Evolving Soft Robots in Tight Spaces

Nick Cheney
Creative Machines Lab
Cornell University
Ithaca, NY, USA
nac93@cornell.edu

Josh C. Bongard
Morphology, Evolution & Cognition Lab.
University of Vermont
Burlington, VT, USA
jbongard@uvm.edu

Hod Lipson
Creative Machines Lab
Cornell University
Ithaca, NY, USA
hod.lipson@cornell.edu

ABSTRACT

Soft robots have become increasingly popular in recent years – and justifiably so. Their compliant structures and (theoretically) infinite degrees of freedom allow them to undertake tasks which would be impossible for their rigid body counterparts, such as conforming to uneven surfaces, efficiently distributing stress, and passing through small apertures. Previous work in the automated deign of soft robots has shown examples of these squishy creatures performing traditional robotic task like locomoting over flat ground. However, designing soft robots for traditional robotic tasks fails to fully utilize their unique advantages. In this work, we present the first example of a soft robot evolutionarily designed for reaching or squeezing through a small aperture – a task naturally suited to its type of morphology. We optimize these creatures with the CPPN-NEAT evolutionary algorithm, introducing a novel implementation of the algorithm which includes multi-objective optimization while retaining its speciation feature for diversity maintenance. We show that more compliant and deformable soft robots perform more effectively at this task than their less flexible counterparts. This work serves mainly as a proof of concept, but we hope that it helps to open the door for the better matching of tasks with appropriate morphologies in robotic design in the future.

Categories and Subject Descriptors: [**Distributed Artificial Intelligence**]:Intelligent Agents, Mobile Agents

General Terms: Algorithms, Design

Keywords: Soft Robot, Generative Encoding, Artificial Life, CPPN-NEAT, HyperNEAT, Multi-Objective, Morphology, Material

GECCO '15, July 11 - 15, 2015, Madrid, Spain

ISBN 978-1-4503-3472-3/15/07…$15.00

DOI: http://dx.doi.org/10.1145/2739480.2754662

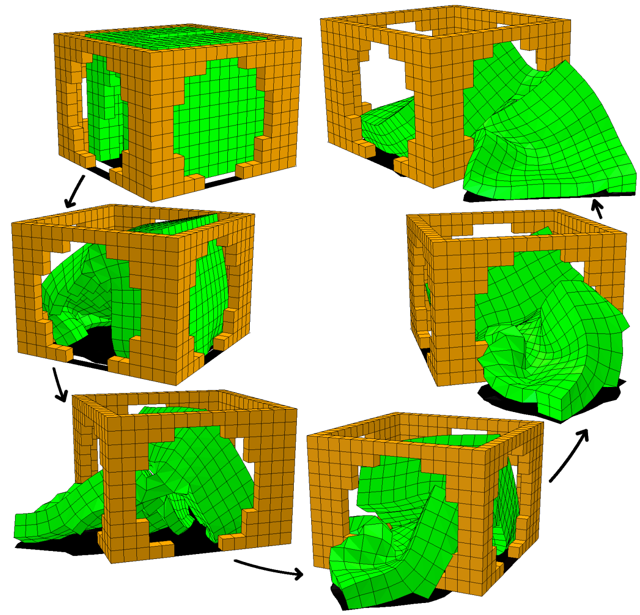

Figure 1: *(counter-clockwise rotating-viewpoint time series starting from top left).* An evolved soft robot reaches through a hole in the side of a cage surrounding it. The width of this aperture is smaller than any of the dimensions of this creature – thus a robot of the same size without a soft body would not be able to squeeze through it.

1. INTRODUCTION

Recent interest and developments in the study of soft robotics [2, 5, 10, 15, 18, 19, 20] have pointed towards a number of potential benefits of using soft material in the design of artificial creatures.

Recent work has also demonstrated the use of evolutionary computation to design effective soft robot bodies [4, 6, 11, 12]. Such an approach holds the potential for significant impact, since the extreme nonlinearities and degrees of freedom apparent in soft robots make their design unintuitive, compared to traditional rigid body robots. The design automation inherent in evolutionary computation removes the prerequisite of an intuitive understanding of these systems for their effective design.

In reviewing the "lessons from biology" that soft robots should inherit, Kim et al. note that "Soft materials are essential to the mechanical design of animals... These soft components provide numerous advantages, helping animals negotiate and adapt to changing, complex environments. They conform to surfaces, distribute stress over a larger volume, and increase contact time, thereby lowering the maximum impact force. Soft materials also lend themselves to highly flexible and deformable structures, providing additional functional advantages to animals, such as enabling entrance into small apertures for shelter or hunting... all of them can squeeze through gaps smaller than their unconstrained body. These are important lessons for building soft robots" [8].

However, up to this point, there has been no attempt to demonstrate the ability of these artificially evolved robots to perform the tasks which their biological counterparts have been evolved for. This is especially important, as Kim et al. make explicit the intuitive notion that "Ultimately it is probably the ecological niche that determines the evolutionary tendency to be stiff or soft" [8].

