# Lamarck's Revenge: Inheritance of Learned Traits Can Make Robot Evolution Better

Jie Luo<sup>1,\*</sup>, Karine Miras<sup>1</sup>, Jakub Tomczak<sup>2</sup>, and Agoston E. Eiben<sup>1</sup>

<sup>1</sup>Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

<sup>2</sup>Eindhoven University of Technology, Eindhoven, The Netherlands

\*j2.luo@vu.nl

# ABSTRACT

Evolutionary robot systems offer two principal advantages: an advanced way of developing robots through evolutionary optimization and a special research platform to conduct what-if experiments regarding questions about evolution. Our study sits at the intersection of these. We investigate the question "What if the 18th-century biologist Lamarck was not completely wrong and individual traits learned during a lifetime could be passed on to offspring through inheritance?" We research this issue through simulations with an evolutionary robot framework where morphologies (bodies) and controllers (brains) of robots are evolvable and robots also can improve their controllers through learning during their lifetime. Within this framework, we compare a Lamarckian system, where learned bits of the brain are inheritable, with a Darwinian system, where they are not. Analyzing simulations based on these systems, we obtain new insights about Lamarckian evolution dynamics and the interaction between evolution and learning. Specifically, we show that Lamarckism amplifies the emergence of 'morphological intelligence', the ability of a given robot body to acquire a good brain by learning, and identify the source of this success: 'newborn' robots have a higher fitness because their inherited brains match their bodies better than those in a Darwinian system.

# Introduction

Evolutionary robotics (ER) is a research field that applies Evolutionary algorithms (EAs) to design and optimize the body, the brain, or both, for simulated and real autonomous robots  $[^{1,2}]$ . It is a promising area with a powerful rationale: as natural evolution has produced successful life forms for practically all possible environmental niches on Earth, it is plausible that artificial evolution can produce specialized robots for various environments and tasks.

Early studies in ER explored the evolution of the controller (brain) only, while the morphologies (bodies) were fixed  $[^{3,4}]$ . A holistic approach - the conjoint evolution of morphology and controller - was introduced by Karl Sims in his seminal work with virtual creatures  $[^{5}]$ . Since then, this approach has been increasingly investigated in the literature  $[^{6-14}]$ . While the inclusion of body evolution was a fundamental step toward complex robotic intelligence, there is (at least) one more layer that should be taken into consideration: learning $[^{15}]$ . Learning allows to fine-tune the coupling between body and brain and also allows for more degrees of freedom that account for environmental changes. The majority of such research consists of applying learning algorithms to the evolvable brains of robots with fixed bodies  $[^{16-25}]$ , but some have also evolved the morphologies  $[^{26-35}]$ .

Unlike classical engineering approaches, based on mathematics and physics, ER is inspired by a less understood mechanism: biological evolution. In biology, experimental research is often slowed down by the fact that evolution requires many generations with large populations of individuals whose lives may last many decades. For this reason, most of the research focuses on organisms whose life cycle is short enough to allow laboratory experiments [<sup>36</sup>]. ER offers a synthetic alternative approach where robots are the evolving entities that enable the experimental testing of hypotheses [<sup>37</sup>]. As John Maynard Smith, one of the fathers of modern theoretical biology, argued: "So far, we have been able to study only one evolving system and we cannot wait for interstellar flight to provide us with a second. If we want to discover generalizations about evolving systems, we will have to look at artificial ones."[<sup>38</sup>]. ER has been already used to study some of the key issues in evolutionary biology, such as the evolution of cooperation, whether altruistic [<sup>39,40</sup>] or not [<sup>41</sup>], the evolution of communication [<sup>42–44</sup>], morphological complexity [<sup>45,46</sup>], and collective swarming [<sup>47</sup>].

One of the most enduring controversial matters in evolutionary biology is Lamarckism, which asserts that adaptations acquired or learned by an individual during its lifetime can be passed onto its offspring [<sup>48</sup>]. Although this theory has been disproven by modern genetics, it is important to note that the concept of Lamarckian evolution is still a subject of debate in the scientific community, and there is no consensus on whether or not it occurs to some extent in nature [<sup>49</sup>]. For instance, one could argue that epigenetic changes [<sup>50</sup>] allow for Lamarckism to occur. Epigenetics studies patterns of gene expression that can be inherited and might remain active for multiple generations. This means that genetic expression regulated throughout the life of an organism can be transmitted to its offspring through temporary modifications to molecular structures of the DNA,

but which do not change the DNA sequence itself. The link between Lamarckism and epigenetics comes from the following reasoning: the behavior and environmental exposure of an organism might induce certain epigenetic changes; these changes can be reflected directly in the phenotypic expression of the offspring; therefore, what the parents experience during their lives might directly affect the phenotype of their offspring.

By simulating Lamarckian evolution in an artificial, non-biological substrate, we can study it in a what-if fashion: What if Lamarck was right and individual traits acquired during a lifetime could be passed on to offspring through inheritance? Empirical data delivered by computer simulations can help study and analyze the evolutionary dynamics and explore the potential benefits and drawbacks of Lamarckian evolution for developing robots. Thus, on the one hand, from the robotics perspective, this can contribute to more advanced evolutionary algorithms that in turn can deliver better robotic systems, perhaps even in less time. On the other hand, from the biological perspective, this can contribute to insights into natural evolution; not Life as we know it, but Life as it could be.

Previous research on artificial evolution combined with learning is mostly limited to Darwinian systems  $[^{14,32,35,51}]$  and the Baldwin effect  $[^{28,29}]$ . The few existing studies on Lamarckian evolution can be divided into three categories. First, disembodied evolution applying an evolutionary algorithm to machine learning techniques  $[^{52-60}]$ . In this category, studies found that the Lamarckian mechanism quickly yields good solutions, accelerating convergence and adapting to dynamic environments but may risk converging to a local optimum and may yield different results in different domains. Second, embodied evolution of controllers for robots with fixed bodies  $[^{22,61,62}]$ . In this category, studies found that Lamarckian evolution is effective in improving the performance of robot controller evolution, and that the learning process reduces the negative impact of the simulation-reality gap. Finally, there is the full-blown case, embodied evolution of morphologies and controllers together in a Lamarckian manner. This is the most complex category that has hardly been studied so far with only two papers (of ourselves) we are aware of  $[^{63,64}]$ . These considered the simplest possible robot task (undirected locomotion, a.k.a. gait learning) and observed the increased efficiency and efficacy of Lamrackian evolution compared to the Darwinian counterpart.

