

On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures

Nick Cheney¹, Josh Bongard², Vytas SunSpiral³ and Hod Lipson⁴

¹Department of Biological Statistics and Computational Biology, Cornell University

²Department of Computer Science, University of Vermont

³Intelligent Systems Division, NASA Ames, SGT Inc.

⁴Department of Mechanical Engineering, Columbia University

nac93@cornell.edu

Abstract

The field of evolved virtual creatures has been suspiciously stagnant in terms of complexification of evolved agents since its inception over two decades ago. Many researchers have proposed algorithmic improvements, but none have taken hold and greatly propelled the scalability of early works. This paper suggests a more fundamental problem with co-evolving both the morphology and control of virtual creatures simultaneously – one cemented in the theory of embodied cognition. We reproduce and explore in greater detail a previous finding in the literature: premature convergence of the morphology (compared to the convergence point of optimizing controllers), and discuss how this finding fits as a symptom of the proposed problem. We hope that this improved understanding of the fundamental problem domain will open the door for further scalability of evolved agents, and note that early findings from our future work point in that direction.

Introduction

In 1994, Karl Sims’ seminal work on “Evolving Virtual Creatures” (Sims, 1994b) created a field of study by that name. This work featured simulated creatures that were able to optimize both their physical layout and their behavioral control strategies for such tasks as terrestrial locomotion, swimming, phototaxis, and competition (Sims, 1994a).

The potential applications of virtual creatures extends beyond their initial contribution to computer graphics and animation, serving as a testbed for the co-optimization of brain-body systems in robotics. With the challenges of continually modifying the morphology of physical robots during the optimization process, the field of Evolutionary Robotics often turns to virtual creatures to optimize morphologies (and their associated controllers) before physical robots are manufactured from the optimized designs (Lund et al., 1997; Funes and Pollack, 1998; Lipson and Pollack, 2000; Nolfi and Floreano, 2002; Doursat et al., 2012; Bongard, 2013).

However, in the two decade lifetime of this field, there have been notable struggles in optimizing creatures, with a very limited ability to extend beyond Sims’ initial works (Geijtenbeek and Pronost, 2012) – despite significant increases in computing power. Many researchers have suggested hypotheses for the cause of this standstill, such as de-

ficiencies in the search algorithms (Hornby, 2006; Lehman and Stanley, 2011; Mouret and Clune, 2015) or genetic encodings (Hornby et al., 2001; Bongard and Pfeifer, 2003). It has also been suggested that the environments/tasks chosen are not complex (or morphologically dependent) enough to necessitate optimization of both the morphology and controller (Auerbach and Bongard, 2014; Cheney et al., 2015). But since we have yet to clearly surpass Sims’ work, each of these hypotheses must be approached with some skepticism.

This work takes note of the particular difficulty in optimizing morphology (Joachimczak et al., 2016) and sets out with the intent of proposing a new hypothesis for the field’s current roadblock. Our hypothesis, unlike many before it, does not rely on more powerful or astute search algorithms to laboriously make our way through the rugged and harsh search space which make optimization of virtual creatures so difficult. Rather, we intend to use our understanding of the behavior of virtual creatures, specifically the theory of embodied cognition, to suggest a fundamental issue in the way that we frame the problem of optimization of virtual creatures – which in turn causes the search landscapes to present such an unpleasant terrain.

The theory of embodied cognition suggests that a fundamental part of the cognitive control process of an individual is being situated (Wilson, 2002). It suggests that the dynamic interactions between a reactive agent and the environment, through sensory-motor feedback loops, are an important driver of behavior (Brooks, 1991), as opposed to cognitivism – the hypothesis that the central functions of mind can be accounted for in terms of the manipulation of symbols according to explicit rules (Anderson, 2003).

This line of reasoning puts an extra emphasis on the morphology of an individual, as it acts as the lens and modulator for all physical communication between that individual’s internal controller and the outside environment (Pfeifer and Bongard, 2006). This work outlines the specific hypothesis that the body’s importance, afforded to it by its role as the connection between internal desires for action and the external consequences of them (as well as external events and the internal sensory observations of them), is understated. With-

out a well established and properly functioning communication channel, the sensory information and motor commands of an individual are ineffective.

