INTERFACE

Research

Author for correspondence:
Nick Cheney
e-mail: nac93@cornell.edu

Scalable co-optimization of morphology and control in embodied machines

Nick Cheney[1,2,3], Josh Bongard[3], Vytas SunSpiral[4] and Hod Lipson[5]

[1]Department of Computational Biology and Biological Statistics, Cornell University, Ithaca, NY, USA
[2]Department of Computer Science, University of Wyoming, Laramie, WY, USA
[3]Department of Computer Science, University of Vermont, Burlington, VT, USA
[4]Intelligent Robotics Group, Intelligent Systems Division, NASA Ames/SGT Inc., Mountain View, CA, USA
[5]Department of Mechanical Engineering, Columbia University, New York, NY, USA

NC, 0000-0002-7140-2213

Evolution sculpts both the body plans and nervous systems of agents together over time. By contrast, in artificial intelligence 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. In psychology, the theory of embodied cognition posits that behaviour 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 behavioural performance. Here, we further examine this hypothesis and demonstrate a technique for 'morphological innovation protection', which temporarily reduces selection pressure on recently morphologically changed individuals, thus enabling evolution some time to 'readapt' to the new morphology with subsequent control policy mutations. We show the potential for this method to avoid local optima and converge to similar highly fit morphologies across widely varying initial conditions, while sustaining fitness improvements further into optimization. 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 behavioural training—while simultaneously providing a test bed to investigate the theory of embodied cognition.

1. Introduction

Designing agile, autonomous machines has been a long-standing challenge in the field of robotics [1]. 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 [2–4]. However, another competing approach is the use of evolutionary algorithms to search for robotic designs and behaviours without presupposing what those designs and behaviours may be. These methods often take inspiration from the evolutionary method 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 [5], as it allows for machines that are specialized for behaviours 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, developmental, or behavioural theory [6–8].

THE ROYAL SOCIETY
PUBLISHING

However, the generalization of design automation to include both the optimization of robot neural controllers and body plans has proved to be problematic. While recent successes have demonstrated the potential of effective optimization for the control policies of agents with fixed morphologies (body plans) [9–11] or—to a lesser extent—the optimization of morphologies for agents with minimal and fixed control policies [12–14], the co-optimization of the two has seen very limited success [15,16].

This inability to perform robot brain–body co-optimization at scale 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 [17], rather than a lack of negative results in unpublished works. Joachimczak *et al.* provide an anecdotal example of premature convergence in the co-optimization of robot brain and body plan (fig. 19 of [18]). Cheney *et al.* [16] further analyse 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 behaviours 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. They support this hypothesis by demonstrating an asymmetry between the ability to optimize the morphology and the controller of their robot, but do not demonstrate the increased effectiveness of an algorithm that addresses this hypothesized co-dependence, as we do here.

2. Background

Attempts to solve this problem of fully automated robot evolution are frequently traced back to the work of Sims [19]. This work introduced the use of evolutionary algorithms to produce goal-directed behaviours and morphologies simultaneously. Despite the advance this work represented, the evolved robots tended to be composed of a small number of components (figures show a mean of 6.042 segments per robot, with each typically controlled by a few neurons). It has been suggested that current computational power should have vastly increased the scale and complexity of robots evolved using Sims' and similar methods. Yet surveys of evolved robots (e.g. [15,20]) 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 [21–23] or genetic encodings [13,24,25], while others pointed to a lack of incentives for complexity in the simple tasks and environments [14,26]. 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 hypothesis, first suggested in [16], that considers the way in which an agent's brain and body plan interact during the optimization process.

In the simulated soft robot platform we employ here, previous works have demonstrated the evolution of effective locomotion behaviour. However, these works did not address the co-optimization of the morphology and control of robots, as control-specific parameters were not optimized in those systems. Behaviours of those robots were instead implicitly specified by the material type and position of cells in the morphology (where different 'muscle' cells simply oscillated counter-phase to one another and with a fixed amplitude) [13,26] or by physically encoding pathways of spiking electrical activity through the placement of conductive and insulating cells in the morphology [27].