This work attempts to provide the first demonstration of an evolved creature for an explicitly "soft-body-oriented" task, by evolving creatures for the task of "entrance into small apertures." In this work, a soft robot bounded by a $11 \times 11 \times 11$ maximum size is placed within a cage of a 15×15 footprint, with height 11. The cage has holes in each side of its sides with a diameter of 10. This aperture may be more restrictive than it first appears, as a square of side length 11 has an area of 121, while a circle of diameter 10 has an area of 78.5 (just 65% of the maximum potential face area of the robot). We evolve soft robots to reach or squeeze through this aperture using the CPPN-NEAT evolutionary algorithm, and demonstrate a variety of effective, creative, and entertaining behaviors (such as Fig. 1).

We hypothesize that softer, more deformable, robots will have an easier time accomplishing these tasks. At either extreme, one could imagine that a (maximum size of 11^3) robot that is a rigid solid would be physically unable to fit through a hole of diameter 10, while a creature composed of that volume of extremely soft material (liquid at the pure extreme) would easily flow through the aperture. Of course, neither of these robots are likely to take on the structure of a lattice of voxels, as our soft robots do (a rigid robot would be likely to include joints, while a flowing liquid would require the free movement of particles), nor would our soft body physics simulator be equipped to handle either of these cases. Thus, we approach this investigation by comparing the ability of more or less compliant soft robots to move through this aperture, and leave the reader to extrapolate to these extreme cases.

Secondarily, this work also demonstrates an example of multi-objective NEAT with speciation. While this is not the focus of the paper, and thus we do not provide comparisons to other examples of multi-objective NEAT without speciation, it is a novel implementation and may be of interest to those hoping to explore added diversity maintenance within multi-objective optimization.

2. BACKGROUND

The notion of soft robots being good at squeezing though small openings has been approached previously by hand designed robots. Sheperd et al. created a molded silicon robot which was able to squeeze under a barrier with a $2cm$ clearance (the maximum dimensions of the robot were $13.6cm \times 5.9cm \times 0.6cm$) [15]. However, the design of the robot's morphology was created by hand, and the robot was controlled manually via tethered pneumatic actuation.

Various examples of evolved soft robots have also been demonstrated [4, 6, 11, 12], however they all focused on the task of locomotion over flat ground. While not to say that soft robots do not hold any advantages for locomotion over a smooth planar surface, the advantages of such an approach are not as inherently apparent as a task in which rigid body robots are unable to perform – such as navigating through an aperture smaller than the robot's body.

3. METHODS

The source code (including a configuration file with parameter values) can be found at: http://git.io/vfSLV

3.1 SIMULATED TASK ENVIRONMENT

In this work (and consistent with [4, 6, 11]), these soft robots are simulated in the soft-body physics simulator VoxCad [7]. Approximating an array of soft voxels as lattice of points connected by simulated beams, this physics engine is capable of efficiently modeling soft bodies, while maintaining physical and quantitative realism. VoxCad creates actuation within these soft robots by employing a sinusoidally varying global temperature. All passive cells (such as the blue support tissue or gold cage voxels) are unaffected by this temperature and remain a constant volume. Active muscle cells (two types: green and red) vary in size as this temperature changes. They do so out of phase to one another, with the green cells contracting then expanding, and the red cells expanding then contracting. The variation in size due to these temperature changes results in a 14% linear contraction/expansion from their baseline size, which results in approximately a 48% volumetric change. Each individual's evaluation period lasted for 20 of these actuation cycles.

New to this study, each soft robot is placed within a cage at the beginning of each simulation for fitness evaluation. The cage has dimensions $15 \times 15 \times 11$, leaving a one voxel gap between in the x and y directions between the edge of the cage and the $11 \times 11 \times 11$ maximum size of the evolved creature. The top is left open for ease of viewing – and no evolved creatures manage to fully escape out of the top of the cage. The cage is simulated to be perfectly rigid and immobile, as well as indestructible – forcing the robot to contort itself and travel through one of the openings in the side of the cage. Each side has an opening which is approximately (rounded to the nearest voxel) a circle of diameter 10 – thus in the 15×11 side face of the cage, there is one voxel above the opening, two voxels to one side of it, and three voxels to the other. This produces a circular opening of area 78.5 (before being discretized to the nearest voxel), which represents about 65% of the area of a full 11×11 face of a soft robot (it's maximum, but not guaranteed size).

A modification to the default operation of VoxCad was necessary to ensure that collisions detection between the cage and the robot ensured the robot was never able to accidentally pass through a part of the cage other than the opening. This modification caused collisions to be calculated between every voxel at every time step (as opposed to just surface voxels in VoxCad's default settings). The expense of this technique varies with the surface area to volume ratio of the evolved creatures, but it could result in as much as a six-fold slow down in simulation speed.

3.2 FITNESS METRICS

At the end of an individual's 20 action cycle evaluation period, one of two metrics were taken to summarize its behavioral outcome. In one configuration, the single farthest voxel from the center of the cage was recorded, and that voxel's position was returned to give the creature's "maximum reach" score. This fitness criteria incentivized the robots to stretch as far out of the cage as possible – and while leaving the cage entirely is one way to maximize this score, it was not necessary to successfully attain a successful reach (e.g. for object manipulation) outside of the cage boundaries.

