Theoretical Computer Science 633 (2016) 19–27

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Topological evolution for embodied cellular automata

Nick Cheney *, Hod Lipson

239 Upson Hall, Cornell University, Ithaca, NY, USA

ARTICLE INFO

Article history:
Received 18 January 2015
Received in revised form 15 May 2015
Accepted 8 June 2015
Available online 11 June 2015

Keywords:
Evolutionary computation
Artificial life
Non-uniform
Cellular automata
Embodied cognition
Topological optimization
Complex systems

ABSTRACT

In this work we introduce a novel method for creating behaviors in cellular automata: optimizing the topology of the cellular substrate while maintaining a single simple update rule. We study the effect of altering the shape of a 3D cellular automaton and local signaling ability of each of its cells on the ability of that automaton as a whole to give rise to emergent locomotion behavior. This system optimizes for the physically embodied interactions between a cellular automaton with an external physically simulated world, rather than optimizing directly for a computational ability internal to the automaton itself. We give each cell in the automaton the ability to have an internal "excited" state, and also the ability to perform a physical action (volumetric contraction and expansion) as a result of that state. We then employ an evolutionary algorithm to optimize for the locomotion ability of the "robot" resulting from the behavior of this embodied automaton. We demonstrate a number of diverse topologies which lead to effective locomotion behaviors in this paradigm. We believe that creating complex behavior from simple rules in a complex substrate not only opens up questions about cellular automata, but also provides insights towards the study of morphological computation and embodied cognition.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Cellular automata are perhaps the quintessential example of emergent complexity (where we define emergence as the complex collective behavior that arises from the interaction of simple individual components [12]). From Wolfram's elementary cellular automata [30] to Conway's game of life [11], these structures take simple rule sets and iterate them over time and space to create fantastically complex behaviors. In elementary cellular automata, (2 states, size 3 neighborhood), the variety of behaviors produced are the result of the 256 (2^{2^3}) possible rules. These rules are typically applied to a homogeneous cellular substrate to produce the patterns we are accustomed to seeing.

In this work, we seek to explore an alternative method of creating variation in these systems. Rather than exploring the behaviors created by a number of different update rules, we take a single simple rule and vary the topology of the substrate upon which it is applied. The simple rule we chose was that of reaction-diffusion. In the binary-state case, the diffusion-like rule we use is the spacial propagation of excitation through the medium. This means that a cell is activated if and only if at least one of its immediate neighbors was active in the previous time step. Once a cell is excited, it is activated for a small period of time, during which it can excite its neighbors. Following this, the cell experiences a refractory period [19], during which it cannot be re-excited. We take the non-linearities produced by this rule to make up the reaction aspect of

* Corresponding author.
E-mail addresses: nac93@cornell.edu (N. Cheney), hod.lipson@cornell.edu (H. Lipson).

http://dx.doi.org/10.1016/j.tcs.2015.06.024
0304-3975/© 2015 Elsevier B.V. All rights reserved.

the rule set. This action potential setup creates waves of excitation (rather than allowing for back-excitation, and thus quick homogenization of the activity across the substrate), and is consistent with excitation in natural systems.

Natural instances of this wave propagation technique through excitable tissue suggest the possibility for complex emergent behaviors from this simple model. The human heart is an example of a natural system which uses the diffusion of electrical signals through an excitable muscle to produce behavior. Unlike skeletal muscle, cardiac muscle has gap junctions, which allow the transfer of signals in this manner. Typically this system is well behaved and electric signals from the sinoatrial node (pacemaker) cause the heart to beat once, then the excitation quickly dissipates. However, in some diseases these signals form chaotic and self-sustaining patterns (causing cardiac arrhythmia) [29], opening up the possibility of rich interactions and complex behavior resulting from this rule set.

In the case of cardiac arrhythmia, it is clearly beneficial to remove the complexity of this behavior. In order to do this, cardiologists may use medication to alter the functional (electrical) properties of the tissue [13]. Though often it is necessary to use catheter ablations to change the structural topology of the heart in order to significantly affect the dynamics [6]. This work draws inspiration from the later method – optimizing the topology of our cellular substrate, rather than altering the way that activation propagates over space and time. In this example, attempting to change the tissue properties would be analogous to modifying the update rules to change the behavior of the system. By employing a substrate which can range in behavior from orderly to chaotic, we hope to find an optimal computational abilities within this spectrum, and perhaps around a the phase transition "at the edge of chaos" [18].

Furthermore, we introduce a novel method of evaluating the behavior of a cellular automaton. We situate the automaton within a physical (in this case simulated) environment. By doing so we are able to produce an actionable physical behavior from the signal information propagation through the cellular automaton. We treat each cell as a muscle cell, which volumetrically contracts and expands upon each stimulation. This actionable behavior in the form of physical movement also defines the "fitness" of a given substrate topology, allowing up to optimize these topologies for maximal displacement using evolutionary computation, as is done in [8]. This differs from previous studies which have optimized these automaton for their internal computational abilities only [26,18,20], as we optimize the automata to produce an emergent behavior in a physical simulator, where the behavior is a result of the automata's interaction with an outside environment. This embodied evaluation method is common within evolutionary robotics, though we are not aware of an evolved robot which fits our definition of a cellular automaton (a grid of cells, each with number of possible states and whose current state is a function of the states of itself and its neighbors).

