

Evolved Electrophysiological Soft Robots

Nicholas Cheney¹, Jeff Clune², and Hod Lipson¹

¹ Creative Machines Lab, Cornell University, Ithaca, NY

² Evolving Artificial Intelligence Lab, University of Wyoming, Laramie, WY
nac93@cornell.edu, jeffclune@uwyo.edu, hod.lipson@cornell.edu

Abstract

The embodied cognition paradigm emphasizes that both bodies and brains combine to produce complex behaviors, in contrast to the traditional view that the only seat of intelligence is the brain. Despite recent excitement about embodied cognition, brains and bodies remain thought of, and implemented as, two separate entities that merely interface with one another to carry out their respective roles. Previous research co-evolving bodies and brains has simulated the physics of bodies that collect sensory information and pass that information on to disembodied neural networks, which then processes that information and return motor commands. Biological animals, in contrast, produce behavior through physically embedded control structures and a complex and continuous interplay between neural and mechanical forces. In addition to the electrical pulses flowing through the physical wiring of the nervous system, the heart elegantly combines control with actuation, as the physical properties of the tissue itself (or defects therein) determine the actuation of the organ. Inspired by these phenomena from cardiac electrophysiology (the study of the electrical properties of heart tissue), we introduce *electrophysiological robots*, whose behavior is dictated by electrical signals flowing through the tissue cells of soft robots. Here we describe these robots and how they are evolved. Videos and images of these robots reveal lifelike behaviors despite the added challenge of having physically embedded control structures. We also provide an initial experimental investigation into the impact of different implementation decisions, such as alternatives for sensing, actuation, and locations of central pattern generators. Overall, this paper provides a first step towards removing the chasm between bodies and brains to encourage further research into physically realistic embodied cognition.

Introduction and Background

The fields of evolutionary robotics and artificial life have seen a great deal of emphasis on embodied cognition in recent years [Cheney et al. (2013); Bongard (2013); Riffel et al. (2013); Auerbach and Bongard (2012); Hiller and Lipson (2012a); Lehman and Stanley (2011); Auerbach and Bongard (2010a,b); Pfeifer et al. (2007); Hornby et al. (2001); Lipson and Pollack (2000)]. There is even a paradigm called *embodied cognition*, which argues that the specifics of the embodiment (such as the morphology) are




Figure 1: Current flowing through an evolved creature. The legend for voltage within each cell (colors) is given in Fig. 3.

vital parts of the resulting behavior of the system: It argues that the co-evolutionary connection between body and brain is more deeply intertwined than the body simply acting as a minimal physical interface between the brain and the environment [Pfeifer and Bongard (2006)].

Recent work in evolutionary robotics has shown that complex behaviors can arise when co-evolving bodies and brains. At one end of the spectrum, Auerbach and Bongard (2010b) demonstrated the evolution of physical structures that had no joints or actuators, and evolved to cover the largest distance in a controlled fall due to gravity. While that work exemplifies the evolution of behavior emerging from morphology alone, it does not co-evolve any actuation or control. Auerbach and Bongard (2010a) then evolved the placement of CPG controlled rotational joints between cellular spheres, thus co-evolving morphology and control.

Cheney et al. (2013) evolved locomoting soft robots made of multiple different materials: two passive voxels of differing rigidity and two actuated voxel types that expanded cyclically via out-of-phase central pattern generators (CPGs). While this work added a variety of soft materials and a new type of actuation, the pairing of muscle types directly to a CPG again reflected a focus on evolving morphology rather than sophisticated neural control.

Many examples in the literature include the co-evolution of a robot morphology with an artificial neural network controller [Sims (1994); Lipson and Pollack (2000); Hornby et al. (2001); Lehman and Stanley (2011)]. These studies (and many more like them) involve what might be called “ghost” networks: artificial neural networks that provide control to the body, yet do not have any physical embodi-




Figure 2: An example of complex electrical wave propagation in cardiac modeling [Fenton et al. (2005)].

ment in the system they control. The state of input nodes to these networks is often set by sensors in the robot and output nodes typically signify behavioral outcomes in the actuators, but the computation is done supernaturally, disjoint from the body itself.

In the age of 3D printing, it is a realistic goal for robots to physically walk out of a printer. It is thus worthwhile to consider designing robots that can be physically realized: i.e., those whose controllers are accounted for by being physically woven into the design of the robot.

