



Investigating Premature Convergence in Co-optimization of Morphology and Control in Evolved Virtual Soft Robots

Alican Mertan^(✉) and Nick Cheney

University of Vermont, Burlington, VT 05401, USA
{alican.mertan,ncheney}@uvm.edu

Abstract. Evolving virtual creatures is a field with a rich history and recently it has been getting more attention, especially in the soft robotics domain. The compliance of soft materials endows soft robots with complex behavior, but it also makes their design process unintuitive and in need of automated design. Despite the great interest, evolved virtual soft robots lack the complexity, and co-optimization of morphology and control remains a challenging problem. Prior work identifies and investigates a major issue with the co-optimization process – fragile co-adaptation of brain and body resulting in premature convergence of morphology. In this work, we expand the investigation of this phenomenon by comparing learnable controllers with proprioceptive observations and fixed controllers without any observations, whereas in the latter case, we only have the optimization of the morphology. Our experiments in two morphology spaces and two environments that vary in complexity show, concrete examples of the existence of high-performing regions in the morphology space that are not able to be discovered during the co-optimization of the morphology and control, yet exist and are easily findable when optimizing morphologies alone. Thus this work clearly demonstrates and characterizes the challenges of optimizing morphology during co-optimization. Based on these results, we propose a new body-centric framework to think about the co-optimization problem which helps us understand the issue from a search perspective. We hope the insights we share with this work attract more attention to the problem and help us to enable efficient brain-body co-optimization.

Keywords: Evolutionary robotics · Soft robotics · Brain-body co-optimization

1 Introduction

Evolving virtual creatures is a field with a long history, starting with Karl Sims' seminal work "Evolving Virtual Creatures" [51] almost 30 years ago. Karl Sims' approach of co-optimizing brain and body then widely adapted and researched, especially in soft robotics [2–6, 10, 20, 38, 40, 54].

The field of soft robotics with volumetric actuation started with [13, 15, 55] and accelerated with the availability of simulators [2, 14, 33, 37]. The compliance and flexibility of the material make soft robots capable of exhibiting complex and unintuitive behaviors, resulting in different abilities of soft robots such as walking [2, 5, 6, 10, 20, 22, 24–26, 38–40, 43, 52], swimming [9, 20], squeezing through obstacles [4], damage recovery and regeneration [16, 17, 28], shape change [48] and self-replication [23] being tested. The highly non-linear and complex nature of soft material dynamics that don't exist in their rigid counterparts [21, 46, 49], provides an increased potential for morphological computation [41, 42] that can be unlocked by a tightly integrated body design and control strategy.

The ability to output complex behavior, however, makes the design process of soft robots counter-intuitive, highly encouraging automated design over manual design. Despite the great interest in the automated brain-body co-optimization [2–6, 9, 10, 13, 15, 20, 22, 26, 35, 38, 40, 53, 54], evolved soft robots struggles to surpass the complexity of the Sims' initial creatures [51]. Prior work identifies and investigates an important phenomenon that hinders the co-optimization process – premature convergence of morphology [3, 20].

As the optimization takes place, two parts of the solution, the brain and the body, become more and more specialized for each other, making the overall performance of the solution very sensitive to changes in either component [3]. In the case of co-optimizing the brain and body in soft robots, this is especially prominent for the morphology of the solution where even small changes in the morphology can drastically reduce the performance of the solution [40]. Commonly used algorithms for brain-body co-optimization don't like solutions with poor performances and focus the search over other parts of the morphology space, making them unable to discover high-performing regions in the search space. To overcome the poor search over morphologies, prior work proposes the use of better search algorithms that reduce the selection pressure of individuals with new body plans to promote more search over morphologies [5, 30], or proposes ways to alleviate the performance decline during the morphological search by using different genetic representations for coordinated changes [54, 57] or by using controllers that are robust to changes in the morphology [40].

Rather than proposing a solution, we provide an investigation of this phenomenon following the previous investigation of Cheney et al. [3]. Specifically, we investigate the use of neural network controllers with proprioceptive observations where we optimize both morphology and control, and heuristically devised fixed controllers where we only optimize the morphology. This allows us to shed light on the effects of control optimization and specialization on the co-optimization process and understand the premature convergence issue. The main contributions of this work are to:

- expand previous investigations towards a more generalized understanding of the brain-body fragile co-optimization phenomenon by focusing on important and missing aspects such as the use of more varied and complex controllers, sensory information, morphology spaces, and environments (Sect. 3).

- provide arguably the most concrete example that shows the existence of high-performing designs that aren’t discovered by the co-optimization process, due to poor search over the morphology space (Sect. 4).
- develop a new body-centric framework to conceptualize the fitness landscape of the co-optimization problem, allowing us to understand the challenge of co-optimization from a search perspective (Sec. 5).
- organize the existing solutions to premature convergence within the proposed framework, helping us to clarify promising research directions (Sect. 5).

We hope these findings and the proposed framework will be useful for further research into enabling better brain-body co-optimization.

2 Experiment Design

2.1 Simulation

We perform our experiments in the Evolution Gym version 1.0.0 (Evogym) [2]. It is an open-source voxel-based soft-body simulator with a suite of benchmark tasks. Voxels are represented as a mass-spring system, and it works in 2D, similar to the simulation engines in [11, 20, 36–40, 43, 54].