From this supposition, we can create a testable hypothesis about the value of the established morphological communication channel. Specifically, control optimization on an existing morphology can be more effective than morphological optimization on a fixed controller – as the latter does not maintain an established communication framework from the controller to the environment (through the morphology). This results in a system which effectively causes large, unintended variations in the behavior of the controller, as its physical interface is constantly being scrambled while optimization seeks to improve the physical shape of the body.

In comparing each of these hypothetical situations to the current state of evolved virtual creatures, we will conclude by discussing a possible connection between this theory of embodied cognition and the lack of effective optimization. Our hope is that such evidence will shed additional light on (at least one of) the problem(s) facing our field, and arm us with the information to help tackle it in future works.

Background

The literature on failed attempts to co-optimize the morphology and control of virtual creatures is sparse. This may be due in part to the bias against publishing negative results (both in submission and acceptance of such findings) (Fanelli, 2011). However, informal conversation with members of the field acknowledge the lack of progress. We note the difficulty of optimizing morphologies in our own virtual creatures (Lipson and Pollack, 2000; Bongard and Pfeifer, 2003; Cheney et al., 2014, 2015) (and unpublished works), but find ourselves grasping for an understanding of why this may be the case.

One clear and concise description of this very problem is expressed in Joachimczak et al. (2016), where they note:

It can also be observed how during the first 100 generations of the evolutionary run, morphological changes occurred very frequently. At generation 125, the overall morphology of the best individual already resembles the best final individual found in the generation 1386 (although its fitness is only 5.07, compared to the 11.15 of the latter). The following generations bring multiple small changes to the morphology of adult form and almost no changes to the larval form. Both stages, however, undergo continuous modifications of their controllers, and it is these alterations that contribute the most to the improvements in fitness. This pattern was also observed in other evolutionary runs: **the final morphology would emerge in the first few hundreds of generations and the remainder of the run would be spent on small tweaks to the bodies and optimization of controllers.** (*emphasis added*)

This notion of premature convergence of morphology is not a stand alone case. At times this premature convergence can be incorrectly interpreted as a positive trait, noted as diversity of results (despite the lack of explicit diversity maintenance), as in Cheney et al. (2013).

In the remainder of this work we set out to reproduce the symptoms described in Joachimczak et al. (2016), where morphology converges prior to control. We seek to further examine and characterize this phenomenon, and describe a theoretical framework which may help to explain its cause.

Methods

Similarly to Joachimczak et al. (2016), we employ soft robots as our instantiation of evolved virtual creatures. We use 3D voxel-based soft robots, following from Cheney et al. (2013), but replace their discrete muscle types and synchronized contractions with voxels which allow individualized phase offsets, consistent with the controllers used in Joachimczak et al. (2016). This allows for behaviors such as propagating waves, which were not possible in Cheney et al. (2013) (but were achieved in Joachimczak et al. (2016) and Cheney et al. (2014)). A global frequency of oscillations is also optimized.

Dual-Network CPPN

We genetically encode the soft robot phenotypes as a network, inspired by the CPPN-NEAT (Stanley, 2007), the algorithm employed by both Cheney et al. (2013) and Joachimczak et al. (2016) (though the later cleverly employs the CPPN alongside development, rather than as an alternative to it). However, this work differs from those two by optimizing two separate networks, one containing only the outputs associated with the physical structure and material placement (“morphology”) of the creatures, while the second network produces only the outputs used to determine the actuation of the muscle voxels (“control”). This allows us to very clearly make variations to either the morphology or the controller, without affecting the genotype of the other¹.

To translate the CPPN genotype to a soft robot phenotype, for each individual voxel in our $7 \times 7 \times 7$ discretized design space, the “presence” output of the morphology network is queried. If the output value (which all span the range $[-1, 1]$) is positive, a voxel is placed there and the “material type” output is queried. If the “material type” output is positive as well, then a the voxel is an active “muscle” cell, otherwise, that voxel is a passive “tissue” cell.

