al., 2014; Cheney et al., 2013; Rieffel et al., 2014; Lessin and Risi, 2015). In addition to enhancing morphological and behavioral diversity, the use of soft materials allows morphologies that more closely mimic biological ones, thus enabling




Figure 1: A soft (a-d) and stiff (e-h) robots evolved to grow towards two lateral light sources. Red voxels expand in response to environmental stimuli, blue ones shrink. While the soft robot only employs expanding voxels and effectively exploits morphological computation, passive dynamics, and the interaction with the environment to solve the task, the stiff one is prevented from doing so due to its unsuitable material properties, and had thus to evolve a more complex and active form of control in order to achieve the same result. See them in action at: <https://youtu.be/Cw2SwPNwcfM>

the investigation of additional aspects of evolution and development that were previously beyond reach. Here we focus on two such aspects relevant to soft-bodied creatures.

The first regards morphological plasticity: the ability to change some aspects of the body during one’s lifetime. Here we investigate environment-mediated morphological development — growth in response to environmental stimuli — referred to henceforth simply as growth. Although it has been shown that morphological growth can provide adaptive advantages for machines (Bongard, 2013), previous work only focussed on rigid-bodied agents and environment-insensitive growth processes. Yet there is evidence that biological development is influenced and driven by the environment. For example, plant roots follow gradients of nutrients in the soil, while human bones and tissues alter their properties in response to mechanical loading (Wolff, 1986). Moreover, when compared to rigid-bodied

creatures, soft ones more naturally allow for some forms of morphological plasticity, that are already within the reach of soft robotics technology as well. Despite that, probably due to the lack of a general understanding of why, when, and how these new capabilities should be exploited, these robots feature to date basic forms of morphological plasticity (Shepherd et al., 2011; Corucci et al., 2015b,c), or no plasticity at all (Calisti et al., 2015; Corucci et al., 2015a; Cacucciolo, Corucci et al., 2014).

The second aspect we explore is the influence of material properties on the evolution and development of adaptive behavior. The behavior of soft-bodied creatures is to a large extent determined by their material properties, yet in soft robotics these are often fixed a priori and for the whole lifetime of an agent, perhaps after a limited number of heuristic tests. Some recently proposed ideas suggest that those properties — and softness in particular — can have implications for the development of intelligent behavior (Pfeifer and Bongard, 2006), but few studies (Nakajima et al., 2013) and theoretical frameworks (Hauser et al., 2011, 2012) elucidate and quantify these implications, typically not embracing an evolutionary and developmental approach.

Here we investigate the evolution and development of soft robots. In this context we hypothesize that with the right combination of geometry, material properties, and growth process, a robot can exploit morphological computation (Paul, 2006; Hauser et al., 2014) better than one for which one or more of these aspects is not well adapted. We test this hypothesis by evolving the body plans and developmental trajectories of simulated soft robots for a phototaxis task. Two variations of such robots are evolved: softer and stiffer ones. We find that the former achieve better performances, despite their evolved growth processes being simpler, according to an information theoretic measure. By interpreting environment-mediated growth as a form of control, it is suggested that the softer robots are in fact exploiting more morphological computation than the stiffer ones. This hypothesis is but one of many that can be tested using the system here introduced. These include the investigation of general relationships among morphology, control, evolution, and development. Such studies could provide a deeper comprehension of biological systems, with potential implications for the development of more complex, autonomous and adaptive machines.

Methods

Simulated task and environment. In this work soft-bodied creatures are simulated in the VoxCad environment (Hiller and Lipson, 2014). A number of changes and new features have been introduced in the simulator in order to enable our experiments.

First, sources of environmental stimuli can now be added to the environment. These sources are characterized by a fixed 3D location, and robot's voxels can sense the distance

from each of them.

Second, the base of each robot is fixed to the ground for the entire simulation. If a robot does not touch the ground at the beginning of the simulation, it is translated along the z axis before the simulation starts, until it does. This is done to put emphasis on growth and deformability, ruling out locomotion strategies to approach the light sources.