Upon examining this premise closer, one may realize that this is an exercise in embodied cognition [2], which states that the body plays a fundamental role in the way which the mind processes information and functions. In this extreme case, the mind of the robot does not exist on it's own, but rather information processing is distributed throughout and within the morphology of the creature. In this sense, the way in which creatures in our system processes information is determined solely through the topology of their body. Given a static environment, one cannot produce changes in the control or behavior of the creature without changing its morphology. By extension of this idea, we imagine this work as an exploration of low-level embodied information processing and morphological computation within the context of robotic design and control.

2. Background

There are, of course, exceptions to the previous notion of cellular automata as homogeneous substrates. Sipper explored non-uniform cellular automata [24–26], looking at one dimensional automata in which each cell could take on a different update rule. This idea of choosing different rule sets at each cell within the automata certainly extends the computational abilities of the automata, and represents a superset of the behaviors created in this paper (if one were to consider the empty cells outside of our optimized topologies as those with rules always leading to the quiescent inactive state). However, as the number of rule sets which must be simultaneously optimized increases, the ability to simultaneously evolve global behavior becomes nearly intractable. To cope with this, Sipper considers a decomposable task (density classification of the initial state) and applies optimization in a completely local manner, receiving performance feedback on the cell and neighborhood levels rather than on the collective behavior of the automata as a whole. In this work, the ability to create richness from just a single rule set helps to keep the problem tractable and (along with our indirect encoding: Section 3.2) allows us to perform optimization on truly global emergent phenomenon, such as the locomotion ability of the entire automata.

Sipper also explored systems in which the automaton could evolve a connection "architecture" in which cells were able to access information from non-neighboring cells [27]. The allowance of information from non-neighboring cells violates the traditional definition of a cellular automaton, in which local interactions collectively produce an emergent global behavior. Though if this assumption is relaxed, it is not surprising that his results suggest allowing such informational "shortcuts" can lead to more efficient computation. Our reaction-diffusion update rule keeps the criteria of local interactions intact.

The idea of using an update rules based on an action potential model is also not novel. Chua and Yang proposed a theory of Cellular Neural Networks, in which a set of neural-network-like nodes are placed in a regular grid, only connected to their immediate neighbors [9]. Gers et al. also demonstrated a cellular automata based neuron model, in which axon and dendrite cells passed signals between spiking neuron cells within a cellular automata. The growth rules of these axons and dendrites were optimized with an evolutionary algorithm to create a substrate that would perform various computations as the signal propagated across the automata [15]. The cardiac modeling work noted above employs a similar type of

cellular spiking behavior as well. However, we believe that we are unique in our use of these update rules for their ability to directly control an embodied physical system who's varied morphology directly interacts with an external (simulated) physical world.

Previous works have also considered cellular automaton to represent swarms of simple embodied robots [3,5], but our work differs from those by using the automata to describe the complex morphology of a single robot, rather than the organization of a swarm of many simple robots.

3. Methods

3.1. Cellular substrate

The cellular automaton upon which this process occurs is a three dimensional grid, which cells of uniform size. Two cells within this grid are said to be connected if they share any side of corner (3-D Moore neighborhood). The discrete states of each cell are: susceptible, excited, and refractory (one could find this similar to a distributed SIRS model, if such an analogy is helpful [4], where individuals are only contagious in the excited stage). In these experiments, the size of the substrate is limited to $10 \times 10 \times 10$ cells for computational efficiency.

3.1.1. Physical simulation

To produce an actionable "robot" from this cellular automaton, we employ the open-source soft-body-physics simulator VoxCad [16]. This simulator takes each of the voxels in our 3D grid, and treats them as cubic cells of varying stiffness, with each cell physically attached to it's immediate neighbors. When a cell is excited (Section 3.1.2), it volumetrically expands then contracts sinusoidally. This distributed actuation collectively causes the robot as a whole to move [8]. In addition to these "muscle" cells which expand and contract, passive soft support tissues and stiffer "bone-like" voxels may also make up the composition of these robots.

3.1.2. Activation propagation

In order to coordinate the behavior of these cells to produce meaningful collective action, this system also requires a control paradigm. In previous work with this substrate, the control paradigm was to have all voxels of a particular material contract and expand synchronously [8]. While this was sufficient to produce some interesting behaviors, it clearly limits the possible behaviors and especially the amount of computation that can be done. To extend the synchronous paradigm towards more flexible control, each voxel cell was also given an "electric" state [7]. Thus, the cells can now transfer signals locally to one another. In this system, the transfer of these signals may be independent or correlated to the cell's mechanical behavior.

In order to allow these signals to propagate far across the substrate, yet also interact in non-linear ways, we use a reaction-diffusion model, similar to the FitzHugh–Nagumo reaction–diffusion model [14,21]. This particular reaction-diffusion model was chosen partially because of its mix between simplicity and complexity. It is able to produce stand-alone propagating waves from a single exciton source, yet can quickly spiral into self-sustaining chaotic regimes with as little as two sources of excitation (or one excitation source and certain initial conditions). The choice of this model also allows us to draw corollaries with cardiac dynamics, as that natural system features signal propagation over an excitable medium, and is often modeled via the FitzHugh–Nagumo reaction–diffusion equations [1].