3. Material and methods

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. This specialization 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.

By rejecting morphological mutations that are immediately detrimental to fitness, traditional methods fail to consider the full long-term potential of the new morphology—as the initially poor fitness may be due to the robot controller's specialization for producing behaviour in the old morphology, rather than any structural features of the new morphology itself. Our proposed method enables new morphological mutations to be temporarily protected from morphological selection pressures while the robot's controller readapts its function to produce behaviour in this new morphology. Thus we judge the benefit of new structural features by their ability to produce fit behaviour only after they are coupled with a well-suited controller—as we believe this to be a better indicator of the potential benefits of this morphological mutation.

3.1. 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 [28]. Indeed the explicit separation between 'brain' and 'body plan' in this set-up draws attention to the difficulties that occur when we frame optimization with this dichotomy. In our soft robot platform, we simply define the presence and material of the physical voxel cells that make up the robot as its morphological parameters, and define the amplitude and phase offset at which each of these cells oscillate (expand and contract) as their controller parameters. We define a genome as two separate neural-network-like functions (one encoding the brain and one encoding the body), as this explicit separation allows evolution to make isolated adaptations to either subsystem. Future works will need to explore the impact of this particular segmentation of robot features into 'morphological' and 'control' parameters.

3.2. Controller readaptation

The most obvious method for modelling 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 behaviour, 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.

3.3. Multi-objective 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 'morphological age' of the robot). 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, this latter robot is said to 'dominate' the former robot in this mutli-objective optimization (on 'age' and 'fitness')—and the former robot is removed from the population.

Inspiration for this technique comes from [21,29], where completely random robots are inserted into the population at each generation of the evolutionary algorithm. The presence of these new random robots may increase the diversity of the population (and potentially lead to more desirable local optima, if these robots are protected from global selection pressure long enough to climb their local fitness peaks). Our proposed method differs from [21,29], as we do not insert new robots into the population, but rather treat robots with newly mutated morphologies as though their controllers had been randomly scrambled (as the behaviour of these controllers are specialized to their previous morphologies) and are in need of protection in the form of temporary reduction of selection pressure while the controllers climb uphill in the fitness landscape of behaviours for that new morphology.

This procedure has the effect of 'protecting' new morphologies with poorly adapted controllers, and will henceforth be referred to as 'morphological innovation protection'. Various other methods also exist for encouraging diversity in an evolving population (e.g. fitness sharing, crowding, random restart parallel hillclimbers [30], novelty [22] or speciation [31]); however, age was chosen for its simplicity of implementation and because it helps to avoid the cost of extended control re-optimization for non-promising morphologies—since the age-pareto optimization allows fitness 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 predefined amount of readaptation).

3.4. Evolutionary algorithm

All optimization is performed by a population-based evolutionary algorithm. All trials follow a (μ, λ)-Evolutionary Strategy [32] 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.

3.5. Genetic encoding for soft robot morphologies

Consistent with prior work studying the co-optimization of robot morphologies [13] and controllers [16], we choose soft robots as our model system due to the open-ended complexity of deformable voxel-based morphologies and distributed controllers. The soft robot morphologies are encoded with a compositional pattern producing network (CPPN) [33]. 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 variety 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). As nearby voxels tend to have similar coordinate inputs, they also tend to produce similar outputs from the network—creating continuous muscle or tissue patches. CPPNs also produce complex 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.

3.6. Soft robot resolution

As the CPPN genetic encoding is a continuous function (mapping the location of a cell to its cell fate) it may be discretized into a phenotype at any resolution (i.e. creating any number of voxels in the morphology, and a unique controller for each voxel), 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.

