



Towards Multi-Morphology Controllers with Diversity and Knowledge Distillation

Alican Mertan
 alican.mertan@uvm.edu
 Neurobotics Lab
 University of Vermont
 Burlington, VT, USA

Nick Cheney
 ncheney@uvm.edu
 Neurobotics Lab
 University of Vermont
 Burlington, VT, USA

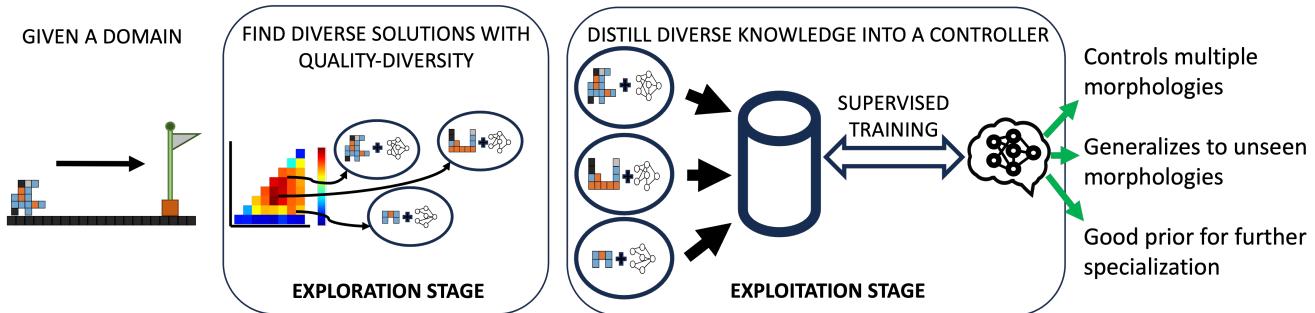


Figure 1: Proposed pipeline for training a single controller capable of controlling multiple morphologies. Given an environment and a task, Quality Diversity algorithms explore the domain and discover effective morphologies and their controllers. These controllers are used as teacher controllers and their knowledge is distilled into a single controller. The distilled controller successfully controls the morphologies of the teacher controllers. Moreover, the distilled controller generalizes well to unseen morphologies in a zero-shot manner and provides a better prior for further specialization on unseen morphologies and tasks.

ABSTRACT

Finding controllers that perform well across multiple morphologies is an important milestone for large-scale robotics, in line with recent advances via foundation models in other areas of machine learning. However, the challenges of learning a single controller to control multiple morphologies make the ‘one robot one task’ paradigm dominant in the field. To alleviate these challenges, we present a pipeline that: (1) leverages Quality Diversity algorithms like MAP-Elites to create a dataset of many single-task/single-morphology teacher controllers, then (2) distills those diverse controllers into a single multi-morphology controller that performs well across many different body plans by mimicking the sensory-action patterns of the teacher controllers via supervised learning. The distilled controller scales well with the number of teachers/morphologies and shows emergent properties. It generalizes to unseen morphologies in a zero-shot manner, providing robustness to morphological perturbations and instant damage recovery. Lastly, the distilled controller is also independent of the teacher controllers – we can distill the teacher’s knowledge into any controller model, making our approach synergistic with architectural improvements and existing training algorithms for teacher controllers.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.
 GECCO '24, July 14–18, 2024, Melbourne, VIC, Australia
 © 2024 Copyright held by the owner/author(s).
 ACM ISBN 979-8-4007-0494-9/24/07.
<https://doi.org/10.1145/3638529.3654013>

CCS CONCEPTS

- Computing methodologies → Evolutionary robotics; Reinforcement learning;
- Theory of computation → Evolutionary algorithms.

KEYWORDS

evolutionary robotics, soft robotics, brain-body co-optimization

ACM Reference Format:

Alican Mertan and Nick Cheney. 2024. Towards Multi-Morphology Controllers with Diversity and Knowledge Distillation. In *Genetic and Evolutionary Computation Conference (GECCO '24)*, July 14–18, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3638529.3654013>

1 INTRODUCTION

Finding controllers that perform well across different morphologies is an important milestone for large-scale robotics. Similar to the ‘foundation models’ that enable progress in other areas of machine learning, such as computer vision or language processing, a foundational multi-morphology controller can facilitate progress in robotics by enabling fine-tuning to downstream tasks with a smaller amount of data (which is important because the best methods for training control models such as evolutionary or reinforcement learning algorithms are data-inefficient). Moreover, models capable of controlling a multitude of robots can enable better brain-body co-optimization by being a good fitness estimator for unseen morphologies [28], hence better specialization and performance can be obtained for a given domain by exploiting multi-morphology

controllers. However, obtaining such controllers, despite being of great interest, remains an open problem. As identified in [13], the field is mostly stuck in the ‘one robot one task’ paradigm – training a new robot and controller from scratch for each new task such as locomotion, object manipulation, climbing, etc.

The challenges that drive the field into the ‘one robot one task’ paradigm are twofold. The first issue is the incompatible control – the differences in action and observation spaces of different robots or tasks [20]. While it is straightforward to have ‘compatible’ controllers through caching trick [16], having wide enough input (and output) layers to accommodate the maximum possible number of sensory inputs (and motor outputs), while simply ignoring unused nodes for morphologies with a smaller number of inputs (outputs), it turns out that such controllers are hard to train to control multiple morphologies, whether reinforcement learning [16] or evolutionary optimization [27] algorithms are used. This hardship of training stems from the complexity of learning many tasks/morphologies.

These two challenges, incompatible control and hardship of training, are intertwined as the controller model one chooses to use as a multi-morphology controller determines the loss landscape and in return, affects its training dynamics. Previous work on obtaining multi-morphology controllers mostly focuses on complexifying controller models by enforcing modularity (aiming for emergent higher-scale control at the robot level [6, 16, 30, 33]) or by employing recent developments in Graph Neural Networks and Transformers (utilizing their ability to deal with arbitrary-sized inputs [13, 20, 47]). These attempts not only solve the incompatibility problem but also, presumably, shapes the loss landscape and offer easier training.

In this work, however, we focus solely on the hardship of training multi-morphology controllers, instead of addressing it indirectly by architectural changes, which makes our approach synergistic with existing work. Our proposed method stems from our search for a ‘simple’ procedure for learning a multi-morphology controller. Inspired by existing literature on knowledge distillation [3, 9, 14, 21, 25, 32, 39, 46] – learning to match input-output patterns of teacher models via supervised learning, we investigate the use of response-based, offline, multi-teacher policy distillation to learn a single multi-morphology controller. We collect a dataset of *(observation, action)* pairs from teacher controllers optimized to control single morphology and then employ supervised learning to train a single controller on this collected dataset.