For each active muscle cell, the control network is queried, and the floating point value of the “phaseOffset” output (again from $[-1, 1]$) is assigned as the relative phase offset of that muscle cell (where 0 is exactly in phase with a global clock, -1 and 1 are synchronized a full phase ahead or behind it, and -0.5 and 0.5 are perfectly out of sync with

¹both source code and resulting data are available upon request

it). Finally, the frequency of this global “clock” oscillator is set using the mean value of the “frequency” output across all voxels (including those not currently expressed in the phenotype). In order to easily allow the full range of possible frequencies to be expressed after averaging, a mean value of -0.5 or lower corresponds to the minimal frequency of 5Hz, while a mean value of 0.5 or higher corresponds to the maximal frequency of 10Hz (with linear scaling between them), despite the continued $[-1, 1]$ range of each individual “frequency” output node. The optimization of the global oscillation speed is intended to allow the muscle actuations to resonate with the natural frequency of a given morphology.

We should note that this encoding does allow for morphological changes to affect the expressed control (as the addition or removal of muscle cells will allow more or less of the underlying phase offset pattern to be expressed in the phenotype). Due to the integrated and embodied nature of control, we believe that such an effect would happen with various definitions of “morphology” and “control” – such as a robot with 6 legs expressing a different number of joint control outputs than a 4 legged robot in the rigid body paradigm. This concept of morphology determining the expression of control may be less about this specific implementation and instead a more general consequence of embodied cognition in a situated creature (Pfeifer and Bongard, 2006).

Physics Simulation in VoxCad

Consistent with Cheney et al. (2013), we employ the open-source soft-body simulator VoxCad (Hiller and Lipson, 2014) as the physics engine which determines the fitness of each creature’s phenotype. In order to normalize the number of actuations per muscle cell across creatures with different actuation frequencies, each individual is evaluated for exactly 20 actuation cycles (following a passive initialization period in which it is allowed to settle on the ground in a relaxed pose – intended to discourage passive falling strategies rather than active locomotion behaviors). This means that two creatures with different actuation frequencies will be simulated for different lengths of time. Following the termination of the simulation, the displacement of the creature’s center-of-mass along the positive x axis is returned to the evolutionary algorithm. All other parameters regarding VoxCad simulation are taken from Cheney et al. (2013).

Evolutionary Algorithm

The optimization of these soft robots takes the form of an evolutionary algorithm. The genotype is a directed acyclic graph, represented in memory as a tree to allow an implementation similar to that of genetic programming. Following from CPPNs (Stanley, 2007), each node in the graph sums its weighted inputs and feeds them through a series of nodes with geometric activation functions (here: sigmoid, sine, absolute value, negative absolute value, square, square root, or negative square root) to arrive at each of its output

value(s). The inputs to this network are Cartesian (x, y, z) and polar (r) coordinates of the voxel in question, along with a bias node. The outputs are interpreted as described above.

Variation and selection follow a $(\mu/\rho + \lambda)$ scheme of $(50/25 + 25)$. Variations may be: the addition/removal of a node to a network, addition/removal of an edge between existing nodes, mutation of the weight associated with an edge, or mutation of a node’s activation function. Each of these variations occurred with equal probability, and each variation occurs to only one network of the phenotype, each with equal probability. Crossover was not considered in this work. Variations to the genotype were only considered valid if they resulted in a phenotypic change in the resulting soft robot. Variations were also disallowed if they resulted in creatures who occupied less than 10% of the available voxels, or employed less than 5% of the available voxels as actuated muscle cells. Selection was rank-based with elitism.

Statistical Reporting

All experimental data below represent the mean values of 30 independent runs lasting for 5000 generations each. P-values are calculated using a Mann-Whitney rank-sum test, as we cannot assume normality of fitness values. Confidence intervals were plotted using bootstrapping of 10,000 samples at the 95% confidence level. Significance values are marked with the following convention: *ns* for $p > 0.05$, * for $p \leq 0.05$, ** for $p \leq 0.01$, and *** for $p \leq 0.001$.

Results

First and foremost, we set out to replicate and examine the results found in Joachimczak et al. (2016), where “the final morphology would emerge in the first few hundreds of generations and the remainder of the run would be spent on small tweaks to the bodies and optimization of controllers.”

By visually inspecting the resulting creatures we find that this implementation appears able to reproduce the phenomenon. Fig. 1 shows the optimization over time of the 10 best performing trials. Notice how conserved the morphologies appear to be over time, with the gross morphology generally emerging at or before the 100 generation mark (middle column). While only the top 10 trials are shown for sake of space, this theme applies generally to all the runs.