3.7. Controllers and their genetic encoding

A unique controller is optimized for each muscle cell in the robot's morphology. For each voxel's controller, two parameters are optimized (the outputs of a separate CPPN with the same inputs as before). One is the phase offset between each individual cell's muscle oscillations and that of the global sinusoidal oscillator (which acts as a central pattern generator), and the second is the frequency of this global clock (since CPPNs do not currently enable global parameters, this is done by averaging the local values at each cell to produce a single global value).

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), defining a robot's behaviour. Passive tissue cells remain at their original size (though they also deform based on their intrinsic compliance).

While this encoding is simple and straightforward, it has the ability to produce complex behaviours, 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 the resonance of soft tissues in appendages).

4

3.8. Physics simulation for evaluation

Once the morphology and controller for a given robot are specified, the fitness (locomotion distance) of that robot is determined by constructing and simulating that robot in the VoxCad soft-body physics simulator [34]. Simulations last for 20 actuation cycles.

3.9. Statistical analysis

All treatments were performed for 30 independent trials. 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. All p-values are generated by a Wilcoxon rank-sum test [35].

4. Results

4.1. The effect of morphological innovation protection on fitness

For the task of locomotion ability (a standard task in evolutionary robotics [20]), 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—i.e. with no 'innovation protection'). In this set-up, the traditional method produces robots with an average fitness of 21.717 voxels.

Additionally, we optimized robots in the same task and environment set-up, 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. This treatment produces significantly more effective robots ($p = 6.067 \times 10^{-6}$), with a mean distance travelled of 31.953 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 [16] is correct, as that work demonstrated the asymmetric difficulty in optimizing the morphology of a robot (as compared to optimizing its controller) and drew the conclusion that this was because the morphology encapsulated the controller (acting as a translator between the 'language' of 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 encoding, it also introduces confounding effects, such as the added population diversity afforded by 'protection' and the added dimensionality of the search space from this protection age—moving search from a single-objective to multi-objective optimization problem.

To tease apart the influence of these two confounds, 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 morphology has 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 temporary reductions in selection pressure. Yet it does not rely on the idea of a broken 'morphological language'. In this treatment, robots locomote 22.049 voxels on average, which fails to show a significant improvement over the single-objective case of no protection ($p = 0.240$), and performs


Figure 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.

significantly worse ($p = 1.211 \times 10^{-4}$) than the protection of morphological innovations.

The fitness trajectories over evolutionary time (figure 1) demonstrate a typical logarithmic fitness improvement over approximately the first 1000 generations, but then show a stagnation for the traditional optimization procedure without innovation protection. The mean fitness value of the treatment without protection shows 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 with morphological innovation protection shows significant improvements over this time ($p = 0.013$) from 25.925 at generation 1000 to 31.953 voxels at generation 5000.

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 between the morphology and controller takes time to become established—and would not introduce fragility into the system before then.

4.2. The effect of morphological innovation protection on population stagnation

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. Figures 2 and 3 represent typical runs, and help to give an intuition of the optimization process. In these figures, each randomly coloured line represents a unique morphology, plotted by its locomotion ability over optimization time.

Note the continued improvement in performance of the most fit individual over optimization time in the case of morphological innovation protection (figure 2)—which is consistent with the trend seen on average in figure 1. This is not seen in the case without innovation protection (figure 3), where the best individual was found before generation 2000—and by generation 1000, fitness has reached 99.6% of its final value.

Consistent with the above observation, we see that the most fit individual in figure 2 changes rapidly in the trial with morphological innovation protection. As each colour in the figure represents a unique morphology, we note that a wide variety of different morphologies hold the title of 'best-so-far'. On average, in runs with morphological innovation protection,

Figure 2. A single optimization trial featuring morphological innovation protection. Each unique morphology is represented by a random colour. 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.


Figure 3. The fitness over time of a single optimization trial with the same initial conditions as figure 2, but without any innovation protection. Again, each unique colour represents a unique morphology. The prevalence of coloured 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 figure 2).

24.179 unique morphologies are the best-so-far at some point in optimization, where the runs without protection show significantly ($p = 1.555 \times 10^{-6}$) less turnover of morphologies, with just 10.115 unique robot body plans doing so.