This particular instantiation features agent based modeling at the voxel cell level. Each cell begins in a state susceptible to excitation. A cell is able to become excited through the transfer of "electric potential" from an already excited neighbor. Once excited, the cell quickly depolarizes (the "reaction") and repolarizes, spreading its electric potential to all of its excitable neighboring cells (the "diffusion"). Once a cell has finished repolarization, it takes on a refractory state (during which it cannot be re-excited) before returning to the susceptible resting state (Fig. 1). This action potential model produces the directional waves of excitation over the substrate, as well as providing the richness for computation.

Since our goal is to be able to define the flow of information through this substrate by changing the substrate itself, we also allow each voxel cell to be conductive or insulating. A cell which is an insulator is unable to accept any of its neighbor's "electrical" signals, and thus cannot become excited (and cannot intrinsically expand or contract). By placing insulating cells throughout the creature, the system is able the define how and where excitation (information) is sent or used throughout the substrate.

Similarly, not all cells which are conductors need be muscle cells capable of expanding and contracting when excited. Passive tissue cells may also be conductive, and pass information from one part of the creature's body to another, without producing any physical action. In this sense the creatures are capable of modifying their internal state, separately from their behavior. Of course to perform an action with that information, the signal must eventually be passed to a conductive muscle cell, which will turn the signal into an action in the form of a cell contraction/expansion.

In this work, the original source of the excitation comes from a pacemaker cell, typically located at the center of the creature. This cell self-excites at a regular interval, mimicking the pacemaker (sinoatrial node) of the human heart. Other methods of self-excitation were explored (such as touch-sensitive sensors that would excite a voxel if it made contact with the ground), but these were not found to provide significantly more interesting behavior in initial studies (as these often excited at a regular pace with the stepping of the creature), and thus will not be the focus of this paper. The notion of a sinusoidal oscillatory input to controller has been used previously in evolutionary robotics [10,17].

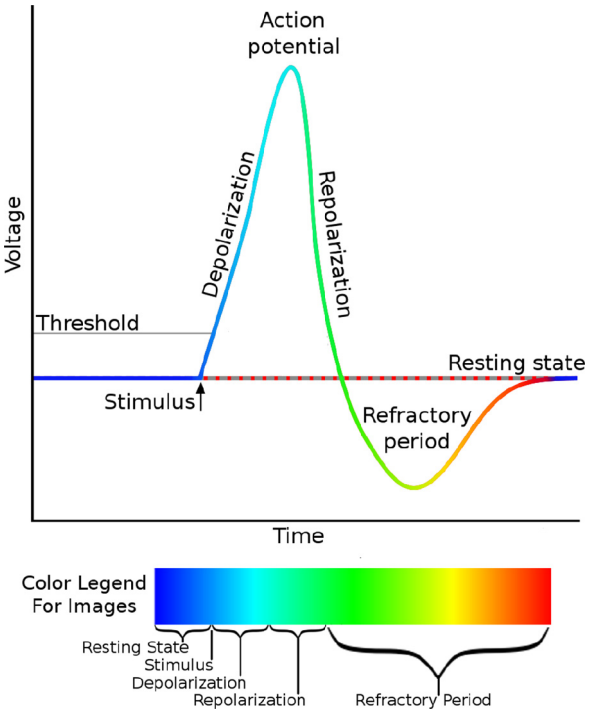


Fig. 1. The action potential model shows the change in voltage over time, corresponding to the stage of susceptibility, excitedness, or recovery of the cell. It is important to note that the color corresponding to each state varies across the spectrum, such that changes in a voxel cell's color indicate its activation and recovery in Figs. 4 and 5, (but not Fig. 2). (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

3.2. Topological optimization

The combination of morphology and control in these creatures results in a locomotion behavior inside this simulated physical world, with some creatures able to move more efficiently than others. Using this locomotion ability as a selection metric, we evolve a population of these creatures to move quickly over flat ground. This is achieved via an evolutionary algorithm which gives individuals with better locomotion abilities a disproportionately larger chance of reproducing and passing their "genes" down to an offspring in the next generation. Over time, this algorithm gradually increases the locomotion abilities of a population of these creatures.

In particular the evolutionary algorithm we employ is CPPN–NEAT [28], which stands for Compositional Pattern Producing Network with NeuroEvolution of Augmenting Topologies. This particular algorithm is chosen largely because it efficiently represents spacial domains, such as robot morphology or cellular automata (Fig. 3). It does so by indirectly representing the voxel phenotype as a network of regular mathematical functions. The regularities in these functions create spatially correlated transformations of basic input gradients (such as the relative x, y, or z position of cell). This can be thought of as analogous to the way in which morphogens are composed and transformed in a developing embryo. These networks thus attempt to capture the large-scale regularities of development, while abstracting away the particular physical processes that cause them.