Indeed, we show that knowledge distillation results in controllers that match the performance of teacher controllers on many different morphologies, without needing any complex architecture or message-passing scheme – just a simple neural network with a single hidden layer that is made compatible with the caching trick can learn to control hundreds of robots for the locomotion task in a matter of minutes on a consumer laptop. We also note that our method is agnostic to the choice of the distilled controller. Once we collect a dataset of *(observation, action)* pairs, it can be used to train any neural network, including the modular or graph-based models developed to alleviate the challenges of incompatible control.

It might seem, however, that optimizing many teacher controllers that specialize in particular morphologies to be able to train a single multi-morphology controller defeats the purpose. For this, we would like to show two important points. First, the teacher controllers are only used to collect *(observation, action)* pairs and then

discarded. They can be as simple as possible for faster optimization on a particular morphology. Moreover, one can use abundant publicly available pre-trained controllers instead of training them from scratch. More importantly, we show that the distilled controller displays emergent properties. Our investigations show that the distilled controller is robust to morphological perturbations and generalizes to unseen morphologies in a zero-shot manner. Moreover, it can be used as a prior for further specialization on unseen morphologies or tasks to speed up the adaptation process. These properties justify the cost of training teacher controllers – we gain more than the sum of teacher controllers by distilling them into a single controller. Most of all, our proposed approach is orthogonal to the previous work developing complex architectures for multi-morphology control – we could use such models to further enhance the emergent properties of the distilled controllers.

It is not obvious, nonetheless, how to select teacher controllers’ morphologies. Indeed, the choice of morphologies to train the multi-morphology controller is an important question that has been ignored in the literature so far. Existing works either heuristically [6, 16, 20, 47] choose the morphologies to train with, or assume that effective morphologies are known prior to the training [13]. Considering that, unlike most prior work, we need experiences *((observation, action)* pairs) of effective morphologies as well to distill that knowledge into a single controller, we resort to Quality Diversity (QD) algorithms [37, 38] for finding effective morphologies with their optimized controllers, similar to [9, 21] where QD is used to create a repertoire of behaviors for a single robot and knowledge distillation is used to distill that knowledge into a single controller. In our case, given that the feature descriptors are based on morphological attributes, a QD algorithm is used to explore the morphology space and optimize controllers for a variety of high-performing morphologies with different trade-offs in their morphological attributes for a given domain. We show that these controllers can be used as teachers to distill into a single controller that is capable of controlling a variety of distinct morphologies that perform well for the given domain, and that generalizes well to unseen morphologies in a zero-shot manner.

Overall, we present the two-stage pipeline illustrated in Fig. 1 that first explores a given domain with QD algorithms to find high-performing morphologies and their respective controllers, and then exploits this knowledge by distilling it into a single controller. The main contributions of this work are to:

- demonstrate the effectiveness of knowledge distillation for multi-morphology controller training – a simple procedure that can train simple models to control multiple morphologies. (Sec. 3)
- propose the use of QD algorithms for the automated discovery of effective morphologies with their optimized controllers for the distillation process. (Sec. 4)
- investigate the capabilities of the distillation process and distilled controllers (Sec. 5) and show that
 - the distillation process is controller agnostic, making it synergistic with existing work that develops complex controller models. (Sec. 5.1)
 - the distilled controllers scale well with the number of teachers. (Sec. 5.2)

- the distilled controllers generalize well to unseen morphologies, justifying the cost of obtaining teacher controllers. (Sec. 5.3)
- the distilled controllers provide a better prior for further specialization, unlocking transfer learning opportunities in the robotics field. (Sec. 5.4)

2 METHODS

Simulation. We use Evolution Gym version 1.0.0 (Evogym) [1] for simulating 2D voxel-based soft bodied robots. The raw observations that the simulation provides are processed to acquire voxel volume ($\in \mathbb{R}$), speed ($\in \mathbb{R}^2$), and material type (one-hot encoded vector v of length 5). We also process the timesteps into a saw wave-shaped time signal by applying $\text{mod } 25$ to it. While the simulation engine is deterministic, we apply a small noise to the observations sampled from $\mathcal{N}(0, 0.01)$ to model sensory noise.

Tasks and performance evaluation. We mainly experiment in the Walker-v0 environment which consists of a flat surface of length 100 in voxels with a locomotion task where robots try to reach the end of the surface. We also experiment with a similar locomotion task that consists of a soft, dynamic surface called BridgeWalker-v0. In both environments, we use the modified reward function from [27, 28] which encourages the robot to finish the tasks as fast as possible. Due to the stochasticity we injected into the system in terms of observation and action noises, we repeat each fitness evaluation multiple times and take the average. Unless otherwise noted, each simulation is repeated 5 times.

Robot representation and control. Robots in Evogym consist of 4 types of materials. There are two materials under active control, one that expands horizontally and one that expands vertically. There are also two passive materials, rigid and elastic, that are under the effects of forces created by active materials and dynamics of the environment. Robots are directly represented as a matrix $R \in T_{H \times W}$ where $T \in \{0, 1, 2, 3, 4\}$ encodes the materials (or lack thereof), and we use $(H, W) = (5, 5)$, following the practice of limiting the robot design space [4, 5, 7, 18, 22, 24, 27, 28, 45]. Robots are controlled by specifying a scalar action ($\in [0.6, 1.6]$) that is used to determine the target length by multiplying the action with the resting length. Controllers are queried every 5th timestep and the last action is repeated for the remaining timesteps to prevent high-frequency dynamics. We apply a small noise to the actions sampled from $\mathcal{N}(0, 0.01)$ to model actuator noise.