While the brains of animals are often a separate module within their bodies, animals also have central and peripheral nervous systems extending throughout their bodies. An extreme example of this is the octopus, which has as much as 90% of its neurons existing outside of its central nervous system [Zullo et al. (2009)]. The distributed and physical layout of the nervous system over space may contribute significantly to neural processing, as the delays and branching in axons (the basis for nerves) are suggested to serve computational functions [Segev and Schneidman (1999)].

Despite the prevalence of embodied, distributed circuitry in nearly all of animal life, the idea of an embodied nervous system has been absent from the field of evolutionary robotics. The sub-field called Evolvable Hardware evolves physical circuits for computer chips [Floreano and Mattiussi (2008)], but such work has not been applied to evolving the circuitry of artificial life organisms. We are unaware of work with virtual creatures that have physically embodied control systems (e.g. where neural circuitry physically runs throughout the body of the creature). We present the first such work in this paper.

We propose a very basic model of electrical signal propagation throughout the body of an evolved creature. This embodied controller is based on electrophysiology (specifically at large scales, such as cardiac electrophysiology, Fig. 2). Electrophysiology is the study of the electrical properties of biological cells and tissues [Hoffman et al. (1960)]. In this model, electrical pulses from a single centralized sinusoidal pacemaker (analogous to the sinoatrial node – the pacemaker in the heart [Brown (1982)]) are propagated through the

electrically conductive tissue of the creature. The location and patterning of this conductive tissue is described by an evolved Compositional Pattern Producing Network (CPPN) genome. Evolution controls the shape of the body and the electrical pathways within it, which both combine to determine the robot’s behavior.

The model involves conductive tissue cells that collect voltage from neighboring cells, causing an action potential (spike) if the collected voltage exceeds the cell’s firing threshold (Fig. 3). Once this threshold is crossed, the cell depolarizes, causing a voltage spike that excites neighboring cells. This voltage spike is followed by a refractory period, during which the cell is temporarily unable to be re-excited.

This model allows for the propagation of information through the body of the creature in the form of electrical signals. The structure of this flow is produced entirely by the topology of the creature and the state of each cell’s direct neighbors. In this sense, the model can be seen as a form of distributed information processing. One could draw similarities between this model and a 3D-grid of neurons, where each neuron receives inputs from, and has outputs to, its immediate neighbors. In this analogy, we are evolving where neurons should exist in the grid, what type of material the neuron is housed in, as well as the material type, if any, of grid locations that do not contain neurons.

The placement of material, which is under evolutionary control, directly determines the resultant behavior of the organism. Cells that actuate will contract and expand as they depolarize (much like the contraction of cardiac muscles), leading to the locomotion behavior of the creature. In order to control the signal flow throughout the creature, insulator cells are allowed, which are unable to accept and pass on the signal. Evolution can also choose not to fill a voxel with material. The morphology of the simulated robot and tissue type at each cell is determined by a CPPN genome.

This model examines the evolution of embodied cognition at a more detailed level of implementation than is typical in the literature – with embodied control circuitry resulting directly from the morphology of the individual creature. While this study only covers the classic problem of locomotion, it is a step towards truly physically embodied robots.

Methods

CPPN-NEAT

The evolutionary algorithm employed in this study is CPPN-NEAT. This algorithm has been previously described in detail (Stanley, 2007, 2006; Auerbach and Bongard, 2010b; Cheney et al., 2013), so it is only briefly described here.

A Compositional Pattern Producing Network (CPPN) [Stanley (2007)] is variation of an Artificial Neural Network (ANN) [McCulloch and Pitts (1943); Floreano and Mattiussi (2008)] where each node can have one of many mathematical functions as an activation function (e.g. sine, cosine, Gaussian, sigmoid, linear, square, or positive square root)




Figure 3: (top) A depiction of an action potential. Notice how the voltage is below the threshold until a stimulus event (such as a pacemaker or neighboring cell spike) pushes the voltage to the threshold value. Once this threshold is met, a voltage spike occurs via a process called depolarization. The cell excites its immediate neighbors during this process. Following the action potential, the cell enters a fixed length refractory period, during which it is physically unable to produce a new action potential. Finally, the cell returns to its resting state, able to start the process again when a new stimulus arrives. (bottom) These phases of the action potential cycle are mapped to the color code used to visualize the soft robots in Figs. 1, 4, and 5. Image licensed via Creative Commons.