It is also interesting to note that the top two final-fitness-achieving runs were the only two to undergo a morphological change between generations 1000 and 5000 (the last two columns). This suggests that creatures to which search immediately converges upon are not optimal, and that better performing solutions may not be that far away in phenotypic space (inferred from the similarities between the top two rows at generations 1000 and 5000), yet such creatures appear to be difficult for this search process to find (inferred by the lack of occurrence before generation 1000 in the top two runs, and at all in the next 8 runs). The idea that each run converges to a local, rather than global, optimum is also




Figure 1: Evolved morphologies at various stages in optimization (voxel color from red to green indicates phase offset of controllers). Each row represents one of the top 10 run (out of 30, order by final fitness). Each column represents a point in time during optimization. Note that morphologies generally lock in before gen 100, often on simple forms.

evident by the fact that the set of final creatures differ from one another, rather than converging to the same form.

This visualization serves as an initial indication that the effect of early convergence is apparent in our setup, as it was in Joachimczak et al. (2016). However, it does not demonstrate that the effect of stagnation is more prominently featured by morphology than controllers, or characterize just how detrimental such an effect may be. These two questions are both approached quantitatively in Figs. 2 and 3.

To quantify how early the morphology converges and how detrimental this may be towards the optimization of virtual creatures, we artificially freeze the morphology after a given amount of time, and only allow control variations to occur after this point. If the resulting fitness value does not show a significant change following a morphology freeze at a given time (compared to optimizing both the morphology and control for the entire optimization process), we can be confident

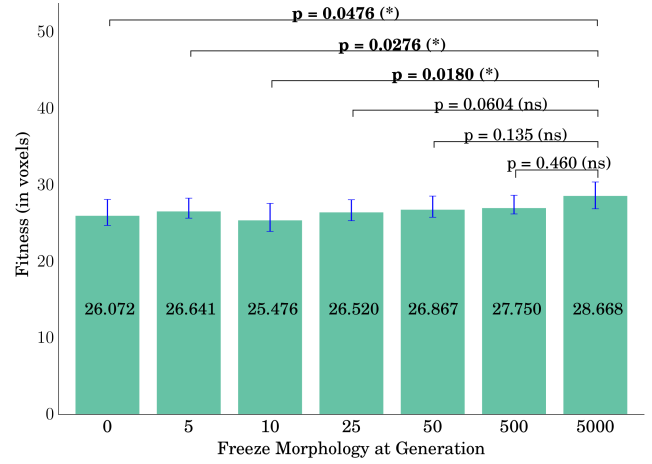


Figure 2: Fitness impact of freezing morphology at various points in optimization. Both morphology and control are optimized up to the freezing point. After it, only control variations are considered for the remainder of the trial. The p-values (and significance markers) reported compare the resulting fitness to that achieved with co-optimization of both morphology and control for the full 5000 generations. Note that morphologies optimized for 25 or more generations show no significant fitness difference, compared to those optimized for all 5000 (noted above in bold).




Figure 3: Fitness impact of freezing control at various points in optimization. Note that controllers with less than 250 generations of optimization (but full morphological optimization) show no significant difference with those optimized for all 5000 gens, suggesting that control mutations continue to provide fitness benefits further into optimization than the morphology variations, which cease to be beneficial to final fitness values much earlier (cf. figure 2, generation 25 – please note the different x-axis compared to that Figure).

that the morphology did not significantly contribute to fitness improvements after that point in optimization time.

Fig. 2 shows the fitness impact of morphology freezes at various times during optimization compared to co-optimization of morphology and control for the entire 5000 generations. We see that full optimization does not show a significant fitness improvement compared to morphologies optimized for 25 generations or more (at the 95% confidence level, as $p \geq 0.0604$ for all freezing points ≥ 25 gens). This means that the morphological variations after generation 25 do not significantly contribute to the fitness of the resulting creatures, suggesting that morphology converges to (near) final forms by generation 25. The visual inspection of these creatures in Fig. 1 does not contradict such a suggestion.

