Unshackling *Unshackling Evolution:*
Evolving Softrobots with a Multitude of Materials

Connor Cassidy

Abstract

While their industrial counterparts that are programmed to execute particular tasks, softbots are designed to be adaptable and organic. Designing these robots is a hard task, partly due to the fact that the field is relatively new and thus, the process of fabrication and design can be costly and time intensive. This project investigates results found in Clune et al.'s paper [1] on the evolution of soft robots via CPPN-NEAT, *Unshackling Evolution*, extending them to account for softbot designs that are more organic in their morphological nature. Previous work found that increasing the quantity of materials available for use via the generative encoding produced more fit individuals in both the short-term and long run. This paper seeks to explore the validity of that argument. One hypothesis is that adding more materials will have a positive effect on fitness but only to a certain extent where a qualitatively negative shift in computation time and fitness occurs with increasing amount of materials added. Aside, another goal of this paper is to characterize a behavior space that results from increasing materials. We find that there is a point in which increasing the amount of materials given to the network actually hinders the overall and highest fitness of a run. We also find that this problem could be solved by discretizing the search space to account for particular locomotive gaits.

1 Introduction

Clune et al's paper presented some interesting developments that could potentially reduce the time and costs for the design and fabrication of soft robots. Using powerful generative-encoding via CPPN-NEAT, they were able to produce virtual creatures with distinct morphologies and locomotive gaits. Fig. 1 shows the variety in distinct, successful locomotive gaits that were evolved using their strategy. Their results are very promising and they provide great insight into the development of these robots, in fact,

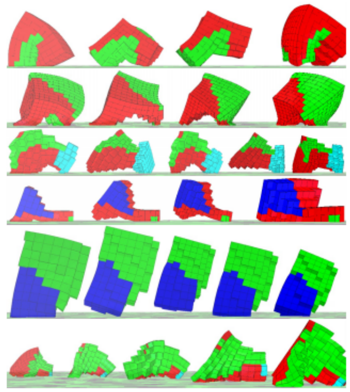

Figure 1: Time series of soft robot gaits found in *Unshackling Evolution*. From top to bottom, they refered to these as L-Walker, Incher, Push-Pull, Jitter, Jumper, and Wings.

after the simulation the robots were reconstructed physically with similar results. Their research showed that the process of evolving these creatures via CPPN-NEAT was much better than opposing strategies such as direct encoding or even human manufacturing and design.

There is one caveat within the research, and that is the obvious lack of material diversity. At most, the generative encoding was deciding at each voxel to choose between four material types (or empty) which could be further reduced to two different types: passive and actuated material voxels. The green and red materials used are actuated materials that go under periods of counter-phase volumetric actuations determined by a coefficient of thermal expansion within the voxel palette. The light and dark blue materials are passive voxels that are softer and stiffer respectively. Even with this limited scope on material type and diversity, they found that an increase in materials produced an increase in performance (See Figure 2).

Clearly, there are limitations that arise from setting a ceiling at four materials. Soft robots aim to produce robots that are more organic in nature in order to increase potential adaptability. One naive hypothesis that may arise is that by letting the CPPN define material types instead of choosing between a palette, we would generate more robust and organic morphologies. One clear issue with this hypothesis is that this may result in a potentially computationally intensive process, while not benefiting from the symmetry that could be acquired from a predefined palette.

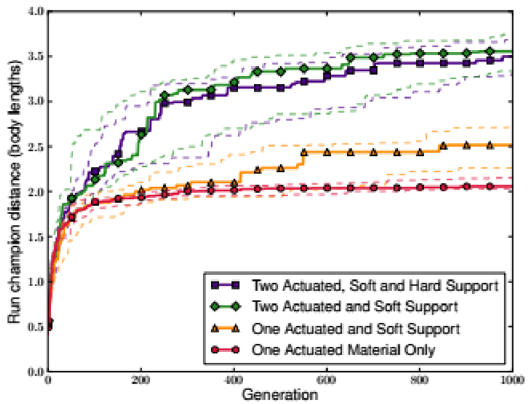

Figure 2: Distance traveled over generations in *Unshackling Evolution* compared with material availibity

Another topic to explore is the characteristics of increasing materials in the pre-defined palette. One hypothesis is that this method will increase computational complexity, but to some extent, this will be weighed out by a spike in fitness. We find that there is a qualitative shift that occurs when increasing the palette size that produces lower best and average fitnesses throughout generations and runs. One potential solution to this problem is explored when seeking to identify locomotive gaits that are produced by the encoding.