Importantly, all previous studies focused on establishing the advantages of Lamarckism without a deeper investigation into why and how Lamarckism delivers such benefits and to date there is hardly any knowledge about the most complex case of morphologically evolvable robots. This latter may be rooted in the difficulty in designing and implementing such a system. Technically speaking, it requires a reversible mapping between (certain segments of) the genotype and the phenotype. In particular, some features of the robot controller must be evolvable (i.e. inheritable) as well as learnable, and the traits acquired by the learning algorithm during the lifetime of a robot must be coded back to the genotype to make them inheritable.

In this work, we investigate the effects of Lamarckism on morphologically evolvable robots. Specifically, we apply a Lamarckian system that acts upon the learning layer –what the parents learn can be inherited by the offspring– while solving the reversible genotype-phenotype mapping problem. We compare the Lamarckian system to a Darwinian system [<sup>35</sup>] in which learning occurs, but learned traits are not inherited. Both systems include body evolution, brain evolution, and learning. All properties and parameters in both systems are the same, except that the inheritance of learned traits is present only in the Lamarckian system. Specifically, we test three hypotheses:

- The Lamarckian system is more effective and efficient than the Darwinian system.
- The Lamarckian system converges into superior bodies faster.
- The Lamarckian system produces better 'newborns' with relatively high performance even before learning takes place.

The main contributions of the present work are twofold: a) a general framework for a Lamarckian robot evolution system with a reversible genotype-phenotype mapping, and b) novel insights into the deeper effects of Lamarckism underlying the increased effectiveness and efficiency occur.



Figure 1. Mean (lines) and maximum (dots) fitness over 30 generations. The bands indicate the 95% confidence intervals in the title is the significance of the comparison between (Sample Mean  $\pm$  t-value  $\times$  Standard Error).



Figure 2. Distribution of newborn fitness before learning (all runs merged). Red dots indicate the mean values. The p-value Lamarckian and Darwinian systems, including all newborns ever born (all generations and all runs merged).

# Results

We arrange the results around different robot features: task performance, morphology and behavior.

#### **Task Performance**

Robots are evolved for a point navigation task, requiring that the robot visits a sequence of target points (See the Methods section for details). Their task ability is used as the fitness function for evolution and as the reward function for the lifetime learning method, cf. Algorithm 1. Figure 1 exhibits the development of fitness over consecutive generations of the Lamarckian and the Darwinian systems. These curves show that the best robots that the Darwinian system produces reach a fitness of 2.5, but the populations produced by the Lamarckian system are significantly better - approximately 25% higher at the end.

Figure 1 also demonstrates the differences in efficiency. The Lamarckian system is more efficient than the Darwinian one, as it finds the best solutions (robots with the highest fitness) much faster. Furthermore, the dotted red lines show that halfway through the run (around generation 14) the Lamarckian system has reached the quality produced by the Darwinian system only by the end of the evolutionary process. This can be seen as significant 'savings' of 2,240,000 evaluations (25 offspring  $\cdot$  16 generations  $\cdot$  280 learning trials  $\cdot$  20 runs).

To investigate more closely what allows the Lamarckian system to be more effective and efficient than the Darwinian, we inspected the fitness of the robots both before and after learning. Figure 3 compares the fitness of the newborns after learning with the fitness of their parents also after they learned: not only are the Lamarckian parents better than the Darwinian parents, but also are the Lamarckian newborns better than the Darwinian newborns. Additionally, Figure 2 shows the fitness distributions of the newborns before learning: the fitness distributions of newborns reach higher values through the Lamarckian system than through the Darwinian system. This holds for almost all generations except a few at the beginning of the search. These observations mean that not only are Lamarckian robots better after they learn, but also better immediately after they are born.

#### **Robot Morphologies**

We analyze the morphological properties of the robots addressing four different aspects: the morphological traits (details about the measures can be found in  $[^{65}]$ ), the morphological similarity between offspring and parents, the morphological diversity at each generation and the morphological intelligence.

#### Morphological traits

To investigate the morphologies generated by the Lamarckian and Darwinian evolution systems consider eight morphological traits to quantitatively analyze the evolved morphologies of all robots. Among these eight traits, only three of them presented significant differences (Figure 4), namely branching, number of limbs, and symmetry. Robots evolved by the Lamarckian system tend to be more symmetric and have more branches and limbs than robots evolved with the Darwinian system. Nevertheless, despite these observed differences, visual inspection of the top bodies does hardly allow for an intuitive differentiation between their shapes (Figure 10). Moreover, a PCA analysis using these same eight traits does not show any difference between the



Figure 3. Average fitness: comparison between newborns (after learning) and their parents (after learning) across generations.



**Figure 4.** Morphological traits over generations. We present the progression of their means averaged over 20 runs for the entire population. Shaded regions denote a 95% confidence interval. The significance level after Bonferroni correction for 6 comparisons is p < 0.006.





**Figure 5.** Tree-edit distance density plots. (a) is the fitness over distance for the Lamarckian system. (b) is for the Darwinian system. The darker the color, the higher density of the robots in that region. The red lines are the regression lines. The correlation efficiency rate is shown in the title of each plot.

**Figure 6.** Tree-edit distance over generation. Shaded regions denote a 95% confidence interval.

morphologies produced by each method. Therefore, although there is evidence for some differences in morphological traits, these differences are marginal (Figure 9).

#### Morphological similarity

In our research, the morphological similarity is calculated as the tree-edit distance of morphological structures between each child and the fittest parent. We use the APTED algorithm, the state-of-the-art solution for computing the tree-edit distance [ $^{66}$ ].

Figure 5 shows the correlation between fitness and distance. For both methods, the correlation between fitness and distance is negative. This means that the most similar the offspring is to the parent, the higher the fitness of the offspring. Importantly, this correlation is even stronger in the case of the Lamarckian system.

Furthermore, 6 shows how the average distance progresses over the generations. With both systems, we see pressure for reducing the distance between the offspring and parent, but this pressure is higher with the Lamarckian system. This effect is logical because it is expected that the brain of a parent would be a better match to a body similar to its own body.

#### Morphological diversity

Morphological diversity is the morphological variety of each population using tree-edit distance. It is calculated as the average distance of the difference between any two robots d(x,y) at each generation.

Figure 7 illustrates a notable trend: the morphological diversity of the Lamarckian system declines at a more rapid rate compared to the Darwinian system. This observation strongly suggests that the Lamarckian system converges into superior bodies at a faster pace.

#### Morphological intelligence

Morphology influences how the brain learns. Some bodies are more suitable for the brains to learn than others. How well the brain learns can be empowered by a better body. Therefore, we define the intelligence of a body as a measure of how well it facilitates the brain to learn and achieve tasks.

