

AUTOMATED DESIGN OF EMBODIED
MACHINES: OPTIMIZATION ALGORITHMS FOR
SOFT ROBOT MORPHOLOGIES AND BEHAVIORS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Nicholas Arthur Cheney

August 2017

© 2017 Nicholas Arthur Cheney

ALL RIGHTS RESERVED

AUTOMATED DESIGN OF EMBODIED MACHINES: OPTIMIZATION
ALGORITHMS FOR SOFT ROBOT MORPHOLOGIES AND BEHAVIORS

Nicholas Arthur Cheney, Ph.D.

Cornell University 2017

The current state of robotics relies largely on hand designed morphologies and controllers. This paradigm of robotics is well suited for controlled and static environments like warehouses or factory floors, but this type of robot often fails to extrapolate to autonomous behaviors in unpredictable and dynamic environments. In contrast, biological animals have evolved to seamlessly interact with the uncertainty of the real world. They accomplish this feat, in part, through specialized and complex morphologies that employ compliant materials. In this work, I explore the interactions of autonomous embodied agents' brains and bodies with each other, and with the outside environment, through the evolution of soft robot morphologies and controllers. These interactions are first explored by evolving robots that perform complex and effective behaviors without high-level controllers in order to demonstrate the potential of morphological computation in compliant bodies. The study of morphological computation is further explored by also demonstrating effective behavior in tasks which are unapproachable with traditional rigid body robots (like squeezing and folding oneself). The focus on morphologically-driven behaviors is extended by evolving soft robots with neural-esque spiking muscles and demonstrating the optimization of physically embodied information pathways, exemplify the continuum between morphologies and controllers in embodied systems. I then turn to the simultaneous optimization of complex morphologies and high-level con-

trollers, using the theory of embodied cognition to hypothesize that the specialization of morphologies and controllers to one another has been hindering the evolution of complex embodied machines. Results here demonstrate that a proposed algorithm for “morphological innovation protection”, which temporarily reduces selection pressure on newly mutated morphologies to enable readaptation of the coupled brain-body systems, produces significantly more fit robots and allows for their sustained optimization over evolutionary time. Generalizing the above methods, the design automation techniques employed here also are applied to problems outside of soft robots – demonstrating the optimization of object topologies towards a desired mechanical resonance. I hope that the work described in this dissertation will help to inform the automated design of embodied machines, like robots, for engineering applications, while also contributing to the fundamental and general understanding of embodied intelligent agents, and their evolution in natural systems.

BIOGRAPHICAL SKETCH

Nick Cheney completed the requirements for a Ph.D. in Computational Biology at Cornell University in 2017, with minors in Cognitive Studies and Behavioral & Evolutionary Neuroscience. He served as a Visiting Fellow at the University of Vermont Complex Systems Center and also as an Affiliated Researcher with the Columbia University Creative Machines Lab from 2015 to 2017. He spent time as a Long-Term Visiting Researcher at the Santa Fe Institute in 2014 and 2015. Nick was also awarded a NASA Space Technology Research Fellowship and served as a Visiting Technologist at NASA Ames from 2013 to 2017.

Prior to his graduate studies, Nick earned a B.S. Magna Cum Laude in Applied Mathematics from the University of Vermont, with concentrations in Complex Systems, Computer Science, Economics, and Business Administration. His undergraduate thesis for the University of Vermont Honors College was titled “Investigations into Morphological Scaffolding”. Before starting graduate school at Cornell, Nick also spent a summer as a Research Assistant with Michigan State University.

In addition to his research, Nick has also contributed to scientific communication and visualization within his field and to the greater public. To date: Nick has given over a dozen invited talks – including competitive and high impact venues such as TEDx – and has won speaking awards for his public research communication. His scientific visualizations have won six peer-reviewed research visualization competitions, and he has co-organized two additional competitions. His work has appeared in popular press outlets such as Discover, NBC, and Popular Science, while his research videos have accumulated over half a million views on YouTube.

This dissertation is dedicated to Arthur H. Cheney Jr.,
who has been waiting long enough for one to call his own.

ACKNOWLEDGEMENTS

Firstly, a big thank you to NASA (Space Technology Research Fellowship #NNX13AL37H) and DARPA (Open Manufacturing Grant #W911NF-12-1-0449) for believing in this research and providing the support to make it happen.

I convey my gratitude to the numerous researchers who have helped to inspire, shape, and support my education and professional journey thus far. This list is far too long to give in its entirety, but includes: Josh Bongard, Peter Dodds, Chris Danforth, Maggie Eppstein, Bob Pepperman Taylor, Karl Sims, Hod Lipson, Jeff Clune, Ken Stanley, Mike Goldstein, Barb Finlay, and Steve Strogatz.

I thank all the friends, roommates, teammates, and labmates who provided endless entertainment and camaraderie. Whether it was dragging me out to the Big Red Barn after work, the Chapter House after a game, countless brewing and grilling sessions, lunchtime hockey, Friday lunches at the statehouse, chats over pho, or just a game night with the gang – you all did your best to keep me from taking life and work too seriously, yet somehow here we are anyway.

But most of all, I need to thank my family, whose support and nurturing has meant more than anything, and is clearly the root cause of any good that I may inadvertently do. I sincerely thank my parents, Stuart and Katy, for believing so strongly in my education and for providing the encouragement and everything necessary to make it happen; to thank my sister, Crystal, for endless laughter and fun; and to thank my grandparents, Art and Millie, for always being there.

I especially need to single out my wife-to-be, Kathryn, and thank her for lending the patience and unselfishness that allowed me to pursue my professional goals; and for the endless amount of support, encouragement, excitement, and silliness that makes me want to change the world, but still be home by five. (Also for being the one to catch my typos – because there were a lot!)

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vi
1 Introduction	1
1.1 Motivation	1
1.2 Background	6
1.3 Contributions	9
1.4 Format of the Following Chapters	12
2 Design Automation of Embodiment-Driven Behaviors in Soft Robots	13
2.1 Introduction	14
2.2 Background	16
2.3 Methods	19
2.3.1 CPPN-NEAT	19
2.3.2 VoxCad	20
2.3.3 Materials	20
2.3.4 GALib	23
2.3.5 Experimental Details	23
2.4 Results	24
2.4.1 Direct vs. Generative Encoding	24
2.4.2 Penalty Functions	26
2.4.3 Material Types	28
2.4.4 Resolution	29
2.5 Discussion	36
2.6 Future Work	38
2.7 Conclusion	40
3 Encouraging Diverse Morphologically-Driven Behaviors with Unique Environments	41
3.1 Introduction	42
3.2 Background	45
3.3 Methods	46
3.3.1 Simulated Task Environment	46
3.3.2 Fitness Metrics	47
3.3.3 CPPN-NEAT	48
3.3.4 Multi-Objective NEAT	51
3.3.5 Run Champions	52
3.3.6 Statistical Analysis	54
3.4 Results	54
3.4.1 Soft/Stiffness Comparison	57

3.4.2	Number of Materials	59
3.4.3	Voxel Penalty	60
3.5	Discussion	62
3.6	Future Work	65
3.7	Conclusion	66
4	Incorporating Neural Information Processing in Morphology-Driven Behaviors	68
4.1	Introduction and Background	69
4.2	Methods	74
4.2.1	CPPN-NEAT	74
4.2.2	Conductive VoxCad	76
4.2.3	Task and Fitness Evaluation	78
4.2.4	Experimental Parameters	79
4.2.5	Statistical Reporting	79
4.3	Results	79
4.3.1	Pacemaker Placement	83
4.3.2	Speed of Pacemaker	84
4.3.3	Touch Sensors	86
4.3.4	Expansion/Contraction Cycle	89
4.4	Discussion	90
4.5	Conclusion	90
5	On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures	91
5.1	Introduction	91
5.2	Background	94
5.3	Methods	96
5.3.1	Dual-Network CPPN	96
5.3.2	Physics Simulation in VoxCad	98
5.3.3	Evolutionary Algorithm	98
5.3.4	Statistical Reporting	99
5.4	Results	100
5.5	Discussion	105
5.5.1	Potential Causes and Limitations	107
5.5.2	Future Work	111
5.6	Conclusion	113
6	Scalable Co-Optimization of Morphology and Control in Embodied Machines	114
6.1	Introduction and Background	115
6.1.1	On the Dichotomy Between Brain and Body Plan	118
6.1.2	Evolving Robot Body Plans and Nervous Systems	119
6.1.3	The Interdependency of Body and Brain	120

6.2	Methods	122
6.2.1	Controller Readaptation	122
6.2.2	Proposed Method: Morphological Innovation Protection	123
6.2.3	Evolutionary Algorithm	124
6.2.4	Genetic Encoding for Soft Robot Morphologies	124
6.2.5	Types of Controllers and their Genetic Encoding	126
6.2.6	Physics Simulation for Evaluation	128
6.2.7	Soft Robot Resolution	129
6.2.8	Morphological Innovation Protection	129
6.2.9	Controller Innovation Protection	130
6.2.10	Morphological Change Threshold	131
6.2.11	Statistical Analysis	131
6.3	Results	132
6.3.1	The Effect of Morphological Innovation Protection on Fitness	132
6.3.2	The Effect of Morphological Innovation Protection on Population Stagnation	136
6.3.3	The Effects of Morphological Innovation Protection on the Progression of Morphologies over Evolutionary Time	141
6.3.4	Generalization to More Complex Implementations	145
6.3.5	The Potential for Convergence Across Initial Conditions	151
6.4	Discussion	157
6.5	Conclusion	162
7	Applications of Morphological Design Automation to Engineering Problems Outside of Robotics	163
7.1	Introduction	164
7.2	Background	167
7.3	Methods	169
7.3.1	Approximating Natural Frequencies with FEM	170
7.3.2	CPPN-NEAT Evolutionary Algorithm	170
7.3.3	Fitness Function	173
7.3.4	Producing Random Target Frequencies	174
7.3.5	Control Treatment	175
7.4	Results	176
7.4.1	Statistical Measures	176
7.4.2	Statistical Data	177
7.4.3	Vibrationally Optimized Beam Examples	182
7.4.4	Fabricated Structures	189
7.5	Discussion	190
7.6	Future Work	191
7.7	Conclusion	192

8	Applications of Morphological Design Automation to Artistic Optimization	193
8.1	Introduction	194
8.2	Background	199
8.3	Results	204
	8.3.1 Target Shapes	206
	8.3.2 Open-Ended Design	206
	8.3.3 Exit Survey Results	207
	8.3.4 Finding Faces	211
8.4	Discussion	212
8.5	Conclusion	217
8.6	Methods	217
	8.6.1 Evolutionary Computation	217
	8.6.2 User Interface and Eye Tracker	219
	8.6.3 Goal-Directed vs. Open-Ended	221
	8.6.4 Control Trials: Mouse Clicking	222
	8.6.5 Subjects	223
9	Argument	224
9.1	Summary of Evidence	224
9.2	Discussion	227
9.3	Future Work	232
9.4	Additional Readings	233
9.5	Closing Remarks	236

CHAPTER 1

INTRODUCTION

1.1 Motivation

Animals are able to effortlessly interact with the world around them, compensating for dynamic environments and unpredictable events like rugged and unstable terrain or other agents. In contrast, the current state of robotics largely consists of machines that are limited to performing repetitive tasks in static and highly controlled environments, like factories or warehouses. The goal of this work is to help understand some of the similarities and differences between biological and artificial agents (both their intrinsic properties and the optimization processes that would allow such properties to come about), and to propose a series of methods which would help allow us to create artificial machines that interact with the world as seamlessly as biological organisms do.

Biological organisms have many advantages over robots that could potentially account for the differences in their abilities. Of these, the complexity and size of the brains of many animals (humans especially) is often emphasized as the cause of their increased intelligence [296, 305, 236]. The emphasis on the brain (rather than the body) as the source of intelligent behavior has been long held, and is often traced to Cartesian dualism – in which Descartes argues that the mind and body are distinct, and that the mind alone can exhibit thought (with or without its body) [224, 242, 77]. This idea is further formalized in philosophy as the mind-body problem [30].

The focus on the brain as the primary source of intelligent behavior and cog-

nition also appears in the study of artificial intelligence, and can be traced back to connectionism (the idea that intelligence is an emergent property arising from the combination of many simple components like neurons) [90, 120] and perhaps even more so from computationalism (the focus on symbolic reasoning and computational programs run on internal mental representations) [231, 100].

However, an alternate hypothesis also exists for the nature and function of cognition in embodied agents. The theory and study of embodied artificial intelligence emphasizes the role that the body plays in shaping the function, efficiency, and thought processes of intelligent agents [26, 27, 276, 121, 3, 49, 221, 219]. Considering (and optimizing for) the role of the body in intelligent behavior is especially relevant for the study of robots – as we desire machines that physically interact with their environments in a seamless manner. Insights gained into intelligent behavior also hold promise to better inform our understanding of the function and evolutionary design of the great diversity and efficiency of body plans found in the natural world.

The benefits of embodiment can take effect in many forms. In the case of morphological computation [222, 220], the organization of the body may allow for coordinated and sophisticated behavior with little or no neural input, as is demonstrated by passive walking robots [59]. Compliant bodies can help to simplify sensorimotor pathways – acting as a high-pass filter for smoothing and integrating sensory information over time, or working as an underactuated muscle to reduce the degrees of freedom necessary for motor commands [219]. It's also been argued that the complex physical interactions of a body with the environment can serve as the source of non-linear bases for reservoir computing [119].

This focus on the role of embodiment also emphasizes more low-level and instinctual “animalistic” behaviors, such as walking and locomotion, where physical interactions between the body and the environment are rich and apparent. As the goal of this work is to understand and design autonomous interactions with uncertain environments, the compromise of focusing solely on low-level behaviors is an acceptable one. Since animals are arguably far more proficient than machines at these tasks, there is still much to be learned in this restricted domain. Though it should be noted that the overarching concepts of embodied cognition, and the morphological influence on learning can also be applied to higher-level tasks as well [27, 306, 3, 219, 255, 22].

I am far from the first to focus on animalistic behavior as the foundation of artificial intelligence. These animal inspired robots, or – as Wilson called them – *animats* have long been studied [307, 198, 281], as they provide a simple and straightforward toybox for the design of autonomous embodied agents.

This idea of a smooth continuum of forms, brain, and behaviors from simple animals all the way up to humans is supported by the observation of highly conserved regularities in the development and evolution of both brains and bodies across many species [95, 50, 33, 233], providing even more confidence that this foray into low-level animalistic behavior will lead to insights that may eventually scale up to human-level behavior.

This work also focuses largely around the idea of automated design. The elegance of an algorithmically designed agent comes ultimately from the evolutionary and developmental design of biological organisms. Though the practical application of automated design should not be understated.

There do exist rare instances of engineered robots with the ability to overcome unpredictable environments – the most notable of these exceptions is the Boston Dynamics Big Dog [234]. However, this robot took tens of millions of dollars and years of manpower to meticulously research and design. For almost all applications, such an investment would prohibit the construction of such a robot. The ability to design a machine that is specialized for any given environment in a matter of minutes or hours rather than years may open many potential doors for the robot design at smaller scales. The ever increasing accessibility of high performance cloud computing [145] and additive manufacturing (e.g. in-home or for-hire 3D-printing services [191, 180, 308, 235]) have resulted in significant reduction in costs for specialized design and manufacturing, making such endeavors economical at small scales.

The added efficiency of design automation also comes from the differences in the inspiration for these bio-inspired robots. For traditional biologically inspired robots (e.g. [295, 143, 234, 158, 111, 310]), inspiration is drawn from an existing biological organism. This approach has the benefits of existing proof-of-concepts and piggybacking on millions of years worth of evolutionary design. However it comes at the cost of having to understand and reverse engineer all the relevant details of a biological organism (including understanding it fully enough to determine which details are not relevant) and recreating similar effects in a different artificial substrate. This painstaking task must also be replicated for each unique pairing of robot design and biological organism.

In contrast to this, the design automation tools in this work draws inspiration from the design algorithm of evolution itself. While understanding and replicating this algorithm is still a large challenge (and the subject of many ongo-

ing works – including this one), the resulting design tools can then be applied to create a wide variety of robots for a wide variety of tasks and environments with little to no additional investments. These robots produced via these tools also have the benefit of being optimized for their specific tasks and environments – reducing the risk of specific behaviors or forms not translating well from natural to man-made environments and tasks, while also allowing for the optimization towards new and unique tasks which biological evolution has never previously encountered.

Thus the focus on automated design specifically thereby addresses a current need in the market of engineered machines. The lack of cheap and distributed robotic design tools is currently the limiting factor for an industry and practice that has seen the barriers of entry come down with the recent popularity of distributed manufacturing and cloud computing. I feel that the design tools demonstrated in this work represent the potential for completing the loop of distributed and in-home design and manufacture of embodied machines (and other engineered objects).

Finally, I hope that the study of the evolution of embodied machines may also feed knowledge from computer science and engineering back into the biological sciences – including evolutionary biology, psychology, neuroscience, anatomy and physiology, and many more.

I draw motivation from a fascination with the evolutionary process itself. The idea that a “blind watchmaker” [72] would be able to create the diversity, complexity, and efficiency of the plants and animals that we see on earth today seems completely unintuitive and awe-inspiring. Despite the potential for catastrophic environmental shocks [251] or mass extinctions from invasive

species takeover [51], despite the cut-throat battle taking place between the genes in our own DNA [73], despite the unlikelihood of finding beneficial mutations through random genome transcription errors [300, 40], and despite the physical world's preference for entropy and chaos [32] life not only finds a way to survive – but it prospers, grows, and complexifies over time.

While it may be impossible to study many parallel controlled and independent replications of biological evolution, and highly unethical to mutate the bodies or brains of living animals to study its effects on their cognition and development, these tasks are straightforward to do in the simulation of evolving and developing embodied agents. It is my hope that these artificial agents can serve as a controlled toybox for asking these theoretical biological questions (e.g. [89, 174, 57]).

1.2 Background

The study of shape and form in biological organisms has long been a fascination for researchers. This dissertation coincides with the centennial of the first edition of Thompson's pioneering work in mathematical biology *On Growth and Form* [287]. Computer scientists have attempted to replicate biological forms as far back as Turing's explorations into *The Chemical Basis of Morphogenesis* [294].

The understanding of the evolutionary design of these forms in nature was widely introduced by Darwin's work *On the Origin of Species* [70], and this understanding has been complemented by additional facets since then [260, 109, 266, 159, 110, 300, 33, 73]. The conceptual framework for a computational instantiation of this idea again dates back to Turing's "genetical or evolutionary

search” [292]. The early implementations of a number of variants upon this idea is credited to a wide array of pioneering works [16, 102, 4, 101, 238, 252], and the popularization of *genetic algorithms* is often credited to Holland [131].

The use of these evolutionary and genetic algorithms to automate the design of robots (deemed *Evolutionary Robotics*) is credited to a burst of activity in the early 90’s [177, 99, 118]. Though the tendency of the field, including these early pioneers, is to only design controllers for hand-specified robot body plans [209, 23].

The first popular instance of extending the ideas of evolutionary robots to include the shape and form of a robot (optimizing its body plan, sensor, and motor placements) was Sim’s *Evolved Virtual Creatures* [263, 264]. While this work was outstanding for its time, the methods employed have not scaled well with the increase of computing power, and many in the field still consider the original results from this pioneering work – now decades old – to be the pinnacle for the evolutionary design of morphologies and controllers in evolutionary robotics, as the scale of evolved morphologies have not obviously increased since this early work [105, 23].

Despite the lack of complexification in the evolved morphologies of embodied machines, in the decades since Sims’ work there has continued to be tremendous achievements in the evolution and understanding of evolved morphologies in the field – often coming at the interfaces of other engineering disciplines.

For example, Lipson and Pollack demonstrated the first example of the automated design and construction of a real-world physical robot through the combined use of morphological design and 3D-printing [181]. The L-Systems

employed by Hornby and Pollack to evolve virtual creature morphologies [139] were employed to design the real-world physical structure of an X-band antenna for the NASA Space Technology 5 spacecraft [182].

Pfeifer et al. review the implications of embodiment on both low-level and high-level behaviors, such as self-stabilization and “how the body shapes the way we think” [224, 219, 223]. Bongard shows how morphological change can accelerate learning and lead to more robust behaviors [22].

Doursat et al. review a subset of evolutionary robotics deemed “morphogenetic engineering” [83, 84]. This subfield consists of self-assembling, swarming, developmental, grammar generating morphologies. Developmental robotics are particularly relevant to the creation of complex morphologies in this work [150, 67, 147, 82, 68] (though – due to computational limitations – the genetic encoding employed here relies on “exploiting regularity without development” [274]).

It should also be noted that the term “embodied evolution” is often used to specify the use of the real world as a simulator for a population of evolving robots – such as in [303, 302] – rather than to refer to the optimization of the morphology of a robot (either in simulation or reality), as I employ (and focus this review on) for this work.

The focus on morphological computation from complex morphologies in this work is also complemented by the presence and optimization of complex materials. This work employs soft robots in particular, as their compliance and (theoretically) infinite degrees of freedom enable rich non-linear physical interactions with the environment. While the prior work on the evolution of mor-

phologies and controllers for soft robotics is limited [241, 240], recent interest in soft robotics has resulted in a number of fully-hand-designed soft morphologies and controllers [258, 167, 259, 157, 213, 288, 17].

1.3 Contributions

This work serves as a contribution to the field of embodied artificial intelligence by exploring the automated design of morphology-driven behaviors and examining the role of morphological computation in these optimized agents. It also explores how morphologically-driven behaviors can be combined with traditional high-level neural controllers, and explores the distinction between the two regimes of robotic control – attempting to blur the line between the two and further expose the false dichotomy of “body” and “brain”. The content behind these arguments are laid out over the following chapters:

- Chapter 2 introduces the problem and methods for the evolution of soft robot morphologies – demonstrating the extreme of minimal control and maximal morphological computation. It does so by evolving “embodiment-driven behaviors” which result solely from the placement of various passive soft tissue and oscillating muscle cells – showing the extent of effective and realistic-appearing behaviors in the absence of a high-level controller.
- Chapter 3 notes that the task optimized for above (locomotion over flat ground) is one that can be effectively solved by soft robots, but that it is not a task that necessarily relies on the complex interactions of the robots’ material properties and the environment to perform. This chapter em-

employs the methods introduced in Chapter 2 for a task specific to soft robots, demonstrating that one of the potential advantages of soft compliant robot bodies is their ability to contort their bodies and squeeze through small openings – a task which is not possible for rigid body robots.

- Chapter 4 extends the notion of morphological computation by further blurring the distinction between high-level information processing and embodiment-driven behaviors in these soft robots. This chapter incorporates low-level computation into the morphology of the robots by allowing explicit information passing (in the form of “electrical impulses”) between neighboring cells, as well as non-linear action-potentials (“neuron-like spikes”) within the muscle cells of the robot. The potential to perform computations similar to those found in a neural network, and to dictate how sensorimotor information and commands are sent to different parts of the robot’s body seeks to mimic the function of high-level control using only the low-level morphology of the robot.
- Chapter 5 attempts to incorporate the optimization of higher-level distributed controllers in addition to the morphological design demonstrated above. This chapter points out some of the potential problems that prevent traditional evolutionary algorithms from solving the coupled co-optimization problem of evolving both the controller and morphology of an embodied agent – and empirically demonstrates the lack of sustained optimization in this type of system. Specifically, I hypothesize that the specialization of the morphology and the controller of an agent to each other creates a fragile coupled system. This specialization means that any mutation to the morphology or controller of an agent (even those which would be beneficial in the long run) break the tight coupling between these two

subcomponents and result in an immediate detriment to fitness – which results in that mutation being discarded from the evolving population.

- Chapter 6 addresses the co-optimization of robot morphologies and controllers, by proposing an algorithm which explicitly accounts for the problems and hypothesized underlying causes proposed in the Chapter 5. Specifically, I address the problem of discarding mutations that break the specialization of controllers to their accompanying morphologies and lead to short term fitness drops. I address this problem by proposing a diversity maintenance mechanism that allows newly mutated morphologies the opportunity to readapt to their new morphologies. I demonstrate that this method leads to more efficient and sustained optimization. This finding provides confidence in both the presence of the previously hypothesized issues, and our ability to confront them – opening the door for effective co-optimization of morphology and control in embodied artificial agents.
- Chapter 7 provides an example of the application of the above method for morphological design to a problem of structural engineering, suggesting the potential for wider generalization of the techniques developed here for tasks outside of robot locomotion. Specifically, we address the problem of resonant frequency tuning in object design – a highly non-linear and unintuitive design problem. I demonstrate the ability to optimize the first ten resonant frequencies of a simulated fixed-free beam by evolving the placement of stiff and soft materials in its construction (a problem that is nearly impossible to solve by hand), and note how the optimization of mechanical resonance of an object may help engineer better objects for structural robustness or energy harvesting applications.

- Chapter 8 exemplifies an additional example application of morphological design – this time showcasing the design of artwork or consumer products. This chapter also demonstrates the use of human-in-the-loop selection mechanisms, and alternative user interface methods like eyetracking – showcasing the extension of the above design techniques to domains in which specifying a set of objective evaluation criteria *a priori* is not possible.

1.4 Format of the Following Chapters

The results of this dissertation are organized around seven papers submitted to peer-reviewed conferences and journals (Chapters 2 through 8). These results are presented in their entirety, and in the order which best represents how they fit into the goals and progression of the automated design of embodied machines – which is not necessarily the order in which they were originally published. Chapter 9 then provides a discussion on the cohesive argument supported by these individual works. That chapter also points the reader towards additional readings in the form of seven other papers which were published alongside the main argument during the course of this dissertation – often in collaboration with others – to complement and support the results included here.

CHAPTER 2
DESIGN AUTOMATION OF EMBODIMENT-DRIVEN
BEHAVIORS IN SOFT ROBOTS

Abstract of Chapter ¹

In 1994 Karl Sims showed that computational evolution can produce interesting morphologies that resemble natural organisms. Despite nearly two decades of work since, evolved morphologies are not obviously more complex or natural, and the field seems to have hit a complexity ceiling. One hypothesis for the lack of increased complexity is that most work, including Sims', evolves morphologies composed of rigid elements, such as solid cubes and cylinders, limiting the design space. A second hypothesis is that the encodings of previous work have been overly regular, not allowing complex regularities with variation. Here we test both hypotheses by evolving *soft robots* with *multiple materials* and a *powerful generative encoding* called a compositional pattern-producing network (CPPN). Robots are selected for locomotion speed. We find that CPPNs evolve faster robots than a direct encoding and that the CPPN morphologies appear more natural. We also find that locomotion performance increases as more materials are added, that diversity of form and behavior can be increased with different cost functions without stifling performance, and that organisms can be evolved at different levels of resolution. These findings suggest the ability of generative soft-voxel systems to scale towards evolving a large diversity of complex, natural, multi-material creatures. Our results suggest that future work that com-

¹Appeared as: Cheney, N., MacCurdy, R., Clune, J., & 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 (pp. 167-174). ACM.

bines the evolution of CPPN-encoded soft, multi-material robots with modern diversity-encouraging techniques could finally enable the creation of creatures far more complex and interesting than those produced by Sims nearly twenty years ago.

2.1 Introduction

In 1994, Karl Sims' evolved virtual creatures showed the potential of evolutionary algorithms to produce natural, complex morphologies and behaviors [263]. One might assume that nearly 20 years of improvements in computational speed and evolutionary algorithms would produce far more impressive organisms, yet the creatures evolved in the field of artificial life today are not obviously more complex, natural, or intelligent. Fig. 2.2 demonstrates an example of similar complexity in robots evolved 17 years apart.




Figure 2.1: An example of a natural looking morphology and behavior evolved by combining a generative encoding with voxel-resolution soft, actuatable materials. The soft robot gallops from left to right across the image with a dog-like gait.

One hypothesis for why there has not been a clear increase in evolved complexity is that most studies follow Sims in evolving morphologies with a limited set of rigid elements [171, 7, 6, 134, 181]. Nature, in contrast, composes organisms with a vast array of different materials, from soft tissue to hard bone, and




Figure 2.2: (left) The scale and resolution of robots evolved by Sims in 1994 [263]. (middle) The scale and resolution at which evolutionary robotics commonly occurs today (from Lehman and Stanley, 2011 [171]). (right) The scale and resolution of robot fabrication techniques (from Lipson and Pollack, 2000 [181]).

uses these materials to create sub-components of arbitrary shapes. The ability to construct morphologies with heterogeneous materials enables nature to produce more complex, agile, high-performing bodies [290]. An open question is whether computational evolution will produce more natural, complex forms if it is able to create organisms out of many material types. Here we test that hypothesis by evolving morphologies composed of voxels of different materials. They can be hard or soft, analogous to bone or soft tissue, and inert or expandable, analogous to supportive tissue or muscle. Contiguous patches of homogeneous voxels can be thought of as different tissue structures.

Another hypothesis is that the encodings used in previous work limited the design space. Direct encodings lack the regularity and evolvability necessary to consistently produce regular morphologies and coordinated behaviors [56, 53, 273, 134], and overly regular indirect encodings constrict the design space by disallowing complex regularities with variation [134, 271, 273]. We test this hypothesis by evolving morphologies with the CPPN-NEAT encoding [271], which has been shown to create complex regularities such as sym-

metry and repetition, both with and without variation (Fig. 2.3). CPPN-NEAT has shown these abilities in 2D images [253] and 3D objects [54] and morphologies [7]. To test the impact of the CPPN encoding, we compare it to a direct encoding.

Overall, we find that evolution does utilize additional materials made available to it; their availability led to a significant amount of diverse, interesting, complex morphologies and locomotion behaviors without hindering performance. Furthermore, the generative encoding produced regular patterns of voxel ‘tissue’, leading to fast, effective locomotion. In contrast, the direct encoding produced no phenotypic regularity and led to poor performance.

Because it is notoriously difficult to quantify attributes such as “impressiveness” and “complexity”, we make no effort to do so here. Instead, we attempt to visually represent the interesting diversity of morphologies and behaviors that evolved once evolution was provided with more materials and a sophisticated encoding. We also demonstrate the ability for this system to scale to higher resolutions and greater material diversity without hindering performance. Finally, we investigate the effects of different fitness functions, revealing that evolution with this encoding and material palette can create different bodies and behaviors in response to different environmental and selective pressures.

2.2 Background

There are many Evolutionary Robotics papers with rigid-body robots [208]. However, few attempts have been made to evolve robots composed of soft materials [241], and most of those attempts are limited to only a few components.

This paucity is due largely to the computational costs of simulating flexible materials and because many genetic encodings do not scale to large parameter spaces [21, 147].




Figure 2.3: (left) Examples of high resolution, complex, natural-looking images evolved with CPPN-NEAT that contain symmetry, repetition, and interesting variation [253]. (right) Examples of CPPN-encoded 3D shapes with these same properties [54].

The CPPN encoding abstracts how developmental biology builds natural complexity, and has been shown to produce complex, natural-appearing images and objects (Fig. 2.3) [253, 54, 271]. Auerbach and Bongard used this generative encoding to evolve robotic structures at finer resolutions than previous work. The systems evolved demonstrated the ability to take advantage of geometric coordinates to inform the evolution of complex bodies. However, this work was limited to rigid building blocks which were actuated by a large number of hinge joints [9, 7, 6], or had no actuation at all [5].

Rigid structures limit the ability of robots to interact with their environments, especially when compared to the complex movements of structures in

biology composed of muscle and connective tissue. These structures, called muscular hydrostats, often display incredible flexibility and strength; examples from biology include octopus arms or elephant trunks [290]. While soft robots can be designed that provide outstanding mobility, strength and reliability, the design process is complicated by multiple competing and difficult-to-define objectives [290]. Evolutionary algorithms excel at such problems, but have historically not been able to scale to larger robotic designs. To demonstrate that evolution can design complex, soft-bodied robots, Hiller and Lipson created a soft-voxel simulator (called VoxCad) [122]. They showed a preliminary result that CPPNs can produce interesting locomotion morphologies, and that such designs can transfer to the real world (Fig. 2.4) [124]. However, this work did



Figure 2.4: A time-series example of a fabricated soft robot, which actuates with cyclic 20% volumetric actuation in a pressure chamber [124]. This proof-of-concept shows that evolved, soft-bodied robots can be physically realized. Current work is investigating soft robot actuation outside of a pressure chamber.

not take advantage of the NEAT algorithm, with its historical markings, speciation, crossover, and complexification over time - which have been shown to greatly improve the search process [275]. Additionally, these preliminary results consisted of only three trials per treatment. Here we conduct a more in-depth exploration of the capabilities of CPPNs when evolving soft robots in VoxCad.

2.3 Methods

2.3.1 CPPN-NEAT

CPPN-NEAT has been repeatedly described in detail [271, 56, 54, 104], so we only briefly summarize it here. A compositional pattern-producing network (CPPN) is similar to a neural network, but its nodes contain multiple math functions (in this paper: sine, sigmoid, Gaussian, and linear). CPPNs evolve according to the NEAT algorithm [271]. The CPPN produces geometric output patterns that are built up from the functions of these nodes. Because the nodes have regular mathematical functions, the output patterns tend to be regular (e.g. a Gaussian function can create symmetry and a sine function can create repetition). In this paper, each voxel has an x , y , and z coordinate that is input into the network, along with the voxel's distance from center (d). One output of the network specifies whether any material is present, while the maximum value of the 4 remaining output nodes (each representing an individual material) specifies the type of material present at that location (Fig. 2.5). This method of separating the presence of a phenotypic component and its parameters into separate CPPN outputs has been shown to improve performance [299]. Robots can be

produced at any desired resolution. If there are multiple disconnected patches, only the most central patch is considered when producing the robot morphology.

2.3.2 VoxCad

Fitness evaluations are performed in the VoxCad soft-body simulator, which is described in detail in Hiller and Lipson 2012 [125]. The simulator efficiently models the statics, dynamics, and non-linear deformation of heterogeneous soft bodies. It also provides support for volumetric actuation of individual voxels (analogous to expanding and contracting muscles) or passive materials of varying stiffness (much like soft support tissue or rigid bone). For visualization, we display each voxel, although a smooth surface mesh could be added via the Marching Cubes algorithm [184, 54].

2.3.3 Materials

Following [123], there are two types of voxels: those that expand and contract at a pre-specified frequency, and passive voxels with no intrinsic actuation, which are either soft or hard. We expand upon [123] to include multiple phases of actuation. Unless otherwise noted, four materials are used: Green voxels undergo periodic volumetric actuations of 20%. Light blue voxels are soft and passive, having no intrinsic actuation, with their deformation caused solely by nearby voxels. Red voxels behave similarly to green ones, but with counter-phase actuations. Dark blue voxels are also passive, but are more stiff and resistant to

deformation than light blue voxels. In treatments with less than 4 materials, voxels are added in the order above (e.g. two material treatments consist of green and light blue voxels).




Figure 2.5: A CPPN is iteratively queried for each voxel within a bounding area and produces output values as a function of the coordinates of that voxel. These outputs determine the presence of voxels and their material properties to specify a soft robot.




Figure 2.6: CPPN-NEAT-encoded soft robots can scale to any resolution. Pictured here are soft robots sampled at voxel resolutions of $5 \times 5 \times 5$ (top-left), $10 \times 10 \times 10$ (top-right), and $20 \times 20 \times 20$ (bottom).

2.3.4 GALib

The direct encoding is from GALib—fully described in [301]—a popular off-the-shelf genetic algorithm library from MIT. In the direct encoding genome, each voxel has its own independent values representing its presence and material outputs. The first value is binary, indicating whether a voxel at that position exists. If the voxel exists, the highest of the material property values determines the type of voxel. Thus, a $10 \times 10 \times 10$ (“ 10^3 ”) voxel soft robot with 4 possible materials would have a genome size of $10^3 \times 5 = 5000$ values.

2.3.5 Experimental Details

Treatments consist of 35 runs, each with a population size of 30, evolved for 1000 generations. Unless otherwise noted, fitness is the difference in the center of mass of the soft robot between initialization and the end of 10 actuation cycles. If any fitness penalties are assessed, they consist of multiplying the above fitness metric by: $1 - \frac{\text{penalty metric}}{\text{maximum penalty metric}}$. For example, if the penalty metric is the number of voxels, an organism with 400 non-empty voxels out of a possible 1000 would have its displacement multiplied by $1 - \frac{400}{1000} = 0.6$ to produce its final fitness value. Other CPPN-NEAT parameters are the same as in Clune and Lipson 2011 [54].

2.4 Results

Quantitative and qualitative analyses reveal that evolution in this system is able to produce effective and interesting locomoting soft robots at different voxel resolutions and using different materials. We also discover that imposing different environmental challenges in the form of penalty functions provides an increased diversity of forms, suggesting the capability to adapt to various selective pressures.

Videos of soft robot locomotion are available at <http://tinyurl.com/EvolvingSoftRobots>. So the reader may verify our subjective, qualitative assessments, we have permanently archived all evolved organisms, data, source code, and parameter settings at the Dryad Digital Repository.

