



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/182721/>

Version: Accepted Version

Article:

Le Goff, Léni, Buchanan Berumen, Edgar, Hart, Emma et al. (2023) Morpho-evolution with learning using a controller archive as an inheritance mechanism. IEEE Transactions on Cognitive and Developmental Systems. 15. pp. 507-517. ISSN: 2379-8920

<https://doi.org/10.1109/TCDS.2022.3148543>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Morpho-evolution with learning using a controller archive as an inheritance mechanism

Léni K. Le Goff, Edgar Buchanan, Emma Hart, Agoston E. Eiben, Wei Li, Matteo De Carlo, Alan F. Winfield, Matthew F. Hale, Robert Woolley, Mike Angus, Jon Timmis, Andy M. Tyrrell

Abstract—The joint optimisation of body-plan and control via evolutionary processes can be challenging in rich morphological spaces in which offspring can have body-plans that are very different from either of their parents. This causes a potential mismatch between the structure of an inherited controller and the new body. To address this, we propose a framework that combines an evolutionary algorithm to generate body-plans and a learning algorithm to optimise the parameters of a neural controller. The topology of this controller is created once the body-plan of each offspring body-plan is generated. The key novelty of the approach is to add an external archive for storing learned controllers that map to explicit ‘types’ of robots (where this is defined with respect to the features of the body-plan). By learning from a controller with an appropriate structure inherited from the archive, rather than from a randomly initialised one, we show that both the speed and magnitude of learning increases over time when compared to an approach that starts from scratch, using three different test-beds. The framework also provides new insights into the complex interactions between evolution and learning, and the role of morphological intelligence in robot design.

Index Terms—Evolutionary robotics, Embodied Intelligence

I. INTRODUCTION

The idea of embodied intelligence — describing the design and behaviours of physical objects situated in the real-world — was first introduced by Brooks in 1991 [1]. Pfeifer and Bongard’s seminal text ‘How the body shapes the way we think’ [2] expanded on the idea that intelligent control is not only dependent on brain, but at the same time both constrained and enabled by the body. Increasingly, artificial evolution approaches have been used in robotics to jointly optimise both the body-plan and controller of a robot to accomplish a desired task. This has the potential advantage of allowing evolution to discover the appropriate balance between morphological and brain complexity and functionality.

However, much of this work has taken place in restricted morphological spaces, for example using regular shaped modules to construct body-plans, in which each module can be individually actuated [3], [4]. If we consider richer spaces which can give rise to complex, irregular robot skeletons with multiple forms of sensing and actuation (e.g. joints and/or

wheels) then not only are more complex, centralised controllers required that link multiple sensors and actuators but in addition, the evolutionary process becomes more challenging: reproduction between two morphologically distinct parents might result in a viable body-plan, but a directly inherited controller is at best unlikely to provide adequate control and, at worst, will not work at all because inputs and outputs do not correspond to the new body-plan.

One approach to address this is to evolve a morphology-independent control mechanism, for example using a compositional pattern producing network (CPPN)[5] to *generate* a controller, thereby enabling direct inheritance of the generator [6]. However, generative methods tend to be computationally expensive. An alternative is to add a learning cycle into the evolutionary loop [7], [8]. This can either improve an inherited controller over an individual’s lifetime – when the inherited controller has an appropriate structure – or learn a new controller from scratch if necessary. Here, we follow the latter approach and propose a novel framework for combining evolution and learning that is capable of the joint optimisation of body and control of robots in a complex morphological space when using controller encodings that do not permit directly inheritance, i.e. when the topology of a child controller does not match the inherited body.

The framework contains a morpho-evolutionary algorithm (MEA) to optimise the body-plan and a learning algorithm to optimise the parameters of the controller. The two optimisation processes are nested: for each body-plan produced with the MEA, the learning process is invoked to optimise its controller. The key novelty of the approach is the addition of an external *controller archive*: this multi-dimensional archive stores the best found controller for a given ‘type’ of robot, where *type* is defined by a vector describing the robot’s morphological features (e.g. number of wheels, number of sensors of Type A, number of sensors of Type B, etc.). If a body-plan is produced that is of the same type as a controller already stored in the archive, the learning process is initiated with this controller, otherwise it starts from scratch. The archive is updated over the generations as better controllers are found. Essentially the archive can be viewed as a form of inheritance, storing successful controllers per robot that can be used to bootstrap *learning* in future generations. Hence the framework is named MELAI: *morpho-evolution with learning using archive inheritance*. Specifically, for the MEA, we use the matrix-based CPPN morpho-evolution (MCME) introduced in our previous work [9] to evolve body-plans. The learning algorithm used is a novelty-driven evolution-strategy, that uses an increasing

L. K. Le Goff and E. Hart are with the School of Computing, Edinburgh Napier University, Scotland, UK, contact l.legoff2@napier.ac.uk

E. Buchanan, W. Li, R. Woolley, M. Angus, A. M. Tyrrell are with the Department of Electronic Engineering, University of York, England, UK

M. F. Hale and A. F. Winfield are with Bristol Robotics Laboratory, University of the West of England, UK

J. Timmis is with School of Computer Science, University of Sunderland, England, UK

M. De Carlo and A. E. Eiben are with Department of Computer Science, Vrije Universiteit Amsterdam, NL

population size (NIP-ES), and was also introduced in our previous work [10]. It learns the weights of a controller specified by an Elman network that has a topology matching the generated child body-plan.

The contributions of the method are two-fold: (1) it offers a novel algorithm for optimisation of both body-plan and controller of robots, that integrates evolution and learning; uniquely using a morpho-evolutionary algorithm for the former and an evolution-strategy for the latter; (2) it proposes the use of an external archive as an efficient mechanism for transferring control knowledge from parents to offspring in situations where offspring are morphologically distinct from their parents. Crucially, this is not direct transfer through regular genetic inheritance from parents to offspring but acts a bootstrap mechanism to initiate learning.

We show the benefits of using an archive as a form of inheritance in terms of increasing the efficiency of the approach (compared to methods that learn from scratch) and provide new insights into the interplay of evolutionary and learning processes. Moreover, as an additional contribution, our results show the emergence of different kind of robots for different tasks.

The rest of the paper is organized as follows: section II analyses the studies related to joint optimisation of the body-plan and controller of robots, then in section III MELAI is explained in detail; the experimental protocol is described in section IV and the results are presented in section V. Finally, sections VI and VII discuss the results and conclude the paper.

II. RELATED WORK

Among the numerous studies in the field of evolutionary robotics, the majority address either the evolution of the body-plan *or* the evolution of the controller. This section focuses on literature which describes methods for the *joint* optimisation of body-plans and controllers. We focus attention on approaches that permit the evolution of offspring that require a controller topology that is different to either parent. That is, methods which evolve changes to morphology that do not impact the topology of controller (i.e. the number of inputs and outputs of a neural controller) are out of scope. For example, this excludes work in which morphological change is restricted to repositioning sensors [11] or altering the length, weight and size of leg-joints, e.g. [12], [13], [14], [15]. We first discuss methods that create controllers that are directly correlated to a specific type of body-plan, followed by morphology-independent methods, i.e. those that are capable of *generating* a controller for any given body-plan.