instead of being limited to a sigmoid activation function. In CPPN-NEAT, a design space is discretized into individual locations (in this case a 3D space is discretized into a $10 \times 10 \times 10$ grid of voxels, for 1000 total voxels). The CPPN is queried once per voxel to determine the phenotypic state at that location (in this case, whether a voxel is present and, if so, the material type). The inputs to the CPPN network for each location are different: specifically, they include one input node for each dimension of the space (here, reporting the x , y , and z values of that location), as well an input that reports the distance (d) from the center to the location. The network also features output nodes for each material property. There are three in this study: one node specifies if a voxel exists at the queried location, the second decides if the material at that location is conductive, and the third decides whether or not the material is an actuated muscle (the latter two only matter if the voxel is present).

Conductive VoxCad

Fitness is evaluated in the VoxCad soft body simulator [Hiller and Lipson (2014)]. Its dynamics have previously allowed the evolution of complex and lifelike behaviors in soft robots, as it can simulate muscle contractions [Cheney et al. (2013)]. Further details about VoxCad can be found in Hiller and Lipson (2012b).

This work adds electrophysiology to VoxCad by adding a simple action-potential model, acting on the scale of a single voxel (analogous to a cell). Each voxel has an immediate membrane potential level (the difference between the electric potential inside and outside the cell), as well as a threshold membrane potential level. In an action-potential model (Fig. 3), a cell’s resting potential is below that of the threshold potential. When the membrane potential reaches its threshold value, the cell depolarizes, causing a spike in the cell’s membrane potential and voltage.

Following the depolarization, the cell hyperpolarizes, dropping the voltage and membrane potential below their original values, as the cell enters a refractory period. During this refractory period, the cell is unable to be depolarized again. In biological cells, the refractory period also consists of a relative refractory period when the cell is able to be depolarized, but only by unusually high voltage levels. For simplicity, we ignore this aspect in our model, and consider only the absolute refractory period, during which depolarization is disallowed. This refractory period means that the current is unable to flow backwards towards recently depolarized cells, causing the unidirectional propagation of action potentials in a wave across the cells.

A cell’s action potential (starting with the beginning of the depolarization phase in Fig. 3) triggers a sinusoidal expansion/contraction of that cell with a maximum amplitude of 39% linear expansion per voxel side.

A given cell may transmit current to any other cell that it is physically touching. In 3D, this rule means that up to 26 neighboring voxels (the “Moore neighborhood”) can be activated by a single voxel. The threshold potential of each cell is set such that it will be excited if, and only if, at least one of its neighboring cells undergoes an action potential that causes that cell’s voltage to spike. The time it takes a cell to excite its neighbor is half of its depolarization period. This delay in excitation means that the electric signal does not instantly activate all contiguously connected cells, but rather spreads outwards in a wave-like pattern of muscle actuation.

Cells may be of any of the following types: empty, conductive muscles, insulating muscles, conductive passive tissue, insulating passive tissue, or a pacemaker cell. Near the center point of the discretized $10 \times 10 \times 10$ design space, a lone pacemaker is placed (cell number 555 out of 1000). Analogous to the sinoatrial node in cardiac electrophysiology, this pacemaker node serves as the source of electric stimulation for the entire system. Insulating cells are similar to the cells explained above, except that they are unable to

accept current from neighboring cells and thus never reach their threshold potential or produce an action potential.

In this model, the refractory period lasts 5 times as long as the depolarization period. This means that at least 5 voxels must separate the leading edges of two serial action potential waves. Since the pacemaker is placed in the center of the $10 \times 10 \times 10$ space, approximately one wave of action potentials would exist at any given time in a setup with a uniform cube of entirely conductive material – where a wave of action potentials would propagate uninterrupted, with a new one starting around the time the first reaches the outer edge of the space. We chose this setup to encourage the evolution of static gaits, which can be more robust and transferable to reality than dynamic gaits [Belter et al. (2008)].

The length of the expansion/contraction period of each node is set equal to the refractory period, such that each cell is guaranteed to be fully returned to its original size before its next actuation cycle begins.