To incentivize creatures to entirely leave the cage (important for locomotion rather than object manipulation), trials of an alternative configuration were conducted where the behavioral fitness score of an individual was total number of voxels the robot was able to move outside of the cage by the end of the trial. The total number of voxels (rather than the proportion of the creature's mass) was used at a method or further incentivize evolution to create large creatures (as robots small enough to walk out of the cage without having to contort and squeeze themselves through the aperture are less interesting for thus study).

In both cases, creatures were also incentivized along a second objective – to maximize their size (Sec. 3.4. This metric was defined simply as the number of voxels from which a creature was composed. However, one could certainly imagine alternative size metrics (such as the diameter or maximal inter-voxel distance of the creature).

3.3 CPPN-NEAT

Consistent with [4, 11], we allow the soft robots to optimize their topology using the CPPN-NEAT evolutionary algorithm [1, 16].

The CPPN encoding represents the voxel phenotype as a network. This network takes a voxel's relative coordinates as inputs, and transforms this information into a material selection for that particular voxel. This transformation takes place by querying each potential voxel (discretized cell in the $11 \times 11 \times 11$ grid of the design space) with the same genotype network.

To query a voxel, the input layer of this network consists of four nodes, encoding the relative (-1 to 1) Cartesian coordinates (x, y, z) and polar radius (r) of that voxel. The network is updated through a series of hidden nodes, and produces real valued numbers for the three output nodes. Then a threshold (all thresholds occur at zero) on the first output node determines whether that potential voxel space contains a solid voxel or is empty space. If there is a voxel, a threshold on the second output determines whether the voxel is a passive support tissue or an active muscle. If the voxel is a muscle cell, the final output node determines which of the two out-of-phase muscle types the cell belongs to.

Thus, similarities in coordinate values for nearby voxels produce gradual changes in the expression of output values (i.e. morphogens) that determine cell fate. This produces global structure in the resulting creatures. Further regularities are produced through the varying activation functions at each hidden node. For example, a node which contains a Gaussian activation function would create a symmetric pattern along the gradient of that node's input values. Similarly, a node with a sinusoidal activation function would create repetition along its input gradient. As these transformations are applied one on top of another, complex shapes quickly emerge [14].

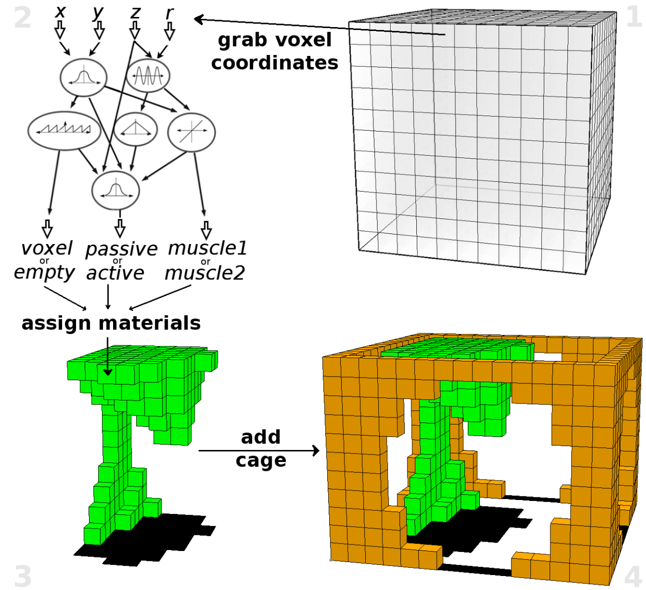

Figure 2: A sketch of the genotype to phenotype decoding. For each voxel in the potential design space, the relative coordinates values are taken and input into the network genotype. The material of each voxel is assigned based on the output of the network for that location. After each voxel has been individually queried for its material properties, the external environment (a cage surrounding the creature – Sec. 3.1) is put in place for the fitness evaluation.

A sketch of this genotype to phenotype decoding is provided in Fig. 2. The authors also note that this study is not directly in regards to the encoding employed, and will not explore the comparison of CPPN-NEAT to alternative encodings.

3.4 MULTI-OBJECTIVE NEAT

While the design flexibility afforded to the CPPN encoding creates a variety of complex forms, it also allows for the evolved creatures to simply produce topologies which would be smaller than the existing aperture, meaning that there would be no need for them to squeeze through a tight opening, and the resulting behavior would not be as interesting. To incentivize large forms which still squeeze through the smaller aperture, we created a multi-objective implementation of the NEAT algorithm. With this implementation, we are able to reward creatures for being large and also for squeezing though the aperture.

While various implementations of multi-objective NEAT have appeared recently [9, 21, 13], we believe that our implantation happens to be the only one which does not require the removal of the NEAT speciation (noted as one of the key features of the NEAT algorithm [17]). Instead of removing speciation and replacing it with a different diversity metric, or considering the added diversity inherent in multi-objective search to be sufficient on its own, our implementation performs a Pareto ranking of individuals within each species. We then perform the traditional tournament selection based on this ranking.