We quantify morphological intelligence by the delta of the learning delta of each method, being the learning delta of the evolved body minus the learning delta of the fixed body, whereas the learning delta, being the fitness value after the parameters were learned minus the fitness value before the parameters were learned.

To verify the presence of morphological intelligence, we conducted an additional experiment. First, we evolved robot brains for fixed bodies, given that these bodies were the same initial (random) bodies produced by the main experiments. Additionally, learning was carried out just like in the main experiments. Second, we calculated the learning delta of each individual as the fitness after learning minus the fitness before learning. Finally, we compared the average learning delta produced by the experiments with fixed bodies (described above) with the learning delta from the main experiments (evolvable bodies).

In Figure 8, we see that the average learning deltas of both methods with evolved bodies grow steadily across the generations which indicates that lifetime learning leads the evolutionary search towards morphologies with increasing learning potential. The average learning delta of fixed bodies, on the other hand, grows so little that it can hardly be seen when included on the same axis of the evolvable bodies experiments. For the Lamarckian system and Darwinian system, the learning delta with evolvable bodies is around 1885% and 1305% higher than when using fixed bodies respectively. Moreover, this delta is around 30% greater for the Lamarckian system.

The plot also illustrates that the delta of learning delta between evolved body and fixed body for each method is growing across generations which demonstrates that this learning delta growth results (specifically) from morphological intelligence, and not simply from the presence of evolution.

#### **Robot Behavior**

To obtain a better understanding of the robots' behaviour, we visualize the trajectories of the 20 best-performing robots from both methods in the last generation across all runs. Figure 11 shows that all robots from the Lamarckian system reached the two target points much earlier than the ones from the Darwinian system. This can be concluded because after reaching the target, they still have time to keep moving further from the target.

# Discussion

This investigation exceeds existing studies about Lamarckism in (simulated) robot evolution systems. It is not limited to evolving brains for fixed bodies, but considers Lamarckism in the most interesting case that has hardly been studied before, where morphologies and controllers both undergo evolution. A key feature of the system is the invertible genotype-phenotype mapping regarding the robot brains. This is an important prerequisite for making learned traits inheritable, because learning always acts on the phenotypes, after 'birth'. If the genotype-phenotype mapping is invertible, then the newly learned traits that were not present in the robot at 'birth' can be coded back to its genotype before it reproduces. In turn, this makes the learned traits inheritable, thus evolvable. Our solution is developed for modular robots whose body configuration is evolvable. While



**Figure 7.** Morphological diversity of two systems averaged over 20 runs at each generation. The P-value shows a significant difference between the means diversity value of the two systems. The bands indicate the 95% confidence intervals.

(a)



**Figure 8.** Progression of the learning delta throughout evolution averaged over 20 runs. The P-value shows a significant difference between the mean learning delta of two systems (evolved body). The bands indicate the 95% confidence intervals.

(b)



Figure 9. Principal Component Analysis (PCA) biplot showing the distribution of samples based on 8 morphological traits in datasets of both methods. Each point represents a robot sample, and the plot displays the first two principal components (Dim1 and Dim2), which explain 43.2% and 30.6% of the total variance, respectively. Furthermore, the biplot displays the variables (morphological traits) as arrows, representing their contribution to the principal components. Traits pointing in similar directions are co-regulated or have similar expression patterns across the samples.



in our current system, the number of different modules is limited, the principle behind our design is generic, and applicable for robots with more different modules, if only the controller architecture can be derived from the morphology, effectively parameterizing the search space of possible robot brains.



**Figure 11.** Trajectories of the best 20 robots from both methods in the point navigation task. The purple square is the starting point. Two yellow circles are the target points which robots aim to go through. The blue lines are the trajectories of robots ending at the green squares.

The first set of our findings reconfirms earlier results about the increased efficiency and efficacy of Lamarckian evolution. Specifically, we showed that the Lamarckian system reaches the top fitness levels of the Darwinian system with just half of the effort (Figure 1). Additionally, it presents a higher overall efficacy: the average fitness of the final populations is 25% higher when using the Lamarckian system than the Darwinian system. Although previous work, from ourselves, has published similar results, the task used in<sup>63,64</sup> was very simple (undirected gait learning). Here we move the front of applicability, showing that Lamarckian evolution is also superior in practically more relevant cases.

We also showed that, although the best morphologies found by the Lamarckian and Darwinian systems are similar to each other (Figure 9, 10), the two systems differ in the morphologies created during evolution. In particular, the Lamarckian system produced offspring more similar to the parents (Figure 5, 6) and converged into superior bodies faster (Figure 7). As a consequence, the learning applied to the brains of the given bodies sped up in finding the optimal brains (Figure 8).

One new insight about Lamarckian evolution was revealed by using the notion of the learning delta, the increase of performance achieved by learning after birth. This is a specific concept for morphological robot evolution combined with learning. In such systems, a 'newborn' robot has a body and a brain, hence its fitness can be measured immediately. Additionally, fitness can be measured also after the learning process and the difference can be calculated. Figure 8 indicates the emergence of morphological intelligence: over the course of evolution, the bodies are becoming better learners in the Darwinian as well as the Lamarckian system. However, this effect was intensified with Lamarckism at 30%, therefore, Lamarckism is deemed superior in the joint evolution of bodies and brains.

Finally, we demonstrated that the newborns produced by the Lamarckian system are better even before the learning process takes place. Despite sounding logical, this observation is not obvious. One reasonable premise to expect Lamarckism to be beneficial is: if parents provide an initial load of 'cognitive resources' to their offspring, this should facilitate the offspring's initial behavior in contrast to starting from scratch. However, such an assumption could turn out not to be true. For instance, it could be that the newborns are not particularly better and that the main role of Lamarckism is acting as a smart initialization operator, which later on leads to better learning. Therefore, the main contribution of the current work is demonstrating where these benefits are coming from. This work contrasts with all previous literature, which explored the conditions in which Lamarckism could be beneficial but did not provide insights about where these benefits come from.

Importantly, one limitation of the present study is the use of small population sizes due to the high computational costs involved. Moreover, we experimented only with simulated robots; applying the Lamarckian system to physical robots and testing their performance in diverse environments would provide valuable insights into the practical applications of our findings. Finally, while Lamarckism has presented benefits in a scenario where the environment is static, it would be interesting to see whether it would still be beneficial when environmental conditions change rapidly or frequently. In such a scenario, Lamarckism could perhaps be counter-productive, leading to over-fitting to a reality that is no longer true when the offspring is born or the opposite. With these challenges in mind, we consider that future work should address Lamarckism in the context of changing environments.