Task and Fitness Evaluation

Following Cheney et al. (2013), we evolve these electrophysiological robots for locomotion over flat ground. This simple task and environment make fitness evaluation easy. Despite its simplicity, the task is a classic problem in the field, and has been repeatedly shown to produce an array of complex morphologies and interesting behaviors [Cheney et al. (2013); Clune et al. (2009, 2011); Auerbach and Bongard (2014); Lehman and Stanley (2011)].

Each creature is simulated for 20 times the length of an expansion/contraction cycle. Its displacement between the starting coordinates and the creature’s final center of mass (in the xy plane) is recorded. In an effort to discourage designs that might excite as many cells as possible, and to encourage designs with sparse spindles of connectivity (similar to the peripheral nervous system), the distance traveled is multiplied by $1 - \frac{(\# \text{ of conductive cells})}{1000}$. Thus the fitness function incentivizes minimizing the amount of conductive tissue and maximizing the distance traveled. While a multi-objective technique may be ideal in finding the optimal tradeoff between these goals, we follow previous CPPN-NEAT research in using this single, multi-part fitness function [Cheney et al. (2013); Auerbach and Bongard (2009)].

Experimental Parameters

Unless otherwise noted, each treatment described below consists of 48 independent runs (with identical initial conditions across treatments). Each run consists of a population size of 30 individuals evolved for 1000 generations. Unless otherwise noted, all other parameters are consistent with Cheney et al. (2013).

Statistical Reporting

Because the data are not normally distributed, all plots show median fitness (thick, center lines) bracketed by two thin

lines that represent 95% bootstrapped confidence intervals of the median [Sokal and Rohlf (1995)]. For the same reason, all p -values are generated with the non-parametric Mann-Whitney-Wilcoxon Rank Sum test, which does not assume normality. Reported p -values compare the distance traveled by the top organism for each of the 48 runs at the final (1000th) generation. Plots report distance traveled, not adjusted fitness (which penalized for the number of conducting voxels as explained previously).

Results

Since this is the first study of evolved electrophysiological robots, there are many unanswered questions regarding the design and implementation of such a system. Many arbitrary design choices were made during the initial implementation. Here, we examine the impact of some of these choices.

As with many explorations in evolved virtual organisms, one of the main goals is complex, natural-appearing behavior. However, there are no satisfactory metrics for the “naturalness” or complexity of evolved behaviors. For this reason, we must rely on our qualitative, subjective assessments. A video of the evolved behaviors can be seen on the “Cornell Creative Machines Lab” Youtube channel, or found directly at this link: <http://goo.gl/CvJp41>. We believe the behaviors are interesting, complex, and lifelike – at least as much as in Cheney et al. (2013) – despite the added challenges of evolving physically embedded control.

We observed that physically instantiated control circuitry can produce both predictable and chaotic behaviors. Fig. 4 shows a simple wave of action potentials propagating outwards from the center of the creature, with little interruption. Fig. 5 reveals the evolution of unpredictable physical dynamics that still produce functional behavior. Notice the multiple “inputs” to a potential self-sustaining circular pathway. Fig. 1 demonstrates a circular actuation pattern of intermediate complexity, due more so to changes in the robot’s shape than to material differences within it. We now turn to more quantitative analyses.

Pacemaker Placement

The placement for the pacemaker was an arbitrary decision made during the design of this new system. In an effort to mimic the midline location of the central nervous system in biology, the pacemaker was placed in the middle of the design space from which the creature was built. Thus action potential waves could propagate out equally in all directions and were not biased in any particular direction of travel. In order to test the effect of this arbitrary choice, a treatment was also performed where the pacemaker was located at the center voxel of the roof of the $10 \times 10 \times 10$ design space – voxel number 955 (where indices increase from the bottom, left hand, nearest corner), as well as a treatment that placed the pacemaker in the top right corner – voxel 999.




Figure 4: An action potential wave propagating across a mostly homogeneous surface. (*left, single robot*): The robot has a large patch of continuous conductive muscle on its back. In this pre-simulation state, cell colors signify the following: orange cells are conductive, blue cells are non-conductive; dark colors (blue or orange) signify muscle cells, while lighter colors (blue or orange) signify non-actuatable cell tissue; the red cell at the bottom is the robot’s pacemaker cell. (*right, 3 × 4 grid of robots*): A progression over time (left to right, top to bottom) shows the wave-like propagation of the action potential phases (color meanings are described in Fig. 3). Note how the action potential emerges from the center, stimulated by the wave propagating out through the conductive tissue from the pacemaker below it. Following the light blue depolarization, the yellow and red phases show the longer lag of the refractory period, following in exactly the same pattern made by the leading edge of the action potential wave. As the wave fully passes, the cells return to their dark resting state and are thus able to spike again with a new action potential when the next wave comes.