This does not mean that optimization as a whole is converged at this point. Improvements from control optimization occurring after the final gross morphology is fixed are noted in Joachimczak et al. (2016). We also see this effect here, with the fitness resulting from control optimization after morphology freezing (26.520) significantly outperforming ($p < 0.001$) the fitness at the time of freezing (21.157).

Fig. 3 shows the impact of the converse treatment, in which the creature’s controller is frozen at a given point in time and only morphological variations are allowed thereafter. This treatment shows that significant differences in resulting fitness values occur for at least 100 generations (at the 95% confidence level, as $p < 0.001$ for freezing points ≤ 100), but not more than 250 generations ($p \geq 0.0565$ for freezing points ≥ 250). The lack of significant difference past 250 generations also points to early convergence of controllers to (near) final levels early in optimization.

However, the significant drop in fitness from control freezing (at times past those when morphological change stops contributing to final fitness values) suggests that this example of virtual creature evolution creates earlier convergence for morphologies than it does for controllers.

This picture is further reinforced when we examine the time of convergence to a final morphology and controller in each run. On average, the convergence to the final (best of run) morphology occurs at generation 558. In comparison the mutation which leads to the best-ever controller occurs significantly later ($p < 0.001$) at generation 2926. Widening our view from only the final successful variations, and considering all individuals who were the top fitness performers at some point during optimization, we see the same story, with controller mutations leading to top performers continuing significantly later than those created by morphological mutations (mean of gen 750 vs. gen 158, $p < 0.001$). The next section will discuss a potential cause for such an effect.

Discussion

The above results suggest that, in this instance of virtual creatures co-evolving morphology and control, we run into a problem of premature convergence, which is especially pro-

nounced with regard to the morphology of the creature. Premature convergence alone could point to issues in any number of aspects of optimization (diversity maintenance, genetic encoding, etc.). However the difference between optimization effectiveness of morphology and control draws our attention towards the theory of embodied cognition.

Let’s revisit the concept of the morphology as the interface between the control architecture and its effect on the environment. This suggests that modifications to an agent’s morphology will not only change the shape of its body, but also change the way in which its control architecture affects the environment, since the commands sent by that controller will now be interpreted differently – as it affects the actuators of a different body layout. Thus mutations to the morphology of a creature will have the effect of also “scrambling” its controller (causing variation in it) as well.

Contrary to the chain reaction effect of morphological mutations, variations which occur to the controller do not affect any part of the morphology’s relationship with the outside environment. While the control signals which the body is receiving may change, these new commands are still executed in the same framework and “language” as previous commands were. The organization and path of information from controller through morphology to environment causes variations in the morphology to propagate upstream (i.e. affecting the controller/morphology interface in addition to the morphology/environment interface), while variations to the controller do not propagate downstream (affecting the control/morphology interface, but not the morphology/environment interactions).

This feature of embodied cognition has the effect of creating larger (and arguably more unpredictable) behavioral changes to similar sized variations to the “morphology” genome than the “controller” genome. This effect would lead to a more rugged fitness landscape in the space of morphologies (for a given controller) than exists in the fitness landscape of controllers (for a given morphology). We would then predict that a more rugged landscape would lead to more local optima and less efficient optimization with quicker convergence to sub-optimal solutions than in less rugged landscape (Kauffman, 1993). This is consistent with what we have experienced thus far with the optimization of morphology converging prior to control.

Potential Causes and Limitations

There are undoubtedly features of this experimental setup which may cause us to overstate (or understate) the importance of embodied cognition compared to other instances. Firstly, this setup employs soft robots, which are notoriously compliant and adaptive to a wider variety of environmental conditions than their rigid body counterparts (Trivedi et al., 2008). Given that adaptability of this robot-environment interface (in our case to unexpected perturbations in control signals), it’s possible that soft robots dampen this effect. In

the extreme, one may conjecture that the soft robot paradigm is so compliant that almost any morphology can adequately move along flat ground. If this is the case, then it would not be surprising that freezing the morphology on an arbitrary shape has little effect on the resulting fitness value. As soft robots are relatively new to the literature, this may explain why this effect has been unnoticed previously.