The transformation from genotype to phenotype occurs by separately querying the genome network for the material type at each potential voxel in the creature. The voxel's coordinates (Cartesian coordinates x, y, z, and d – for radial distance from the center of the space) are fed as inputs to the network. These inputs are propagated down through the network, leading to output nodes which describe the material properties of that voxel. One output denotes whether the voxel is present in the creature – defining it's overall shape. A second output defines the voxels physical properties (whether it is a muscle or a soft supportive tissue). The third output node in the network determines the electrical signaling properties of the voxel (whether it is conductive or insulating). Thus each cell may be one of four possible configurations: empty (not part of the creature), a conductive muscle, a conductive passive tissue, or a non-conductive tissue cell (note that the muscular physical property of a non-conductive cell is irrelevant since it can never be excited). A sketch of this genotype to phenotype mapping is shown in Fig. 2.

4. Results

Since we are not aware of a previous experiment which optimized the topology of a cellular automaton with a constant update rule in order to change its behavior, we aim primarily to demonstrate a proof by construction.

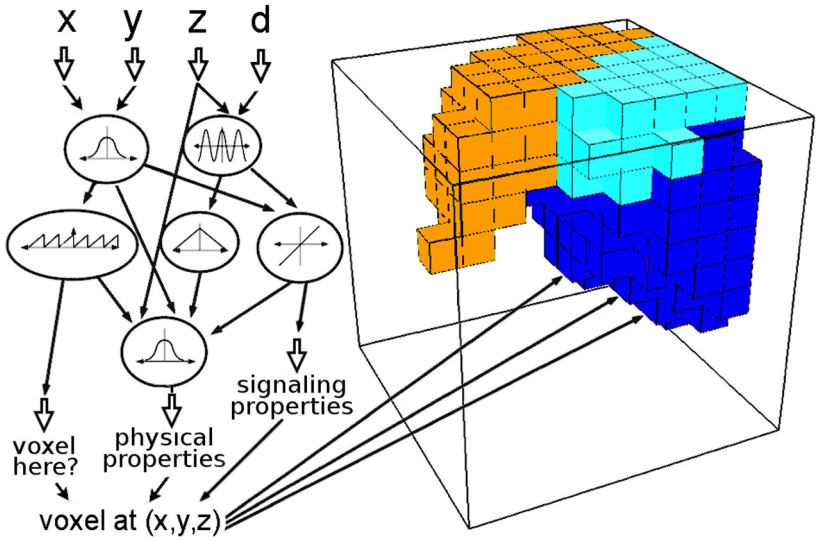


Fig. 2. This diagram demonstrates the transformation from a Compositional Pattern Producing Network genotypel *(left)* to a 3D virtual creature phenotype *(right)*. The querying of the network is done for each voxel (with that voxel's coordinates serving as the input variables) to determine the voxel cell properties at each location within the $10 \times 10 \times 10$ voxel design space. The voxel colors in this figure represent the material types (orange for muscle, dark blue for conductive tissue, and light blue for insulating tissue). These material colors are unrelated to the activation color legend in Fig. 1, and the authors apologize for the overlap of color maps. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

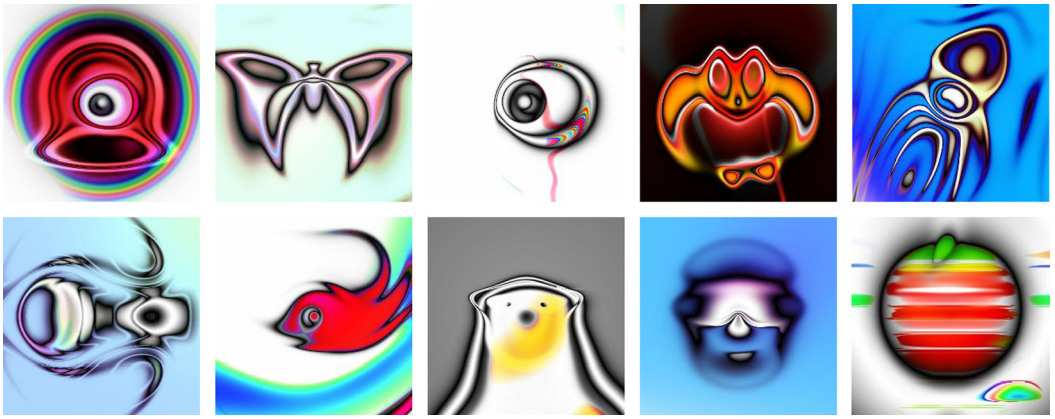


Fig. 3. Example images evolved with the CPPN–NEAT encoding we employ help to demonstrate the richness and organization possible in our system [23].

Fig. 4 demonstrates a number of different topologies which lead to unique behaviors of the resulting creatures, upon the iteration of our reaction-diffusion rule set. The variety of topologies (and thus control strategies) that came about within the context of the same update rule suggests the possibility of this technique to provide to a much richer possibility space for the design of cellular automata, compared to those which use the same update rules on the traditional regular and homogeneous substrates. As this type of dynamic behavior is best demonstrated via video, a variety of resulting behaviors may be viewed at: tinyurl.com/EvolvingSoftRobots2.