Third, differently from other works adopting VoxCad (Cheney et al., 2015, 2013; Methenitis et al., 2015), in this experiment there is no fast-twitch actuation mechanism (i.e. the fast control based on an oscillating global signal is disabled). A distributed growth mechanism has been implemented instead, acting at a slower time scale (more below).

The task is inspired by plants, and consists in performing stationary phototaxis: growing towards static sources present in the environment. While reaching a single source is not a particularly difficult task, simultaneously pointing toward multiple ones becomes more challenging, as it requires the ability to evolve modular, branching structures. The specific growth mechanism is detailed in the next sections, as well as the underlying developmental paradigm.

Developmental paradigm. Different approaches have been adopted in the literature in order to model developmental processes (Stanley and Miikkulainen, 2003). Those can be roughly classified based on the level of abstraction with respect to the biological phenomenon they try to capture. Among high-level abstractions we find grammar-based approaches (Rieffel et al., 2014; Hornby and Pollack, 2002) and CPPN-based ones (Stanley, 2007; Cheney et al., 2013; Auerbach and Bongard, 2014). Lower-level abstractions, broadly referred as *cell chemistry* approaches (Doursat et al., 2013; Joachimczak et al., 2014; Bongard and Pfeifer, 2003), model finer details of developmental processes, such as gene regulatory dynamics.

Despite many achievements in cell chemistry, a drawback of these approaches lies in their complexity. On the other hand, when adopting a high-level perspective, the risk exists of overlooking potentially useful aspects of development. As an example, CPPNs and grammatical encodings neglect both the unfolding over time of biological developmental processes as well as the interaction of the creature with the environment during those processes.

The approach proposed here is based on CPPNs, but empowers them by joining their ability to capture the formation of regular patterns (Stanley, 2006) with an environment-mediated developmental stage that unfolds over time. While the implications of such a choice deserve to be thoroughly investigated in future work, we note that this approach enables potentially interesting feedback loops during development: growth is guided by environmental stimuli, but modifies in turn the sensory information the creature will experience next. Also, as mutation effects may arise at different points during development, the ability to enact changes later




Figure 2: Different attributes of the robot can be “painted” by different CPPNs. CPPN1 dictates the geometry of the robot, while CPPN2 determines its growth properties. In the current system red voxels expand in response to environmental stimuli, while blue ones shrink.

in development may allow for smoother fitness gradients than all-or-nothing mutations (Hinton and Nowlan, 1987).

Time scales. The proposed developmental paradigm embeds all three time scales experienced by living systems: the evolutionary/phylogenetic time scale, the developmental/ontogenetic time scale, and the sensorimotor dynamics timescale (Pfeifer and Bongard, 2006). The sensorimotor timescale is here represented by the interaction of the body with the environment during growth (e.g. gravity, collisions, detected light levels, etc.).

The genotype. The encoding here adopted is based on CPPNs (Stanley, 2007). Designed to capture the formation of regular patterns in developmental systems without modeling development per se (Stanley, 2006), CPPNs are networks that convolve incoming spatial information to produce outputs that tend to exhibit symmetry, repetition, and repetition with variation.

In order to enforce a distinction between geometric and growth properties, here we adopt two different CPPNs (Fig. 2) that are queried for each voxel of a cubic workspace. They receive the same inputs: the 3D location of the voxel (x, y, z), the polar radius (d), and a constant bias (b).

The first CPPN, determining the creature’s geometrical structure, has a single output $o \in [-1, 1]$ that dictates whether a voxel should be empty ($o < 0$) or filled ($o \geq 0$). Differently from similar setups (Cheney et al., 2013), evolution is provided with a single material in the experiments here reported, that is assigned to all non-empty voxels. A different stiffness can be specified for this material in different evolutionary runs, though all voxels in each run share a global and constant stiffness parameter.

The second CPPN determines the growth properties of each voxel. It also has a single output $\in [-1, 1]$ henceforth referred to as the *growth rate* (g_{rv}), whose role is described in detail in the next section.