2 Methods

2.1 CPPN-NEAT

CPPN-NEAT is a powerful generative encoding developed by Kenneth O. Stanley [2] that uses both compositional pattern-producing networks (CPPNs) and neural evolution of augmented topologies (NEAT). NEAT introduces complexification to topologies that evolve over time. CPPN-NEAT is an extension of standard NEAT that allows it to evolve CPPNs that activate with compositions of canonical, symmetric functions (e.g. sinusoidal, sigmoid, Gaussian) instead of just a single activation across network nodes such as sigmoid.

2.2 VoxCad and Voxelyze

VoxCad and voxelyze are simulators written by Hiller and Lipson in C++. Voxelyze is the underlying software behind VoxCad, which is just a GUI for the voxelyze simulator. Materials were defined based on density, coefficient of friction, and coefficient of thermal expansion (other parameters are available such as Poissons ratio and elasticity). Pre-defined materials for the palette were chosen to be symmetrical nature.

2.3 Experiment

Treatments consisted of 5 runs for 300 generations with a population size of 30. Fitness is measured as is in *Unshackling Evolution*: the difference in the center of mass between initialization and after the end of 10 actuation cycles. There were four treatments, one with the original ceiling of four materials, another with 2 additional actuation cycles, one with 4 more density types that had two more coefficients of friction, and another where the CPPN is let to vary actuations, frictions, and densities of material types on the fly.

3 Results

Adding more actuation cycles to the pre-defined palette increased fitness by only a very small margin, while computation time increased by almost a second per calculation. Including the additional densities caused a small drop in fitness. No interesting morphologies arises from adding more materials or letting the CPPN tweak materials. The latter was a computationally intensive process which produced results that were worse than using a predefined palette.

3.1 4 Materials Revisited

Reproducing the results from *Unshackling Evolution* resulted in pretty similar fitness characteristics, the average fitness for the 5 treatments peaked at about a distance of 3.26. This metric is a bit smaller quantitatively than the original result due to the fact that fewer generations were evolved (only 300 generations were evolved for each treatment in this project, while 1000 were evolved in the original paper). The output also had the same qualitative characteristics as in the original experiment.

3.2 6 Materials

The addition of two actuated materials produced individuals with fitnesses that peaked only a bit higher than those with four materials on materials on average. There was a positive difference of .12 (only a 3% increase) between the average highest overall fitness of the two treatments. The materials palette was extended to include an orange material fixed at a 10% actuation cycle (at half an actuation of the red material), with a yellow material at a counter phase actuation cycle. The characteristic softbot evolved was a galloping cube like creature that used only 5 of the 6 materials.

3.3 10 Materials

This treatment adds two density types to the passive and soft materials that were described in the original experiment. There was a qualitatively negative difference of .23 between the average highest overall fitness of this treatment and the original four material treatment. The CPPN barely took advantage of using these new densities. Creatures produced started to look a bit messier.

3.4 Multi-material CPPN

Here is when we see the biggest computational and qualitative shift in treatments. When given the full palette, the CPPN fails to create proper gaits or organic morphologies. The activation functions would generally make materials and actuations that were symmetric, but did not mimic any organic gaits. The was a qualitatively negative difference of .78 between the highest overall fitness of this treatment and the original four material treatment.

3.5 Complexity

The result of average fitnesses through generations can be seen below in Figure 3. Additionally, the time to complete this treatments on average is shown in Figure 4. The shifts described above in terms of fitness should be easily visible. Note that aside from providing negative shifts in fitness, the ten material and multimaterial treatments also took considerably more time to compute.

3.6 Behavior Space Charaterization

Although increasingly many materials resulted in higher computation times and worse performances for fitness on average, there is still a lot that we can