The types of locomotion behaviors which were evolved also spanned from simple and organized to complex and difficult to interpret. Fig. 5 exemplifies each extreme. The creature at the top of the figure shows a homogeneous conductive muscle, which smoothly diffuses the electrical signal out from the center of the topology. This orderly diffusion pattern demonstrates no non-linearities. The creature at the bottom of the figure exemplifies one with many dedicated conductive channels evolved to carefully route the electric signals throughout the creatures. The particular channel shown on the front face of the creature demonstrates wave splitting and wave collision of the electric pulse, examples of non-linearities which could, and evidently do, arise.

To demonstrate that these behaviors have been, in some sense, optimized by our evolutionary algorithm, Fig. 6 exemplifies the increase in locomotion abilities over time. The rise in distance traveled over time (i.e. the difference in locomotion ability from generation 1 to generation 1000) shows that by changing the topology of our substrate only, we are able to drastically improve the behavior of the overall system towards some particular goal (with locomotion being our goal in this

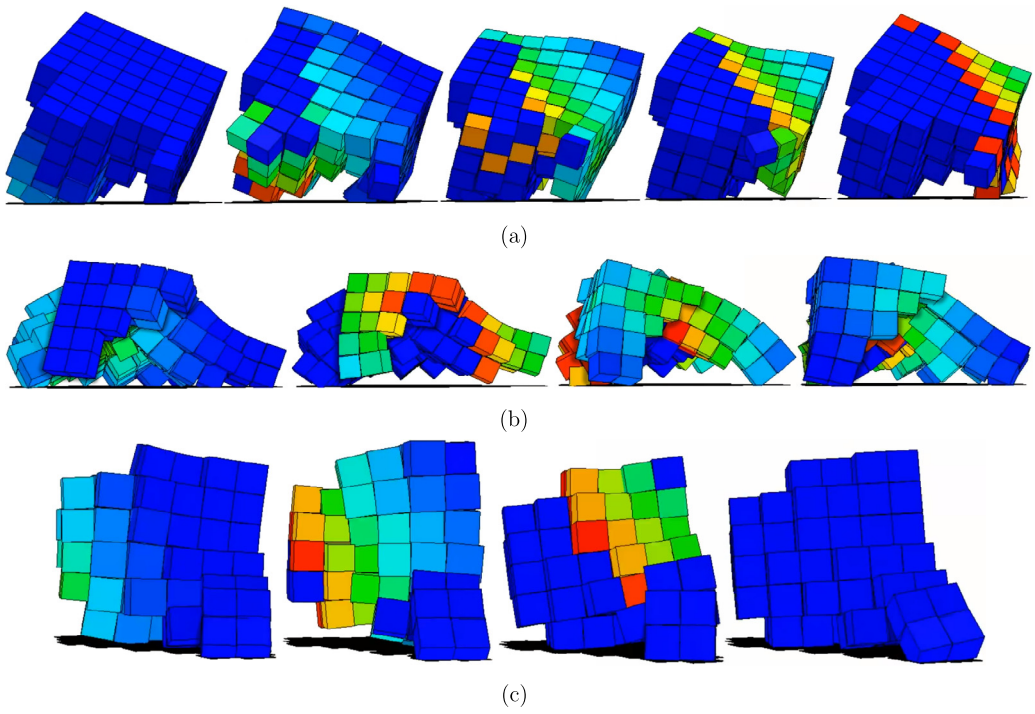


Fig. 4. (*See Fig. 1 for color map.*) A diversity of behavior came about through different evolved substrate topologies following the same update rule. *(a)* A front-back stepping motion kicks the front "leg" back to pull the creature forward then pushes the appendage forward to plant the it again. *(b)* A tripod gait springs forward with the front leg, while stabilizing and pushing off with the side and back appendages. *(c)* A small, non-actuated, side appendage is used to push off and propel the main body of the creature forward while it contracts its body from back to front. Videos of these gaits and others can be found at: tinyurl.com/EvolvingSoftRobots2.

scenario). The difference between the relative increase in the curves indicates the dependence on initial conditions (and in some sense additional perturbations) of the system, as the different curves correspond to different locations of the excitation added to the system by the pacemaker. The plot shows that a system which includes signals originating from the center of the creature are more easily optimized for locomotion abilities than systems in which the source of signals in located along an edge or in a corner of the design space.

5. Discussion

By demonstrating a variety of behaviors produced from the same update rules applied to different topologies (Fig. 4) and showing that these topologies can be optimized to produce more effective behavior along some predefined criteria (Fig. 6), we believe that we have demonstrated a novel and significant expansion on the way in which behaviors can be created within a cellular automaton. This serves as a proof by construction that altering the topology of the substrate upon which a set of update rules is applied can dramatically change the behavior which results from it.

We have also shown that this implementation can produce globally emergent behavior. This represents a significant improvement on previous non-uniform cellular automata optimization which focus on decomposable talks [24–26] – which are, by definition, less emergent. The task of automata locomotion exemplifies such emergence, as the function of any single cell gives little to no indication of the locomotion abilities of all the cells interacting together.