2.4.1 Direct vs. Generative Encoding

The CPPN-NEAT generative encoding far outperforms the direct encoding (Figure 2.8), which is consistent with previous findings [56, 53]. The most stark difference is in the regularity of the voxel distributions (compare Figs. 2.1, 2.6, 2.12, 2.13 to Fig. 2.7). CPPN-NEAT soft robots consist of homogeneous patches of materials akin to tissues (e.g. one large patch of muscle, another patch of bone, etc.). The direct encoding, on the other hand, seems to randomly assign a material to each voxel. These homogeneous tissue structures are beneficial because similar types of voxels can work in a coordinated fashion to achieve the locomotion objective. For example, all the voxels in one large section of green voxels will expand at the same time, functioning as muscle tissue. This global




Figure 2.7: A representative example of a soft robot evolved with a direct encoding. Note the lack of regularity and organization: there are few contiguous, homogeneous patches of one type of voxel. Instead, the organism appears to be composed of randomly distributed voxels. The resolution is the default 10^3 .

coordination leads to jumping, bounding, stepping, and many other behaviors. In the direct encoding, each voxel works independently from—and often at odds with—its neighboring voxels, preventing coordinated behaviors. Instead, final organisms appear visually similar to those at initialization, and performance barely improves across generations (Figure 2.8).

Another reason for the success of the CPPN-NEAT encoding is one of the key

properties of the NEAT algorithm: it starts with CPPN networks that produce simple geometric voxel patterns and *complexifies* those patterns over time [271].

2.4.2 Penalty Functions

To explore performance under different selective or environmental pressures, we tested four different penalty regimes. All four require the soft robot to move as far as possible, but have different restrictions. In one environment, the soft robots are penalized for their number of voxels, similar to an animal having to work harder to carry more weight. In another, the soft robots are penalized for their amount of actuatable material, analogous to the cost of expending energy to contract muscles. In a third treatment, a penalty is assessed for the number of connections (adjoining faces between voxels), akin to animals that live in warm environments and overheat if their surface area is small in comparison to their volume. Finally, there is also the baseline treatment in which no penalties are assessed.

While a cost for actuated voxels does perform significantly worse than a setup with no cost ($p = 1.9 \times 10^{-5}$ comparing final fitness values), all treatments tend to perform similarly over evolutionary time (Fig. 2.9). This rough equivalence suggests that the system has the ability to adapt to different cost requirements without major reductions in performance. However, drastically different types of body-plans and behaviors evolved for the different fitness functions. There are differences in the proportions of each material found in evolved organisms, indicating that evolution utilizes different material distributions to fine tune morphologies to various environments (Fig. 2.10). For example, when

no penalty cost is assessed, more voxels are present ($p < 2 \times 10^{-13}$). When there is a cost for the number of actuated voxels, but not for support tissue, evolution uses more of these inert support materials ($p < 0.02$).

More revealing are the differences in behaviors. Fig. 2.11 categorizes locomotion strategies into several broad classes, and shows that different task requirements favor different classes of these behaviors. To limit subjectivity in the categorization process, we made clear category definitions, as is common in observational biology, and provide an online archive of all organisms for reader evaluation (see Sec. 7.4).

Fig. 2.12 displays the common locomotion strategies and Fig. 2.11 shows how frequently they evolved. They are described in order of appearance in Fig. 2.12. The L-Walker is named after the “L” shape its rectangular body forms, and is distinguished by its blocky form and hinge-like pivot point in the bend of the L. The Incher is named after its inchworm like behavior, in which it pulls its back leg up to its front legs by arching its back, then stretches out to flatten itself and reach its front legs forward. Its morphology is distinguished by its sharp spine and diagonal separation between actuatable materials. The Push-Pull is a fairly wide class of behaviors and is tied together by the soft robot’s powerful push with its (often large) hind leg to propel itself forward, which is usually coupled with a twisting or tipping of its front limb/head to pull itself forward between pushes. The head shape and thinner neck region are surprisingly common features. Next, the Jitter (or Bouncer) moves by bouncing its (often large) back section up and down, which pushes the creature forward. It is distinguished by its long body and is often composed mainly of a single actuatable material. The Jumper is similar in that it is often comprised of a single

actuatable material, but locomotes in an upright position, springing up into the air and using its weight to angle its jumping and falling in a controlled fashion to move forward. The Wings is distinguished by its unique vertical axis of rotation. It brings its arms (or wings) in front of it, then pushes them down and out to the sides, propelling its body forward with each flapping-like motion. Fig. 2.13 demonstrates other, less-common behaviors that evolved.

These example locomotion strategies display the system's ability to produce a diverse set of morphologies and behaviors, which likely stems from its access to multiple types of materials. Our results suggest that with even more materials, computational evolution could produce even more sophisticated morphologies and behaviors. Note that different behaviors show up more frequently for different task settings (Fig. 2.11), suggesting the ability of the system to fine tune to adapt to different selective pressures.

2.4.3 Material Types

To meet its full potential, this system must scale to arbitrarily large numbers of materials and resolutions. We first explore its ability to compose soft robots out of a range of materials by separately evolving soft robots with increasing numbers of materials (in the order outlined in Sec. 2.3.3). Adding a second, and then a third, material significantly improved performance (Fig. 2.14, $p < 2 \times 10^{-6}$), and adding a further hard, inert material did not significantly hurt performance (Fig. 2.14, $p = 0.68$). This improved performance suggests that CPPN-NEAT is capable of taking advantage of the increase in morphological and behavioral options. This result is interesting, as one might have expected a

drop in performance associated with the need to search in a higher dimensional space and coordinate more materials.

2.4.4 Resolution

This system also is capable of scaling to higher resolution renderings of soft robots, involving increasing numbers of voxels. Fig. 2.6 shows example morphologies evolved at each resolution. The generative encoding tended to perform roughly the same regardless of resolution, although the computational expense of simulating large numbers of voxels prevented a rigorous investigation of the effect of resolution on performance. Faster computers will enable such research and the evolution of higher-resolution soft robots.




Figure 2.8: The best individuals from 35 independent runs with a direct or generative encoding. Note how the generative encoding sees large improvements early in evolution, while it is exploring new locomotion types. It then settles on specific types and gradually improves coordination, timing, etc., to exploit a given strategy. The direct encoding is unable to produce globally coordinated behavior to develop new locomotion strategies, resulting in very minor improvements as it exploits its initial random forms. Here, and in all figures, thick lines are medians $\pm 95%$ bootstrapped confidence intervals.




Figure 2.9: Performance is mostly unaffected by different selection pressures (i.e. fitness functions).




Figure 2.10: The amount of each material that evolved for different cost functions, revealing the system’s ability to adapt material distributions to different environments. For example, without a cost, evolution used more voxels to produce actuation ($p < 2 \times 10^{-13}$). With a cost for actuated voxels, evolution tends to use more inert support tissue ($p < 0.02$).




Figure 2.11: Common behaviors evolved under different cost functions, summed across all runs. These behaviors are described in Sec. 2.4.2 and visualized in Fig. 2.12. Some behaviors occur more frequently under certain selective regimes. For example, the L-Walker is more common without a voxel cost, while Jitter, Jumper, and Wings do not evolve in any of the no cost runs.




Figure 2.12: Time series of common soft robot behaviors as they move from left to right across the image. From top to bottom, we refer to them as L-Walker, Incher, Push-Pull, Jitter, Jumper, and Wings. Fig. 2.11 reports how frequently they evolved.




Figure 2.13: Time series of other evolved strategies. (top) Opposite leg stepping creates a traditional animal walk or trot. (middle) A trunk-like appendage on the front of the robot helps to pull it forward. (bottom) A trot, quite reminiscent of a galloping horse, demonstrates the inclusion of stiff material to create bone-like support in longer appendages.

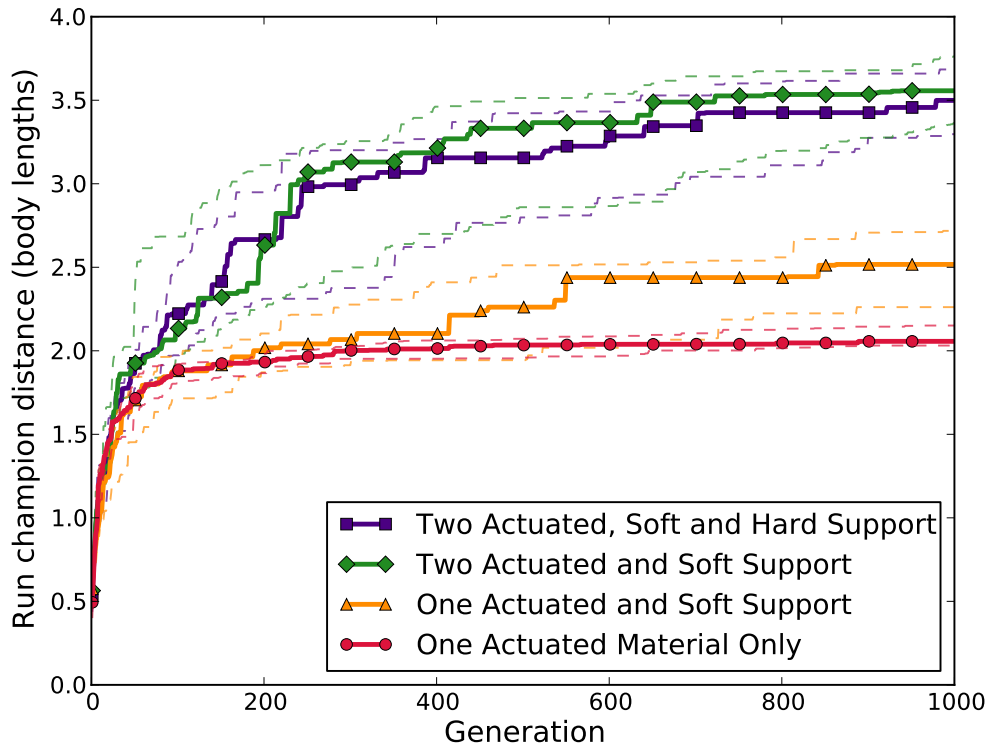


Figure 2.14: The number of materials also affects performance. With only one, only simple behaviors like Jumping or Bouncing are possible, so performance peaks early and fails to discover new gaits over time. Upon adding a second material, more complex jumping and L-Walker behavior develops. When a second actuatable material is added, most behavior strategies from Fig. 2.12 become possible. Adding a stiff support material broadens the range of possible gaits, but is only rarely taken advantage of (such as in the bottom gallop of Fig. 2.13) and thus has a minimal impact on overall performance. These observational assessments may be verified, as all evolved organisms are available online (Sec. 7.4)

2.5 Discussion

The results show that life-like, complex, interesting morphologies and behaviors are possible when we expand the design space of evolutionary robotics

to include soft materials that behave similarly to organic tissue or muscle, and search that design space with a powerful generative encoding like CPPN-NEAT. Our preliminary experiments suggest that soft robotics at the voxel resolution will someday provide complex and breathtaking demonstrations of lifelike artificial forms. Soft robotics will also showcase the ability of evolutionary design because human intuitions and engineering fare poorly in such entangled, non-linear design spaces.

We challenged multiple scientists to design fast, locomoting soft robots by hand, using the same resolution and materials. While the sample size is not sufficient to report hard data, all participants (both those with and without engineering backgrounds) were unable to produce organisms that scored higher than the evolved creatures. Participants noted the surprising difficulty of producing efficient walkers with these four materials. This preliminary experiment supports the claim that systems like the CPPN-NEAT generative encoding will increasingly highlight the effectiveness of automated design relative to a human designer.

This work shows that the presence of soft materials alone is not sufficient to provide interesting and efficient locomotion, as soft robots created from the direct encoding performed poorly. Our results are consistent with work evolving rigid-body robots that shows that generative encodings outperform direct encodings for evolutionary robotics [139, 161, 56, 53]. Unfortunately, there have been few attempts to evolve robot morphologies with CPPN-NEAT [5], and there is no consensus in the field of a proper measurement of “complexity”, “interestingness”, or “natural” appearance, so we cannot directly compare our soft robots to their rigid-body counterparts. However, we hope that the reader

will agree about the potential of evolved soft robots upon viewing the creatures in action [<http://tinyurl.com/EvolvingSoftRobots>].

2.6 Future Work

The ability to evolve complex and intricate forms lends itself naturally to other questions in the field. Auerbach and Bongard have explored the relationship between environment and morphology with rigid robots in highly regular environments [7]. Because our system allows more flexibility in robot morphology and behavior, it may shed additional, or different, light on the relationship between morphology, behavior, and the environment. Preliminary results demonstrate the ability of this system to produce morphologies well suited for obstacles in their environments (Fig. 2.15).

While our research produced an impressive array of diverse forms, it did use a target-based fitness objective, which can hinder search [312]. Switching to modern techniques for explicitly generating diversity, such as the MOLE algorithm by Mouret and Clune [202, 55] or algorithms by Lehman and Stanley [171], has the potential to create an incredibly complex and diverse set of morphologies and behaviors.

Additionally, we are currently pursuing methods to minimize the need for expensive simulations and to evolve specific material properties instead of having a predefined palette of materials. These avenues are expected to allow increased complexity and diversity in future studies.

The HyperNEAT algorithm [272], which utilizes CPPNs, has been shown




Figure 2.15: An example of a soft robot that has evolved “teeth” to hook onto the obstacle rings in its environment and propel itself across them.

to be effective for evolving artificial neural network controllers for robots [56, 169, 53]. The same encoding from this work could thus co-evolve robot controllers and soft robot morphologies. Bongard and Pfeifer have argued that such body-brain co-evolution is critical toward progress in evolutionary robotics and artificial intelligence [219].

Soft robots have shown promise in multiple areas of robotics, such as gripping [128] or human-robot interaction [245]. The scale-invariant encoding and soft actuation from this work has potential in these other areas of soft robotics as well.

In order to compare different approaches, the field would benefit from general, accepted definitions and quantitative measures of complexity, impressive-

ness, and naturalness. Such metrics will enable more quantitative analyses in future work.

2.7 Conclusion

In this work we investigate the difficult-to-address question of why we as a field have failed to substantially improve upon the work of Karl Sims nearly two decades ago. We show that combining a powerful generative encoding based on principles of developmental biology with soft, biologically-inspired materials produces a diverse array of interesting morphologies and behaviors. The evolved organisms are qualitatively different from those evolved in previous research with more traditional rigid materials and either direct, or overly regular, encodings. The CPPN-NEAT encoding produces complex, life-like organisms with properties seen in natural organisms, such as symmetry and repetition, with and without variation. Further, it adapts to increased resolutions, numbers of available materials, and different environmental pressures by tailoring designs to different selective pressures without substantial performance degradation. Our results suggest that investigating soft robotics and modern generative encodings may offer a path towards eventually producing the next generation of impressive, computationally evolved creatures to fill artificial worlds and showcase the power of evolutionary algorithms.

CHAPTER 3
ENCOURAGING DIVERSE MORPHOLOGICALLY-DRIVEN
BEHAVIORS WITH UNIQUE ENVIRONMENTS

Abstract of Chapter¹

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 design of soft robots has shown examples of these squishy creatures performing traditional robotic tasks 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.

¹Appeared as: Cheney, N., Bongard, J., & Lipson, H. (2015). Evolving Soft Robots in Tight Spaces. In Proceedings of the 2015 annual Conference on Genetic and Evolutionary Computation (pp. 935-942). ACM.

3.1 Introduction

Recent interest and developments in the study of soft robotics [31, 48, 189, 258, 277, 278, 289] 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 [43, 124, 197, 240]. 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” [157].

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.




Figure 3.1: (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.

make explicit the intuitive notion that “Ultimately it is probably the ecological niche that determines the evolutionary tendency to be stiff or soft” [157].

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. 3.1).

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.

3.2 Background

The notion of soft robots being good at squeezing through 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 $13.6\text{cm} \times 5.9\text{cm} \times 0.6\text{cm}$) [258]. 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 [43, 124, 197, 240], 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.

3.3 Methods

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

3.3.1 Simulated Task Environment

In this work (and consistent with [43, 124, 197]), these soft robots are simulated in the soft-body physics simulator VoxCad [127]. 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.

3.3.2 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 fit-

ness 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 as a method to 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 this study).

In both cases, creatures were also incentivized along a second objective – to maximize their size (Sec. 3.3.4). 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).

3.3.3 CPPN-NEAT

Consistent with [43, 197], we allow the soft robots to optimize their topology using the CPPN-NEAT evolutionary algorithm [5, 271].

The CPPN encoding represents the voxel phenotype as a network. This network takes a voxel’s relative coordinates as inputs, and transforms this infor-




Figure 3.2: 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. 3.3.1) is put in place for the fitness evaluation.

mation 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 [253].

A sketch of this genotype to phenotype decoding is provided in Fig. 3.2. 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.

3.3.4 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 through the aperture.

While various implementations of multi-objective NEAT have appeared recently [172, 297, 250], we believe that our implementation 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 [275]). 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 dif-

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 multi-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.

3.3.5 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. 3.3.4), 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.

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.

3.3.6 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.

3.4 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 from 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 3.3) 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




Figure 3.3: (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.

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. 3.4 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 creature 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. 3.5, 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 through an opening smaller than the width of its body.

3.4.1 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. 3.1. 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.

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




Figure 3.4: (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.

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.

3.4.2 Number of Materials

A surprising result from Figs. 3.1, 3.3, and 3.4 is that the creatures commonly appear to be made out of only a single material. This is unlike previously published results [43], where a figure is presented to demonstrate the variety and consistency of the material compositions for locomotion over flat ground.

Fig. 3.5, 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 multi-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.

3.4.3 Voxel Penalty

Also in slight disagreement with previously work in this area [43], 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 dis-




Figure 3.5: (*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.)

tance 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.

3.5 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.

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 vice 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 [43] 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. 3.3) 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 result 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 peri-

ods. 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).

3.6 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 [43]. 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 [38]. 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.

3.7 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 through, 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 counter-

parts. 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.

CHAPTER 4
INCORPORATING NEURAL INFORMATION PROCESSING
IN MORPHOLOGY-DRIVEN BEHAVIORS

Abstract of Chapter ¹

The embodied cognition paradigm emphasizes that both bodies and brains combine to produce complex behaviors, in contrast to the traditional view that the only seat of intelligence is the brain. Despite recent excitement about embodied cognition, brains and bodies remain thought of, and implemented as, two separate entities that merely interface with one another to carry out their respective roles. Previous research co-evolving bodies and brains has simulated the physics of bodies that collect sensory information and pass that information on to disembodied neural networks, which then processes that information and return motor commands. Biological animals, in contrast, produce behavior through physically embedded control structures and a complex and continuous interplay between neural and mechanical forces. In addition to the electrical pulses flowing through the physical wiring of the nervous system, the heart elegantly combines control with actuation, as the physical properties of the tissue itself (or defects therein) determine the actuation of the organ. Inspired by these phenomena from cardiac electrophysiology (the study of the electrical properties of heart tissue), we introduce *electrophysiological robots*, whose behavior is dictated by electrical signals flowing through the tissue cells of soft robots. Here we describe these robots and how they are evolved. Videos and images of these

¹Appeared as: Cheney, N., Clune, J., & Lipson, H. (2014). Evolved Electrophysiological Soft Robots. In the Proceedings of the International Conference on the Synthesis and Simulation of Living Systems (Vol. 14, pp. 222-229).

robots reveal lifelike behaviors despite the added challenge of having physically embedded control structures. We also provide an initial experimental investigation into the impact of different implementation decisions, such as alternatives for sensing, actuation, and locations of central pattern generators. Overall, this paper provides a first step towards removing the chasm between bodies and brains to encourage further research into physically realistic embodied cognition.

4.1 Introduction and Background

The fields of evolutionary robotics and artificial life have seen a great deal of emphasis on embodied cognition in recent years [[43, 23, 239, 7, 124, 171, 9, 5, 223, 138, 181]]. There is even a paradigm called *embodied cognition*, which argues that the specifics of the embodiment (such as the morphology) are vital parts of the resulting behavior of the system: It argues that the co-evolutionary connection between body and brain is more deeply intertwined than the body simply acting as a minimal physical interface between the brain and the environment [[219]].

Recent work in evolutionary robotics has shown that complex behaviors can arise when co-evolving bodies and brains. At one end of the spectrum, [5] demonstrated the evolution of physical structures that had no joints or actuators, and evolved to cover the largest distance in a controlled fall due to gravity. While that work exemplifies the evolution of behavior emerging from morphology alone, it does not co-evolve any actuation or control. [9] then evolved the placement of CPG controlled rotational joints between cellular spheres, thus co-

evolving morphology and control.

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

Many examples in the literature include the co-evolution of a robot morphology with an artificial neural network controller [[263, 181, 138, 171]]. These studies (and many more like them) involve what might be called “ghost” networks: artificial neural networks that provide control to the body, yet do not have any physical embodiment in the system they control. The state of input nodes to these networks is often set by sensors in the robot and output nodes typically signify behavioral outcomes in the actuators, but the computation is done supernaturally, disjoint from the body itself.

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

While the brains of animals are often a separate module within their bodies, animals also have central and peripheral nervous systems extending throughout their bodies. An extreme example of this is the octopus, which has as much as 90% of its neurons existing outside of its central nervous system [[316]]. The distributed and physical layout of the nervous system over space may con-




Figure 4.1: Current flowing through an evolved creature. The legend for voltage within each cell (colors) is given in Fig. 4.3.

tribute significantly to neural processing, as the delays and branching in axons (the basis for nerves) are suggested to serve computational functions [[254]].

Despite the prevalence of embodied, distributed circuitry in nearly all of animal life, the idea of an embodied nervous system has been absent from the field of evolutionary robotics. The sub-field called Evolvable Hardware evolves physical circuits for computer chips [[98]], but such work has not been applied to evolving the circuitry of artificial life organisms. We are unaware of work with virtual creatures that have physically embodied control systems (e.g. where neural circuitry physically runs throughout the body of the creature). We present the first such work in this paper.

We propose a very basic model of electrical signal propagation throughout the body of an evolved creature. This embodied controller is based on electrophysiology (specifically at large scales, such as cardiac electrophysiology, Fig. 4.2). Electrophysiology is the study of the electrical properties of biological cells and tissues [[129]]. In this model, electrical pulses from a single centralized sinusoidal pacemaker (analogous to the sinoatrial node – the pacemaker in the heart [[28]]) are propagated through the electrically conductive tissue of the creature. The location and patterning of this conductive tissue is described by

an evolved Compositional Pattern Producing Network (CPPN) genome. Evolution controls the shape of the body and the electrical pathways within it, which both combine to determine the robot's behavior.

The model involves conductive tissue cells that collect voltage from neighboring cells, causing an action potential (spike) if the collected voltage exceeds the cell's firing threshold (Fig. 4.3). Once this threshold is crossed, the cell depolarizes, causing a voltage spike that excites neighboring cells. This voltage spike is followed by a refractory period, during which the cell is temporarily unable to be re-excited.

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

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




Figure 4.2: An example of complex electrical wave propagation in cardiac modeling [[93]].

determined by a CPPN genome.

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

4.2 Methods

4.2.1 CPPN-NEAT

The evolutionary algorithm employed in this study is CPPN-NEAT. This algorithm has been previously described in detail [271, 274, 5, 43], so it is only briefly described here.

A Compositional Pattern Producing Network (CPPN) [[271]] is variation of an Artificial Neural Network (ANN) [[196, 98]] where each node can have one of many mathematical functions as an activation function (e.g. sine, cosine, Gaussian, sigmoid, linear, square, or positive square root) instead of being limited to a sigmoid activation function. In CPPN-NEAT, a design space is discretized into individual locations (in this case a 3D space is discretized into a $10 \times 10 \times 10$ grid of voxels, for 1000 total voxels). The CPPN is queried once per voxel to determine the phenotypic state at that location (in this case, whether a voxel is present and, if so, the material type). The inputs to the CPPN network for each location are different: specifically, they include one input node for each dimension of the space (here, reporting the x , y , and z values of that location), as well an input that reports the distance (d) from the center to the location. The network also features output nodes for each material property. There are three in this study: one node specifies if a voxel exists at the queried location, the second decides if the material at that location is conductive, and the third decides whether or not the material is an actuated muscle (the latter two only matter if the voxel is present).

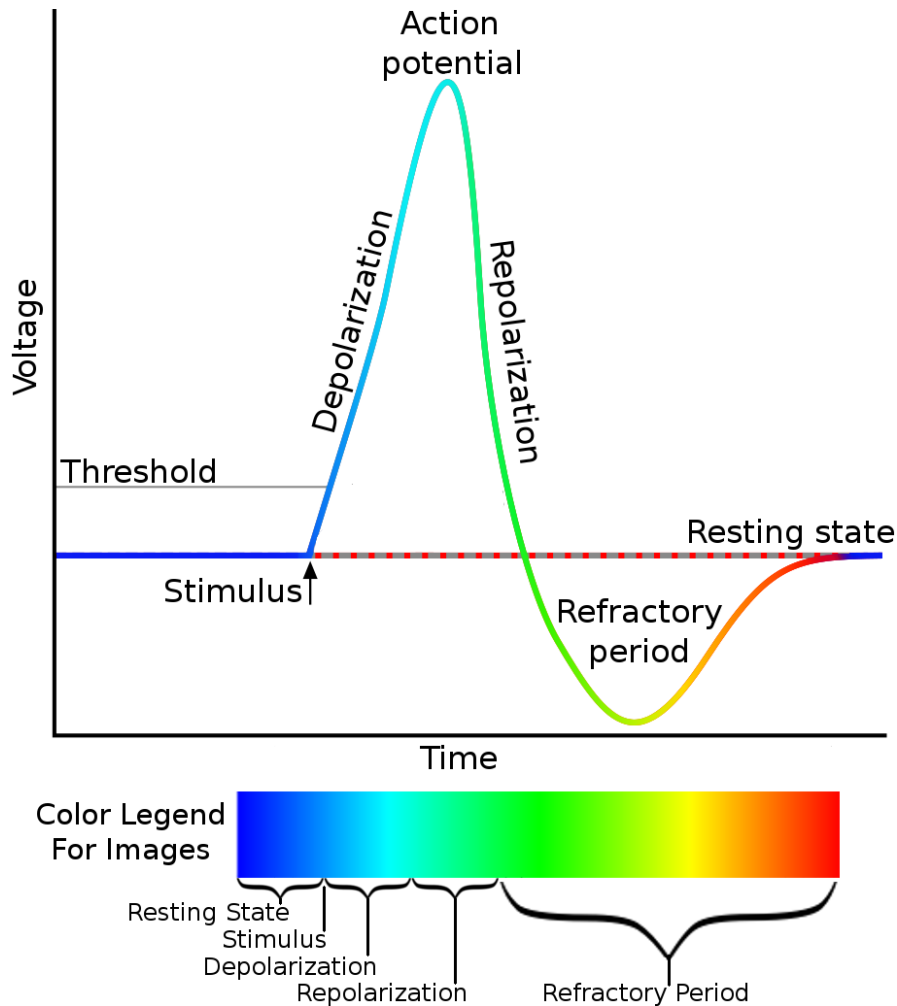


Figure 4.3: (*top*) A depiction of an action potential. Notice how the voltage is below the threshold until a stimulus event (such as a pacemaker or neighboring cell spike) pushes the voltage to the threshold value. Once this threshold is met, a voltage spike occurs via a process called depolarization. The cell excites its immediate neighbors during this process. Following the action potential, the cell enters a fixed length refractory period, during which it is physically unable to produce a new action potential. Finally, the cell returns to its resting state, able to start the process again when a new stimulus arrives. (*bottom*) These phases of the action potential cycle are mapped to the color code used to visualize the soft robots in Figs. 4.1, 4.4, and 4.5. Image licensed via Creative Commons.

4.2.2 Conductive VoxCad

Fitness is evaluated in the VoxCad soft body simulator [[127]]. Its dynamics have previously allowed the evolution of complex and lifelike behaviors in soft robots, as it can simulate muscle contractions [[43]]. Further details about VoxCad can be found in [125].

This work adds electrophysiology to VoxCad by adding a simple action-potential model, acting on the scale of a single voxel (analogous to a cell). Each voxel has an immediate membrane potential level (the difference between the electric potential inside and outside the cell), as well as a threshold membrane potential level. In an action-potential model (Fig. 4.3), a cell's resting potential is below that of the threshold potential. When the membrane potential reaches its threshold value, the cell depolarizes, causing a spike in the cell's membrane potential and voltage.

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

A cell's action potential (starting with the beginning of the depolarization

phase in Fig. 4.3) triggers a sinusoidal expansion/contraction of that cell with a maximum amplitude of 39% linear expansion per voxel side.

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

Cells may be of any of the following types: empty, conductive muscles, insulating muscles, conductive passive tissue, insulating passive tissue, or a pacemaker cell. Near the center point of the discretized $10 \times 10 \times 10$ design space, a lone pacemaker is placed (cell number 555 out of 1000). Analogous to the sinoatrial node in cardiac electrophysiology, this pacemaker node serves as the source of electric stimulation for the entire system. Insulating cells are similar to the cells explained above, except that they are unable to accept current from neighboring cells and thus never reach their threshold potential or produce an action potential.

In this model, the refractory period lasts 5 times as long as the depolarization period. This means that at least 5 voxels must separate the leading edges of two serial action potential waves. Since the pacemaker is placed in the center of the $10 \times 10 \times 10$ space, approximately one wave of action potentials would exist at any given time in a setup with a uniform cube of entirely conductive material –

where a wave of action potentials would propagate uninterrupted, with a new one starting around the time the first reaches the outer edge of the space. We chose this setup to encourage the evolution of static gaits, which can be more robust and transferable to reality than dynamic gaits [[18]].

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

4.2.3 Task and Fitness Evaluation

Following [43], we evolve these electrophysiological robots for locomotion over flat ground. This simple task and environment make fitness evaluation easy. Despite its simplicity, the task is a classic problem in the field, and has been repeatedly shown to produce an array of complex morphologies and interesting behaviors [[43, 53, 56, 10, 171]].

Each creature is simulated for 20 times the length of an expansion/contraction cycle. Its displacement between the starting coordinates and the creature's final center of mass (in the xy plane) is recorded. In an effort to discourage designs that might excite as many cells as possible, and to encourage designs with sparse spindles of connectivity (similar to the peripheral nervous system), the distance traveled is multiplied by $1 - \frac{(\# \text{ of conductive cells})}{1000}$. Thus the fitness function incentivizes minimizing the amount of conductive tissue and maximizing the distance traveled. While a multi-objective technique may be ideal in finding the optimal tradeoff between these goals, we follow previous CPPN-NEAT research in using this single, multi-part fitness function [[43, 8]].

4.2.4 Experimental Parameters

Unless otherwise noted, each treatment described below consists of 48 independent runs (with identical initial conditions across treatments). Each run consists of a population size of 30 individuals evolved for 1000 generations. Unless otherwise noted, all other parameters are consistent with [43].

4.2.5 Statistical Reporting

Because the data are not normally distributed, all plots show median fitness (thick, center lines) bracketed by two thin lines that represent 95% bootstrapped confidence intervals of the median [[268]]. For the same reason, all p -values are generated with the non-parametric Mann-Whitney-Wilcoxon Rank Sum test, which does not assume normality. Reported p -values compare the distance traveled by the top organism for each of the 48 runs at the final (1000th) generation. Plots report distance traveled, not adjusted fitness (which penalized for the number of conducting voxels as explained previously).

4.3 Results

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

As with many explorations in evolved virtual organisms, one of the main




Figure 4.4: An action potential wave propagating across a mostly homogeneous surface. (*left, single robot*): The robot has a large patch of continuous conductive muscle on its back. In this pre-simulation state, cell colors signify the following: orange cells are conductive, blue cells are non-conductive; dark colors (blue or orange) signify muscle cells, while lighter colors (blue or orange) signify non-actuable cell tissue; the red cell at the bottom is the robot's pacemaker cell. (*right, 3×4 grid of robots*): A progression over time (left to right, top to bottom) shows the wave-like propagation of the action potential phases (color meanings are described in Fig. 4.3). Note how the action potential emerges from the center, stimulated by the wave propagating out through the conductive tissue from the pacemaker below it. Following the light blue depolarization, the yellow and red phases show the longer lag of the refractory period, following in exactly the same pattern made by the leading edge of the action potential wave. As the wave fully passes, the cells return to their dark resting state and are thus able to spike again with a new action potential when the next wave comes.




Figure 4.5: A more complex electrophysiological robot. (*left robot*): Contrary to Fig. 4.4, this creature shows complex patterning of the orange conductive tissue within the insulating blue tissue. (*right three robots*): As they unfold over time (left to right), the action potential waves in this robot produce a highly fractured, counter-intuitive actuation pattern that involves electrical signals flowing through long, sparse connective corridors and around corners (an explanation of the colors is provided in Figs. 4.3 and 4.4). The result is an interesting and unexpected behavioral pattern wherein the creature mashes and spins the left side of its body, which is separated from the larger, right side of its body by a large, oddly shaped internal cavity. Despite this bizarre behavior, it effectively locomotes. This behavior and others can be viewed on Youtube at: <http://goo.gl/CvJp41>.

goals is complex, natural-appearing behavior. However, there are no satisfactory metrics for the “naturalness” or complexity of evolved behaviors. For this reason, we must rely on our qualitative, subjective assessments. A video of the evolved behaviors can be seen on the “Cornell Creative Machines Lab” Youtube channel, or found directly at this link: <http://goo.gl/CvJp41>. We believe the behaviors are interesting, complex, and lifelike – at least as much as in [43] – despite the added challenges of evolving physically embedded control.

We observed that physically instantiated control circuitry can produce both predictable and chaotic behaviors. Fig. 4.4 shows a simple wave of action potentials propagating outwards from the center of the creature, with little interruption. Fig. 4.5 reveals the evolution of unpredictable physical dynamics that still

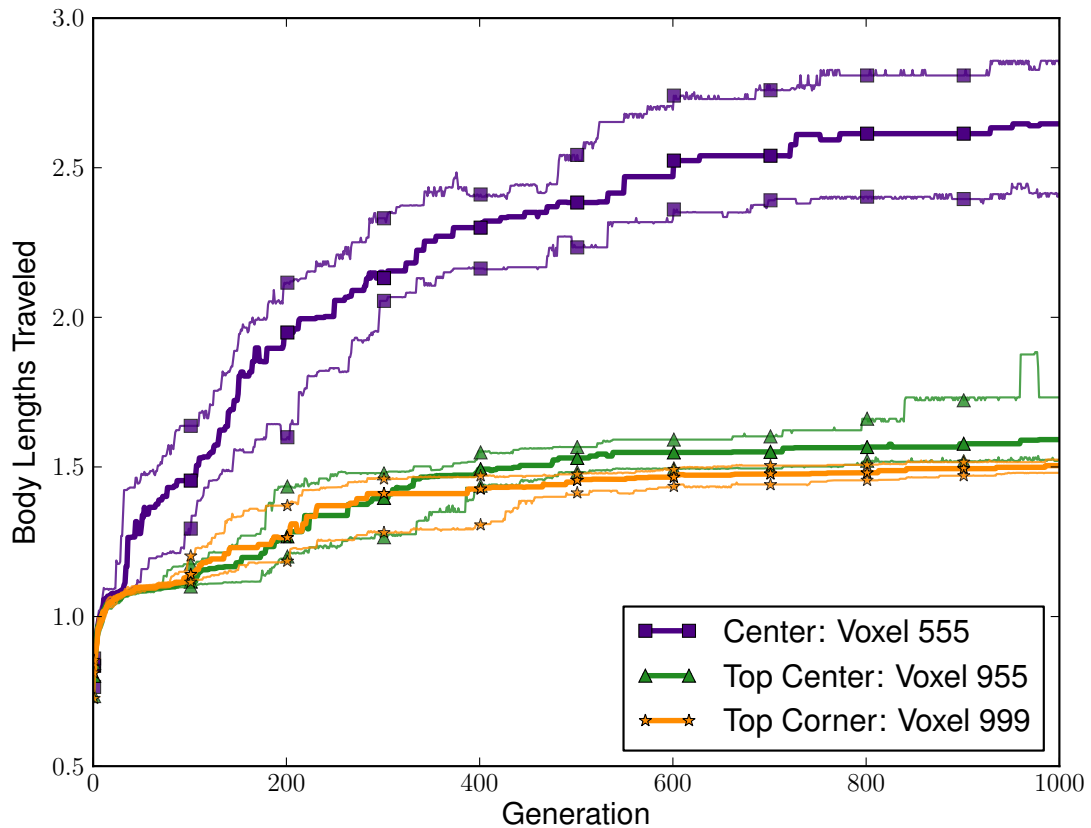


Figure 4.6: The effect of the placement of the central pattern generator (CPG) on the evolved speed. In one treatment, the CPG is placed at the top corner of the $10 \times 10 \times 10$ design space (voxel 999). This treatment performs slightly, but significantly ($p = 3.43 \times 10^{-4}$), better than another treatment that places the CPG at the center of the top plane of the bounding box (voxel 955). Outperforming both of these ($p < 4.91 \times 10^{-11}$) is the baseline treatment in which the CPG is always placed as close to the center of the bounding box as possible.

produce functional behavior. Notice the multiple “inputs” to a potential self-sustaining circular pathway. Fig. 4.1 demonstrates a circular actuation pattern of intermediate complexity, due more so to changes in the robot’s shape than to material differences within it. We now turn to more quantitative analyses.