As shown in Fig 6, the placement of this pacemaker significantly affects performance. While a central location (baseline treatment) shows significant advantages compared to the top-center and top-corner pacemaker locations ($p = 4.91 \times 10^{-11}$ and 7.16×10^{-16} , respectively), a statistically significant difference is also demonstrated between the two less-different treatments: the top-center location outperformed the top-corner location ($p = 3.43 \times 10^{-4}$). These results show that the pacemaker location can have a clear effect on the evolved behaviors. Future work shall place the exact location under evolutionary optimization.

Speed of Pacemaker

Another implementation decision was the low-frequency pacemaker to allow for static gaits. The increased stability and robustness of static gaits is appealing, and this may allow better transferability to physical robots (Belter et al.,




Figure 5: A more complex electrophysiological robot. (*left robot*): Contrary to Fig. 4, this creature shows complex patterning of the orange conductive tissue within the insulating blue tissue. (*right three robots*): As they unfold over time (left to right), the action potential waves in this robot produce a highly fractured, counter-intuitive actuation pattern that involves electrical signals flowing through long, sparse connective corridors and around corners (an explanation of the colors is provided in Figs. 3 and 4). The result is an interesting and unexpected behavioral pattern wherein the creature mashes and spins the left side of its body, which is separated from the larger, right side of its body by a large, oddly shaped internal cavity. Despite this bizarre behavior, it effectively locomotes. This behavior and others can be viewed on Youtube at: <http://goo.gl/CvJp41>.

2008). However, animals often employ dynamic gaits when there is an incentive for speed (as there is here). The tradeoff between these two is not known in this system. To examine this tradeoff, we compared three different treatments. First, the baseline treatment includes a pacemaker with the relatively slow pace of 40 beats per second (BPS). Since the baseline evaluation period is half a second, this results in 20 electrical pulses from the pacemaker per trial. A second treatment explores the increased potential for dynamic gaits at the maximum pacemaker speed of 80 BPS (the limit is due to the fixed length of the refractory period). In this faster treatment, each individual cell contracts at the same rate as before, but the pacemaker is now exciting cells as soon as their refractory period ends, instead of waiting (the length of an additional actuation cycle) before sending another pulse into the system. This system uses twice the amount of energy, producing 40 action potential waves in the same half second. In a third treatment, the faster paced (80 BPS) pacemaker is evaluated for half its normal time length, resulting in 20 beats per evaluation. This treatment allows a fair comparison of pacemakers in terms of distance traveled per “beat”, rather than per unit time.

Unsurprisingly, the faster pacemaker evaluated for the full half second outperforms both the slower pacemaker evaluated for the same time period and the faster pacemaker evaluated for the shorter evaluation time ($p < 10^{-16}$ for both, Fig. 7). Interestingly, the frequency of the pacemaker has no significant effect on the distance traveled ($p = 0.51$ at generation 1000), suggesting that any disparity between the faster and slower gaits was not realized in simulation (with the number of beats held constant). Testing this result in the transfer to physical robots is a subject for future work.




Figure 6: The effect of the placement of the central pattern generator (CPG) on the evolved speed. In one treatment, the CPG is placed at the top corner of the $10 \times 10 \times 10$ design space (voxel 999). This treatment performs slightly, but significantly ($p = 3.43 \times 10^{-4}$), better than another treatment that places the CPG at the center of the top plane of the bounding box (voxel 955). Outperforming both of these ($p < 4.91 \times 10^{-11}$) is the baseline treatment in which the CPG is always placed as close to the center of the bounding box as possible.

Touch Sensors

Another implementation decision was the use of pacemakers as the primary drivers of the system. While pacemakers, also known as central pattern generators, are biologically motivated [Ijspeert et al. (2007)], we could instead ask evolution to generate its own cadence. To provide an alternative to the pacemaker, we tested a treatment with touch sensors in lieu of a steady internal signal.