Each voxel of the robot can have a different type that determines its behavior and is represented with different colors visually. There are two types of active materials that can actuate volumetrically, either **horizontally** or **vertically**, and they are under the direct control of the controller. These materials are controlled by specifying the target length $a \in [0.6, 1.6]$ (in the form of a multiplier of the resting length) at each time step where materials gradually expand/contract to achieve the target length. We query the controller every 5th timestep (referred to as the effective timestep) and repeat the queried action until the next effective time step to limit high-frequency dynamics that might lead to unstable behavior. Additionally, there are two types of passive materials that are not under direct control, and they differ in their elasticity, one being **rigid** and the other being **elastic**. Designing a robot is simply choosing the existence and material type of voxels in a grid layout. While Evogym allows for specifying whether neighboring voxels are connected to each other, we omit this feature and assume that neighboring voxels are always connected to each other.

The simulation engine provides various proprioceptive and environmental observations. We only use proprioceptive observations of volume, speed, and material properties of voxels and a global time signal.

2.2 Task and Environments

We use two environments with a locomotion task that are presented in the Evogym benchmark suite, namely Walker-v0 and BridgeWalker-v0 [2]. Locomotion is the most used task for voxel-based soft robotics research [3, 5–7, 10, 11, 16, 17, 24–28, 35, 36, 38–40, 43, 52–54]. We adopt the modified fitness function in [40], $R(r, T) = \Delta p_x^r + \mathbf{I}(r) + \sum_{t=0}^T -0.01 + 5$, which rewards robots r for

moving in the positive x direction with the term Δp_x^r , rewards them for finishing the task with the indicator function $\mathbf{I}(p_x^r)$, and encourages them to finish the task faster by applying a negative penalty at each time step with the term $\sum_{t=0}^T -0.01$. The last term, $+5$, is equal in magnitude to the maximum penalty and is used to shift the reward to be positive for ease of analysis.

Environments often consist of a flat [3, 5–7, 10, 16, 17, 24–28, 35, 36, 38–40, 52–54] or uneven [6, 11, 43, 52] surfaces, supposedly harder task. We experiment with both locomotion on a flat surface (Walker-v0) and an uneven, dynamic surface (BridgeWalker-v0) to understand the relationship between the environment and the brain-body co-optimization process.

2.3 Robot Design and Controller Strategies

Robot Representation and Design. Following [2, 38, 40], we use a direct representation for robot design where the existence of voxels and their material type encoded in a matrix $\in M_{h \times w}(T)$, where $T \in \{0, 1, 2, 3, 4\}$. While indirect encodings are commonly used in prior work that works in 3D [3–6, 9, 10, 20, 22, 25], working in 2D makes it possible to use direct encoding which is shown to be comparable to complex indirect encodings [2] and allows us to control for the change in morphology during mutation. Following the common practice of limiting the morphology space [3, 5, 10, 24, 35, 38, 40, 54], we experiment with two morphology spaces, $(h, w) \in \{(5, 5), (7, 7)\}$, to provide more instances for investigation of premature convergence, as encouraged by previous investigation [3].

Controller Design and Model. We compare a learnable controller and a fixed, non-sensing controller. The fixed controller allows us to create a scenario where fragile co-adaptation of design and controller cannot occur due to the lack of optimization on the controller part and helps us investigate premature convergence that we observe with learnable controllers. Specifically, we use a modular/decentralized control strategy as a learnable control¹ as they are shown to help with the brain-body co-optimization [40].

Learnable Controller. One of the common choices in control design is the use of modular control strategy [18, 36, 38, 40, 43]. They are considered to be more versatile and robust to changes in the robot morphology [59], which helps with the co-optimization process [40], and are compatible with any robot morphology.

The modular controller consists of a shared neural network model (single hidden layer with 32 neurons) assigned to each active voxel, observing a local patch around it and only determining the behavior of a single voxel it is assigned. The controller observes a 3×3 window centered around the active voxel. The observations to the controller consist of the proprioceptive information (volume, speed, and material type) of the voxels in this observation window and timestep mod 2 as the time signal to have a fair comparison with the fixed controller.

¹ We repeated every experiment with the most commonly used global/centralized control strategy [11, 18, 36, 39, 40, 52, 53] and didn’t observe any qualitatively different results. For the sake of space and simplicity, we omit the global controller and only present the results with the modular controller.

Fixed Controller. One of the simplest control strategies used in the literature is the fixed actuation following a global signal [3–6, 9, 22, 24–28, 53]. While the learnable controller performs a nonlinear mapping from proprioceptive observations to actions, fixed controller lacks any sensory input. Active voxels under the control of the fixed controller alternate between expanding and contracting maximally at every effective timestep. Therefore, the controller has zero learnable parameters and doesn’t require any control optimization, hence the name fixed. Having a fixed controller that doesn’t require any control optimization allows us to create a scenario where fragile co-adaptation doesn’t occur, helping us understand premature convergence with learnable controllers better.