	Description
$D_v^{(t)}$	Linear dimension of voxel v at time t
$s_v^{(t)}$	Scaling factor of voxel v at time t
D_{vn}	Nominal dimension of voxel v (set to 1)
g_a	Growth amplitude $\in (0, 1)$ (equal \forall voxel, set to 0.5)
N	Number of environmental sources
i	A specific source
$s_{vi}^{(t)}$	Influence of source i on the scaling factor $s_v^{(t)}$
$s_{v,\min}$	Lower bound for $s_v^{(t)}$ (set to 0.1)
g_{rv}	Growth parameter of voxel v ($\in [-1, 1]$)
$\bar{d}_i^{(t)}$	Distance of voxel v from source i , normalized by voxel size (the latter being set to 0.01)

Table 1: Description of parameters appearing in Eq. 1

Environment-mediated morphological development.

We simulate environment-mediated growth by enabling voxels to change volume in response to environmental stimuli (i.e. distance from each light source). This choice is motivated by the fact that localized volumetric changes are easy to achieve with currently-available soft robots (e.g. exploiting pneumatic actuation). A localized change in stiffness is also within the reach of current technology (Majmudar et al., 2007), making this kind of plasticity the next candidate to be integrated into our system. Topological modifications are, on the other hand, more difficult to achieve in real soft robots, as they require adding or removing material. Nevertheless, attempts have been made in this direction as well (Brodbeck et al., 2012).

The growth process is governed by equations and parameters reported in Eq. 1 and in Tab. 1:

\forall voxel v , at each time step t

$$\begin{aligned}
 D_v^{(t)} &= s_v^{(t)} \cdot D_{vn} \\
 s_v^{(t)} &= \max \left[g_a \cdot \tanh \left(\sum_{i=1}^N s_{vi}^{(t)} \right) + 1, s_{v,\min} \right] \\
 s_{vi}^{(t)} &= g_{rv} \cdot \bar{d}_i^{(t)}
 \end{aligned} \tag{1}$$

The *growth rate* parameter g_{rv} determines the quality and the extent of the localized volumetric change for each voxel. When $g_{rv} > 0$ the voxel will expand when close to a source, when $g_{rv} < 0$ it will shrink. When g_{rv} is exactly zero, the voxel is insensitive to environmental stimuli. The greater the magnitude of g_{rv} for a particular voxel, the more pronounced will be its volumetric variation, that is also modulated by the distance from the sources. Each voxel can experience a considerable modification due to development: having set

$g_a = 0.5$ entails a 50% linear contraction/expansion with respect to the nominal size, resulting in a $\sim 238\%$ variation in volume. The parameter $s_{v,\min}$ ensures that the voxel does not shrink below a given size (here the 10% of the nominal size), for stability of simulations. The quantity $\Delta s_{v,\max}$ (not reported in Eq. 1 for ease of reading) dictates the maximum allowed $\Delta s_v = |s_v^{(t)} - s_v^{(t-1)}|$ between two subsequent time steps, as follows:

$$\text{if } (\Delta s_v > \Delta s_{v,\max}) \text{ then:}$$

$$s_v^{(t)} = s_v^{(t-1)} + \text{sign}(\Delta s_v) \cdot \Delta s_{v,\max}$$

In addition to influencing the stability of the simulation, this parameter (set to 0.0005 in our experiments) regulates the speed of the growth process: the higher $\Delta s_{v,\max}$, the more rapid the growth. The value for $\Delta s_{v,\max}$ was selected in such a way that development acts over a slower time scale with respect to the typical sensorimotor dynamics (such as those generating locomotion or grasping behavior), thus implementing the developmental time scale.

Development is based on distributed sensing and actuation: each voxel senses the distance from all the sources and acts accordingly. Nevertheless, coordinated behavior emerges, for at least two reasons: first, nearby voxels experience similar sensory stimuli, and second, CPPNs tend to produce patches of tissue with homogeneous or smoothly varying growth parameters.