A naïve approach to avoiding a potential mis-match between a controller and body-plan is to evolve only the body-plan and then learn a new controller with the correct topology from scratch for each child body-plan. The work of Gupta *et al* [16] follows this approach by using evolution for the body-plans and reinforcement learning to optimize the controllers for simulated robots composed from articulated 3D rigid parts connected via motor actuated hinge joints. Learning starts from a randomly initialised controller for each body-plan, and uses a distributed implementation across multiple CPU

to minimise computational cost. Lia *et al.* [17] also propose a nested optimisation process, with the aim of finding the best morphology for a walker micro-robot. Bayesian optimisation is used to learn the controller. Despite Bayesian optimisation being a well known method to be sample efficient, it only works well for a small parameter space.

Instead of learning from scratch, an alternative approach is to use a morphology-independent control generation mechanism that can generate a control parameters for any given body-plan. For example, Cheney *et al* [18], [19] evolve soft-robots built from voxels, in which each voxel has a parameterised local controller. Both body-plans and controller parameters are outputs of two separate compositional pattern producing networks (CPPN) [5], both of which are evolved via the well known neuro-evolution with augmenting topology (NEAT) algorithm[20]. However, due the distributed nature of the controller, the variety of possible behaviours for the same body-plan is limited. Sims [21] also use a decentralised form of control in which a genotype is represented as nested graph: the graph specifies morphological nodes, each of which contains another graph specifying the neural circuitry for that node. More recent work has achieved a similar effect with the use of Lindenmayer-Systems (L-systems) decoding, for example in [22] and [23], [24].

In the latter work, an additional learning mechanism is applied to improve the inherited brains of newborn robots: the authors show that learning not only influences the morphology of the resulting robots but that also, the capacity to learn increases over generations. Jelisavcic *et al.* [25] also employ a learning mechanism. Their genome carries a pool of CPPNs which can be used to specify the weights of a controller generated to match the child body-plan. Differently to the work of Cheney *et al* which encodes a single CPPN that undergoes evolution, here a child inherits a subset of CPPNs from each parent, then a learning algorithm (HyperNEAT) is applied to the inherited pool to evolve a new pool. The process is thus Lamarckian.

To summarise, in the context of combined optimisation of body and control, on the one hand the literature has shown that using generative encodings (with and without additional learning) can mitigate the issues arising regarding inheritance of controllers that might not be applicable to a new child body-plan. However, these methods often require many evaluations to converge [6] and add additional hyper-parameters which may be difficult to optimise. On the other hand, neural controller encodings which are explicitly tied to a body-plan can be rapidly optimised as they only require weight optimisation rather than topology. Although they often cannot be inherited, this can be addressed by learning a controller from scratch, e.g. as in [16], although at the expense of ignoring any previously learned knowledge.

In this paper, we choose to use a fixed structure neural network for reasons of efficiency, motivated by the goal of eventually evolving directly in hardware. As in previous works, we use a learning algorithm to optimise a controller that has a fixed structure that matches the new body-plan [10]. However, in order to avoid starting from scratch for each body-plan as in previous work, we introduce a novel method for

storing past solutions that can be accessed by the learning algorithm to bootstrap learning. This takes the form of an archive that stores the weights of a controller for each ‘type’ of robot that has previously been encountered as described in the previous section. This archive or ‘brain pool’ is dynamically composed and adapted during the evolutionary process.

Note that the term *archive* should not be confused with other uses of the word in the wider evolutionary literature. For example, archives are commonly used in multi-objective optimisation to either drive the population toward the Pareto front or to maintain the population diversity, that is, they directly interact with an evolving population. Our approach has more in common with methods which try and enhance a search process by re-using past experience gained when solving related problems. For example, Louis and McDonnell [26] maintain a store of past solutions from similar instances which are periodically injected into an evolving population. However the approach is only applicable if instances share structural properties, hence cannot be applied to controllers with different topologies. Feng *et al* [27] attempted to reuse structured common knowledge captured in the optimized solutions of past search experiences in a form independent of solution representation, however their specific implementation is tailored to combinatorial optimisation. Here, we draw inspiration from [26] in maintaining an archive of past solutions, but store solutions with different types, corresponding to different controller topologies. In addition, the archive does not interact directly with the EA controlling body-plan evolution, but is used to initialise a learning algorithm which searches for an improved controller.

III. METHODS

A. Algorithm Description

Morpho-evolution with learning using an archive inheritance (MELAI) is an algorithm with an optimisation algorithm nested inside a second optimisation algorithm. As illustrated in figure 1, the body-plans are optimised with an evolutionary algorithm, then for each body-plan, a learning process is used to optimise their controller.

The first optimisation algorithm or morpho-evolution algorithm (MEA) uses a generative encoding to produce the robot’s body-plan, based on our previously work described in [9]. This is a matrix-based CPPN morpho-evolution denoted MCME.

The second optimization algorithm (learning) optimises the parameters of a controller with a fixed size neural network structure. Therefore, the number of parameters is fixed. The novelty-driven increasing population evolutionary strategies (NIP-ES) algorithm [10] is used for learning.

A detailed description of both MCME and NIP-ES is given in supplementary materials. Although the instantiation of MELAI described in this paper uses MCME and NIP-ES, the framework itself is general in that any kind of MEA or learning algorithm could be used.

In the rest of the paper, the *fitness* indicates the value used by the MEA for selection and *task-performance* for the value assessing the quality of a behaviour.

Inheritance of controllers from parents to children is challenging for MELAI as previously noted, since children might




Fig. 1. Diagram illustrating the MELAI algorithm. MELAI has two nested optimisation processes. As main process, a morpho-evolution algorithm, shown in black, divided in four main steps: *computation* of the fitness values, *update* the controller archive with the best ones from the current population, *Selection, mutation, and recombination*, and finally send the new population of new body-plans to the learning process. For each body-plans, a controller is learned, shown in blue. The learning process can either start from a random controller, a inherited controller, or a controller from the archive. Then, the learning process run until reaching an ending condition.

have different body-plan configurations than their parents. One way to address this issue is to learn the controller for each robot from scratch as in [16]. However, this has a number of disadvantages, including the fact that previously learned information from past learning cycles is wasted. In order to address this issue MELAI introduces three initialisation options:

- 1) Select a controller from the archive with the same number of sensors and actuators if one exists.
- 2) Start from a randomly initialised controller.
- 3) Direct controller inheritance if the parent and child share the same number and type of actuators and sensors.

In this paper, the third option of direct inheritance is not considered because the encoding and the morphological space used in MCME make it unlikely that a parent and child will share the same number and type of actuators and sensors. Thus, the benefit of direct inheritance will be negligible. This enables the experiments to focus directly on determining the benefit of the archive.

