Interoceptive robustness through environment-mediated morphological development



Figure 1: A single robot grows calluses as it walks, in response to pressure on its feet (youtu.be/@cmwpcxSUWI).

ABSTRACT

Typically, AI researchers and roboticists try to realize intelligent behavior in machines by tuning parameters of a predefined structure (body plan and/or neural network architecture) using evolutionary or learning algorithms. Another but not unrelated longstanding property of these systems is their brittleness to slight aberrations, as highlighted by the growing deep learning literature on adversarial examples. Here we show robustness can be achieved by evolving the geometry of soft robots, their control systems, and how their material properties develop in response to one particular interoceptive stimulus (engineering stress) during their lifetimes. By doing so we realized robots that were equally fit but more robust to extreme material defects (such as might occur during fabrication or by damage thereafter) than robots that did not develop during their lifetimes, or developed in response to a different interoceptive stimulus (pressure). This suggests that the interplay between changes in the containing systems of agents (body plan and/or neural architecture) at different temporal scales (evolutionary and developmental) along different modalities (geometry, material properties, synaptic weights) and in response to different signals (interoceptive and external perception) all dictate those agents' abilities to evolve or learn capable and robust strategies.

GECCO '18, July 15–19, 2018, Kyoto, Japan

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5618-3/18/07...\$15.00 https://doi.org/10.1145/3205455.3205529

CCS CONCEPTS

Computing methodologies → Mobile agents;

KEYWORDS

Soft robotics.

ACM Reference Format:

Sam Kriegman, Nick Cheney, Francesco Corucci, and Josh C. Bongard. 2018. Interoceptive robustness through environment-mediated morphological development. In *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3205455.3205529

1 INTRODUCTION

A major characteristic of life is that three broad time scales are relevant to it: evolution, development and physiological functioning. Engineered systems, in marked contrast, often employ an evolutionary or learning algorithm to improve their behavior over time, but rarely employ morphological development; any changes to the physical layout are made in between evaluations [9, 26, 34], if they are made at all.

Two notable exceptions are modular robots [42], which may reconfigure their bodies by adding and removing discrete structures, and soft robots [33], which may continuously alter the local volumes of different parts of their bodies while behaving. Others [4] have approximated topological change in rigid bodies by extending outward and angling downward appendages using a combination of linear and rotary actuators, thus simulating limb growth.

Several computational but embodied models of *prenatal* development have been reported in the literature [5, 12, 14, 15, 28]. As implied, cellular growth therein occurred prior to any physiological

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

functioning. Thus, these studies included change during only two of the three time scales relevant to life: evolutionary and behavioral change, but not postnatal developmental change.

The most common argument in favor of development is that some aspects of the environment are unpredictable, so it is advantageous to leave some decisions up to development rather than specifying them genetically. Although self evident, it remains to determine which mechanisms of development should be instantiated in robots to realize plastic, adaptive, and useful machines.

Naturally, the performance of an evolved system depends on its capacity for evolutionary improvement: its evolvability. Development can, under certain conditions, smooth the search space evolution that operates in, thus increasing evolvability. This process, known as the Baldwin effect [2, 13], starts with an advantageous characteristic acquired during the development of individuals, such as the callouses in Fig. 1. This can create a new gradient in the evolutionary search space, rewarding descendants that more rapidly manifest the trait during their lifetimes [20, 24] and retain it through the remainder of their lifetime [25]. Assuming such mutations exist and can be naturally selected [24], following the gradient requires incrementally reducing development in the manifold of the search space that can express variations on the trait [40].

However, fitness landscapes that evolution climbs, and development sometimes smooths, tend not to remain static in realistic settings. On this vacillating landscape, when the best thing to do does not remain the same, a highly evolvable but non-robust system will need to keep starting over from scratch every time the conditions change. Computational and engineered systems provide countless examples of systems with nearly perfect performance in a controlled environment, such as a factory, but who turn out to be (often comically) brittle to slight changes in their internal structure, such as damage, or their external environment such as moving on to new terrain or transferal from simulation to reality [1, 7, 17, 23, 29, 37].

Although generally absent from engineered systems (but see [3, 11]), the canonical form of robustness is seen to some extent in all organisms, and it comes from the act of development itself. For example, a plant that grows according to a fixed program will capture less light than a plant that grows toward sunlight. But there is another, more subtle form of robustness that we will refer to as 'intrinsic robustness' because it is a property of a system's structure rather than of the process by which it may change.