2.4 Optimization Algorithm

To optimize the design and control of the soft robots, we adopt the use of evolutionary algorithms, similar to [3, 5, 20, 38, 40, 44, 54, 56, 57]. Evolutionary algorithms allow us to simultaneously make improvements in both design and control, supposedly reducing the computational cost of two-level optimization approaches where outer loop does design optimization and the inner loop optimizes the control for a given design, as in [2].

In particular, we use age-fitness Pareto optimization (AFPO) [47] with truncation selection. Individuals’ ages are increased at every generation, and we inject a random individual at each generation with the age of 0 for diversity. Recombination wasn’t considered and the offspring are created through mutation only. Individuals created by mutation inherit the age of their parent.

Following [2, 40], we use a mutation operator that creates new designs by changing each voxel type of a robot with a 10% probability from/to an empty voxel. To mutate in the control space, we add noise sampled from $\mathcal{N}(0, 0.1)$ to all learnable controller parameters.

When creating offspring through mutation, we either mutate the body plan or the controller of the solution. Following [5, 40], there is a 50% probability of choosing either component of the solution for the mutation. Note that the fixed controller doesn’t have any controller parameters. Yet, for a fair comparison, individuals with fixed controllers go through control mutation with the same probability, effectively creating a copy of the same solution.

3 Co-optimization

The main problem we are interested in is the co-optimization of morphological design and control, and understanding the premature convergence of morphology. To compare the effects of different control strategies, we evolve two populations of solutions: one with learnable and one with fixed controllers. While with the learnable controller, we have the co-optimization of morphology and control, the fixed controller allows us to test the case where we only optimize the morphology, allowing us to investigate the effect of controllers on the premature convergence issue. We experiment with evolving these populations under 4




Fig. 1. Fitness over time plots (left) and distributions of the best solutions (right) for the co-optimization experiments. The simplest controller, the fixed controller, significantly outperforms the learnable controller in all experimental settings. Moreover, populations with fixed controllers converge faster. Left: Solid lines show the best fitness found at each generation, averaged across 10 runs. Shaded regions show the 95% confidence intervals. Right: Each data point is plotted, as well as the mean values which are marked with dark red. Horizontal lines indicate statistically different results. ***: $P < 0.001$, **: $P < 0.005$, *: $P < 0.05$ (Color figure online)

different settings by varying the environment, Walker-v0 and BridgeWalker-v0, and the morphology space, $(h, w) \in \{(5, 5), (7, 7)\}$. Each setting is named as $\{\text{Environment name}\}\{(h, w)\}$, e.g. **W5** for the setting Walker-v0, (5, 5). For each setting, we repeat the experiment 10 times and use the Wilcoxon Rank Sum test [58] to report P -values. Each run consists of a population size of 16 individuals evolved for 10000 generations with the AFPO algorithm [47].

Figure 1 left, illustrates the fitness over evolutionary time plots for each experiment. Solid lines show the fitness of the best solution found at each generation, averaged across 10 runs. Shaded regions show the 95% bootstrapped confidence interval. Additionally, Fig. 1 right illustrates the distribution of the performance of the final solutions. The results demonstrate that the populations with fixed controllers find better solutions faster (all $P < 0.05$ for achieving 85% of final performance) and find significantly better solutions (all $P < 0.005$). These results are consistent across different environments and morphology spaces. While individuals with fixed controllers are, admittedly, incapable of observing their environment, hence unlikely to scale to more complex problems, *fixed controllers nevertheless provide a strong baseline for performance for simple tasks such as locomotion, where a lot of research is still being done.*

The success of fixed controller might seem interesting, yet it is a well-established phenomenon that it is difficult to optimize body plan and control together [3, 5, 20, 40]. Having two interdependent parts of a solution creates a challenging optimization problem where the parts become more and more spe-

cialized for each other as the optimization takes place, making them unamenable to change without breaking the solution’s performance – fragile co-adaptation. This is especially true for morphology, as it is the interface between the controller and environment and thus changing morphology also scrambles the controller. It makes it harder to search the morphology space, resulting in premature convergence of morphology [3, 20]. On the other hand, the lack of parameters for the fixed controller turns the co-optimization of the morphology and control problem into only the optimization of the morphology, eliminating the issue of fragile co-adaptation. The results demonstrating superior performance for fixed controllers suggest that *fixed controller doesn’t suffer from premature convergence and allows optimization process to find body plans that outsource complexity to the dynamics of the system, an example of morphological computation* [42].

Lastly, we note that the fixed controller consistently outperforms the more complex learnable controller in the BridgeWalker-v0 environment in two different morphology spaces. In light of these results, *we conjecture that the locomotion over uneven surface doesn’t present fundamentally harder/different optimization problem than locomotion on flat surface for brain-body co-optimization*.

4 Analysis

Our experiments demonstrate the inferior performance of co-optimization with learnable controllers compared to morphology optimization with fixed controllers. In this section, we perform a series of analyses to understand how the differences in controllers affect the co-optimization. Since we expect learnable controllers to be capable of learning to output the simple behavior that the fixed controllers exhibit, the differences in performance should arise from the different controllers’ effects on the search over the morphology space. As prior work strongly emphasizes the issue of premature convergence as the major hurdle in the brain-body co-optimization [3, 20, 40], we specifically focus on this issue in our analysis.