The question of how temporary reduced selection pressure (via the morphological-age dimension) of morphological innovation protection may help to improve overall fitness and continued optimization may be best demonstrated

Figure 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 colour of the morphologies represents their fitness values.

in the pop-out box for figure 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 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), morphological innovation protection does not expect this new robot to immediately outperform its parent, and 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. We will call this phenomenon a morphological 'overtake'. As the fitness of the parent 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 continuing throughout this run (as the blue child then overtakes the red parent, and green overtakes blue in the pop out of figure 2). On average across all runs, we see morphological overtakes significantly


Figure 5. The progression of morphologies over evolutionary time for the top 10 performing runs with no innovation protection. The colour legend and column snapshot times are identical to figure 4.

($p \leq 6.939 \times 10^{-10}$) more often in runs with morphological innovation protection (an average of 76.714 overtakes in the first 5000 generations) than without any innovation protection (where there are only 1.432 overtakes). We also see more morphological overtakes in the morphological innovation protection treatment than in the trials with controller innovation protection (where there are just 1.333 morphological overtakes on average).

We could also consider controller 'overtakes' occurring in the treatment for controller innovation protection. In that case, an overtake would consist of an initial mutation to the controller of a robot that leads to worse fitness than its parent, which is then followed by subsequent iterations of morphological readaptation, eventually overtaking the fitness of its parent. Interestingly, the number of morphological overtakes in the morphological innovation protection treatment is not significantly different ($p = 0.533$) than the number of controller overtakes in the controller innovation protection treatment (74.542 times on average). Combined with the finding in figure 1 (that morphological innovation protection leads to more fit robots than controller innovation protection), 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.

4.3. The effects of morphological innovation protection on the progression of morphologies over evolutionary time

Visual inspection of actual morphologies over evolutionary time further supports the proposed method's improved optimization efficiency and ability to escape local optima. Figure 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 colour of each robot). Also note how the final morphology of some robots (e.g. runs 24, 27, 18, 11, and 16) results in identical morphologies, despite starting from a variety of initial conditions and not finding this convergent morphology until hundreds or thousands of generations into the optimization process.

In contrast with the sustained turnover of morphologies shown above, figure 5 shows the snapshots of the 10 best runs in the treatment without innovation protection. Note how the fitness/colours of the robots tend to show little change over the evolutionary process, mirroring the stagnation shown in figure 1. While convergence of the final morphologies is present here as well, the gross morphologies found here (variants of the full cube with no appendages) are found early on in the search—and often provided in the set of random initial morphologies. In this treatment, gross morphological changes tend to be absent after generation 50 (just 1% into the full 5000 generations).

The differences between figures 4 and 5 suggest that the traditional method without morphological innovation protection tends to converge prematurely to morphologies early in the evolutionary search, while morphological innovation protection may allow search to escape these local optima and converge to 'more global' optima.

4.4. The potential for convergence across initial conditions

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 rather than the lower resolution morphologies ($5^3 = 125$ voxels) employed in the previous experiments.

The increased number of 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 algorithm. As the number of cells increases, 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'. This parameter is specific to the voxel-based soft robot implementation employed in this work. 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 higher-resolution morphologies, we find that resetting the 'morphological age' of an individual only after a mutation that changes more than 20% of their voxels produces optimal results. The value was found via a parameter sweep (of minimum age-resets of 0%, 5%, 10%, 20%, 30%, 40% and 50% of voxels changed). We also investigated a minimum change threshold for controller innovation protection and found that its benefit falls as the threshold increases (showing

optimal performance with no threshold), so we ignore the threshold for minimum controller change here.

In this high-resolution setting, basic morphological innovation protection (with no threshold) travels significantly farther (60.728 voxels) than the case with no protection (32.575 voxels, $p = 1.208 \times 10^{-4}$). Basic morphological innovation protection (with no threshold) also travels farther than the case with controller innovation protection (35.375 voxels, $p = 0.018$). While morphological innovation protection with a threshold (60.728 voxels) did achieve higher average fitness than morphological innovation protection without a threshold (43.340 voxels), this difference was not significant at the $p = 0.05$ level ($p = 0.063$).