Controller models. Throughout the work, we experiment with 3 different controllers – ‘Global FC’, ‘Global Tx’, and ‘Modular FC’, modeled by different neural network architectures – ‘FC’ stands for fully connected and ‘Tx’ stands for transformer, and belonging 2 different control paradigms – global indicates a centralized controller where observations from all voxels are concatenated and consumed at once by the controller to output actions for each voxel, similar to the ones used in [10, 23, 26, 27, 43, 44], and modular indicates a shared, decentralized controller that observes a local neighborhood and output action for a single voxel, similar to the

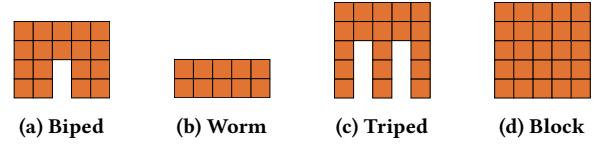


Figure 2: Experimented morphologies to show the effectiveness of knowledge distillation for training multi-morphology controller.

ones used in [16, 23, 24, 27, 28, 35]. All controllers are made compatible through the use of caching trick and are not conditioned on the morphology explicitly.¹

Optimization algorithms. To show the ineffectiveness of joint training on fixed morphologies, we use the reinforcement learning algorithm Twin Delayed Deep Deterministic policy gradients algorithm (TD3) [12] (starting from the clean implementation of [15]) and the evolutionary optimization algorithm Age-Fitness Pareto Optimization (AFPO) [40] (with population size of 16). After demonstrating the ineffectiveness of both approaches, we introduce our approach where we use the Map-Elites algorithm [29] as a Quality Diversity algorithm to create an archive of diverse and effective morphologies with their optimized controllers where the two feature descriptors are the number of existing voxels and the number of active voxels. In each generation of the Map-Elites algorithm, we create 16 new solutions from 16 randomly chosen solutions from the map. The offspring are created by mutation only. Following [1, 27, 28], robot representations are mutated by changing the material of each voxel with a 10% probability and controllers are mutated by adding a noise sampled from $\mathcal{N}(0, 0.1)$. The new solutions are created by mutating either the morphology or the controller, chosen with a 50% probability, as in [5, 27, 28].

Statistical testing To measure the statistical significance for distilled controllers, we calculate the relative performance of the distilled controllers compared to teacher controllers and apply one-sample t-test [41] with the null hypothesis of the population mean of 1. Where we compare two samples, we use the Wilcoxon Rank Sum test [48]. All comparisons are done with the $p = .05$ threshold.

3 KNOWLEDGE DISTILLATION FOR MULTI-MORPHOLOGY CONTROL

Optimizing controllers capable of controlling multiple different morphologies is a challenging problem [13, 16, 27]. Here, we provide a naive attempt (referred to as joint training) at learning a controller (Global FC) for controlling the four predetermined and fixed morphologies shown in Fig. 2, both with evolutionary optimization (AFPO [40]) and reinforcement learning (TD3 [12]). Following [27, 36], we use the minimum performance among all morphologies as the fitness for the evolutionary algorithm to help avoid specialization to a subset of morphologies while ignoring others.

Fig. 3 demonstrates the joint training trajectories with both algorithms, as well as isolated training for each morphology individually. Both algorithms struggle to optimize a multi-morphology controller during joint training. While the evolutionary algorithm converges

¹Details of controller architectures can be found in our code repository: mertan-a/towards-multi-morphology-controllers

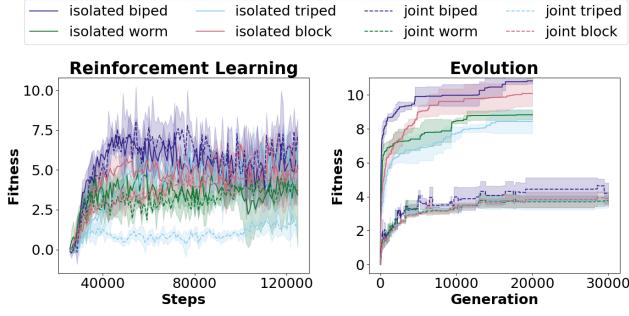


Figure 3: Training trajectories of isolated training on each morphology individually (solid lines) vs. the performance of each morphology during joint training on all (dashed lines). Lines show the mean values and the shaded areas show the standard errors, calculated over 3 repetitions. Reinforcement Learning (left) can find solutions that work well for multiple morphologies but ignore others. Evolutionary algorithms (right) find solutions that perform similarly on all morphologies, but they exhibit sub-optimal performance.

to a sub-optimal but similar-performing solution for all morphologies, reinforcement learning finds solutions that work better for some morphologies, ignoring others. Yet both algorithms fail to find controllers that match the performance of single-morphology controllers each trained in isolation.

To overcome the challenging training of multi-morphology controllers, here we propose the use of response-based, offline, multi-teacher knowledge distillation [3, 14], also known as policy distillation in reinforcement learning [32, 39], for the training of a single general controller capable of controlling multiple robots with different morphologies. We evolve single-morphology controllers for our experimented morphologies in Fig. 2 and use their experiences to create a dataset of $(\text{observation}, \text{action})$ pairs. This dataset is then used for supervised learning to train a single multi-morphology controller to minimize the error in recreating the correct action for a given observation regardless of which robot the action is sampled from. Once we have the dataset, the training of the multi-morphology controller can be done in an offline supervised fashion. In our experiments, we model the controllers by neural networks and use the gradient descent algorithm Adam [17] to train them.

To demonstrate the effectiveness of this approach, we replay the best controllers found by the evolutionary algorithm during isolated training of each morphology at each repetition, $C_m^r, m \in \{\text{Biped, Worm, Triped, Block}\}, r \in \{1, 2, 3\}$, 100 times and collect the $(\text{observation}, \text{action})$ pairs into 81 (3^4 , all possible combinations of champions for each morphology) datasets $D^i, i \in [1..81]$. Using these datasets, we distill 81 multi-morphology controllers, all with the same controller model as the teacher controllers – Global FC, by training the controllers on the datasets for 100,000 steps with mini-batches of size 128 and learning rate of 0.001, in a supervised manner where the loss is the mean squared error between the estimated and ground truth actions. All hyperparameters are chosen heuristically.

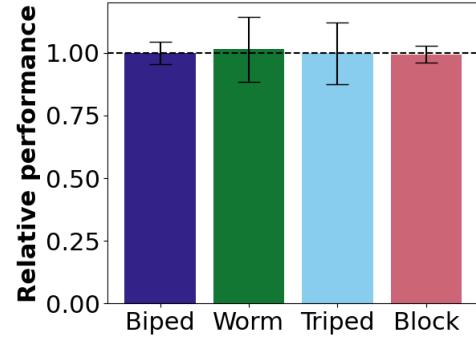


Figure 4: Performance of the multi-morphology controller relative to teacher single-morphology controllers for each experimented morphology ($\text{mean} \pm \text{SE}$). Here, and in all figures, the dotted line marks equal performance. Knowledge distillation can successfully train a single controller to control multiple morphologies as well as controllers specifically optimized for individual morphologies.