As noted, the learning algorithm used is NIP-ES, first described in [10]. The core of this method is a co-variance matrix adaptation evolutionary strategy (CMA-ES) algorithm in which a normal multivariate distribution (MVND) is used to sample a new population at each iteration. When using a controller from the archive, it is used to provide the starting mean of the MVND and thus the starting population is sampled in the surrounding of the parameters of this controller. When

starting from scratch, CMA-ES starts from a random mean.

The learning process stops when one of the ending conditions is reached:

- *A satisfactory solution is reached*: the learning algorithm finds a controller with a task-performance value above a certain threshold.
- *The maximum number of evaluations is reached*: each optimisation process has a maximum number of updates. For the MEA, this parameter is the number of generations and for the learning process is the number of evaluations. The values of these parameters have to be chosen according to the difficulty of the task and environment but also according to the constraints of the system on which the algorithm is running. In this study, a constraint of 100000 maximum evaluations is used. Given this overall budget, an additional choice that must be made is to decide how to divide it between the MEA and the learning process.
- *The performance of the robot stays very low during a trial period*. The trial period is defined by a fixed number of evaluations (50 in all experiments). If a robot has not moved (i.e. very low performance) by the end of this period then the learning process stops.

Finally, several fitness functions can be used for the MEA. The most natural fitness function to use will be the best task-performance value found during the learning. This is the one used in this paper. However, the learning process produces additional data, such as for instance the task-performance of every evaluations, behavioural descriptors, and in the case of NIP-ES, novelty scores, which could also be exploited by the MEA.




Fig. 2. Diagram illustrating the update of the controller archive. A body-plan with 2 actuators of type x and 3 sensors of type y has a new controller output of the learning process. If the cell corresponding to 2 actuators of type x and 3 sensors of type y is not empty, the new controller is compared with the stored one. The new controller replace the stored one if its task-performance is greater.

B. Controller Archive

The controller archive stores the best controller found for different ‘types’ of robot. The archive is ordered by the number of actuators and sensors. Each type of actuator and sensor constitutes a dimension of a grid cell. A cell contains the best controller found for a body-plan with the corresponding

number of actuators and sensors of each type. The cell stays empty until a body-plan of its category is generated by the MEA. For instance, let us consider a morphological space with one type of actuator and sensor as shown in figure 2. After, the learning process has ended for a body-plan with 2 actuators and 3 sensors, its controller’s task-performance value is compared with the stored one in the corresponding cell and replaces it if its task-performance value is greater. If the cell is empty the new controller is added to the cell.

We consider two types of actuator (wheel and joint), and one type of sensor. So, a ‘type’ of robot in this case is defined by a tuple (num_sensors, num_wheels, num_joints). The controller archive can be considered as a new form of inheritance. All the behavioural knowledge from past generations is stored in a common archive to be used by future generations. In this way, new child robots can leverage the learned behaviours of their ancestors.

IV. EXPERIMENTS

A. Experimental protocol

The experiments presented in this article aim to answer the following questions:

1. To what extent does using a controller-archive for inheritance improve effectiveness and efficiency when compared to learning from scratch?
2. How much the learning process is important and How much should the total evaluation budget be most effectively shared between the MEA and the learning process?

Experiments are conducted with and without the controller archive to answer question 1. In this way, the benefit of the controller archive in MELAI can be isolated. In the results section, the variant of MELAI without controller archive is called morpho-evolution with learning (MEL).

For all experiments, a fixed budget of 100000 evaluations is shared between the two optimisation processes. This has the objective of studying any trade-offs in resource allocation between the two components of the framework to answer question 2. Parameter values tested are the following : [100,40], [150,30], [200,20], [400,10], [800,5], where the first value corresponds to the number of evaluations for each body-plan during the learning phase and the second to the number of generations of the MEA. Also, these variants of MELAI are compared with a ground-truth algorithm in which the learning process is replaced by random sampling using latin hypercube sampling. This variant is called Morpho-Evolution with Latin Hypercube Sampling (MELHS). MELHS runs for 40 generations and for each body-plans 100 random controllers are sampled. These experiments are conducted only with the controller archive.

All experiments feature a population of 25 body-plans for the MEA. The hyper-parameters used for the experiments are given in the supplementary materials. All experiments are conducted in the three environments described in section IV-D and shown in figure 4. Twenty replicates are performed for each experiment. The source code to run these experiments and their data are available here : *the code and data will be provided if the paper is accepted.*




Fig. 3. Active and passive components used for the experiments shown in this paper. The active components are the wheel, joint and sensor. The passive component is the caster wheel.

B. Body-plans

The body-plans evolved using the MEA described have two main features: the skeletons of the body-plans can have complex and widely differing shapes and the body-plans can have different numbers of components. The components can be active or passive where the active ones interact with the controller. The different component types are shown in Figure 3. Each body-plan in this paper is generated by an encoded CPPN [5]. For more information on the body-plan decoding refer to the supplementary materials.

The controller has a different interaction with each active component:

- Each *wheel* takes one output from the controller which translates to the speed of the rotational movement of the wheel.
- Each *joint* takes one output from the controller which translates to frequency of the oscillatory movement of the joint. Oscillatory control for a joint tends to show better results for locomotion. Previous works using central pattern generator networks (CPG), which output oscillatory signals, shows good results for locomotion [ref]
- Each *sensor* gives two inputs to the controller where the first input is binary for the detection of a beacon and the second input is the distance from the closest obstacle. The detection of a beacon uses a simulated IR sensor and the distance measure uses a simulated time-of-flight sensor.

All of these components have been designed to match the physical ones which are used in the ARE project [28].

To have an accurate simulation, the weight of the body-plan has to be estimated. The weight of a robot is the sum of its components weights and of his skeleton. To estimate the weight of the skeleton, which has various shapes, the density of the plastic used in reality is used to estimate the weight of a voxel. Then to obtain the skeleton weight it is enough to multiply the voxels weight by its number of voxels. The robots have weights between 500g and 2kg depending on the size of the skeleton.

C. Controllers

The controller used in this study is a modified version of an Elman network [29]. An Elman network is a recurrent neural network with two hidden layers (see the figure in section C in Appendix II of the supplementary materials). The first hidden

layer is fully connected to the input and output layers. Then, each neuron is forward connected to one neuron of the second hidden layer, called the context layer. Neurons in the context layer (context units) are recursively connected to themselves, and the context units are also fully backward connected to the hidden layer. Each neuron has a sigmoid function as their activation function.

The context layer act like a short term memory and allow the network to process real numbers sequences such as time series [29]. Therefore an Elman network are more efficient as a controllers for a navigation task then a simple feed forward network[6].

The input and output layers is the only part of the network changing. For each body-plan, each Elman network have a number of inputs and outputs corresponding to its body-plans number of sensors and actuators. The hidden and context layers have a fixed structure for all the body-plans. Thus, transferring directly a trained network from its body-plan to another is not possible, unless they have the same number of sensors and actuators.