It is also noteworthy to mention that other implementation decisions were made to provide a bias towards the behavioral objective (squeezing through the narrow aperture) above the static objective (being larger in size). This implementation decision was made following the assumption that producing new individuals who were able to move farther (but were smaller) was a more difficult task than producing new individuals who were larger (but not able to move as effectively). This stems from the observation that any individual who adds one or more additional voxel to an existing fit phenotype would fulfill the later category, while not all individuals who were smaller (and likely a very small subset of them, who happened to be coordinated enough to improve their movement and behavior) would prove to be Pareto optimal in the former scenario. This intentional bias was instantiated by favoring locomotion behavior over size when comparing two individuals on the same Pareto front, and also by using the behavioral objective as the single objective required for NEAT's fitness sharing between species.

We should make explicit that this implementation of mulit-objective NEAT is not the focus of this paper. Thus no claims or comparisons relative to other implementations are presented, nor will the results section of this paper provide any quantitative support for any of the implementation decisions made above. Future work is needed to investigate multi-objective speciation in a variety of task scenarios, in order to make claims of its suitability and potential advantages or disadvantages within them.

3.5 RUN CHAMPIONS

Since this work relies on multi-objective optimization, the best resulting creatures from each of the 30 independent runs will fall along a Pareto front (on behavioral performance and size). While this variety is generally beneficial, it makes comparison between trials and treatments more difficult. In order to simplify the comparison, we provide a more specific definition of the optimal robot we seek to create.

Consistent with our preference for behavioral outcomes over size outcomes (Sec. 3.4), we seek to optimize along the single objective of behavioral performance (reaching or moving). However, moving through an aperture becomes trivial if the size of the robot is less than that of the aperture, so we place a strict size constraint on the robots we consider for run champions. This constraint relates not to the volume of the evolved robots, but deals explicitly with a 2D slice of the robot – which must fit through the 2D aperture of the cage.

In an effort to ensure that the evolved morphologies actually do have a full 11×11 face (and thus have to deform or compress themselves to "squeeze" through the aperture on the side of the cage. In all comparisons below, we consider only robots who have at least one 2D slice that spans the maximum 11 voxel width (at some point along the face) in both directions. Those robots who do not have at least one slice (along the Cartesian coordinate axes) that meets this criteria are thrown out and not considered in the analysis below. While we realize that this is only a proxy, and not an exact match, for the criteria of needing to constrict oneself to squeeze through the cage aperture, we believe it to be a good first pass approximation – and informal visual inspection of evolved topologies supports this belief.

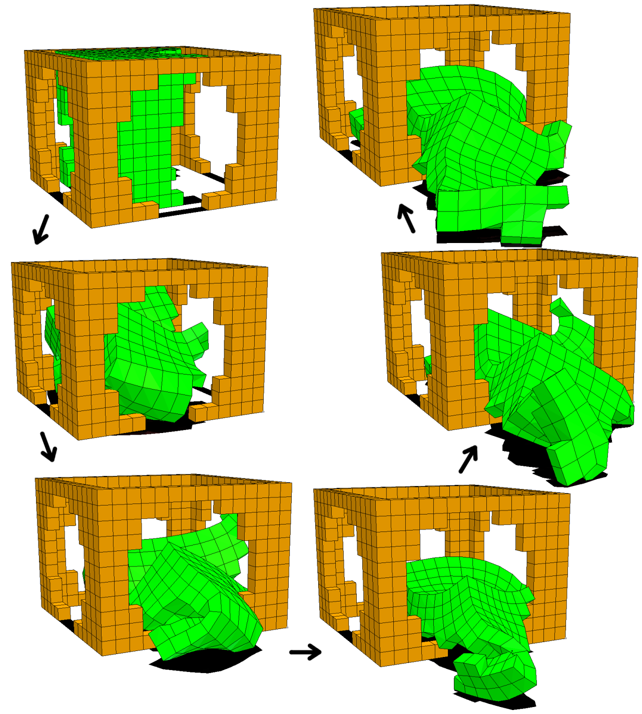

Figure 3: ***(counter-clockwise side-view time series starting from top left)*** **This creature, evolved for reaching, writhes back and forth in an effort to unroll itself and produce a long "arm" which reaches out from its cage.**

At the end of each run, the evolved robot which meets this size criteria threshold and demonstrates the farthest movement (or reaching) performance, is considered to be the best individual (the "champion") of that run. It is worth noting that this thresholding process does not take place during the actual optimization, but simply performed in post hoc analysis.

3.6 STATISTICAL ANALYSIS

Any statistical values reported below come from the Mann-Whitney U test, since normality of the distributions cannot be assumed. Since the distributions are unknown (and in an effort to help inform the skew of the distributions) both mean and median values are reported.

4. RESULTS

Since this work serves only as a first-pass proof of concept and demonstration of soft robots evolved to reach or squeeze through a small aperture, we primarily seek to demonstrate that is possible to successfully evolve soft robots for this task. For the case of reaching as far as possible outside of the cage, the best evolved robots of each run are able to reach an average length of 15.69 voxels (1.43 times their original body length of 11) outside of the cage (standard deviation: 3.21 voxels, median: 15.18). The fact that their farthest point is more than one original body length form the outside edge of the cage should not be an indication that these creatures frequently were able to completely exit the cage in the allotted time, as often they would unfold and spread out (e.g. Fig 3) to reach a significant distance while still remaining partially inside the cage.