As a side note, in a more difficult set-up (where the robots need to optimize a closed-loop neural network controller), trials with the morphological change threshold significantly outperform all other treatments (travelling 31.313 voxels—while trials with no protection, controller protection, and morphological innovation protection with no threshold travel 13.941, 15.356, and 17.224 voxels, respectively; all $p \leq 7.287 \times 10^{-8}$).

Visualization of these soft robot morphologies over time in the trials with a 20% morphological-change threshold is depicted in figure 6 and compared to the case with no protection in figure 7. The ability of evolution to converge to the same high-performing morphology across many independent trials, despite starting from different initial conditions, suggests that the inclusion of thresholded morphological innovation protection may help to escape local optima around these starting conditions and find 'more global' optima in this search space.

In the case without any protection, search stagnates quickly and again appears unable to escape the local optima near its initial conditions (figure 7).

Interestingly, the low resolution soft robot implementation (figures 4 and 5) does not benefit from the inclusion of a threshold. As presumably the more discrete construction of these robots mean that all mutations are large enough to create a meaningful 'morphological change'.

5. Discussion

The above results demonstrate a new method for entire robot evolution (i.e. both brain and body plan) that is more scalable in terms of continued optimization for longer periods of time and better resulting fitness than the traditional evolutionary methods. 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 [16].

The hypothesis from [16] that the fragile co-optimization of brain and body plan caused by specialization of one subcomponent to the other is consistent with the findings above. Furthermore, the benefits of temporarily reduced selection pressure provided by morphological innovation protection suggest 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 to help evolution to rate proposed solutions based on their long-term potential rather than on immediate fitness impact. As was shown here, using morphological innovation protection for this purpose can help to reduce premature convergence in the search space and stagnation at suboptimal values.

Figure 6. The progression of morphologies over evolutionary time in the setting with 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. Colour represents 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 final morphology.

Figure 7. The progression of high-resolution morphologies over evolutionary time for the top 10 performing runs with no innovation protection. The colour legend and column snapshot times are identical to figure 6.

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 arising from the psychological theory of embodied cognition [36] and uses this intuition to propose a method to escape local optima in the fitness landscape.

5.1. Analogies to biological systems

The foundational motivation of this method, that brain and body are tightly coupled in creating behaviour, is well appreciated in biomechanics and neuromuscular control. In a review of this area, Nishikawa et al. [37] state that 'the inevitable coupling between neural information processing and the emergent mechanical behaviour of animals is a central theme in neurobiology today. Such 'neuromechanical' approaches ask how mechanical systems may offload some tasks of the neural system; how size, shape, structural properties and even the physics of the medium may determine how the neural system functions to control movement.' They note that 'From a design perspective, it seems clear that algorithms for controlling a system must take into account the details of how that system works.' They provide numerous examples of the coupling between physical structure/layout and dynamic control in nature, such as how the hind-limb muscles of the death-head cockroach appear to have similar properties and positioning, yet 'their mechanical functions during dynamic contractions differ considerably' [37,38]. They also suggest examples of the converse, how physical embodiment affects control in natural systems, such as noting that the physical lag of insect visual systems increases the gain on the responses to such stimuli,

resulting in a resonant-like behaviour [37,39]. A review of 'why neuroscience needs biomechanics' by Tytell et al. echos this idea that 'together, the nervous system, body, external environment, and sensory systems form a set of distributed, nested feedback loops' [40]. They note that control changes, such as muscle activations that lengthen or stiffen that muscle, cause substantial changes in the kinematics and dynamics of the structure's motion and behaviour—and that these subsystems should not be considered independent of one another.

