

Limits of Lamarckian Evolution Under Pressure of Morphological Novelty

1st Jed R Muff

Department of Computer Science
Vrije Universiteit
Amsterdam, The Netherlands
0009-0004-9934-6403

2nd Karine Miras

Department of Computer Science
Vrije Universiteit
Amsterdam, The Netherlands
0000-0003-4942-3488

3rd A.E. Eiben

Department of Computer Science
Vrije Universiteit
Amsterdam, The Netherlands
0000-0002-3106-4213

Abstract—Lamarckian inheritance has been shown to be a powerful accelerator in systems where the joint evolution of robot morphologies and controllers is enhanced with individual learning. Its defining advantage lies in the offspring inheriting controllers learned by their parents. The efficacy of this option, however, relies on morphological similarity between parent and offspring. In this study, we examine how Lamarckian inheritance performs when the search process is driven toward high morphological variance, potentially straining the requirement for parent-offspring similarity. Using a system of modular robots that can evolve and learn to solve a locomotion task, we compare Darwinian and Lamarckian evolution to determine how they respond to shifting from pure task-based selection to a multi-objective pressure that also rewards morphological novelty. Our results confirm that Lamarckian evolution outperforms Darwinian evolution when optimizing task-performance alone. However, introducing selection pressure for morphological diversity causes a substantial performance drop, which is much greater in the Lamarckian system. Further analyses show that promoting diversity reduces parent-offspring similarity, which in turn reduces the benefits of inheriting controllers learned by parents. These results reveal the limits of Lamarckian evolution by exposing a fundamental trade-off between inheritance-based exploitation and diversity-driven exploration.

Index Terms—Evolutionary Robotics, Lamarckian Evolution, Limits of Lamarckian Evolution, Morphological Evolution, Evolution and Learning

I. INTRODUCTION

Evolutionary robotics aims to automatically design robot morphologies and control systems through artificial evolution, mimicking natural selection to discover effective solutions without extensive human engineering [1]. A central challenge in this field is the joint optimization of body plans and controllers, a problem compounded by the vast search space and the need for each candidate robot to undergo costly learning or adaptation during its ‘lifetime’ before its evolutionary fitness can be assessed.

Lamarckian evolution offers a compelling mechanism to accelerate this process by allowing offspring to inherit learned traits from their parents. In evolutionary robotics, this typically takes the form of *warm starting*: initializing an offspring’s controller with parameters learned by its parent during lifetime learning, rather than starting from scratch [2]. A substantial body of work has shown that such Lamarckian mechanisms can significantly speed up learning and improve evolutionary

performance across a range of robotic systems and learning algorithms [2], [3], [4], [5], [6], [7], [8], [9].

However, the effectiveness of Lamarckian inheritance relies on a critical assumption: that offspring are sufficiently similar to their parents for inherited controllers to remain compatible. Jelisavcic et al. [2] demonstrated that the benefits of controller inheritance strongly depend on parent-offspring morphological similarity, with performance gains diminishing as morphological divergence increases. In many prior studies, this assumption is implicitly satisfied, as evolution tends to converge toward relatively homogeneous morphologies over time.

Another important factor shaping the limits of Lamarckian evolution is environmental variability. Learning typically adapts agents to local, transient conditions, whereas evolution operates on longer time scales and targets more persistent environmental structure. As a result, inheriting a controller adapted to a previous environment can be counterproductive when conditions change. This intuition is supported by recent work showing that following certain environmental transitions, particularly shifts from easier to more challenging tasks, Lamarckian populations can temporarily underperform Darwinian ones [9]. These findings suggest that Lamarckian inheritance is most effective in settings where disruptive forces, such as abrupt environmental change, are limited.

Taken together, this evidence indicates that Lamarckian evolution performs best in regimes characterized by relative stability, both in morphology and environment. We interpret this as suggesting that Lamarckism’s strength relies on the absence of forces that introduce substantial transgenerational variability. Such forces may include environmental non-stationarity, as well as evolutionary pressures that explicitly violate the parent-offspring similarity assumption under which controller inheritance is effective.