# **Methods**

## Robot Morphology (Body)

#### **Body Phenotype**

The phenotype of the body is a subset of RoboGen's 3D-printable components [ $^{67}$ ]: a morphology consists of one core component, one or more brick components, and one or more active hinges. The phenotype follows a tree structure, with the core module being the root node from which further components branch out. Child modules can be rotated 90 degrees when connected to their parent, making 3D morphologies possible. The resulting bodies are suitable for both simulation and physical robots through 3D printing.

### Body Genotype

The phenotype of bodies is encoded in a Compositional Pattern Producing Network (CPPN) which was introduced by Stanley [ $^{68}$ ] and has been successfully applied to the evolution of both 2D and 3D robot morphologies in prior studies as it can create complex and regular patterns. The structure of the CPPN has four inputs and five outputs. The first three inputs are the x, y, and z coordinates of a component, and the fourth input is the distance from that component to the core component in the tree structure. The first three outputs are the probabilities of the modules being a brick, a joint, or empty space, and the last two outputs are the probabilities of the module being rotated 0 or 90 degrees. For both module type and rotation the output with the highest probability is always chosen; randomness is not involved.

The body's genotype to phenotype mapping operates as follows: The core component is generated at the origin. We move outwards from the core component until there are no open sockets(breadth-first exploration), querying the CPPN network to determine the type and rotation of each module. Additionally, we stop when ten modules have been created. The coordinates of each module are integers; a module attached to the front of the core module will have coordinates (0,1,0). If a module would be placed on a location already occupied by a previous module, the module is simply not placed and the branch ends there. In the evolutionary loop for generating the body of offspring, we use the same mutation and crossover operators as in MultiNEAT (https://github.com/MultiNEAT/).

#### **Robot Controller (Brain)**

#### **Brain Phenotype**

We use Central Pattern Generators (CPGs)-based controller to drive the modular robots, which has demonstrated their success in controlling various types of robots, from legged to wheeled ones in previous research. Each joint of the robot has an associated CPG that is defined by three neurons: an  $x_i$ -neuron, a  $y_i$ -neuron and an *out<sub>i</sub>*-neuron. The change of the  $x_i$  and  $y_i$  neurons' states with respect to time is obtained by multiplying the activation value of the opposite neuron with the corresponding weight  $\dot{x}_i = w_i y_i$ ,  $\dot{y}_i = -w_i x_i$ . To reduce the search space we set  $w_{x_i y_i}$  to be equal to  $-w_{y_i x_i}$  and call their absolute value  $w_i$ . The resulting activations of neurons  $x_i$  and  $y_i$  are periodic and bounded. The initial states of all x and y neurons are set to  $\frac{\sqrt{2}}{2}$  because this leads to a sine wave with amplitude 1, which matches the limited rotating angle of the joints.

To enable more complex output patterns, connections between CPGs of neighbouring joints are implemented. An example of the CPG network of a "+" shape robot is shown in Figure 12. Two joints are said to be neighbours if their distance in the morphology tree is less than or equal to two. Consider the  $i_{th}$  joint, and  $\mathcal{N}_i$  the set of indices of the joints neighbouring it,  $w_{ij}$  the weight of the connection between  $x_i$  and  $x_j$ . Again,  $w_{ij}$  is set to be  $-w_{ji}$ . The extended system of differential equations becomes equation 1.

$$\dot{x}_{i} = w_{i}y_{i} + \sum_{j \in \mathcal{N}_{i}} w_{ji}x_{j}, \qquad \dot{y}_{i} = -w_{i}x_{i} \qquad (1) \qquad out_{(i,t)}(x_{(i,t)}) = \frac{2}{1 + e^{-2x_{(i,t)}}} - 1 \qquad (2)$$

Because of this addition, x neurons are no longer bounded between [-1,1]. For this reason, we use the hyperbolic tangent function (*tanh*) as the activation function of *out<sub>i</sub>*-neurons (equation 2).

#### Brain Genotype

In biological organisms, including humans, not all genes are actively expressed or used at all times. Gene expression regulation allows cells to control which genes are turned on (expressed) or off (silenced) in response to various internal and external factors. Inspired by this, we utilize a fixed size array-based structure for the brain's genotypic representation to map the CPG weights. It is important to notice that not all the elements of the genotype matrix are going to be used by each robot. This means that their brain's genotype can carry additional information that could be exploited by their children with different morphologies.

The mapping is achieved via direct encoding, a method chosen specifically for its potential to enable reversible encoding in future stages. Every modular robot can be represented as a 3D grid in which the core module occupies the central position and each module's position is given by a triple of coordinates. When building the controller from our genotype, we use the coordinates of the joints in the grid to locate the corresponding CPG weight. To reduce the size of our genotype, instead of the



**Figure 12.** An example of a "+" shape robot and its brain phenotype (CPG network). In our design, the topology of the brain is determined by the topology of the body. The red rectangle is a single CPG which controls a corresponding hinge.



**Figure 13.** Brain genotype to phenotype mapping of a "+" shape robot. The left image (brain phenotype) shows the schema of the "+" shape robot with the coordinates of its joints in the 2D body grid. The right image (brain genotype) is the distance 2 neighbour of the joint at (1,0). The coordinates reported in the neighbourhood are relative to this joint. The CPG weight of the joint is highlighted in purple and its 2-distance neighbors are in blue.

3D grid, we use a simplified 3D in which the third dimension is removed. For this reason, some joints might end up with the same coordinates and will be dealt with accordingly.

Since our robots have a maximum of 10 modules, every robot configuration can be represented in a grid of  $21 \times 21$ . Each joint in a robot can occupy any position of the grid except the center. For this reason, the possible positions of a joint in our morphologies are exactly  $(21 \cdot 21) - 1 = 440$ . We can represent all the internal weights of every possible CPG in our morphologies as a 440-long array. When building the phenotype from this array, we can simply retrieve the corresponding weight starting from a joint's coordinates in the body grid.

To represent the external connections between CPGs, we need to consider all the possible neighbours a joint can have. In the 2-dimensional grid, the number of cells in a distance-2 neighbourhood for each position is represented by the Delannoy number D(2,2) = 13, including the central element. Each one of the neighbours can be identified using the relative position from the joint taken into consideration. Since our robots can assume a 3D position, we need to consider an additional connection for modules with the same 2D coordinates.

To conclude, for each of the 440 possible joints in the body grid, we need to store 1 internal weight for its CPG, 12 weights for external connections, and 1 weight for connections with CPGs at the same coordinate for a total of 14 weights. The genotype used to represent the robots' brains is an array of size  $440 \times 14$ . An example of the brain genotype of a "+" shape robot is shown in Figure 13.