In response to selection for the total number of voxels outside of the cage, the best of each run in the baseline conditions were able to move a mean of 140.30 voxels outside of their enclosure (median: 120.50). These results showed a great deal of variation (standard deviation: 98.77 voxels), with the most fit creature (on the movement objective) able to move 611 (of its 1130 voxels) out of the cage in the allotted time. The most fit creature of the run champions on the size objective was able to fill all 1331 of its voxels, but only move 60 of them out of the cage.

Perhaps a more compelling case for the ability to design squishy creatures who are able to navigate through their enclosures is the visual inspection of the resulting evolved soft robots. We believe that the diversity of forms and strategies, the effectiveness of these creatures towards their behavioral goals, as well as the clear necessity and use of physical deformation presented in the following examples, more strongly supports the existence of effective and successful evolutionary design than the values cited above. Here are a couple more examples of this:

The top of each cage is intentionally left open – not for creatures to climb out of, but for viewing purposes. Fig. 4 shows how this view of a reaching robot clearly demonstrates an example of a creature which spans the entire frame of the cage, yet is able to deform itself and fold its body in upon itself to provide additional reaching opportunities. While this creatures would be able to fit through the aperture in its deformed state, it was simply incentivized for reaching distance – and one of its "supports" in the upright starting position provides an excellent "arm" for effortlessly reaching out through the aperture.

The case of a robot squeezing its entire body through the aperture is demonstrated in Fig. 5, as this creature comes from a trial in which the entire mass of the robot had to be moved (instead of simply considering its farthest reaching voxel). This example clearly demonstrates a creature squeezing itself and using its pliability to fit through an aperture smaller than the width of its body. The frames of the time series in which the robot is moving through the aperture clearly show the sides of the robot curled and folded back in upon themselves, creating a narrow enough girth to fit through the opening. This would not be possible without the deformability of the soft materials. To our knowledge, this creature represents the first evolved robot to fit though an opening smaller than the width of its body.

4.1 SOFT/STIFFNESS COMPARISON

We find it intuitive, and take for granted, that rigid body robots of equal size and shape would be not able to navigate through the openings of the cages presented here. This conjecture may not be more clearly demonstrated than in the case of Fig. 1. This creature fills up nearly the entire potential voxel space (leaving just a thin strip of empty voxels separating its front and back segments) and spans the full 11 voxels wide on each face. However one may attempt to turn or twist this robot, it would not be physically possible to fit it through one of the apertures without deforming the creature. Despite this, the time series of this creature clearly and simply demonstrates how it is able to deform itself, folding and twisting the flat face of its front half to squeeze it through the tight aperture and reach out of its cage.

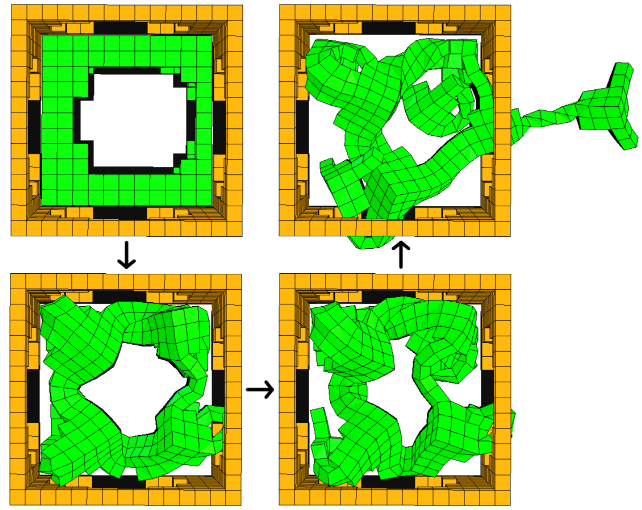

Figure 4: ***(counter-clockwise top-view time series starting from top left)*** **This evolved soft robot demonstrates a pure reaching behavior, where the majority of its body stays within the cage, but a dedicated arm reaches out through the aperture. This robot also exemplifies a highly deformable structure. In its starting configuration, its thin frame spans the entire inside of the cage, yet once it is given the chance to deform, it folds in upon itself to produce an entangled and complex morphology. Unlike its initial configuration (which a rigid body robot would have to stay in), the deformed morphology allows enough flexibility to position itself next a hole in the cage and reach out through it.**

However, the implementation of a completely rigid robot in this context (a voxel array with no explicit joints) makes little sense. Thus a fair comparison cannot be made between fully rigid and soft – so our comparison will be between two soft robots of varying stiffness. The above result provides the stiff(er) creatures, which have a Young's modulus (measuring tensile elasticity) of 10 megapascals (MPa). We also test a treatment which employs muscle and tissues that are an order of magnitude more elastic (Young's modulus of 1 MPa). For contextual grounding, this is approximately the range of values corresponding to the least elastic and the most elastic silicone rubbers. The more elastic (i.e. softer) robots are able to move, on average, 174.43 voxels outside of the cage (standard deviation: 69.18, median: 155.50). According to the Mann Whitney U test, this is a significant increase from the stiffer treatment (with mean of 140.30, standard deviation: 98.77, median: 120.50), with a one-sided p-value of 0.0002.