D. Task and environments

Two tasks are used in the experiments presented in this paper: exploration and photo-taxis.

a) *Exploration task*: In the *exploration task*, the robot has to visit the most zones in a limited time. The zones are squares of same size composing a grid. The task-performance is computed by counting the number of zones visited and dividing the count by the total number of zones. The grid is 8 by 8 with cells of 25 cm sides, so the total number of zones is 64. The evaluation lasts 30 seconds and takes place in the *obstacles* environment (see figure 4). In this task, there is not any target performance value.

b) *Photo-taxis task*: In the *photo-taxis task*, the robot starts at one point and has to reach a target where a beacon is placed. The robot has first to find the beacon in the arena and then go toward it. As the beacon is detected using a simulated IR sensor, the robot can not see it when it is occluded by an obstacle. The robot is evaluated three times with the target at a different positions. The task performance is then the average of the task performance obtained in each evaluation.

The task-performance function is the normalised distance between the final position (p_f) of the robot (at the end of the evaluation) and the position of the beacon (p_b) (see equation 1). This distance is subtracted to one to have a function to maximise. The distance is normalized by the length of the diagonal of the arena. As the arenas are squares of two by two metres the diagonal measures $D = \sqrt{2^3} \approx 2.83$

$$F = 1 - \frac{\|p_f - p_b\|}{D} \quad (1)$$

The success threshold used to stop the learning process is equal to 0.95 for this task. This value correspond to a circle with a radius of 14 cm around the target.

c) *Environments*: Three different environments (figure 4) are used in the experiments. They are all square of 2 metres sides and have tiled floor.

The environments have been designed to be reproducible in reality. The tiles are spaced with a small gap of 1 millimetre which corresponds to the floor of our real arenas and the walls in the obstacles environments have feet to hold them standing. These constraints in design introduce small irregularities.




Fig. 4. The three environments used for the experiments in this paper: obstacles, escape room and arena.

V. RESULTS

This section is split in four parts: experiments related to the efficiency and effectiveness of the algorithm, the controller archive dynamics, the influence of learning, and the robots diversity. The first part focuses on comparing the algorithm with and without archives by measuring the quality of the solutions produced and the efficiency of both variants. The second part looks into the controller stored in the archive, their number and quality. The third part studies the influence of learning by comparing MELAI with different learning budget and a variant without learning. Finally, the last part looks into the influence of the task and environment over the type of robots generated.

Where it is relevant a statistical test is conducted. The test used is the Mann-Whitney U under the null hypothesis. For the series over the generations or the number of evaluations the distributions of the last generations is tested. All the experiments have been replicated 20 times.

A. Efficiency and effectiveness

Three measures are used to assess the efficiency and effectiveness of MELAI: (a) the best fitness for each population, (b) the best and average initial task-performance of the learning process of the population (c) total number of evaluations. The best fitness (a) is calculated for each population after the learning has finished. The initial task-performance (b) is the lowest task-performance from the first iteration of NIP-ES. The number of evaluations (c) used during one generation is the sum of the number of evaluations used by the learning process for each body-plan in the population.

The best fitness and the initial task-performance is a measure of the performance (effectiveness) of the complete

MELAI algorithm while the number of evaluations measures the efficiency of the algorithm.

Figure 5 shows the plots of the best fitness (first row) and the best and average initial task-performance (second row) over the generations. For the exploration task and the arena with three targets, MELAI achieves better performances than MEL, which corresponds to more zones visited for the exploration and smaller distance from the target for the photo-taxis. With the exploration task, MELAI generates robots able to visit between 26 and 28 zones over 64 where with MEL, the best robots visits only 22 and 24 zones. In the arena, MELAI find solution which reaches the success threshold (about 0.14 meter from the target) in a fewer generation (about 4 generations) than MEL (about 6 generations). Also the variance over the replicates is very low with MELAI while MEL still have some replicates over the success threshold after the 12th generation. However both MELAI and MEL produce similar results on the escape room.

More interestingly, on the three environments, the best and average initial task-performance (see second row of figure 5) of MELAI is above the one obtained by MEL. This shows that starting from a controller from the archive provides a better start for the learning algorithm.

Another benefit of the controller archive is with respect to efficiency (figure 6). On the photo-taxis task, the total number of evaluations used per generation decreases over time for both algorithms (MEL and MELAI). The number of evaluations used by MELAI decrease faster and lower than MEL (see first row of figure 6). This dynamic does not appear on the exploration task because the learning algorithm does not have a target performance value for this task. So, when possible, transferring the controllers through the generation speed up the learning. In other words, the learning process increases its efficiency over generations when the archive is used.

Moreover the difference between the initial and best performance values (learning delta) stays constant over the generation for both MELAI and MEL. As the archive allows the learning to start from better solution, MELAI can reach better solution after learning and in a shorter time for the photo-taxis task. But starting from a better solution does not allow a wider learning delta. In the supplementary material, readers can find additional figures plotting the learning delta over the initial task-performance and the learning delta over the generations, which support the above interpretation.

B. Controller archive dynamics

To analyse the dynamics associated with using the controller archive, two metrics related to the controllers stored in the archive are monitored: their average and best task-performance value (first row of figure 7) and their number (second row of figure 7). The number of controllers in the archive corresponds to the number of ‘types’ of body-plan generated by the MEA (according to the 3-dimensional descriptor used).

The average and best task-performance values of the controllers stored in the archive shown in figure 7 follow the expected dynamic. It reaches quickly a plateau. The average task-performance value reached is as expected an average of

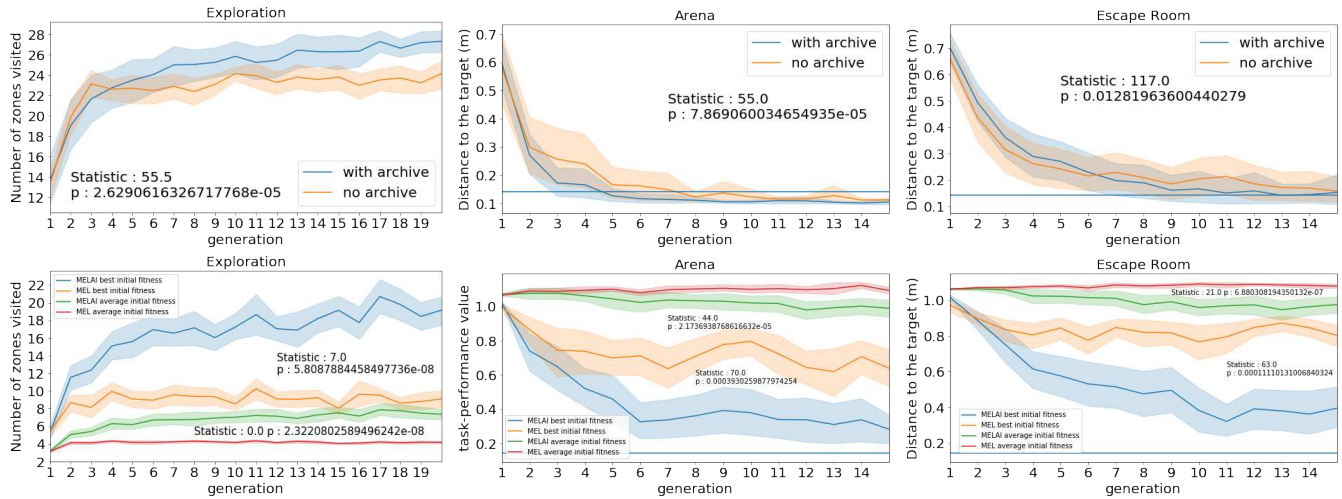


Fig. 5. Measures of effectiveness of the algorithms. In the first row, the best fitness values over the generations and in the second row the best and average of the *initial* task-performance values over the generations. The values are plotted for MEL and MELAI. These experiments have been conducted with a budget of 200 evaluations per body-plans and 20 generations. The coloured areas correspond to the confidence interval around the mean. Difference between the distributions of the last generation is significant when a p-value and its critical value is indicated. The significant test is the Mann-Whitney U test.