The recombination operator for the brain genotype is implemented as a uniform crossover where each gene is chosen from either parent with equal probability. The new genotype is generated by essentially flipping a coin for each element of the parents' genotype to decide whether or not it will be included in the offspring's genotype. In the uniform crossover operator, each gene is treated separately. The mutation operator applies a Gaussian mutation to each element of the genotype by adding a value, with a probability of 0.8, sampled from a Gaussian distribution with 0 mean and 0.5 standard deviation.

#### **Evolution+Learning systems**

The complete integrated process of evolution and learning is illustrated in Figure 14, while Algorithm 1 displays the pseudocode. With the yellow highlighted code, it is the Lamarckian learning mechanism, without it is the Darwinian learning mechanism. Note that for the sake of generality, we distinguish two types of quality testing depending on the context, evolution or learning. Within the evolutionary cycle (line 2 and line 14) a test is called an evaluation and it delivers a fitness value. Inside the learning

cycle which is blue highlighted, a test is called an assessment (line 11) and it delivers a reward value. This distinction reflects that in general the notion of fitness can be different from the task performance, perhaps more complex involving more tasks, other behavioral traits not related to any task, or even morphological properties.

Algorithm 1 Evolution+Learning

- 1: INITIALIZE robot population (genotypes + phenotypes with body and brain)
- 2: EVALUATE each robot (evaluation delivers a fitness value)
- 3: while not STOP-EVOLUTION do
- 4: SELECT parents; (based on fitness)
- 5: RECOMBINE+MUTATE parents' bodies; (this delivers a new body genotype)
- 6: RECOMBINE+MUTATE parents' brains; (this delivers a new brain genotype)
- 7: CREATE offspring robot body; (this delivers a new body phenotype)
- 8: CREATE offspring robot brain; (this delivers a new brain phenotype)
- 9: INITIALIZE brain(s) for the learning process; (in the new body)
- 10: while not STOP-LEARNING do
- 11: ASSESS offspring; (assessment delivers a reward value)
- 12: GENERATE new brain for offspring;
- 13: end while
- 14: EVALUATE offspring with learned brain; (evaluation delivers a fitness value)

16: SELECT survivors / UPDATE population

17: end while

#### Evolution loop

For the outer evolutionary loop, we use a variant of the well-known  $(\mu + \lambda)$  selection mechanism to update the population. The bodies of the robots are evolved with sexual reproduction while the brains of the robots are evolved with asexual reproduction.

Body - sexual reproduction: The body of every new offspring is created through recombination and mutation of the genotypes of its parents. Parents are selected from the current generation using binary tournaments with replacement. We perform two tournaments in which two random potential parents are selected. In each tournament the potential parents are compared, the one with the highest fitness wins the tournament and becomes a parent.

Brain - asexual reproduction: The brain genotype of the best-performing parent is mutated (without recombination) before being inherited by its offspring. This choice is based on preliminary experiments that indicated that asexual brain reproduction is the better method, as it resulted in robots with higher fitness.

#### Learning loop

For the inner learning loop which is to search in the space of brain configurations and fine-tune the parameters, we have chosen Reversible Differential Evolution (RevDE) as a learner, because in a recent study on modular robots [ $^{69}$ ], it was demonstrated that RevDE [ $^{70,71}$ ], an altered version of Differential Evolution, performs and generalizes well across various morphologies. This algorithm works as follows:

- 1. Initialize a population with  $\mu$  samples (*n*-dimensional vectors),  $\mathscr{P}_{\mu}$ .
- 2. Evaluate all  $\mu$  samples.
- 3. Apply the reversible differential mutation operator and the uniform crossover operator.
- *The reversible differential mutation operator*: Three new candidates are generated by randomly picking a triplet from the population,  $(\mathbf{w}_i, \mathbf{w}_j, \mathbf{w}_k) \in \mathscr{P}_{\mu}$ , then all three individuals are perturbed by adding a scaled difference in the following manner:

$$\mathbf{v}_{1} = \mathbf{w}_{i} + F \cdot (\mathbf{w}_{j} - \mathbf{w}_{k})$$

$$\mathbf{v}_{2} = \mathbf{w}_{j} + F \cdot (\mathbf{w}_{k} - \mathbf{v}_{1})$$

$$\mathbf{v}_{3} = \mathbf{w}_{k} + F \cdot (\mathbf{v}_{1} - \mathbf{v}_{2})$$
(3)

<sup>15:</sup> UPDATE brain genotype



**Figure 14.** Evolution + Learning framework. This is a general framework for optimizing robots via two interacting adaptive processes. The evolutionary loop (left) optimizes robot morphologies and controllers simultaneously using genotypes that encode both morphologies and controllers. The learning loop (yellow box inside the Evaluation step of the evolutionary loop) optimizes the controller for a given morphology. Note that in general the fitness measure used within the evolutionary loop need not be the same as the quality measure used inside the learning method. With the red lines, it is the Lamarckian learning mechanism which allows the phenotype of the brain coded back to the genotype and pass it to the next generation. Without the red lines, it is the Darwinian learning mechanism. cf. Algorithm 1

where  $F \in R_+$  is the scaling factor. New candidates  $y_1$  and  $y_2$  are used to calculate perturbations using points outside the population. This approach does not follow the typical construction of an EA where only evaluated candidates are mutated.

The uniform crossover operator: Following the original DE method [<sup>72</sup>], we first sample a binary mask  $\mathbf{m} \in \{0, 1\}^D$  according to the Bernoulli distribution with probability *CR* shared across *D* dimensions, and calculate the final candidate according to the following formula:

$$\mathbf{u} = \mathbf{m} \odot \mathbf{w}_n + (1 - m) \odot \mathbf{w}_n \tag{4}$$

Following general recommendations in literature [<sup>73</sup>] to obtain stable exploration behaviour, the crossover probability CR is fixed to a value of 0.9 and according to the analysis provided in [<sup>70</sup>], the scaling factor F is fixed to a value of 0.5.

- 4. Perform a selection over the population based on the fitness value and select  $\mu$  samples.
- 5. Repeat from step (2) until the maximum number of iterations is reached.

As explained above, we apply RevDE here as a learning method for 'newborn' robots. In particular, it will be used to optimize the weights of the CPGs of our modular robots for the tasks during the Infancy stage. The initial population of X = 10 weight vectors for RevDE is created by using the inherited brain of the given robot. Specifically, the values of the inherited weight vector are altered by adding Gaussian noise to create mutant vectors and the initial population consists of nine such mutants and the vector with the inherited weights.

#### Task and Fitness function

Point navigation requires feedback (coordinates) from the environment passing to the controller to steer the robot. The coordinates are used to obtain the angle between the current position and the target. If the target is on the right, the right joints are slowed down and vice versa.