Fig. 4 shows the performance of the distilled multi-morphology controllers relative to the corresponding single-morphology controllers, averaging over 81 cases. Knowledge distillation from single-morphology controllers results in multi-morphology controllers that achieve near-perfect teacher-level performance. Moreover, we see that in some cases the distilled multi-morphology controller outperforms the teacher single-morphology controller, as indicated by the error bars. These results demonstrate the effectiveness of knowledge distillation for the training of compatible controllers, without resorting to complex architectures or training schemes.

4 QUALITY DIVERSITY FOR DOMAIN EXPLORATION

In the previous section, we experiment with heuristically chosen fixed morphologies and show that knowledge distillation can be utilized successfully to train a single controller capable of controlling multiple robot morphologies, without the need for any specialized architecture. Now the question is, how should we choose which morphologies to use for the training of the teacher single-morphology controllers, given a domain – an environment and a task?

As opposed to using heuristically chosen morphologies [16, 20, 30], we propose to utilize the QD algorithms [37, 38] to evolve distinct high-performing solutions representing trade-offs in a feature space, exploring the solution space for a given domain. Defining the feature descriptors based on the morphology of the robot, we can utilize QD algorithms to evolve different robot-controller pairs with varying morphologies, optimized for a particular environment and task. We prefer the QD algorithm over alternatives, as it allows finer control over the diversity of morphologies by explicitly describing feature descriptors and bins, and results in greater diversity [31]. Knowledge distillation can then be applied to create a single controller capable of controlling a variety of different robots, exhibiting behavior optimized for the domain.

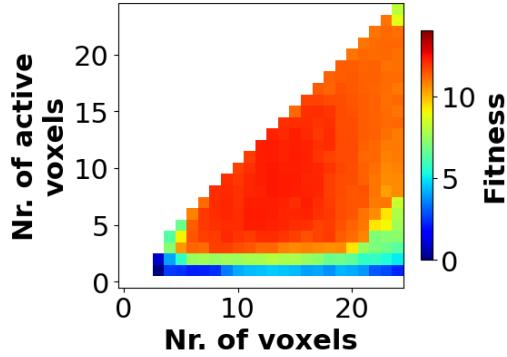


Figure 5: An example map produced by the MAP-Elites algorithm. Each cell corresponds to a robot-controller pair with its fitness shown as the color of the cell. The X-axis differentiates bins in the map by the number of total voxels present in the robot, while the y-axis stratifies robots by their number of active voxels. MAP-Elites successfully evolves a variety of high-performing robots.

In particular, we experiment with the MAP-Elites algorithm [29], similar to [31], where the feature descriptors for robots are the number of existing voxels and the number of active voxels. Fig. 5 shows the map produced by the MAP-Elites algorithm after 20,000 generations of evolution. We see that the MAP-Elites algorithm explores the morphology space for the given task and evolves a map with 296 wide-ranging unique morphologies that can locomote as effectively as the ones in similar works [27, 28].

Having produced this map for a given task or domain, we start experimenting with knowledge distillation to create a multi-morphology controller. We experiment with two heuristically chosen criteria for the selection of teacher robots: fitness and morphology.

First, we experiment with selecting different robots as teachers based on their fitness values. We order the solutions by their fitness values and distill a multi-morphology controller for the top 10% of solutions, which are shown in Fig. 6 (left). The distilled multi-morphology controller successfully controls 29 slightly different morphologies representing different trade-offs between the number of voxels and the number of active voxels, and its performance is statistically indistinguishable (at $p = 0.05$ -level) from the single-morphology controllers as shown in Fig. 6 (right).

To test the whether distillation process can result in controllers capable of controlling maximally different morphologies, we choose 35 individuals from the map that are spread across the feature space, as shown in Fig. 7 (left). The performance of the distilled controller compared to teacher controllers on each experimented morphology can be seen in Fig. 7 (right). While the controller is capable of matching the performance of the teacher controllers in 28 out of 35 morphologies ($p > .32$), it performs worse than the teachers in 5 cases and exceeds the teacher performance in 2 cases. Surprisingly, the cases where the distilled controller fails to achieve teacher-level performance occur when the teachers' performances are low (color represents the fitness of the original controller).

²See the distilled controller in action.

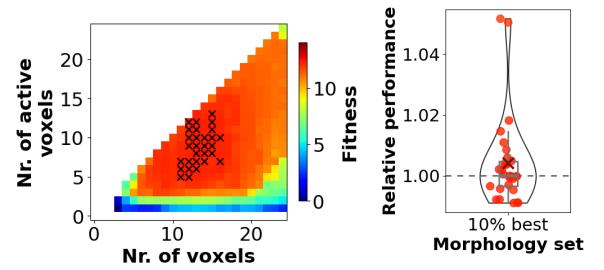


Figure 6: (left) Top 10% of individuals (29 in total) used as teachers for distilling a multi-morphology controller, marked with an x. (right) Performance of the distilled multi-morphology controller on each trained morphology compared to their original teacher controllers across 10 runs with noise. Each data point is plotted and its color represents the fitness of the original controller. The mean point is labeled with an x. The distilled controller achieves almost perfect performance, matching the performances of teacher single-morphology controllers.

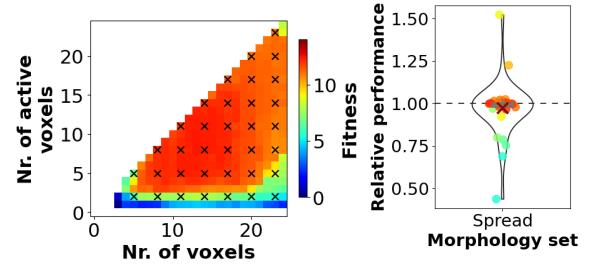


Figure 7: (left) 35 individuals maximally spread across the feature space, marked with an x. (right) Performance of the distilled multi-morphology controller on each trained morphology compared to their original teacher controllers across 10 runs with noise. The distilled controller matches the performance of teacher controllers in most cases. The cases where there is a performance drop occur where the teacher controller does not perform well on the morphology.²

Overall, our experimentation shows that QD algorithms can be utilized to discover effective morphologies, as well as controllers optimized to control them. Subsequently, the knowledge that resides in the map can be distilled into a single controller, resulting in a controller capable of controlling a diverse set of morphologies.