Optimization. A multi-objective implementation of NEAT (Cheney et al., 2015) has been adopted. Before performing selection, pareto ranking is applied, according to three objectives: the order in which sorting is performed determines the relative importance of each of them. The objectives are listed below from the most to the least important:

1. Minimize the distance from each of the sources
2. Minimize the number of employed voxels
3. Minimize the age of each individual

The first objective selects for phototaxis. This is implemented by minimizing the sum $\sum_{i=1}^N d_{\min,i}$, where N is the number of sources and $d_{\min,i}$ is the minimum distance between the robot and the i -th source. The second objective selects for smaller robots. The first two objectives are antagonistic as it is easier, in general, for larger robots to be closer to the sources (even if they do not grow at all). The combination of the two objectives thus selects for robots that exploit the growth process and their deformability to accomplish the task. The third objective helps maintain diversity in the population (Schmidt and Lipson, 2011).

Morphological computation and control complexity. Morphological computation (Hauser et al., 2014) has been defined as "computation obtained through the interaction of

physical forms" (Paul, 2006). When it comes to robotics and embodied cognition (Pfeifer and Bongard, 2006), the idea is that part of the computation needed to perform a task can take place (implicitly or explicitly) not only in the brain/controller, but also within the body itself, provided that it has suitable characteristics. It has been argued that this property can alleviate the computational burden of the brain, simplifying the controller and achieving a more balanced brain-body trade-off (Pfeifer and Bongard, 2006; Paul, 2006; Hauser et al., 2011) that could hold the key to more intelligent, effective and natural behaviors. Many examples have been described in the literature (Pfeifer and Bongard, 2006). It is often postulated that systems benefiting from morphological computation tend to exploit the interaction and dynamical coupling with the environment in a beneficial way, e.g. leveraging passive dynamics in place of active control.

We hypothesize that material properties can affect evolution's ability to exploit morphological computation. To test this, we here define morphology as the robot's shape and material properties, and its 'controller' as the distributed growth mechanism which achieves phototaxis. We defend this latter definition as, like control, growth here closes the sensation-action feedback loop (although over a slower time scale).

For our purposes, we can thus define morphological computation as a property that simplifies the growth controller by exploiting in a beneficial way the interplay between material, geometric, and growth properties through the dynamical interaction with the environment.

To measure the extent of morphological computation in a given robot, we define control complexity as follows:

$$H(g_c) = - \sum_{i=1}^n p_i \log_2 p_i \quad (2)$$

where:

$$g_c = \{g_{rv} \forall \text{ voxel } v\}$$

$$p_i = \int_{x_i}^{x_{i+1}} p(x) dx \quad (3)$$

$$x_i = -1 + 0.02i \quad i = 0 \dots 100$$

The real-valued random variable g_c is associated to the *growth rate* quantity (Eq. 1, Tab. 1), embracing all parameters that collectively shape the growth trajectory of a given robot. The quantity $H(g_c)$ is the Shannon entropy (Shannon, 1948) of such a variable, whose probability density function $p(x)$ is discretized using $n = 101$ uniform bins (Eq. 3).

The control complexity of g_c thus corresponds to the number of bits that are necessary to describe the pattern of growth parameters: the higher this number, the more complex the controller. Consider two robots (r_1, r_2) and their associated growth controllers (g_{c1}, g_{c2}). We will state that a difference $\Delta H = H(g_{c2}) - H(g_{c1}) > 0$ between the

two controllers indicates that g_{c2} is more complex than g_{c1} . Moreover, if the two robots happen to score the same fitness, we will argue that r_1 better exploits morphological computation, as it requires simpler control to produce an equally effective behavior. It should be made clear that we are not providing here a general information theoretic metric to capture morphological computation (Zahedi and Ay, 2013), but rather a proxy to measure its effect in our setting.

Experiments Populations composed of 30 individuals are allowed to evolve for a maximum of 1500 generations. The maximum evaluation time for each individual is 3.5s (simulation time, wall time is higher). The simulation is stopped earlier if the robot settles into a static conformation before the allocated time elapses. The growth process starts after the first 0.5 seconds, which usually allows the initial shape to settle into an equilibrium position.