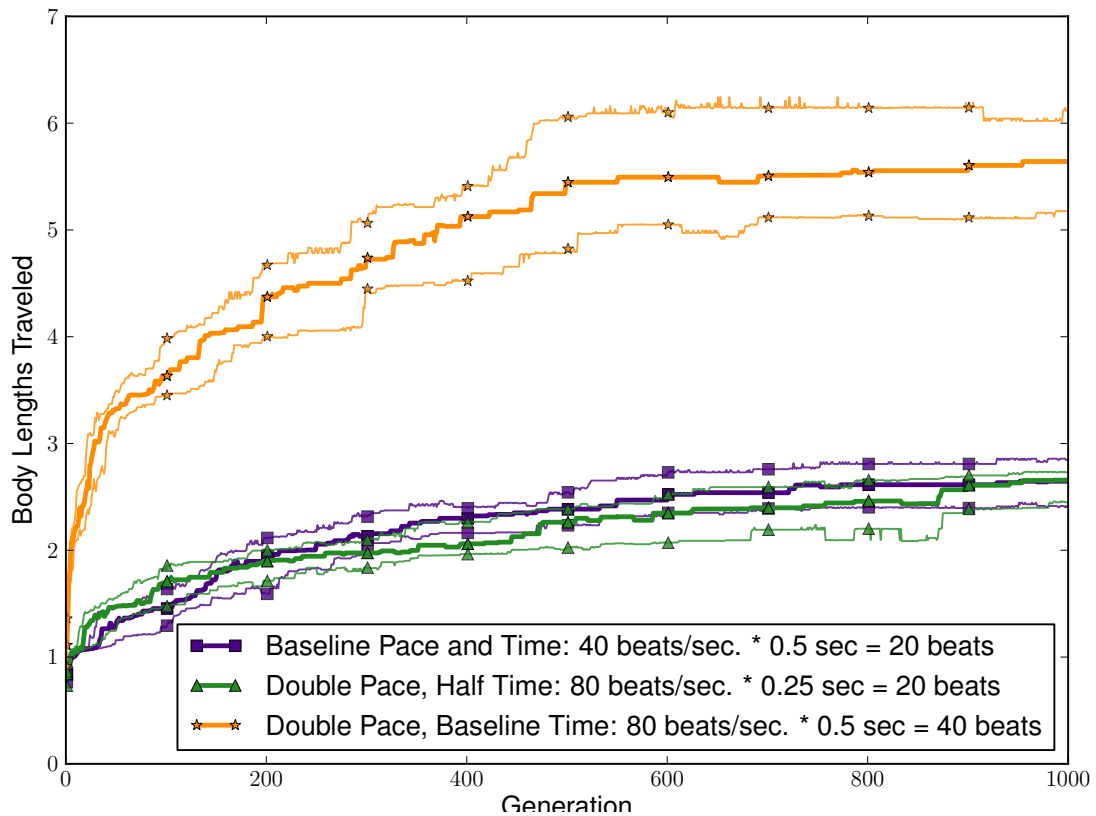


Figure 4.7: The effect of faster pacemakers (CPGs). It is not surprising that a faster CPG (80 beats per second) travels farther when evaluated for longer, or when compared to a slower CPG ($p < 10^{-16}$). However, when the comparison is made according to distance per beat (half time/full speed against half speed/full time – both producing a total of 20 beats) there is no difference in their performance at Generation 1000 (p -value = 0.51), suggesting that CPG speed does not greatly affect evolved locomotion speed.

4.3.1 Pacemaker Placement

The placement for the pacemaker was an arbitrary decision made during the design of this new system. In an effort to mimic the midline location of the central nervous system in biology, the pacemaker was placed in the middle of the design space from which the creature was built. Thus action potential waves

could propagate out equally in all directions and were not biased in any particular direction of travel. In order to test the effect of this arbitrary choice, a treatment was also performed where the pacemaker was located at the center voxel of the roof of the $10 \times 10 \times 10$ design space – voxel number 955 (where indices increase from the bottom, left hand, nearest corner), as well as a treatment that placed the pacemaker in the top right corner – voxel 999.

As shown in Fig 4.6, the placement of this pacemaker significantly affects performance. While a central location (baseline treatment) shows significant advantages compared to the top-center and top-corner pacemaker locations ($p = 4.91 \times 10^{-11}$ and 7.16×10^{-16} , respectively), a statistically significant difference is also demonstrated between the two less-different treatments: the top-center location outperformed the top-corner location ($p = 3.43 \times 10^{-4}$). These results show that the pacemaker location can have a clear effect on the evolved behaviors. Future work shall place the exact location under evolutionary optimization.

4.3.2 Speed of Pacemaker

Another implementation decision was the low-frequency pacemaker to allow for static gaits. The increased stability and robustness of static gaits is appealing, and this may allow better transferability to physical robots [18]. However, animals often employ dynamic gaits when there is an incentive for speed (as there is here). The tradeoff between these two is not known in this system. To examine this tradeoff, we compared three different treatments. First, the baseline treatment includes a pacemaker with the relatively slow pace of 40 beats

per second (BPS). Since the baseline evaluation period is half a second, this results in 20 electrical pulses from the pacemaker per trial. A second treatment explores the increased potential for dynamic gaits at the maximum pacemaker speed of 80 BPS (the limit is due to the fixed length of the refractory period). In this faster treatment, each individual cell contracts at the same rate as before, but the pacemaker is now exciting cells as soon as their refractory period ends, instead of waiting (the length of an additional actuation cycle) before sending another pulse into the system. This system uses twice the amount of energy, producing 40 action potential waves in the same half second. In a third treatment, the faster paced (80 BPS) pacemaker is evaluated for half its normal time length, resulting in 20 beats per evaluation. This treatment allows a fair comparison of pacemakers in terms of distance traveled per “beat”, rather than per unit time.

Unsurprisingly, the faster pacemaker evaluated for the full half second outperforms both the slower pacemaker evaluated for the same time period and the faster pacemaker evaluated for the shorter evaluation time ($p < 10^{-16}$ for both, Fig. 4.7). Interestingly, the frequency of the pacemaker has no significant effect on the distance traveled ($p = 0.51$ at generation 1000), suggesting that any disparity between the faster and slower gaits was not realized in simulation (with the number of beats held constant). Testing this result in the transfer to physical robots is a subject for future work.

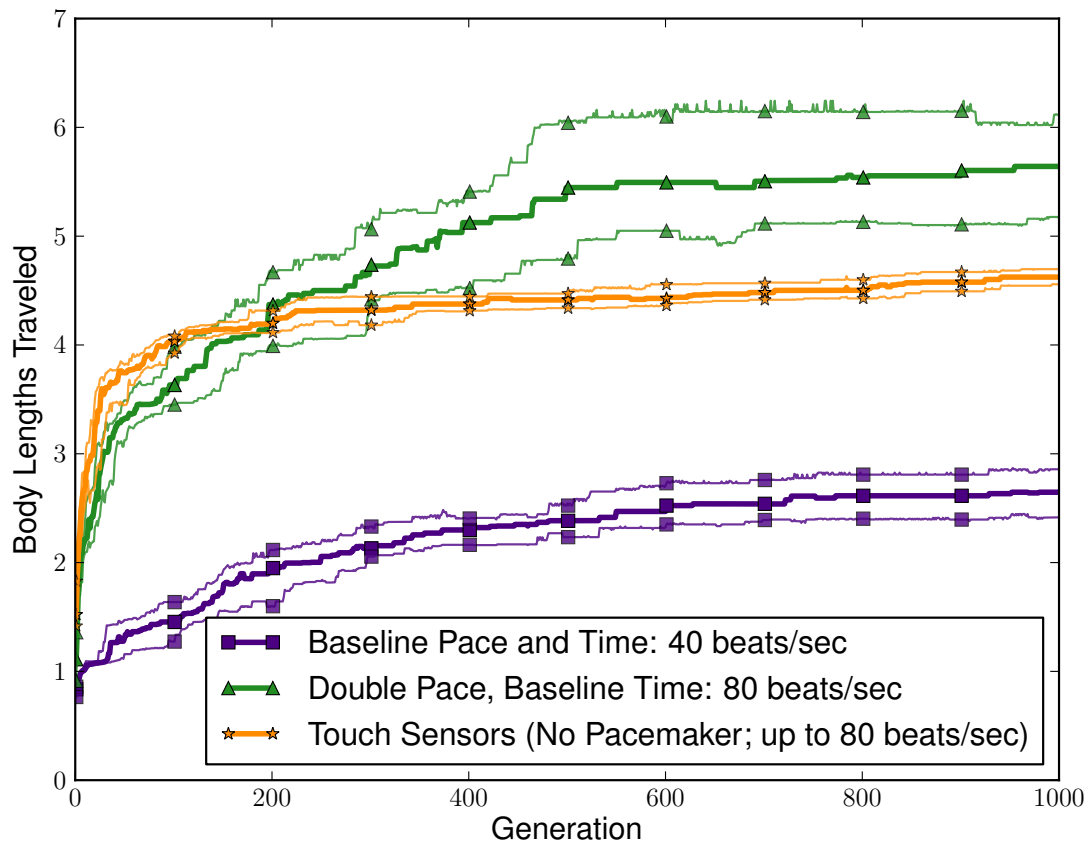


Figure 4.8: The performance of touch sensors vs. central pattern generators. The touch sensor treatment produces an expected number of beats with the upper bound set by the faster (80 beat/sec) CPG. Despite early evolvability leading to a statistically significant advantage in the first 150 generations, in later generations the touch sensor setup is unable to produce creatures that travel as far as the faster CPG setup ($p = 1.27 \times 10^{-7}$ at Gen. 1000). Artificially throttled, the slow CPG is unable to compete with either ($p < 10^{-16}$).

4.3.3 Touch Sensors

Another implementation decision was the use of pacemakers as the primary drivers of the system. While pacemakers, also known as central pattern generators, are biologically motivated [[142]], we could instead ask evolution to generate its own cadence. To provide an alternative to the pacemaker, we tested

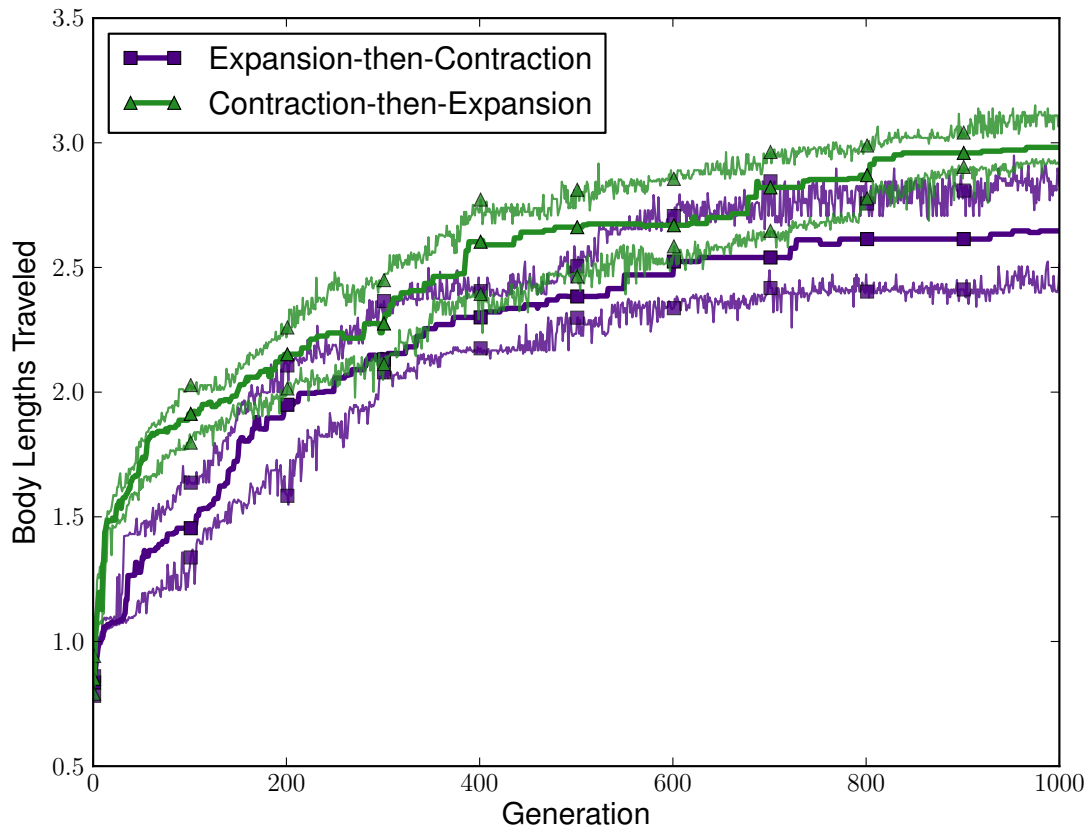


Figure 4.9: Unlike the regularly occurring actuation cycles of [43], the electrophysiological actuations in this paper do not have a necessary order: either expansion or contraction can occur first. It turns out that performance is significantly higher when muscles contract first and then expand, rather than vice versa ($p = 1.94 \times 10^{-3}$).

a treatment with touch sensors in lieu of a steady internal signal.

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

the ground and fully recovered from their prior depolarization.

Thus, the upper bound on the number of action potentials that the touch sensors could produce is that of an 80 BPS pacemaker (the 80 BPS pacemaker fires again as soon as exiting the refractory period, where the touch sensors do so only if also touching the ground at that time – the slower 40 BPS pacemaker waits the length of one cycle before firing again). To reach this upper bound, touch sensors would have to be touching the ground exactly at the time when they completed their refractory period, and thus it is likely that this ceiling would not be reached in all cases. For a comparison, Fig. 4.8 shows the median distance traveled over evolutionary time plotted against that of the slower pacemaker (40 BPS) and the faster pacemaker (80 BPS) described above, and evaluated for the baseline half-second evaluation time. It is not surprising that the slower pacemaker falls behind the pack here, as it is handicapped by a throttle on its only source of action potentials compared to the faster pacemaker and the touch sensors ($p < 10^{-16}$). The tighter race is between the touch sensor and the faster pacemaker. In the early stages (< 150 generations), the robots with touch sensors significantly outperform robots with a pacemaker. However, in the later stages of evolutionary optimization, the touch sensor treatment shows modest gains compared to the continued innovation of the pacemaker treatment, with the pacemaker treatment significantly outperforming it at the end of the run ($p = 1.27 \times 10^{-7}$). The relatively low level of improvement in the touch sensor treatment in the later stages of evolution may suggest the premature convergence on local optima. The multiple points of origin for action potential waves, and thus wave collisions, may have also had an effect. An additional issue that could have hindered performance in this treatment is the upward propagation of signals from touch sensors on the ground, versus outwards expanding waves

from the center of the organism.

4.3.4 Expansion/Contraction Cycle

In the soft robot evolution system described by [43], regular, quickly repeating, and coupled out-of-phase sinusoidal action cycles defined the expansion and contraction of cells. In this model, which does not feature the same complementary muscle types, the question of actuation cycle is not entirely clear. In an attempt to explore this, here we test the effectiveness of contraction-then-expansion phase cycles against expansion-then-contraction cycles (Fig. 4.9). These treatments take place on the baseline (slow) pacemaker setup, as to not allow continuous and quickly repeating expansion/contraction cycles, but rather to have a break between actuations. Despite the same number of beats (and thus the same amount of overall expansion and contraction) in both setups, the contraction-then-expansion setup performs significantly better ($p = 1.94 \times 10^{-3}$). While the reason for this difference is not entirely clear, it may be due, in part, to a larger continuous expansion period from the trough of the sine wave to its peak (continuous expansion from minimum to maximum size) in the contraction-then-expansion treatment. In contrast, the expansion-then-contraction setup includes a full-cycle length pause in the middle of its expansionary period. This explanation would suggest that more locomotion tends to come from pushing than pulling, which is in line with our observations from viewing videos of the evolved behaviors.

4.4 Discussion

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

4.5 Conclusion

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

CHAPTER 5
ON THE DIFFICULTY OF CO-OPTIMIZING MORPHOLOGY
AND CONTROL IN EVOLVED VIRTUAL CREATURES

Abstract of Chapter ¹

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.

5.1 Introduction

In 1994, Karl Sims’ seminal work on “Evolving Virtual Creatures” [263] created a field of study by that name. This work featured simulated creatures that

¹Appeared as: Cheney, N., Bongard, J., SunSpiral, V., & Lipson, H. (2016). On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures. In the Proceedings of the International Conference on the Synthesis and Simulation of Living Systems (pp. 226-234).

were able to optimize both their physical layout and their behavioral control strategies for such tasks as terrestrial locomotion, swimming, phototaxis, and competition [264].

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 [185, 103, 181, 210, 83, 23].

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 [105] – despite significant increases in computing power. Many researchers have suggested hypotheses for the cause of this standstill, such as deficiencies in the search algorithms [135, 171, 203] or genetic encodings [137, 24]. 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 [10, 39]. 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 [149] 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 vir-

tual 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 [306]. It suggests that the dynamic interactions between a reactive agent and the environment, through sensory-motor feedback loops, are an important driver of behavior [27], 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 [3].

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 [219]. 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. Without 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.

5.2 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) [92]. 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 [181, 24, 38, 39] (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 [149], 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 [43].

In the remainder of this work we set out to reproduce the symptoms described in [149], 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.

5.3 Methods

Similarly to [149], we employ soft robots as our instantiation of evolved virtual creatures. We use 3D voxel-based soft robots, following from [43], but replace their discrete muscle types and synchronized contractions with voxels which allow individualized phase offsets, consistent with the controllers used in [149]. This allows for behaviors such as propagating waves, which were not possible in [43] (but were achieved in [149] and [38]). A global frequency of oscillations is also optimized.

5.3.1 Dual-Network CPPN

We genetically encode the soft robot phenotypes as a network, inspired by the CPPN-NEAT [271], the algorithm employed by both [43] and [149] (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 other².

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

²both source code and resulting data are available upon request

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 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 [219].

5.3.2 Physics Simulation in VoxCad

Consistent with [43], we employ the open-source soft-body simulator VoxCad [127] 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 [43].

5.3.3 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 [271], 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.

5.3.4 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 $p \leq 0.001$.

5.4 Results

First and foremost, we set out to replicate and examine the results found in [149], 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. 5.1 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 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 [149]. However, it does not




Figure 5.1: 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.




Figure 5.2: 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).

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. 5.2 and 5.3.




Figure 5.3: 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).

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

Fig. 5.2 shows the fitness impact of morphology freezes at various times during optimization, and its significance compared to co-optimization of morphology and control for the entire 5000 generations. We see that full optimization does not show a significant improvement in fitness 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. 5.1 does not contradict such a suggestion.

This does not mean that optimization as a whole is converged at this point. Fitness improvements from the control optimization occurring after the final gross morphology is fixed are noted in [149]. 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 value at the time of freezing (21.157).

Fig. 5.3 shows the impact of the converse treatment, in which the creature's controllers are frozen at a given point in time and only morphological variations are allowed thereafter. This treatment show that significant differences in resulting fitness values occur for at least 100 generations (at the 95% confidence level, as $p < 0.001$ for all freezing points ≤ 100), but not more than 250 generations ($p \geq 0.0565$ for all 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.

5.5 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 pronounced 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, while variations to the controller do not propagate downstream.

This feature of embodied cognition has the effect of creating larger (and arguably more unpredictable) behavioral changes to similar sized variations in genotypic space. 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 [154]. This is consistent with what we have experienced thus far with the optimization of morphology converging prior to control.

5.5.1 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 [290]. 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 5.4 visually shows the continued con-

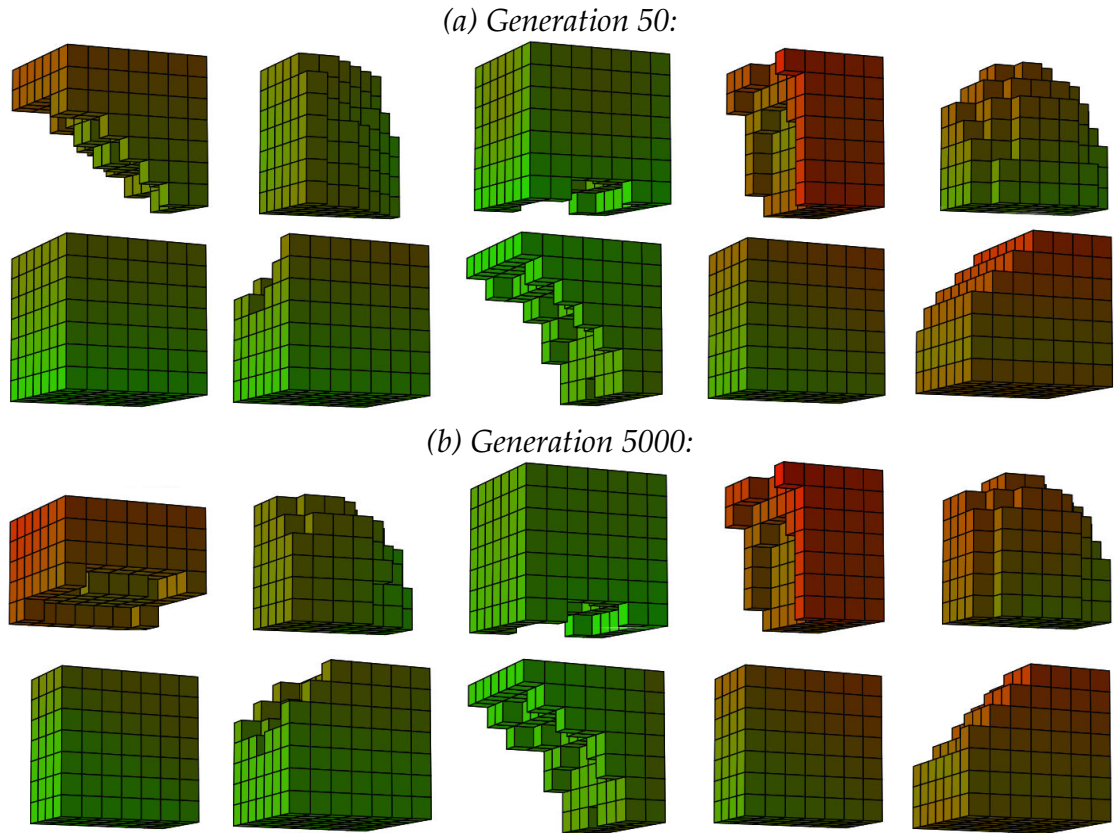


Figure 5.4: Stagnation shown in the top 10 morphologies under the distance/energy fitness treatment. Note the similarity in gross morphologies from gen 50 (top) to gen 5000 (bottom). The top performing creature shows the largest change between these points, with the new morphology arriving from a mutation at gen 53. Also note the variance and complexity in forms, compared to Fig. 5.1, suggesting the added morphological dependence of this fitness function.

vergence 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. 5.1.

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 $3^{343} = 4.5 * 10^{163}$ 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 [263, 171].

In attempting to reproduce the work from [149], 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; limited in practice only by machine precision) search space for control. 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 [43]. This allows just two offset control states (implemented by rounding the continuous phase offset values) to create a search space of size $2^{343} = 1.8 * 10^{103}$ (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 [43] should be considered to be entirely 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 [220] 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.

5.5.2 Future Work

The results shown in this work are specific only to this instance and experimental setup. Thus, many more instances of this approach (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 [263]; 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 both 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 morpho-

logical 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 conformation that it's possible and that the understanding gained from this current work may contribute to future improvements.

5.6 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 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 a 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.

CHAPTER 6
SCALABLE CO-OPTIMIZATION OF
MORPHOLOGY AND CONTROL IN EMBODIED MACHINES

Abstract of Chapter ¹

Evolution sculpts both the body plans and nervous systems of agents together over time. In contrast, in AI and robotics, a robot's body plan is usually designed by hand, and control policies are then optimized for that fixed design. The task of simultaneously co-optimizing the morphology and controller of an embodied robot has remained a challenge – as evidenced by the little improvement upon early techniques over the decades since their introduction. Embodied cognition posits that behavior arises from a close coupling between body plan and sensorimotor control, which suggests why co-optimizing these two subsystems is so difficult: most evolutionary changes to morphology tend to adversely impact sensorimotor control, leading to an overall decrease in behavioral performance. Here, we further examine this hypothesis and demonstrate a technique for “morphological innovation protection”, which reduces selection pressure on recently morphologically-changed individuals, thus enabling evolution some time to “readapt” to the new morphology with subsequent control policy mutations. This treatment tends to yield individuals that are significantly more fit than those that existed before the morphological change and increases evolvability. We also show the potential for this method to avoid local optima and show fitness increases further into optimization, as well as the potential for convergence to similar highly fit morphologies across widely varying ini-

¹To appear as: Cheney, N., Bongard, J., SunSpiral, V., & Lipson, H. (in review). Scalable Co-Optimization of Morphology and Control in Embodied Machines

tial conditions. While this technique is admittedly only the first of many steps that must be taken to achieve scalable optimization of embodied machines, we hope that theoretical insight into the cause of evolutionary stagnation in current methods will help to enable the automation of robot design and behavioral training.

6.1 Introduction and Background

Designing agile, autonomous machines has been a long-standing challenge in the field of robotics [223]. Animals, including humans, have served as examples of inspiration for many researchers, who meticulously and painstakingly attempt to reverse engineer the biological organisms that navigate even the most dynamic, rugged, and unpredictable environments with relative ease [45, 234, 311]. However, another competing approach is the use of evolutionary algorithms to search for robotic designs and behaviors without presupposing what those designs and behaviors may be. These methods often take inspiration from evolution itself, rather than the exact specifications of any given organism produced by it.

The use of an evolutionary algorithm for automated design comes with many benefits: It removes the costly endeavour of determining which traits of a given organism are specific to its biological niche, and which are useful design features that can provide the same beneficial functions, if instantiated in a machine. It can yield machines that do not resemble any animals currently found on earth [166], as it allows for machines that are specialized for behaviors and environments that differ from those of the model organism. Additionally, the

optimization process can serve as a controlled and repeatable test-bed for the study of evolutionary and developmental theory [173, 174, 22].

However the approach of bio-inspired optimization also presents challenges that have yet to be overcome, causing the scale and complexity of evolved robots and virtual agents to pale in comparison to their biological counterparts. It is unclear what aspects of (or omissions from) current evolutionary algorithms are preventing robotic optimization to scale the way that biological evolution has demonstrated.

The generalization of design automation to include both the optimization of robot neural controllers and body plans has proven to be problematic [105, 40]. While recent successes have demonstrated the potential of effective optimization for the control policies of agents with fixed morphologies [106, 66, 175] or – to a lesser extent – the optimization of morphologies (body plans) for agents with minimal and fixed control policies [58, 43, 10], the co-optimization of the two has seen very limited success [40].

The optimization of the shape and general morphology of a robot is of great importance to the goal of autonomous robotics, as biological animals appear to rely heavily on the adaptation of their body plans to effectively interact with the natural world. This is seen in the adaptation of specialized physical traits on an evolutionary time-scale (like the specialized beak shape and size of Darwin's Finches [69, 164]). It is also seen through experience-dependent growth over an animal's lifetime (such as Wolff's Law [309] or Davis' Law [71], the body's response to stress and loading on bone and soft tissues, respectively).

Other reasons to evolve robot body plans include the phenomena of embod-

ied cognition (the way in which the presence and organization of perceptual and motor systems impact the frames and processing of higher level cognitive functions [306, 3, 267]) or morphological computation (where certain higher level cognitive functions, like preprocessing, are subsumed and solved by the organization of the body [222, 220, 119]). Creating body plans by hand that exploit these properties is extremely difficult to do because they are non-intuitive. Thus, enabling evolutionary methods to discover such morphologies is desirable.

The inability to perform robot optimization at scale, and specifically that traditional evolutionary robotics methods tend to become trapped in local optima, has been experienced and noted informally by many researchers involved with robot optimization, yet published rarely. Thus the lack of publication is presumably because the field lacks incentives for the publication of negative results [79, 92], rather than a lack of negative results in unpublished works. Joachimczak et al. [148] provide an anecdotal example of premature convergence in the co-optimization of robot brain and body plan (Fig. 19 of that paper). Cheney et al. [40] further analyze the phenomenon of premature convergence in embodied machines and suggest that traditional evolutionary algorithms are hindered in this setting primarily in their ability to perform continued optimization on the morphology of the robot. They hypothesized that the premature convergence may be due to an effect of embodied cognition, in which an individual's body plan and brain have an incentive to specialize their behaviors to complement one another. This specialization makes improvements to either subsystem difficult without complementary changes in the other (a highly unlikely event given current algorithms) and thus results in an embodied agent which is fragile with respect to design perturbations.

6.1.1 On the Dichotomy Between Brain and Body Plan

It should be noted that the dichotomy between “brain” and “body plan” is certainly a false one, as computation is known to occur at various levels throughout an entire organism [15]. Indeed, our rationale for the algorithm proposed below will rely heavily on the notion that an agent’s morphology and control are not independent, and the explicit separation of “brain” and “body plan” merely draws attention to the difficulties that occur when we frame optimization with this dichotomy.

We want to make clear from the start, that this algorithm is only an improvement on the current (and clearly incomplete) methods for the evolution of embodied machines – and is not intended to be the end-all solution. The end goal, from our current prospective, involves complex mechanisms for interwoven morphogenesis and neurogenesis, which adaptively and continually mold a complete agent from a combination of genetically-encoded patterns and experience-dependent mechanisms, much the way that biological organisms are formed. In these biological agents the distinction between “brain” and “body plan” is blurred, and organisms are genetically encoded as a single “embodied agents”. But until we have the understanding and engineering ability to optimize robots with complex embodiment and controllers in such a way, let the algorithm proposed here serve as a short term solution to just some of the many unsolved problems in the field of evolutionary robotics [23].

6.1.2 Evolving Robot Body Plans and Nervous Systems

Attempts to solve this problem of fully-automated robot evolution are frequently traced back to the work of Sims [263]. This work introduced the use of evolutionary algorithms to produce goal-directed behaviors and morphologies simultaneously. Despite the advance this work represented, the evolved robots tended to be composed of a relatively small number of components (actual values not published, but figures show a mean of 6.042 objects per robot, and each segment typically controlled by a few neurons). It has been suggested that in the decades since this time increases in computational power following Moore’s Law [248] should have vastly increased the scale and complexity of robots evolved using Sims’ and similar methods [43]. Yet vast increases in scale have not been the case empirically, as evolved robots reported in [105, 23] fail to exhibit any significant increase in size or complexity.

A wide range of hypotheses for the lack of scalability have been proposed. Some focused on a lack of efficient evolutionary search algorithms [135, 171, 203] or genetic encodings [138, 24, 43], while others pointed to a lack of incentives for complexity in the simple tasks and environments of previous work [10, 39]. Yet attempts to evolve robots using methods designed to overcome these challenges have yet to obviously surpass Sims’ work in terms of complexity and scale. This work investigates a different hypotheses, first suggested in [40], that considers the way in which an agent’s brain and body plan interact during the optimization process.

This work is related to co-evolutionary methods (such as [162, 264, 52, 80, 156, 212]) in that there are multiple components being evolved. These classical co-evolutionary systems (like predator-prey) typically call for two separate

populations being evolved as generalists against one another. However, for the case of brain-body co-optimization in evolutionary robotics, it is often desirable to have the body and brain of a robot specialized to one another, and developed simultaneously from a single genome, and in a single population (as in [263, 181, 138, 9, 24] – and as is the case in natural systems). This combined brain-body genome causes particular optimization issues which we hypothesize below.

6.1.3 The Interdependency of Body and Brain

Specifically, we investigate the notion that the specialization of brain and body plan to one another during evolution creates a fragile co-dependent system that is not easily amenable to change. The specialization of each subcomponent to the other creates local optima in the search space and premature convergence to suboptimal designs. In this paper, we explore a direct solution to the problem of fragile coupled systems: explicitly readapting one subsystem (e.g. the body plan or the brain) after each evolutionary perturbation to the other. The proposed method differs from a traditional evolutionary algorithm, which evaluates the fitness of a newly proposed variation immediately (i.e. with no readaptation), and uses only this valuation of fitness to determine the long term potential of that variation.

Consider a hypothetical, partly-evolved robot with a partially optimized body plan (for example a quadrupedal form) and a partially optimized controller (for example the legs swing forward and back through the sagittal plane). Suppose that this controller has co-adapted to the morphology during evolution

such that for each step forward, the robot contracts a muscle near its hip with enough force to swing one of its anterior legs forward with just enough force to land in front of the body and successfully take a step. Consider further a variation of this morphology proposed by an evolutionary algorithm in which the new robot possesses longer legs but the controller has not changed.

If this machine were being evolved for rapid locomotion, having longer legs and being able to take longer strides would be beneficial and this variation should be more successful than the original design. However, during evaluation of this new robot, the original controller applies the same amount of force to the now longer leg, failing to move it, and thus frustrating the robot's ability to walk in a coordinated fashion. Current evolutionary methods would treat this robot as the recipient of a detrimental mutation and remove it from the population.

In this example, a mutation to one subsystem (the body plan) could be beneficial to the robot's descendants, but since the immediate impact of the change is detrimental, the mutated robot suffers a decrease in fitness and does not survive to produce offspring. If all such variations are considered to be detrimental and rejected (regardless of their long term potential), then the evolutionary algorithm has prematurely converged to a local optima in the search space because there appear to be no better alternatives in the immediate neighborhood of the current design.

However, the newly proposed morphology would have resulted in a robot which outperformed its predecessor, if coupled with a controller suited to that morphology. We can determine that the newly proposed morphology is superior by suppressing mutations to that body plan and allowing readaptation of

its controller to properly coordinate behavior.

Herein lies the foundation for our proposed algorithm: readapt each controller until the new proposed morphology is more fit than its predecessor.

6.2 Methods

6.2.1 Controller Readaptation

The most obvious method for modeling controller readaptation would be to protect any lineage that has recently experienced a mutation to the body plan by allowing it to undergo several generations of evolutionary change restricted to the control subsystem. If any member of the lineage achieves higher fitness than the pre-mutation ancestor during that time period, the descendant is retained. Otherwise, the new morphological variant dies out.

However, it is unclear how to set the time period for this protection *a priori*. Surely the amount of time a controller takes to readapt to a new morphology depends on many specific features of the complexity, genetic encoding, desired behavior, and current ability level of the robot (which changes over optimization time). Determining the correct value of this parameter would require a full parameter sweep over various values of readaptation time for each new combination of brain, body, and environment. If our goal is simply to optimize a robot, then searching for this value in each unique optimization scenario is intractable.

6.2.2 Proposed Method: Morphological Innovation Protection

In response to the unintuitive nature of the optimal value for readaptation length, our proposed approach is free of this parameter. Descendants of robots that experience morphological mutations are retained in the population and the number of generations that have elapsed since that mutation occurred are tracked (referred to as the “age” of the morphology). If two individuals are found in the population such that the latter robot exhibits better performance on the desired task and has experienced fewer generations since a morphological mutation than the former robot, the former robot is removed from the population. In effect, the latter robot has exhibited an ability to not only recover from its ancestor’s morphological mutation but improve upon it (and others in the population). The concept of tracking age using it as an optimization objective are borrowed from [135, 249]. The major difference here being that the age refers to the length of time that a subsystem of the agent (e.g. the morphology) has remained unvaried, rather than (the original definition of) the total time since a random individual was introduced to the population.

This procedure has the effect of “protecting” new morphologies with poorly adapted controllers, and thus we will henceforth refer to this procedure as “morphological innovation protection.” This protection is a form a diversity maintenance, though reduced selection pressure for newly mutated morphologies. Various other methods exist for encouraging diversity (e.g. fitness sharing, crowding, random restart parallel hillclimbers [107], novelty [171], speciation [275]), however age was chosen for its simplicity of implementation, it’s parallels to multi-timescale learning in biology, and because it helps to avoid the cost of extended control re-optimization for non-promising morphologies

– since the the age-pareto optimization allows comparisons between all new “child” morphologies that have had equal readaptation time, even if they are not yet fully readapted (rather than making comparisons only after a set amount of readaptation).