Developmental change produces intrinsically robust systems because they evolved from designs that had to maintain adequate performance along additional dimensions of change [4, 24]. Through morphological development specifically, evolution is compelled to maintain designs that are capable across a series of body plans, with different sensor-motor contingencies; and the ability to tolerate such perturbations can become inherited to some extent in descendants' behaviors [4] and morphologies [24], even when their developmental flexibility is reduced or completely removed by canalization or fabrication.

And yet, despite the ubiquity of morphological development in nature, and the adaptive advantages it seemingly confers, there are only a handful of cases reported in the literature in which a simulated robot's mechanical structure was allowed to change while it was behaving (e.g. [4, 22, 24, 25, 39]), all of which modeled morphological development as a genetically predetermined process: the environment could not influence the way in which development unfolded.

Assuming that an engineered system is capable of local morphological change in response to environmental signals, it is unclear how it should do so, beyond the examples of morphological plasticity observed in nature. Examples include Wolff's law [30]-bone grows in response to particular mechanical loading profiles-and Davis' law-soft tissue increases in strength in response to intermittent mechanical demands. One can envisage other such laws that are not known to occur in biology but could be helpful in a specific artificial system, such as end effectors softening in response to pressure, which might enhance their ability to safely manipulate irregular or delicate objects [6]. Indeed, the genesis of the work presented here is one such anecdotal example given in [10], where a single robot, subjected to an abrupt doubling in gravity, stiffened its body in reaction to the increased pressure. However, whether or how it could provide a behavioral advantage, nor whether pressure is the best interoceptive signal to developmentally respond to, was not investigated.

As a step towards a more adequate picture, we introduce here a simple form of a developmental feedback mechanism: Genetic systems dictate how organisms develop in response to interoceptive stimuli, and development alters the kinds of interoceptive conditions the organism experiences. More specifically, at every time step, the proposed model of closed-loop development:

- (1) 'listens to' load signatures generated from movement; and, in response,
- (2) modifies the robot's rigidity,

which will change the way it distributes load and generates movement at the next time step.

Optimizing a system that may form a continuum of rigid and soft components—and in which this admixture may change over time—is extremely nonintuitive and underexplored. Thus, a study of the adaptive properties of such systems—and how they can best be optimized to render useful work—is initiated here.

2 METHODS

We evolved locomotive machines constructed from voxels with heterogeneous stiffness. Like many organisms [30], the robot's material stiffness progressively changes in response to mechanical loading incurred as the robot behaves. This ontogenetic change occurs independently at each voxel according to an evolved local rule.

Physical simulation.

The soft-matter physics engine *Voxelyze* [19] is used to calculate the movement of robots resulting from their interaction with a virtual 3D terrestrial environment. Each robot is simulated for 25 times the length of an expansion/contraction cycle (a total of five seconds). The displacement between the starting coordinates and the agent's final center of mass (in the *xy* plane) is recorded.¹

¹github.com/skriegman/2018-gecco contains the source code necessary for reproducing the results reported in this paper.

Interoceptive robustness through environment-mediated morphological development

Heavy materials.

Materials are simulated to have increased mass relative to those used by [8, 9, 18]. Constructed from heavier materials, many previously mobile robots become crushed under their own weight and require stiffer material to support locomotion with the same geometry. However, we also restricted actuation amplitude as materials grow stiffer to better approximate the properties of real materials with different stiffnesses. This creates an interesting and realistic trade-off: the stiffest material can easily support any body plan but cannot move on its own (like a skeleton without muscle), whereas the softest material can readily elicit forward movement in smaller bodies but cannot support many larger and potentially faster-moving body plans, such as those with narrow supporting limbs. Thus, a robot must carefully balance support with actuation.

Quad-CPPN encoding: $\mathbb{C}_1, \mathbb{C}_2, \mathbb{C}_3, \mathbb{C}_4$

Following [9], robot physiology is genetically encoded by a Compositional Pattern Producing Network (CPPN) [35], a scale-free mapping that biases search toward symmetrical and regular patterns which are known to facilitate locomotion.