The touch sensors, like the pacemaker, are capable of producing an electrical signal. However, they do so in response to contact with the ground, rather than in a regular rhythm. In this treatment, all conductive cells have this touch-sensing ability and produce an action potential when in contact with the ground if not in the refractory period. Thus waves of action potentials propagate outwards from the touch sensors only when they are both in contact with the ground and fully recovered from their prior depolarization.

Thus, the upper bound on the number of action potentials that the touch sensors could produce is that of an 80 BPS pacemaker (the 80 BPS pacemaker fires again as soon as exiting the refractory period, where the touch sensors do so only if also touching the ground at that time – the slower 40 BPS pacemaker waits the length of one cycle before firing again). To reach this upper bound, touch sensors would have to be touching the ground exactly at the time when they com-

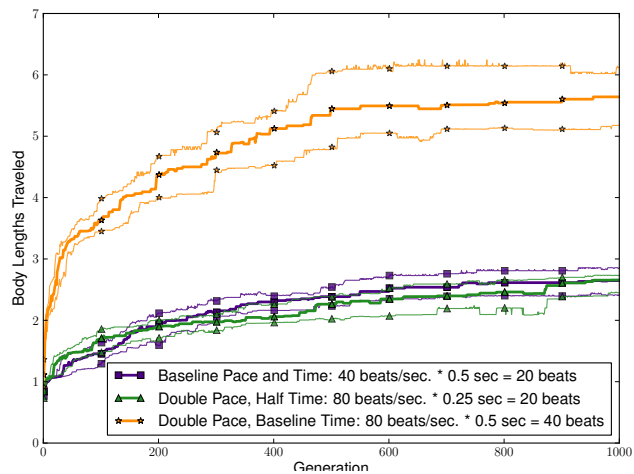


Figure 7: The effect of faster pacemakers (CPGs). It is not surprising that a faster CPG (80 beats per second) travels farther when evaluated for longer, or when compared to a slower CPG ($p < 10^{-16}$). However, when the comparison is made according to distance per beat (half time/full speed against half speed/full time – both producing a total of 20 beats) there is no difference in their performance at Generation 1000 (p -value = 0.51), suggesting that CPG speed does not greatly affect evolved locomotion speed.

pleted their refractory period, and thus it is likely that this ceiling would not be reached in all cases. For a comparison, Fig. 8 shows the median distance traveled over evolutionary time plotted against that of the slower pacemaker (40 BPS) and the faster pacemaker (80 BPS) described above, and evaluated for the baseline half-second evaluation time. It is not surprising that the slower pacemaker falls behind the pack here, as it is handicapped by a throttle on its only source of action potentials compared to the faster pacemaker and the touch sensors ($p < 10^{-16}$). The tighter race is between the touch sensor and the faster pacemaker. In the early stages (< 150 generations), the robots with touch sensors significantly outperform robots with a pacemaker. However, in the later stages of evolutionary optimization, the touch sensor treatment shows modest gains compared to the continued innovation of the pacemaker treatment, with the pacemaker treatment significantly outperforming it at the end of the run ($p = 1.27 \times 10^{-7}$). The relatively low level of improvement in the touch sensor treatment in the later stages of evolution may suggest the premature convergence on local optima. The multiple points of origin for action potential waves, and thus wave collisions, may have also had an effect. An additional issue that could have hindered performance in this treatment is the upward propagation of signals from touch sensors on the ground, versus outwards expanding waves from the center of the organism.

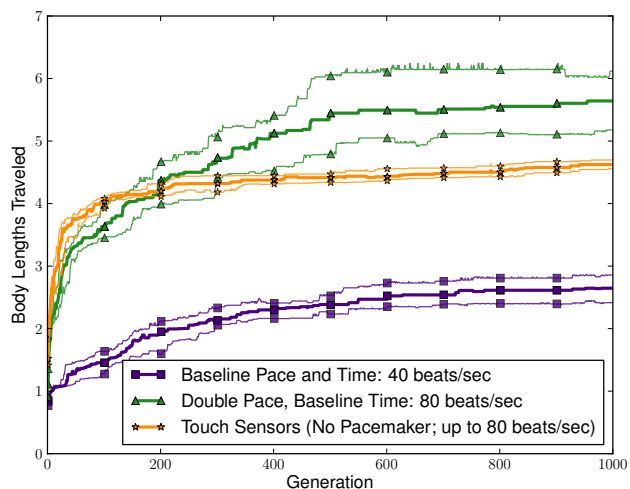


Figure 8: The performance of touch sensors vs. central pattern generators. The touch sensor treatment produces an expected number of beats with the upper bound set by the faster (80 beat/sec) CPG. Despite early evolvability leading to a statistically significant advantage in the first 150 generations, in later generations the touch sensor setup is unable to produce creatures that travel as far as the faster CPG setup ($p = 1.27 \times 10^{-7}$ at Gen. 1000). Artificially throttled, the slow CPG is unable to compete with either ($p < 10^{-16}$).