6.2.3 Evolutionary Algorithm

All optimization is performed by a population-based Evolutionary Algorithm, inspired by the popular algorithm NEAT (NeuroEvolution of Augmenting Topologies) [275] (importantly without speciation). All trials follow a (μ, λ) -Evolutionary Strategy [20] with $\mu = 25$ parents and $\lambda = 25$ mutants for a population size of 50. Trials last for 5000 generations. Crossover was not considered in this work. Mutation had a 50% chance of creating a variation to either the morphology or the controller of a given robot, but not both. Other ratios of morphology:controller mutations were considered (1:99, 20:80, 50:50, 80:20, and 99:1), but none showed a significant effect on resolving the premature convergence and resulting fitness in preliminary trials without innovation protection.

6.2.4 Genetic Encoding for Soft Robot Morphologies

The soft robot morphologies are encoded with a Compositional Pattern Producing Network (CPPN) [271], consistent with prior work on soft robot evolution in [40] and [43]. The CPPN encoding produces the cell fate of each voxel in the robot through a type of neural network that takes each cell’s geometric location (x, y, z Cartesian coordinates and r radial polar coordinate) and outputs a vari-

ety of “morphogens” (in this work, there is one to determine whether a cell is present in that location and one to determine whether a present voxel should be a muscle or a passive tissue cell). Since nearby voxels tend to have similar coordinate inputs, they also tend to produce similar outputs from the network – creating continuous muscle patches. However, CPPNs produce particularly interesting geometric patterns, as the activation functions at each node can take on a variety of functions (here: *sigmoid*, *sine*, *absolute value*, *negative absolute value*, *square*, *negative square*, *square root*, and *negative square root*). These functions tend to produce regular patterns and features across the coordinate inputs (for example: an *absolute value* node with an x input would produce left-right symmetry, or a *sine* node with a y input would produce front-to-back repetition).

This network is optimized to produce high performing morphologies by iterating through various proposed perturbations to it. These include the addition or removal of a node, or edge to the network, as well as the mutation of the weight of any edge or the activation function at each node.

As the CPPN is an indirect and generative encoding, the genotype (neural network) is independent of the scale of the phenotype (soft robot being built), since the network is a continuous function that is discretized only to be queried once per voxel. This means that the optimized network is scale free, and can be applied to build robots with any number of voxels (i.e. any resolution or overall size) – which is limited only by available physical resources (for real world robots) or computation resources (in the case of simulated robots here).

6.2.5 Types of Controllers and their Genetic Encoding

In order to control the robots (determine when to contract each muscle cell – and with what speed and force), two different controller strategies were considered for different experimental treatments: a simpler open-loop oscillatory controller, and a more complex closed-loop neural network controller. For each of these control strategies, a unique controller was optimized for each muscle cell in the robot’s morphology. However, since each parameter in the controller is also optimized by a CPPN (a separate CPPN from the one which produces the robot morphology) – which is a scale free encoding, the number of unique controllers being optimized is not substantially more costly than optimizing a single global controller. All controllers output a value between -1 and 1 at each time step, which corresponds to a linear change in each dimension of a muscle cell ($\pm 14\%$ of its original length, or $\pm 48\%$ of its original volume). Passive tissue cells remain at their original size (though they also deform based on their intrinsic compliance).

The encoding of the morphology and the controller of the robot into two separate CPPN networks emphasizes the false dichotomy of robot brains and body plans. However, this explicit separation allows us to make changes specific and isolated to either the brain or the body. This is necessary for the proposed algorithm, as controller readaptation requires iterating through controller changes without affecting the morphology.

Open-Loop Controllers

Open-loop controllers take as input a global sinusoidal oscillator, but no sensory information from the environment. In this scenario, two parameters are optimized. One is the phase offset between each individual cell's muscle oscillations and that of the global clock (which acts as a central pattern generator, or CPG) – the second is the frequency of this global clock (since CPPNs don't currently enable global parameters, this is done by averaging the local values for each individual cell to produce a single global value). While this encoding is simple and straightforward, it has the ability to produce complex behaviors, such as multiple patches of muscle groups that are in sync, counter-sync, or any real valued phase offset from each other. It also has the ability to produce gradually varying sweeps of phase offset, resulting in propagating waves of excitation across large muscle groups. Furthermore, the optimization of the global frequency is able to produce oscillation speeds which are fine tuned to the properties of individual morphologies (such as optimizing to maximize the resonance of soft tissues in appendages).

Closed-Loop Controllers

Closed-loop neural network controllers also use this global clock as input (now with a fixed frequency of 10Hz), but additionally include sensory information in the form of touch sensors in all cells (which output 1 when in contact with the ground, and 0 otherwise), two localized proprioceptive sensor modalities – one which returns that cell's muscle output at the previous time step, and one that returns the average muscle contraction of neighboring cells (those in direct contact with the cell in question) – as well as a bias node. These five inputs are

fed to a two-node hidden layer with recurrent connections, before being passed to the single output node (that dictates current muscle contraction levels for that cell). In all, this network contains 16 edges, each of which have one parameter to determine the presence/absence of that connection [299] and one to determine the weight of that synapse (discretized to either -1 or 1 for their excitatory or inhibitory connections) – resulting in 32 binary free parameters (or 32 output nodes for the controller CPPN) in this setting.

6.2.6 Physics Simulation for Evaluation

Once the morphology and controller for a given robot are specified, the fitness (locomotion speed) of that robot is determined by constructing and simulating that robot in the VoxCad soft-body physics simulator [127]. Simulations last for 20 actuation cycles (which may be a variable amount of time, depending on the length of the globally optimized frequency – though this method of normalizing for the number of “steps” taken leads to a more fair comparison than normalizing by the amount of time per simulation). In the case of closed-loop control the number of oscillations refer to cycles of the CPG input node, which may or may not correspond to cycling of the actual muscle contractions. Following the evaluation, the total displacement for the center of mass of the robot is returned to the evolutionary algorithm as the individual’s fitness.

6.2.7 Soft Robot Resolution

As the CPPN genetic encoding is a continuous function, it may be discretized into a phenotype at any resolution (i.e. number of voxels), and in practice this resolution is only limited by computational resources (as more elements are more computationally expensive to simulate). In the default lower-resolution treatment, this discretization occurs over a $5 \times 5 \times 5$ space. The higher-resolution robots use phenotypes created at a $10 \times 10 \times 10$ resolution. Note that the distance values are noted in absolute number of voxels, and voxels are held at a constant size regardless of resolution – one centimeter in our simulations, so higher fitness values tend to be produced by phenotypes of higher resolution, as more muscle mass is available to those individuals.

6.2.8 Morphological Innovation Protection

In our newly proposed method we set the “morphological age” to zero for each new “child” that was the result of a morphological mutation to a current individual in the population. This means that for an individual to have a large value in their age objective, that individual must have been the result of a large number of successive controller mutations. This setup thus allows a simple comparison method for individuals who have had similar amounts of controller (re)adaptation to their current morphologies – as a dominated individual would have to have been out-competed in fitness ability by a morphology that is paired with a controller that is less-well adapted to it. This diversity maintenance mechanism encourages the exploration of new peaks in the rugged landscape of brain-body plans – with the implicit assumption that unique morphologies

correspond to peaks in this landscape.

The method of age resets corresponding to morphological mutations to existing individuals differs from the prior technique (e.g. [135, 249]) of inserting completely random individuals, as it allows the improvement in the fitness levels of age-zero individuals over time. In the case of traditional age-fitness optimization, age-zero individuals are drawn from the same distribution of fitness values for random genotypes regardless of when they are created. However, in the case of morphological innovation protection, age-zero individuals are not random and inherit many of the properties of their parents – meaning that they show higher fitness values over time. To demonstrate this empirically, we performed linear regression on the age-zero individuals from Fig. 6.2 (containing morphological innovation protection), which showed a significant ($p < 0.001$) increase in fitness over time (from 18.715 at generation 0 to 22.916 at generation 5000; $r^2 = 0.211$). This confirms a major difference between our proposed technique and the standard approach of age-based diversity maintenance through the introduction of random individuals.

6.2.9 Controller Innovation Protection

The procedure for controller innovation is similar to that of morphological innovation protection in all regards, except that the “protection age” is only reset to zero when a child results from a mutation to the controller (rather than a mutation to the morphology). This means that individuals with large values for the protection age objective are those who have had the same controller for a large amount of time and have thus made many attempts at morphological

(re)adaptations to that controller.

6.2.10 Morphological Change Threshold

In treatments with the minimum morphological change parameter, an individual's "morphological age" is only set to zero if a morphological change occurs that affects at least the proportion of voxels dictated by the threshold. This proportion is relative to the total possible number of voxels (1000 or 125 in the experiments above), rather than the number of voxels present in any particular robot (to help minimize the number of small and non-functional morphological changes for robots composed of few voxels).

6.2.11 Statistical Analysis

All treatments were performed for 30 independent trials, with random seeds consistent between treatments. All plots show mean values averaged across the most fit individual of 30 trials for each condition with shaded areas representing 95% bootstrapped confidence intervals of this average, and all p-values are generated by a Wilcoxon Rank-Sum Test [304].

6.3 Results

6.3.1 The Effect of Morphological Innovation Protection on Fitness

For the task of locomotion ability (a standard task in evolutionary robotics [23]), we optimize robots for the maximum distance they travel with 20 oscillations of their muscles (this distance is measured in voxel lengths – corresponding to one centimeter in our simulation). We first optimize the robots using the traditional method of “greedy” fitness evaluations for our selection criteria (where immediate locomotion ability determines survival in the population of candidate morphologies). In this setup, the traditional method produces robots with an average fitness of 21.717 voxels (with 95% bootstrapped confidence (CI) interval of 19.457 to 22.426 voxels).

Additionally, we optimized robots in the same task and environment setup, but this time using “morphological innovation protection” for our selection method – in which individuals can only be out-competed by those with equal or lesser amounts of controller (re)adaptation to their current morphologies. The treatment with morphological innovation protection produces significantly more effective robots ($p = 6.067 * 10^{-6}$), with a mean distance travelled of 31.953 voxels (and 95% CI of 28.157 to 36.511 voxels).

The increase in fitness shows that morphological innovation protection is a more effective way of optimizing robots, yet it does not conclusively demonstrate that the intuition of [40] is correct, as that work demonstrated the asymmetric difficulty in optimizing the morphology of a robot (as compared to op-

timizing its controller) and drew the hypotheses that this was due to the fact that the morphology encapsulated the controller (acting as a translator between the “cognitive” functions and the outside environment). While the above experiment does help to support the intuition that the controller must readapt to a new “morphological language” [40], it also introduces confounding features, such as the added population diversity afforded by “protection” and the added dimensionality of the search space from its this protection age – moving search from a single-objective to multi-objective optimization problem.

To tease apart the influence of these two confounds from the hypothesized mechanism, we present a treatment where the controllers of the robot undergo an equivalent protection to which the morphologies did in the above experiment. In this treatment, individuals can only be out competed by others whose morphologies have had equal or lesser amounts of readaptation to their newly mutated controllers – deemed “controller innovation protection”. This condition provides the potential advantages of multi-dimensional search and the added diversity (from reduced selection pressure [11]) of a protection mechanism. Yet it does not rely on the idea of a broken “morphological language” proposed by [40], which suggests the key role of the body as the interface between the brain and the environment (and thus the need to explicitly protect it). Under the condition of “controller innovation protection”, the robots are able to locomote 22.049 voxels on average (95% CI of 20.726 to 22.641 voxels), which fails to show a significant improvement over the single-objective case of no protection ($p = 0.240$), and performs significantly worse ($p = 1.211 * 10^{-4}$) than the protection of morphological innovations.




Figure 6.1: The fitness impact (distance travelled, in voxels) over optimization time (in generations) for various types of brain/body plan protection mechanisms. Values plotted represent the mean value of 30 independent trials, with 95% bootstrapped confidence intervals denoted by colorized regions. Note that at the end of optimization time (5000 generations), the “morphological innovation protection” (readapting controllers to new morphologies before evaluating their long term potential) significantly outperform all other treatments ($p < 0.001$), while the “controller innovation protection” does not perform significantly better ($p = 0.240$) than the case with no protection. This suggests that the added efficacy of protecting morphological innovations goes beyond the effects of an added dimension to the search space, and further suggests an fundamental asymmetry between the morphology and controller while optimizing an embodied robot.

Since morphological innovation protection and controller innovation protection both have the added dimension of “protection age”, the fact that they differ significantly from each other (and that controller innovation protection does not show significant improvements over the baseline case of no innovation protection) suggests that the fitness benefits of the morphological innovation protection are due to more than simply the added dimensionality of the search space and the added population diversity from reduced selection pressure. Such a result supports the idea of an asymmetry between morphological and control optimization, and demonstrates that the most effective solution to embodied robotic optimization includes the protection of morphological innovations.

The full comparison of these fitness values over optimization time are shown in Fig. 6.1. The visual inspection of these evolutionary trajectories demonstrate a typical logarithmic fitness improvement over the first 1000 generations or so, but then show a stagnation for the traditional optimization procedure without innovation protection, while the fitness values of the treatment with morphological innovation protection contrast these results by demonstrating sustained improvement further into the optimization period. The mean fitness values of the treatment without protection show no significant improvement ($p = 0.085$) from generation 1000 to 5000 (with average fitness values of 20.988 and 21.717 voxels, respectively). Contrary to this, the treatment which includes morphological innovation protection shows a significant improvement over this time, improving significantly ($p = 0.013$) from 25.925 at generation 1000 to 31.953 voxels at generation 5000.

Somewhat interestingly, the improvement in the controller innovation protection treatment is also significant ($p = 0.017$) over the same period of time,

through upon visual inspection this appears to be due to tight confidence interval bands more so than drastic fitness changes as the value improves from a mean of 21.385 to 22.049 voxels.

Additionally, the rapid improvement in the controller innovation protection and no protection cases during the first 1000 generations does not contradict the hypotheses of a fragile “morphological language”, as the coupling and dependency of the morphology and controller through this language would take time to be originally established – and would not introduce fragility into the system until it was established.

6.3.2 The Effect of Morphological Innovation Protection on Population Stagnation

The early stagnation of traditional evolutionary robotics without protection is indicated by the flatline in fitness value in Fig. 6.1, and suggests the prevalence of local optima in this space. The treatments which feature innovation protection of some sort appear to stave off stagnation to a greater degree, with visually noticeable results and significantly higher fitness values resulting from trials with morphological innovation protection. However, these averaged statistics provide little mechanistic insight into why and how protection is able to overcome the local optima.

Perhaps more telling than the average locomotion ability at the end of optimization time is the examination of the optimization process within each individual run. Figs. 6.2 and 6.3 represent typical runs, and help to give an intu-

ition of the optimization process. In these figures, each colored line represents a unique morphology, plotted by its locomotion ability over optimization time.

Note that these runs outlast the 5000 generation total from Fig 6.1 – as all trials were optimized for 5 days of walltime, before being truncated for comparisons with other runs. Some trials reached as far as 10,000 generations, while the minimum number completed was 5014 generations, thus 5000 generations was chosen as the cut off for comparisons of treatment averages – but figures for single trials show their full length.

The first thing to note is the continued improvement in performance of the most fit individual over optimization time in the case of morphological innovation protection (Fig. 6.2) – which is consistent with the trend seen on average in Fig. 6.1. This is not seen in the case without innovation protection (Fig. 6.3), where the best individual was found well before generation 2000 – and by generation 1000, fitness has reached 99.6% of its final value.

Consistent with the above observation, we also see that the most fit individual in Fig 6.2 changes rapidly, continually turning over in the trial with morphological innovation protection. As each color in the figure represents a unique morphology, we also see that a wide variety of different morphologies hold the title of “best-so-far”. On average, runs with morphological innovation protection 24.179 unique morphologies are the best-so-far at some point in optimization, where the runs without protection show significantly ($p = 1.555 * 10^{-6}$) less turnover, with just 10.115 unique robot body plans doing so.

The question of how the reduced selection pressure of morphological innovation protection may help to improve overall fitness and continued opti-




Figure 6.2: A single optimization trial featuring morphological innovation protection. Each unique morphology is represented by a random color. Note the continual improvement in the locomotion ability over optimization time, and the continual turnover in morphology of the top performing individuals. The pop out in this figure highlights an example of an “overtake” where the new child morphology (in red) initially performs worse than its parent morphology (in teal), only to outperform that previous morphology after successive control optimization to both the parent and child. In traditional greedy methods this initial drop in performance would signify a poor solution and thus remove the child morphology from the population – leading to stagnation. The example in the pop-out supports the idea that the continued improvement in the morphological innovation protection treatment is due in part to the ability to properly recognize the long(er) term potential of initially detrimental mutations and allow for such overtake events.




Figure 6.3: A single optimization trial with the same initial conditions as Fig. 6.2, but without any innovation protection. Again, each unique color represents a unique morphology. Note the early improvement in locomotion ability during the initialization period (the first couple hundred generations), but the stagnation which occurs immediately afterwards. Also note the prevalence of colored dots filling the space underneath the best performing individual, which represent new morphologies which initially performed worse than the current best individual and were thus rejected and thrown out of the population (as opposed to those individuals which were protected and eventually led to “overtakes” in Fig. 6.2). This treatment – where morphology and controllers are attempted to be optimized simultaneously and without innovation protection – represents the currently employed method in the field of evolutionary robotics [23].

mization may be best demonstrated in the pop out box for Fig. 6.2. Here we see the current best morphology, in teal. This morphology was unable to improve on itself for some time, as we see its fitness value (*y-axis*) flatlining. This “parent” morphology has a “child”, a new proposed variation of its morphology, highlighted in red. As the original fitness value of this morphology (its leftmost point, as the *x-axis* represents optimization time) falls below its parent, this individual empirically shows worse performance than its parent – and thus would not be considered as a viable solution in a traditional evolutionary method. However, since this new morphology does not have a controller that is well adapted to it (as the controller is specialized for the previous morphology, in teal), we should not expect it to outperform its parent. Thus morphological innovation protection keeps this individual in consideration as one which could hold long term potential, but does not show immediate promise.

Indeed, we see that after a number of controller optimization iterations later (occurring in equal amounts to both the parent and child during this intermediate period), the child morphology (in red) overtakes the parent morphology (in teal) – achieving higher fitness and demonstrating that it did indeed hold a better long term potential than its parent, despite the immediate drop in fitness. As the fitness of the parent (which has had more time to specialize its controller to its morphology) is outperformed by the child (which has had less time to fine tune its controller to its morphology), we assume that the parent is unlikely to be the most promising robot body plan in the long run, and thus remove it from the population.

We see this trend of “overtakes” – where children start out with worse fitness than their parents, but eventually outperform them – continuing through-

out this run (as the blue overtakes the red, and the green overtakes the blue in the pop out of Fig. 6.2). Unsurprisingly, we see morphological overtakes significantly ($p \leq 6.939 * 10^{-10}$) more often in runs that explicitly protect morphological innovations (an average of 76.714 overtakes in the first 5000 generations) than without any innovation protection (where there are only 1.432 overtakes) and the case with controller innovation protection (where there are just 1.333 morphological overtakes on average).

Interestingly, there is not a significant difference ($p = 0.533$) between the number of total controller overtakes in the controller innovation protection treatment (where a “child” is a robot with a new *controller* variation, that readapts its morphology to catch back up to its “parent” controller – which happens 74.542 times on average) and the number of morphology overtakes in the morphological innovation protection treatment. Combined with the finding in Fig. 6.1 that morphological innovation protection outperforms the other two treatments, this suggests a greater potential for the relative importance of morphological overtakes over controller overtakes – and again reinforces the asymmetry between morphologies and controllers from an optimization perspective.

6.3.3 The Effects of Morphological Innovation Protection on the Progression of Morphologies over Evolutionary Time

In the above example, we saw continued improvement throughout the length of optimization and search covering a larger number of unique morphologies. This suggests that evolution with morphological innovation protection is better

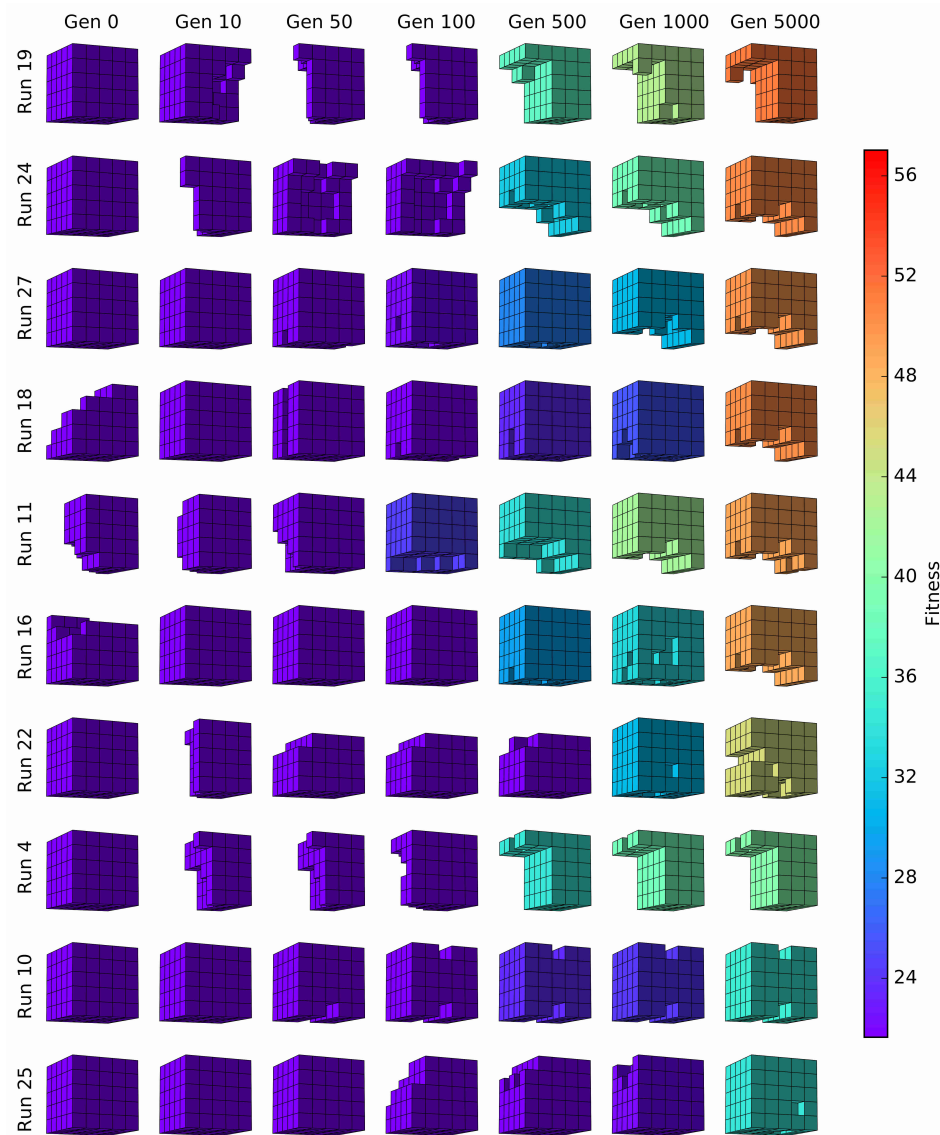


Figure 6.4: The progression of morphologies over evolutionary time with morphological innovation protection. Rows represent the top 10 (out of 30) performing runs at generation 5000, while columns represent snapshots of the morphology at various points during the optimization process. Note how some of the runs converge upon the same morphology (a front and back legged robot), despite starting from varying initial conditions. The color of the morphologies represent their fitness values. Note the progression from cooler to warmer colors over optimization time, as fitness values increase.




Figure 6.5: The progression of morphologies over evolutionary time with no innovation protection. Rows represent the top 10 performing runs in this treatment, while columns represents the progression of optimization over evolutionary time. The color of the morphologies represent their fitness values, and share the same color mapping from fitness to color as Fig 6.4. Note how neither the improvement in fitness (color), nor the change in morphologies occurs over evolutionary time, as the final morphologies found at generation 5000 closely resemble the state of their run as early as generation 50 – just 1% of the way into optimization time.

able to escape local optima and find more optimal solutions.

The improved optimization efficiency of morphological innovation protection is further supported upon visual inspection of the trajectory of morphological innovations over evolutionary time. Fig. 6.4 shows the current best individual at various points over evolutionary time for each of the top 10 runs in the treatment with morphological innovation protection. Note how the fitness values of the robots increase over time (from left to right, and indicated by the color of each robot). Also note how the final morphology of some robots (e.g. runs 24, 27, 18, 11, and 16) result in identical morphologies, despite starting from a range of starting morphologies and not finding this convergent morphology until hundreds or thousands of generations into the optimization process.

In contrast to the sustained turnover of morphologies shown above, Fig. 6.5 shows the snapshots of the 10 best runs in the treatment *without* innovation protection. Notice how the colors of the robots tend to show little change over the evolutionary process, mirroring the stagnation shown in Fig. 6.1. While convergence of the final morphologies is present here as well, the gross morphologies found here (variants of the a full cube with no appendages) are found early on in the search. These morphologies are often provided in the random initial morphologies at generation 0 (an artifact of this genetic encoding's tendency to start with simple shapes and complexify them over time), or found early on in search (by generation 10). In this treatment, gross morphological changes tend to be absent after generation 50 (just 1% into the full 5000 generations).

The differences between Figs. 6.4 and 6.5 suggests that the traditional method without morphological innovation protection tends to converge prematurely to morphologies that are present early on in the evolutionary search.

While the inclusion of morphological innovation protection may allow search to escape local optima and converge to “more global” optima. Since the common final morphology found in runs 24, 27, 18, 11, and 16 of Fig. 6.4 is not the best found in that treatment (as it is outperformed by run 19), it is obviously not the global optima. But the fact that a diversity of random initial conditions converged to this final morphology (and that the converged upon morphology is more fit than those initial conditions) suggests that morphological innovation protection is better able to search over a larger basin than treatments without it (i.e. are less sensitive to initial conditions and can better escape local optima). The convergence of morphologies across varied initial conditions is even further pronounced, and more visually impressive, in Fig. 6.10 below.

6.3.4 Generalization to More Complex Implementations

The results above demonstrate the effectiveness of the proposed algorithm on one specific instance of the optimization of embodied machines. However, the wide application of this algorithm also requires the demonstration of its effectiveness as the complexity of these machines scales up.

To explore the question of scale, we apply morphological innovation protection to the evolution of robots with higher resolution morphologies (up to $10^3 = 1000$ voxels) and also those with closed-loop neural network controllers. These both represent more complex instantiations than the lower resolution morphologies ($5^3 = 125$ voxels) and open-loop phase-offset controllers employed in the previous experiments.

In this particular robot morphology and encoding, the increased number of




Figure 6.6: The resulting average fitness value of the best individual at generation 5000 across 30 independent trials for various values of the minimum change threshold. These runs employed the more complex closed-loop neural networks controllers and higher resolution morphologies. The line in green represents the minimum morphological change for an “age reset” in the case of morphological innovation protection – resulting in the highest fitness performance when only changes of more than 40% of the potential voxels in a robot were mutated triggered a morphological age reset. The blue line represent the case of the minimum threshold for controller changes in the case of controller innovation protection. Here, the optimal value of the threshold parameter (triggering an age reset only when at least 5% of the potential synapses in a given robot are mutated) did not show significant improvement over the case with no minimum controller change threshold ($p = 0.237$). Thus minimum controller change thresholds were not considered in this work.

voxel cells that make up each robot allows for greater expressiveness and finer details in its morphology. However this also presents a challenge for the above




Figure 6.7: The mean fitness over evolutionary time of the best individual for 30 independent runs in the most complex setting tested here – with closed-loop neural network controllers and high-resolution morphologies. The average performance of evolution with morphological innovation protection with a 40% threshold outperforms all other treatments shown here (mean fitness of 13.909, all $p \leq 1.949 * 10^{07}$). Interestingly, without the threshold, the treatment with morphological innovation protection (mean fitness of 17.019), controller innovation protection (fitness of 16.914), or no protection do not show any significant difference (all $p \geq 0.0668$).

algorithm. As the total number of voxel cells increase, the effect of changing a single voxel (the minimum morphological variation) is reduced. In the extreme, the concept of readapting controllers since the last “morphological change” is less straightforward – as increasingly small changes can modify minor details of the morphology without affecting its overall function.

To help address the problem of non-functional morphology changes, we introduce a parameter to represent the minimum percentage of voxels that must be varied in order to qualify as a “gross morphological change”. It is important to note that this parameter is specific to the voxel-based soft robot implementation employed in this work – and thus the optimal setting of this parameter is not of great importance for its generalization outside of this soft robot encoding. But the general concept of a threshold for the minimum morphological change is a universal concept that could be applied to any robot instantiation, as necessary.

In the case of robots with neural network controllers and higher resolution morphologies, we find that resetting the “morphological innovation protection age” of an individual only after a mutation that changes more than 40% of their voxels produces optimal results. The 40% value was found via a parameter sweep, shown in Fig. 6.6. Interestingly, the benefit of controller innovation protection falls as the threshold for “controller innovation protection age” increases (showing optimal performance with small or no thresholds), so we ignore the threshold for minimum controller change here.

The increased efficacy of morphological innovation protection with this threshold is reinforced by viewing the fitness traces over time in Fig. 6.7, which show a significant improvement in fitness values at the end of optimization time over both the case without morphological protection ($p = 6.780 * 10^{-7}$) and with morphological protection without a threshold ($p = 5.225 * 10^{-9}$). Figs. 6.8 and 6.9 visually present the progression of morphologies over evolutionary time. Note how the presences of morphological innovation protection with this threshold enables new – more efficient – morphologies to be found even thousands of

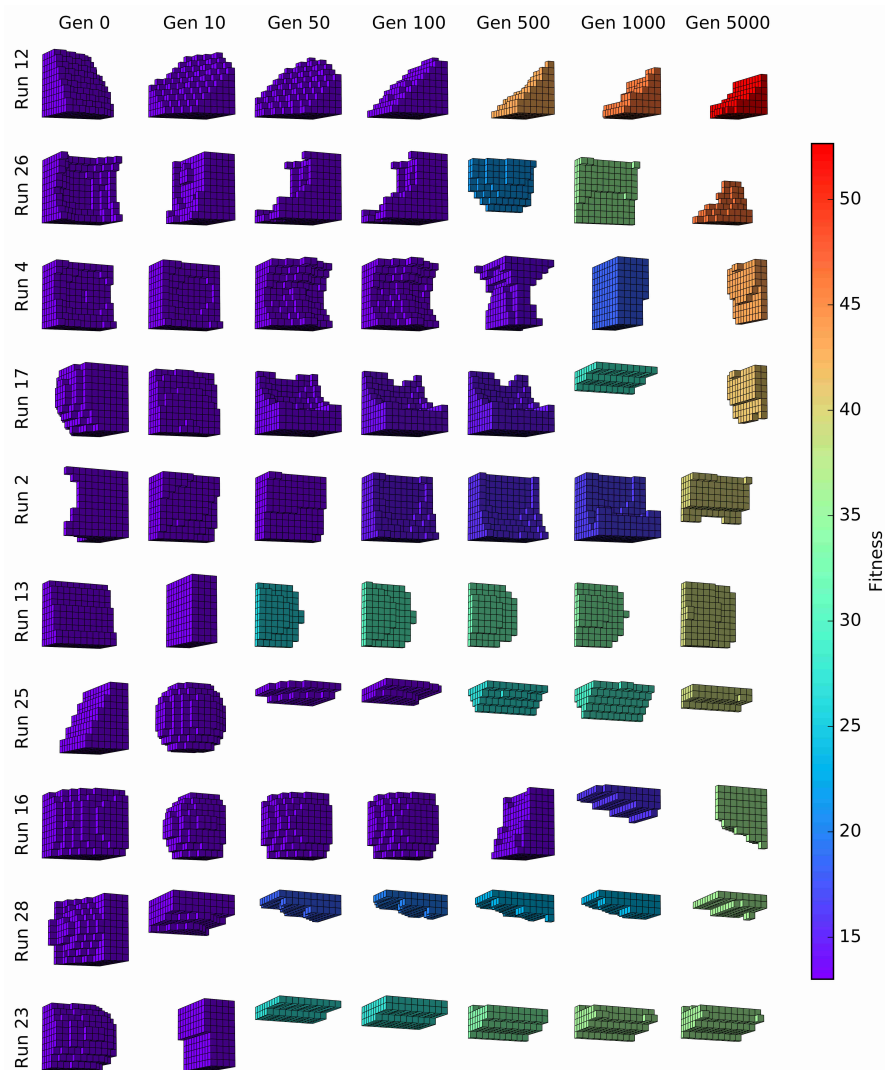


Figure 6.8: The progression of morphologies over evolutionary time in the setting with closed-loop neural network controllers and high-resolution morphologies and evolution with morphological innovation protection with a 40% threshold. The rows represent the top 10 (of 30) independent trials, while the columns represent the progression over evolutionary time. Color represent the fitness values of the robot (their locomotion speed), with warmer values depicting more fit individuals. Note the continued increases in fitness and change in morphologies late into optimization time. Also note that while exact convergence isn't as clear as other other experimental setups, many independent runs find the strategy of flat sheet morphologies which fall to lay flat on the ground and undulate to locomote.

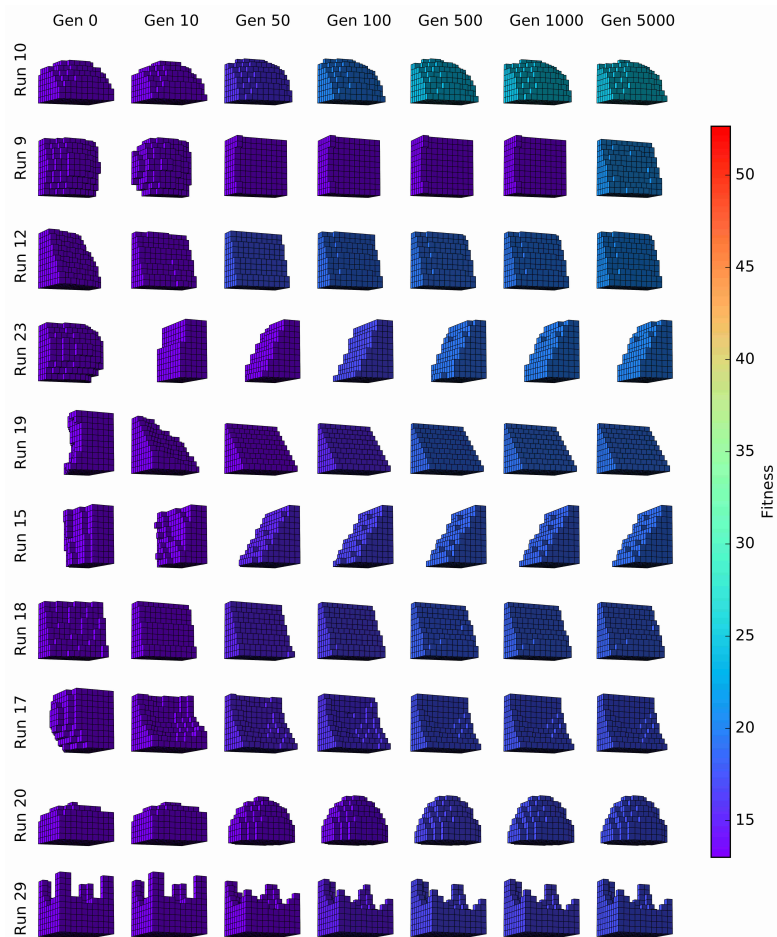


Figure 6.9: The progression of morphologies over evolutionary time in the setting with closed-loop neural network controllers and high-resolution morphologies and traditional evolution without morphological innovation protection. The rows represent the top 10 independent trials, while the columns represent the progression over evolutionary time. Color represent the fitness values of the robot, with the fitness to color mapping matching that of Fig. 6.8. Note that final fitness values and morphologies at the end of optimization time (gen. 5000) rarely change the state of their runs just 1% of the way into this optimization length (gen. 50). This suggest that the search space has many local optima, as there is one nearby each of the random starting conditions – and that traditional evolutionary robotics methods (without innovation protection) are unable to escape these local optima, cause premature convergence to suboptimal morphologies.

generations into optimization time.

6.3.5 The Potential for Convergence Across Initial Conditions

The case of high resolution soft robots with the phase offset controllers (in Figs. 6.10 and 6.11) is perhaps even more interesting, as they show extreme convergence of morphologies across varying initial conditions. Similarly to above, we find the optimal threshold value with a parameter sweep, this time resulting in an optimal value of 20%. That the threshold in this scenario differs from the optimal 40% value in the case above is interesting to note, as it suggests that the optimal threshold value for minimum morphological changes may be dependent on both the implementation of the morphology and that of the controller. The differences in threshold values for different controllers makes intuitive sense, as some controllers (like the open-loop distributed phase offset oscillating controller) may be more robust to changes in morphologies than other (like the closed-loop neural network controllers), and thus better able to generalize across closely related morphologies.