A first set of experiments is performed to qualitatively assess the overall ability of the system to evolve effective robots. To this end, 10x10x10 and 8x8x8 robots have been evolved in several environments, differing in the number (from 1 to 4) and position of environmental sources (Fig. 3). Material stiffness is set to $E = 5 \text{ MPa}$, corresponding to a rather soft material (comparable to rubber). Five runs were performed for each configuration.

A second set of experiments is then performed with 6x6x6 robots, characterized by different stiffness values ($E_1 = 500 \text{ MPa}$, $E_2 = 5 \text{ MPa}$) evolving in the same environment (2 laterally placed sources). Twenty independent runs were performed for each treatment. Reported confidence intervals are computed with a bootstrapping method, while p-values are the result of the Mann-Whitney U test.

The code used to produce these results is publicly available at: <https://goo.gl/cA2lu0>. A video showing some of the creatures in action is available at: <https://youtu.be/Cw2SwPNwcfM>

Results and Discussion

A sample of the fittest morphologies evolved in preliminary trials is reported in Fig. 3 and in the accompanying video. Symmetry and modularity can be observed, with the latter property evidenced by the emergence of relatively independent appendages. As these features are ubiquitous in natural systems, their presence here may suggest the potential for more competent and scalable virtual creatures. Moreover, these properties appear to be selected for by our task and environment, as this level of morphological regularity is not common in similar settings (Cheney et al., 2013, 2015; Methenitis et al., 2015).

It can be noted that the best individuals from these runs tend to only exploit expanding (red) tissue, and not shrinking (blue) voxels (Fig. 3). Evolved creatures appear able to leverage their passive deformability and interaction with the environment in order to solve the task, rather than requiring




Figure 4: Average fitness over 20 independent runs. Softer robots have an evolutionary advantage over stiffer ones in this task/environment.

differential expansion and contraction to point towards light sources. For example, cantilevers spontaneously evolve (Fig 3b). Like human-built cantilevered structures, these robots distribute stresses across themselves with a minimum of support structure. Moreover, the curvature needed to point towards the two lateral sources spontaneously emerge from the passive interaction of the expanding body with the environment (and with gravity, in particular), rather than through internal actuation of the creature. This corresponds, intuitively, to the idea of morphological computation.

Given the scarce presence of highly-fit robots exploiting more complex forms of control — based on the combination of shrinking and expanding voxels — we hypothesize that it may be easier for evolution to discover solutions based on morphological computation rather than explicit control, provided that material properties allow it to do so. This would confer, in general, an evolutionary advantage to robots that have the “right” material properties for a given task/environment. This hypothesis is tested with the second set of experiments.

Geometry, materials, growth, and morphological computation. Fig. 4 reports the evolution of robots with stiffer (E_1) or softer (E_2) material properties, optimized to simultaneously approach two lateral light sources placed on opposite sides of the creature. Results show that softer robots have an evolutionary advantage over stiffer ones in this particular task/environment, which may be attributed to evolution’s ability to better exploit morphological computation instead of developing more complex and active forms of control to regulate their shape. This can be qualitatively observed in the comparison of two highly-fit robots reported in Fig. 1 (see them in action in the accompanying video). With the softest material, evolution produced a passive cantilever, taking advantage of gravity to unfold the shape towards the two sources during growth. Under the effect of gravity, the