5 INVESTIGATING ABILITIES OF DISTILLED CONTROLLERS

5.1 Model Independence

So far we have experimented with the same controller model for the distilled controller as the teacher controller – Global FC. However, once we collect the dataset of $\{observation, action\}$ pairs, we can train any controller architecture to do the mapping. It allows us to use deeper or more complex architectures only for the distillation

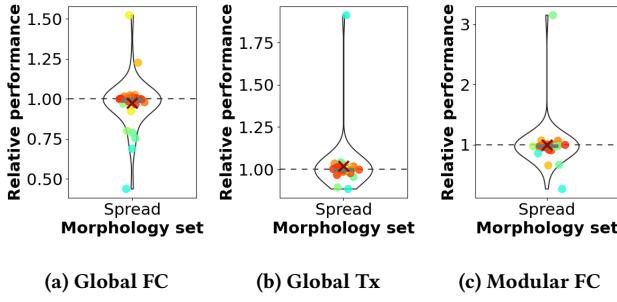


Figure 8: The same dataset acquired from individuals that are spread across the feature space distilled into three different controllers. (left) The same controller as the teacher controller achieves the same level of performance as the teachers for most of the cases. (middle) We distill into a more capable transformer controller and show that it is capable of achieving teacher-level performance in all cases, even surpassing teacher performance in one case. (right) We process the dataset to change the global observations to local observations and distill a modular controller that achieves a similar multi-morphology performance as the global controller. Please note the varied y-axis scales across subfigures.

part of the process where we use supervised learning while using simpler architectures that can be effectively trained with evolution during the QD process. Moreover, the distilled controller is not limited to the same control paradigm as the teachers either. For instance, the dataset can be pre-processed to turn global observations into local observations to distill a shared decentralized/modular controller (similar to the ones used in [16, 23, 24, 27, 28, 35]) while the teachers are centralized/global controllers (similar to the ones used in [10, 23, 26, 27, 43, 44]).

To display controller model independence, we use the 35 individuals that are maximally spread to the feature space, shown in Fig. 7 (left), as teachers and distill a transformer-based global controller ("Global Tx") as well as a fully-connected neural network based modular controller ("Modular FC"). All controllers achieve near-teacher-level performance on almost all morphologies, as shown in Fig. 8. Moreover, their performances show variation, indicating that the type of controller affects the distillation performance. This creates an opportunity for incorporating existing work that designs compatible multi-morphology architectures [6, 13, 16, 20, 47] into our approach to improve multi-morphology performance further.

5.2 Scaling

To test how many different teacher individuals can be distilled into a single multi-morphology controller, we experiment with using increasing numbers of individuals from the map as teachers. Specifically, we experiment with using the top 10, 40, 75, 100% of individuals in the order of their fitness as teachers, which results in 29, 118, 222, and 296 teachers, respectively.

When we distill into all of the introduced multi-morphology controllers (Global FC, Global Tx, Modular FC), we see that all of the distilled controllers achieve performances closer to teachers, as

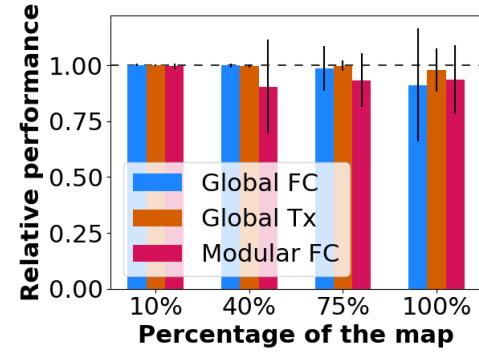


Figure 9: Performance of the distilled controllers with an increasing number of teachers with different morphologies. While different distilled controllers scale differently, they all achieve similar performances compared to teachers.

shown in Fig. 9. As the number of teachers increases, the performance for the distilled controllers tends to slightly decrease and the error bars grow. However, we see that they show different scaling behaviors, demonstrating the possibility that a bigger or more complex controller architecture can achieve better performance. While we experiment with heuristically chosen architectures as a proof of concept, one can treat the architecture as a hyperparameter and optimize it for a particular domain.

5.3 Generalization to Unseen Morphologies

We have shown above that (1) we can evolve a diverse set of morphologies with their respective controllers for the locomotion task optimized for different trade-offs in the feature space, and (2) these individuals can be used as teachers to distill into a single multi-morphology controller. The distilled controller is capable of controlling a diverse set of morphologies with near teacher-level performance and scaling up to the full map that covers 296 distinct morphologies. Here, we examine whether distilled controllers can generalize to unseen morphologies.

To test the generalization of the distilled controller to unseen morphologies, we choose two random individuals from the map, shown in Fig. 10 (left), and create a list of morphologies that transition from the morphology of one individual to the other, by changing one voxel at a time. We discard any morphology if it is not connected or if it exists in the map. This list of morphologies contains a number of unseen morphologies between the chosen individuals. We also note that this process can create unseen morphologies that are missing a number of voxels from seen morphologies and can be considered amputation scenarios. Therefore testing controllers on these morphologies also provides a way for us to test the damage recovery abilities of the controllers. As opposed to methods like [2, 8, 19] where the system spends time to figure out how to recover from damage, we would be measuring instant (i.e. zero-shot) damage recovery through the robustness of the controller.

To this end, we test distilled controllers (trained on the full map) on this set, as well as the single controller from the map was trained on the body plan closest to new morphologies. Fig. 10 (right) shows

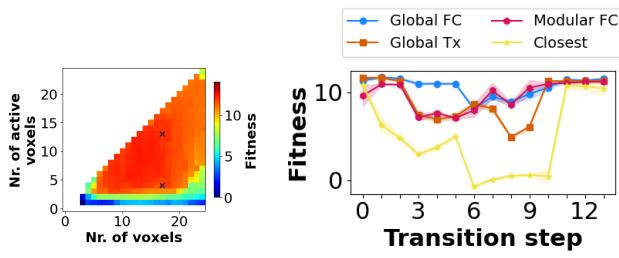


Figure 10: Using the pair of individuals marked on the map (left), we create a list of unseen morphologies in between the two morphologies of the selected individuals and measure the performance of the distilled controllers (right). As a baseline, we also measure the performance of the controller with the closest morphology from the map. Solid lines show the mean values of 10 fitness evaluations and shady regions show the standard errors. Distilled controllers generalize well to unseen morphologies, never performing worse than the closest controller from the map and outperforming them in most cases (transition steps 1 to 10).