Expansion/Contraction Cycle

In the soft robot evolution system described by Cheney et al. (2013), regular, quickly repeating, and coupled out-of-phase sinusoidal action cycles defined the expansion and contraction of cells. In this model, which does not feature the same complimentary muscle types, the question of actuation cycle is not entirely clear. In an attempt to explore this, here we test the effectiveness of contraction-then-expansion phase cycles against expansion-then-contraction cycles (Fig. 9). These treatments take place on the baseline (slow) pacemaker setup, as to not allow continuous and quickly repeating expansion/contraction cycles, but rather to have a break between actuations. Despite the same number of beats (and thus the same amount of overall expansion and contraction) in both setups, the contraction-then-expansion setup performs significantly better ($p = 1.94 \times 10^{-3}$). While the reason for this difference is not entirely clear, it may be due, in part, to a larger continuous expansion period from the trough of the sine wave to its peak (continuous expansion from minimum to maximum size) in the contraction-then-expansion treatment. In contrast, the expansion-then-contraction setup includes a full-cycle length pause in the middle of its expansionary period. This explanation would suggest that more locomotion tends to come from pushing than pulling, which is in line with our observations from viewing videos of the evolved behaviors.

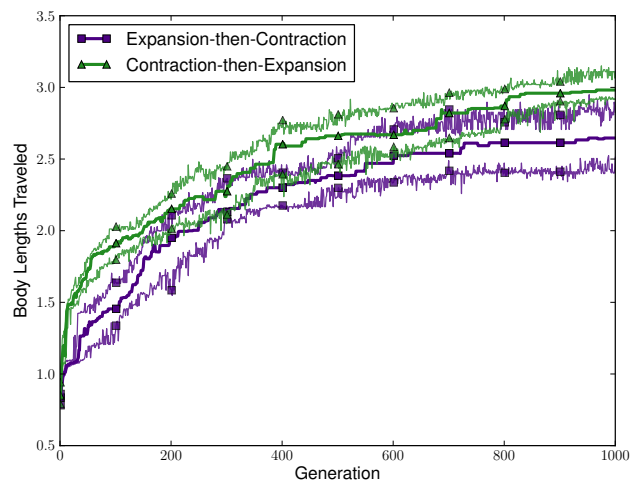


Figure 9: Unlike the regularly occurring actuation cycles of Cheney et al. (2013), the electrophysiological actuations in this paper do not have a necessary order: either expansion or contraction can occur first. It turns out that performance is significantly higher when muscles contract first and then expand, rather than vice versa ($p = 1.94 \times 10^{-3}$).

Discussion

This work reduces the separation between bodies and brains in research into embodied cognition. We did so by embedding the control systems into the physical simulation of the robot’s morphology. Perhaps most interesting about this work is that the complex and interesting behaviors are the direct result of the morphology of the creatures, as the control is woven directly into the structure of the organisms. In this work the size of the creatures was limited for computational reasons, but in future work we plan to explore larger design spaces. We also plan to test different ways of implementing electrophysiological robots and to challenge them to perform more difficult tasks.

Conclusion

We have introduced electrophysiological robots, which are inspired by the electrical properties of cardiac tissue. The behavior of these robots is governed by electrical signals flowing through the evolved cells of soft robots. We described these robots and how they are evolved, including the evolution of interesting behaviors, despite the added challenge of physically embedded control structures. We also provided an initial experimental investigation into different implementation decisions, such as alternatives for sensing, actuation, and central pattern generator locations. We believe that this paper provides a first step towards removing the gulf between brains and bodies to encourage further research into physically realistic embodied cognition.

Acknowledgments

This work was supported by DARPA Open Manufacturing Grant W911NF-12-1-0449 and NASA Space Technology Research Fellowship NNX13AL37H for Nicholas Cheney.