In order to further explore this facet, we produced an alternative fitness function which explicitly selects for shape (adding a term to minimize the number of actuated voxels or “energy”). In the extreme this would produce creatures with minimal muscle cells, though since actuated cells directly contribute to locomotion ability, a complex trade-off creates an incentive for specialized energy-efficient morphologies. Another way incentivize to specialized morphologies would be to evaluate the robot in a more complex (and morphologically dependent) task environment than flat ground.

Performing the same “freezing” tests on creatures evolved under the alternative fitness criteria, we see that that freezing morphology continues to show a non-significant effect on fitness at times when control freezing produces a significant fitness drop (e.g. gen 50). Fig 4 visually shows the continued convergence to final gross morphologies (with morphologies at gen 50 generally mirroring those found at gen 5000), as well as the added morphological dependence of the task – as the morphologies demonstrated here visually appear more complex than the more fully occupied shapes in Fig. 1.

In this treatment, we also see the final controllers appearing significantly later (gen 2968) than the final morphologies (gen 419, $p < 0.001$). This is also seen in the average best-so-far individuals, with those produced by control mutations continuing to appear significantly later on average (gen 709) than those produced by morphological variations (gen 119, $p < 0.001$). This data suggests that while the original task was not as “morphologically dependent” as others, the findings still hold in a scenario which puts more of an emphasis on morphological optimization.

A second aspect which may contribute to this effect is the size of the search space. These runs use robots of size $7 \times 7 \times 7$. As each of these voxels can have one of three states (empty, actuated, or passive) which results in $3^{343} = 4.5 \times 10^{163}$ distinct morphological phenotypes. It’s possible that the difficulty in searching the morphology space is due in part to its size. This could explain why this effect was not seen sooner (as previous work in evolutionary robots heavily favors legged morphologies with low degrees of freedom). This phenotype is indirectly encoded, but generative in different ways than previous work evolving morphology (Sims, 1994b; Lehman and Stanley, 2011).

In attempting to reproduce the work from Joachimczak et al. (2016), we optimize phase offset and frequency for an oscillating actuation as the control parameters. These values are encoded by floating point numbers, and thereby create a continuous (theoretically infinite) search space for control.

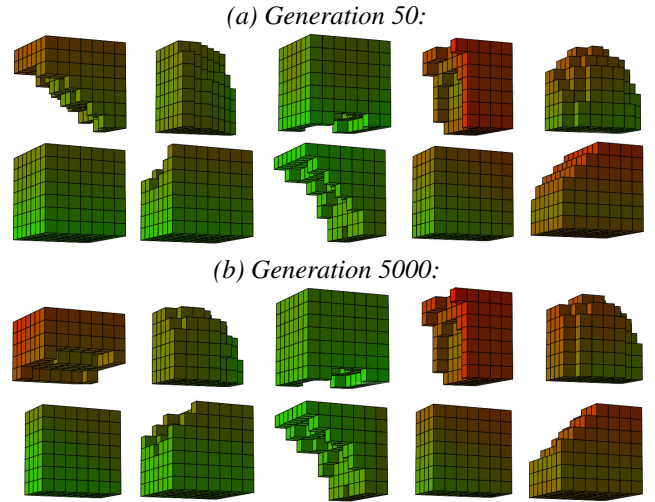


Figure 4: Stagnation shown in the top 10 morphologies under the distance/energy fitness treatment. Note the similarity in gross morphologies from gen 50 (top) to gen 5000 (bottom). The top performing creature shows the largest change between these points, with the new morphology arriving from a mutation at gen 53. Also note the variance and complexity in forms, compared to Fig. 1, suggesting the added morphological dependence of this fitness function.

The concept of discrete physical cells creating a morphology and real valued control parameters (such as neuronal synapse weights) fits biologically – but the differing search spaces give us pause from an optimization perspective.

To create a similar scenario where the size of the controller search space was smaller than that of the morphology, we borrow the two distinct “muscle type” system from Cheney et al. (2013). This allows just two offset control states (implemented by rounding the continuous phase offset values) to create a search space of size $2^{343} = 1.8 \times 10^{103}$ (smaller than the morphology space). In this set of trials, we see the above effect disappear, and morphology no longer appears to be more difficult to optimize than “control”. Here, the final morphological innovation of each run, on average, occurs at generation 665, while control innovation continues only to gen 795 – an insignificant difference ($p = 0.149$). Similarly, the point at which freezing morphology causes a non-significant difference in resulting fitness values no longer occurs before that of controller freezing.