Each point on a $10 \times 10 \times 10$ lattice is queried by its cartesian coordinates in 3D space and its radial distance from the lattice center. An evolved CPPN takes these coordinates as input and returns a single value which is used to set some property of that point in the workspace. We used four independent CPPNs to separately encode: Geometry, Stiffness, Development, and Actuation.

\mathbb{C}_1 Geometry.

The geometry of a robot is specified by a bitstring that indicates whether material is present (1) or absent (0) at each lattice point in the workspace, as dictated by \mathbb{C}_1 . The robot's geometric shape is taken to be the largest contiguous collection of present voxels.

\mathbb{C}_2 Stiffness.

Young's modulus is often used as an approximate measure of material stiffness. Robots are typically constructed of materials such as metals and hard plastics that have moduli in the order of $10^9 - 10^{12}$ pascals (Pa), whereas soft robots (and natural organisms) are often composed of materials with moduli in the order of $10^4 - 10^9$ Pa [31].

Here, voxels may have moduli in the range $10^4 - 10^{10}$ Pa. The robot's congenital stiffness is set at each voxel by \mathbb{C}_2 , but may be changed by development.

C₃ Development.

The robot's stiffness distribution k can change progressively during its lifetime t in response to localized engineering stress σ or pressure p. The rate of change α_i is specified at the i^{th} voxel by \mathbb{C}_3 , with possible values in 0 ± 10 . We compare three developmental variants.

None:
$$\frac{dk_i}{dt} = 0$$
 (1)

Stress: $\frac{dk_i}{dt} = \alpha_i \cdot \frac{d\sigma_i}{dt}$

Pressure:
$$\frac{dk_i}{dt} = \alpha_i \cdot \frac{dp_i}{dt}$$
 (3)

\mathbb{C}_4 Actuation.

Robots are 'controlled by' volumetric actuation: a sinusoidal expansion/contraction of each voxel with a maximum amplitude of 50% volumetric change.² However, linear damping ξ is implemented into the system such that the stiffest material does not actuate (Eq. 5). The phase difference ϕ_i of each voxel is determined by \mathbb{C}_4 , which offsets its oscillation relative to a central pattern generator.

Prior to actuation, each voxel has a resting length of one centimeter. This length is periodically varying (f = 5 Hz) by approximately 14.5% ($A \approx 0.145$ cm), but damped by ξ . The instantaneous length of the *i*th voxel is thus:

$$\psi_i(t) = 1 + A \cdot \sin(2\pi f t + \phi_i) \cdot \xi(k_i), \qquad (4)$$

where:

$$\xi(k_i) = \frac{k_{\max} - k_i}{k_{\max} - k_{\min}} \,. \tag{5}$$

Evolution.

We employed a standard evolutionary algorithm, Age-Fitness-Pareto Optimization [32], which uses the concept of Pareto dominance and an objective of age (in addition to fitness) intended to promote diversity among candidate designs.

We performed twenty independent evolutionary trials with different random seeds (1-20); in each trial, a population of 24 robots was evolved for five thousand generations. Every generation, the population is first doubled by creating modified copies of each individual in the population. The age of each individual is then incremented by one. Next, an additional random individual (with age zero) is injected into the population (which now consists of 49 robots). Finally, selection reduces the population down to its original size (24 robots) according to the two objectives of fitness (maximized) and age (minimized).

Mutations add/remove/alter a particular node/link of a CPPN. They are applied by first selecting which networks to mutate, with the possibility to select all four, and then choosing which operations to apply to each.

Hypothesis testing.

We use bootstrapping to construct hypothesis tests. All *P*-values are reported with Bonferroni correction for multiple (typically three) comparisons. We adopt the following convention: '***' for P < 0.001; '**' for P < 0.01; and '*' for P < 0.05.

3 RESULTS

Videos of all sixty run champions (pictured in Figs. 2, 3, 4) can be seen at **goo.gl/T5wZNQ**.

Evolvability.

We first investigated whether environment-mediated morphological development affected evolvability (Fig. 5A). At the termination of an evolutionary trial, we only consider the most fit individual—the run champion—from each of the twenty independent trials (Figs. 2, 3, 4). After correcting for three comparisons, there was not enough evidence to reject the null hypothesis—that there is no

(2)

 $^{^2{\}rm The}$ scare quotes are intended to highlight the fact that actuation policies no more dictate the movement of a robot than its geometry.

difference between adaptive and nonadaptive robots, in terms of mean displacement—at the 0.05 level.