Additionally, the softer treatment also scores better on the second objective for total voxel size. The softer robots (mean: 905.08, standard deviation: 266.20, median: 971.63 voxels) were significantly larger than the stiffer robots (mean: 718.74, standard deviation: 306.13, median: 612.26 voxels), with a one-sided p-value of 0.0107. Presumably this made it more difficult for them to fit through the opening (though both treatments were subjected to the post hoc size thresholding noted above, so neither one should be able to do so effortlessly).

This suggests that, at least for intermediate stiffnesses, softer robots are more effective at squeezing through a small aperture.

4.2 NUMBER OF MATERIALS

A surprising result from Figs. 1, 3, and 4 is that the creatures commonly appear to be made out of only a single material. This is unlike previously published results [4], where a figure is presented to demonstrate the variety and consistency of the material compositions for locomotion over flat ground.

Fig. 5, which shows a creature rewarded for moving its entire body through the aperture, does exemplify multiple materials. Perhaps one might assume that this task is more closely related to locomotion over flat ground, and thus more likely to match the previous result of mult-material evolved soft robots noted previously for that task.

However, the statistical data does not support this assumption. Run champions on the reaching task have, on average, 1.23 different materials (the median number of materials is 1 for all tasks). In comparison, run champions from the movement task tended to have even less material diversity, with an average value of 1.20 materials per individual. Not surprisingly, there is no statistically significant difference between material composition of robots evolved for the two different tasks ($p = 0.76$).

An analysis of the the entire population (not just run champions) shows average number of different materials per individual of 1.36 and 1.37 (of reaching and movement, respectively). While the difference is not statistically significant ($p = 0.19, 0.09$), the fact that the run champions have a mean towards a lower level of material diversity – compared to a randomly drawn individual during the optimization process – may suggest that this task actually incentivizes creatures to be of a single material to effectively and efficiently solve both the reaching and traveling tasks.

4.3 VOXEL PENALTY

Also in slight disagreement with previously work in this area [4], is the notion that larger creatures with more voxels are less able to move around. Cheney et al. show a weak tie between robots composed of more voxels and more distance traveled across flat ground. However in this scenario, more voxels lead to worse performance metrics. This relationship is true of the reaching task ($slope = -11.99, r^2 = 0.70$) as well as the movement task ($slope = -5.26, r^2 = 0.51$). This relationship is not surprising, as larger creatures have the obvious disadvantage of struggling to fit through the narrow aperture. It is interesting to note that the movement task in this work is more closely related to the task of locomotion over flat ground (than is reaching), perhaps helping to explain why larger sized robots correlate less strongly with negative behavioral outcomes in the movement task.

5. DISCUSSION

The results presented above clearly exemplify soft robots evolved to reach or squeeze through apertures smaller than their own body, and are the first of their kind to do so. We find the pictorial representations of these behaviors to be more clarifying and convincing than the accompanying statistics, and are amazed at the power of this evolutionary process to design creatures which achieve these tasks in such creative and unintuitive ways.

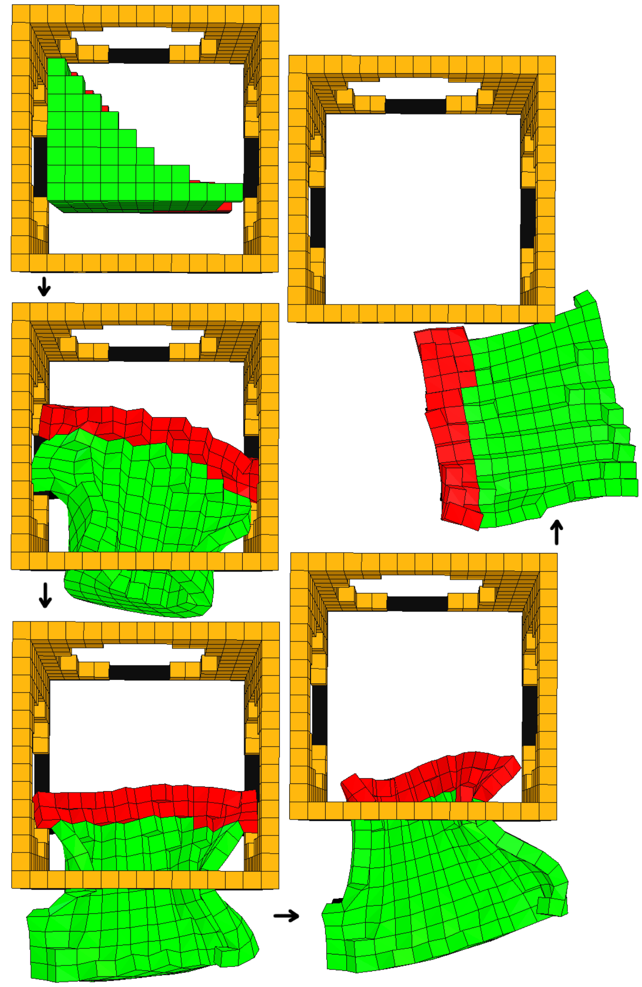

Figure 5: *(counter-clockwise top-view time series starting from top left)* This example of a multi-material soft robot squeezes entirely through the aperture to escape the cage – as it is rewarded for the movements of all its voxels. Notice in the lower three frames how the creature's body is clearly wider than the opening of the aperture, yet it is able to squeeze and roll/fold itself up to fit through the tight opening. This is a prime example of the abilities afforded by the deformable bodies of soft robots. (Note: To fully escape the enclosure, this creature required more than the allotted 20 actuation cycles, and was thus only able to do so in post hoc analysis.)