In this work, we focus on the latter. Specifically, we hypothesize that *explicitly selecting for morphological diversity will reduce the benefits of Lamarckian evolution, while leaving the benefits of Darwinian evolution largely unaffected*. Although parent-offspring similarity is advantageous in any evolutionary system, it is particularly critical in the Lamarckian case, where inherited controllers must remain compatible with the offspring’s body.

To test this hypothesis, we present a systematic study comparing Darwinian and Lamarckian evolution under differing evolutionary pressures. We evaluate four conditions: Darwinian versus Lamarckian inheritance, each combined with either a simple locomotion objective or an objective that additionally rewards morphological diversity. By analyzing population fitness and learning dynamics, our goal is to determine whether the benefits of Lamarckian inheritance critically depend on morphological similarity, or whether they persist under evolutionary pressures that actively promote morphological change.

The code for this work is available at: <https://tinyurl.com/mpfak3xt>. Videos showcasing interesting evolved individuals under various experimental conditions can be viewed here: <https://tinyurl.com/3rj3w27k>.

II. RELATED WORK

Lamarckian inheritance has been extensively studied in evolutionary robotics as a mechanism for accelerating learning through controller transfer. Early and contemporary work has consistently shown that initializing offspring controllers with parameters learned by their parents can lead to faster convergence and improved performance in static environments [5], [10], [11]. More recent studies have extended these findings to modular robots and complex co-evolutionary settings, further demonstrating the potential of Lamarckian evolution to reduce lifetime learning costs [2], [4].

A few studies have also examined Lamarckian mechanisms under non-stationary conditions. Luo et al. [9] investigated environmental change and showed that although Lamarckian populations may initially suffer after abrupt task transitions, they often recover more quickly and achieve higher long-term performance than Darwinian populations. Similar trade-offs between rapid exploitation and adaptability have been observed in earlier evolutionary computation work, where Lamarckian inheritance was found to reduce robustness in dynamic environments [11].

Beyond evolutionary robotics, Lamarckian principles have been revisited in modern learning systems. Applications include neural architecture search [12] and soft robotic co-design [13], where variants of weight inheritance and transfer learning are used to mitigate mismatches between inherited controllers and changing morphologies or architectures. These approaches demonstrate that Lamarckian ideas remain relevant, but often rely on mechanisms specifically designed to preserve functional continuity across structural change.

In parallel, a substantial body of work in evolutionary robotics has focused on promoting morphological diversity. Novelty search and quality-diversity algorithms explicitly encourage exploration of diverse body plans, often leading to more creative, adaptable, or evolvable designs [14], [15], [16]. However, morphological variation can disrupt learned behaviors, and prior studies have shown that evolvability depends on both environmental structure [17] and intrinsic morphological properties [18]. These approaches typically




Fig. 1: ARIEL framework phenotype space including three components, two passive and one active which is controlled by the brain. On the right demonstrates what a single individual would look like in real life.

assume Darwinian inheritance and do not incorporate lifetime learning or controller transfer.

Despite extensive work on Lamarckian inheritance and on morphological diversity in isolation, their interaction has received little direct attention. Existing studies suggest that Lamarckian benefits rely on parent–offspring similarity, but it remains unclear whether these benefits persist when evolution actively promotes morphological diversity. Our work addresses this gap by explicitly combining diversity pressure with Lamarckian inheritance, enabling a systematic evaluation of warm starting under conditions where parent–offspring similarity is intentionally weakened.

III. METHODS

A. Robot Platform and Task Definition

We use the ARIEL framework (previously known as Revolve) similar to that used in previous works [2], [8], [9] for simulated modular robots, implemented in the MuJoCo physics simulator. Each robot consists of a fixed core module and a variable number of rigid brick modules and active hinge joints. Figure 1 shows what each component looks like in real life and in simulation and how these parts can be combined together to make an individual.