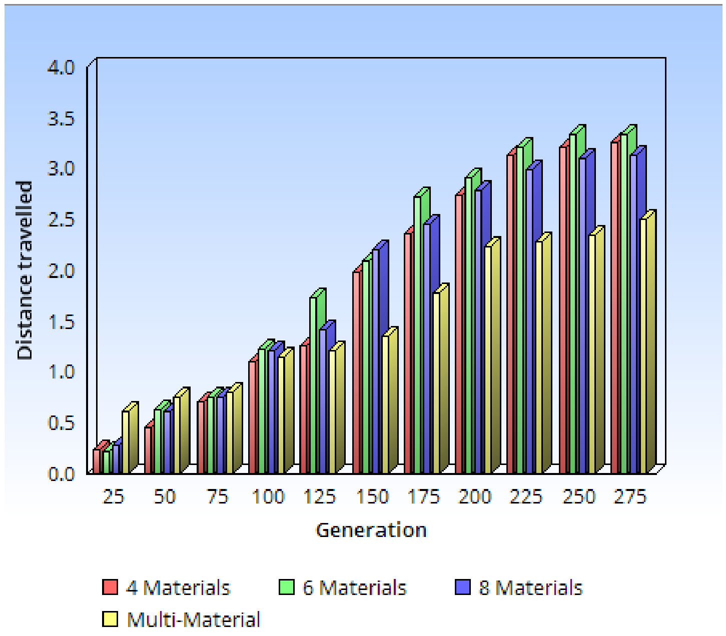

Figure 3: Distance traveled over generations compared with material availibity

take from this. It is important to note how the fitness landscape changes during these shift in material quantity. Intuitively, we can note that as long as reuse the palette from a treatment with less materials, a treatment with additional materials can result in similar morphologies. For example, in run 3 of treatment 2 (6 materials) the highest fitness was achieved by a push-pull robot with the original four materials.

Clearly, the algorithm is able to adapt to incorporate the new materials, but it does not provide a mechanism in which we are assured new gaits or different morphologies. By characterizing the search space, it is easy to see what happens when new materials are added. When new materials are added to the search space, we do not change the previous landscape. Instead, the space is expanded to include the new morphologies and potential gaits that can arise from including additional materials. Due to the nature of its complexity, the landscape was not characterized, but this intuition provides a reason behind the failure of multi-material evolution. Simply put, more materials create a broader landscape and thus more local optima to potentially get stuck in. Also note that the addition of more materials does

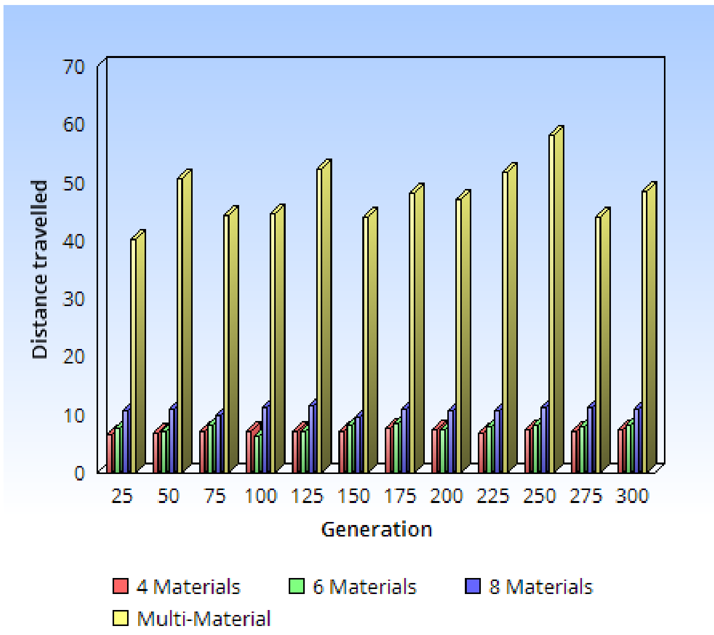

Figure 4: Average computation time for each treatment in a 300 generation run

not assure that these materials will be used by the network, thus creating a more computationally intensive process with same or maybe even worse results.

4 Discussion

It is clear now that the addition of materials is not always beneficial. We see that the addition of more materials can actually hinder the process of softbot evolution. The reason for this is that when adding more materials, we are not given an assurance as to what the resulting morphologies and gaits will be. When the network is given full control of materials to choose from, a large bottleneck occurs from the creation of "garbage" materials that may not even be used. There are two harmful side effects when giving full control of the material choice to the CPPN, and that a potential memory leak or an even larger bottleneck that may occur during garbage collection.

The biggest take away here is that having some human input before the actual evolution process creates better and more organic creatures. This is due to the fact that although the CPPN may produce patterns that are symmetric, we are not always guaranteed locomotive success from this symmetry. One ways to potentially curb this problem is to have locomotive properties or morphologies that are organic and later evolve the species' materials based on this initial guess. [3]

5 Future Work

In the future, we hope that a clear discretization of the search space will help us better understand morphologies and gaits that are possible during the evolution of soft robots. By knowing one of these two things, we eliminate the chance of getting stuck in local optima while also assuring that the materials that we want to use or the locomotive properties we want to achieve are reached without additional, unnecessary computations.

6 Acknowledgements

Thanks to Lisa Meeden for helping with oversight in this paper. Also, special thanks to Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson for inspiring this project.

References

[1] Cheney, MacCurdy, Clune, Lipson. Unshackling Evolution: Evolving Soft Robots with Multiple Materials and a Powerful Generative Encoding *Proceedings of GECCO 2013.*

[2] Kenneth O. Stanley. Compositional pattern producing networks: A novel abstraction of development Genetic programming and evolvable machines, Vol. 8 2007.

[3] Davis Knox, John Rieffel. Scalable Co-Evolution of Soft Robot Properties and Gaits