The case of primarily single material creatures is slightly puzzling if one approaches the results expecting similar robots to those which tend to be evolved for locomotion on flat ground. But a couple potential causes for the single material beasts may be suggested. For the case of the reaching task, only the farthest (and not the average) distance was incentivized. A long arm which contains muscles of both types would have a part of the arm contract as the rest of it extended (and vise versa) – potentially leading to hindered reaching abilities. Since the two muscle types contract with opposite sinusoidal actuations, and only the reach at the last time step was recorded for fitness purposes, the material composition of the robot was originally assumed not to affect its behavior – as the simulations stopping point was chosen to come at the end of a temperature cycle, where both muscle groups were supposed to be their original size (and the same size as the support tissue). However, due to non-zero simulation time steps, it is unlikely that this stopping point occurred when the voxels were all exactly the same volume.

The case of single materials in creatures incentivized for movement follows a similar logic of opposing muscle groups. These creatures were implicitly rewarded for contracting as tightly as possible in order to squeeze their entire body through one of the apertures. Once part way through the aperture, expanding again is not a significant problem, since the soft body will simply deform around the edge of the opening. Thus, creatures with a continuous muscle type are afforded the advantage of being able to constrict themselves more compactly, with relatively little penalty for ballooning to a much larger size afterward.

The opposing selection pressures towards single material creatures in this work and towards multi-material composition in previous results [4] suggest that squeezing through an aperture and locomotion over flat ground require different muscle contraction patterns. Additionally, the fact that these creatures require more than 20 actuation cycles to produce a behavior which has them entirely exit their cage suggests that there is room for significant gains in efficiency. We believe that this points towards the potential (and perhaps necessity) of a closed loop controller for this task. The feedback of sensory information (regarding position in or out of the enclosure, as well as normal force applied against/from the structure) certainly hold the potential to encourage specific pulling/pushing motions, which would effectively and efficiently squeeze the soft robot through the opening. Furthermore, once the creature escaped the enclosure, and the tactile feedback from the cage is no longer apparent, it would be free to employ a movement pattern specifically tailored towards locomotion over flat ground.

It is also unclear to what extent these traditional locomotion behaviors are incentivized in the results above. One may imagine that a creature (once it has unfolded to lay on the ground, as in Fig. 3) would benefit from the ability to move along the surface, as doing so would push its farthest voxel further away from the cage and also pull its back end further out of the cage (the selected behaviors in both the treatments). However, one may also imagine that such a significant amount of selection pressure involves navigating around the cage, that any morphology alteration which benefits locomotion at the cost of additional interference with the structure would results in poor fitness. This is especially true when one considers the relatively short evaluation time of 20 actuation cycles (relative in comparison to the time it must take for a creature to fully escape the cage, given by the fact that creatures here do not tend to do so). In such a short period, the robots may not be on the ground (and at the edge of the an opening) long enough to make significant gains by crawling out of it – even if they happen to evolve the ability to do so. Extended trials of 100 actuation cycles were performed for some of the run champions to explore this idea, but no significant improvements were noted. This is not especially surprising, given that these creatures had little (or no) selection pressure to evolve behaviors which would continue to be effective past the end of their evaluation periods. To ensure that this is the case, the creatures would need to be evolved with 100 actuation cycle lifetimes for the entire optimization process (a treatment not possible with our current implementation and the limitations imposed by our Advanced Supercomputing Division).

6. FUTURE WORK

As this serves primarily as a demonstration and proof of concept, there is much future work left to be done.

This work included a number of implementation decisions which should be explored in a rigorous manner. The size of the cage, size of the aperture, shape of the aperture, shape of the cage, size of the contractions/expansions, rate of contractions, number of materials, stiffness of materials, length of evaluations, and many additional parameters of the evolutionary algorithm itself were all set arbitrarily or based on previously published parameter values [4]. A rigorous exploration into this parameter space would likely lead to greater information and understanding of the system.

Another improvement to this system would be the implementation of a more sophisticated controller than VoxCad's global temperature. This could take the form of a high-level neural network controller, as is common in the field, or the form of low level morphological computation, as was previously implemented in this system [3]. While the deformability of soft robots plays a crucial part in the ability of biological creatures to fit through small openings, so too does their control – an aspect currently missing from the initial implementation of this system.

Additionally, the implementation of multi-objective NEAT with speciation was not compared to prior multi-objective implementations of NEAT without it. Such a comparison would need to be made if one desired to make any claim of efficiency or desirability for this implementation.

A natural next step for future work, and perhaps the most exciting and open-ended avenue of future work, is to investigate the evolution of other behaviors which soft robots are particularly well suited for. These could include the aforementioned conformation to surfaces or shock absorption, or could take the form of higher-level behaviors relating to the brain-body interactions rather than the body-environment interactions in these brain-body-environment systems. The intersection of body shape, material, control, and environment are rarely studied in conjunction – yet each of these aspects plays a significant role in the behavior of an embodied system, and are desired in future work.