Fig. 6. The number of evaluations per generation. The values are plotted for MEL and MELAI. The coloured areas correspond to the confidence interval around the mean. Difference between the distributions of the last generation is significant when a p-value and its critical value is indicated. The significant test is the Mann-Whitney U test.

the worst possible task-performance value and the best task-performance value reachable. Of course, the worst possible task-performance value is different for each environment and task. The controller archive accumulates controllers through the generations: some controllers are replaced over time by higher performing versions, while others may never be updated if the type of body-plan they belong to is not selected. On the other hand the best task-performance in the archive corresponds to the best task-performance in the population (see figure 5).

The accumulation of controllers is shown in the second row of figure 7. The number keep increasing and never reach a plateau but slows down. This means that the MEA is capable of continuing to produce body-plans with new combinations of sensors and actuators. This is not surprising as the number of different possible combination of sensors, joints and wheels using MCME is 1024.

C. Influence of learning

Besides the controller archive, MELAI has a complexity as it combined evolution of morphology and learning of behaviours. To have some insights on the interaction between these two optimisation processes, MELAI have been run with five different budgets: [800,5],[400,10],[200,20], [150,30], and

[100,40] which correspond respectively to [*learning budget, number of generation*]. All the variants has a population of 25 body-plans. Thus each variant tests a different total number of body-plans, for instance, the variant [400,10] tests 2500 body-plans. However, all the variants have the same total number of evaluations of 100000. These variants are compared with MELHS as ground-truth. All these complementary experiments have been conducted on the exploration task.

Figure 8 shows the best task-performances (number of zones visited) over the number of evaluations. As expected, the advantage of using learning is clear. The best individuals produced by MELHS are visiting in average between 8 and 9 zones while all the variants of MELAI reaches between 23 and 28 in average. Also the quality of the solutions is increasing very slowly starting from 8 and reaching barely 9 at the end. With MELAI all the variants displays a neat learning curve.

On the other hand, the difference between the different budget on MELAI is small. The variant [800,5] is suboptimal comparing with the others. The variant [100,40] is the fastest to reach a satisfactory solution with in average 25 zones visited after around 15000 evaluations. Generally, reducing the learning budget speed up the process. This is due to the generational aspect of MELAI, smaller the learning budget faster the generations. Ultimately, apart from [800,5] all the




Fig. 7. Plots of three metrics over the generations related to the controllers in the archive. First row: the average and best task-performance. Second row: the number of controllers in the archive. The coloured areas correspond to the confidence interval around the mean.




Fig. 8. Best fitness over the number of evaluations on the exploration task for six variants: $[evaluations, generations]$ of $[800,5]$, $[400,10]$, $[200,20]$, $[150,30]$, and $[100,40]$ and MELHS. The coloured areas correspond to the confidence interval around the mean.

variants converges to similar task-performances.

D. Robot diversity

Finally, figure 9 shows the distribution of number of wheels, joints and sensors of the robots with the highest task-performance over the generations. These show that the majority of successful robots have between three and five wheels for all the environments and both tasks. While, the majority of successful robots have no joints. This was expected as the three environments have a flat floor. Interestingly, for the photo-taxis task the best type of robot features at least one sensor, while for the exploration task the majority of the best solutions does not feature sensors. Indeed, to reach the target three time at different location, the robot needs a sensor. On the contrary, blind robots can easily visits a lot of zones.

This result is not surprising given that the main optimisation process in MELAI is an MEA. Evolution is most likely to proceed along the ‘easiest’ path that enables it to maximise the fitness function. In this case, this corresponds to robots

with only wheels. This type of robot is easier to control and therefore it is easier to learn a controller for them than robots with joints and sensors, even though the latter may be more efficient. Sensors emerge only if there are necessary like in the photo-taxis task.

Therefore MELAI is able to produce different ‘type’ of robots depending on the task. In the supplementary material, the reader can find a bar plot of the ‘type’ of robots produced over the whole population and all the generations. It shows that the algorithm generate more robot with sensors on the photo-taxis tasks.

Pictures of successful robots are shown in figure 10 as examples.

VI. DISCUSSION

Advancing previous work in the domain of body-brain evolution, we proposed a method to evolve robots in a rich morphological space that includes a variety of sensors and actuators and can realise skeletons with diverse forms and

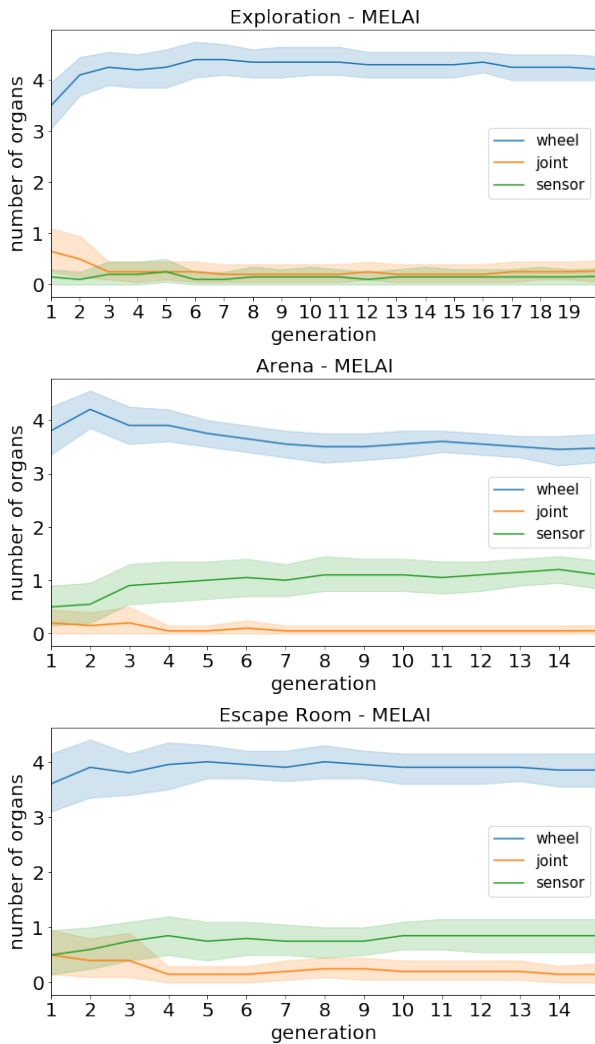


Fig. 9. For each of the three environments, distribution of the number of wheels, joints and sensors over the best robots of each generation. The coloured areas correspond to the confidence interval and the solid curves to the central tendency.