The 20% threshold value again shows a significant improvement over the treatment with no morphological innovation protection ($p = 1.208 \cdot 10^{-5}$), shown in Fig. 6.12. Though interestingly, the morphological change threshold does not provide a significant fitness improvement over morphological innovation protection with no threshold in this setting ($p = 0.156$). Also note, in Fig. 6.12, how the fitness curves for this thresholded value appear to flatten out later in optimization, as fitness gains slow down. The explanation for this may come in Fig. 6.10, where most of the runs depicted have converged to the same gross




Figure 6.10: The progression of morphologies over evolutionary time in the setting with open-loop phase-offset controllers and high-resolution morphologies and evolution with morphological innovation protection with a 20% threshold. The rows represent the top 10 (of 30) independent trials, while the columns represent the progression over evolutionary time. Color represent the fitness values of the robot (their locomotion speed), with warmer values depicting more fit individuals. Note the convergence of all 10 of these runs to the same morphology at the end of optimization time (and many of these runs find this gross morphology by generation 500). This convergence occurs despite a variety of different initial conditions across these trials.




Figure 6.11: The progression of morphologies over evolutionary time in the setting with open-loop phase-offset controllers and high-resolution morphologies and evolution without morphological innovation protection. The rows represent the top 10 independent trials, while the columns represent the progression over evolutionary time. Color represent the fitness values of each robot, and is consistent with that in Figs. 6.10. Note the lack of high fitness values and lack of morphological change, with most runs finding their final morphology by generation 50.

morphology by generation 500, and only make fitness improvements from controller changes thereafter.

The ability of evolution to converge to the same high-performing morphology across many independent trials, despite starting from different initial conditions suggests that (in this particular soft robot implementation) the inclusion of thresholded morphological innovation protection is able to escape the local optima around these starting conditions and find “more global” optima in this search space. The lack of this convergence in Fig. 6.8 (the case of closed-loop neural network controllers) may suggest a differently structured search space or simply the inefficiency of the algorithm to find such optima in the allotted time with a more complex controller to optimize.

In the case without any protection, search stagnates quickly and again appears unable to escape the local optima near its initial conditions (Fig. 6.11).

Interestingly, the low resolution soft robot implementation employed for the primary experiments above does not benefit from the inclusion of a threshold. Fig. 6.13 shows the fitness traces over time for this soft robot implementation, with no significant improvement shown for a 20% threshold in this setting. As the only difference between this instance and Fig. 6.12 is the resolution of the robot, it is likely the case that non-functional morphological changes are less prevalent when the robot is composed of fewer and discrete subcomponents.

The inclusion of a threshold for the minimum morphological change appears to hold potential for the generalization and efficiency of this proposed algorithm. Though the complete characterization and understanding of this parameter in the soft robot implementation used here, as well as other robot

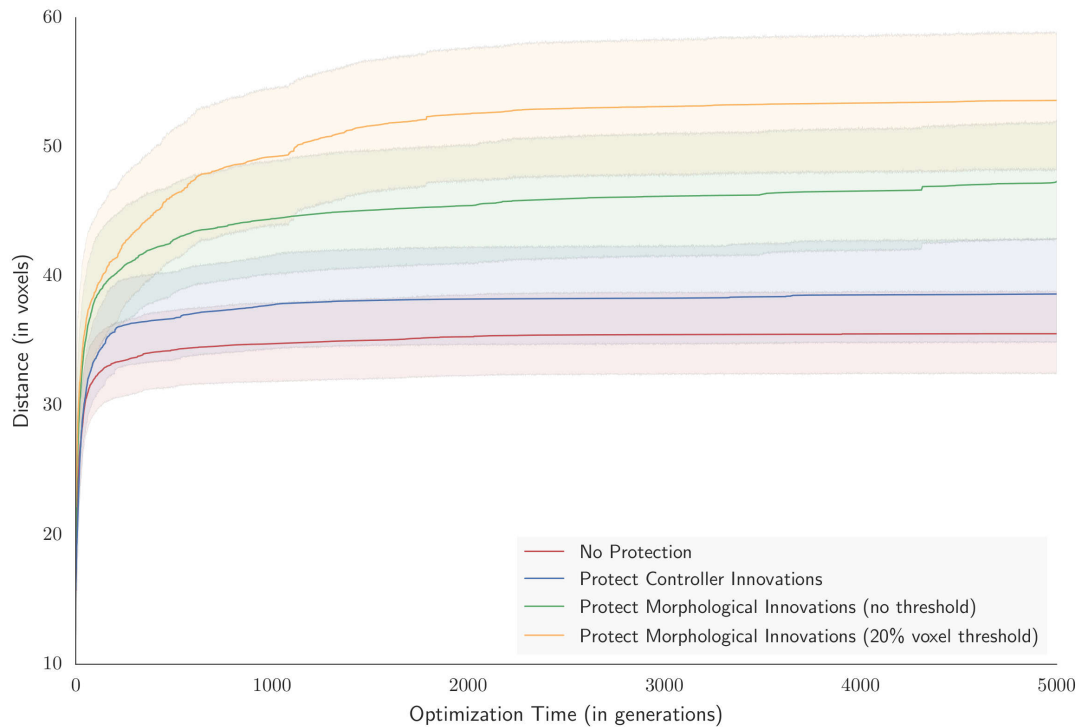


Figure 6.12: The mean fitness over evolutionary time of the best individual for 30 independent runs with open-loop phase-offset controllers and high-resolution ($10 \times 10 \times 10$) morphologies. The average performance of evolution with morphological innovation protection with a 20% threshold (mean fitness of 53.334) outperforms the trials without any innovation protection (mean fitness of 35.559, $p = 1.208 \times 10^{-5}$) or with controller innovation protection (mean fitness of 38.632, $p = 3.274 \times 10^{-4}$). However having the optimal 20% threshold for minimum morphological changes still does not significantly outperform the case with no threshold at generation 5000 (mean fitness of 47.335, $p = 0.156$). The usage of morphological innovation protection with no threshold does significantly outperform both the no protection and controller innovation protection treatments though (both $p \leq 2.436 \times 10^{-3}$).




Figure 6.13: The mean fitness values over optimization time for the lower-resolution ($5 \times 5 \times 5$) robots with open-loop phase offset controllers. This figure is identical to Fig. 6.1, plus the inclusion of a treatment with morphological innovation protection and a 20% minimum morphological change threshold. This threshold does not show significantly higher fitness (32.519) than the case of morphological innovation protection without a threshold (31.425, $p = 0.859$). This differs from the case of high-resolution robots and open-loop phase offset controllers in Fig. 6.12, suggesting that the resolution of the morphology may be related to the benefits of a threshold.

implementations, is required in future work.

6.4 Discussion

The above results demonstrate a new method for entire robot (brain and body plan) evolution that is more scalable in terms of continued optimization for longer periods of time, and better resulting fitness than the traditional evolutionary method (without innovation protection). This method for “morphological innovation protection” helps prevent premature convergence to the many local optima which appear to be present in the rugged search space of robot morphologies and controllers [40].

The hypothesis from [40] that the fragile co-optimization of brain and body plan caused by specialization of one sub-component to the other is consistent with the findings above. This work also reveals that there is an asymmetry between the brain and body plan: protecting innovations to the morphology leads to more effective optimization than protecting innovations to the controller.

The benefits of the temporarily reduced selection pressure provided by morphological innovation protection suggests that the long-term potential and immediate fitness impact of a morphological mutation are not always correlated. Thus, we require a form of diversity maintenance, such as innovation protection, which helps evolution to rate proposed solutions based on long-term potential rather than on immediate fitness impact. As was shown here, this protection can help to reduce premature convergence in the search space and stagnation at suboptimal values.

We should point out that many other diversity maintenance methods already exist (such as fitness sharing [246], novelty [170], speciation [275], or binning/niching [203]). Yet we believe the proposed method here to be unique from these other forms of diversity maintenance, as we reduce selection pressure only for individuals with new morphologies (and not those with new controllers) to allow them to adapt. It could be argued that this takes place also by parallel (or random restart) hillclimbers [188] that are instantiated with a diverse set of morphologies, or certain forms of speciation – but such methods would not allow for the direct competition between evolving individuals, based on their current level of controller-to-morphology specialization, and thus may not perform as efficiently (though future work is needed to compare the proposed method against all such competitors systematically).

We believe this to be the first example of a design automation algorithm for robotics that considers the interdependence of neural controllers and body plans (specifically arising from embodied cognition [219]) and to use this intuition to propose a method to escape local optima in the fitness landscape of embodied machines.

But despite the significant improvement to our ability to simultaneously optimize the brain and body plan of embodied robotics, there is much work still to be done. Firstly, the proposed method was only applied to one class of robot. This class may actually represent the simplest form of brain-body co-optimization because the distributed sensing, actuation, and information processing that the cellular soft robot paradigm was designed to possess [43, 38] helps to blur the line between physical interactions of the morphology with the environment and information processing of a controller.

In the case of centralized controllers and robots composed of rigid components, topological rather than parametric changes to the cognitive architecture would be required for control readaptation if morphological mutations add or remove physical components. Future work should explore the effect of morphological innovation protection in such a paradigm, where there is the potential for morphological changes to more drastically change the function of the robot – and thus for readaptation to those morphologies to play an even more critical role in optimization.

In these experiments, the genotype encoding of the soft robots was modularized such that one part of the genome dictated the shape and material properties of the robot and a separate part encoded the actuations in the form of volumetric deformations of the voxels during behavior. In these robots (and those with even more complex phenotypes), one can envision different splits as to what constitutes morphology and control – and what qualifies as a “morphological change” in this proposed algorithm. It would be of interest to investigate the effect of various splitting points of this dichotomy on the efficacy of morphological innovation protection.

Furthermore, the very idea of the split was noted above to be a false dichotomy, as information processing and physical processes happen throughout the agent. Rather than various mutation operators which affect only the genotypes of the “brain” or the “body plan”, and phenotype mappings which take information only from one module of the genotype, it is necessary in our artificial agents (as is the case for biological organisms) that a developmental process with a richness of crosstalk and feedback loops be the mechanism responsible for co-adapting the brain and body plan as they grow together.

Engineering the optimal way to create such a developmental process, and understanding how such processes can impede or increase evolvability in biological populations, remains an open problem. However, the method introduced here may serve as one step in this direction – and provide a more effective way to evolve the brains and body plans of embodied machines until a developmental method that enables rapid readaptation during the growth process is available.

The biological analogs of the procedure of “protecting morphological innovations” are not entirely clear, but some rough analogies can be drawn. For example, the fact that brains learn and adapt at a much faster rate than bodies grow fits into this paradigm, as does the even slower change of gross morphological features over evolutionary time. In this way, the readaptation of controller strategies for varying body plans is built into the neuroplasticity of the brain. The speed at which neurogenesis and morphogenesis occur is certainly constrained by the energetic resources of each process, but herein lies a potential benefit for the automated optimization of behavior and form in machines. The basic idea of “protection” is simply a diversity maintenance measure resulting in temporarily reduced selection pressure on specific individuals within an evolving population.

One could also imagine periods of evolutionary time when an entire species is under relatively little morphological selection pressure before an environmental shock suddenly reapplies that pressure. Also possible are periods of a single individual’s lifetime when selection pressure varies: For example, human infants may not be selected highly for their locomotion speed, as their parents tend to physically carry and protect them while their brain and body develops.

The work reported here could then be viewed from the perspective of learned versus innate behaviors (exploring the benefits of neuroplasticity and development), or used to pose questions about how a varying level of selection pressure could help lead to more complex behaviors, such as upright walking.

We should also note that the results presented here do not explicitly show an increase in the complexity of evolved agents with and without this method. They do, however, show that this method allows for increases in sustained optimization, leading to higher fitness values in the CPPN-based soft robot domain. This domain is more complex in its shape (e.g. number of components) and material properties (e.g. compliant bodies) than previous benchmarks [263, 264], and this method has been shown to cope with the difficulties in co-optimizing body and brain that arise in such a domain.

We do not believe that the results stated here are (at a high level) restricted to this particular domain. The above algorithm is simple to implement (requiring only: an age counter, a check for variations in brain and/or body for each mutation, and –optionally – a criterion for the minimal gross morphological change), and thus we believe it will be widely applicable. Future work will test this supposition.

Due to the recent interest in co-optimization of neural network topology and weights [94, 199, 160], we should also note that the domain of this work – an agent’s controller embodied within its morphology – is closely related to that of neural network’s weights embodied within its topology. Future work will show whether the method proposed here will show similar gains in the design of neural network topologies.

6.5 Conclusion

We demonstrate an example of a robot design automation algorithm that considers the interdependence of neural controllers and body plans (due to the theory of embodied cognition) on the optimization process. We use this intuition to temporarily reduce selection pressure on newly mutated robot morphologies, thus allowing the agents to readapt their controllers and better escape local optima in the fitness landscape. We have shown that this technique – deemed “morphological innovation protection” – produces evolutionary optimization which delays premature convergence and stagnation, and results in more efficient evolved robots. We showcase the ability of this technique to escape local optima in the search space by demonstrating the convergence to a similar morphology across many independent trials from randomly initial conditions. While we hope that this technique will be surpassed in the future by a developmental process with feedback loops between the body and brain, we propose the above algorithm as a short term improvement over the current techniques for the co-optimization of morphology and control in virtual creatures.

CHAPTER 7
APPLICATIONS OF MORPHOLOGICAL DESIGN AUTOMATION
TO ENGINEERING PROBLEMS OUTSIDE OF ROBOTICS

Abstract of Chapter ¹

Natural frequency tuning is a vital engineering problem. Every structure has natural frequencies, where vibrational loading at nearby frequencies excite the structure. This causes the structure to resonate, oscillating until energy is dissipated through friction or structural failure. Examples of fragility and distress from vibrational loading include civil structures during earthquakes or aircraft rotor blades. Tuning the structure's natural frequencies away from these vibrations increases the structure's robustness. Conversely, tuning towards the frequencies caused by vibrations can channel power into energy harvesting systems. Despite its importance, natural frequency tuning is often performed ad-hoc, by attaching external vibrational absorbers to a structure. This is usually adequate only for the lowest ("fundamental") resonant frequencies, yet remains standard practice due to the unintuitive and difficult nature of the problem. Given Evolutionary Algorithms' (EA's) ability to solve these types of problems, we propose to approach this problem with the EA CPPN-NEAT to evolve multi-material structures which resonate at multiple desired natural frequencies without external damping. The EA assigns the material type of each voxel within the discretized space of the object's existing topology, preserving the object's shape and using only its material composition to shape its frequency response.

¹Appeared as: Cheney, N., Ritz, E., & Lipson, H. (2014). Automated Vibrational Design and Natural Frequency Tuning of Multi-Material Structures. In Proceedings of the 2014 annual Conference on Genetic and Evolutionary Computation (pp. 1079-1086). ACM.



Figure 7.1: (*left*) Resonant vibrations are a common source of fragility in natural objects. Here a glass fails in the presence of an acoustic wave at resonance. (*right*) On a larger scale, structural failure through periodic loading of the Broughton Suspension Bridge caused its collapse in 1831 due to the resonance of the soldiers marching in lockstep [115].

7.1 Introduction

In engineering mechanics, the response of a structure to vibrational loads is of acute interest and importance. Every object will exhibit some motion when excited with any periodic load – such as vibrations. But vibrations at certain frequencies will excite certain objects with greater intensity. The natural frequencies of each object dictate the frequencies of vibrations which exhibit particularly intense responses. In many cases, this type of loading is harmful to the system – larger responses put more stress on the structure. The energy trapped in the structure through these oscillating motions is not efficiently dissipated into the surrounding environment in these cases, leading to weakening or even failure over time (Fig. 7.1).

Large arch dams in earthquake prone areas are one example of antagonistic periodic loading, where oscillations at resonance encourage cracking and rupture through brittle concrete [1]. Wind turbine blades are another instance of this, where rotation at a resonant frequency may cause tip deflection that stresses the blades and lowers their effectiveness [117]. In aerospace applications, or other vehicles producing massive propulsive loads, the vibrations caused can be particularly damaging [65].

On the other hand, some applications may benefit from increased oscillatory response. Oscillations of larger magnitude contain more energy, which are ideal for power harvesting mechanisms such as microelectromechanical systems (MEMS) like piezoelectric microcantilevers. In this case, the lack of energy dissipation from structures with resonant frequencies close to the frequency of the vibrations enhances the efficiency of these energy harvesting systems [35, 257].

However, the process of tuning structures to have specific resonant frequencies remains largely unintuitive – changes to the material properties at one point in a structure often lead to non-linear effects on its vibrational response, both near and far from the point of change. Thus existing methods usually change the shape of the structure, often involving the addition of mass dampers to the system or requiring significant structural modifications. In cases where weight and size are at a premium, such as aerospace or remote sensing, these solutions are unsatisfying. Additionally, many of these damping strategies are only useful for the first (“fundamental”) resonant frequency.

In many cases, important structures must interface into a pre-existing system, be robust (free from fragile additions), or serve aesthetic purposes. In

existing structures the topology has often already been optimized, such as airfoils tuned for aerodynamic efficiency. In these cases, it would be advantageous to perform vibrational optimization of structures which preserves the overall shape of the object while also giving it desirable vibrational properties.

Muti-material design involves the use of different material types within a given shape to produce overall object properties outside of those available to the given shape composed of a single material. By optimizing a structure's vibrational response through multi-material design within an existing envelope, we open new doors towards the study and implementation of vibrational optimization of these fixed topology structures. Additionally, with current advances in multi-material additive manufacturing, we now have the ability to specify the placement and interwovenness of individual material droplets with vastly different properties during the manufacturing of a given structure, making such designs physically realizable today.

In this study, we optimize the two dimensional projection of a fixed-free cantilevered beam, with the first ten natural frequencies optimized to reproduce a randomly chosen resonant frequency profile. This is a particularly difficult and unintuitive problem because the material properties at each voxel are coupled (often non-linearly, and non-locally) with the material properties at every other voxel to produce the vibrational response of the object as a whole. Furthermore, this static topology responds differently to vibrations at different frequencies, making ad-hoc tuning of more than one or two natural frequencies exceptionally challenging.

Due to the unintuitive nature of the problem, we use evolutionary computation to traverse this design space – specifically the evolutionary algorithm

CPPN-NEAT. We choose this because the Compositional Pattern Producing Network (CPPN) genome provides a compact and evolvable representation of the discretized physical design space necessary for this problem. We use this EA to optimize the placement of two materials (stiff and soft) at each discretized voxel of the structure's original shape envelope. We optimize towards randomly chosen frequency profiles (which specify the first ten natural frequencies of a structure), and show the promise of this approach to become an automated design platform for structural vibration optimization going forward.

7.2 Background

Controlling system performance through the frequency domain is a classical idea [144], but conventional engineering strategies have remained largely unchanged since the early 20th century. In many cases, harmful vibrations are attenuated by directing energy away from the most sensitive parts of the system, to another auxiliary system. This is done through the addition of tuned vibrational absorbers (TVA) [140, 280, 286]. However these systems come with the tradeoff of increasing the mass and complexity of the original structure. This typically involves augmenting the system with spring-mass element or small cantilever with a first resonant frequency tuned to that of the undesired excitation, where the undesirable energy is contained within oscillations of this auxiliary structure until it is dissipated through friction.

This process is not only inefficient, but fails to fully explore the design space - as the attachment location and parameters of these devices are often chosen through intuition and physical guess-and-check iterations by the engineers, and

are thus biased by the engineer's assumptions and training. Additionally, the amount of energy that can be absorbed by these types of systems is limited. As the TVA becomes large, its own dynamics begin coupling with the original system. This not only produces harmful vibrations within the original structures, but also changes its resonance profile and further complicates design. Furthermore, volume constraints of the system, such as inside the fairing of a rocket in launch, inherently restrict the size and shape of a TVA.

Some work has examined changing the system's topology to have a desired frequency profile without the addition of other components. Many of these strategies optimize the placement or parameter settings of a few predetermined basic structures, such as rods or trusses [114, 270, 75]. Duhring et al. studied the automatic design of structures with desired natural frequencies using homogenization [87]. However, this method could only optimize for a single frequency band, trying to maximize or minimize the frequency response between a set of frequency borders. This strategy also involved the generation of structures with complex and unpredictable shapes, rather than optimizing the response of an existing structure's topology. Du and Olhoff used topological optimization to automate the design of a voxelized plate structure with a binary material array to minimize sound power flow [86]. However, their designs were also constrained to the optimization of just a single frequency.

Our system expands upon these strategies by allowing us to optimize an arbitrarily large number of natural frequencies of a structure. Additionally, any number of material values can be used for any geometry. Since our search of the design space is topology-preserving; the final product will still have the shape of the original design, ensuring it will have the same functionality and maintain

previous topology optimization, with only natural frequencies changed.

7.3 Methods

Our method involves tuning the natural frequencies of a structure, which for lightly damped systems well approximates the resonant frequencies (where vibrational energy resonates to create sustained oscillations). First, the user specifies a design geometry. In many practical cases this is a pre-existing object shape, in this study we simply use a 2D projection of a cantilever beam, fixed on one end. Next the user produces a list of n desired natural frequencies. These frequencies are again dependent on the specific application; vibration in the environment can be found by measuring excitation loads with an accelerometer and applying the Discrete Fourier Transform [85]. Here, we select the frequency profile randomly (Sec. 7.3.4). Finally a selection of materials from which to construct the object is necessary. These depend on the additive manufacturing capabilities and supplies available to the user, and in this case is simply represented by two idealized materials, one an order of magnitude stiffer than the other. The structure to be optimized is meshed in a uniform voxelized grid. In this study, the beam is discretized into a 40x10x1 set of uniform cubic voxels. Each of these voxels is assigned a material according to the phenotype described by its associated Compositional Pattern Producing Network (CPPN) genome (Sec. 7.3.2). The natural frequencies of the structure are calculated (Sec. 7.3.1), and used to determine the individual's fitness (Sec. 7.3.3). Individuals who best minimize the error between the calculated and desired natural frequency profile are disproportionately favored to reproduce (and are subject to both genetic mutations and crossover in this process), creating the next generation to again iterative this

evolutionary process.

7.3.1 Approximating Natural Frequencies with FEM

Determining the harmonic behavior of an object is equivalent to solving for the eigenvalues of the matrix representing the FEM mesh of the structure. The shape and material of the system will govern the frequency at which each of the n desired natural frequencies lie. There is one natural frequency per degree of freedom of the system, but typically the number of nodes required for accurate simulation far exceeds the number of natural frequencies in an engineering range of interest. In this case of the 40x10 discretized voxels, the 400 voxels are approximated by 1301 nodes, with 8 nodes per quadrilateral element (many nodes are shared between adjacent elements). This is more than adequate to approximate the first 10 natural frequencies. These are computed using a generalized conjugate residual method, to a residual error $< 10^{-8}$. For this computation we employ Elmer, a popular open source finite element software for multiphysical problems, developed and maintained by the CSC - IT Center for Science [186].

7.3.2 CPPN-NEAT Evolutionary Algorithm

CPPN-NEAT has been repeatedly described in detail [271, 56, 54, 104], so we only briefly summarize it here. A Compositional Pattern Producing Network (CPPN) is similar to a neural network, but its nodes contain one of multiple mathematical functions (sine, cosine, Gaussian, sigmoid, linear, square, or pos-




Figure 7.2: The CPPN genome iterates through each voxel in the discretized design space, placing either a soft or stiff voxel at each location to produce the phenotype structure it encodes (Sec. 7.3.2)

itive square root). CPPNs evolve according to the NEAT algorithm, which is largely based on: complexification of genomes over time, speciation within the genotypic space for diversity maintenance, and tournament selection within species [271].

The CPPN produces a spacial output pattern that is built up from these functions' geometric transformations of the input gradients (changing values of each input coordinate over the space). Because the nodes have regular mathematical

functions, the output patterns tend to be regular (e.g. a Gaussian function can create symmetry and a sine function can create repetition). In this paper, each voxel has x , y , and z coordinates, polar coordinates (r and θ) in each of the x , y , and z planes, and a measure of the voxel's distance from the center of the search space (d), which are all input in the range $[-1, 1]$, describing the relative location of each voxel in this geometric design space. The single output of this network is interpreted as either a stiff material present in that voxel's location for a value above zero, or a soft material present for a value less than or equal to zero.

Alternate experiments were conducted with the value of the output node ($[-1, 1]$) representing a mixture of the stiff and soft materials with a compliance representative of an interpolation between the maximally stiff or soft materials at each endpoint. While the real valued compliance affords more flexibility in the design space, the problem of physically mixing materials during the additive manufacturing process is not yet a solved problem [126].

In order to produce a structure from a CPPN, each voxel in the discretized design space is iterated through. At each iteration through this space, the voxel's coordinate are input into the network, which is then undergoes its update function to produce an output value. The value out this output determines the type of voxel which is placed at this location. A positive output stipulates the placement of a stiff voxel (denoted by to a Young's Modulus (E) of 10 gigapascal (GPa), and Poisson's ratio of 0.3), while a negative output value produces a soft voxel at this location (resulting in a Young's Modulus of 1 GPa, and the same Poisson's ratio of 0.3). After iterating through each voxel within the design space, the CPPN genome has produced a phenotypic description of the structure, which can then be analyzed to find its natural frequencies, or sent to

a 3D printer for fabrication. This iterative process is outlined in Fig. 7.2.

7.3.3 Fitness Function

The quality (fitness) of each structure's ability to match a desired frequency profile of n frequencies is described by:

$$\frac{1}{\sum_{i=1}^n \left(\frac{n-i+1}{n}\right) E(i)^2}$$

Where n is the number of frequencies being optimized, and $E(i)$ is the relative error in matching the i^{th} frequency. The linear weighting factor $\frac{n-i+1}{n}$ is to place more weight on the primary frequencies than the later ones, which contribute less significantly to the behavior of real world systems. For example: for frequency 1 of 10, $\frac{n-i+1}{n} = \frac{10-1+1}{10} = 1$, thus the first frequency contributes to the fitness penalty with a weight of 1. For frequency 6 of 10, this weighting factor becomes $\frac{10-6+1}{10} = 0.5$, so the 6th frequency is discounted such that its relative error is only counted by half of that belonging to the first frequency. Similarly, the 10th frequency has a discount factor of 0.1, meaning that its relative error contributes to the overall fitness only one tenth the amount which the error of the 1st frequency does.

The motivation for using the weighted sum of the squared errors is simply a standard practice to negate the sign of the error terms, and to more heavily penalize larger deviations from the desired frequencies.

7.3.4 Producing Random Target Frequencies

n	Full Soft	Full Stiff	Rand. #1	Rand. #2	Rand. #3	Rand. #4
1	0.782	1.596	1.215	1.062	1.044	1.315
2	4.666	9.549	7.994	7.279	6.858	6.613
3	11.973	24.894	18.592	16.059	17.447	19.177
4	12.220	25.087	15.594	17.749	21.404	20.691
5	22.080	45.475	30.026	29.250	33.929	38.461
6	33.513	69.209	49.484	52.985	52.274	49.015
7	35.868	74.574	51.903	59.555	55.862	58.415
8	45.964	95.121	65.639	76.686	59.337	69.871
9	59.072	122.436	84.398	79.505	77.530	97.132
10	59.598	123.911	105.121	87.085	85.164	97.485

Figure 7.3: The first 10 natural frequencies (in MHz) for beams with all soft voxels or all stiff voxels gives the boundaries for the creation of 4 random frequency target profiles (Sec 7.3.4).

In order to optimize a structure to match a given frequency profile, such a profile must exist within the limits of the structure's realizable material properties. Thus to produce random frequency profiles to be used as targets for the optimization process, we use the following:

$$f_n = \frac{f_n^{\text{stiff}} + f_n^{\text{soft}}}{2} + \frac{f_n^{\text{stiff}} - f_n^{\text{soft}}}{2} * \text{rand}(-0.5, 0.5)$$

For frequency number $n = 1, 2, \dots, 10$. Where f_n^{stiff} is the n^{th} frequency of a beam fully populated with maximally stiff voxels (and conversely with soft voxels for f_n^{soft}). Thus the above equation produces a random number from the middle 50% of the range between these minimum and maximum frequencies. Much like the fact that a primary frequency above or below that of the fully stiff or soft beams is not physically realizable with the chosen materials, the frequencies at

(or near) the edge of this allowable range necessitate beams that are (almost) entirely filled with stiff/soft voxels, and thus may or may not be realizable in a case where more than a single frequency is being optimized for. For this proof of concept we conservatively constrain ourselves to the middle 50% of this range, where we feel a balance of both materials is likely to produce a vast array of frequency responses from the relative positions of voxels rather than the proportion of voxel types themselves. Future work will experimentally examine how optimization success drops off as these goals become less and less physically realizable with the expansion of the allowable target frequency range.

After all 10 frequencies are chosen, the values are then sorted from smallest to largest, as natural frequencies must occur in a monotonically increasing order. The 10 target frequencies for four random seeds are given in Fig 7.3.4, as well as those of the beams fully populated with stiff or soft voxels.

7.3.5 Control Treatment

In order to test the validity of the evolved structures, a control method was devised. To isolate the effects of the optimization towards natural frequency matching structures, the control groups consist of the same 32 independent runs of 30 CPPN genomes evolving for 1000 generations. This negates any natural advantage that the genomic representation of the CPPN might have in this domain, as well as the complexification of the NEAT algorithm over time. In the control setup, however, no preferential reproduction is afforded to those individuals who more closely match the desired frequency profile, but rather this selection happens at random. At the end of the 1000 generations of this evolu-

tionary drift, the resulting structure are still compared to each of the four Random Frequency Profiles and their effectiveness is measured in the same manner as with the experimental treatments.

7.4 Results

All treatments below consist of 32 independent runs, with populations of 30 individuals evolved for 1000 generations.

7.4.1 Statistical Measures

The data resulting from the control conditions pass the Shapiro-Wilk test for normality [256] (with $p > 0.0329$ for all 10 frequencies), and thus error bars are used to describe this data in the following plots. However, the experimental treatments routinely fail the Shapiro-Wilk test for normality, with many of the 10 frequencies for each random frequency profile falling above the $p = 0.05$ confidence cutoff for the normality of the distribution. Given that the shape these distributions are unknown, we employ bootstrapping to produce 95% confidence intervals to graphically describe the experimental data [88]. Given that at least one distribution for any statistical tests will be an experimental treatment, we employ the Mann-Whitney U test, as it does not require the assumption of normality of the data, yet performs almost as well as the student's t-test on normally distributed data (such as the control data) [194].

7.4.2 Statistical Data

Optimization for Random Frequency Profiles

We optimize the material makeup of our 40x10x1 voxel beam to set its natural frequencies to resonate with Random Frequency Profile #1. This is done for 32 independent trials, using a "soft" material (Young's Modulus(E) of 1 gigapascal(GPa), and Poisson's ratio of 0.3), and a "hard" material (Young's Modulus of 10 GPa, and Poisson's ratio of 0.3) as the material library. We compare the relative error for the task of matching with Random Frequency Profile #1 with the control condition for each of the $n = 10$ desired frequencies, resulting in 10 different p-values. A one-sided U test is employed in order to test the hypothesis that optimized structures will produce lower errors than the control structures. All of these 10 measures fall below the 95% confidence threshold, as $\max(\text{p-values}) < 1.008 * 10^{-5}$, showing statistical significance that our system can effectively optimize all of the first 10 frequencies of this desired frequency profile.

In order to test the sensitivity of this analysis on this randomly generated target frequency profile, three additional desired frequency sets are randomly produced. These each consist of $n = 10$ desired natural frequencies, but each of these frequencies differs from those in Random Frequency Profile #1, so a comparison coupled by desired frequency (in MHz) is no longer possible. One could imagine instead coupling by frequency number, and comparing the accuracy of $n = 1$ for each of the four frequency profiles, then comparing $n = 2$ for all four, and so on. However, actual frequencies between desired profiles can greatly differ: even when comparing with the equivalent frequency number (for exam-

ple 15.594 MHz for Random Profile #1 is the 4th desired natural frequency of the system, which has the value 21.404 MHz for Random Profile #3 – a value 37% higher). Additionally, with the highly non-linear nature of the frequency tuning domain, the effects of differences in other frequencies may play a major role in the overall structural evolution strategy and influence desired frequencies that do happen to be closely coupled by both n and frequency value (in MHz). Instead, we attempt to compare “apples-to-apples” but assessing each of the optimized structures against control structures evolved through a random walk. These random control structures are then compared to the same target frequencies for each of the Random Frequency Profiles. As seen in Fig 7.4, each of these four independently generated random frequency profiles show a statistically significant improvement over their associated control treatments (with the max p-value of the Mann-Whitney U test $> 1.008 * 10^{-5}$).

10 Materials vs. 2 Materials

Fig. 7.5 shows the effect of intermediate materials. In this setup, we optimize structures via the same method as in Sec. 7.4.2, where the CPPN genome specifies the material present in each voxel of the structure by again choosing from one of two materials – either stiff ($E=10$ GPa) or soft ($E=1$ GPa). To explore the impact of this assumption of two available materials, we compare this to another treatment in which there are 10 material choices, representing materials of intermediate compliances (between and including the endpoints of the fully stiff and fully soft voxels as before). Materials 2-9 are assigned a Young’s Modulus (E) of 2- 9 GPa, respectively, with a constant Poisson’s ratio of 0.3. The results of a two-sided Mann-Whitney U test show that 8 of the 10 desired frequencies

showed no significant difference ($\min(\text{p-value}) > 0.0745$). The exceptions to this rule were the frequencies of 15.594 MHz ($n = 4$), which had a p-value ; 0.029, and 105.121 MHz ($n = 10$) with p-value ; 0.0005. In the case of the final ($n = 10$) frequency of the group, the 10 material structures actually performed worse than their 2 material counterparts. Thus, in general, we believe that using 10 intermediate materials does not typically produce significantly better optimized structures than those made with just the two extreme materials.

Optimization for the First Natural Frequency

In the course of the work, we also explored other fitness function setups. Of particular interest is the one which rewards only (or disproportionately more) for the matching of the first natural frequency to its desired target. This is of particular interest because the optimization of the first natural frequency alone is considered state of the art in frequency tuning via external damping, and we would like to show that our approach also can be compared directly to traditional approaches. In this case, the following data comes from a quadratic-scaling fitness function. This is a variation on the fitness function described in Sec. 7.3.3, with the $\frac{n-i+1}{n}$ term squared, such that the primary natural frequency now accounts for 100 times as much weight as the 10th natural frequency does (as opposed to 10 times more in the linear scaling scenario above). In this scenario (not plotted here), the 32 independent runs shows a mean of 0.106% error for the primary natural frequency, with a 95% bootstrapped confidence interval of [0.104%, 0.110%] error. One could also imagine this trend being even more extreme should the entire weight be placed on the primary frequency.




Figure 7.4: Examination of the effect of particular random frequency profile targets on the resulting optimization process. Structures are optimized to match (*red stars*) Random Frequency Profile #1, (*blue squares*) #2, (*green circles*) #3, or (*magenta crosses*) #4. While it is not possible to compare the responses to the same frequencies across treatments here, since no two treatments optimize towards the same frequency, we can compare each Random Profile to a control treatment (without selection pressure), plotted as *black triangles*. In all 10 frequencies for all 4 runs, the optimized structures outperform their control counterparts against the same frequencies (with all p-values $< 1.008 \times 10^{-5}$), suggesting the ability of this technique to effectively evolve structures for various frequency profiles.




Figure 7.5: Examination of the effect of number of materials used on the performance of optimized structures. Both data sets use the same random seeds and target frequency (Random Frequency Profile #1). (*red stars*) Optimized structures from a combination of the maximally stiff or maximally soft voxels only (2 materials). (*blue circles*) Each voxel can take on one of 10 intermediate material properties (including the two endpoints in the previous treatment). A two-sided Mann-Whitney U test lacks significant at the 95% confidence level with p-values ≤ 0.0745 for all frequencies except: $n = 4$ (15.594 MHz) with p-value ≤ 0.029 , and $n = 10$ (105.121 MHz) with p-value ≤ 0.0005 . This suggests that for most frequencies (besides $n = 4, 10$) the addition of intermediate material properties often does not significantly impact the frequency matching ability of the resulting structures.

7.4.3 Vibrationally Optimized Beam Examples

Fig. 7.6 shows a representative sample of interesting beams evolved to match Random Frequency Profile #1. The top end has a fixed boundary condition, where the bottom is free; white and black voxels correspond to the soft and stiff materials, respectively. These beams tend to feature a long vertical primary structures (or two connected shorter beams spanning the length of the beam), often with horizontal protrusions acting as secondary structures.

Fig. 7.7 shows beams which are stereotypical of the structures evolved in response to Random Frequency Profile #2. These structures tend to exhibit greater curvature, exhibiting circular patterns from the center of the beam, which presumably rely heavily on the gradients caused by the polar coordinate inputs to the CPPNs. These structures seem more likely than Fig. 7.6 to include material discontinuities and have multiple floating sections of continuous material instead of a single one than spans the length of the beam.