Qualitative Analysis of Body Plans. As we discussed earlier, one of the biggest challenges of brain-body co-optimization is the premature convergence of the morphology [3, 20, 40]. While optimizing a controller for a given body plan is relatively easy, search over the morphology space often prematurely converges to a local optimum. To compare experimented controllers in terms of their effect on the search over the morphology space, we start by looking at the body plans of the best solutions found at each run, i.e. run champions, under each setting.

First, we try to measure the run champions’ convergence to a particular body plan, which indicates a better search over the morphology space as we expect a successful search to be able to exploit high-performing regions of the space across runs [3]. To this end, we perform t-SNE dimensionality reduction [34] of the optimized morphologies. We visualize the clusters in Fig. 2 and measure the average intra-cluster distance, to quantify the variation in found morphologies for a given controller paradigm. We find that run champions with fixed controllers




Fig. 2. t-SNE plots of the run champions’ body plans and average intra-cluster distances in the embedding space. The populations with the fixed controller converge better to a body plan across different runs and morphology spaces.

converge to more similar solutions across different runs and morphology spaces in two different environments. We also note that the learnable run champions show less diversity in the BridgeWalker-v0 environment, showing that the environment also plays a role in shaping the fitness landscape and affecting which body plans can be easily found during the co-optimization process.

To qualitatively analyze the robot body plans and gaits, we present Fig. 3, where we show the body plans of the run champions, and Fig. 4, where we display 15 snapshots during the lifetime of select individuals. In our experimental setup, fixed controllers (Fig. 4b) consistently find similar body plans with near identical gaits. We do not observe this level of converge across runs that co-optimize morphology with a learnable controller (Fig. 4a). This trend holds across all experimental settings.

Quantitative Analysis of Body Plans. Our qualitative analysis of body plans shows that champions with the learnable controller show more diversity while fixed run champions are grouped closer in the morphology space. The diversity of solutions found with learnable controllers can be interpreted as a positive feature. However, the literature strongly suggests that it is a sign of poor search over the morphology space [3]. Although it is possible that the search space contains many distinct near-optimal solutions, the high-performing body plans of fixed run champions imply the existence of better-performing body plans that aren’t discovered in runs with the learnable controller.

Therefore in this section, we try to definitively answer our main question: have learnable controllers failed to find high-performing body plans that were discovered by fixed controllers? To answer this question, we take the morphologies of the run champions found in runs with fixed controllers and no co-optimization, then optimize a learnable controller for each body plan. We suggest that by pairing the same learned closed-loop controller from the co-optimization paradigm to the morphologies found when optimized with the fixed controller, we will be able to fairly compare how well each optimization setup is able to find



(a) L-W5 (b) L-B5 (c) L-W7 (d) L-B7 (e) F-W5 (f) F-B5 (g) F-W7 (h) F-B7

Fig. 3. Morphologies of the run champions under each experimental setting. The left half of the figure shows the run champions with learnable controllers, and the right half shows the ones with fixed controllers. We see diverse body plans evolve with the learnable controller, especially in (a). On the other hand, we see the same form of the mostly active bottom with two vertical apparatus at the front and back, resembling a head and a tail, evolved with the fixed controller in (e). For the BridgeWalker-v0 environment, individuals with learnable controllers show similar body plan features in (b) and (d), e.g. upright posture, r or T shape, individuals with fixed controller usually consists of a thin, horizontal body with active materials and forward apparatus resembling a leg in (f–h).



Fig. 4. Timelapse images of three run champions' behavior from each of the two tasks of the (5,5) morphology space, co-optimized brain and body (a) and morphology-only optimization with fixed control (b). Learnable controllers demonstrate more diversity, especially in the Walker-v0 environment. Most of the individuals with learnable controllers (a) find a bipedal (rows 1, 2, 4–6) or monopedal gait (row 3). Fixed controllers (b) show one common behavior on flat ground (rows 1–3), consisting of an active bottom with two vertical apparatus at the front and back, resembling a head and a tail. When the muscle at the bottom of the robot contracts, vertical apparatuses help the bottom part form an arch, creating front and back legs to locomote. In the bridge environment (rows 4–6) often includes a thin, horizontal body with active materials, a forward apparatus resembling a leg, and some upper body (presumably used for balancing). Individuals with this body plan throw themselves forward by actuating their muscles in phase and pulling themselves forward with their forward apparatus, achieving bipedal locomotion.




Fig. 5. Comparison of performances on the body plans of learnable and fixed run champions with learnable and fixed controllers. The first and last columns in each plot show the results of the main co-optimization experiment, and the second and third columns show the results when the learnable controller is optimized from scratch on body plans of learnable run champions and fixed run champions, respectively. In 3 out of 4 experimental settings, the body plans found by the fixed controller during the co-optimization experiments outperform the body plans found by the learnable controller when the learnable controller is optimized to control them. It demonstrates the failure of search over the morphology space during co-optimization with learnable controllers since they failed to discover high-performing body plans of fixed champions.

high-performing morphologies in the search space. We take the body plans of the run champions with fixed controllers and start new evolutionary runs from scratch, with 16 individuals having the same fixed body plan as the corresponding run champion, and optimize only the controller for 5000 generations.