Fig. 10. Examples of successful robots in each of the three environments

sizes. Hence, a considerably more diverse range of body-plans can be produced in this space than in previous work that either uses modular systems [24] or spaces in which the components have common control mechanisms [16]. The richer space increases the likelihood that a controller produced via evolutionary operators will not match a new body-plan. Additionally, the diversity of actuators in particular increases the difficulty of learning a useful controller. Although using a generative (morphology-independent) encoding can address the inheritance issue, the time-complexity associated with these methods can be prohibitive when working with physical robots.

To address this we proposed the use of an external archive that stores a learned controller associated with a ‘type’ of robot. As described in the results section, we have demonstrated that the archive improves significantly the quality of the solution and the efficiency in evolving a body-plan capable of solving a task. The role of the archive and the various components of the framework used in leading to this result are discussed below.

Each cell of the controller archive stores the best controller learned during an individual lifetime. The archive thus represents a history of knowledge that was learned in previous generations and can be passed to future generations. It therefore acts as a novel form of *inheritance*. In the sense that it stores information learned during an individual lifetime, it shares characteristics with Lamarckian artificial evolutionary systems [25]. Note that the tuple defining a ‘type’ is deliberately simple. However, it should be clear that many different body-plans can be mapped to each single cell, given that for any given combination of sensors and actuators, the robot-skeleton they are attached to can vary enormously in shape and size, and the configuration of components can also vary. Given this variation, it might be expected that inheriting a controller of the correct ‘type’ would not necessarily bring much benefit. However, it is clear from figure 5 that shows the best and initial task-performance, that starting from a controller from the archive brings a significant advantage. Interestingly, this suggests that there is some generalisation of controllers across a range of body-plans. The results shown in the second row of figure 5 show that inheriting from the archive bootstraps the learning process, and the size of this effect increases in magnitude as the generations progress.

It would of course be possible to define each cell using a higher degree of granularity, although there it is reasonable to assume there is a balance to be struck in not making the archive too granular (which at the extreme would map every robot to an individual cell). Another way to approach this would be to store multiple controllers per cell, and either try them all, select one at random, or use a clustering or species system to select the most suitable one. Also, increasing the granularity of the archive means finding a morphological descriptor meaningful in term of control. Finding a suitable morphological descriptor is not an easy task. In one of our previous works [9], different morphological descriptors have been studied for novelty search. Among them was a symmetry descriptor and a simplest one in which each components of the robot was counted including the sensors and actuators. They

are both suitable to separate different type of control. The results shows that the most simple was better for producing diverse body-plan with novelty search.

Recall that the framework consists of two components: an MEA that learns body-plans and a learning algorithm based on an evolutionary strategy that learns controllers. The former is selected for its ability to explore a diverse space of plans and based on previous work [9]. The selection of NIP-ES as the learner is deliberate in that this algorithm demonstrates high *exploration* capabilities. This is essential as the learner might have to start from scratch if no controller is available in the archive, or a selected controller might not be well adapted to a new body. In contrast previous work which has used learning as a mechanism to enhance a controller selected by evolution (e.g. [24]) can afford to be much more exploitative.

Given the importance of the learning loop just discussed when jointly optimising body-plans and controllers, it is natural then to discuss how a computational budget should be balanced between the outer evolutionary loop and the inner learning loop. The results shown in figure 8 shed some light on this by varying the budget assigned to the learning from 100 to 200 evaluations. The smaller learning budget delivers a faster bootstrap in both environments. It is also clear that using the archive results in lower variance, particularly noticeable when using the smallest learning budgets. However, the budget of 200 evaluations gives more consistent results over the environments and the two variants. Also, 200 evaluations is the necessary minimum budget to have NIP-ES to its full potential [10] (see supplementary materials). So, the choice of budget is dependant on the learning algorithm used in MELAI. Also, it is worth remarking that the decision regarding how to split this budget is influenced by whether one is working in simulation or on physical robots: in simulation, generating a body-plan has negligible cost whereas in reality, producing a physical robot can take weeks [17]. In contrast, evaluations are cheap in both environments hence this may influence the choice.

Finally, many design choices made in this work are aiming to match the present method with our physical system [28]. In particular, NIPES and the controller archive reduce the number of evaluations needed to reach a satisfactory solution. Of course, the present framework cannot be applied directly on the real robotic platform. Indeed, on the photo-taxis task in the arena, MELAI needs about 15000 evaluations shared among 100 body-plans tested to reach a robots fulfilling the task. In future work, hybrid methods using both simulated and real robots will be investigated.

VII. CONCLUSION

This paper proposed a new framework MELAI for the joint optimisation of body-plans and controllers in a diverse and complex morphological space. The framework intertwines an evolutionary algorithm MCME for evolving body-plans and an evolution strategy NIP-ES for learning individual controllers. Its key novelty is in the use of an external archive for storing learned controllers for different ‘types’ of robot. This acts as a novel form of inheritance and is shown to bring benefits with respect to efficiency, bootstrapping the learning process

and leading to increased rates and magnitude of learning over generations. It provides new insights into the complex interactions between evolution and learning, and the role of morphological intelligence in robot design. The work provides a foundation for moving towards applying the framework to evolve robots completely in hardware: in such a space, increasing the efficiency of the evolutionary cycle is key for reasons that include time, cost of materials, and wear and tear on robotic parts.

ACKNOWLEDGMENT

This work is funded by EPSRC ARE project, EP/R03561X, EP/R035733, EP/R035679, and the Vrije 610 Universiteit Amsterdam.

APPENDIX I: ALGORITHMS DESCRIPTION

A. Matrix-based CPPN morpho-evolution

1) *Body-plan decoding*: The body-plan decoding is a variation of the one proposed in the work of Buchanan et al.[9]. The CPPN genome in this paper has four inputs and five outputs. Three inputs represent the x, y and z coordinates of a cell in a 3D matrix to queried and a fourth input represents the distance from the cell of the matrix to the centre of the matrix. Each of the outputs defines the presence or absence of the skeleton and/or component of each type.

The genome decoding takes place in four steps:

- 1) All the cells in the 3D matrix are queried to generate the skeleton of the robot.
- 2) A repair mechanism makes changes to the skeleton to meet the printing restrictions. Some of the restrictions include: make sure there is only one piece of skeleton and the skeleton is connected to the base of the head organ.
- 3) All the cells on the surface of the skeleton are queried to generate the organs. For this, the four outputs of the organs are taken. The output with the highest value defines the organ to be place on the cell.
- 4) A second repair mechanism removes colliding organs.

The decoding used in this paper has the additional feature of generating multi-segmented robots and it works as follows. The position of each skeleton voxel is queried in CPPN (Figure 11.1). If the component generated is a joint (Figure 11.2) then a cuboid skeleton is generated at the other end of the joint (Figure 11.3), The position of each face of cuboid is queried to the same CPPN and components are generated (Figure 11.4). The work of Hale et al.[28] described how the physical multi-segmented robot are assembled in the robot fabricator.