Figs. 7.8 and 7.9 display structures optimized for the remaining target frequency profiles. In Fig. 7.8, the results of optimization towards Random Frequency Profile #3 display simple vertical structures, often with a repetition or separation point along the mid-line between top and bottom halves of the structures. While in Fig. 7.9, structures evolved for Random Frequency Profile #4 display a distinct pattern of stiff material near the center of the beam, with mostly soft voxels at the top and bottom edges of the structures. These unique structural types, recurring frequently within treatments, but rarely for treatments optimizing towards another random frequency profile, provide further evidence of the algorithm's ability to produce multi-material placements suited for specific vibrational optimization problems.

Alternatively, Fig. 7.10 showcases the beams created from a process without the selection pressure to optimize towards a specific frequency profile. These structures appear more pixelated, featuring isolated single, or small groups of, voxels. These structures tend to display very broad divisions, such as a single fuzzy line splitting stiff and soft sides of a beam. The presence of these structures implies that the CPPN itself complexifying over time is not in itself sufficient to produce the crisp lines, solid patches, and features such as symmetry and repetition found in Figs. 7.6-7.9. But suggests that selection towards specific frequencies, combined with the apparent vibrations benefit of such features for these specific cases, is required to produce these motifs.




Figure 7.6: Six examples of stereotypical structures evolved for Random Frequency Profile #1. Black pixels represent stiff material, while white pixels represent soft material. Optimized structures often consist of one or two large continuous regions. They often span all or nearly all of the design space allocated in both width and length. They also often display minor branching structures off of this primary long-spanning structure.




Figure 7.7: Six examples of stereotypical structures evolved for Random Frequency Profile #2. These structures appear to favor sweeping curvature in their designs, often with repeating or symmetric motifs. This group shows more use of thinly (or non-)connected areas of stiff material than group #1.




Figure 7.8: Typical structures evolved to match the natural frequencies of Random Frequency Profile #3 display very simple structures with long vertical motifs, often displaying symmetry or repetition along the midpoint of the vertical axis, where a thin strip of stiff voxels connects the top and bottom features.




Figure 7.9: The top six performing structures optimized for Random Frequency Profile #4. This frequency profile tends to drive the evolution of structures with stiff voxels concentrated around the center, with near-symmetric "wings" stretching vertically from the primary horizontal structure.




Figure 7.10: Six examples of stereotypical structures from the control group that features a random walk through the genotypic space in lieu of selection pressure towards high performing structures. The resulting structures appear to be much noisier, with many unconnected single pixels. This suggests that the continuous shapes from the previous trials are preserved because of their vibrational advantages.

7.4.4 Fabricated Structures

While these structures are evolved in simulation, the optimized beams are easily produced via additive manufacturing. Fig. 7.11 shows two optimized structures in simulation, and their fabricated counterparts. While the simulated materials in this study were not modeled after any real-world materials, and thus the printed structures were not tested for their actual natural frequencies, future and ongoing work will demonstrate both of these features.




Figure 7.11: Optimized designs printed on an Objet Connex 500. Shown here are voxels of 0.9mm for visualization purposes, though the resolution on this 3D printer would allow voxels as small as 16 microns.

7.5 Discussion

The above describes a system in which vibrational properties of an existing topological design are optimized through the placement of soft and stiff voxels according to an evolved Compositional Pattern Producing Network. The culmination of the results presented previously implies that this new system is indeed capable of producing structures which are optimized to match a predetermined set of natural frequencies. Furthermore, the system is fairly robust to the specific natural frequencies, and their relations to each other – at least within the limited space we explored between range of frequencies existing in single material (all stiff or all soft) beams. We compared the results of these beams to those created from a random evolutionary walk to show the ability of the system to optimize for many natural frequencies simultaneously, since there have not existed previous studies which optimize as many as 10 natural frequencies which we could compare our results to. Existing structural optimization methods such as homogenization allow for the optimization of one or few frequencies, but rely heavily on spatially-local gradients or properties which do not capture the highly coupled relationships across a structure that occur when optimizing for multiple frequencies. The employment of an evolutionary algorithm helps to remove the top-down design from this difficult problem, and the use of a CPPN genome creates correlated global mutations which evidently is helpful in optimizing the spacial coupling which made this problem previously unsolvable. In order to compare to the state of the art single frequency optimization, we show that our method can move an object’s natural frequency to the order of a tenth of a percent error, without having to affect it’s existing topology or augment it with dampers.

Arguably the most interesting of all, are the examples of optimized topologies in Figs. 7.6-7.9. The differences between the strategies in designing these structures, at such a basic level that it is immediately obvious to the eye, display the flexibility and creativity of design which evolutionary algorithms have previously made their acclaim [19]. The ability to produce structures that are inherently and consistently different in their fundamental layout and design strategy, attuned to each new scenario (as opposed to minor variations on a preexisting paradigm) provide the hope that future work with this system will demonstrate the ability to produce effective and realizable solutions specifically designed for the many real-world applications of vibrational analysis.

7.6 Future Work

Given the novelty of this work, the potential for future work is nearly limitless. We are currently upgrading this system to produce 3D structures which include boundary conditions outside of classical beams, running experiments on a larger number of frequencies and on a less restricted set of desired frequency profiles, and exploring the effect of finer resolution on this design paradigm.

It is also not yet clear how the system will approach demands to push natural frequencies away from wide frequency bands, rather than pushing them towards specific targets (though these two problems are highly related). The fitness function employed in this work was chosen logically, though not rigorously, and further experiments will take place to provide evidence of its effectiveness in relation to others. Specifically, this applies to multi-objective fitness functions. One could easily imagine other desirable function properties of a real

world system, such as the minimization of deflection in a fixed-free beam, which needs to be optimized in parallel with the structures vibrational properties.

Given the ability the system's consistent ability to match all 10 frequencies, it is unknown where the limits of this paradigm fall. In future work, we push this system to failure, exploring its limitations on number of frequencies optimized.

7.7 Conclusion

In this work, we approached the tuning of an object's natural frequencies without the use of external damping or changing the shape envelope of the existing object topology. This problem is of vital engineering importance to produce parts, objects, and structures which are robust to the structural weakening and eventual failures caused by vibrations in many domains such as civil or aerospace engineering. The inverse case is also an important problem, where systems that collect energy and drive oscillations of a piezoelectric beam represent a significant advance in the efficiency of energy harvesting. Despite the difficulty of this problem, we show here that we are capable of producing structures which can optimally place their natural frequencies to match one of multiple desired resonant frequency profiles of 10 frequencies. We do so by optimizing the placement of multiple materials within an existing topology with the evolutionary algorithm CPPN-NEAT. The demonstrated ability to optimize many frequencies simultaneously, as well as the fundamental differences in structures optimized across multiple target frequency profiles show promise for this technique to soon be a design automation platform for the vibrational optimization of important real-world structures.

CHAPTER 8
APPLICATIONS OF MORPHOLOGICAL DESIGN AUTOMATION
TO ARTISTIC OPTIMIZATION

Abstract of Chapter ¹

Can we design without being aware of what we want? Design tools today, both manual and automated, require a designer to make explicit design choices. There is, however, a potentially much larger space of interesting designs that is left unexplored, as designers may not be able to deliberately recognize as valuable. Here, we introduce an implicit method that attempt to access this latent space through subconscious interaction. We use an eye tracker to capture user interest, and use this information to breed designs through artificial evolution. We demonstrate a variety of 3D shapes designed without any explicit selection, a process that we call Inadvertent Design. Results show that the process can work effectively, and an exit survey of participants shows that they found this technology to be enjoyable and that it can aid their ability to explore an alternative space and find novel interesting designs, while aiding their creativity. We suggest that this technology could help tap into design talent that may be inaccessible through explicit methods, and help leverage time lost to gaming and entertainment.

Can we design without being aware of what we want? Design tools today, both manual and automated, require a designer to make explicit design choices. There is, however, a potentially much larger space of interesting designs that is

¹To appear as: Cheney, N., Clune, J., Yosinski, J. & Lipson, H. (in review). Inadvertent Design: Accessing the Subconscious Design Space Through Eye Tracking.

left unexplored, as designers may not be able to deliberately recognize as valuable. Here, we introduce an implicit method that attempt to access this latent space through subconscious interaction. We use an eye tracker to capture user interest, and use this information to breed designs through artificial evolution. We demonstrate a variety of 3D shapes designed without any explicit selection, a process that we call Inadvertent Design. Results show that the process can work effectively, and an exit survey of participants shows that they found this technology to be enjoyable and that it can aid their ability to explore an alternative space and find novel interesting designs, while aiding their creativity. We suggest that this technology could help tap into design talent that may be inaccessible through explicit methods, and help leverage time lost to gaming and entertainment.

8.1 Introduction

Can we design without knowing what we want? Most design tools today require a designer to know what they are looking for, before they can create it: to explicitly specify primitive shapes, constraints, goals, or preferences. There may, however, be a much larger space of possible good designs that designers are simply unaware of – a space that is largely inaccessible to current “explicit design methodologies. Here, we demonstrate a method that begins to explore the space of unconscious design preferences, one that is based on the principle “I’ll know it when I see it.” We call this process Inadvertent Design.

This premise may not be as farfetched as it first appears. Advertising and product recommendation rely strongly on this premise, where suggested prod-

ucts routinely make companies enormous sums of money [229, 247, 46, 216]. Yet if the consumer had been consciously aware that they had wanted to buy a product before it was suggested, would they not have simply purchased it on their own? We are all aware of the presence of subconscious reactions in personalized marketing and product suggestions, but why isn't there a complementary method for personalized design?

Perhaps great design is heralded as a stroke of genius, that a great idea suddenly appears in the mind of a brilliant designer with no concept of where it came from or how it got there. This makes creative design an intimidating concept to approach procedurally with artificial intelligence. But this need not be the case. Much of design is actually iterative and formulaic. One such example on a grand scale is the great diversity and effectiveness of life on earth. Perhaps this is why we once assumed that a brilliant creator must have hatched the ideas for each of these creatures. Yet we now know that the life we see around us is the outcome of a very formulaic process: the combination of heritable traits, genetic variations, and natural selection that we know as evolution. We also see examples of this formula routinely used on smaller scales in animal breeding [108, 230, 291].

Optimization algorithms that employ this recipe, deemed "evolutionary computation", are able to autonomously produce solutions to a problem, given that one is able to specify a preference between potential solutions based on how well they solve the problem [12, 13, 217]. In the use of evolutionary computation for artistic design, it is common for the user to be in-the-loop and specify a rating of each potential solution proposed by the algorithm based on its appeal – a technique called Interactive Evolution [72, 166, 262, 284]. Though in the




Figure 8.1: A sketch of an interactive design system via eye tracking. A webcam and gaze tracking software (stand-alone hardware or embedded in the television, smartphone, or computer monitor) measures the user's fixation towards various objects on screen, inferring the user's preference for each one. This information is used to produce a new set of suggested objects (shown here as a object shapes on a black background – as provided to users in this study, but these objects may also be embedded within traditional media or online content) allowing the user to inadvertently design their own custom products from the comfort of their own couch or desk and without any design training and little to no conscious effort.

case of designing for subconscious preferences we describe in the opening, a user may not be able to specify an explicit preference for one form over another. Their preference towards a certain design lives outside of their conscious decision making process. In order to produce design solutions which access these subconscious preferences, one would require a method for the subconscious to indicate a preference for one solution over another. In this work, we propose a method to enable this: the use of eye tracking to capture the user's preferences and drive the process of Inadvertent Design.

Eye tracking for preference inference provides distinct advantages over traditional computer interfaces (keyboard or mouse):

(1) There is less physical effort required. A user does not have to make a physical response (like a mouse click) to provide feedback to the system. All feedback is inferred from simply looking at the screen. The eye movements which are measured occur continuously, and may take place at an entirely subconsciously level in the absence of explicit intent.

(2) Since there is no need to physically press a button to show preference for one design over another, this also means that there is no need for the user to make an explicit decision in ranking the potential solutions – thus decreasing the required mental load. While the user may choose to select a design by intentionally focusing on one of the proposed designs, the system infers all preference relations on its own and does not require the user to formally make a decision (a task that users become worse at as the amount of information or alternatives increases [146, 190]).

(3) Since no physical interaction or mental decision point are required, this

process can take place without the conscious or explicit intent of the user. Simply by looking at a screen with various alternatives, we conjecture that the system should be able to infer their preferences between design alternatives (even if they are within the context of a larger scene – though we do not formally test this supposition here). In this work we do test the design abilities of users without explicit design intentions, simply by sitting in front of (and looking at) a screen. Hence we refer to this process as “Inadvertent Design”.

(4) The type of feedback related by one’s eye movements may be different (and provide a different set of preference information) from the feedback one consciously chooses to provide (as in the case in subconscious priming [155, 153, 193]).

(5) Eye tracking may also have an advantage over other non-haptic interfaces, like EEG control [205, 237], in that it is familiar and unobtrusive. Users of all ages are already used to passively interacting with television screens for hours on end [60, 243, 195, 298].

Based on these advantages, we foresee a great potential for the use of eye tracking in product design. With the growing popularity of front facing cameras/webcams on computers, cell phones, and televisions (and the ability to perform eye tracking with these devices [36, 29, 244]) we see the potential for this technology to be accessible to a large number of users. The commercial/marketing application of such an approach, in which users could design customized products inadvertently from product suggestions embedded within regular shows, games, or commercials – only to be shown these completed designs at a later time – provides additional motivation for future work beyond the scope of this paper (Fig. 8.1).

8.2 Background

An optimization process which mimics the main tenants of natural evolutionary design was proposed as early as 1950 by Alan Turing [293], and became popular with Holland's 1975 work on Genetic Algorithms [130]. Interactive evolution is often traced back to Dawkins' description of "Biomorphs" in his 1986 book, *The Blind Watchmaker* [72]. In 1991, Karl Sims demonstrated interactive evolution for computer graphics, in which a human-in-the-loop algorithm allowed users to be presented with various suggested "artworks", and then provide their preference for one over another [261].

In fact, Sim's work was the first to hint at the possibility for inadvertent design. This work extended the traditional computer interface of evolutionary design to create various museum exhibits (*Genetic Images* starting in 1993, and *Galápagos* starting in 1997). In these exhibits, multiple monitors were used to display the potential images suggested by the algorithm, with step sensors on the ground in front of each screen. An image's quality ("fitness") was judged on the number of times the step sensor in front of it was pressed. If the viewer might inadvertently tend to step on the sensors in front of the images they preferred because they stood in front of such images longer or more frequently, this work could be considered the first to employ a fundamental tenant of inadvertent evolution: that the potential designs are competing with one another over the scarce resource of users' visual attention, and that natural selection weeds out designs which fail to attract or entertain the users. Sims' work did instructed users to decide which images they liked the best and explicitly step on the sensors in front of them (relatively small sensors, which may not have been stepped on accidentally by everyone viewing that image), so this work fails to be truly

“inadvertent” in its ability to work outside of conscious decision making. But we must remark at its ingenuity nonetheless.

More recent work [253] has also explored the interface of these interactive evolutionary systems through the creation of online exhibits, where users can click through a number of images and highlight the ones they wish to select. Such a system performs the selection in a very explicit manner, but the website format allows for the design process to reach anyone with an internet connection. This interactive website framework was also ported to the design and fabrication of 3D shapes[54]. While these works clearly require explicit decisions to be made regarding design preferences, their ability to reach a near endless number of users through the web is a necessity for the wide use and acceptance of an inadvertent design system.

Xu et al. also demonstrated an example of shape evolution for computer graphics [314], though their work differs from the previous examples (and our proposed work) as it relies on pre-segmented shapes and variation operators on these individual segments rather than affecting pattern generators that affect the entire shape in a bottom-up manner.

As one may imagine, these human-in-the-loop evaluations (for hundreds or thousands of iterations) become quite time intensive, causing the user become fatigued and lose attention or quit. This may lead to insufficient amounts of time and low effort decision making being provided to the design process. There have been numerous efforts to relieve this fatigue by predicting user preferences and offloading some of this interactivity to a machine proxy or computational substitute for the human user, often interweaving human evaluations with those of the computer counterpart [215, 152, 112, 187, 136]. Though

these works failed to address user fatigue in terms of effort required to provide each piece of feedback. While we know that staring at a television screen for extended periods of time can lead to fatigue in the forms of eye strain, headaches, dizziness, and nausea [206], this has evidentially not been a serious enough deterrent to prevent consumers from watching hours of television each day [60, 243, 195]

The use of eye tracking relies on the concept of users' selection preferences correlating to their viewing (implied by Sims' museum setups). While gaze may not be a perfect match for user intention or preference, there is reason to suggest that the two are highly correlated. Lohse showed that consumers spent 54% more time viewing yellow pages ads of businesses they ended up choosing, over those they did not [183]. Pieters and Warlop found that consumers make more saccades to their preferred products (i.e. look to them more frequently) and fixate on them for longer [226].

Aside from attention towards preferred objects, it has been demonstrated that novel or unexpected stimuli can also draw increased visual attention [151, 279]. While increased attention towards a novel stimuli might seem as though it would distract from the type of design which the user is ultimately interested in, it has been shown that rewarding novel designs is also a viable means for achieving desired results (and can perform better than simple goal-directed search in some systems) [170]. It should be noted that rewarding both goal-directed and novel designs together (as we conjecture is the case in this system) can be more effective than either search driver alone [172]. This may be especially true in creative optimization domains like product design.

Perhaps a fundamental question underlying the use of eye tracking regards

the existence of subconscious preferences in decision making. It may first be important to simply define our interpretation of the term “subconscious.” We consider conscious behavior to be the single, rational, and purposeful behavior in focus in one’s mind, as opposed to subconscious processes which are often parallel, instinctual, and automatic behaviors not at the forefront of one’s attention. We also note that most complex behaviors can consist of sub-routines with both forms of attention, and may even transition between the two (such as a driver zoned-out and on “autopilot” called into focusing explicitly on their route once they realize they have made a wrong turn). We direct the reader to [76] for an in-depth discussion on “the status of cognition and consciousness in consumer behavior theory”.

A large literature exist regarding behaviors which do not “rely on deliberative mental operations” [91] in consumer choice and preferences. Most notably, studies of subconscious priming effects in psychology demonstrate altered decision making from internal states or cues which users are not explicitly aware of [155, 153, 193]. Similarly, the subconscious emotional reactions (or automatic reactions to low level sensory stimuli) has become an important aspect of industrial and product design [211, 78, 165]. We seek to take these implicit user preferences and employ them to drive automated design.

It should also be noted that a major tenant of this approach is the ability to place these design tools in the hands of everyday users. This would not be possible if each user was required to purchase an eye tracker for each input device (cell phone, computer, television) they owned. Recent work [178, 315] has demonstrated gaze tracking using simple cameras (such as webcams) without the need to shine infared light on the user’s eyes (as is the case with traditional

eye tracking). This is especially important with the prevalence of front facing cameras on smartphones and webcams on computers and televisions.

Despite the promise of such a design interface, we are not aware of a previous attempt to implement an open-ended interactive evolutionary setup for product shape design based on gaze tracking. The general idea for such a system was considered by Pallez et al., though they did not implement a system with an eye tracker capable of capturing inadvertent reactions from the user, but rather “simulated such equipment with the help of a mouse” and explicitly “ask[ed] the user to move the mouse to where he is looking” [214]. By definition, this rules out the possibility of inadvertent design. They also did not suggest applying this interface to a shape design paradigm, but rather a version of the one-max optimization problem [44], where the user must maximize the value of 24 bits. Holmes and Zanker provide an example of an interactive evolutionary setup which did employ an eye tracker [132], however the goal of this system was simply to optimize a single parameter (the length of a fixed-area 2D rectangle) rather than performing complex shape design in an open-ended design possibility space.

Suggestive design interfaces that provide assistance to novice users seeking to create complex designs and employ neither eye tracking nor evolutionary computation have also been described previously. Igarashi and Hughes provide a suggestive interface for 3D drawing which hinted at possible additional components which users may desire when creating a larger scene, but their work allowed for only a fixed number of predetermined geometric relations (20 in their prototype), rather than open-ended complexification of entire shapes [141]. Talton et al. created a modeling tool which would suggest a num-

ber of high quality alternatives to a user based on designs of previous users. However, their algorithm was designed for use on parametric design spaces, and required many other users to provide feedback on the same parametrization system, limiting the novelty of the designs suggested [285]. Chaudhuri and Koltun created a data-driven approach to design suggestion. This approach had the user input a vague query shape, upon which the algorithm would search a large database of existing related shapes to suggest additions to the original query design. This approach allows for a novel combination of existing components, but not the bottom-up design of those basic components and individual features or the generation of entire shapes at once [37].

8.3 Results

To explore the concept of inadvertent design and provide a proof-of-concept demonstration for this method of exploring a new part of the design space, we perform an experimental treatment in which users are asked to sit in front of a computer equipped with an eyetracker and design 3D objects based on a series of instructions (see Methods for details). Users are randomly assigned to either complete the task with a traditional (mouse-based) interface, or our experimental eyetracking user interface. Following their design trials, users are asked a short survey about their experience with their randomly assigned platform.




Figure 8.2: *(top left)* The final iteration of designs for a run in which the user was directed to find a small blue oval via eye tracking inputs. *(top right)* The final iteration of a run in which the user was directed to find a small yellow cone via traditional mouse inputs. *(bottom left)* An example of a user converging to a common shape, despite not originally having it as a target before starting an open-ended design run. Some of the designs in this iteration are reminiscent of a computer monitor. *(bottom right)* An example of a set of evolved objects demonstrating the diversity and variety of unfamiliar shapes that a user might encounter within an open-ended design run – despite not intentionally directing search towards any of them.

8.3.1 Target Shapes

Users and first given three unique target shapes to design, starting from three random initial conditions. In analyzing the shapes designed during these trials, we did not verify the users' self-reported success ratings in reaching the goal shapes. Nor did we provide our own rating of the match/complexity of the shapes they produced. We felt that the users' egocentric evaluations of their designs (reported below) was the most pertinent valuation of their own creations (as, in the case of individualized product design, each consumer would be the end user of the products which he or she designs). We also did not have an objective measure to compare the users' designs to their prompted goals, nor specific instances of target shapes shared by all users (since they were given a verbal description rather than a visual template). We did visually inspect many of the runs, with Fig. 8.2 exemplifying successful matching of the target shapes for each input method.

8.3.2 Open-Ended Design

Each user was then tasked with an open-ended trial in which they were instructed to use the same interface type, but now used to simply create any sort of shape they found interesting, rather than specific target shapes. Given that this experiment was even more free-form, it is difficult to place ratings on the resulting designs. As expected, the user could create a variety of novel shapes – from simple solid objects to complex shapes and abstract mathematical art. Fig. 8.2 demonstrates a variety of novel designs which are not reminiscent of existing shapes, but are novel creations of this particular user. More unexpect-

edly, this free-form design process may stumble upon existing shapes. Fig. 8.2 also shows an example of a user creating the general outline of a computer monitor. The user was not instructed to focus on objects of this shape before the run began, and iterations earlier in the run do not suggest that it was goal from the start, but rather something that was found and exploited later in the design process. One could think of this as a random walk through the design space entering a subset of that space with shapes that were reminiscent enough for the user to envision a computer monitor – at which time the user may have been drawn towards this goal and purposefully switched strategies and attempted to produce this specific object. We also do not discount any potential subconscious priming effects on the user, who performed this process on a computer monitor and sitting in room with many other computer monitors.

8.3.3 Exit Survey Results

At the end of the experiment, each user was also asked general questions about their interaction with the system. They were asked if the system: helped them to produce interesting shapes, helped them produce the goal shapes requested by the prompts, provided them with novel design suggestions, aided their overall creativity, and was generally enjoyable to use.

Fig. 8.3 shows the results of this survey. Values shown here are on a 1-5 Likert scale (with 1 meaning "strongly disagree and 5 corresponding to "strongly agree"). *P*-values denote the difference (via Wilcoxon rank-sum test) between the user reported values and a "neural" response of 3 on the Likert scale. For both set of users, those who were assigned to either the mouse clicking (*m*)




Figure 8.3: The mean values reported on the exit survey for users assigned to either the mouse-clicking or eyetracking treatments (error bars represented 95% bootstrapped confidence intervals). The values given are on a Likert scale (with 5 corresponding to "completely agree" and 1 corresponding to "completely disagree") and are compared to the neutral response (Likert scale value of 3) denoted by the dashed line.

and eyetracking (e) treatments, noted that the system suggested novel design to them ($p_m = 0.004$, $p_e = 0.001$). Users of the eyetracking treatment also reported that the system helped them to find interesting shapes ($p_e = 0.047$), aid their creativity ($p_e = 0.017$), and was enjoyable to use ($p_e = 0.047$). Users who were assigned to the mouse clicking interface, on average, did not report any of these statements to be true ($p_m = 0.631$, $p_m = 0.337$, and $p_m = 0.631$, respec-

tively). Interestingly, neither group found that this interface significantly helped them to reach the specific goals provided in the prompt ($p_m = 0.337$, $p_e = 0.691$), suggesting that the positive benefits of this system may have been largely felt during the open-ended design.

Users were also given a free response section to leave comments about their interactions with the system. Their reviews mirrored the generally positive feedback from the Likert scale rating. Users commented about the eye tracking interface: “It was really neat for the technology to be able to recognize my shape preferences through only my eye sight. I did enjoy using it.”, “ I really liked the experience. The eye technology tools has great potential for design, etc. It was very easy to use.”, “I think it’s amazing how the program detects what I am looking at and then takes it a step further by suggesting interesting shapes. Yes, I liked using it.”, “It was very interesting, but I was not sure if it was capturing my ideas at all times, sometimes it was right on and other times it was a little off, but it prompted suggestions and ideas I had not thought about”, and “ I thought the system was rather intelligent in design. It was enjoyable using it for a short period of time, but the twenty minute task put a strain on my eyes.”. Besides the comment on fatigue, other negative feedback on the eye tracking setup included users who wanted more time at each iterations and were “not ready for the next screen of shapes because I was looking around for the best fit.”. Other users were generally confused by the hidden back-end algorithm of the system, stating that they “don’t know what it’s doing”, or simply that they weren’t able change the shapes as they desired and thus “didn’t really like using the technology.” This is not surprising, as we chose not to explain the concept of an evolutionary algorithm or any of the technical details of the back-end system to the users during the instructional period (as we were concerned that intri-




Figure 8.4: The results of a trial of eye-tracking-based design with the target shape of a face. This trial consisted of three separate runs from random restarts. Note the diversity and abstract qualities/interpretations of the faces produced.

cate knowledge of the underlying algorithms might interfere with the notion of “inadvertent design”).

The mouse-click users were slightly more critical, with just two users providing only positive comments “I think it’s pretty cool! I did like using it.” and “I thought it was interesting. I did like using it and seeing the projected shapes that I otherwise would not have thought of.” Regarding the negative feedback and constructive criticism, one user noted that “Often the shapes were too similar.” Another user “didn’t particularly enjoy using it”, citing the fatigue factor as well as the inability to retain desirable features of some of the designs that had been previously selected. Multiple users noted that the open-ended section was more enjoyable than the target driven trials, saying “it was more fun towards the end” and noted that the creative open-ended section made them “more familiar with how it works to design shapes.”




Figure 8.5: The results of a trial of mouse-click-based design from the same user as Fig. 8.4. This trial consisted of three separate runs from random restarts. Despite the random restarts, note the similarity of results produced when the user explicitly makes the decision of which suggested designs to choose at each iteration.

8.3.4 Finding Faces

Following these user trials, we hypothesized that some forms may be innately more attention-grabbing for the human visual system than others. In particular, we chose to explore the use of faces as the target object. We chose faces for the target shape for not only for their range of possible variation and interpretation, but also because we suspect that potential users (and humans in general) may be subconsciously primed to see faces in abstract contexts [113].

Rather than apply for additional IRB approval for human subject testing, these later trials were performed using one of the authors (N.C.) as the user. As a result, the following results are anecdotal and highly biased. But we still find them to be intriguing, and to inform a potentially interesting study to be executed on a large population of randomly chosen subjects.

First, this user was given three separate runs (from random restarts) using the eye-tracking interface. Fig. 8.4 exemplifies the six most diverse faces found when analyzing the user's saved iterations post hoc. The wide variety of "creative" interpretations of a "face", as well as the diversity of shape, size, and

color within this interpretation should be noted.

Next, the user performed the same task with the mouse clicking interface. The six most diverse faces from the same number of runs (three random restarts) are shown in Fig. 8.5. Note the similarity of the designs, even between runs. In fact, upon first seeing these images, a co-author not involved in the testing procedure initially assumed that these designs were all created from the same run, rather than from three distinct random starting conditions. To dig slightly deeper into this user's ability to show such consistent convergence across runs, we examined the path (and the major "stepping stones" along it) from a random starting point to the resulting designed face (Fig. 8.6). We found that consistently this user found a large solid block early in the design (as the employed evolutionary algorithm begins with simple shapes before "complexifying" them over time). The user then was able to carve out a cavity within the center of the block which would result in holes appearing on one or more faces. The shape, size, and location of these holes were then modified over successive iterations to create the faces seen in Fig. 8.5. One potential hypothesis for the differences in outcomes between these two interface methods is discussed in the next section.

8.4 Discussion

The users' similar evaluations from self-reported feedback regarding the ability to achieve target shapes and interesting free-form designs suggest that that this new eye tracking user interface can achieve results similar to the established methods that require mouse-controlled explicit user feedback (previously shown to create complex designs in 2D shapes [253] and 3D objects [54]).




Figure 8.6: To gain insight into the way the user from Fig. 8.5 was able to produce similar results despite random restarts, we examine the stepping stones (key design stages) from a plain block to a face. All three of the mouse-driven iterations from this trial repeated the same general strategy of producing two holes in a solid block then modifying their size, shape, and position in order to produce a design reminiscent of a face.

The self-reported enjoyment and creative assistance and suggestion of novel and interesting shapes with the eyetracking interface suggest its future potential to be employed for creative object design, specific to each individual user and with potentially less fatigue than traditional interactive design methods. These survey results help to demonstrate the effectiveness of this proof-of-concept implementation. The written user feedback also provides a strong indication that the eye tracking system is well accepted and is at least as user friendly as the traditional physical interface. Given that one of the goals of this interface is to reduce user fatigue, dissatisfaction, and drop-out, we place a premium on such user feedback.

In both the eye tracking and mouse clicking (Fig. 8.2), users exemplified the ability to achieve prespecified targets for simple shapes. However, the exit survey suggested that these targeted trials were not as fruitful as the free-form creative trials. This sentiment is not particularly surprising, as it was noted above that this is a deceptive domain, meaning that goal directed optimization may be frustrating (or impossible for more complex targets) [313]. It is also not

surprising that the users did not achieve significantly different results between the mouse clicking and eye tracking interfaces, during the goal-directed portion of the user trials as this task instructed users to specifically focus on designing certain and very specific types of shapes – and thus acts as new interface for traditional “explicit” design, rather than employing truly inadvertent design.

Search results in deceptive domains may also suggest that open-ended design benefits from the exploration of interesting or novel design spaces, even if they are not moving towards a particular target [313]. While this does not necessarily explain the users’ preference for open-ended design, one may or may not choose to attribute this towards an innate desire for new and exciting designs or simply to a lack of frustration at not struggling to find a target shape. It does, however, suggest the potential for this interactive design (and especially eye tracking design) to create interesting shapes. We believe this is achieved in the eye-tracking-driven designs shown in Fig. 8.2.

Perhaps the most interesting result is that of the face design. It was initially surprising that the sets of images produced by the different methods were so dissimilar – though we can suggest possible explanations for this. Firstly, it’s important to note that this is one anecdotal and bias case, and not necessarily representative of all users – especially those not familiar with the system and its underlying algorithm. We should note that the eye tracking trials took place first and thus the user may have been more familiar with a necessary trajectory required to consistently reach faces upon reaching the mouse click interface (and thus possibly more able to perform their desired selections with more accuracy). It is also possible that a noisy signal from the eye tracker (compared to precise mouse clicks) caused undesired selection – though based on initial

tests (see Methods) we have no reason to believe that selection errors are due to the eyetracking hardware itself. As the mouse click trials took place after the eyetracking trials, additional poor selections from user fatigue may have been a confounding factor – though it is also possible that user fatigue reduced the creativity of the later trials more so than affecting the ability to precisely choose a desired shape.

Our hypothesized explanation for the effect seen in Figs. 8.4 and 8.5 does not contradict these previous suggestions, but complements them by noting that the mouse click interface also attempts to give the user more precise control over the search path which is taken (i.e. which set of objects the user will be shown next). By requiring that the user perform a conscious and explicit decision about which path to take, the user may be more likely to recognize a promising “stepping stone” (intermediate design which they know will put them on a path that leads to a face) from a previous run and make the risk-averse decision to exploit it, rather than searching for a novel way to produce a (likely different-looking) face. In contrast, the eye tracking setup is hypothesized to minimize reward for reoccurring designs (as the user is presumably habituated to repeated shapes [25], gains less information by looking at it again, and thus spends less time doing so). Assuming the user is also risk averse (as most decision makers tend to be in measured certain scenarios [81]), the absence of an explicit decision point also does not force the user to make a actual decision between a known successful path (the “stepping stone” that this user knows leads to a face) and a novel path may or may not lead to something better. Thus we suggest that the successful incentivization of novelty and diversity during search (and thus the potential for success [168, 170, 171]) of the eye tracking approach may come from its potential to employ subconscious/or inadvertent design.

On the actual 3D designs created by this system, we should note that the types of shapes that were created are largely dependent on the particular encoding we selected (CPPN-NEAT). If a different encoding was implemented instead (as was the case in Sims [262] and Dawkins' [72] work), this same interface could potentially create very different looking designs. Considering the idea of this system as a general purpose design framework, which could be employed for fabricated physical products, animated character, 2D artwork, or other potential uses: the underlying eye tracking methodology of this design tool (the primary contribution of this paper) is independent of the underlying design encoding, and thus independent of the particular types of shapes one might want to create with it. Comparing different encodings or performing a meta-optimization of the encoding via this eyetracking interface would be an interesting direction for future work.

Additional future work should: Explore the analysis of the designed shapes (both for their complexity and relationship to target shapes). Combine the eye-tracking user interface with other non-haptic input devices such as EEG headsets to explore the potential for multi-modal effects (e.g. informing whether a fixation on an object is due to a positive or negative reaction). Embed object design and evaluation within larger storylines and distractors (such as media entertainment) to explore the potential to further reduce fatigue and focus more subconscious preferences. Also the further testing of these findings on larger and unbiased sets of subjects, and on mainstream devices (e.g. built-in cameras on smartphones and televisions) would help to better demonstrate the readiness and wide applicability of this technique.

8.5 Conclusion

We have introduced a novel design interface for open-ended interactive shape evolution: eye tracking for inadvertent design. We believe that this paradigm has the potential to unlock a previously unexplored design space consisting of designs conforming to subconscious preferences. We laid out a number of theoretical advantages of such a paradigm, based on our current knowledge of subconscious visual phenomenon and theory within evolutionary optimization. We showed that, in a small pilot study, this new technology does not show any significant drop-off in self-reported achievement or satisfaction, when compared to the traditional interactive evolution interface. We demonstrated examples of successful designs via this new interface for simple target-driven design, free-form design, and a combination of the two (face design). We noted the differences, and potential advantages of the eye tracking paradigm, for an anecdotal instance of this combination task. Finally, we suggested abundant opportunities for future work regarding this new paradigm of inadvertent design.

8.6 Methods

All source code is available at: <https://github.com/ncheney/eyetracking>

8.6.1 Evolutionary Computation

Consistent with [261, 253, 54], we employ an evolutionary algorithm to produce the set of shapes shown to the user at each iteration. An evolutionary algo-




Figure 8.7: A schematic of the genotype to phenotype relationship in CPPN-NEAT, where the network is decoded into a physical object. At each cell within the design space (black wireframe box), the cell's Cartesian coordinates (x, y, z) and distance from center (d) are input into the network to produce the material and color outputs for that cell. The resulting shape is smoothed using the marching cubes algorithm to produce a continuous surface for the object.

rithm is an optimization method which abstracts the main tenants of evolutionary processes in biology: heritable traits, genetic variations, and competition over scarce resources. We choose this search method as it serves as a heuristic method for non-convex problems which are unsuited for many traditional optimization methods [192] – such as the often-deceptive and highly-multi-modal search spaces of shape design.

In our interactive evolutionary setup, the scarce resource which designs are competing for is the visual attention of the user. At each iteration of the design process, a number of suggested designs are shown on screen and the relative

amount of time the user spends fixated on each of the proposed designs is provided as feedback to the evolutionary algorithm. The algorithm uses this information as the reproductive “fitness” of the design, giving shapes who gather the most looking time a disproportionately larger chance of creating a modified version of themselves as one of the designs shown to the user in the next iteration.