A robot is spawned at the centre of a flat arena  $(10 \times 10 \text{ m2})$  to reach a sequence of target points  $P_1, ..., P_N$ . In each evaluation, the robot has to reach as many targets in order as possible. Success in this task requires the ability to move fast to reach one target and then quickly change direction to another target in a short duration. A target point is considered to be

reached if the robot gets within 0.01 meters from it. To keep runtimes within practically acceptable limits, we set the simulation time per evaluation to be 40 seconds which allows robots to reach at least 2 targets  $P_1(1,-1)$ ,  $P_2(0,-2)$ .

The data collected from the simulator is the following:

- The coordinates of the core component of the robot at the start of the simulation are approximate to  $P_0(0,0)$ ;
- The coordinates of the robot, sampled during the simulation at 5Hz, allowing us to plot and approximate the length of the followed path;
- The coordinates of the robot at the end of the simulation  $P_T(x_T, y_T)$ ;
- The coordinates of the target points  $P_1(x_1, y_1)$ ...  $P_n(x_n, y_n)$ .
- The coordinates of the robot, sampled during the simulation at 5Hz, allow us to plot and approximate the length of the path *L*.

The fitness function is designed to maximize the number of targets reached and to minimize the path length.

$$F = \sum_{i=1}^{k} dist(P_i, P_{i-1}) + (dist(P_k, P_{k-1}) - dist(P_T, P_k)) - \omega \cdot L$$
(5)

where *k* is the number of target points reached by the robot at the end of the evaluation, and *L* is the path travelled. The first term of the function is a sum of the distances between the target points the robot has reached. The second term is necessary when the robot has not reached all the targets and it calculates the distance travelled toward the next unreached target. The last term is used to penalize longer paths and  $\omega$  is a constant scalar that is set to 0.1 in the experiments. E.g., a robot just reached 2 targets, the maximum fitness value will be  $dist(P_1, P_0) + (dist(P_2, P_1) - dist(P_2, P_2)) - 0.1 \cdot L = \sqrt{2} + \sqrt{2} - 0.2 \cdot \sqrt{2} \approx 2.54$  (*L* is shortest path length to go through  $P_1$  and  $P_2$  which is equal to  $2 \cdot \sqrt{2}$ ).

#### **Experimental setup**

We use a Mujoco simulator-based wrapper called Revolve2 (https://github.com/ci-group/revolve2) to run experiments. For the  $(\mu + \lambda)$  selection in the outer evolutionary loop, we set  $\mu = 50$  and  $\lambda = 25$ . The evolutionary process is terminated after 30 generations. Therefore, we perform  $(25 + 25 \cdot 30)$  robots= 775 fitness evaluations for each evolutionary loop.

For the inner learning loop, we apply RevDE on each robot body resulting in 280 extra fitness evaluations. This number is based on the learning assessment from RevDE for running 10 initial samples with 10 iterations using  $(\mu + \lambda)$  selection. The first iteration contains 10 samples, and from the second iteration onwards each iteration creates 30 new candidates, resulting in a total of  $10+30 \cdot (10-1) = 280$  evaluations.

In our research, the fitness measure to drive evolution and the performance measure to drive learning are the same by design. Thus, we use the same test procedure, simulating one robot for 40 simulated seconds for the point navigation task, for the evolutionary as well as the learning trials.

To sum up, for running the experiments we perform  $775 \cdot 280 \cdot 2 = 434,000$  fitness evaluations which amount to  $434,000 \cdot 40/60/60 = 4,822$  hours of simulated time. To get a robust assessment of the performance all the experiments are repeated 20 times independently. In practice, it takes about 4.25 days to run 2 runs in parallel on 2 64-core processors. The experimental parameters we used in the experiments are described in Table 1 in the Supplementary Information section.

## Data availability statement

The code for replicating this work and carrying out the experiments is available online: https://shorturl.at/epzGS. A short video providing a visual overview of our research is available at https://shorturl.at/ahETW.

# References

- 1. Nolfi, S., Floreano, D. & Floreano, D. D. Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines (MIT press, 2000).
- 2. Doncieux, S., Bredeche, N., Mouret, J.-B. & Eiben, A. E. Evolutionary robotics: what, why, and where to. *Front. Robotics* AI 2, 4 (2015).
- **3.** Floreano, D. & Mondada, F. Evolutionary neurocontrollers for autonomous mobile robots. *Neural networks* **11**, 1461–1478 (1998).
- 4. Capi, G., Pojani, G. & Kaneko, S.-I. Evolution of task switching behaviors in real mobile robots. In 2008 3rd International Conference on Innovative Computing Information and Control, 495–495, DOI: 10.1109/ICICIC.2008.261 (2008).
- 5. Sims, K. Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, 15–22 (ACM, 1994).
- 6. Auerbach, J. E. & Bongard, J. C. On the relationship between environmental and morphological complexity in evolved robots. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, 521–528 (ACM, 2012).
- Weel, B., Crosato, E., Heinerman, J., Haasdijk, E. & Eiben, A. E. A Robotic Ecosystem with Evolvable Minds and Bodies. In 2014 IEEE International Conference on Evolvable Systems, 165 – 172 (IEEE Press, Piscataway, NJ, 2014).
- 8. Lipson, H., SunSpiral, V., Bongard, J. C. & Cheney, N. On the difficulty of co-optimizing morphology and control in evolved virtual creatures. *Artif. Life* 226–233 (2016).
- 9. Cheney, N., Bongard, J., SunSpiral, V. & Lipson, H. Scalable co-optimization of morphology and control in embodied machines. *J. Royal Soc. Interface* **15** (2018).
- Miras, K. & Eiben, A. E. Effects of environmental conditions on evolved robot morphologies and behavior. In *Proceedings* of the Genetic and Evolutionary Computation Conference, GECCO '19, 125–132, DOI: 10.1145/3321707.3321811 (Association for Computing Machinery, New York, NY, USA, 2019).
- Hockings, N. & Howard, D. New biological morphogenetic methods for evolutionary design of robot bodies. *Front. Bioeng. Biotechnol.* 8, DOI: 10.3389/fbioe.2020.00621 (2020).
- Stensby, E. H., Ellefsen, K. O. & Glette, K. Co-optimising robot morphology and controller in a simulated open-ended environment. In *Applications of Evolutionary Computation: 24th International Conference, EvoApplications 2021, Held as Part of EvoStar 2021, Virtual Event, April 7–9, 2021, Proceedings*, 34–49, DOI: 10.1007/978-3-030-72699-7\_3 (Springer-Verlag, Berlin, Heidelberg, 2021).
- 13. Medvet, E., Bartoli, A., Pigozzi, F. & Rochelli, M. Biodiversity in evolved voxel-based soft robots. *Proc. 2021 Genet. Evol. Comput. Conf.* 129–137 (2021).
- 14. Nygaard, T. F., Martin, C. P., Torresen, J., Glette, K. & Howard, D. Real-world embodied ai through a morphologically adaptive quadruped robot. *Nat. Mach. Intell.* 1, 12–19 (2021).
- 15. Nolfi, S. & Floreano, D. Learning and evolution. Auton. robots 7, 89-113 (1999).
- 16. Floreano, D. & Mondada, F. Evolution of plastic neurocontrollers for situated agents. In *Proc. of The Fourth International Conference on Simulation of Adaptive Behavior (SAB), From Animals to Animats* (ETH Zürich, 1996).
- 17. Harvey, I., Husbands, P., Cliff, D., Thompson, A. & Jakobi, N. Evolutionary robotics: the sussex approach. *Robotics autonomous systems* **20**, 205–224 (1997).
- **18.** Schembri, M., Mirolli, M. & Baldassarre, G. Evolution and learning in an intrinsically motivated reinforcement learning robot. In *Proceedings of the 9th European Conference on Advances in Artificial Life*, vol. 4648, 294–303 (2007).
- 19. Sproewitz, A., Moeckel, R., Maye, J. & Ijspeert, A. J. Learning to move in modular robots using central pattern generators and online optimization. *The Int. J. Robotics Res.* 27, 423–443 (2008).
- **20.** Bellas, F., Duro, R. J., Faiña, A. & Souto, D. Multilevel darwinist brain (mdb): Artificial evolution in a cognitive architecture for real robots. *IEEE Transactions on autonomous mental development* **2**, 340–354 (2010).
- **21.** Khazanov, M., Jocque, J. & Rieffel, J. Evolution of locomotion on a physical tensegrity robot. In *ALIFE 14: The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, 232–238 (MIT Press, 2014).
- 22. Ruud, E. L., Samuelsen, E. & Glette, K. Memetic robot control evolution and adaption to reality. 2016 IEEE Symp. Ser. on Comput. Intell. SSCI 2016 (2017).