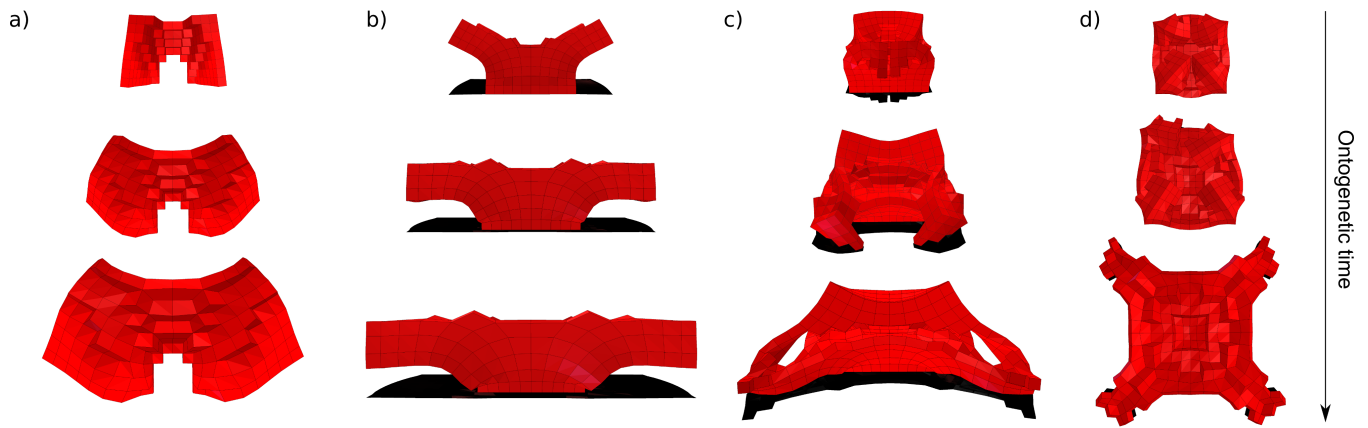


Figure 3: A small sample of the growing soft creatures, evolved in environments featuring: a-c) two lateral sources d) four sources placed at the corners. a,d) top view, b,c) front view. Most of the highly fit soft robots only exploit expanding tissue.

structure passively deforms and achieves an effective curvature, being able to sustain its own weight and point towards the sources at the same time. On the other hand, the stiffer robot fights gravity holding the two rigid arms horizontally, achieving the curvature needed to direct its appendages towards the sources through the antagonistic action of shrinking and expanding voxels in the central part of the body.

Further analyses suggest the generality of these observations in this task/environment. Figure 5 shows that stiffer robots use substantially more shrinking voxels in general. Their growth processes involve more active control in the sense that the appendages are pulled as well as pushed to achieve proximity to the environmental sources. This suggests that for stiffer robots, simpler strategies in which curvature is achieved passively in response to weight are either harder to find in the search space or are not viable at all.

Given the fitness benefit of expanding voxels, enlarging the volume of the robots and allowing them to approach the sources more closely, we would expect shrinking voxels only to be employed when necessary to control the direction of evolved appendages. Thus, the presence of more shrinking voxels in the stiffer robots suggest that they are unable to perform passive pointing from their material properties alone, as exemplified in the softer robots (Fig. 1).

The intuitive considerations regarding control complexity and morphological computation are also confirmed by an information theoretic analysis of the evolved robots: the average control complexity $H(g_c)$ (the global entropy across g_c) is significantly higher for the stiffer robots than for the softer ones (Fig. 6). In other words, stiffer robots employ more complex and active controllers. This again suggests that their morphologies are unable to perform the task to the same level without control.

In summation, softer robots better exploit morphological computation in this particular task/environment, achieving better performances (Fig. 4) with simpler controllers (Fig.




Figure 5: Stiffer robots tend to employ significantly more shrinking voxels than softer ones ($p < 0.002$), in the attempt to actively control the shape.

5, 6). This confers an evolutionary advantage to them over stiffer robots, if evolved alongside.

It should be noted that these results do not mean that softer is *always* better. It would be possible to define a task/environment that confers an advantage to stiffer robots (e.g. grow towards sources placed mid-air, where it is easier for stiffer robots to sustain their weight). What has been demonstrated is that for a specific task/environment, material properties can have a pronounced effect on evolution's ability to exploit morphological computation, i.e. to produce well adapted morphologies that exploit the interaction with the environment in a beneficial way.