The genotype which encodes the heritable traits of the designs (and variation within them) is defined by the evolutionary algorithm CPPN-NEAT [271]. This genetic encoding abstracts the concept of developmental regularity by employing geometric transformations over the gradients of the design space’s axial coordinates. The set of transformations which create a phenotype (design) is represented as a neural network with various geometric activation functions. The network’s functional transformation of the network’s inputs (coordinates of a given cell) to outputs (material properties of that cell), is applied iteratively to each cell within a discretized design area. Through this method, different activations functions create different shape properties over the design space (such as nodes with gaussian activation functions creating symmetry or sinusoidal activation functions creating repetition). This process is best described visually in Fig. 8.7.

8.6.2 User Interface and Eye Tracker

Subjects are seated directly in front of a 20 inch monitor, placed at eye level. An Miramatrix S1 Eye Tracker is placed directly below the monitor. Miramatrix notes the performance limits of this device as accuracy within 0.5-1.0 degrees

of the user's gaze direction, and a data rate of 60 Hz. Users then perform a calibration step in which they track the position of a blue and white dot as it moves to 9 positions across a black background screen.

Users are then briefly shown the camera feed from the eye tracker. This eye tracker operates by shining an infrared light towards the subjects' eyes. This light produces strong reflections of the subjects' pupil, as well as a corneal reflection (which appears as a small, sharp, glint outside of the pupil. The relative position of these two reflections is used to triangulate the gaze direction of the user [228]. The users are shown these camera feeds, and are able to explore the positions and postures which move their eye outside of the camera's capture range. If their eye are ever outside of the camera's range during the experiment, the program will pause until they return to a valid position.

During the experiment, users were presented with a set of written instructions which introduced them to the system and informed them of the goal of their particular trial. The full text of these instructions can be found at <https://goo.gl/3dURq1>. At each iteration within a trial, the user is shown 15 objects (arranged in a 3 by 5 grid) against a full-screen black background. The total amount of time the user spends looking at each grid cell is recorded. As soon as one of the cells accumulates 1.5 seconds of looking time, the iteration ends. At this point, the on-screen design in each cell is assigned the fitness value equal to the total amount of looking time it collected.

In preliminary testing on the authors, objects were highlighted when the computer believed that the user's gaze was directed towards its cell in the 3×5 grid. In this small sample, no instances were found in which the object that the user was looking towards was not the one highlighted. Thus we have no reason

to believe that selection error played a significant role in the results below. In order to minimize distractions, no highlighting took place during experimental trials.

8.6.3 Goal-Directed vs. Open-Ended

Some trials feature “goal” or “target” driven optimization, where the user has a specific object they desire to create. We contrasted this with a more free-form and open-ended type of search where the user doesn’t have a target shape *a priori*. To examine the differences between these types of search, we had each user perform two distinct trials.

In the first trial, users were told to design specific objects assigned randomly from a set of shape descriptors. These descriptors included size (large or small), color (red, blue, or yellow), and shape (oval, cube, or cone). Each user was asked to complete three such runs with three random and unique target shapes, each time starting from a random initial set of starting shapes. The user would end each run by quitting (pressing the “Q” key on the keyboard), and were instructed to do so when they felt they had completed the task and designed their target object.

Following a short break, the user was instructed to use the same interface method they had before, but this time to do so without trying to produce a specific object. Instead they were given instructions to play around with the system and see if they could produce something interesting. This open-ended trial lasted for 20 minutes. Users again had the option to quit a run at any time, but were simply started on a new run from a different initial configuration if

they quit before the time period had expired.

In both of these trials, users were told to save any interesting (or successful) sets of images by pressing the "S" key. Doing so captured a screen shot of the objects in the current iteration. The final set of objects from each run were automatically saved whenever the user quit a run as well. In addition to the images saved in user selected iterations, the "genomes" of all designs were saved, such that any object could be recreated from its initial blueprints, even if the user had not chosen to record it.

Following the design trials, users were given a brief survey asking if the system helped them to design interesting shapes, helped them to reach their target design goals, suggested novel shapes they would not have otherwise thought of, or aided their creativity. They were also asked if they enjoyed using this technology, and if they felt that they had control of what shapes appeared on the screen – as well as basic demographic information.

8.6.4 Control Trials: Mouse Clicking

In an effort to isolate the effects of the eye tracking interface for design, we randomly assigned half of the users to a control condition where they used a physical interface driven by mouse clicks. Users went through the same introductory procedure (including eye tracking calibration) and were given instructions to do the same general tasks, but this group of users were told to click on objects they preferred. They advanced from one iteration to the next by left-clicking on their preferred object (with no 1.5 second time-out). Users were able to share fitness over multiple shapes by selecting more than one object in each iteration.

They did this by right-clicking on as many objects as they desired to highlight a preferred set before left-clicking on one to progress iterations. In this setup, since only two indications of user preference were possible, all selected objects were given maximal fitness values, while all non-selected objects were assigned a minimal fitness value.

8.6.5 Subjects

The testing protocol for human subjects was approved by our institutional review board, as Protocol #1209003270. Participants were recruited by informational fliers. Outside of the preliminary tests performed by the authors, there were a total of 18 participants (10 male, 8 female), who were undergraduate or graduate students in a wide variety of disciplines. Subjects were monetarily compensated for their time. The completion of the testing procedure took half an hour to an hour.

CHAPTER 9

ARGUMENT

9.1 Summary of Evidence

- Chapter 2 serves as an introduction to the problem statement and methods of CPPN-based morphological design for soft robotics. As the robots evolved in this chapter do not optimize “controllers” (simply using open-loop oscillators with fixed amplitude, phase, and frequency), the resulting forms represent behaviors which stem from the shape and form of the robot’s body plan. Thus, this chapter represents an extreme case of morphological computation in our voxel-based soft robot setup, laying the groundwork to explore the gradual incorporation of control complexity into the optimization of an embodied machine. This chapter also provides the experimental setup which leads to the most straightforward optimization of robot morphologies – as the added complexity of co-optimizing controllers for these robots is not present. Since computation is unable to be shared by – or offloaded to – a controller, this chapter also results in arguably the most complex and interesting set of morphologies.
- Chapter 3 emphasizes the abilities specific to soft robots – noting that locomotion over flat ground does not explicitly require a compliant morphology. To demonstrate the benefits of the embodiment of this soft robot platform, the task of squeezing a robot’s body through a small hole is undertaken. The aperture in this example is smaller than the width of the robot, meaning that the robot must bend and fold itself in order to fit through the opening. The production of robots which are able to complete this task

provides one instance of evidence for the importance of designing robots from soft materials, and helps to emphasize the relationship between the need for material complexity in the robot and the challenges apparent in that robot's task and environment.

- Chapter 4 takes the first step towards relaxing the strict assumption of morphology-driven behaviors by allowing the morphology to produce embodied channels for information flow (where “electrical” impulses can propagate sensory information or oscillations from a centralized pacemaker throughout the morphology via physical interfaces between cells). These “neural pathways” create an embodied peripheral nervous system, blurring the boundaries between the optimization of the robot's “controller” and “body plan”.
- Chapter 5 introduces the optimization of “high-level” distributed controllers in addition to morphological design. Here, the frequency of muscle oscillations and the phase offset of each individual voxel are put under evolutionary control. This work shows that when a robot's morphology and controller are co-optimized as a coupled system of two sub-components, not only does controller optimization fail to take place, but the morphological optimization that previously occurred in the absence of controller optimization is also stifled. I hypothesize that the stagnation in co-optimization is due to a specialization between subcomponents that creates a local optima in the search space each time a controller and morphology specialize to one another.
- Chapter 6 further investigates the idea of specialization between morphology and controller in an evolving robot, and proposes an algorithm to address the fragility in this coupled system. The proposed algorithm di-

rectly addresses the problem of specialization by providing a “protected” re-adaptation period for the controller each time the robot’s morphology is mutated. This proposed method results in more effective and sustained optimization of embodied machines, supporting the previous hypothesis of a fragile specialization between robot “bodies” and “brains” and providing a viable solution to the problem of co-optimizing morphologies and controllers in embodied machines.

- Chapter 7 provides an example of the generalization of CPPN/voxel-based shape optimization by applying the same methods used for the evolution of the morphologies of soft robots to the design of micro-structures for multi-material 3D-printed objects. Specifically, we optimize the placement of voxels of two different 3D-printed materials to create patterns which have mechanical properties that differ from those of the original materials – in our case changing the mechanical resonance of an object by optimizing its natural frequencies towards desired frequency values. The extrapolation of the optimization approach from earlier chapters to the design of structures outside of soft robot morphologies helps to suggest the broader applicability of the methods developed here.
- Chapter 8 showcases another example of shape optimization outside of the core application of soft robotics. In this instance, design automation is applied to the production of artistic (or consumer) shapes and objects via interactive evolution. The inclusion of user interaction during the fitness evaluation process allows the system to optimize designs towards subjective evaluation criteria – extending the above work to include domains where the declaration of an objective fitness function is impractical. Furthermore this chapter demonstrates the novel inclusion of an eyetracker as

the user interface for this subjective evaluation – allowing user feedback to be provided in a way which help to minimize fatigue and increase user enjoyment and creativity. As this interface does not require the user to make explicit decisions, but simply infers preference from looking times instead, it may allow this design system to capture and optimize towards subconscious preferences as well.

9.2 Discussion

The above chapters focus around a central theme: the role of morphology in embodied artificial intelligence. More specifically, the role of morphology in creating complex and effective behaviors – as it supplements and compliments higher-level controllers.

This subject is rooted in the physiological study of brains and behavior. The supplementing of high-level control with low-level morphology-driven behaviors is known as morphological computation [222, 220]. Additionally, the complementary nature of the body and brain – as the body helps to shape the organization, learning, and information processing of the brain – is referred to as embodied cognition [3, 49, 221, 219]. Both of these phenomenon play a crucial role in this study of embodied artificial intelligence, and especially in the design of algorithms to address the optimization of behaviors in this setting.

This investigation into morphologically-driven behaviors begins with the extreme of robots with evolved behaviors that are due entirely to their morphologies – as they possess no high-level controllers and each voxel of the robot only has access to inputs from a global open-loop oscillator. In chapter 2, the or-

ganization of the two types of oscillating muscle types, and the soft passive tissues surrounding those muscle cells determine how the robot interacts with its environment (the physical force of gravity and the normal force of the ground) to result in behavior.

The case of purely morphological computation presented here closely mimics the purely embodied behaviors of a passive walking robot [59]. Though in the case presented here, the robot is able to have richer interactions with its environment due to its compliance (such as the folding and squeezing in chapter 3). The use of a generative and scale-free encoding also allows for robots of arbitrary complexity (subject to limitations on computational resources for the physical simulation of these arbitrarily complex robots). Furthermore, these robots are designed in an automated fashion, meaning that careful understanding of the physical interactions at play, and engineering of morphologies which precisely account for these physical interactions are not necessarily – as is the case in hand designed walking robots.

The demonstration of purely morphologically-driven behaviors is an excellent exercise in morphological computation, and provides great insight into the causes of behavior. However, its practical use cases are limited by the fact that morphological changes tend to take place at slower time scales than neural controller changes (if morphological change is present at all in an artificial embodied machine). Thus for reactionary and sensory-dependent behaviors, the incorporation of neural controllers is desirable to make the resulting robots more practical for autonomous interaction with unpredictable environments.

However, the incorporation of spiking neural controllers does not necessitate a dedicated high-level controller. In chapter 4, the line between high-level

neural controllers and low-level morphological behaviors is blurred by allowing information flow (through “electric” conductance and spikes) to propagate through neighboring cells of the physical embodiment of the robot. As these spikes convey information about oscillations from a central pattern generator, and optionally also about touch sensors, this mechanism acts similarly to a neural network controlling sensorimotor behavior of the robot. However, instead of nodes and edges of a dedicated neural network, this information is propagated through the evolved conductive pathways in the muscle and tissues themselves.

The proposal of low-level embodied “neural-esque” processing helps to support the argument of [15], that intelligent and cognitive processes may occur in a distributed fashion throughout agents with no dedicated “brain”. This work also helps to generalize the findings of chapter 2 to include sensory information processing for more dynamic and reactive behaviors without compromising its focus on the morphologically-embedded drivers of behaviors. Biologically, the presence of conductive pathways here shares commonalities with both the peripheral nervous system and the growth of neural pathways in the brain through synaptic plasticity and reorganization [283].

Relaxing the strict assumption of morphologically-driven behavior further, I turn to the investigation of the simultaneous optimization of robot morphologies and high-level controllers. High-level controllers, such as neural networks, are powerful drivers of behavior used for robotic applications today, but they tend to be optimized for fixed and simple morphologies of agents [208, 176, 53, 227, 200, 175]. Their inclusion in the evolution of robots is of great importance to the goal of optimizing autonomous behavior in embodied machines, as the interactions between all three subcomponents of brain-

body-environment systems can lead to cognitive systems and adaptive behaviors [47, 265].

However, it turns out that the optimization of morphologies and controllers is not as straightforward as it might appear (as demonstrated in chapter 5). Rather than sustained optimization of morphologies and controllers, we find that the co-optimization of robotic morphology and control quickly stagnates, falling into local optima early during the evolutionary process.

To help rationalize and explain this premature convergence, I turn to the theory of embodied cognition. Specifically to the idea that the organization and function of the brain is specific to the morphology it is embodied within – and that the organization of the brain is determined in part by the learning that occurs within this morphology.

Applying this theory to the case of morphology and controller co-optimization, it follows that a robot controller will behave differently when placed in a different morphology. Specifically for our interest in optimization, a controller will not be well adapted to a morphology which is has not been optimized within, and thus will not behave optimally. For example a controller optimized for a quadruped morphology is unlikely to continue to produce effective behavior if the sensory information it receives and the motor commands it produces are to/from a hexapod or biped robot morphology. Thus this controller would immediately produce a low fitness value regardless if a mutation from a quadruped body plan to either a biped or hexapod might result in a more promising morphology in the long-run for the task at hand.

One might know that the hexapod or biped were indeed better morpholo-

gies for the task of locomotion in a given environment by optimizing a controller specifically to those morphologies and evaluating its fitness. Thus the solution to the problem of evaluating the long term potential of a proposed morphological variation is as simple as that: readapt controllers for each new morphological change before evaluating their long term potential and performing selection on them. This is done efficiently using the method introduced above for “morphological innovation protection” and results in a temporary decrease in selection pressure on robots with newly mutated morphologies while those robots readapt their controllers to their new morphologies.

The results demonstrated by this method (in chapter 6) show an increase in the sustained evolution of robot morphologies and controllers, resulting in higher fitness values than the traditional approach. This not only introduces a methodology for co-optimizing complex morphologies with high-level neural network controllers, but it also demonstrates the importance of understanding the embodied nature of sensorimotor behavior on designing optimization algorithms. The deep appreciation for the role of embodied cognition in creating behavior led directly to this proposed algorithm. In fact, this proposed method was simply apparent when considering the theory of embodied cognition and noting the dependence of the controller on the morphology of the body. The lack of such a method in the decades since Sims’ pioneering works on *Evolved Virtual Creatures* [263, 264] demonstrates the need for a greater appreciation of the psychological theory behind embodied artificial intelligence.

9.3 Future Work

The study of developmental psychology [225, 96, 14, 97] might further suggest that our reliance on genomes which encode fixed body plans and controllers is an oversimplification that also limits our ability to evolve complex and effective robots. As part of the developmental process, the body takes shape in response to sensory information that is either invariant or situationally dependent [64, 71, 309]. The brain of an agent also learns in response to the organization, movement, and resulting stimulation of the body [116].

The inclusion of both a plastic morphology and controller would help to alleviate some of the need for explicit readaptation, as proposed in “morphological innovation protection”. The specialization and fragility would be less extreme in a coupled system that was capable of continuous co-adaptation of each sub-component to the other through a closed-loop developmental process (sensory and state dependant, with complex feedback loops between the growth and development of the body and brain). However, this level of neurogenesis and morphogenesis is far from the current state of evolutionary robotics and should be explored in future work. Until the implementation a fully adaptive developmental encoding of embodied machines exist, “morphological innovation protection” may serve as a temporary solution to this problem.

One potential solution for the problem of lifetime neural adaptation is the use of reinforcement learning [282]. Deep reinforcement learning employs large scale neural networks to learn sensorimotor policies from previous actions [200, 201, 179, 218]. Many animals, including humans, learn sensorimotor coordination and action from experience – and adjust this learnt behavior as our

bodies and tasks change over time. This biological process may share some similarities with computational reinforcement learning algorithms [282, 2, 74, 133], suggesting that lifetime reinforcement learning of control may pair well with the evolutionary optimization of gross morphologies.

Additionally, development occurs in the morphology of animals as well. This developmental process has been shown to help minimize the phenotypic impact of genetic mutations in evolving robots [24], to allow controllers to more efficiently solve challenging tasks by gradually adapting to changing morphologies [22], and to allow for complex morphologies to arise through interactions of developing cells to each other and to a physical environment [147]. The evolution of development (*evo-devo*) [232, 207, 33, 204, 34, 233] enables evolution to guide the gradual growth and adaptation of morphologies to their environments, a feature which may allow controllers to more closely track the small changes that occur in morphologies over a lifetime. Compared to the larger shocks caused by mutations to static morphologies, gradually developing morphologies may also help to reduce the need for explicit readaptation periods of their associated controllers.

9.4 Additional Readings

In addition to the primary argument made above, extensions and complementary works were also carried out alongside this dissertation – often in collaboration with others. These additional works are not included here, so we also refer the reader to those original works – as they complement and enhance the core argument of this dissertation in the following ways:

- The exploration of the evolution of soft robot morphologies in tasks that are well suited for soft robots in particular (an extension of Chapter 3) is described in [62]. This work investigated the task of swimming behaviors, a task for which we conjecture that soft robots are particularly well suited for because the interaction of their compliant bodies and the viscosity of the water-based environment provide the potential for rich morphological computation. In addition to this comment on the role of environmental complexity in morphological computation, this work helps to demonstrate the wide applicability of the methods described in Chapter 2 to other tasks and environments.
- The notion of morphological computation cited throughout the above chapters is further explored in [63]. In this work, soft robots are optimized yet another task – this time reaching towards a target. Evolution is given just two materials: one that slowly grows larger and one that slowly shrinks. We show that evolved robots composed of a soft material are able to take advantage of rich interactions with gravity, creating a shape which uses passive dynamics to point towards the target. This passive shape is composed homogeneously of growing cells, maximizing the amount of reaching the robot can do as it stretches out towards the target. In contrast to this, robots evolved from stiffer, less compliant, materials are unable to employ passive dynamics and must actively create arrangements of growing and shrinking to act as opposing muscle groups in order to point towards the targets. This additional use of shrinking voxels is presumed to be necessary, as their presence inherently diminishes the reaching ability of the robot. This trend, where softer robots able to use less shrinking voxels and rely more on passive dynamics is repeatedly observed, providing

evidence for a measured increase in the morphological computation of soft robots. While this particular metric is limited in scope, the quantification of morphological computation has been elusive, and this work provides an excellent example of its objective measurement.

- A review and reframing of the work reported in the above further readings is presented in [61]. This work frames the slow growth demonstrated above in the context of development in plants and animals. This work compares robot morphologies evolved across different task environments. It also introduces the use of environment-mediated growth for our evolved soft robots, providing a strong example of the evolution of an adaptive developmental process which is able to readapt the morphology of a robot to perform effectively in multiple different environments. This work represents the first example of a soft robot which changes its body on three unique time-scales: evolution over many generations of a population, development over an individual's lifetime, and the rapid contraction and expansion of muscles in real time.
- The notion of tri-time-scale morphological changes is explored in detail in [163]. This work seeks to explore the limits of the advantages of lifetime development by introducing a minimal developmental model, in which an "open-loop" developmental process grows (or shrinks) from an evolved infant shape to an evolved adult shape without any sensory modulation. We found that the inclusion of even this minimal developmental process proved advantages over an evolutionary process during which only a fixed static shape was evolved. We hypothesize that the "parameter sweep" that takes place as development changes over a variety of sizes may produce a gradient around isolated fitness peaks in the search space –

which are particularly helpful when the search space does not already provide straightforward gradients for evolution to follow (in the case where it is deceptive or may include large fitness plateaus).

- The blurring of the line between morphological computation in the body and neural process in the brain (explored in the electrically conductive and spiking muscles in Chapter 4) is reframed in [42]. This editorial compares the paradigm of soft robot evolution with these spiking muscles to an inverted cellular automata – in which the update rules are held constant and the topology of the substrate is optimized for a given behavior, rather than the reverse for traditional cellular automata.
- The notion of sustained, long-term evolution of morphologies and controllers of autonomous agents (in Chapters 2, 5, and 6) is further explored [269]. This work seeks to create a theory about the minimal criteria necessary for such open-ended evolution. In particular, this work investigates how the strictness of survival for reproduction impacts the sustained evolution of an agent. We find that survival criterion that are too strict discourage diversity and lead to population stagnation, while survival criterion that are too lenient diverge wildly to random behavior. However, at a critical value at this phase transition, we find evidence for sustained open-ended evolution of morphologies and controllers in autonomous agents.

9.5 Closing Remarks

This dissertation examined the role of morphology in the evolution of autonomous behavior in embodied machines. This started from the extreme of

entirely morphology-driven behaviors, and demonstrated the gradual introduction of controllers to these complex morphologies – first with low-level control, then with the full co-optimization of morphology and high-level neural controllers. This work both contributes to, and relies on, the understanding of morphological computation and embodied cognition to explain and optimize the robot behaviors seen above. While the goal of fully seamless integration of autonomous robots into the unpredictable and dynamic environments of the real world still remains a challenge, I hope that this work serves as a contribution that helps to move the automated design of complex embodied machines in the right direction by contributing to our understanding of intelligent and embodied behaviors and their evolutionary design.

BIBLIOGRAPHY

- [1] Ahmed M Abdel-Ghaffar and Ronald F Scott. Analysis of earth dam response to earthquakes. *Journal of the Geotechnical Engineering Division*, 105(12):1379–1404, 1979.
- [2] John Robert Anderson. Learning and memory. 2000.
- [3] Michael L Anderson. Embodied cognition: A field guide. *Artificial intelligence*, 149(1):91–130, 2003.
- [4] Russell W Anderson and Michael Conrad. Hans j. bremermann: A pioneer in mathematical biology. *BioSystems*, 34(1-3):1–10, 1995.
- [5] J. E. Auerbach and J. C. Bongard. Evolving CPPNs to grow three-dimensional physical structures. In *Proc. of the Genetic & Evolutionary Comp. Conf.*, pages 627–634. ACM, 2010.
- [6] J. E. Auerbach and J. C. Bongard. On the relationship between environmental and mechanical complexity in evolved robots. In *Proc. of the Artificial Life Conf.*, pages 309–316, 2012.
- [7] J. E. Auerbach and J. C. Bongard. On the relationship between environmental and morphological complexity in evolved robots. In *Proc. of Genetic & Evol. Comp. Conf.*, pages 521–8. ACM, 2012.
- [8] Joshua Auerbach and Josh C Bongard. How robot morphology and training order affect the learning of multiple behaviors. In *IEEE Congress on Evol. Comp., 2009*, pages 39–46. IEEE, 2009.
- [9] Joshua E Auerbach and Josh C Bongard. Dynamic resolution in the co-evolution of morphology and control. In *Artificial Life XII: Proceedings of*

the Twelfth International Conference on the Synthesis and Simulation of Living Systems, number EPFL-CONF-191277, pages 451–458. MIT Press, 2010.

- [10] Joshua E Auerbach and Josh C Bongard. Environmental influence on the evolution of morphological complexity in machines. *PLoS computational biology*, 10(1):e1003399, 2014.
- [11] Thomas Back. Selective pressure in evolutionary algorithms: A characterization of selection mechanisms. In *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*, pages 57–62. IEEE, 1994.
- [12] Thomas Back, David B Fogel, and Zbigniew Michalewicz. *Handbook of evolutionary computation*. IOP Publishing Ltd., 1997.
- [13] Thomas Bäck, Ulrich Hammel, and Hans-Paul Schwefel. Evolutionary computation: Comments on the history and current state. *Evolutionary computation, IEEE Transactions on*, 1(1):3–17, 1997.
- [14] Paul B Baltes. Theoretical propositions of life-span developmental psychology: On the dynamics between growth and decline. *Developmental psychology*, 23(5):611, 1987.
- [15] Frantisek Baluska and Michael Levin. On having no head: Cognition throughout biological systems. *Frontiers in Psychology*, 7:902, 2016.
- [16] Nils Aall Barricelli. Symbiogenetic evolution processes realized by artificial methods. *Methodos*, 9(35-36):143–182, 1957.
- [17] Nicholas W Bartlett, Michael T Tolley, Johannes TB Overvelde, James C Weaver, Bobak Mosadegh, Katia Bertoldi, George M Whitesides, and

- Robert J Wood. A 3d-printed, functionally graded soft robot powered by combustion. *Science*, 349(6244):161–165, 2015.
- [18] Dominik Belter, Andrzej Kasinski, and Piotr Skrzypczynski. Evolving feasible gaits for a hexapod robot by reducing the space of possible solutions. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 2673–2678. IEEE, 2008.
- [19] Peter John Bentley and David W Corne. *Creative evolutionary systems*. Morgan Kaufmann, 2002.
- [20] Hans-Georg Beyer and Hans-Paul Schwefel. Evolution strategies—a comprehensive introduction. *Natural computing*, 1(1):3–52, 2002.
- [21] J. Bongard. Evolving modular genetic regulatory networks. In *Evolutionary Computation, 2002. Proc. of the 2002 Congress on*, volume 2, pages 1872–1877. IEEE, 2002.
- [22] J. Bongard. Morphological change in machines accelerates the evolution of robust behavior. *Proc. of the National Academy of Sciences*, 108(4):1234–1239, 2011.
- [23] Josh C Bongard. Evolutionary robotics. *Comm. of the ACM*, 56(8):74–83, 2013.
- [24] Josh C Bongard and Rolf Pfeifer. Evolving complete agents using artificial ontogeny. In *Morpho-functional Machines: The new species*, pages 237–258. Springer Japan, 2003.
- [25] Mark E Bouton. *Learning and behavior: A contemporary synthesis*. Sinauer Associates, 2007.

- [26] Rodney A Brooks. Elephants don't play chess. *Robotics and autonomous systems*, 6(1-2):3–15, 1990.
- [27] Rodney A Brooks. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.
- [28] HILARY F Brown. Electrophysiology of the sinoatrial node. *Physiological Reviews*, 62(2):505–530, 1982.
- [29] Andreas Bulling and Hans Gellersen. Toward mobile eye-based human-computer interaction. *Pervasive Computing, IEEE*, 9(4):8–12, 2010.
- [30] Mario Bunge. The mind-body problem. In *Matter and Mind*, pages 143–157. Springer, 2010.
- [31] Marcello Calisti, Michele Giorelli, Guy Levy, Barbara Mazzolai, B Hochner, Cecilia Laschi, and Paolo Dario. An octopus-bioinspired solution to movement and manipulation for soft robots. *Bioinspiration & biomimetics*, 6(3):036002, 2011.
- [32] Sadi Carnot. Reflections on the motive power of fire, and on machines fitted to develop that power. *Paris: Bachelier*, 1824.
- [33] Sean B Carroll. *Endless forms most beautiful: The new science of evo devo and the making of the animal kingdom*. Number 54. WW Norton & Company, 2005.
- [34] Sean B Carroll. Evo-devo and an expanding evolutionary synthesis: a genetic theory of morphological evolution. *Cell*, 134(1):25–36, 2008.
- [35] Vinod R Challa, MG Prasad, Yong Shi, and Frank T Fisher. A vibration en-

- ergy harvesting device with bidirectional resonance frequency tunability. *Smart Materials and Structures*, 17(1):015035, 2008.
- [36] Michael Chau and Margrit Betke. Real time eye tracking and blink detection with usb cameras. *Boston University Computer Science*, 2215(2005-2012):1–10, 2005.
- [37] Siddhartha Chaudhuri and Vladlen Koltun. Data-driven suggestions for creativity support in 3d modeling. In *ACM Transactions on Graphics (TOG)*, volume 29, page 183. ACM, 2010.
- [38] Nicholas Cheney, Jeff Clune, and Hod 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.
- [39] Nick Cheney, Josh Bongard, and Hod Lipson. Evolving soft robots in tight spaces. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 935–942. ACM, 2015.
- [40] Nick Cheney, Josh Bongard, Vytas SunSpiral, and Hod Lipson. On the difficulty of co-optimizing morphology and control in evolved virtual creatures. In *ALIFE XV, The Fifteenth International Conference on the Synthesis and Simulation of Living Systems*, volume 15, pages 226–233. MIT Press.
- [41] Nick Cheney, Jeff Clune, Jason Yosinski, and Hod Lipson. Hands-free evolution of 3d-printable objects via eye tracking. *arXiv preprint arXiv:1304.4889*, 2013.
- [42] Nick Cheney and Hod Lipson. Topological evolution for embodied cellular automata. *Theoretical Computer Science*, 633:19–27, 2016.

- [43] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proc. of the 15th Genetic and Evol. Comp. Conf.*, pages 167–174. ACM, 2013.
- [44] Chihyung Derrick Cheng and Alexander Kosorukoff. Interactive one-max problem allows to compare the performance of interactive and human-based genetic algorithms. In *Genetic and Evolutionary Computation–GECCO 2004*, pages 983–993. Springer, 2004.
- [45] Joel Chestnutt, Manfred Lau, German Cheung, James Kuffner, Jessica Hodgins, and Takeo Kanade. Footstep planning for the honda asimo humanoid. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 629–634. IEEE, 2005.
- [46] Kwok-Wai Cheung, James T Kwok, Martin H Law, and Kwok-Ching Tsui. Mining customer product ratings for personalized marketing. *Decision Support Systems*, 35(2):231–243, 2003.
- [47] Hillel J Chiel and Randall D Beer. The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in neurosciences*, 20(12):553–557, 1997.
- [48] Kyu-Jin Cho, Je-Sung Koh, Sangwoo Kim, Won-Shik Chu, Yongtaek Hong, and Sung-Hoon Ahn. Review of manufacturing processes for soft biomimetic robots. *International Journal of Precision Engineering and Manufacturing*, 10(3):171–181, 2009.
- [49] Ron Chrisley. Embodied artificial intelligence. *Artificial Intelligence*, 149(1):131–150, 2003.

- [50] B Clancy, RB Darlington, and BL Finlay. Translating developmental time across mammalian species. *Neuroscience*, 105(1):7–17, 2001.
- [51] Miguel Clavero and Emili Garcia Berthou. Invasive species are a leading cause of animal extinctions. *TRENDS in Ecology and Evolution*, 20(3):110–110, 2005.
- [52] Dave Cliff and Geoffrey F Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *European Conference on Artificial Life*, pages 200–218. Springer, 1995.
- [53] J. Clune, B.E. Beckmann, C. Ofria, and R.T. Pennock. Evolving coordinated quadruped gaits with the HyperNEAT generative encoding. In *Proc. of the IEEE Congress on Evol. Comp.*, pages 2764–71, 2009.
- [54] J. Clune and H. Lipson. Evolving three-dimensional objects with a generative encoding inspired by developmental biology. In *Proc. of the European Conf. on Artificial Life*, pages 144–148, 2011.
- [55] J. Clune, J-B. Mouret, and H. Lipson. The evolutionary origins of modularity. *Proc. of the Royal Society B*, 280(20122863), 2013.
- [56] J. Clune, K. O. Stanley, R. T. Pennock, and C. Ofria. On the performance of indirect encoding across the continuum of regularity. *IEEE Trans. on Evol. Comp.*, 15(4):346–67, 2011.
- [57] Jeff Clune, Dusan Misevic, Charles Ofria, Richard E Lenski, Santiago F Elena, and Rafael Sanjuán. Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS computational biology*, 4(9):e1000187, 2008.

- [58] Steve Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085, 2005.
- [59] Steven H Collins, Martijn Wisse, and Andy Ruina. A three-dimensional passive-dynamic walking robot with two legs and knees. *The International Journal of Robotics Research*, 20(7):607–615, 2001.
- [60] George Comstock et al. *Television and human behavior*. ERIC, 1978.
- [61] Francesco Corucci, Nick Cheney, Sam Kriegman, Cecilia Laschi, and Josh C. Bongard. Evolutionary developmental soft robotics as a framework to study plant and animal intelligence. *submitted to Frontiers in Robotics and AI*, 2017.
- [62] Francesco Corucci, Nick Cheney, Hod Lipson, Cecilia Laschi, and Josh Bongard. Evolving swimming soft-bodied creatures. In *ALIFE XV, The Fifteenth International Conference on the Synthesis and Simulation of Living Systems, Late Breaking Proceedings*, page 6, 2016.
- [63] Francesco Corucci, Nick Cheney, Hod Lipson, Cecilia Laschi, and Josh Bongard. Material properties affect evolutions ability to exploit morphological computation in growing soft-bodied creatures. In *ALIFE XV, The Fifteenth International Conference on the Synthesis and Simulation of Living Systems, Late Breaking Proceedings*, pages 234–241, 2016.
- [64] Susan J Crawford-Young. Effects of microgravity on cell cytoskeleton and embryogenesis. *International journal of developmental biology*, 50(2-3):183–191, 2003.

- [65] Fred EC Culick. Acoustic oscillations in solid propellant rocket chambers. *Astronautica Acta*, 12(2):113–126, 1966.
- [66] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.
- [67] Sylvain Cussat-Blanc, Hervé Luga, and Yves Duthen. From single cell to simple creature morphology and metabolism. In *ALIFE*, pages 134–141, 2008.
- [68] Sylvain Cussat-Blanc, Jonathan Pascalie, Sébastien Mazac, Hervé Luga, and Yves Duthen. A synthesis of the cell2organ developmental model. In *Morphogenetic Engineering*, pages 353–381. Springer, 2012.
- [69] Charles Darwin. *The origin of species*. Lulu. com, 1872.
- [70] Charles Darwin and William F Bynum. The origin of species by means of natural selection: or, the preservation of favored races in the struggle for life. 2009.
- [71] Henry Gassett Davis. *Conservative surgery*. Appleton, 1867.
- [72] Richard Dawkins. *The blind watchmaker: Why the evidence of evolution reveals a universe without design*. WW Norton & Company, 1986.
- [73] Richard Dawkins. *The selfish gene*. Oxford university press, 2016.
- [74] Peter Dayan and Bernard W Balleine. Reward, motivation, and reinforcement learning. *Neuron*, 36(2):285–298, 2002.
- [75] André de Ávila Borges, Leticia Miguel, and Leandro Miguel. Size and shape optimization of structures by harmony search. 2013.

- [76] C Derbaix and Pierre Vanden Abeele. Consumer inferences and consumer preferences. the status of cognition and consciousness in consumer behavior theory. *International Journal of Research in Marketing*, 2(3):157–174, 1985.
- [77] René Descartes and John Cottingham. *René Descartes: Meditations on first philosophy: With selections from the objections and replies*. Cambridge University Press, 2013.
- [78] PMA Desmet and P Hekkert. Special issue editorial: Design & emotion. *International Journal of Design*, 3(2):1–6, 2009.
- [79] Kay Dickersin. The existence of publication bias and risk factors for its occurrence. *Jama*, 263(10):1385–1389, 1990.
- [80] Ulf Dieckmann, Paul Marrow, and Richard Law. Evolutionary cycling in predator-prey interactions: population dynamics and the red queen. *Journal of theoretical biology*, 176(1):91–102, 1995.
- [81] Thomas J Dohmen, Armin Falk, David Huffman, and Uwe Sunde. *Are risk aversion and impatience related to cognitive ability?* CESifo Working Paper Series, 2009.
- [82] René Doursat, Carlos Sánchez, Razvan Dordea, David Fourquet, and Taras Kowaliw. Embryomorphic engineering: emergent innovation through evolutionary development. In *Morphogenetic Engineering*, pages 275–311. Springer, 2012.
- [83] René Doursat, Hiroki Sayama, and Olivier Michel. *Morphogenetic engineering: toward programmable complex systems*. Springer, 2012.
- [84] René Doursat, Hiroki Sayama, and Olivier Michel. A review of morphogenetic engineering. *Natural Computing*, 12(4):517–535, 2013.