- 23. Schaff, C., Yunis, D., Chakrabarti, A. & Walter, M. R. Jointly learning to construct and control agents using deep reinforcement learning. In 2019 International Conference on Robotics and Automation (ICRA), 9798–9805 (IEEE, 2019).
- 24. Hwangbo, J. et al. Learning agile and dynamic motor skills for legged robots. Sci. Robotics 4, eaau5872 (2019).
- 25. Le Goff, L. K. *et al.* Sample and time efficient policy learning with CMA-ES and Bayesian Optimisation. In *The 2020 Conference on Artificial Life*, January, 2020 (2020).
- **26.** Lehman, J. & Stanley, K. O. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, 211–218 (2011).
- 27. Jelisavcic, M. et al. Real-world evolution of robot morphologies: A proof of concept. Artif. life 23, 206–235 (2017).
- 28. Kriegman, S., Cheney, N. & Bongard, J. How morphological development can guide evolution. Sci. reports 8, 1–10 (2018).
- 29. Aguilar, L., Bennati, S. & Helbing, D. How learning can change the course of evolution. *Plos one* 14, e0219502 (2019).
- 30. Howard, D. et al. Evolving embodied intelligence from materials to machines. Nat. Mach. Intell. 1, 12–19 (2019).
- **31.** Miras, K., De Carlo, M., Akhatou, S. & Eiben, A. E. Evolving-controllers versus learning-controllers for morphologically evolvable robots. In *Applications of Evolutionary Computation*, 86–99 (2020).
- **32.** Gupta, A., Savarese, S., Ganguli, S. & Fei-Fei, L. Embodied Intelligence via Learning and Evolution. *Nat. Commun.* **12** (2021).
- **33.** Goff, L. K. L. & Hart, E. On the challenges of jointly optimising robot morphology and control using a hierarchical optimisation scheme. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 1498–1502 (2021).
- 34. Hart, E. & Le Goff, L. K. Artificial evolution of robot bodies and control: on the interaction between evolution, learning and culture. *Philos. Transactions Royal Soc. B* 377, 20210117 (2022).
- 35. Luo, J., Stuurman, A., Tomczak, J. M., Ellers, J. & Eiben, A. E. The effects of learning in morphologically evolving robot systems. *Front. Robotics AI* 5 (2022).
- Wiser, M. J., Ribeck, N. & Lenski, R. E. Long-term dynamics of adaptation in asexual populations. *Science* 342, 1364–1367 (2013).
- **37.** Long, J. Darwin's devices: What evolving robots can teach us about the history of life and the future of technology (Basic Books (AZ), 2012).
- 38. Maynard Smith, J. Byte-sized evolution. Nature 355, 772–773 (1992).
- **39.** Montanier, J.-M. & Bredeche, N. Surviving the tragedy of commons: emergence of altruism in a population of evolving autonomous agents. In *European Conference on Artificial Life* (2011).
- **40.** Waibel, M., Floreano, D. & Keller, L. A quantitative test of hamilton's rule for the evolution of altruism. *PLoS biology* **9**, e1000615 (2011).
- **41.** Solomon, M., Soule, T. & Heckendorn, R. B. A comparison of a communication strategies in cooperative learning. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, 153–160 (2012).
- **42.** Floreano, D., Mitri, S., Magnenat, S. & Keller, L. Evolutionary conditions for the emergence of communication in robots. *Curr. biology* **17**, 514–519 (2007).
- **43.** Mitri, S., Floreano, D. & Keller, L. The evolution of information suppression in communicating robots with conflicting interests. *Proc. Natl. Acad. Sci.* **106**, 15786–15790 (2009).
- 44. Wischmann, S., Floreano, D. & Keller, L. Historical contingency affects signaling strategies and competitive abilities in evolving populations of simulated robots. *Proc. Natl. Acad. Sci.* **109**, 864–868 (2012).
- **45.** Bongard, J. Morphological change in machines accelerates the evolution of robust behavior. *Proc. Natl. Acad. Sci.* **108**, 1234–1239 (2011).
- **46.** Auerbach, J. E. & Bongard, J. C. Environmental influence on the evolution of morphological complexity in machines. *PLoS computational biology* **10**, e1003399 (2014).
- 47. Olson, R. S., Hintze, A., Dyer, F. C., Knoester, D. B. & Adami, C. Predator confusion is sufficient to evolve swarming behaviour. *J. The Royal Soc. Interface* **10**, 20130305 (2013).
- **48.** Zirkle, C. The early history of the idea of the inheritance of acquired characters and of pangenesis. *Transactions Am. Philos. Soc.* **35**, 91–151 (1946).