These results could be taken to suggest one of the following. Either the search space is sufficiently smooth prior to development (actuation and support are not as antagonistic as envisaged), or the proposed developmental mechanism is an insufficient smoothing mechanism (successful ways to change stiffness as a linear function of stimuli are sparse in the search space).

Geometric diversity.

We investigate morphological diversity next by employing the Hausdorff distance d_H as a metric to compare the similarity of two robot geometries, A and B. For each voxel in A, the closest voxel in Bis identified, according to euclidean distance d. Similarly, for each voxel in B, the closest voxel in A is identified. The Hausdorff metric is the larger of these two distances. Formally,

$$d_H(A,B) = \max\{\sup_{a \in A} \inf_{b \in B} d(a,b), \sup_{b \in B} \inf_{a \in A} d(a,b)\}.$$
 (6)

Informally, two robots are close in the Hausdorff distance if every voxel of either robot is close to some voxel of the other robot.

We calculated the Hausdorff distance between each of the $\binom{20}{2}$ = 190 possible pairings of the 20 run champions (Fig. 5B). Because $d_H(A, B)$ depends on the orientations of *A* and *B*, we rotate *B* in the *xy* plane (0, 90, 180, and 270 degrees) and the *yz* plane (0 and 90 degrees), and select the rotation that creates the smallest $d_H(A, B)$.

We found the evolved body shapes of pressure-adaptive robots to be more diverse than those of stress-adaptive robots (P < 0.001). We did not find a significant difference, at the 0.05 level, between adaptive and nonadaptive treatments using this particular measure of morphological diversity.

Across all three treatments, there appear on visual inspection to be three types of geometries (Figs. 2-4): a Π robot with wide posterior and anterior legs; a Γ robot whose legs meet perpendicularly; and a Υ robot that connects a (mainly cylindrical) leg perpendicularly to the center of a 10 × 10 vertical plane. Depending on how one counts, the Υ species can be seen in at most one nonadaptive robot (Fig. 2, run 19), two stress-adaptive robots (Fig. 3, runs 7 and 16), and six pressure-adaptive robots (Fig. 4, runs 2, 6, 7, 9, 12, and 16). Pressure-adaptive robots have more diversity by virtue of more Υ robots.

Interoceptive robustness.

To investigate the relative robustness (if any) across the three treatments, in the following experiment, development was manually removed from the stress- and pressure-adaptive run champions. We then tested the sensitivity of the resulting reduced robots to their evolved congenital stiffness distribution (Fig. 5C). To do so, we replaced the evolved network dictating material stiffness, \mathbb{C}_2 , with a random number generator that draws from the same range of possible stiffness ($10^4 - 10^{10}$ Pa). That is to say, we 'built' the evolved run champions without any errors in the specifications of geometry and actuation, but completely ignored the evolved specifications of their material stiffness, replacing them instead with random noise. We then calculated the relative fitness

$$R = F_{\text{test}} / F_{\text{train}} , \qquad (7)$$

run 1 run 2 run 3 run 4 run 5 run 6 run 7 run 8 run 9 run 10 run 11 run 12 run 13 run 14 run 15 run 16 run 17 run 18 run 19 run 20 210 10⁹ 104 6¹⁰ 810

Non-developmental champions (Eq. 1)

Figure 2: Run champions colored by congenital stiffness, which ranges from 10^4 to 10^{10} Pa. After settling under gravity, robots move toward the right-hand side of the page.

Interoceptive robustness through environment-mediated morphological development

GECCO '18, July 15-19, 2018, Kyoto, Japan



Figure 3: Run champions colored by congenital stiffness which can change during operation (ontogeny) in response to *engineering stress*.

Figure 4: Run champions are colored by congenital stiffness which can change during operation (ontogeny) in response to *pressure*.