Conclusions and Future Work

In this paper a novel system to study the evolution and development of soft-bodied creatures has been presented. The system is able to evolve robots that exhibit desirable morphological properties such as symmetry, modularity, and ex-




Figure 6: Stiffer robots exhibit more complex growth controllers than softer ones ($p < 0.005$), yet the latter achieve better performances (Fig. 4). It is argued that this difference is due to morphological computation, strictly connected to the material properties of robots in the two treatments.

exploitation of morphological computation. More specifically, it was shown that certain combinations of geometry, material properties, and environment-mediated growth make it more or less difficult for evolution to discover phenotypes that exploit morphological computation. Despite being a fundamental aspect of soft robotics, the interplay among these properties remains to date largely unexplored. Results also suggest that arbitrarily fixing even one of these dimensions can make it difficult for evolution to produce effective behaviors. Ideally, all of them should be put under evolutionary and/or developmental control, so that an optimal combination can be discovered. Future work will be directed towards exploring the potential evolutionary advantage of morphological plasticity, as well as possible benefits in terms of adaptivity and robustness. The environmental influence during development deserves special attention as well: its potential benefits for organisms as well as adaptive machines will be investigated. Another major topic that can now be studied is the general relationship between evolution, development and adaptive behavior. We believe that many interesting questions can be answered using the approach described here, and could help shed light on biological questions, while simultaneously contributing to engineering fields such as soft robotics.

Acknowledgments

This work was supported by the Vermont Advanced Computing Core, the National Science Foundation through grants PECASE-0953837 and INSPIRE-1344227, and NASA Space Technology Research Fellowship #NNX13AL37H for N. Cheney. Thanks to Dr. Marcin Szubert for proofreading and useful discussion.

References

- Auerbach, J. E. and Bongard, J. C. (2014). Environmental influence on the evolution of morphological complexity in machines. *PLoS Comput Biol*, 10(1):e1003399.
- Bongard, J. C. (2013). Evolutionary robotics. *Communications of the ACM*, 56(8):74–83.
- Bongard, J. C. and Pfeifer, R. (2003). Evolving complete agents using artificial ontogeny. In *Morpho-functional Machines: The New Species*, pages 237–258. Springer.
- Brodbeck, L., Wang, L., and Iida, F. (2012). Robotic body extension based on hot melt adhesives. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 4322–4327. IEEE.
- Cacucciolo, V.*, Corucci, F.*, Cianchetti, M., and Laschi, C. (2014). Evolving optimal swimming in different fluids: A study inspired by batoid fishes. In *Biomimetic and Biohybrid Systems*, pages 23–34. Springer.
- Calisti, M., Corucci, F., Arienti, A., and Laschi, C. (2015). Dynamics of underwater legged locomotion: modeling and experiments on an octopus-inspired robot. *Bioinspiration & Biomimetics*, 10(4):046012.
- Cheney, N., Bongard, J., and Lipson, H. (2015). Evolving soft robots in tight spaces. In *Proceedings of the 2015 annual conference on Genetic and Evolutionary Computation*, pages 935–942. ACM.
- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th annual conference on Genetic and Evolutionary Computation*, pages 167–174. ACM.
- Corucci, F., Calisti, M., Hauser, H., and Laschi, C. (2015a). Evolutionary discovery of self-stabilized dynamic gaits for a soft underwater legged robot. In *Proceedings of the 17th International Conference on the Advanced Robotics*, pages 337–344. IEEE.
- Corucci, F., Calisti, M., Hauser, H., and Laschi, C. (2015b). Novelty-based evolutionary design of morphing underwater robots. In *Proceedings of the 2015 annual conference on Genetic and Evolutionary Computation*, pages 145–152. ACM.
- Corucci, F., Calisti, M., Hauser, H., and Laschi, C. (2015c). Shaping the body to shape the behavior: a more active role of the morphology in the brain-body trade-off. *Late Breaking Proceedings of the 13th European Conference on Artificial Life*, pages 7–8.
- Doursat, R., Sayama, H., and Michel, O. (2013). A review of morphogenetic engineering. *Natural Computing*, 12(4):517–535.