The locomotion task requires robots to maximize forward displacement on a flat terrain over a 30-second period. The evaluation proceeds in two phases: first, a 5-second settling phase where the robot is spawned above ground and the brain sends no commands, allowing the robot to settle; then, a 30-second active phase where the brain controls the robot’s actuators. The fitness function is calculated as the distance traveled by the robot’s core module from its starting position along the positive x-axis, with two penalty terms applied. First, a penalty is applied based on the robot’s initial height (measured before the settling phase begins) to discourage designs that exploit high starting positions. Second, a penalty of 0.005 is applied for each instance in which the hinge part

touches the ground during evaluation, steering the evolutionary process away from behaviors that exploit physics glitches in simulation. If the hinge part touches the ground more than 200 times during a single evaluation, the fitness is set to -1 regardless of distance traveled, as this indicates physics exploitation for forward propulsion.

$$f = \begin{cases} -1 & \text{if } c_{\text{hinge}} > 200 \\ d - l - 0.005 \cdot c_{\text{hinge}} & \text{otherwise} \end{cases} \quad (1)$$

where f is the fitness value, d is the Euclidean distance traveled by the robot's core module, l is the penalty based on the initial height, c_{hinge} is the number of contacts between the hinge parts and the ground.

B. Evolutionary Algorithm

We use the $(\mu + \lambda)$ evolution strategy as outlined in Algorithm 1. Some important details to highlight include the following. 80% of the population uses the mutation operator, while 20% uses the crossover operator based on initial parameter tuning. The initial population P_0 consists of 10 random generated small trees (depth 2), 10 medium trees (depth 3), and 10 large trees (depth 4) randomly generated. After crossover, the offspring undergo the mutation operator. The weight inheritance of the artificial neural network during crossover comes from the closest parent – determined using the tree edit distance regarding the body. This ensures that the offspring acquire the brain that is most likely to work for its new body. Although fitness is evaluated once per individual, novelty scores are recomputed at each generation for all individuals in the population and offspring pool.

C. Representation and Reproduction

1) *Body*: The body is encoded as a tree structure rooted at the core, with nodes corresponding to modules and edges representing physical connections (Figure 1). This tree simultaneously defines the body genotype and is evolved using standard tree-based genetic operators, namely mutation and crossover. To ensure morphological feasibility and computational tractability, several constraints are imposed. First, the total number of body parts is capped at 25, with at most 12 active hinges. If a genetic operation produces a morphology exceeding this hinge limit, excess hinges are pruned using a breadth-first traversal of the tree. These constraints limit morphological complexity while maintaining reasonable simulation times. Second, a minimum of four hinges is required; morphologies that do not satisfy this condition are assigned a fitness of zero, preventing computational resources from being spent on infeasible or degenerate solutions.

To ensure physical realizability, all generated tree genotypes, whether produced through random initialization, mutation, or crossover, undergo a pruning procedure to eliminate colliding components. This pruning is implemented via breadth-first search traversal of the genotype tree. During traversal, the morphology is incrementally constructed using the Python-FCL collision detection library, with each prospective body part being evaluated for spatial interference with all

Algorithm 1: $(\mu + \lambda)$ Evolution Strategy. Optional components for **Lamarckian inheritance** and **novelty search** are shown as shaded boxes; these blocks are included only in the corresponding algorithm variants.

Parameters: $\mu=30, \lambda=30, n_{\text{mut}}=24, n_{\text{cross}}=6$

- 1 $P_0 \leftarrow$ Random population of μ tree genotypes
- 2 EvaluatePopulation(P_0)
- 3 **ComputeNovelty**(P_0)
- 4 **for** $g = 1 \rightarrow n_{\text{gen}}$ **do**
- 5 $O_{\text{mut}} \leftarrow$ Mutate(TournamentSelect(P_{g-1}, n_{mut}))
- 6 $O_{\text{cross}} \leftarrow$ Crossover(TournamentSelect($P_{g-1}, 2n_{\text{cross}}$))
- 7 $O_g \leftarrow O_{\text{mut}} \cup O_{\text{cross}}$
- 8 **for** $c \in O_g$ **do**
 $\perp$ InheritWeights($c, \text{ClosestParent}(c)$)
- 9 EvaluatePopulation(O_g)
- 10 **UpdateNoveltyArchive**($P_{g-1} \cup O_g$)
 UpdateFitnessWithNovelty($P_{g-1} \cup O_g$)
- 11 $P_g \leftarrow \text{Top}_{\mu}(P_{g-1} \cup O_g)$