Figure 5: Means (with 95% C.I.) for various statistics of the run champions, at generation 5000: (A) Fitness as the final displacement of a robot, measured by voxel-length units; (B) Diversity as the pairwise Hausdorff distances of robot geometries; (C) Robustness as the relative fitness (testing fitness divided by training fitness) after development is removed and a random stiffness distribution is introduced into the champion's body (Eq. 7); (D) Mean, taken across the body, of relative lifetime change in stiffness, as a measure of the lack of canalization (Eq. 8; lower bars indicate more canalization); (E) Variance, taken across the body, of relative lifetime change in stiffness, as a measure of heterogeneity/nonuniformity in developmental reactions (Eq. 9); (F) Variance, taken across the body, of the coefficients/gain of developmental reactivity (Eq. 10).

where F_{train} is the fitness achieved using the evolved stiffness and F_{test} is the fitness when tested with a random stiffness distribution. We repeated this process ten times for each run champion, each time drawing a new random stiffness distribution.

We found that, compared to nonadaptive robots, reduced stressadaptive robots were more robust to this (extreme) discrepancy between training and testing stiffness distributions (P < 0.01). These results are consistent with the found correlation between development and robustness [4, 24, 28]. However, the results here indicate that this correlation is contingent on the kind of environmental signal the developing agent responds to: there was no difference between pressure-adaptive and nonadaptive robots in this regard, at the 0.05 level.

This implies that by behaving interoceptively with respect to engineering stress, robots evolved the ability to ameliorate large deviations from their expected material properties, but by behaving interoceptively with respect to pressure, robots did not evolve this character. Because development was manually removed beforehand, robustness in our case was not a matter of changing one's body, as in the example of plant growth [36]; rather, it is an intrinsic property of structure (geometries and actuation patterns) educed from ancestors who changed in response to one particular internal state (stress), but not from those who responded to another (pressure).

The difference in robustness between nonadaptive robots and stress-adaptive, but not pressure-adaptive robots, could be due in part to the fact that there are simply more pressure-adaptive Υ robots than stress-adaptive Υ robots. While Υ robots tend to be more fit than Π and Γ robots ($P_{\Pi} < 0.05$; $P_{\Gamma} < 0.05$), they also appear to exploit their material properties to a greater degree, and are thus more sensitive to changes in its constitution, compared to Π and Γ robots ($P_{\Pi} < 0.05$; $P_{\Gamma} < 0.01$).

The Υ robot generates movement by pushing off its posterior leg, which must be rigid enough to support itself as well as propel forward the center portion of its anterior wall (e.g. run 12 in Fig. 4). Interoceptive robustness through environment-mediated morphological development

The robot loses kinetic energy, which is stored as elastic strain energy in the spring-like voxels between the wall's center and edge. The most strain is present in the dorsal portion of the anterior wall. The springs recoil, restoring kinetic energy and generating forward motion. If the posterior leg is too soft, or the dorsal anterior wall too rigid, the Υ robot can suffer a large drop in performance.

Differences in geometry, however, shed no light on why (the reduced) stress-adaptive robots are more robust than nonadaptive robots: the level of significance (P < 0.01) does not change after removing the only nonadaptive robot that could possibly be classified as Υ (Fig. 2, run 19). Thus we continue our investigation by analyzing how stress and pressure might differentially affect the rate of developmental reactions.

Canalization.

One indication of canalization [24, 40] in our system is given by the magnitude of α_i in each voxel, as defined by Eqs. 2 and 3. However, this is but one of two necessary ingredients for a developmental reaction: it indicates bodywide responsiveness to *potential* stimuli, but ignores the *actual* stimulus.

Thus, as proxy for canalization, we measured the amount of morphological change in reaction to local stimulus, during evaluation. More precisely, we recorded the mean, across the body, of relative lifetime change in stiffness

$$M_{\text{body}} = \frac{1}{\#\gamma} \sum_{i \in \gamma} \left| k_i^+ / k_i^\circ - 1 \right|, \tag{8}$$

where k_i° is the congenital stiffness, k_i^+ is the final stiffness, and $\gamma = \{i : g_i = 1\}$ contains the coordinates *i* of each voxel g_i present in the (bit array) geometry which has cardinality $\#\gamma$ (total voxels). Less change—lower M_{body} —indicates more canalization.

On average, voxels in stress-adaptive robots change their relative stiffness less than voxels in pressure-adaptive robots (P < 0.001) (Fig. 5D). In other words, developmental reactions are canalized to a greater extent in stress-adaptive robots. It follows, then, that the treatment with increased robustness was also the treatment with increased canalization.