7. CONCLUSION

This work presents the first case of a soft robot evolved to perform a task specifically suited for soft robots. Specifically, this entailed designing creatures to reach or squeeze through a small opening – a task explicitly noted in the literature as one which soft robots are advantageous for. The interpretation of a "small" opening was one in which an equally sized and fully rigid robot would be unable to pass though, further supporting the claim that this task is one for which soft robots are better suited than rigid robots – and thus represents a task for which soft robots should be explicitly designed for. It was found that softer robots were better suited for this task than their (slightly) more rigid counterparts. In optimizing these robots, we also demonstrate a novel implementation of multi-objective NEAT, which relaxes the previous requirement of the removal of diversity maintenance through speciation. While this work serves primarily as an existence proof for evolved squishy robots squeezing through tight spaces, we believe this work also serves as a statement that soft robots should be designed for the tasks in which they excel, and thus we hope that this work opens up a host of questions and future possibilities along this avenue.

8. ACKNOWLEDGMENTS

This work was funded by NASA Space Technology Research Fellowship #NNX13AL37H. Nick would also like to thank the Santa Fe Institute and Ken Stanley for resources, discussions, and feedback during the time of this writing.

9. REFERENCES

[1] J. E. Auerbach and J. C. Bongard. Evolving cppns to grow three-dimensional physical structures. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 627–634. ACM, 2010.

[2] M. Calisti, M. Giorelli, G. Levy, B. Mazzolai, B. Hochner, C. Laschi, and P. Dario. An octopus-bioinspired solution to movement and manipulation for soft robots. *Bioinspiration & biomimetics*, 6(3):036002, 2011.

[3] N. Cheney, J. Clune, and H. Lipson. Evolved electrophysiological soft robots. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 222–229, 2014.

[4] N. Cheney, R. MacCurdy, J. Clune, and H. Lipson. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 167–174. ACM, 2013.

[5] K.-J. Cho, J.-S. Koh, S. Kim, W.-S. Chu, Y. Hong, and S.-H. Ahn. Review of manufacturing processes for soft biomimetic robots. *International Journal of Precision Engineering and Manufacturing*, 10(3):171–181, 2009.

[6] J. Hiller and H. Lipson. Automatic design and manufacture of soft robots. *Robotics, IEEE Transactions on*, 28(2):457–466, 2012.

[7] J. Hiller and H. Lipson. Dynamic simulation of soft multimaterial 3d-printed objects. *Soft Robotics*, 1(1):88–101, 2014.

[8] S. Kim, C. Laschi, and B. Trimmer. Soft robotics: a bioinspired evolution in robotics. *Trends in biotechnology*, 31(5):287–294, 2013.

[9] J. Lehman, K. O. Stanley, and R. Miikkulainen. Effective diversity maintenance in deceptive domains. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 215–222. ACM, 2013.

[10] C. Majidi. Soft robotics: a perspectiveâĂŤcurrent trends and prospects for the future. *Soft Robotics*, 1(1):5–11, 2014.

[11] G. Methenitis. Evolution of soft robots by novelty search. 2014.

[12] J. Rieffel, D. Knox, S. Smith, and B. Trimmer. Growing and evolving soft robots. *Artificial life*, 20(1):143–162, 2014.

[13] J. Schrum and R. Miikkulainen. Evolving multimodal behavior with modular neural networks in ms. pac-man. In *Proceedings of the 2014 conference on Genetic and evolutionary computation*, pages 325–332. ACM, 2014.

[14] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, and K. O. Stanley. Picbreeder: evolving pictures collaboratively online. In *Proc. of the 26th SIGCHI Conf. on Human Factors in Computing Systems*, pages 1759–1768. ACM, 2008.

[15] R. F. Shepherd, F. Ilievski, W. Choi, S. A. Morin, A. A. Stokes, A. D. Mazzeo, X. Chen, M. Wang, and G. M. Whitesides. Multigait soft robot. *Proceedings of the National Academy of Sciences*, 108(51):20400–20403, 2011.

[16] K. O. Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162, 2007.

[17] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.

[18] A. Stilli, H. A. Wurdemann, and K. Althoefer. Shrinkable, stiffness-controllable soft manipulator based on a bio-inspired antagonistic actuation principle. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 2476–2481. IEEE, 2014.

[19] Y. Sugiyama and S. Hirai. Crawling and jumping by a deformable robot. *The International journal of robotics research*, 25(5-6):603–620, 2006.

[20] B. A. Trimmer, A. E. Takesian, B. M. Sweet, C. B. Rogers, D. C. Hake, and D. J. Rogers. Caterpillar locomotion: a new model for soft-bodied climbing and burrowing robots. In *7th International Symposium on Technology and the Mine Problem*, volume 1, pages 1–10. Mine Warfare Association Monterey, CA, 2006.

[21] W. van Willigen, E. Haasdijk, and L. Kester. Evolving intelligent vehicle control using multi-objective neat. In *Computational Intelligence in Vehicles and Transportation Systems (CIVTS), 2013 IEEE Symposium on*, pages 9–15. IEEE, 2013.