- Hauser, H., Fuchslin, R. M., and Pfeifer, R. (2014). *Opinions and Outlooks on Morphological Computation*. Zürich.
- Hauser, H., Ijspeert, A. J., Fuchslin, R. M., Pfeifer, R., and Maass, W. (2011). Towards a theoretical foundation for morphological computation with compliant bodies. *Biological cybernetics*, 105(5-6):355–370.
- Hauser, H., Ijspeert, A. J., Fuchslin, R. M., Pfeifer, R., and Maass, W. (2012). The role of feedback in morphological computation with compliant bodies. *Biological cybernetics*, 106(10):595–613.
- Hiller, J. and Lipson, H. (2012). Automatic design and manufacture of soft robots. *Robotics, IEEE Transactions on*, 28(2):457–466.
- Hiller, J. and Lipson, H. (2014). Dynamic simulation of soft multimaterial 3d-printed objects. *Soft Robotics*, 1(1):88–101.
- Hiller, J. D. and Lipson, H. (2010). Evolving amorphous robots. In *Proceedings of the 12th International Conference on the Synthesis and Simulation of Living Systems*, pages 717–724.
- Hinton, G. E. and Nowlan, S. J. (1987). How learning can guide evolution. *Complex systems*, 1(3):495–502.
- Hornby, G. S. and Pollack, J. B. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial life*, 8(3):223–246.
- Joachimczak, M., Suzuki, R., and Arita, T. (2014). Fine grained artificial development for body-controller co-evolution of soft-bodied animats. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 239–246.
- Joachimczak, M. and Wróbel, B. (2012). Co-evolution of morphology and control of soft-bodied multicellular animats. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pages 561–568. ACM.
- Kim, S., Laschi, C., and Trimmer, B. (2013). Soft robotics: a bioinspired evolution in robotics. *Trends in biotechnology*, 31(5):287–294.
- Lessin, D. and Risi, S. (2015). Soft-body muscles for evolved virtual creatures: The next step on a biomimetic path to meaningful morphological complexity. In *European Conference on Artificial Life*.
- Majmudar, T., Sperl, M., Luding, S., and Behringer, R. P. (2007). Jamming transition in granular systems. *Physical review letters*, 98(5):058001.
- Methenitis, G., Hennes, D., Izzo, D., and Visser, A. (2015). Novelty search for soft robotic space exploration. In *Proceedings of the 2015 annual conference on Genetic and Evolutionary Computation*, pages 193–200. ACM.
- Nakajima, K., Hauser, H., Kang, R., Guglielmino, E., Caldwell, D. G., and Pfeifer, R. (2013). A soft body as a reservoir: case studies in a dynamic model of octopus-inspired soft robotic arm. *Front. Comput. Neurosci*, 7(10.3389).
- Paul, C. (2006). Morphological computation: A basis for the analysis of morphology and control requirements. *Robotics and Autonomous Systems*, 54(8):619–630.
- Pfeifer, R. and Bongard, J. (2006). *How the body shapes the way we think: a new view of intelligence*. MIT press.
- Rieffel, J., Knox, D., Smith, S., and Trimmer, B. (2014). Growing and evolving soft robots. *Artificial life*, 20(1):143–162.
- Rus, D. and Tolley, M. T. (2015). Design, fabrication and control of soft robots. *Nature*, 521(7553):467–475.
- Schmidt, M. and Lipson, H. (2011). Age-fitness pareto optimization. In *Genetic Programming Theory and Practice VIII*, pages 129–146. Springer.
- Shannon, C. (1948). A mathematical theory of communication, bell system technical journal 27: 379-423 and 623–656. *Mathematical Reviews (MathSciNet): MR10, 133e*.
- Shepherd, R. F., Ilijevski, F., Choi, W., Morin, S. A., Stokes, A. A., Mazzeo, A. D., Chen, X., Wang, M., and Whitesides, G. M. (2011). Multigait soft robot. *Proceedings of the National Academy of Sciences*, 108(51):20400–20403.
- Sims, K. (1994). Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22. ACM.
- Stanley, K. O. (2006). Exploiting regularity without development. In *Proceedings of the AAAI Fall Symposium on Developmental Systems*, page 37. AAAI Press Menlo Park, CA.
- Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162.
- Stanley, K. O. and Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130.
- Wolff, J. (1986). *The law of bone remodelling*. Springer Science & Business Media.
- Zahedi, K. and Ay, N. (2013). Quantifying morphological computation. *Entropy*, 15(5):1887–1915.