The biological analogues of the proposed method for reduced selection pressure on new morphological innovations are also interesting. In biological evolution, new morphological innovations are thought to open up previously inaccessible niches [41], potentially leading to rapid innovations in the style of a punctuated equilibrium [42]. With little competition in these new niches, individuals with these new morphological innovations may have the potential to experience temporarily reduced selection pressure until competition within that niche builds. Morphological innovations may potentially increase the degrees of freedom within a structural system [43], which can also reduce the selective pressure on individual aspects of the morphology and lead to further evolutionary changes [44].

The notion that an individual's brain learns to control its body during its lifetime through sensorimotor exploration [45] also fits well within the framework of this proposed algorithm. If gross morphological changes tend to occur more slowly than the rate of learning in the brain (e.g. morphological changes through evolutionary adaptations or gradual lifetime growth), then the bias seen in these robotic results towards greater success when controllers rapidly readapt to more gradual morphological changes (rather than vice versa, where morphologies rapidly adapt to more gradual controller changes) may be consistent with the rates of change in their biological analogues as well.

5.2. Future work

Despite the significant improvement to our ability to simultaneously optimize the brain and body plan of embodied robots, there is much work still to be done. The proposed method was only applied to one class of robot. This class may represent a fairly straightforward version of brain–body co-optimization, as the distributed sensing and actuation paradigm was designed to possess helps to blur the line between physical interactions of the morphology with the environment and information processing of a controller [13,27]. 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 (e.g. limbs, joints, sensors). 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.

We do not believe that the methods introduced here are 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 [46,47], we should also note that our domain—an agent's controller embodied within its morphology—is closely related to that of a neural network's weights embodied within its topology. Future work will show whether the method proposed here will show similar optimization gains in the design of neural network topologies as well.

6. 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 random 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.

References

1. Pfeifer R, Lungarella M, Iida F. 2007 Self-organization, embodiment, and biologically inspired robotics. *Science* **318**, 1088–1093. (doi:10.1126/science.1145803)

2. Chestnutt J, Lau M, Cheung G, Kuffner J, Hodgins J, Kanade T. 2005 Footstep planning for the Honda ASIMO humanoid. In *Proc. 2005 IEEE Int. Conf. on Robotics and Automation, Barcelona, Spain, 18–22 April 2005*, pp. 629–634. (doi:10.1109/ROBOT.2005.1570188)

3. Raibert M, Blankespoor K, Nelson G, Playter R, Team TB. 2008 Bigdog, the rough-terrain quadruped robot. In *Proc. 17th World Congress of the International Federation of Automatic Control, Seoul, Korea*, vol. 17, pp. 10 822–10 825.

4. Wood RJ. 2008 The first takeoff of a biologically inspired at-scale robotic insect. *IEEE. Trans. Robot.* **24**, 341–347. (doi:10.1109/TRO.2008.916997)

5. Langton CG, Taylor C, Farmer JD, Rasmussen S. 1989 *Artificial life*. Redwood City, CA: Addison-Wesley Publishing Company.

6. Lenski RE, Ofria C, Collier TC, Adami C. 1999 Genome complexity, robustness and genetic interactions in digital organisms. *Nature* **400**, 661–664. (doi:10.1038/23245)

7. Lenski RE, Ofria C, Pennock RT, Adami C. 2003 The evolutionary origin of complex features. *Nature* **423**, 139–144. (doi:10.1038/nature01568)

8. Bongard J. 2011 Morphological change in machines accelerates the evolution of robust behavior. *Proc. Natl Acad. Sci. USA* **108**, 1234–1239. (doi:10.1073/pnas.1015390108)

9. Geijtenbeek T, van de Panne M, van der Stappen AF. 2013 Flexible muscle-based locomotion for bipedal creatures. *ACM Trans. Graphics (TOG)* **32**, 206. (doi:10.1145/2508363.2508399)

10. Cully A, Clune J, Tarapore D, Mouret J-B. 2015 Robots that can adapt like animals. *Nature* **521**, 503–507. (doi:10.1038/nature14422)

11. Levine S, Finn C, Darrell T, Abbeel P. 2016 End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.* **17**, 1334–1373.