However in this scenario, the line between “morphology” and “control” becomes very blurry. In practice, a two-oscillator-actuation system can be viewed as the placement of cells of these two types (a “morphological” concept) more so than the fitting of phase offset parameters to a predefined placement of muscles (a “control” concept). Thus one could easily argue that the two discrete-phase-offset system from Cheney et al. (2013) should be considered to be en-

tirely morphological optimization, with little to no control to be optimized (as is argued in that paper), and thus immune from our embodied cognition argument.

This is representative of a larger “problem” of this CPPN oscillating actuation setup: that there may not be a clean distinction between “morphology” and “control” to be made, and such divides may be arbitrary labeling. In our example, one could argue that the output node denoting if a cell is actuated or passive should belong in the “morphology” CPPN, as it denoted the placement of different types of cells (“muscles” or “tissues”). But another person could argue equally well that this output belongs in the “control” CPPN, since it does not change the actual shape or stiffness of the creature, and only informs where actuations do or do not occur.

The point here is that virtual creatures are situated and embodied, and thus ideas like embodied cognition or morphological computation (Pfeifer and Gómez, 2009) suggest that there isn’t a clear cut distinction or dualism between two separate pieces (the body and the brain), but rather a single integrated and embodied agent. Therefore we need to consider the tight coupling and interdependencies of the “morphology” and “control” and consider holistic effects whenever we attempt to modify a single part of the system.

Future Work

The results shown in this work are specific only to this instance and experimental setup. Thus, many more instances of this approach (artificially separating morphology from control and freezing each to measure their independent impact on fitness) would need to be attempted on different experimental setups to extrapolate from this single instance. This should ideally include different: morphological encodings (such as the generative block encodings used by Sims (1994b); control architectures (perhaps complexifying to neural nets rather than simplifying to discrete oscillations as we did in our follow up tests – or employing closed-loop control, which may help controllers to adapt to new morphologies); evolutionary algorithms (especially those with a strong emphasis on diversity); tasks (increasing environmental complexity); and/or scales (as increased scales of a cellular creature closer approximate a “continuous” morphology – which comes with various benefits and costs).

Regarding the distinction between “morphology” and “control”, this work necessarily chooses a logical splitting point between the two: representing CPPN outputs that dictate placement of voxels as “morphology” and outputs that dictate voxel size changes as “control”. But this distinction is far from black and white. Future work should explore various groupings of outputs into the categories of “morphology” and “control” (or any grouping names), and examine the effect that such distinctions produce on these results.

The central issue to this paper can be viewed as a problem stemming from the dynamic coupling of control on morphology, with different morphologies creating hills and

valleys in the fitness landscape of controllers. As in any multi-modal landscape, diversity maintenance during search is crucial. This includes diversity coming from crossover (omitted here), or from any existing diversity maintenance method. However, informed by this paper, we would be wise to notice that since hills and valleys of this landscape may be caused *by* the morphology and *onto* the controller, diversity maintenance would do best to focus on protection of diversity within morphologies if it were to encourage the morphological variations (despite their adverse effects on control).

The most important future work would involve potential solutions to this problem. Initial results regarding future work already suggest that our understanding of embodied cognition, and the finding of especially poor mutation success for morphological variations, can inform improved search methods. Specifically, results employing a multi-timescale model, in which morphological mutations are given time to re-adapt their controllers to their new situated forms (and thus conform themselves to their new morphological “communication channels”, thereby “unscrambling” the detrimental effects of the morphological mutation) before the value of these morphological variations are evaluated, shows an improved ability for optimization of virtual creatures compared to traditional methods. This is exactly the type of diversity maintenance that focuses on protecting innovations to the morphology specifically.

Ideally, further algorithmic improvements will occur from embracing the fundamental theory of embodied cognition, but the positive initial results noted here provide confirmation that it’s possible and that the understanding gained from this current work may contribute to future improvements.