References

- Auerbach, J. and Bongard, J. C. (2009). How robot morphology and training order affect the learning of multiple behaviors. In *IEEE Congress on Evol. Comp., 2009*, pages 39–46. IEEE.
- Auerbach, J. E. and Bongard, J. C. (2010a). Dynamic resolution in the co-evolution of morphology and control. In *ALife XII*.
- Auerbach, J. E. and Bongard, J. C. (2010b). Evolving CPPNs to grow three-dimensional physical structures. In *Proc. of the Genetic & Evolutionary Comp. Conf.*, pages 627–634. ACM.
- Auerbach, J. E. and Bongard, J. C. (2012). On the relationship between environmental and morphological complexity in evolved robots. In *Proc. of Genetic & Evol. Comp. Conf.*, pages 521–8. ACM.
- Auerbach, J. E. and Bongard, J. C. (2014). Environmental influence on the evolution of morphological complexity in machines. *PLoS computational biology*, 10(1):e1003399.
- Belter, D., Kasinski, A., and Skrzypczynski, P. (2008). Evolving feasible gaits for a hexapod robot by reducing the space of possible solutions. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 2673–2678. IEEE.
- Bongard, J. C. (2013). Evolutionary robotics. *Comm. of the ACM*, 56(8):74–83.
- Brown, H. F. (1982). Electrophysiology of the sinoatrial node. *Physiological Reviews*, 62(2):505–530.
- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proc. of the 15th Genetic and Evol. Comp. Conf.*, pages 167–174. ACM.
- Clune, J., Beckmann, B., Ofria, C., and Pennock, R. (2009). Evolving coordinated quadruped gaits with the HyperNEAT generative encoding. In *Proc. of the IEEE Congress on Evol. Comp.*, pages 2764–71.
- Clune, J., Stanley, K. O., Pennock, R. T., and Ofria, C. (2011). On the performance of indirect encoding across the continuum of regularity. *IEEE Trans. on Evol. Comp.*, 15(4):346–67.
- Fenton, F. H., Cherry, E. M., Karma, A., and Rappel, W.-J. (2005). Modeling wave propagation in realistic heart geometries using the phase-field method. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 15(1):013502.
- Floreano, D. and Mattiussi, C. (2008). *Bio-inspired artificial intelligence: theories, methods, and technologies*. MIT Press.
- 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. (2012a). Automatic design and manufacture of soft robots. *IEEE Trans. on Rob.*, 28(2):457–466.
- Hiller, J. D. and Lipson, H. (2012b). Dynamic simulation of soft heterogeneous objects. *ArXiv:1212.2845*.
- Hoffman, B. F., Cranefield, P. F., and Johnston, F. D. (1960). Electrophysiology of the heart.
- Hornby, G. S., Pollack, J. B., et al. (2001). Body-brain co-evolution using l-systems as a generative encoding. In *Proc. of the Genetic and Evolutionary Comp. Conf.*, pages 868–875.
- Ijspeert, A. J., Crespi, A., Ryczko, D., and Cabelguen, J. M. (2007). From swimming to walking with a salamander robot driven by a spinal cord model. *Science*, 315(5817):1416–1420.
- Lehman, J. and Stanley, K. O. (2011). Evolving a diversity of virtual creatures through novelty search and local comp. In *Proc. of 13th Genetic & Evol. Comp. Conf*, pages 211–8. ACM.
- Lipson, H. and Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974–978.
- McCulloch, W. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of mathematical biology*, 5(4):115–133.
- Pfeifer, R. and Bongard, J. C. (2006). *How the body shapes the way we think: a new view of intelligence*. MIT press.
- Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *science*, 318(5853):1088–1093.
- Rieffel, J., Knox, D., Smith, S., and Trimmer, B. (2013). Growing and evolving soft robots. *Artificial Life*, (Early Access):1–20.
- Segev, I. and Schneidman, E. (1999). Axons as computing devices: basic insights gained from models. *Journal of Physiology-Paris*, 93(4):263–270.
- Sims, K. (1994). Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372.
- Sokal, R. and Rohlf, F. (1995). *Biometry: the principles and practice of statistics in biological research*. WH Freeman.
- 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.
- Zullo, L., Sumbre, G., Agnisola, C., Flash, T., and Hochner, B. (2009). Nonsomatotopic organization of the higher motor centers in octopus. *Current Biology*, 19(19):1632–1636.