Figure 5 compares the performance of the learnable controller on the body plans of run champions found during co-optimization with learnable and morphology/only optimization with fixed controllers. For comparison, we also show the original performances of the fixed run champions with the fixed controller. In 3 out of 4 experimental settings (W5, W7, and B5; all $P < 0.05$), the learnable controller achieves significantly better performance on the body plans of fixed champions, which definitively demonstrates the existence of better body plans than the ones found during co-optimization with the learnable controller. *To the best of our knowledge, this is the most concrete example to date showing the failure of brain-body co-optimization to find specific and demonstrably-reachable high-performing regions of the morphology space.*

It is also possible that simply the act of re-training the controller once the body plan has already been found and fixed produces a much different optimization trajectory than co-optimizing both at once. To account for this we also

take the run champions found when co-optimizing morphology with a learnable controller, freeze the resulting morphology and re-train a learnable controller on it from scratch (emulating the setup described above). Across all conditions, this resulted in statistically indistinguishable performance from the original co-optimized robots (Fig. 5, first versus second columns, all $P > 0.44$), demonstrating that the controllers found during co-optimization were well already converged to their current body plan. This finding also provides further evidence consistent with the notion that the struggle to co-optimize brain-body systems is due to sub-par morphologies more so than sub-par controllers.

For completeness, we also measured the performance of learnable run champions’ body plans when they were controlled by the fixed controller. As expected, they performed very poorly (mean performance in W5: 4.88, W7: 6.31, B5: 2.64, B7: 2.47) and most solutions couldn’t even reach the end of the environment. This suggests that the fixed controllers are not inherently better than the learnable controllers, but benefit from enabling search to discover more effective body plans, and especially body plans that are well adapted to the given control policy.

Lastly, contrary to what we expect, we observe that the learnable controller, when it is trained on fixed run champions’ body plans, doesn’t always achieve the same level of performance as fixed controllers do (Fig. 5, third versus last columns, $P < 0.05$ in W5, W7, and B7), which we interpret as the challenge of achieving very simple behaviors with complex neural network controllers when they are optimized through fitness alone. Future work could consider experimenting with more complex tasks and environments to understand how the co-optimization process unfolds when the complexity of the task/environment better matches the complexity of the controller.

5 Discussion

Our co-optimization experiment and the following analysis produced a number of important results that can be seen in Fig. 5. Most importantly, we have shown that optimizing learnable controllers on the high-performing body plans found with fixed controllers, resulted in better solutions compared to the solutions found with learnable controllers during co-optimization in most experimental settings (first vs middle columns, $P < 0.05$ in 3 out of 4), which clearly demonstrates the failure of the search over the morphology space with learnable controllers during co-optimization, providing more evidence for the premature convergence of morphology as the main challenge of brain-body co-optimization.

While Cheney et al. [3] provides an embodied cognition perspective into this phenomenon, here we try to understand it from an optimization perspective. The standard intuition conceives a solution space where each solution has a single corresponding fitness creating a landscape over the solution space. However rugged it may be, we hope solutions with similar fitness values are grouped together in the solution space at some resolution, which allows us to search the space with heuristic search methods. This is certainly applicable to the brain-body co-optimization problem as well, where the brain and body together create the solution space. But instead of visualizing a standard landscape with

one fitness value per brain-body pair, we could alternatively conceptualize two solution domains: one a fitness landscape over controllers and another a fitness landscape over body plans. In doing so, we notice a peculiarity about brain-body solution space. While small changes in one part of the solution, that being the brain, can result in solutions with similar fitness and thus produce a typical and relatively smooth fitness landscape, changes to a robot’s body almost always result in a severe fitness drop [40]. In light of this insight, we advocate for a new morphology-centric way of conceptualizing the fitness landscape that the co-optimization process runs on. *Rather than imagining a fixed fitness value per each solution candidate, we can conceptualize the fitness landscape as a surface over the morphology space with a range of possible fitness values per body plan, where the actual fitness depends on the controller.*

This new view allows us to understand the premature convergence of morphology from a search perspective. Body plans that we found early in the co-optimization process accrue more controller optimization steps, climbing up to their maximum fitness value – an example of a first-mover advantage [31]. Conversely, body plans we found later on during the search via mutation, however promising these body plans may be, are at a disadvantage: they can’t compete with first-movers until a controller is optimized for them to provide a better assessment of their maximum fitness, but typically can’t survive long enough to accrue control optimization if they can’t compete with first movers initially.

Indeed, Cheney et al. [3] investigate this by designing a smaller search space for the control which reduces the range of possible fitness values per body plan, effectively limiting the first-movers advantage, and showing that the premature convergence of morphology disappears. We take this even further by using a fixed controller which collapses the range of possible fitness values per body plan into a single value, turning the fitness landscape into a familiar format that we can search effectively. The success of the fixed controller shows that our algorithms are adequate in searching the standard landscape, but confronted with the peculiar co-optimization landscape that is unforgiving to changes in the body plan, they fail and get stuck in local optima.