To get a sense of the consistency of developmental reactions, as they occur *across the body* of evolved robots, we also recorded the spatial variance of this relative lifetime change

$$V_{\text{body}} = \operatorname{Var}_{i \in \gamma} \left(\left| k_i^+ / k_i^\circ - 1 \right| \right).$$
(9)

By this measure, stress-adaptive robots exhibit more uniform reactions than pressure-adaptive robots (P < 0.05) (Fig. 5E).

Taken together, then, we may say that the developmental reactions of stress-adaptive robots are more uniform in space (lower V_{body} ; Fig. 5E), and more canalized in magnitude (lower M_{body} ; Fig. 5D) than those of pressure-adaptive robots. Pressure-adaptive robots therefore experience larger and more localized changes in stiffness during their lifetime.

There are two possibilities that could explain this more localized change in stiffness in the pressure-adaptive robots. One possibility is that there is greater variance among the α_i in the pressure-adaptive robots. The alternative is that there is greater variance in the application of pressure throughout the body. To test the first possibility, we first normalized α_i in pressure- and stress-adaptive robots by the differing ranges of α that evolved in the pressure-adaptive (-5.36)

to 5.63) and stress-adaptive (-10.00 to 6.42) robots. Then we took the variance of α_i across the body of each run champion, individually:

$$V_{\text{gain}} = \operatorname{Var}_{i \in \gamma} \tilde{\alpha}_i$$
, where $\tilde{\alpha}_i = \frac{\alpha_i - \alpha_{\min}}{\alpha_{\max} - \alpha_{\min}}$. (10)

We found no evidence to support the hypothesis that α_i in pressureadaptive robots vary more (or less) in space than those in stressadaptive robots (Fig. 5F).

Therefore, because there is no difference in the variation of α_i (Fig. 5F), and because α_i cannot change during operation (Eqs. 2 and 3), it follows that pressure was generally much more localized within the bodies of the pressure-adaptive robots than stress within the bodies of the stress-adaptive robots. In other words, the entire body plan encountered stress, but only a small portion of the body encountered appreciable pressure. (An example of this localized response to localized pressure can be seen in Fig. 1.) We hypothesize that this global spread of stress is the likely cause of increased robustness in the stress-adaptive robots (Fig. 5C).

4 **DISCUSSION**

Building systems that are robust in the face of changing environmental conditions is a grand challenge in robotics and AI. The brittleness of current systems is exemplified by the growing literature on adversarial examples [1, 29, 37], and the fact that almost all practical robots are confined to the perfectly flat floors they clean, or the hermetic factories built around their work. Robustness is not unknown in human-engineered systems, but it is relatively rare; in nature it is everywhere, and one of the reasons is that in nature organisms develop: They constantly change not just their cognitive architectures but the morphologies that contain them and mediate with the external world.

It has been shown for rigid robots [4] that morphological development can in some cases increase robustness since it exposes evolution to richer sensory information: the robot must maintain locomotion while changing its body. Soft robots have much greater potential in this domain: If soft, there are more ways that morphology can change, so by definition the increase in breadth in sensorimotor experiment induced by development will be even greater than that for developing yet rigid machines. Toward this goal, by allowing material stiffness to be plastic, we have here investigated a heretofore unexplored dimension of morphological change (stiffness) not available to rigid robots.

Advances in materials science and 3D printing promise new engineered systems—protean machines—that may continuously morph in response to changing environmental signals. Simply put, if a robot always changes its strategy along many morphological and neural modalities, it is more difficult to fool with a static adversarial example or a new task environment. Little to no analysis has been conducted, however, into how such systems should respond to environmental stimuli in order to adapt their functions in the face of changing environmental conditions.

In initiating such a study here, we have shown that it is not just a matter of reacting to *any* stimulus: different types of developmental feedback loops elicit different *evolved* properties. We observed that if one modality (stiffness) responds to one particular internal signal

(engineering stress) but not another (pressure), robots evolved structure that intrinsically buffered large deviations from their expected material properties.