While this work serves as a simple proof of concept for varying topology, many extensions and variations beg additional future work. For example, an extension of these results may imply that we can learn more about different update rules themselves by varying the topology upon which they are applied. In a system where robustness to failure is important, one might imagine testing a number of different rule sets on many different topologies which vary in the size, distribution, or quantity of empty cells. Thus depending on the differences in resulting behavior for different sets of "knocked-out" empty cells, one may be able to quantify the robustness of a given set of rules in general. Furthermore, one could imagine co-evolving a set of topologies which would be indicators of a certain threshold of robustness, when applied across a wide variety of rule sets (optimizing for a novel level of average "disruption" across rule sets, compared to other topologies).

One could also imagine a number of applications of this sort of topological design. When dealing with hardware implementations of cellular automata (with cells of an automaton ranging in complexity from materials with certain conductive properties, to combinations of embedded integrated circuits) one could imagine that altering the functionality of each cell (how it "updates" over time) might be far more difficult or costly than simply altering the system's topology (and thus connectivity of cells). While the space of designs through topological design is often still limited by the particular update rule

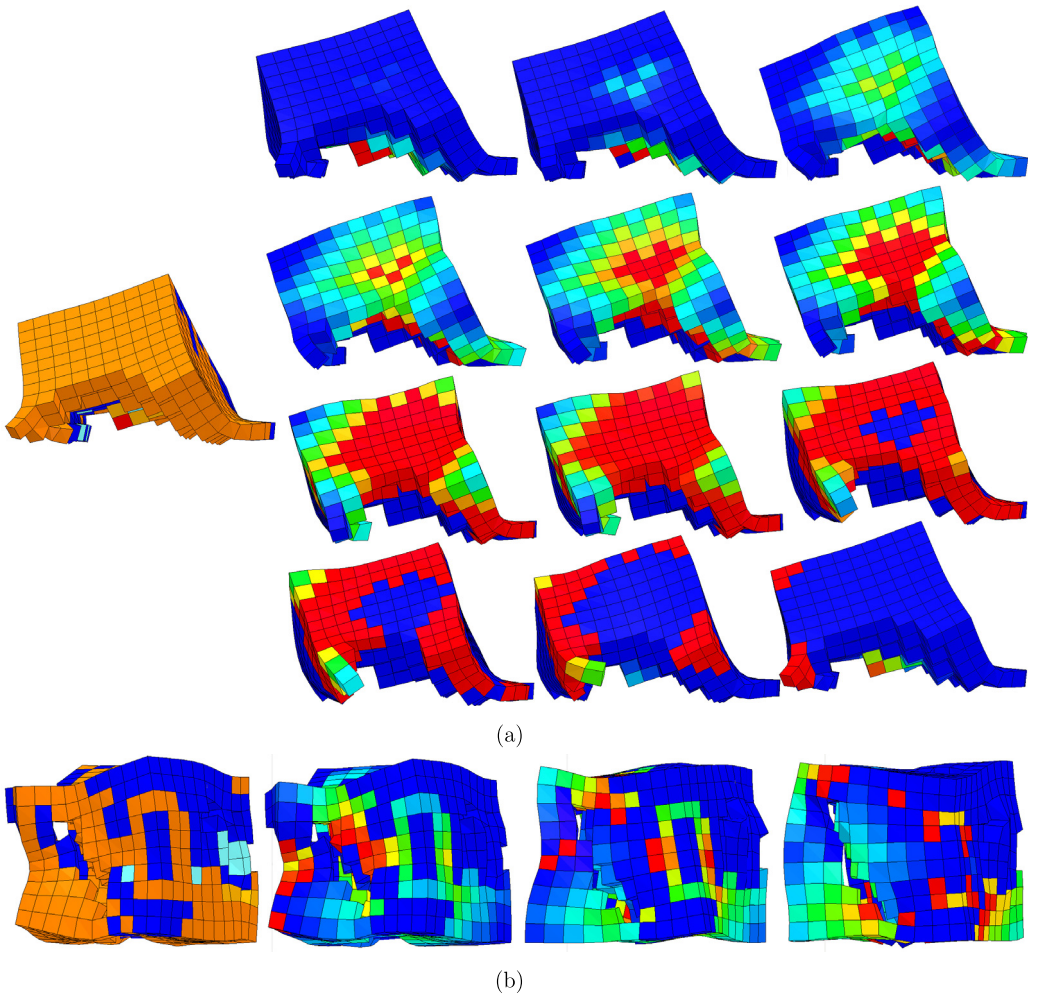


Fig. 5. (*See Fig. 2 for static material type color map of left side of figure, and Fig. 1 for iterated excitation phase color map of right side.*) The types of behaviors evolved varied from organized and simple to intricate and complex. *(a, left single)* The front side of the top creature shows a simple homogeneous muscle patch (in orange). *(a, right* 4×3*)* Once the update rules are iterated on this simple topology, we see the electrical signal originate from the pacemaker under the muscle at the center of this creature, and notice the signal diffuse across the muscle patch in an orderly fashion as the figures shows progressive snapshots of the creature's state over time (from left to right, top to bottom). *(b, left single)* In more complex creatures, specific conductive pathways were evolved (orange circle) within insulating tissue (blue). *(b, right* 1×3*)* When signals are propagated through such pathways they often result in non-linearities. Exemplified here are two simple demonstrations of non-linearities. First, the signal splits to form two independent streams as it hits a T-junction. Following this, it converges on the other side of the circle to crash them back together into a single stream. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

it is coupled with, the idea of topological design may be particularly appealing towards design automation of such systems, given the relative success of this proof of concept.