Luckily, this new understanding of the fitness landscape allows us to put recently proposed solutions into perspective: *if the search over the morphology space is failing due to our inability to effectively measure the true potential of body plans, then (1) we could give our algorithms more resources to assess new body plans [5], and/or (2) we could try to have a better initial assessment for new body plans by alleviating the performance decline during the morphological search by using different genetic representations for coordinated changes [54,57] or by using controllers that are robust to changes in the morphology [40].* These two ways are orthogonal to each other and we believe further research on these avenues holds the potential for unlocking successful brain-body co-optimization.

It also enables us to draw connections to other fields. For example, considering the optimization of the controller for a given morphology to be an approximation of the true potential (i.e. maximum) fitness of that morphology. Thus our framework connects brain-body co-optimization with a wealth of literature on surro-

gate proxy fitness models [19, 50], on approaches to optimization under uncertainty such as for the multi-arm bandit problems [1, 29]. In robot brain-body co-optimization the computation expense of fitness evaluations makes it unreasonable to train individual controllers from scratch to ideally and independently determine the fitness of each morphology. This is also the case for training the weights for a deep neural network architecture, and this framing draws analogies to the notion of sharing controllers for all possible body-plans/architectures in the fitness landscape, deemed supernet weight-sharing algorithms in the Neural Architecture Search literature [32, 45]. It also suggests the importance of methods to enable effective transfer of controllers from one morphology to another when sharing these fitness approximators across disparate body plans via approaches like meta-learning, transfer learning, and few-shot learning [12, 40, 60, 61].

Moreover, the success of body optimization with fixed controllers (Fig. 5, first vs last columns, all $P < 0.005$) brings the formulation of “brain” and “body” into question. Formulating the control of the robot in ways that it assigns a single or smaller range of fitness values unlocks the body plan optimization by collapsing the fitness landscape and limiting the first-movers advantage. Solutions where robots’ materials react to their environment in simpler ways and complex behavior emerges by exploiting the dynamics of the body and interaction with the environment, such as [4, 6, 8–10], can be further investigated. Especially, scaling this methodology to solve more complex tasks that would normally assumed to require complex closed-loop decision making should be investigated.

6 Conclusion

We have compared learnable neural network controllers and heuristically designed fixed controllers in the problem of co-optimization of brain and body (while for the latter, it is just the optimization of the body) for the locomotion task in multiple environments and morphology spaces that vary in complexity. Our experiments demonstrated the existence of high-performing regions that weren’t found during the co-optimization process. Based on our analysis with the fixed controller, we developed a new morphology-centric understanding of the fitness landscape which explains the premature convergence from a search perspective and clarifies previously proposed solutions. We hope this new framework will help us systematically tackle the challenges of brain-body co-optimization.

Acknowledgements. This material is based upon work supported by the National Science Foundation under Grant No. 2008413, 2239691, 2218063. Computations were performed on the Vermont Advanced Computing Core supported in part by NSF Award No. OAC-1827314.

References

1. Agrawal, S., Goyal, N.: Analysis of Thompson sampling for the multi-armed bandit problem. In: Conference on Learning Theory, p. 39-1. JMLR Workshop and Conference Proceedings (2012)

2. Bhatia, J., Jackson, H., Tian, Y., Xu, J., Matusik, W.: Evolution gym: a large-scale benchmark for evolving soft robots. *Adv. Neural. Inf. Process. Syst.* **34**, 2201–2214 (2021)
3. Cheney, N., Bongard, J., Sunspiral, V., Lipson, H.: On the difficulty of co-optimizing morphology and control in evolved virtual creatures. In: *Proceedings of the Artificial Life Conference 2016*, pp. 226–233. MIT Press, Cancun (2016). <https://doi.org/10.7551/978-0-262-33936-0-ch042>
4. Cheney, N., Bongard, J., Lipson, H.: Evolving soft robots in tight spaces. In: *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pp. 935–942. ACM, Madrid (2015). <https://doi.org/10.1145/2739480.2754662>
5. Cheney, N., Bongard, J., SunSpiral, V., Lipson, H.: Scalable co-optimization of morphology and control in embodied machines. *J. R. Soc. Interface* **15**(143), 20170937 (2018). <https://doi.org/10.1098/rsif.2017.0937>
6. Cheney, N., MacCurdy, R., Clune, J., Lipson, H.: Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding, p. 8 (2014)
7. Corucci, F., Cheney, N., Kriegman, S., Bongard, J., Laschi, C.: Evolutionary developmental soft robotics as a framework to study intelligence and adaptive behavior in animals and plants. *Front. Robot. AI* **4**, 34 (2017). <https://doi.org/10.3389/frobt.2017.00034>
8. Corucci, F., Cheney, N., Lipson, H., Laschi, C., Bongard, J.: Material properties affect evolutions ability to exploit morphological computation in growing soft-bodied creatures. In: *Proceedings of the Artificial Life Conference 2016*, pp. 234–241. MIT Press, Cancun (2016). <https://doi.org/10.7551/978-0-262-33936-0-ch043>
9. Corucci, F., Cheney, N., Lipson, H., Laschi, C., Bongard, J.C.: Evolving swimming soft-bodied creatures, p. 2 (2016)
10. Cheney, N., Clune, J., Lipson, H.: Evolved electrophysiological soft robots. In: *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pp. 222–229. The MIT Press (2014). <https://doi.org/10.7551/978-0-262-32621-6-ch037>
11. Ferigo, A., Iacca, G., Medvet, E., Pigozzi, F.: Evolving Hebbian learning rules in voxel-based soft robots. Preprint (2021). <https://doi.org/10.36227/techrxiv.17091218.v1>
12. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: *International Conference on Machine Learning*, pp. 1126–1135. PMLR (2017)
13. Hiller, J., Lipson, H.: Automatic design and manufacture of soft robots. *IEEE Trans. Rob.* **28**(2), 457–466 (2012). <https://doi.org/10.1109/TRO.2011.2172702>
14. Hiller, J., Lipson, H.: Dynamic simulation of soft multimaterial 3D-printed objects. *Soft Rob.* **1**(1), 88–101 (2014). <https://doi.org/10.1089/soro.2013.0010>
15. Hiller, J.D., Lipson, H.: Evolving amorphous robots. In: *Alife*, pp. 717–724. Citeseer (2010)
16. Horibe, K., Walker, K., Berg Palm, R., Sudhakaran, S., Risi, S.: Severe damage recovery in evolving soft robots through differentiable programming. *Genet. Program Evolvable Mach.* **23**(3), 405–426 (2022)
17. Horibe, K., Walker, K., Risi, S.: Regenerating soft robots through neural cellular automata. In: Hu, T., Lourenço, N., Medvet, E. (eds.) *EuroGP 2021. LNCS*, vol. 12691, pp. 36–50. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-72812-0_3
18. Huang, W., Mordatch, I., Pathak, D.: One policy to control them all: shared modular policies for agent-agnostic control. In: *International Conference on Machine Learning*, pp. 4455–4464. PMLR (2020)