- **49.** Burkhardt, R. W. Lamarck, evolution, and the inheritance of acquired characters. In *Genetics*, 793–805, DOI: https://doi.org/10.1534/genetics.113.151852 (2013).
- 50. Bossdorf, O., Richards, C. L. & Pigliucci, M. Epigenetics for ecologists. Ecol. letters 11, 106–115 (2008).
- **51.** Goff, L. K. L. *et al.* Morpho-evolution with learning using a controller archive as an inheritance mechanism. *http://arxiv.org/abs/2104.04269* (2021).
- **52.** Mingo, J. M., Aler, R., Maravall, D. & de Lope, J. Investigations into lamarckism, baldwinism and local search in grammatical evolution guided by reinforcement. *Comput. Informatics* **32**, 595–627 (2013).
- **53.** Whitley, D., Gordon, V. S. & Mathias, K. Lamarckian evolution, the baldwin effect and function optimization. In *Parallel Problem Solving from Nature—PPSN III: International Conference on Evolutionary Computation The Third Conference on Parallel Problem Solving from Nature Jerusalem, Israel, October 9–14, 1994 Proceedings 3, 5–15 (Springer, 1994).*
- 54. Holzinger, A. *et al.* Darwin, lamarck, or baldwin: Applying evolutionary algorithms to machine learning techniques. In 2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), vol. 2, 449–453, DOI: 10.1109/WI-IAT.2014.132 (2014).
- 55. Mikami, S., Wada, M. & Kakazu, Y. Combining reinforcement learning with ga to find co-ordinated control rules for multi-agent system. In *Proceedings of IEEE International Conference on Evolutionary Computation*, 356–361 (IEEE, 1996).
- 56. Ku, K. W. C., Mak, M. W. & Siu, W. C. Adding learning to cellular genetic algorithms for training recurrent neural networks. *IEEE Transactions on Neural Networks* 10, 239–252 (1999).
- 57. Houck, C. R., Joines, J. A., Kay, M. G. & Wilson, J. R. Empirical investigation of the benefits of partial lamarckianism. *Evol. Comput.* 5, 31–60 (1997).
- 58. Castillo, P. A. et al. Lamarckian evolution and the baldwin effect in evolutionary neural networks (2006). cs/0603004.
- **59.** Zhang, C., Chen, J. & Xin, B. Distributed memetic differential evolution with the synergy of lamarckian and baldwinian learning. *Appl. Soft Comput.* **13**, 2947–2959, DOI: https://doi.org/10.1016/j.asoc.2012.02.028 (2013).
- **60.** Elsken, T., Metzen, J. H. & Hutter, F. Efficient multi-objective neural architecture search via lamarckian evolution. *arXiv* preprint arXiv:1804.09081 (2018).
- Nishiwaki, Y., Mukosaka, N., Tanev, I. & Shimohara, K. On the effects of epigenetic programming on the efficiency of incremental evolution of the simulated khepera robot. *J. Robotics, Netw. Artif. Life* 2, 108, DOI: 10.2991/jrnal.2015.2.2.9 (2015).
- **62.** Grefenstette, J. J. Lamarckian learning in multi-agent environments. In *International Conference on Genetic Algorithms* (1991).
- 63. Jelisavcic, M., Glette, K., Haasdijk, E. & Eiben, A. Lamarckian evolution of simulated modular robots. *Front. Robotics AI* 6, 9 (2019).
- **64.** Jelisavcic, M., Miras, K. & Eiben, A. Morphological attractors in darwinian and lamarchian evolutionary robot systems. In 2018 IEEE Symposium Series on Computational Intelligence (SSCI), 859–866 (IEEE, 2018).
- **65.** Miras, K., Haasdijk, E., Glette, K. & Eiben, A. Search space analysis of evolvable robot morphologies. In *International Conference on the Applications of Evolutionary Computation*, 703–718 (Springer, 2018).
- Pawlik, M. & Augsten, N. Tree edit distance: Robust and memory-efficient. *Inf. Syst.* 56, DOI: 10.1016/j.is.2015.08.004 (2015).
- **67.** Auerbach, J. E. *et al.* Robogen: Robot generation through artificial evolution. In *Proceedings of the 14th International Conference on the Synthesis and Simulation of Living Systems, ALIFE 2014*, 136–137 (2014).
- **68.** Stanley, K. O. Compositional pattern producing networks: A novel abstraction of development. *Genet. Program. Evolvable Mach.* **8**, 131–162 (2007).
- **69.** van Diggelen, F., Ferrante, E. & Eiben, A. E. Comparing lifetime learning methods for morphologically evolving robots. In *GECCO '21: Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 93–94 (2021).
- **70.** Tomczak, J. M., Weglarz-Tomczak, E. & Eiben, A. E. Differential Evolution with Reversible Linear Transformations. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*, 205–206 (2020). 2002.02869.
- **71.** Weglarz-Tomczak, E., Tomczak, J. M., Eiben, A. E. & Brul, S. Population-based parameter identification for dynamical models of biological networks with an application to saccharomyces cerevisiae. *Processes* **9**, 98 (2021).

- 72. Storn, R. M. Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces. In *Journal of Global Optimization*, 131–141 (1997).
- 73. Pedersen, M. Good Parameters for Differential Evolution. Evolution 1–10 (2010).

# Author contributions statement

- J.L.: Design, implementation, and execution of experiments. Data analysis, data visualization and paper writing.
  - K.M.: Experimental design, data analysis, contribution to paper writing.
  - J.T.: The manuscript reviewing.
  - A.E.E.: Design, Data analysis, contribution to paper writing.

# **Competing interests**

The authors declare no competing interests.

# **Supplementary Information**

Parameters	Value	Description
Population size	50	Number of individuals per generation
Offspring size	25	Number of offspring produced per generation
Mutation	0.8	Probability of mutation for individuals
Crossover	0.8	Probability of crossover for individual's body
Generations	30	Termination condition for each run
Learning trials	280	Number of the evaluations performed by RevDE on each robot
$\mu$	10	RevDE - Population size
Ν	30	RevDE - New candidates per iteration
λ	10	RevDE - Top-sample size
F	0.5	RevDE - Scaling factor
CR	0.9	RevDE - Crossover probability
Iterations	10	RevDE - Number of iterations
Evaluation time	40	Duration of evaluation in seconds
Tournament size	2	Number of individuals used in the parent selection - (k-tournament)
$\lambda/\mu$	0.5	The ratio used in the survivor selection - $(\mu + \lambda)$
Repetitions	20	Number of repetitions per experiment

## Table 1. Main experiment parameters