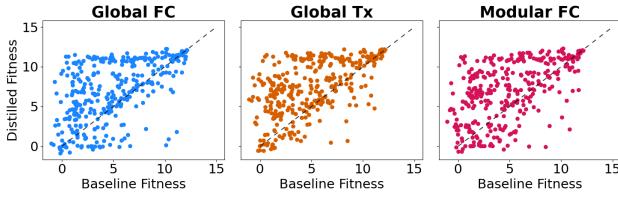


Figure 11: Performances of the distilled controllers plotted against the baseline controllers on 357 unseen morphologies. Each data point is plotted. The distilled controllers outperform the baseline in most cases (all $p < .001$), showing the emergent ability to instantly adapt to unseen morphologies.

a typical example of this process. While the performances of the closest controllers from the map drop as the morphology changes, distilled controllers perform better than the baseline. Moreover, their performance differs based on the controller, demonstrating the possibility that a specialized controller architecture or control strategy can be even more effective in its generalization capabilities.

We repeat this process with 30 randomly chosen individual pairs and report the performances of distilled and baseline controllers on each unseen morphology (357 in total) in Fig. 11. Points above the dotted line are the cases where the distilled controllers outperformed the controller of the most similar morphology from the map (the number of such cases is 250 for the Global FC, 267 for Global Tx, and 246 for Modular FC). Overall, the distilled controllers perform approximately 1.5 times better than the closest controller on average, outperforming the baseline (all $p < .001$ with the alternative hypothesis that the sample distribution mean is greater than 1) and showing ability to adapt to unseen morphologies instantly.

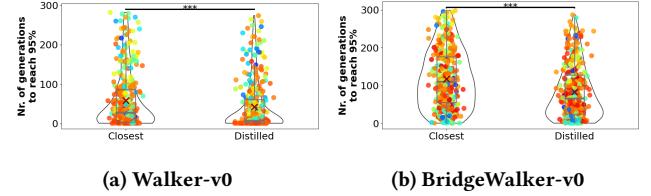


Figure 12: Number of generations to achieve 95% of the end performance for the distilled controller and the baseline controller, lower is better. The distilled controller converges slightly faster compared to the baseline controllers. This trend also held in the time to reach 90% or 99% of the end performance (all $p < .01$).

5.4 Rapid Finetuning

The above results show impressive instant generalization to new morphologies interpolated within the map. But what about adaptation to these morphologies via further controller optimization? We are interested in testing how good of a prior the distilled controllers are for further specialization on unseen morphologies or on different tasks. Ideally, we expect the distilled controller to perform better as a starting point for further specialization, allowing us to have foundational models that enable rapid adaptation to downstream tasks, whether they are unseen morphologies or different objectives (akin to initializing an image classifier with ResNets pre-trained on ImageNet). This can help mitigate the cost of training for many robotics applications and accelerate the development in the field.

To see how good of a prior the distilled controller is for further specialization on unseen morphologies and tasks, we run 300 generations of AFPO [40] where all individuals in the starting population are initialized with the Global FC multi-morphology controller distilled from the full map. For each unseen morphology that we sampled in the previous experiment (357 in total), we fix the morphology and finetune the distilled controller on that morphology performing the original Walker-v0 task. To assess wider generalization, in a second condition, we also take that same morphology and controller originally learned on the Walker-v0 task and finetune it on the BridgeWalker-v0 environment (a slightly different locomotion task). As a baseline, for each unseen morphology, we find the controller of the robot with the most similar morphology from the map and initialize the evolutionary run with that controller instead of the distilled controller. We note that both the distilled controller and the baseline controllers have the same architecture.

While the performances of the run champions do not show any statistically significant differences between the baseline and the distilled controller in both environments, we see that runs with the distilled controller achieve 90%, 95%, and 99% of their respective final performances faster compared to runs with the baseline controllers (Fig. 12). These results demonstrate that the distilled controller provides a better starting point for further specialization, enabling rapid finetuning to downstream tasks.

6 DISCUSSION

The first part of our pipeline consists of QD algorithms for the discovery of high-performing diverse morphologies and their controllers. While this solves the issue of finding good morphologies to train a distilled model on, which has not been addressed in previous work, it also increases the computational cost of our pipeline greatly, since we end up optimizing a controller for each morphology in the map. While we believe our method has an advantage in this regard since we can use simple controllers in the QD phase and then distill their knowledge into any complex model, we leave the comparison of our method to methods such as [13, 16, 30] in terms of number of simulation steps for future work.

In our QD experiments, we observed that the number of migrations (i.e. solutions that move from one cell of the MAP-Elites map to another via a mutation to their morphology) is very limited. This is in line with the existing literature in brain-body co-optimization that demonstrates mutation to morphologies is often detrimental to the performance and results in ineffective search over the morphology space [4, 5, 27, 28]. The inability of solutions to migrate to new morphologies is concerning as goal-switching and creating stepping stones are core principles that make QD algorithms, and especially MAP-Elites, effective and important. We believe that having a relatively high-resolution map was helpful in our case, as it reduces the difference between morphologies in neighboring cells. However, we are doubtful how well QD algorithms will scale to more complex morphology spaces where having a higher-resolution map would be infeasible. In future work, we are going to investigate this phenomenon and how QD algorithms behave on the problem of brain-body co-optimization.

The second part of our pipeline is knowledge distillation as an alternative for the training of multi-morphology controllers (given that there exist teacher controllers). Our investigation shows that knowledge distillation can train simple controllers for multiple morphologies that would be untrainable from scratch without the capacity of a much larger network. Crucially, the distilled controllers show emergent properties. The most important of them is the generalization to unseen morphologies in a zero-shot manner. The generalization ability allows distilled controllers to be robust to perturbations to the morphology of the robot. In this sense, we consider our work as the successor of the "Resilient Machines" [2] and "Robots that can adapt like animals" [8]. In the former, the adaptation to morphological changes happens through continuous self-modeling that happens in an evolutionary time scale, and in the latter, the adaptation occurs in a faster time scale through intelligent search over a behavior repertoire. In our case, the distilled controller is already capable of controlling multiple morphologies (including seen and unseen morphologies) and can adapt instantly. In the case of a failure, one can work backward in the methods to find or optimize a controller that can recover from the damage.

Moreover, we are interested in utilizing distilled controllers to enhance the search over the morphology space. The literature on brain-body co-optimization points out the ineffective search over the morphology space due to fragile co-adaptation of body and brain as a major challenge for brain-body co-optimization [4, 28, 34]. Recently, the investigation of [28] indicates that not being able to estimate the maximum performance of morphology without fully

optimizing a controller may be a critical part of what makes the search ineffective. The distilled multi-morphology controller that generalizes to unseen morphologies can alleviate this issue and enable a better search over the morphology space.