Returning to the inspiration from our natural system (the human heart), it is worthwhile to note that the most complicated non-linear behavior – the chaotic self-sustaining rotors which arise in cardiac modeling of arrhythmia (mentioned in Section 1) – were not evolved here. We conjecture that this is due in part to the nature of our experimental setup. Since we provide a regularly self-exciting pacemaker cell, the need to create self-sustaining patterns of excitation does not arise. It is also unclear if a single point of excitation is sufficient to produce such complexity in our control paradigm without simultaneously optimizing initial conditions, and thus future work with this paradigm will employ multiple sources of excitation. Similarly, our task (locomotion over flat ground) is fairly straightforward and incentives regularly paced actions with little required computation or complexity. Future work is required to explore experimental setups which allow and incentivize the rich interactions found in the chaotic state of such systems, and to analyze the resulting low-level signaling and physical behaviors.

Finally we return to the comment on embodied cognition. We noted above that embodied cognition states that the body plays a fundamental role in the way which the mind processes information and functions. This is typically interpreted as an implementation where the body modulates the way in which information is collected, and that the organization of the brain is thus made to properly synthesize and compute the information as it sees it through the lens of its particular body.

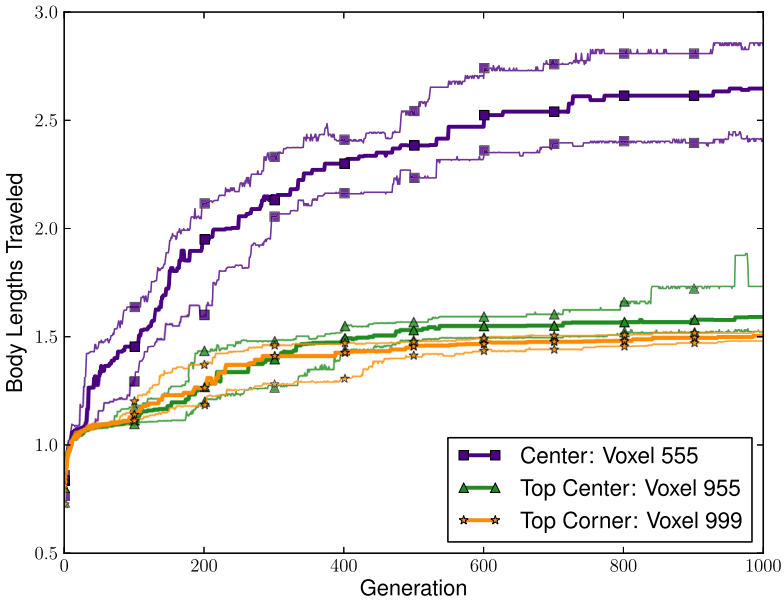


Fig. 6. A plot showing the relative ability of the system to optimize locomotion behavior given different locations of the pacemaker cell, which acts as the source of excitation to the creatures. The significant improvement (95% confidence interval plotted with the median value across 30 independent trials) between the locomotion abilities at generation 1 and 1000 show that we can optimize the behavior of a 3D cellular automaton by optimizing only its topology and keeping its update rules constant. However, the difference in the relative ability to do so in different experimental setups (since the pacemaker is the originator of excitation, we can think of this loosely as different initial conditions) suggests that certain setups may be better than others, and additional future work is required to explore what makes a system such as this more or less evolvable. Since the center placement of the pacemaker significantly outperforms the placement of the pacemaker along an edge or corner, our default condition has pacemaker placed at the center of the design space, unless otherwise noted. (Voxel numbers are counted left to right, back to front, bottom to top, out of a possible 1000.)

In our experimental setup, the information in the system was not only collected by the body (though its connection to the pacemaker – or touch sensors, in alternate experiments not listed in detail here), but the body is also what performs computations with it. We see this computation through the splitting or colliding of information pathways in Fig. 5, and more generally in the ability to optimize non-muscular signaling pathways within these creatures. Since the creature do not have a separate "brain" for these pathway to lead to, their use suggests that some sort of information processing is done directly within (and by) the morphologies of these creatures to achieve their behaviors – as is suggested by the theory of morphological computation [22].

6. Conclusion

We have proposed and demonstrated a novel method for creating richness and altering behavior of cellular automata. Rather than varying the update rules of each cell, we alter the topology of the substrate – thereby modifying the way in which information can be passed throughout the automaton. We show that altering topology can create a diversity of behaviors for a given update rule. We also demonstrate physical embodiment to cellular automata, allowing them to interact with an external environment and produce emergent behavior within the context of that environment – creating behavioral metrics upon which we can optimize the topology of the automaton. As this is a novel approach to optimization of behaviors in cellular automata, many questions are opened and left unanswered for future work, but these authors are excited for the potential and application of such a technique.