19. Jin, Y.: Surrogate-assisted evolutionary computation: recent advances and future challenges. *Swarm Evol. Comput.* **1**(2), 61–70 (2011)
20. Joachimczak, M., Suzuki, R., Arita, T.: Artificial metamorphosis: evolutionary design of transforming, soft-bodied robots. *Artif. Life* **22**(3), 271–298 (2016). <https://doi.org/10.1162/ARTL.a.00207>
21. Kim, S., Laschi, C., Trimmer, B.: Soft robotics: a bioinspired evolution in robotics. *Trends Biotechnol.* **31**(5), 287–294 (2013)
22. Kriegman, S., Blackiston, D., Levin, M., Bongard, J.: A scalable pipeline for designing reconfigurable organisms. *Proc. Natl. Acad. Sci.* **117**(4), 1853–1859 (2020). <https://doi.org/10.1073/pnas.1910837117>
23. Kriegman, S., Blackiston, D., Levin, M., Bongard, J.: Kinematic self-replication in reconfigurable organisms. *Proc. Natl. Acad. Sci.* **118**(49), e2112672118 (2021). <https://doi.org/10.1073/pnas.2112672118>
24. Kriegman, S., Cheney, N., Bongard, J.: How morphological development can guide evolution. *Sci. Rep.* **8**(1), 13934 (2018). <https://doi.org/10.1038/s41598-018-31868-7>
25. Kriegman, S., Cheney, N., Corucci, F., Bongard, J.C.: Interoceptive robustness through environment-mediated morphological development. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 109–116. ACM, Kyoto (2018). <https://doi.org/10.1145/3205455.3205529>
26. Kriegman, S., et al.: Scale invariant robot behavior with fractals. In: *Robotics: Science and Systems XVII. Robotics: Science and Systems Foundation* (2021). <https://doi.org/10.15607/RSS.2021.XVII.059>
27. Kriegman, S., et al.: Scalable sim-to-real transfer of soft robot designs. In: *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, pp. 359–366. IEEE (2020)
28. Kriegman, S., Walker, S., Shah, D.S., Levin, M., Kramer-Bottiglio, R., Bongard, J.: Automated shapeshifting for function recovery in damaged robots. In: *Robotics: Science and Systems XV. Robotics: Science and Systems Foundation* (2019). <https://doi.org/10.15607/RSS.2019.XV.028>
29. Kuleshov, V., Precup, D.: Algorithms for multi-armed bandit problems. *arXiv preprint arXiv:1402.6028* (2014)
30. Lehman, J., Stanley, K.O.: Evolving a diversity of virtual creatures through novelty search and local competition. In: *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation - GECCO 2011*, p. 211. ACM Press, Dublin (2011). <https://doi.org/10.1145/2001576.2001606>
31. Lieberman, M.B., Montgomery, D.B.: First-mover advantages. *Strategic Manage. J.* **9**(S1), 41–58 (1988)
32. Liu, H., Simonyan, K., Yang, Y.: DARTS: differentiable architecture search. In: *International Conference on Learning Representations* (2019). <https://openreview.net/forum?id=S1eYHoC5FX>
33. Liu, S., Matthews, D., Kriegman, S., Bongard, J.: Voxcraft-sim, a GPU-accelerated voxel-based physics engine (2020)
34. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**(11) (2008)
35. Marzougui, D., Biondina, M.: A comparative analysis on genome pleiotropy for evolved soft robots, p. 4 (2022)
36. Medvet, E., Bartoli, A., De Lorenzo, A., Fidel, G.: Evolution of distributed neural controllers for voxel-based soft robots. In: *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, pp. 112–120. ACM, Cancún (2020). <https://doi.org/10.1145/3377930.3390173>