12. Collins S, Ruina A, Tedrake R, Wisse M. 2005 Efficient bipedal robots based on passive-dynamic walkers. *Science* **307**, 1082–1085. (doi:10.1126/science.1107799)

13. Cheney N, MacCurdy R, Clune J, Lipson H. 2013 Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proc. 15th Annual Conf. on Genetic and Evolutionary Computation, Amsterdam, The Netherlands, 6–10 July 2013*, pp. 167–174. New York, NY: ACM. (doi:10.1145/2463372.2463404)

14. Auerbach JE, Bongard JC. 2014 Environmental influence on the evolution of morphological complexity in machines. *PLoS Comput. Biol.* **10**, e1003399. (doi:10.1371/journal.pcbi.1003399)

15. Geijtenbeek T, Pronost N. 2012 Interactive character animation using simulated physics: a state-of-the-art review. In *Computer graphics forum*, vol. 31, pp. 2492–2515. Oxford, UK: Blackwell Publishing.

16. Cheney N, Bongard J, SunSpiral V, Lipson H. 2016 On the difficulty of co-optimizing morphology and control in evolved virtual creatures. In *ALIFE XV, 15th Int. Conf. on the Synthesis and Simulation of Living Systems*, vol. 15, pp. 226–233. Cambridge, MA: MIT Press.

17. Fanelli D. 2011 Negative results are disappearing from most disciplines and

countries. *Scientometrics* **90**, 891–904. (doi:10.1007/s11192-011-0494-7)

18. Joachimczak M, Suzuki R, Arita T. 2016 Artificial metamorphosis: evolutionary design of transforming, soft-bodied robots. *Artificial Life* **22**, 271–298. (doi:10.1162/ARTL_a00207)

19. Sims K. 1994 Evolving virtual creatures. In *Proc. 21st Annual Conf. on Computer Graphics and Interactive Techniques*, pp. 15–22. New York, NY: ACM. (doi:10.1145/192161.192167)

20. Bongard JC. 2013 Evolutionary robotics. *Commun. ACM* **56**, 74–83.

21. Hornby GS. 2006 Alps: the age-layered population structure for reducing the problem of premature convergence. In *Proc. 8th Annual Conf. on Genetic and Evolutionary Computation*, pp. 815–822. New York, NY: ACM. (doi:10.1145/1143997.1144142)

22. Lehman J, Stanley KO. 2011 Evolving a diversity of virtual creatures through novelty search and local competition. In *Proc. 13th Annual Conf. on Genetic and Evolutionary Computation, Dublin, Ireland, 12–16 July 2011*, pp. 211–218. New York, NY: ACM. (doi:10.1145/2001576.2001606)

23. Mouret J-B, Clune J. 2015 Illuminating search spaces by mapping elites. (http://arxiv.org/abs/1504.04909)

24. Hornby GS, Pollack JB. 2001 Body-brain co-evolution using l-systems as a generative encoding. In *Proc. 3rd Annual Conf. on Genetic and Evolutionary Computation, San Francisco, CA, USA, 7–11 July 2001*, pp. 868–875. San Francisco, CA: Morgan Kaufmann.

25. Bongard JC, Pfeifer R. 2003 Evolving complete agents using artificial ontogeny. In *Morpho-functional machines: the new species*, pp. 237–258. Tokyo, Japan: Springer.

26. Cheney N, Bongard J, Lipson H. 2015 Evolving soft robots in tight spaces. In *Proc. 2015 Annual Conf. on Genetic and Evolutionary Computation, Madrid, Spain, 11–15 July 2015*, pp. 935–942.

New York, NY: ACM. (doi:10.1145/2739480.2754662)

27. Cheney N, Clune J, Lipson H. 2014 Evolved electrophysiological soft robots. In *ALIFE 14: 14th Conf. on the Synthesis and Simulation of Living Systems*, vol. 14, pp. 222–229. Cambridge, MA: MIT Press.