In a similar vein, being able to control a multitude of modular morphologies unlocks a potential for adaptation and functionality through morphological changes. A distilled general controller can enable ideas such as damage recovery through shape-shifting, and reconfiguration to perform different functions similar to the ones in [19, 49, 50]. We would like to investigate the ways of exploiting general controllers in these ways in future work.

In future work, we would like to examine the generalization ability of the distilled controllers when we use more complex compatible controllers [13, 16, 20, 47]. Moreover, we used heuristically chosen parameters for the distillation process where we trained the distilled controller for a fixed number of steps. We would like to investigate the use of a held-out validation set of morphologies to maximize the generalization abilities via for early stopping or meta-learning of controllers [11, 42]. We also experimented with heuristically chosen ways of selecting teachers and assumed that a multi-morphology controller distilled from the full map would be the best for their emergent properties such as generalization. However, it is not clear how many teacher robots we need to effectively train a distilled general controller, or which set of pre-trained robots make the best teachers. Future work should investigate how the selection of teachers affects the distilled controllers' performance and how they should be selected. Lastly, we used the Evogym simulator [1], which is a 2D voxel-based soft robot simulator. The applicability of our method in different types of robots/domains, in 3D, and in real robot scenarios should be investigated in future work to show the generality of our approach.

7 CONCLUSION

We present a pipeline for creating compatible controllers that can control multiple morphologies. Given an environment and a task, we explore with Quality Diversity algorithms [37, 38] to find morphologies representing different trade-offs and individual controllers optimized to control these morphologies. After the exploration phase, we exploit these specialized single-morphology controllers as teacher controllers and distill their behaviors into a single controller that can control multiple morphologies via supervised knowledge distillation. We find that distilled controllers generalize well to unseen morphologies, and increase adaptation efficiency as a prior for further finetuning of the behavior. We hope that this approach will contribute to future progress in multi-robot general controllers and brain-body co-optimization.

ACKNOWLEDGMENTS

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] Jagdeep Bhatia, Holly Jackson, Yunsheng Tian, Jie Xu, and Wojciech Matusik. 2021. Evolution gym: A large-scale benchmark for evolving soft robots. *Advances in Neural Information Processing Systems* 34 (2021), 2201–2214.

[2] Josh Bongard, Victor Zykov, and Hod Lipson. 2006. Resilient Machines Through Continuous Self-Modeling. *Science* 314, 5802 (Nov. 2006), 1118–1121. <https://doi.org/10.1126/science.1133687>

[3] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. 2006. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. 535–541.

[4] Nicholas Cheney, Josh Bongard, Vytas SunSpiral, and Hod Lipson. 2016. On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures. In *Proceedings of the Artificial Life Conference 2016*. MIT Press, Cancun, Mexico, 226–233. <https://doi.org/10.7551/978-0-262-33936-0-ch042>

[5] Nick Cheney, Josh Bongard, Vytas SunSpiral, and Hod Lipson. 2018. Scalable co-optimization of morphology and control in embodied machines. *Journal of The Royal Society Interface* 15, 143 (June 2018), 20170937. <https://doi.org/10.1098/rsif.2017.0937>

[6] David Johan Christensen, Ulrik Pagh Schultz, and Kasper Stoy. 2013. A distributed and morphology-independent strategy for adaptive locomotion in self-reconfigurable modular robots. *Robotics and Autonomous Systems* 61, 9 (Sept. 2013), 1021–1035. <https://doi.org/10.1016/j.robot.2013.05.009>

[7] Creative Machines Lab, Cornell University, Ithaca, NY, Nicholas Cheney, Jeff Clune, and Hod Lipson. 2014. Evolved Electrophysiological Soft Robots. In *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*. The MIT Press, 222–229. <https://doi.org/10.7551/978-0-262-32621-6-ch037>

[8] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (May 2015), 503–507. <https://doi.org/10.1038/nature14422> arXiv: 1407.3501.

[9] Maxence Falldor, Félix Chalumeau, Manon Flageat, and Antoine Cully. 2023. Map-eldites with descriptor-conditioned gradients and archive distillation into a single policy. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 138–146.

[10] Andrea Ferigo, Giovanni Iacca, Eric Medvet, and Federico Pigozzi. 2021. *Evolving Hebbian Learning Rules in Voxel-based Soft Robots*. preprint. <https://doi.org/10.36227/techrxiv.17091218.v1>

[11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *arXiv:1703.03400 [cs]* (July 2017). <http://arxiv.org/abs/1703.03400> arXiv: 1703.03400.

[12] Scott Fujimoto, Herke Hoof, and David Meger. 2018. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*. PMLR, 1587–1596.

[13] Agrim Gupta, Linxi Fan, Surya Ganguli, and Li Fei-Fei. 2022. Metamorph: Learning universal controllers with transformers. *arXiv preprint arXiv:2203.11931* (2022).

[14] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).

[15] Shengyi Huang, Rousslan Fernand Julian Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty, Kinal Mehta, and João G.M. Araújo. 2022. CleanRL: High-quality Single-file Implementations of Deep Reinforcement Learning Algorithms. *Journal of Machine Learning Research* 23, 274 (2022), 1–18. <http://jmlr.org/papers/v23/21-1342.html>

[16] Wenlong Huang, Igor Mordatch, and Deepak Pathak. 2020. One Policy to Control Them All: Shared Modular Policies for Agent-Agnostic Control. *arXiv:2007.04976 [cs, stat]* (July 2020). <http://arxiv.org/abs/2007.04976> arXiv: 2007.04976.

[17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[18] Sam Kriegman, Nick Cheney, and Josh Bongard. 2018. How morphological development can guide evolution. *Scientific Reports* 8, 1 (Dec. 2018), 13934. <https://doi.org/10.1038/s41598-018-31868-7>

[19] Sam Kriegman, Stephanie Walker, Dylan S. Shah, Michael Levin, Rebecca Kramer-Bottiglio, and Josh Bongard. 2019. Automated Shapeshifting for Function Recovery in Damaged Robots. In *Robotics: Science and Systems XV*. Robotics: Science and Systems Foundation. <https://doi.org/10.15607/RSS.2019.XV.028>

[20] Vitaly Kurin, Maximilian Igl, Tim Rocktäschel, Wendelin Boehmer, and Shimon Whiteson. 2021. My Body is a Cage: the Role of Morphology in Graph-Based Incompatible Control. <http://arxiv.org/abs/2010.01856> arXiv:2010.01856 [cs, stat].