- [85] Charles H Dowding. *Blast vibration monitoring and control*, volume 297. Prentice-Hall Englewood Cliffs, 1985.
- [86] Jianbin Du and Niels Olhoff. Minimization of sound radiation from vibrating bi-material structures using topology optimization. *Structural and Multidisciplinary Optimization*, 33(4-5):305–321, 2007.
- [87] Maria B Dühring, Jakob S Jensen, and Ole Sigmund. Acoustic design by topology optimization. *Journal of sound and vibration*, 317(3):557–575, 2008.
- [88] Bradley Efron and Robert Tibshirani. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical science*, pages 54–75, 1986.
- [89] Santiago F Elena and Richard E Lenski. Evolution experiments with microorganisms: the dynamics and genetic bases of adaptation. *Nature Reviews Genetics*, 4(6):457–469, 2003.
- [90] Jeffrey L Elman. *Rethinking innateness: A connectionist perspective on development*, volume 10. MIT press, 1998.
- [91] Hans-Peter Erb, Antoine Bioy, and Denis J Hilton. Choice preferences without inferences: Subconscious priming of risk attitudes. *Journal of Behavioral Decision Making*, 15(3):251–262, 2002.
- [92] Daniele Fanelli. Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90(3):891–904, 2011.
- [93] Flavio H Fenton, Elizabeth M Cherry, Alain Karma, and Wouter-Jan Rappel. Modeling wave propagation in realistic heart geometries using the phase-field method. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 15(1):013502, 2005.

- [94] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017.
- [95] Barbara L Finlay and Richard B Darlington. Linked regularities in the development and evolution of mammalian brains. *Science*, 268(5217):1578, 1995.
- [96] John H Flavell. The developmental psychology of jean piaget. 1963.
- [97] John H Flavell. Cognitive development: Past, present, and future. *Developmental psychology*, 28(6):998, 1992.
- [98] Dario Floreano and Claudio Mattiussi. *Bio-inspired artificial intelligence: theories, methods, and technologies*. MIT Press, 2008.
- [99] Dario Floreano and Francesco Mondada. Automatic creation of an autonomous agent: Genetic evolution of a neural network driven robot. In *Proceedings of the third international conference on Simulation of adaptive behavior: From Animals to Animats 3*, number LIS-CONF-1994-003, pages 421–430. MIT Press, 1994.
- [100] Jerry A Fodor. *The modularity of mind: An essay on faculty psychology*. MIT press, 1983.
- [101] David B Fogel. *Evolutionary computation: the fossil record*. Wiley-IEEE Press, 1998.
- [102] Alex S Fraser. Simulation of genetic systems by automatic digital computers i. introduction. *Australian Journal of Biological Sciences*, 10(4):484–491, 1957.

- [103] Pablo Funes and Jordan Pollack. Evolutionary body building: Adaptive physical designs for robots. *Artificial Life*, 4(4):337–357, 1998.
- [104] J. Gauci and K. O. Stanley. Generating large-scale neural networks through discovering geometric regularities. In *Proc. of the Genetic & Evolutionary Computation Conf.*, pages 997–1004. ACM, 2007.
- [105] Thomas Geijtenbeek and Nicolas Pronost. Interactive character animation using simulated physics: A state-of-the-art review. In *Computer Graphics Forum*, volume 31, pages 2492–2515. Wiley Online Library, 2012.
- [106] Thomas Geijtenbeek, Michiel van de Panne, and A Frank van der Stappen. Flexible muscle-based locomotion for bipedal creatures. *ACM Transactions on Graphics (TOG)*, 32(6):206, 2013.
- [107] V Scott Gordon and Darrell Whitley. Serial and parallel genetic algorithms as function optimizers. In *ICGA*, pages 177–183, 1993.
- [108] YOSHI GOTO, Masaei Kakizaki, and Naoyoshi Masaki. Production of spontaneous diabetic rats by repetition of selective breeding. *The Tohoku journal of experimental medicine*, 119(1):85–90, 1976.
- [109] Stephen Jay Gould. *Ontogeny and phylogeny*. Harvard University Press, 1977.
- [110] Stephen Jay Gould. *The structure of evolutionary theory*. Harvard University Press, 2002.
- [111] Andrzej Grzesiak, Ralf Becker, and Alexander Verl. The bionic handling assistant: a success story of additive manufacturing. *Assembly Automation*, 31(4):329–333, 2011.

- [112] Zhenyu Gu, Ming Xi Tang, and John Hamilton Frazer. Capturing aesthetic intention during interactive evolution. *Computer-Aided Design*, 38(3):224–237, 2006.
- [113] Stewart Guthrie. *Faces in the Clouds*. Oxford University Press, 1993.
- [114] Zdenek Hadas, Jiri Kurfurst, Cestmir Ondrusek, and Vladislav Singule. Ai based optimization for vibration energy harvesting apps. *Microsystem technologies*, 18(7-8):1003–14, 2012.
- [115] Crispin Hales and Shayne Gooch. *Managing Eng Design*. Springer, 04.
- [116] Carla Hannaford. *Smart moves: Why learning is not all in your head*. ERIC, 1995.
- [117] MH Hansen. Improved modal dynamic wind turbines avoid stall-induced vibrations. *Wind Energy*, 6(2):179–95, 2003.
- [118] Inman Harvey, Phil Husbands, Dave Cliff, Adrian Thompson, and Nick Jakobi. Evolutionary robotics: the sussex approach. *Robotics and autonomous systems*, 20(2-4):205–224, 1997.
- [119] Helmut Hauser, Auke J Ijspeert, Rudolf M Fuchslin, Rolf Pfeifer, and Wolfgang Maass. Towards a theoretical foundation for morphological computation with compliant bodies. *Biological cybernetics*, 105(5-6):355–370, 2011.
- [120] Donald Olding Hebb. *The organization of behavior: A neuropsychological theory*. Psychology Press, 2005.
- [121] Horst Hendriks-Jansen. *Catching ourselves in the act: Situated activity, interactive emergence, evolution, and human thought*. MIT Press, 1996.

- [122] J. D. Hiller and H. Lipson. Multi-material topological optimization of structures and mechanisms. In *Proc. of the Genetic & Evolutionary Computation Conf.*, pages 1521–1528. ACM, 2009.
- [123] J. D. Hiller and H. Lipson. Evolving amorphous robots. *Artificial Life XII*, pages 717–724, 2010.
- [124] J. D. Hiller and H. Lipson. Automatic design and manufacture of soft robots. *IEEE Trans. on Rob.*, 28(2):457–466, 2012.
- [125] J. D. Hiller and H. Lipson. Dynamic simulation of soft heterogeneous objects. *ArXiv:1212.2845*, 2012.
- [126] Jonathan Hiller and Hod Lipson. Tunable digital material properties for 3d voxel printers. *Rapid Prototyping Journal*, 16(4):241–7, 2010.
- [127] Jonathan Hiller and Hod Lipson. Dynamic simulation of soft multimaterial 3d-printed objects. *Soft Robotics*, 1(1):88–101, 2014.
- [128] Shigeo Hirose and Yoji Umetani. The development of soft gripper for the versatile robot hand. *Mechanism and Machine Theory*, 13(3):351–359, 1978.
- [129] Brian Francis Hoffman, Paul Frederic Cranefield, and Franklin D Johnston. *Electrophysiology of the heart*. 1960.
- [130] John H Holland. *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. U Michigan Press, 1975.
- [131] John H Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press, 1992.

- [132] Tim Holmes and Johannes Zanker. Eye on the prize: using overt visual attention to drive fitness for interactive evolutionary computation. In *Proc. of the 10th annual conference on Genetic and Evolutionary Computation*, pages 1531–1538. ACM, 2008.
- [133] Clay B Holroyd and Michael GH Coles. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychological review*, 109(4):679, 2002.
- [134] G. S. Hornby, H. Lipson, and J. B. Pollack. Generative representations for the automated design of modular physical robots. *IEEE Trans. on Robotics and Automation*, 19(4):703–719, 2003.
- [135] Gregory S Hornby. 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, 2006.
- [136] Gregory S Hornby and Josh C Bongard. Accelerating human-computer collaborative search through learning comparative and predictive user models. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*, pages 225–232. ACM, 2012.
- [137] Gregory S Hornby, Hod Lipson, and Jordan B Pollack. 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, 2001.
- [138] Gregory S Hornby and Jordan B Pollack. Body-brain co-evolution using l-systems as a generative encoding. In *Proceedings of the 3rd Annual Con-*

- ference on Genetic and Evolutionary Computation*, pages 868–875. Morgan Kaufmann Publishers Inc., 2001.
- [139] Gregory S Hornby and Jordan B Pollack. Evolving L-systems to generate virtual creatures. *Computers&Graphics*, 25(6):1041–8, 2001.
- [140] Carl Q Howard, Colin H Hansen, and Anthony C Zander. Noise reduction of a rocket payload fairing using tuned vibration absorbers with translational and rotational dofs. In *Proc. Australian Acoustical Society*, pages 165–71, 2005.
- [141] Takeo Igarashi and John F Hughes. A suggestive interface for 3d drawing. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 173–181. ACM, 2001.
- [142] A J Ijspeert, A Crespi, D Ryczko, and J M Cabelguen. From swimming to walking with a salamander robot driven by a spinal cord model. *Science*, 315(5817):1416–1420, 2007.
- [143] Auke Jan Ijspeert, Alessandro Crespi, Dimitri Ryczko, and Jean-Marie Cabelguen. From swimming to walking with a salamander robot driven by a spinal cord model. *science*, 315(5817):1416–1420, 2007.
- [144] Daniel J Inman. Engineering vibrations. *Prentice-Hall*, 2010.
- [145] Keith R Jackson, Lavanya Ramakrishnan, Krishna Muriki, Shane Canon, Shreyas Cholia, John Shalf, Harvey J Wasserman, and Nicholas J Wright. Performance analysis of high performance computing applications on the amazon web services cloud. In *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*, pages 159–168. IEEE, 2010.

- [146] Jacob Jacoby, Donald E Speller, and Carol A Kohn. Brand choice behavior as a function of information load. *Journal of Marketing Research*, pages 63–69, 1974.
- [147] M. Joachimczak and B. Wróbel. Co-evolution of morphology and control of soft-bodied multicellular animats. In *Proc. of the Genetic & Evolutionary Computation Conf.*, pages 561–568. ACM, 2012.
- [148] Michał Joachimczak, Reiji Suzuki, and Takaya Arita. Artificial metamorphosis: evolutionary design of transforming, soft-bodied robots. *Artificial life*, 2016.
- [149] Michal Joachimczak, Reiji Suzuki, and Takaya Arita. Artificial metamorphosis: Evolutionary design of transforming, soft-bodied robots. *Artificial Life*, in press: 2016.
- [150] Michal Joachimczak and Borys Wróbel. Evo-devo in silico-a model of a gene network regulating multicellular development in 3d space with artificial physics. In *ALIFE*, pages 297–304, 2008.
- [151] Julian S Joseph and Lance M Optican. Involuntary attentional shifts due to orientation differences. *Perception & Psychophysics*, 58(5):651–665, 1996.
- [152] Raffi Kamalian, Ying Zhang, Hideyuki Takagi, and Alice M Agogino. Reduced human fatigue interactive evolutionary computation for micromachine design. In *Machine Learning and Cybernetics, 2005. Proceedings of 2005 International Conference on*, volume 9, pages 5666–5671. IEEE, 2005.
- [153] Johan C Karremans, Wolfgang Stroebe, and Jasper Claus. Beyond vicarys fantasies: The impact of subliminal priming and brand choice. *Journal of Experimental Social Psychology*, 42(6):792–798, 2006.

- [154] Stuart A. Kauffman. *The origins of order: Self organization and selection in evolution*. Oxford University Press, 1993.
- [155] Aaron C Kay, S Christian Wheeler, John A Bargh, and Lee Ross. Material priming: The influence of mundane physical objects on situational construal and competitive behavioral choice. *Organizational Behavior and Human Decision Processes*, 95(1):83–96, 2004.
- [156] Nattavut Keerativuttitumrong, Nachol Chaiyaratana, and Vara Varavithya. Multi-objective co-operative co-evolutionary genetic algorithm. In *International Conference on Parallel Problem Solving from Nature*, pages 288–297. Springer, 2002.
- [157] Sangbae Kim, Cecilia Laschi, and Barry Trimmer. Soft robotics: a bioinspired evolution in robotics. *Trends in biotechnology*, 31(5):287–294, 2013.
- [158] Sangbae Kim, Matthew Spenko, Salomon Trujillo, Barrett Heyneman, Daniel Santos, and Mark R Cutkosky. Smooth vertical surface climbing with directional adhesion. *IEEE Transactions on robotics*, 24(1):65–74, 2008.
- [159] Charles B Kimmel, William W Ballard, Seth R Kimmel, Bonnie Ullmann, and Thomas F Schilling. Stages of embryonic development of the zebrafish. *Developmental dynamics*, 203(3):253–310, 1995.
- [160] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, page 201611835, 2017.

- [161] Maciej Komosinski and Adam Rotaru-Varga. Comparison of different genotype encodings for simulated three dimensional agents. *Artificial Life*, 7(4):395–418, 2001.
- [162] John R Koza. Genetic evolution and co-evolution of computer programs. *Artificial life II*, 10:603–629, 1991.
- [163] Sam Kriegman, Nick Cheney, Francesco Corucci, and Josh C. Bongard. A minimal developmental model can increase evolvability in soft robots. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 2017.
- [164] David Lack. *Darwin's finches*. CUP Archive, 1947.
- [165] Donald A Laird. How the consumer estimates quality by subconscious sensory impressions. *Journal of Applied psychology*, 16(3):241, 1932.
- [166] Christopher G Langton et al. *Artificial life*. Addison-Wesley Publishing Company Redwood City, CA, 1989.
- [167] Cecilia Laschi, Matteo Cianchetti, Barbara Mazzolai, Laura Margheri, Maurizio Follador, and Paolo Dario. Soft robot arm inspired by the octopus. *Advanced Robotics*, 26(7):709–727, 2012.
- [168] Marco Laumanns, Lothar Thiele, Kalyanmoy Deb, and Eckart Zitzler. Combining convergence and diversity in evolutionary multiobjective optimization. *Evolutionary computation*, 10(3):263–282, 2002.
- [169] S. Lee, J. Yosinski, K. Glette, H. Lipson, and J. Clune. Evolving gaits for physical robots with the HyperNEAT generative encoding: the benefits of simulation., 2013.

- [170] Joel Lehman and Kenneth O Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pages 329–336, 2008.
- [171] Joel Lehman and Kenneth O Stanley. Evolving a diversity of virtual creatures through novelty search and local comp. In *Proc. of 13th Genetic & Evol. Comp. Conf.*, pages 211–8. ACM, 2011.
- [172] Joel Lehman, Kenneth O Stanley, and Risto Miikkulainen. Effective diversity maintenance in deceptive domains. In *Proceeding of the fifteenth annual conference on Genetic and evolutionary computation conference*, pages 215–222. ACM, 2013.
- [173] Richard E Lenski, Charles Ofria, Travis C Collier, and Christoph Adami. Genome complexity, robustness and genetic interactions in digital organisms. *Nature*, 400(6745):661–664, 1999.
- [174] Richard E Lenski, Charles Ofria, Robert T Pennock, and Christoph Adami. The evolutionary origin of complex features. *Nature*, 423(6936):139–144, 2003.
- [175] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40, 2016.
- [176] FW Lewis, Suresh Jagannathan, and A Yesildirak. *Neural network control of robot manipulators and non-linear systems*. CRC Press, 1998.
- [177] M Anthony Lewis, Andrew H Fagg, and Alan Solidum. Genetic programming approach to the construction of a neural network for control of a walking robot. In *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, pages 2618–2623. IEEE, 1992.

- [178] Dongheng Li and Derrick Parkhurst. Open-source software for real-time visible-spectrum eye tracking. In *Proceedings of the COGAIN Conference*, volume 17, 2006.
- [179] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [180] Hod Lipson and Melba Kurman. *Fabricated: The new world of 3D printing*. John Wiley & Sons, 2013.
- [181] Hod Lipson and Jordan B Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974–978, 2000.
- [182] Jason D Lohn, Derek S Linden, Gregory S Hornby, and William F Kraus. Evolutionary design of an x-band antenna for nasa’s space technology 5 mission. In *Antennas and Propagation Society International Symposium, 2004. IEEE*, volume 3, pages 2313–2316. IEEE, 2004.
- [183] Gerald L Lohse. Consumer eye movement patterns on yellow pages advertising. *Journal of Advertising*, 26(1):61–73, 1997.
- [184] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *ACM Siggraph*, volume 21, pages 163–169, 1987.
- [185] Henrik Hautop Lund, John Hallam, and Wei-Po Lee. Evolving robot morphology. In *Evolutionary Computation, 1997., IEEE International Conference on*, pages 197–202. IEEE, 1997.

- [186] Mikko Lyly, Juha Ruokolainen, and Esko Järvinen. Elmer—a finite element solver for multiphysics. *CSC-report on scientific computing*, 2000:156–159, 1999.
- [187] AT Machwe, IC Parmee, et al. Introducing machine learning within an interactive evolutionary design environment. In *DS 36: Proceedings DESIGN 2006, the 9th International Design Conference, Dubrovnik, Croatia*, 2006.
- [188] Samir W Mahfoud. A comparison of parallel and sequential niching methods. In *Conference on genetic algorithms*, volume 136, page 143, 1995.
- [189] Carmel Majidi. Soft robotics: a perspective current trends and prospects for the future. *Soft Robotics*, 1(1):5–11, 2014.
- [190] Naresh K Malhotra. Information load and consumer decision making. *Journal of consumer research*, pages 419–430, 1982.
- [191] Evan Malone and Hod Lipson. Fab@ home: the personal desktop fabricator kit. *Rapid Prototyping Journal*, 13(4):245–255, 2007.
- [192] Kim-Fung Man, Kit-Sang Tang, and Sam Kwong. Genetic algorithms: concepts and applications. *IEEE Transactions on Industrial Electronics*, 43(5):519–534, 1996.
- [193] Naomi Mandel. Shifting selves and decision making: The effects of self-construal priming on consumer risk-taking. *Journal of Consumer Research*, 30(1):30–40, 2003.
- [194] Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, 18(1):50–60, 1947.

- [195] Amy E Mark, William F Boyce, and Ian Janssen. Television viewing, computer use and total screen time in canadian youth. *Paediatrics & child health*, 11(9):595, 2006.
- [196] W.S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of mathematical biology*, 5(4):115–133, 1943.
- [197] Georgios Methenitis. Evolution of soft robots by novelty search. 2014.
- [198] Jean-Arcady Meyer and Stewart W Wilson. *From animals to animats*. MIT Press Cambridge, MA, 1991.
- [199] Risto Miikkulainen, Jason Liang, Elliot Meyerson, Aditya Rawal, Dan Fink, Olivier Francon, Bala Raju, Arshak Navruzyan, Nigel Duffy, and Babak Hodjat. Evolving deep neural networks. *arXiv preprint arXiv:1703.00548*, 2017.
- [200] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [201] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [202] J-B. Mouret and J. Clune. An algorithm to create phenotype-fitness maps. *Proc. of the Artificial Life Conf.*, pages 593–594, 2012.
- [203] Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

- [204] Gerd B Müller. Evo–devo: extending the evolutionary synthesis. *Nature Reviews Genetics*, 8(12):943–949, 2007.
- [205] Gernot R Muller-Putz, Rafal Scherer, Christa Neuper, and Gert Pfurtscheller. Steady-state somatosensory evoked potentials: suitable brain signals for brain-computer interfaces? *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, 14(1):30–37, 2006.
- [206] Paul Muter, Susane A Latrémouille, William C Treurniet, and Paul Beam. Extended reading of continuous text on television screens. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 24(5):501–508, 1982.
- [207] H Frederik Nijhout. Development and evolution of adaptive polyphenisms. *Evolution & development*, 5(1):9–18, 2003.
- [208] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In *Artificial Life IV*, pages 190–197. Cambridge, MA: MIT Press, 1994.
- [209] Stefano Nolfi and Dario Floreano. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press, 2000.
- [210] Stefano Nolfi and Dario Floreano. Synthesis of autonomous robots through evolution. *Trends in cognitive sciences*, 6(1):31–37, 2002.
- [211] Don Norman. Emotion & design: attractive things work better. *interactions*, 9(4):36–42, 2002.
- [212] Mohammad Nabi Omidvar, Xiaodong Li, Yi Mei, and Xin Yao. Cooperative co-evolution with differential grouping for large scale optimization. *IEEE Transactions on Evolutionary Computation*, 18(3):378–393, 2014.

- [213] Cagdas D Onal and Daniela Rus. Autonomous undulatory serpentine locomotion utilizing body dynamics of a fluidic soft robot. *Bioinspiration & biomimetics*, 8(2):026003, 2013.
- [214] Denis Pallez, Philippe Collard, Thierry Baccino, and Laurent Dumercy. Eye-tracking evolutionary algorithm to minimize user fatigue in iec applied to interactive one-max problem. In *Proceedings of the 2007 GECCO conference companion on Genetic and evolutionary computation*, pages 2883–2886. ACM, 2007.
- [215] Ian C Parmee, Dragan Cvetković, Andrew H Watson, and Christopher R Bonham. Multiobjective satisfaction within an interactive evolutionary design environment. *Evolutionary Computation*, 8(2):197–222, 2000.
- [216] Michael J Pazzani and Daniel Billsus. Content-based recommendation systems. In *The adaptive web*, pages 325–341. Springer, 2007.
- [217] Carlos Andrés Pena-Reyes and Moshe Sipper. Evolutionary computation in medicine: an overview. *Artificial Intelligence in Medicine*, 19(1):1–23, 2000.
- [218] Xue Bin Peng, Glen Berseth, and Michiel van de Panne. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 35(4):81, 2016.
- [219] R. Pfeifer and J. C. Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.
- [220] Rolf Pfeifer and Gabriel Gómez. Morphological computation—connecting brain, body, and environment. In *Creating Brain-Like Intelligence*, pages 66–83. Springer, 2009.

- [221] Rolf Pfeifer and Fumiya Iida. Embodied artificial intelligence: Trends and challenges. In *Embodied artificial intelligence*, pages 1–26. Springer, 2004.
- [222] Rolf Pfeifer, Fumiya Iida, and Gabriel Gómez. Morphological computation for adaptive behavior and cognition. In *International Congress Series*, volume 1291, pages 22–29. Elsevier, 2006.
- [223] Rolf Pfeifer, Max Lungarella, and Fumiya Iida. Self-organization, embodiment, and biologically inspired robotics. *science*, 318(5853):1088–1093, 2007.
- [224] Rolf Pfeifer and Christian Scheier. *Understanding intelligence*. MIT press, 2001.
- [225] Jean Piaget and Margaret Cook. *The origins of intelligence in children*, volume 8. International Universities Press New York, 1952.
- [226] Rik Pieters and Luk Warlop. Visual attention during brand choice: The impact of time pressure and task motivation. *International Journal of Research in Marketing*, 16(1):1–16, 1999.
- [227] Dean A Pomerleau. *Neural network perception for mobile robot guidance*, volume 239. Springer Science & Business Media, 2012.
- [228] Alex Poole and Linden J Ball. Eye tracking in hci and usability research. *Encyclopedia of human computer interaction*, 1:211–219, 2006.
- [229] Alexandrin Popescul, David M Pennock, and Steve Lawrence. Probabilistic models for unified collaborative and content-based recommendation in sparse-data environments. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 437–444. Morgan Kaufmann Publishers Inc., 2001.

- [230] TG Pottinger and TR Carrick. Modification of the plasma cortisol response to stress in rainbow trout by selective breeding. *General and comparative endocrinology*, 116(1):122–132, 1999.
- [231] Hilary Putnam. Brains and behavior. *Readings in philosophy of psychology*, 1:24–36, 1980.
- [232] Rudolf A Raff. Evo-devo: the evolution of a new discipline. *Nature Reviews Genetics*, 1(1):74–79, 2000.
- [233] Rudolf A Raff. *The shape of life: genes, development, and the evolution of animal form*. University of Chicago Press, 2012.
- [234] Marc Raibert, Kevin Blankespoor, Gabriel Nelson, Rob Playter, and T Bigdog Team. Bigdog, the rough-terrain quadruped robot. In *Proceedings of the 17th world congress*, volume 17, pages 10822–10825. Proceedings Seoul, Korea, 2008.
- [235] Thierry Rayna and Ludmila Striukova. From rapid prototyping to home fabrication: How 3d printing is changing business model innovation. *Technological Forecasting and Social Change*, 102:214–224, 2016.
- [236] Simon M Reader and Kevin N Laland. Social intelligence, innovation, and enhanced brain size in primates. *Proceedings of the National Academy of Sciences*, 99(7):4436–4441, 2002.
- [237] Brice Rebsamen, Cuntai Guan, Haihong Zhang, Chuanchu Wang, Cheeleong Teo, Marcelo H Ang Jr, and Etienne Burdet. A brain controlled wheelchair to navigate in familiar environments. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, 18(6):590–598, 2010.

- [238] Ingo Rechenberg. Evolutionsstrategien. In *Simulationsmethoden in der Medizin und Biologie*, pages 83–114. Springer, 1978.
- [239] John Rieffel, Davis Knox, Schuyler Smith, and Barry Trimmer. Growing and evolving soft robots. *Artificial Life*, (Early Access):1–20, 2013.
- [240] John Rieffel, Davis Knox, Schuyler Smith, and Barry Trimmer. Growing and evolving soft robots. *Artificial life*, 20(1):143–162, 2014.
- [241] John Rieffel, Frank Saunders, Shilpa Nadimpalli, Harvey Zhou, Soha Hassoun, Jason Rife, and Barry Trimmer. Evolving soft robotic locomotion in PhysX. In *Proc. of the Genetic & Evolutionary Computation Conf.*, pages 2499–2504. ACM, 2009.
- [242] Howard Robinson. Dualism. 2003.
- [243] Alan M Rubin. Ritualized and instrumental television viewing. *Journal of communication*, 34(3):67–77, 1984.
- [244] Javier San Agustin, Henrik Skovsgaard, Emilie Mollenbach, Maria Barret, Martin Tall, Dan Witzner Hansen, and John Paulin Hansen. Evaluation of a low-cost open-source gaze tracker. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications*, pages 77–80. ACM, 2010.
- [245] Siddharth Sanan, J. Moidel, and Christopher G Atkeson. A continuum approach to safe robots for physical human interaction. In *Int'l Symposium on Quality of Life Technology*, 2011.
- [246] Bruno Sareni and Laurent Krahenbuhl. Fitness sharing and niching methods revisited. *IEEE transactions on Evolutionary Computation*, 2(3):97–106, 1998.

- [247] J Ben Schafer, Joseph A Konstan, and John Riedl. E-commerce recommendation applications. In *Applications of Data Mining to Electronic Commerce*, pages 115–153. Springer, 2001.
- [248] Robert R Schaller. Moore’s law: past, present and future. *IEEE spectrum*, 34(6):52–59, 1997.
- [249] Michael Schmidt and Hod Lipson. Age-fitness pareto optimization. In *Genetic Programming Theory and Practice VIII*, pages 129–146. Springer, 2011.
- [250] Jacob Schrum and Risto 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.
- [251] Peter Schulte, Laia Alegret, Ignacio Arenillas, José A Arz, Penny J Barton, Paul R Bown, Timothy J Bralower, Gail L Christeson, Philippe Claeys, Charles S Cockell, et al. The chicxulub asteroid impact and mass extinction at the cretaceous-paleogene boundary. *Science*, 327(5970):1214–1218, 2010.
- [252] Hans-Paul Schwefel. *Numerische optimierung von computer-modellen mittels der evolutionsstrategie*, volume 1. Birkhäuser, Basel Switzerland, 1977.
- [253] 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.
- [254] Idan Segev and Elad Schneidman. Axons as computing devices: basic

- insights gained from models. *Journal of Physiology-Paris*, 93(4):263–270, 1999.
- [255] Lawrence Shapiro. *Embodied cognition*. Routledge, 2010.
- [256] Samuel Sanford Shapiro and Martin B Wilk. An analysis of variance test for normality. *Biometrika*, 52(3):591–611, 1965.
- [257] Dongna Shen, Jung-Hyun Park, Jyoti Ajitsaria, Song-Yul Choe, Howard C Wickle III, and Dong-Joo Kim. The design, fabrication and evaluation of a mems pzt cantilever with an integrated si proof mass for vibration energy harvesting. *Journal of Micromechanics and Microengineering*, 18(5):055017, 2008.
- [258] Robert F Shepherd, Filip Ilievski, Wonjae Choi, Stephen A Morin, Adam A Stokes, Aaron D Mazzeo, Xin Chen, Michael Wang, and George M Whitesides. Multigait soft robot. *Proceedings of the National Academy of Sciences*, 108(51):20400–20403, 2011.
- [259] Robert F Shepherd, Adam A Stokes, Jacob Freake, Jabulani Barber, Phillip W Snyder, Aaron D Mazzeo, Ludovico Cademartiri, Stephen A Morin, and George M Whitesides. Using explosions to power a soft robot. *Angewandte Chemie International Edition*, 52(10):2892–2896, 2013.
- [260] George Gaylord Simpson. *Tempo and mode in evolution*. Number 15. Columbia University Press, 1944.
- [261] Karl Sims. *Artificial evolution for computer graphics*, volume 25. ACM, 1991.
- [262] Karl Sims. Interactive evolution of equations for procedural models. *The Visual Computer*, 9(8):466–476, 1993.

- [263] Karl Sims. Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372, 1994.
- [264] Karl Sims. Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372, 1994.
- [265] Andrew C Slocum, Douglas C Downey, and Randall D Beer. Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In *From animals to animats 6: Proceedings of the sixth international conference on simulation of adaptive behavior*, pages 430–439, 2000.
- [266] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge university press, 1982.
- [267] Linda Smith and Michael Gasser. The development of embodied cognition: Six lessons from babies. *Artificial life*, 11(1-2):13–29, 2005.
- [268] R.R. Sokal and F.J. Rohlf. *Biometry: the principles and practice of statistics in biological research*. WH Freeman, 1995.
- [269] LB Soros, Nick Cheney, and Kenneth O Stanley. How the strictness of the minimal criterion impacts open-ended evolution. In *ALIFE XV, The Fifteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 208–215, 2016.
- [270] Bret K Stanford and Peter D Dunning. Optimal topology of aircraft rib and spar structures under aeroelastic loads.
- [271] K. O. Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, 8(2):131–162, 2007.

- [272] K. O. Stanley, D. B. D'Ambrosio, and J. Gauci. A hypercube-based encoding for evolving large-scale neural networks. *Artificial Life*, 15(2):185–212, 2009.
- [273] K. O. Stanley and R. Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.
- [274] Kenneth O Stanley. Exploiting regularity without development. In *Proceedings of the AAAI Fall Symposium on Developmental Systems*, page 37. AAAI Press Menlo Park, CA, 2006.
- [275] Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.
- [276] Luc Steels and Rodney Brooks. *The artificial life route to artificial intelligence: Building embodied, situated agents*. L. Erlbaum Associates Inc., 1995.
- [277] Agostino Stilli, Helge A Wurdemann, and Kaspar 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.
- [278] Yuuta Sugiyama and Shinichi Hirai. Crawling and jumping by a deformable robot. *The International journal of robotics research*, 25(5-6):603–620, 2006.
- [279] Christopher Summerfield and Tobias Egner. Expectation (and attention) in visual cognition. *Trends in cognitive sciences*, 13(9):403–409, 2009.
- [280] JQ Sun, M Re Jolly, and MA Norris. Passive, adaptive and active tuned vibration absorbers - a survey. *Journal of mechanical design*, 117:234, 1995.

- [281] Richard S Sutton. Reinforcement learning architectures for animats. In *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 288–296, 1991.
- [282] Richard S Sutton and Andrew G Barto. *Introduction to reinforcement learning*, volume 135. MIT Press Cambridge, 1998.
- [283] Thomas Sutula, He Xiao-Xian, Jose Cavazos, and Grayson Scott. Synaptic reorganization in the hippocampus induced by abnormal functional activity. *Science*, 239(4844):1147–1151, 1988.
- [284] Hideyuki Takagi. Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE*, 89(9):1275–1296, 2001.
- [285] Jerry O Talton, Daniel Gibson, Lingfeng Yang, Pat Hanrahan, and Vladlen Koltun. Exploratory modeling with collaborative design spaces. *ACM Transactions on Graphics-TOG*, 28(5):167, 2009.
- [286] YS Tarng, JY Kao, and EC Lee. Chatter suppression in turning operations with a tuned vibration absorber. *Journal of materials processing technology*, 105(1):55–60, 2000.
- [287] Darcy Wentworth Thompson et al. On growth and form. *On growth and form.*, 1942.
- [288] Michael T Tolley, Robert F Shepherd, Bobak Mosadegh, Kevin C Galloway, Michael Wehner, Michael Karpelson, Robert J Wood, and George M Whitesides. A resilient, untethered soft robot. *Soft Robotics*, 1(3):213–223, 2014.

- [289] Barry A Trimmer, Ann E Takesian, Brian M Sweet, Chris B Rogers, Daniel C Hake, and Daniel 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.
- [290] D. Trivedi, C.D. Rahn, W.M. Kier, and I.D. Walker. Soft robotics: Biological inspiration, state of the art, and future research. *Applied Bionics and Biomechanics*, 5(3):99–117, 2008.
- [291] Lyudmila Trut. Early canid domestication: The farm-fox experiment foxes bred for tamability in a 40-year experiment exhibit remarkable transformations that suggest an interplay between behavioral genetics and development. *American Scientist*, 87(2):160–169, 1999.
- [292] Alan M Turing. Intelligent machinery, a heretical theory. *The Turing Test: Verbal Behavior as the Hallmark of Intelligence*, 105, 1948.
- [293] Alan M Turing. Computing machinery and intelligence. *Mind*, pages 433–460, 1950.
- [294] Alan Mathison Turing. The chemical basis of morphogenesis. *Bulletin of mathematical biology*, 52(1-2):153–197, 1990.
- [295] Ozgur Unver, Ali Uneri, Alper Aydemir, and Metin Sitti. Geckobot: A gecko inspired climbing robot using elastomer adhesives. In *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 2329–2335. IEEE, 2006.
- [296] Leigh Van Valen. Brain size and intelligence in man. *American Journal of Physical Anthropology*, 40(3):417–423, 1974.

- [297] WH van Willigen, Evert Haasdijk, and LJHM 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.
- [298] Elizabeth A Vandewater, Victoria J Rideout, Ellen A Wartella, Xuan Huang, June H Lee, and Mi-suk Shim. Digital childhood: electronic media and technology use among infants, toddlers, and preschoolers. *Pediatrics*, 119(5):e1006–e1015, 2007.
- [299] Phillip Verbancsics and Kenneth O Stanley. Constraining connectivity to encourage modularity in HyperNEAT. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 1483–1490. ACM, 2011.
- [300] Andreas Wagner. *Arrival of the Fittest: Solving Evolution’s Greatest Puzzle*. Penguin, 2014.
- [301] M. Wall. Galib: A c++ library of genetic algorithm components. *Mechanical Engineering Department, Massachusetts Institute of Technology*, 87, 1996.
- [302] Richard A Watson, Sevan G Ficici, and Jordan B Pollack. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1–18, 2002.
- [303] Richard A Watson, SG Ficiei, and Jordan B Pollack. Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on*, volume 1. IEEE, 1999.

- [304] Frank Wilcoxon and Roberta A Wilcox. *Some rapid approximate statistical procedures*. Lederle Laboratories, 1964.
- [305] Lee Willerman, Robert Schultz, J Neal Rutledge, and Erin D Bigler. In vivo brain size and intelligence. *Intelligence*, 15(2):223–228, 1991.
- [306] Margaret Wilson. Six views of embodied cognition. *Psychonomic bulletin & review*, 9(4):625–636, 2002.
- [307] Stewart W Wilson. Knowledge growth in an artificial animal. In *Adaptive and Learning Systems*, pages 255–264. Springer, 1986.
- [308] Marco Wirth and Frederic Thiesse. Shapeways and the 3d printing revolution. 2014.
- [309] Julius Wolff. *The law of bone remodelling*. Springer Science & Business Media, 2012.
- [310] Robert Wood, Radhika Nagpal, and Gu-Yeon Wei. Flight of the robobees. *Scientific American*, 308(3):60–65, 2013.
- [311] Robert J Wood. The first takeoff of a biologically inspired at-scale robotic insect. *IEEE transactions on robotics*, 24(2):341–347, 2008.
- [312] B. G. Woolley and K. O. Stanley. On the deleterious effects of a priori objectives on evolution and representation. In *Proc. of the Genetic and Evolutionary Computation Conf.*, pages 957–964. ACM, 2011.
- [313] Brian G Woolley and Kenneth O Stanley. Exploring promising stepping stones by combining novelty search with interactive evolution. *arXiv preprint arXiv:1207.6682*, 2012.

- [314] Kai Xu, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen. Fit and diverse: set evolution for inspiring 3d shape galleries. *ACM Transactions on Graphics (TOG)*, 31(4):57, 2012.
- [315] Songhua Xu, Hao Jiang, and Francis Lau. User-oriented document summarization through vision-based eye-tracking. In *Proceedings of the 14th international conference on Intelligent user interfaces*, pages 7–16. ACM, 2009.
- [316] Letizia Zullo, German Sumbre, Claudio Agnisola, Tamar Flash, and Binyamin Hochner. Nonsomatotopic organization of the higher motor centers in octopus. *Current Biology*, 19(19):1632–1636, 2009.