28. Baluska F, Levin M. 2016 On having no head: cognition throughout biological systems. *Front. Psychol.* **7**, 902. (doi:10.3389/fpsyg.2016.00902)

29. Schmidt M, Lipson H. 2011 Age-fitness pareto optimization. In *Genetic programming theory and practice VIII* (eds R Riolo, T McConaghy, E Vladislavleva), pp. 129–146. Berlin, Germany: Springer.

30. Gordon VS, Whitley D. 1993 Serial and parallel genetic algorithms as function optimizers. In *Proc. 5th Int. Conf. on Genetic Algorithms*, pp. 177–183. San Francisco, CA: Morgan Kaufmann.

31. Stanley KO, Miikkulainen R. 2002 Evolving neural networks through augmenting topologies. *Evol. Comput.* **10**, 99–127. (doi:10.1162/106365602320169811)

32. Beyer H-G, Schwefel H-P. 2002 Evolution strategies—a comprehensive introduction. *Nat. Comput.* **1**, 3–52. (doi:10.1023/A:1015059928466)

33. Stanley KO. 2007 Compositional pattern producing networks: a novel abstraction of development. *Genet. Program. Evol. Mach.* **8**, 131–162. (doi:10.1007/s10710-007-9028-8)

34. Hiller J, Lipson H. 2014 Dynamic simulation of soft multimaterial 3D-printed objects. *Soft Robot.* **1**, 88–101. (doi:10.1089/soro.2013.0010)

35. Wilcoxon F, Wilcox RA. 1964 *Some rapid approximate statistical procedures*. London, UK: Lederle Laboratories.

36. Pfeifer R, Bongard JC. 2006 *How the body shapes the way we think: a new view of intelligence*. Cambridge, MA: MIT press.

37. Nishikawa K et al. 2007 Neuromechanics: an integrative approach for understanding motor control. *Integr. Comp. Biol.* **47**, 16–54. (doi:10.1093/icb/icm024)

38. Ahn AN, Meijer K, Full RJ. 2006 *In situ* muscle power differs without varying *in vitro* mechanical properties in two insect leg muscles innervated by the same motor neuron. *J. Exp. Biol.* **209**, 3370–3382. (doi:10.1242/jeb.02392)

39. Harris RA, O'Carroll DC, Laughlin SB. 1999 Adaptation and the temporal delay filter of fly motion detectors. *Vision Res.* **39**, 2603–2613. (doi:10.1016/S0042-6989(98)00297-1)

40. Tytell E, Holmes P, Cohen A. 2011 Spikes alone do not behavior make: why neuroscience needs biomechanics. *Curr. Opin. Neurobiol.* **21**, 816–822. (doi:10.1016/j.conb.2011.05.017)

41. Simpson. GG. 1944 *Tempo and mode in evolution*, vol. 15. New York, NY: Columbia University Press.

42. Jay Gould S, Eldredge N. 1993 Punctuated equilibrium comes of age. *Nature* **366**, 223–227. (doi:10.1038/366223a0)

43. Vermeij GJ. 1973 Adaptation, versatility, and evolution. *Syst. Zool.* **22**, 466–477. (doi:10.2307/2412953)

44. Lauder GV. 1990 Functional morphology and systematics: studying functional patterns in an historical context. *Annu. Rev. Ecol. Syst.* **21**, 317–340. (doi:10.1146/annurev.es.21.110190.001533)

45. Thelen E. 1995 Motor development: a new synthesis. *Am. Psychol.* **50**, 79–95. (doi:10.1037/0003-066X.50.2.79)

46. Fernando C, Banarse D, Blundell C, Zwols Y, Ha D, Rusu AA, Pritzel A, Wierstra D. 2017 Pathnet: evolution channels gradient descent in super neural networks. (http://arxiv.org/abs/1701.08734)

47. Miikkulainen R et al. 2017 Evolving deep neural networks. (http://arxiv.org/abs/1703.00548)