[21] Valentin Macé, Raphaël Boige, Félix Chalumeau, Thomas Pierrot, Guillaume Richard, and Nicolas Perrin-Gilbert. 2023. The quality-diversity transformer: Generating behavior-conditioned trajectories with decision transformers. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 1221–1229.

[22] Dries Marzougui and Matthijs Biondina. 2022. A Comparative Analysis on Genome Pleiotropy for Evolved Soft Robots. (2022), 4.

[23] Eric Medvet, Alberto Bartoli, Andrea De Lorenzo, and Giulio Fidel. 2020. Evolution of distributed neural controllers for voxel-based soft robots. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. ACM, Cancún Mexico, 112–120. <https://doi.org/10.1145/3377930.3390173>

[24] Eric Medvet, Alberto Bartoli, Federico Pigozzi, and Marco Rochelli. 2021. Biodiversity in evolved voxel-based soft robots. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lille France, 129–137. <https://doi.org/10.1145/3449639.3459315>

[25] Eric Medvet and Giorgia Nadizar. 2024. GP for Continuous Control: Teacher or Learner? The Case of Simulated Modular Soft Robots. In *Genetic Programming Theory and Practice XX*. Springer, 203–224.

[26] Eric Medvet, Giorgia Nadizar, and Federico Pigozzi. 2022. On the Impact of Body Material Properties on Neuroevolution for Embodied Agents: the Case of Voxel-based Soft Robots. (2022), 9.

[27] Alican Mertan and Nick Cheney. 2023. Modular Controllers Facilitate the Co-Optimization of Morphology and Control in Soft Robots. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Lisbon Portugal, 174–183. <https://doi.org/10.1145/3583131.3590416>

[28] Alican Mertan and Nick Cheney. 2024. Investigating Premature Convergence in Co-optimization of Morphology and Control in Evolved Virtual Soft Robots. In *European Conference on Genetic Programming (Part of EvoStar)*. Springer.

[29] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv:1504.04909 [cs, q-bio]* (April 2015). <http://arxiv.org/abs/1504.04909> arXiv: 1504.04909.

[30] Giorgia Nadizar, Eric Medvet, Kathryn Walker, and Sebastian Risi. 2023. A Fully-distributed Shape-aware Neural Controller for Modular Robots. (2023).

[31] Jørgen Nordmoen, Frank Veenstra, Kai Olav Ellefsen, and Kyrré Glette. 2021. Map-eldites enables powerful stepping stones and diversity for modular robotics. *Frontiers in Robotics and AI* 8 (2021), 639173.

[32] Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. 2015. Actor-mimic: Deep multitask and transfer reinforcement learning. *arXiv preprint arXiv:1511.06342* (2015).

[33] Deepak Pathak, Christopher Lu, Trevor Darrell, Phillip Isola, and Alexei A Efros. 2019. Learning to Control Self-Assembling Morphologies: A Study of Generalization via Modularity. (2019), 11.

[34] Rolf Pfeifer and Josh Bongard. 2006. *How the body shapes the way we think: a new view of intelligence*.

[35] Federico Pigozzi, Yujin Tang, Eric Medvet, and David Ha. 2022. Evolving Modular Soft Robots without Explicit Inter-Module Communication using Local Self-Attention. *arXiv:2204.06481 [cs]* (April 2022). <https://doi.org/10.1145/3512290.3528762> arXiv: 2204.06481.

[36] Joshua Powers, Sam Kriegman, and Josh Bongard. 2018. The effects of morphology and fitness on catastrophic interference. In *The 2018 Conference on Artificial Life*. MIT Press, Tokyo, Japan, 606–613. https://doi.org/10.1162/isal_a_00111

[37] Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley. 2016. Quality Diversity: A New Frontier for Evolutionary Computation. *Frontiers in Robotics and AI* 3 (July 2016). <https://doi.org/10.3389/frobt.2016.00040>

[38] Justin K. Pugh, L. B. Soros, Paul A. Szerlip, and Kenneth O. Stanley. 2015. Confronting the Challenge of Quality Diversity. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation* (Madrid, Spain) (GECCO '15). Association for Computing Machinery, New York, NY, USA, 967–974. <https://doi.org/10.1145/2739480.2754664>

[39] Andrei A Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. 2015. Policy distillation. *arXiv preprint arXiv:1511.06295* (2015).

[40] Michael D Schmidt and Hod Lipson. 2010. Age-fitness pareto optimization. (2010), 2.

[41] Student. 1908. The probable error of a mean. *Biometrika* (1908), 1–25.

[42] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. 2019. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 403–412.

[43] Jacopo Talamini, Eric Medvet, Alberto Bartoli, and Andrea De Lorenzo. 2019. Evolutionary Synthesis of Sensing Controllers for Voxel-based Soft Robots. (2019), 8.

[44] Jacopo Talamini, Eric Medvet, and Stefano Nichele. 2021. Criticality-Driven Evolution of Adaptable Morphologies of Voxel-Based Soft-Robots. *Frontiers in Robotics and AI* 8 (2021). <https://www.frontiersin.org/article/10.3389/frobt.2021.673156>

[45] Fabio Tanaka and Claus Aranha. 2022. Co-evolving morphology and control of soft robots using a single genome. <http://arxiv.org/abs/2212.11517> arXiv:2212.11517 [cs].

[46] Yee Teh, Victor Baptst, Wojciech M Czarnecki, John Quan, James Kirkpatrick, Raia Hadsell, Nicolas Heess, and Razvan Pascanu. 2017. Distral: Robust multitask reinforcement learning. *Advances in neural information processing systems* 30 (2017).

[47] Tingwu Wang, Renjie Liao, Jimmy Ba, and Sanja Fidler. 2018. Nervenet: Learning structured policy with graph neural networks. In *International conference on learning representations*.

[48] Frank Wilcoxon and Roberta A Wilcox. 1964. Some rapid approximate statistical procedures. (*No Title*) (1964).

[49] Victor Zykov, Efstrathios Mytilinaios, Bryant Adams, and Hod Lipson. 2005. Self-reproducing machines. *Nature* 435, 7039 (2005), 163–164.

[50] Victor Zykov, Efstrathios Mytilinaios, Mark Desnoyer, and Hod Lipson. 2007. Evolved and Designed Self-Reproducing Modular Robotics. *IEEE Transactions on*

Robotics 23, 2 (2007), 308–319. <https://doi.org/10.1109/TRO.2007.894685>