37. Medvet, E., Bartoli, A., De Lorenzo, A., Seriani, S.: 2D-VSR-sim: a simulation tool for the optimization of 2-D voxel-based soft robots. *SoftwareX* **12**, 100573 (2020). <https://doi.org/10.1016/j.softx.2020.100573>
38. Medvet, E., Bartoli, A., Pigozzi, F., Rochelli, M.: Biodiversity in evolved voxel-based soft robots. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 129–137. ACM, Lille (2021). <https://doi.org/10.1145/3449639.3459315>
39. Medvet, E., Nadizar, G., Pigozzi, F.: On the impact of body material properties on neuroevolution for embodied agents: the case of voxel-based soft robots, p. 9 (2022)
40. Mertan, A., Cheney, N.: Modular controllers facilitate the co-optimization of morphology and control in soft robots. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 174–183. ACM, Lisbon (2023). <https://doi.org/10.1145/3583131.3590416>
41. Pfeifer, R., Bongard, J.: *How the body shapes the way we think: a new view of intelligence* (2006)
42. Pfeifer, R., Iida, F.: *Morphological computation: connecting body, brain, and environment*, p. 5 (2009)
43. Pigozzi, F., Tang, Y., Medvet, E., Ha, D.: Evolving modular soft robots without explicit inter-module communication using local self-attention. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 148–157 (2022)
44. Pontes-Filho, S., Walker, K., Najarro, E., Nichele, S., Risi, S.: A single neural cellular automaton for body-brain co-evolution, p. 4 (2022)
45. Ren, P., et al.: A comprehensive survey of neural architecture search: challenges and solutions. *ACM Comput. Surv.* **54**(4) (2021). <https://doi.org/10.1145/3447582>
46. Rus, D., Tolley, M.T.: Design, fabrication and control of soft robots. *Nature* **521**(7553), 467–475 (2015)
47. Schmidt, M.D., Lipson, H.: Age-fitness pareto optimization, p. 2 (2010)
48. Shah, D.S., Powers, J.P., Tilton, L.G., Kriegman, S., Bongard, J., Kramer-Bottiglio, R.: A soft robot that adapts to environments through shape change. *Nat. Mach. Intell.* **3**(1), 51–59 (2021)
49. Shepherd, R.F., et al.: Multigait soft robot. *Proc. Natl. Acad. Sci.* **108**(51), 20400–20403 (2011)
50. Shi, L., Rasheed, K.: A survey of fitness approximation methods applied in evolutionary algorithms. In: Tenne, Y., Goh, C.-K. (eds.) *Computational Intelligence in Expensive Optimization Problems*. ALO, vol. 2, pp. 3–28. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-10701-6_1
51. Sims, K.: Evolving virtual creatures. In: *Proceedings of the 21st annual conference on Computer graphics and interactive techniques - SIGGRAPH 1994*, pp. 15–22. ACM Press, Not Known (1994). <https://doi.org/10.1145/192161.192167>
52. Talamini, J., Medvet, E., Bartoli, A., Lorenzo, A.D.: Evolutionary synthesis of sensing controllers for voxel-based soft robots, p. 8 (2019)
53. Talamini, J., Medvet, E., Nichele, S.: Criticality-driven evolution of adaptable morphologies of voxel-based soft-robots. *Front. Robot. AI* **8** (2021). <https://www.frontiersin.org/article/10.3389/frobt.2021.673156>
54. Tanaka, F., Aranha, C.: Co-evolving morphology and control of soft robots using a single genome. In: *2022 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1235–1242. IEEE (2022)
55. Trimmer, B.A.: New challenges in biorobotics: incorporating soft tissue into control systems. *Appl. Bionics Biomech.* **5**(3), 119–126 (2008). <https://doi.org/10.1080/11762320802617255>

56. Veenstra, F., Glette, K.: How different encodings affect performance and diversification when evolving the morphology and control of 2D virtual creatures. In: The 2020 Conference on Artificial Life, pp. 592–601. MIT Press (2020). https://doi.org/10.1162/isal.a_00295
57. Veenstra, F., Olsen, M.H., Glette, K.: Effects of encodings and quality-diversity on evolving 2D virtual creatures, p. 4 (2022)
58. Wilcoxon, F., Wilcox, R.A.: Some rapid approximate statistical procedures. (No Title) (1964)
59. Yim, M., et al.: Modular self-reconfigurable robot systems [grand challenges of robotics]. *IEEE Robot. Autom. Mag.* **14**(1), 43–52 (2007). <https://doi.org/10.1109/MRA.2007.339623>
60. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? In: *Adv. Neural Inf. Process. Syst.* **27** (2014)
61. Zhao, Y., Wang, L., Tian, Y., Fonseca, R., Guo, T.: Few-shot neural architecture search. In: *International Conference on Machine Learning*, pp. 12707–12718. PMLR (2021)