Pressure and stress bear distinct mechanical load signatures which in turn stimulated very different developmental reactions. Intriguingly, increased robustness was correlated with increased canalization: developmental reactions with stress were canalized to a greater degree than those with pressure. Although developmental reactions with pressure did not afford the evolution of robustness here, it did increase evolutionary divergence: pressureadaptive robots evolved more diverse (congenital) shapes than stress-adaptive robots. Our work here suggests there may be other developmental feedback loops that could be made available to evolution that would lead to more diverse and robust robots.

For our purposes, 'morphology' is a robot body, but the concepts here could equally be applied to non-embodied systems, such as the architectures of deep neural networks [27, 41]. One could define internal neural processes such as node sharpening [16], Hebbian learning, or neurotransmitter diffusion [21, 38] as interoceptive signals to which the neural network developmentally responds in a structural manner, such as adding or removing neurons. Meanwhile, at a faster time scale, synaptic weights might be tuned in response to exteroceptive signals such as gradients of a loss function. Finally, such a network could be placed inside a robot which itself is experiencing morphological change.

5 ACKNOWLEDGMENTS

This work was supported by Army Research Office award W911NF-16-1-10304 and DARPA contract HR0011-18-2-0022. The computational resources provided by the UVM's Vermont Advanced Computing Core (VACC) are gratefully acknowledged.

REFERENCES

- Anish Athalye and Ilya Sutskever. 2017. Synthesizing robust adversarial examples. arXiv preprint arXiv:1707.07397 (2017).
- [2] J Mark Baldwin. 1896. A new factor in evolution. The american naturalist 30, 354 (1896), 441–451.
- [3] Josh Bongard, Victor Zykov, and Hod Lipson. 2006. Resilient machines through continuous self-modeling. Science 314, 5802 (2006), 1118–1121.
- [4] Josh C Bongard. 2011. Morphological change in machines accelerates the evolution of robust behavior. Proceedings of the National Academy of Sciences 108, 4 (2011), 1234–1239.
- [5] Josh C Bongard and Rolf Pfeifer. 2001. Repeated Structure and Dissociation of Genotypic and Phenotypic Complexity in Artificial Ontogeny. Proceedings of The Genetic and Evolutionary Computation Conference (GECCO 2001) (2001), 829–836.
- [6] Eric Brown, Nicholas Rodenberg, John Amend, Annan Mozeika, Erik Steltz, Mitchell R Zakin, Hod Lipson, and Heinrich M Jaeger. 2010. Universal robotic gripper based on the jamming of granular material. *Proceedings of the National Academy of Sciences* 107, 44 (2010), 18809–18814.
- [7] Jennifer Carlson and Robin R Murphy. 2005. How UGVs physically fail in the field. IEEE Transactions on robotics 21, 3 (2005), 423–437.
- [8] Nick Cheney, Jeff Clune, and Hod Lipson. 2014. Evolved electrophysiological soft robots. In ALIFE, Vol. 14. 222–229.
- [9] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. 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. ACM, 167–174.
- [10] Francesco Corucci, Nick Cheney, Sam Kriegman, Josh Bongard, and Cecilia Laschi. 2017. Evolutionary developmental soft robotics as a framework to study intelligence and adaptive behavior in animals and plants. *Frontiers in Robotics* and AI 4 (2017), 34.
- [11] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521 (2015), 503–507.
- [12] Frank Dellaert and Randall D Beer. 1996. A Developmental Model for the Evolution of Complete Autonomous Agents. In Proceedings of the fourth international

conference on simulation of adaptive behavior.