Acknowledgements

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

References

[1] R.R. Aliev, A.V. Panfilov, A simple two-variable model of cardiac excitation, Chaos Solitons Fractals 7 (3) (1996) 293–301.
[2] M.L. Anderson, Embodied cognition: a field guide, Artificial Intelligence 149 (1) (2003) 91–130.
[3] G. Beni, J. Wang, Swarm intelligence in cellular robotic systems, in: Robots and Biological Systems: Towards a New Bionics?, Springer, 1993, pp. 703–712.
[4] F. Brauer, C. Castillo-Chavez, Mathematical Models in Population Biology and Epidemiology, Springer, 2011.
[5] Z. Butler, K. Kotay, D. Rus, K. Tomita, Cellular automata for decentralized control of self-reconfigurable robots, in: Proc. of the ICRA 2001 Workshop on Modular Robots, 2001, pp. 21–26, Citeseer.

[6] H. Calkins, J. Brugada, D.L. Packer, R. Cappato, S.-A. Chen, H.J. Crijns, R.J. Damiano, D.W. Davies, D.E. Haines, M. Haissaguerre, et al., Hrs/ehra/ecas expert consensus statement on catheter and surgical ablation of atrial fibrillation, Europace 9 (6) (2007) 335–379.
[7] N. Cheney, J. Clune, H. Lipson, Evolved electrophysiological soft robots, in: ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems, vol. 14, 2014, pp. 222–229.
[8] N. Cheney, R. MacCurdy, J. Clune, H. Lipson, Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding, in: Proc. of the 15th Genetic and Evol. Comp. Conf, ACM, 2013, pp. 167–174.
[9] L.O. Chua, L. Yang, Cellular neural networks: applications, IEEE Trans. Circuits. Syst. 35 (10) (1988) 1273–1290.
[10] J. Clune, B.E. Beckmann, C. Ofria, R.T. Pennock, Evolving coordinated quadruped gaits with the hyperneat generative encoding, in: IEEE Congress on Evolutionary Computation, CEC'09, IEEE, 2009, pp. 2764–2771.
[11] J. Conway, The game of life, Sci. Am. 223 (4) (1970) 4.
[12] J.P. Crutchfield, M. Mitchell, The evolution of emergent computation, Proc. Natl. Acad. Sci. USA 92 (23) (1995) 10742–10746.
[13] M.J. Cutler, D. Jeyaraj, D.S. Rosenbaum, Cardiac electrical remodeling in health and disease, Trends Pharmacol. Sci. 32 (3) (2011) 174–180.
[14] R. FitzHugh, Impulses and physiological states in theoretical models of nerve membrane, Biophys. J. 1 (6) (1961) 445–466.
[15] F. Gers, H. De Garis, M. Korkin, Codi-1bit: a simplified cellular automata based neuron model, in: Artificial Evolution, Springer, 1998, pp. 315–333.
[16] J. Hiller, H. Lipson, Dynamic simulation of soft multimaterial 3D-printed objects, Soft Robot. 1 (1) (2014) 88–101.
[17] G.S. Hornby, J.B. Pollack, Evolving l-systems to generate virtual creatures, Computers & Graphics 25 (6) (2001) 1041–1048.
[18] C.G. Langton, Computation at the edge of chaos: phase transitions and emergent computation, Phys. D 42 (1) (1990) 12–37.
[19] K. Lucas, On the refractory period of muscle and nerve, J. Physiol. 39 (5) (1909) 331–340.
[20] M. Mitchell, J.P. Crutchfield, P.T. Hraber, Evolving cellular automata to perform computations: mechanisms and impediments, Phys. D 75 (1) (1994) 361–391.
[21] J. Nagumo, S. Arimoto, S. Yoshizawa, An active pulse transmission line simulating nerve axon, Proc. IRE 50 (10) (1962) 2061–2070.
[22] R. Pfeifer, G. Gómez, Morphological computation-connecting brain, body, and environment, in: Creating Brain-Like Intelligence, Springer, 2009, pp. 66–83.
[23] J. Secretan, N. Beato, D.B. D'Ambrosio, A. Rodriguez, A. Campbell, K.O. Stanley, Picbreeder: evolving pictures collaboratively online, in: Proc. of the 26th SIGCHI Conf. on Human Factors in Computing Systems, ACM, 2008, pp. 1759–1768.
[24] M. Sipper, Non-uniform cellular automata: evolution in rule space and formation of complex structures, in: Artificial Life IV, vol. 2, 1994, pp. 394–399.
[25] M. Sipper, Studying artificial life using a simple, general cellular model, Artif. Life 2 (1) (1994) 1–35.
[26] M. Sipper, Co-evolving non-uniform cellular automata to perform computations, Phys. D 92 (3) (1996) 193–208.
[27] M. Sipper, E. Ruppin, Co-evolving architectures for cellular machines, Phys. D 99 (4) (1997) 428–441.
[28] K.O. Stanley, Compositional pattern producing networks: a novel abstraction of development, Genet. Program. Evol. Mach. 8 (2) (2007) 131–162.
[29] J.N. Weiss, A. Garfinkel, H.S. Karagueuzian, Z. Qu, P.-S. Chen, Chaos and the transition to ventricular fibrillation a new approach to antiarrhythmic drug evaluation, Circulation 99 (21) (1999) 2819–2826.
[30] S. Wolfram, Statistical mechanics of cellular automata, Rev. Modern Phys. 55 (3) (1983) 601.