Fig. 11. Generation of multi-segmented robots. (1) The main skeleton is generated first. (2) A joint is place on the surface of one of the voxels. (3) a cuboid skeleton with 4 cm side is generated at the other end of the joint. (4) The CPPN is queried to generate components at each side of the cuboid.

The algorithm evolving the CPPN is the neuro-evolution of augmenting topology (NEAT) [20] which is a generational EA using a generative encoding to evolve both the topology and the weights of the network. In this work, we use the implementation of NEAT from the MultiNEAT library¹.

2) *Manufacturability restrictions*: Each component in the body-plan has to meet the same manufacturability criteria introduced in the work of Buchanan et al.[9]. If an component

fails any of the manufacturability tests then the component is removed from the final body-plan phenotype.

The physical head organ has 8 electrical connections for components, therefore only up to 8 active components can be connected to head skeleton at any time. The joints offer the option to electrically daisy chain one more active component. In total, a body-plan can have up to 16 active components.

The size of the skeleton connected to the head component can be as big as 23cm x 23 cm 23 cm.




Fig. 12. Diagram describing NIP-ES algorithm

B. NIP-ES

The novelty-driven increasing population evolutionary strategies (NIP-ES) is a learning algorithm introduced in one of our previous works [10]. Primarily, this algorithm was designed to find a solution using as few evaluations as possible. NIP-ES is a custom version of the increasing population co-variance matrix adaptation evolutionary strategies (IPOP-CMA-ES) proposed by Auger and Hansen [30]. NIP-ES is an algorithm of the CMA-ES family [31], [32] in which a multivariate normal distribution (MVND) is used to sample a set of solutions to be evaluated. This set of solutions is equivalent to the population of an EA.

An iteration of a CMA-ES consists in three steps (see figure 12):

- (1) *update* the co-variance matrix and mean of the MVND;
- (2) *sample* a new population;
- (3) *evaluate* the population.

The magnitude of change of the co-variance matrix is controlled by a parameter called sigma step. This parameter is similar to the learning rate in reinforcement learning. From an initial value (σ_0), the sigma step decreases at each iteration and so, the adaptation of co-variance matrix slows down. Among, the different updates function of the co-variance matrix existing in the CMA-ES family, naturally, NIP-ES uses the one of IPOP-CMA-ES implemented in libcmaes² and used in the benchmark of CMA-ES by Hansen [33].

In IPOP-CMA-ES can restart under certain conditions. After a restart, the MVND's parameters along with the sigma step

¹<http://www.multineat.com/>

²<https://github.com/CMA-ES/libcmaes>

are reinitialised and the population size is increased by a factor two. As shown in figure 12, NIP-ES features the same restart mechanism but differs in the stopping conditions. Two conditions can trigger the restart of NIP-ES:

- *Best task-performance stagnation* : Over a window of 20 iterations if the standard deviation of the best task-performance values are below a threshold (τ_1);
- *Low behavioural diversity* : If the standard deviation of the populations behavioural descriptors are below a threshold (τ_2). In this paper, the behavioural descriptor of an individual is its final position in the arena.

These two conditions aim at detecting when the algorithm get stuck in a local optimum.

Finally, NIP-ES' fitness function is a weighted sum of two objectives: the task-performance value and the behavioural novelty score (see equation 2). The novelty score measures how much the behaviour of an individual is new in comparison with the other individual in the population and past individuals stored in an archive [34]. The novelty score is computed by averaging the distances between the individual and its 15 nearest neighbours in the population and the archive. The archive of past individuals is updated at each iteration by adding randomly a part of the population and individuals with a novelty score above a threshold.

$$F = \eta * S + (1 - \eta) * r \quad (2)$$

The objectives are weighted with a novelty ratio (η), the novelty score (S) is multiplied by the novelty ratio (η) and the task-performance value (r) by the opposite novelty ratio (see equation 2). The novelty ratio starts at one and then decreases by a fix decrements (η_d). When the algorithm restart the novelty ratio is reinitialized at one. NIP-ES starts with a pure exploratory behaviour to slowly transitions to a exploitative behaviours.

In the context of this work, NIP-ES has three stopping conditions:

- if an individual get a task-performance value above a success threshold (τ_S);
- if the maximum budget of evaluations is reached. The budget can be exceeded when the size of the last population is greater than the number of evaluations remaining.
- if after a trial periods of 50 iterations is passed with getting the minimal task-performance value.

This last condition was introduced for MELAI to detect when a body-plan does not have the minimum capability required to solve the task. In a navigation task, it is simply by detecting if the robot does not move after a certain simulation duration.

NIP-ES is constituted of cycles by starting to explore for solutions and then exploit the most promising ones. And after each restart the exploration power of the algorithm increases by doubling the population's size. By starting with a small population and only increasing it if necessary, NIP-ES tends to use the minimum necessary number of evaluations [10].

With the hyper-parameters used in this paper, these cycles are roughly of 20 iterations, therefore with a starting population of 10, the minimum budget to have NIP-ES working in

TABLE I
HYPER-PARAMETERS OF NIPE-ES

Initial sigma step (σ_0)	1
Initial population size	10
Initial novelty ratio (η_0)	1
Novelty ratio decrements (η_d)	0.05
Sparseness number of nearest neighbors	15
Novelty threshold to add to archive	0.9
Probability to add to archive	0.4
Simulation time	60 seconds
Best task-performance stagnation threshold (τ_1)	0.05
Low behavioural diversity threshold (τ_2)	0.1
Trial period (number of iteration)	50
Success threshold (τ_S)	0.95

its full potential is 200. This estimation is due to the fact that the first stopping criterion can be triggered every 20 iterations and the second criterion is the least probable to be triggered.

The hyper-parameters and their values for each experiments of this article are listed in table I.

APPENDIX II: COMPLEMENTARY PLOTS AND FIGURES

C. Elman Network




Fig. 13. Diagram describing the Elman network structure. Solid lines correspond to forward connections and dashed lines to backward connections.

D. Learning Delta over the initial task-performance

The learning delta is the difference between the the initial and best task-performance. Figure 14 shows that at the best, the learning delta decreases linearly when the initial task-performance increases. Which means that starting from a better initial solution does not increase the learning delta but allow in most cases to reach better solution. Also, there is not a clear difference between with and without archive on these plots. Only, MELAI reaches higher initial task-performance on the exploration task and on the photo-taxis task, the distribution for MEL is more packed around the mean (0.6 of initial task-performance).