- [13] Daniel C Dennett. 2003. The Baldwin effect: A crane, not a skyhook. Evolution and learning: The Baldwin effect reconsidered (2003), 60–79.
- [14] René Doursat. 2009. Organically grown architectures: Creating decentralized, autonomous systems by embryomorphic engineering. In Organic computing. Springer, 167–199.
- [15] Peter Eggenberger. 1997. Evolving morphologies of simulated 3D organisms based on differential gene expression. Procs. of the Fourth European Conf. on Artificial Life (1997), 205–213.
- [16] Robert M French. 1994. Dynamically constraining connectionist networks to produce distributed, orthogonal representations to reduce catastrophic interference. *network* 1111 (1994), 00001.
- [17] Ian Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. arXiv preprint arXiv:1312.6211 (2013).
- [18] Jonathan Hiller and Hod Lipson. 2012. Automatic Design and Manufacture of Soft Robots. *IEEE Transactions on Robotics* 28, 2 (April 2012), 457–466.
- [19] Jonathan Hiller and Hod Lipson. 2014. Dynamic simulation of soft multimaterial 3d-printed objects. Soft Robotics 1, 1 (2014), 88–101.
- [20] Geoffrey E Hinton and Steven J Nowlan. 1987. How learning can guide evolution. Complex systems 1, 3 (1987), 495–502.
- [21] Phil Husbands, Tom Smith, Nick Jakobi, and Michael O'Shea. 1998. Better living through chemistry: Evolving GasNets for robot control. *Connection Science* 10, 3-4 (1998), 185–210.
- [22] Maciej Komosinski. 2003. The Framsticks system: versatile simulator of 3D agents and their evolution. *Kybernetes* 32, 1/2 (2003), 156–173.
- [23] Sylvain Koos, Jean-Baptiste Mouret, and Stéphane Doncieux. 2013. The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation* 17, 1 (2013), 122–145.
- [24] Sam Kriegman, Nick Cheney, and Josh Bongard. 2017. How morphological development can guide evolution. arXiv preprint arXiv:1711.07387 (2017).
- [25] Sam Kriegman, Nick Cheney, Francesco Corucci, and Josh C Bongard. 2017. A minimal developmental model can increase evolvability in soft robots. In Proceedings of the Genetic and Evolutionary Computation Conference. ACM, 131– 138.
- [26] Hod Lipson and Jordan B Pollack. 2000. Automatic design and manufacture of robotic lifeforms. *Nature* 406, 6799 (2000), 974–978.
- [27] Risto Miikkulainen, Jason Liang, Elliot Meyerson, Aditya Rawal, Dan Fink, Olivier Francon, Bala Raju, Arshak Navruzyan, Nigel Duffy, and Babak Hodjat. 2017. Evolving Deep Neural Networks. arXiv preprint arXiv:1703.00548 (2017).
- [28] Julian Francis Miller. 2004. Evolving a self-repairing, self-regulating, French flag organism. In *Genetic and Evolutionary Computation Conference*. Springer, 129–139.
- [29] Anh Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 427–436.
- [30] Christopher Ruff, Brigitte Holt, and Erik Trinkaus. 2006. Who's afraid of the big bad Wolff?: "Wolff's law" and bone functional adaptation. American journal of physical anthropology 129, 4 (2006), 484–498.
- [31] Daniela Rus and Michael T Tolley. 2015. Design, fabrication and control of soft robots. Nature 521, 7553 (2015), 467–475.
- [32] Michael Schmidt and Hod Lipson. 2011. Age-fitness pareto optimization. In Genetic Programming Theory and Practice VIII. Springer, 129–146.
- [33] Robert F Shepherd, Filip Ilievski, Wonjae Choi, Stephen A Morin, Adam A Stokes, Aaron D Mazzeo, Xin Chen, Michael Wang, and George M Whitesides. 2011. Multigait soft robot. Proceedings of the National Academy of Sciences 108, 51 (2011), 20400–20403.
- [34] Karl Sims. 1994. Evolving virtual creatures. In Proceedings of the 21st annual conference on Computer graphics and interactive techniques. ACM, 15–22.
- [35] Kenneth O. Stanley. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines* 8, 2 (2007), 131–162.
- [36] Sonia E Sultan. 2000. Phenotypic plasticity for plant development, function and life history. *Trends in plant science* 5, 12 (2000), 537–542.
- [37] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199 (2013).
- [38] Roby Velez and Jeff Clune. 2017. Diffusion-based neuromodulation can eliminate catastrophic forgetting in simple neural networks. *PloS one* 12, 11 (2017), e0187736.
- [39] Jeffrey Ventrella. 1998. Designing emergence in animated artificial life worlds. In Virtual Worlds. Springer, 143–155.
- [40] Conrad H Waddington. 1942. Canalization of development and the inheritance of acquired characters. *Nature* 150 (1942), 563–565.
- [41] Barret Zoph and Quoc V Le. 2016. Neural architecture search with reinforcement learning. arXiv preprint arXiv:1611.01578 (2016).
- [42] Victor Zykov, Efstathios Mytilinaios, Bryant Adams, and Hod Lipson. 2005. Robotics: Self-reproducing machines. *Nature* 435, 7039 (2005), 163–164.