Conclusion

We have examined a specific example of co-evolving morphology and control in virtual creatures. In this example, morphology prematurely converges: converging quicker than control, showing a lack of fitness benefits after as little as 25 of the 5000 generations, and with “optimal” final morphologies emerging significantly sooner than final controllers. We have suggested a theoretical basis, founded in the concept of embodied cognition, that could explain such an obstacle and is consistent with the results we present. While there is plenty of work still to be done to solidify this theory, we conclude by suggesting future work based from our newly proposed understanding, and note its striking potential in early initial results. We hope this work will help to explain the difficulty we face in scaling the complexity of evolved virtual creatures, and will help inspire (combined with other efforts) a solution to our current stagnation.

Acknowledgments

Thanks to NASA Space Technology Research Fellowship #NNX13AL37H to N. Cheney for support, Steve Strogatz for feedback on drafts, and Kathryn Miller for copy editing.

References

- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial intelligence*, 149(1):91–130.
- Auerbach, J. E. and Bongard, J. C. (2014). Environmental influence on the evolution of morphological complexity in machines. *PLoS Comput Biol*, 10(1):e1003399.
- Bongard, J. C. (2013). Evolutionary robotics. *Communications of the ACM*, 56(8):74–83.
- Bongard, J. C. and Pfeifer, R. (2003). Evolving complete agents using artificial ontogeny. In *Morpho-functional Machines: The New Species*, pages 237–258. Springer.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial intelligence*, 47(1):139–159.
- Cheney, N., Bongard, J., and Lipson, H. (2015). Evolving soft robots in tight spaces. In *Proceedings of the 2015 on Genetic and Evolutionary Computation Conference*, pages 935–942. ACM.
- Cheney, N., Clune, J., and Lipson, H. (2014). Evolved electrophysiological soft robots. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 222–229.
- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 167–174. ACM.
- Doursat, R., Sayama, H., and Michel, O. (2012). *Morphogenetic engineering: toward programmable complex systems*. Springer.
- Fanelli, D. (2011). Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90(3):891–904.
- Funes, P. and Pollack, J. (1998). Evolutionary body building: Adaptive physical designs for robots. *Artificial Life*, 4(4):337–357.
- Geijtenbeek, T. and Pronost, N. (2012). Interactive character animation using simulated physics: A state-of-the-art review. In *Computer Graphics Forum*, volume 31, pages 2492–2515. Wiley Online Library.
- Hiller, J. and Lipson, H. (2014). Dynamic simulation of soft multimaterial 3d-printed objects. *Soft Robotics*, 1(1):88–101.
- Hornby, G. S. (2006). Alps: the age-layered population structure for reducing the problem of premature convergence. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 815–822. ACM.
- Hornby, G. S., Lipson, H., and Pollack, J. B. (2001). Evolution of generative design systems for modular physical robots. In *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, volume 4, pages 4146–4151. IEEE.
- Joachimczak, M., Suzuki, R., and Arita, T. (in press: 2016). Artificial metamorphosis: Evolutionary design of transforming, soft-bodied robots. *Artificial Life*.
- Kauffman, S. A. (1993). *The origins of order: Self organization and selection in evolution*. Oxford University Press.
- Lehman, J. and Stanley, K. O. (2011). Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 211–218. ACM.
- Lipson, H. and Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974–978.
- Lund, H. H., Hallam, J., and Lee, W.-P. (1997). Evolving robot morphology. In *Evolutionary Computation, 1997., IEEE International Conference on*, pages 197–202. IEEE.
- Mouret, J.-B. and Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.
- Nolfi, S. and Floreano, D. (2002). Synthesis of autonomous robots through evolution. *Trends in cognitive sciences*, 6(1):31–37.
- Pfeifer, R. and Bongard, J. (2006). *How the body shapes the way we think: a new view of intelligence*. MIT press.
- Pfeifer, R. and Gómez, G. (2009). Morphological computation—connecting brain, body, and environment. In *Creating Brain-Like Intelligence*, pages 66–83.
- Sims, K. (1994a). Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372.
- Sims, K. (1994b). Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22. ACM.
- Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162.
- Trivedi, D., Rahn, C. D., Kier, W. M., and Walker, I. D. (2008). Soft robotics: Biological inspiration, state of the art, and future research. *Applied Bionics and Biomechanics*, 5(3):99–117.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic bulletin & review*, 9(4):625–636.