E. Type of robots produced by MELAI




Fig. 14. Scatter plots of the learning delta over initial task-performance. All the values are given with the normalised task-performance

REFERENCES

- [1] R. A. Brooks, "Intelligence without reason," in *Proceedings of the 12th International Joint Conference on Artificial Intelligence - Volume 1*, ser. IJCAI'91. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1991, p. 569–595.
- [2] R. Pfeifer and J. Bongard, *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.
- [3] S. Kriegman, D. Blackiston, M. Levin, and J. Bongard, "A scalable pipeline for designing reconfigurable organisms," *Proceedings of the National Academy of Sciences*, vol. 117, no. 4, pp. 1853–1859, 2020. [Online]. Available: <https://www.pnas.org/content/117/4/1853>
- [4] V. Vujovic, A. Rosendo, L. Brodbeck, and F. Iida, "Evolutionary developmental robotics: Improving morphology and control of physical robots," *Artificial Life*, vol. 23, no. 2, pp. 169–185, 2017, pMID: 28513207.
- [5] K. O. Stanley, "Compositional pattern producing networks: A novel abstraction of development," *Genetic programming and evolvable machines*, vol. 8, no. 2, pp. 131–162, 2007.
- [6] L. K. Le Goff, E. Hart, A. Coninx, and S. Doncieux, "On pros and cons of evolving topologies with novelty search," in *Artificial Life Conference Proceedings*. MIT Press, 2020, pp. 423–431.
- [7] A. Eiben, N. Bredeche, M. Hoogendoorn, J. Stradner, J. Timmis, A. Tyrrell, and A. Winfield, "The triangle of life: Evolving robots in real-time and real-space," in *Proc. of the 12th European Conference on the Synthesis and Simulation of Living Systems (ECAL 2013)*, P. Lio, O. Miglino, G. Nicosia, S. Nolfi, and M. Pavone, Eds. MIT Press, 2013, pp. 1056–1063.

- [8] A. E. Eiben and E. Hart, "If it evolves it needs to learn," in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*, ser. GECCO '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1383–1384.
- [9] E. Buchanan, L. K. Le Goff, W. Li, E. Hart, A. E. Eiben, M. De Carlo, A. F. Winfield, M. F. Hale, R. Woolley, M. Angus *et al.*, "Bootstrapping artificial evolution to design robots for autonomous fabrication," *Robotics*, vol. 9, no. 4, p. 106, 2020.
- [10] L. K. Le Goff, E. Buchanan, E. Hart, A. E. Eiben, W. Li, M. de Carlo, M. F. Hale, M. Angus, R. Woolley, J. Timmis *et al.*, "Sample and time efficient policy learning with cma-es and bayesian optimisation," in *Artificial Life Conference Proceedings*. MIT Press, 2020, pp. 432–440.
- [11] G. Buason, N. Bergfeldt, and T. Ziemke, "Brains, bodies, and beyond: Competitive co-evolution of robot controllers, morphologies and environments," *Genetic Programming and Evolvable Machines*, vol. 6, no. 1, pp. 25–51, 2005.
- [12] T. F. Nygaard, E. Samuelsen, and K. Glette, "Overcoming initial convergence in multi-objective evolution of robot control and morphology using a two-phase approach," in *European Conference on the Applications of Evolutionary Computation*. Springer, 2017, pp. 825–836.
- [13] T. F. Nygaard, C. P. Martin, E. Samuelsen, J. Torresen, and K. Glette, "Real-world evolution adapts robot morphology and control to hardware limitations," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2018, pp. 125–132.
- [14] T. F. Nygaard, C. P. Martin, D. Howard, J. Torresen, and K. Glette, "Environmental adaptation of robot morphology and control through real-world evolution," *arXiv preprint arXiv:2003.13254*, 2020.
- [15] J. Juárez-Guerrero, S. Muñoz-Gutiérrez, and W. M. Cuevas, "Design of a walking machine structure using evolutionary strategies," in *SMC'98 Conference Proceedings. 1998 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No. 98CH36218)*, vol. 2. IEEE, 1998, pp. 1427–1432.
- [16] A. Gupta, S. Savarese, S. Ganguli, and L. Fei-Fei, "Embodied intelligence via learning and evolution," *arXiv preprint arXiv:2102.02202*, 2021.
- [17] T. Liao, G. Wang, B. Yang, R. Lee, K. Pister, S. Levine, and R. Calandra, "Data-efficient learning of morphology and controller for a microrobot," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2488–2494.
- [18] N. Cheney, J. Bongard, V. SunSpiral, and H. Lipson, "On the difficulty of co-optimizing morphology and control in evolved virtual creatures," in *Artificial Life Conference 2016*, 2016, p. 226.
- [19] —, "Scalable co-optimization of morphology and control in embodied machines," *Journal of The Royal Society Interface*, vol. 15, no. 143, p. 20170937, 2018.
- [20] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [21] K. Sims, "Evolving 3d morphology and behavior by competition," *Artificial life*, vol. 1, no. 4, pp. 353–372, 1994.
- [22] G. S. Hornby, H. Lipson, and J. B. Pollack, "Generative representations for the automated design of modular physical robots," *IEEE transactions on Robotics and Automation*, vol. 19, no. 4, pp. 703–719, 2003.
- [23] K. Miras, E. Haasdijk, K. Glette, and A. Eiben, "Search space analysis of evolvable robot morphologies," in *International Conference on the Applications of Evolutionary Computation*. Springer, 2018, pp. 703–718.
- [24] K. Miras, M. De Carlo, S. Akhatou, and A. Eiben, "Evolving-controllers versus learning-controllers for morphologically evolvable robots," in *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*. Springer, 2020, pp. 86–99.
- [25] M. Jelisavcic, K. Glette, E. Haasdijk, and A. Eiben, "Lamarckian evolution of simulated modular robots," *Frontiers in Robotics and AI*, vol. 6, p. 9, 2019.
- [26] S. J. Louis and J. McDonnell, "Learning with case-injected genetic algorithms," *IEEE Transactions on Evolutionary Computation*, vol. 8, no. 4, pp. 316–328, 2004.
- [27] L. Feng, Y.-S. Ong, A.-H. Tan, and I. W. Tsang, "Memes as building blocks: a case study on evolutionary optimization+ transfer learning for routing problems," *Memetic Computing*, vol. 7, no. 3, pp. 159–180, 2015.
- [28] M. F. Hale, M. Angus, E. Buchanan, W. Li, R. Woolley, L. K. Le Goff, M. De Carlo, J. Timmis, A. F. Winfield, E. Hart *et al.*, "Hardware design for autonomous robot evolution," in *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2020, pp. 2140–2147.
- [29] J. L. Elman, "Finding structure in time," *Cognitive science*, vol. 14, no. 2, pp. 179–211, 1990.

- [30] A. Auger and N. Hansen, "A restart cma evolution strategy with increasing population size," in *2005 IEEE congress on evolutionary computation*, vol. 2. IEEE, 2005, pp. 1769–1776.
- [31] N. Hansen, "The cma evolution strategy: a comparing review," *Towards a new evolutionary computation*, pp. 75–102, 2006.
- [32] —, "The cma evolution strategy: A tutorial," *arXiv preprint arXiv:1604.00772*, 2016.
- [33] —, "Benchmarking a bi-population cma-es on the bbob-2009 function testbed," in *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers*, 2009, pp. 2389–2396.
- [34] J. Lehman and K. O. Stanley, "Abandoning objectives: Evolution through the search for novelty alone," *Evolutionary computation*, vol. 19, no. 2, pp. 189–223, 2011.