Tree Body Representation

Digital Twin




Fig. 2: Representation used in this work. The body employs a tree-based representation, while the brain uses an ordered array of parameters that maps to an artificial neural network (ANN). Note that the ANN depicted in this diagram is illustrative only and does not reflect the actual ANN architecture implemented in this paper (see Section III-C2).

previously placed components. When a collision is detected, the offending node and all of its descendant nodes are removed from the tree, ensuring that only physically valid morphologies are retained for evaluation.

2) *Brain*: The brain phenotype is an artificial neural network (ANN) with three hidden layers containing 32, 16, and 32 neurons, respectively. This architecture was determined through hyperparameter optimization conducted on a reference gecko morphology with 6 hinges. The input layer consists of $9 + h$ neurons, where h denotes the number of hinges in the morphology. These inputs encode the core's Euler orientation (3 values), linear velocity (3 values), angular velocity (3 values), and the current joint positions for all h hinges. The output layer contains h neurons, each specifying the target position of its corresponding joint.

This phenotype is represented in genotype space as one long ordered array of parameters. The total number of trainable parameters (weights and biases) in the network is given by:

$$P = (n_0 \times n_1 + n_1) + (n_1 \times n_2 + n_2) + (n_3 \times n_3 + n_3) + (n_3 \times n_4 + n_4) \quad (2)$$

For the architecture used in this study, where $n_0 = 9 + h$, $n_1 = 32$, $n_2 = 16$, $n_3 = 32$, and $n_4 = h$, this simplifies to the following.

$$P = 1344 + 33h \quad (3)$$

3) *Learning*: Weight optimization is performed using the Covariance Matrix Adaptation Evolution Strategy (CMA-ES). CMA-ES optimizes parameters by iteratively sampling candidate solutions from a multivariate normal distribution defined by a mean vector μ , a global step size σ , and a covariance matrix $\mathbf{C}$. After evaluating sampled candidates, the mean is updated toward higher performing solutions, while the covariance matrix is adapted to capture correlations between parameters that lead to successful search directions. The step size controls the overall exploration scale and is adapted online to balance exploration and exploitation. Together, $(\mu, \sigma, \mathbf{C})$ define a learned search distribution that encodes both promising regions of the parameter space and parameter dependencies.

We use the default parameters provided by the python-cmaes library with a population size of 20 and 1000 function evaluations per individual. Preliminary testing on handmade designs indicated that 1000 function evaluations were sufficient with minimal learning time.

D. Lifetime Inheritance

In CMA-ES, optimization proceeds by adapting a search distribution parameterized by a mean vector μ , a global step size σ , and a covariance matrix $\mathbf{C}$. When an offspring is created, its morphology may differ from that of its parent. In this case, we transfer and transform the parent's learned distribution parameters to account for the morphological change and to provide a warm start for optimization. This Lamarckian inheritance allows offspring to reuse information about promising regions of the parameter space and parameter correlations, in contrast to Darwinian evolution, where each offspring reinitializes $(\mu, \sigma, \mathbf{C})$.

1) *Covariance Matrix*: The covariance matrix $\mathbf{C}$ encodes correlations between controller parameters that have proven beneficial during the parent's lifetime. When morphological changes alter the neural controller's architecture (e.g., added or removed sensors or actuators), a correspondence mapping aligns overlapping parent and offspring parameters. The correspondence mapping defines a deterministic alignment between parent and offspring controller parameters based on shared morphological components, identifying which parameters represent the same functional elements across morphologies. The correlations between the shared parameters are preserved, as encoded by the CMA-ES covariance matrix $\mathbf{C}$, while the new parameters are initialized as uncorrelated dimensions:

$$C_o[i, j] = \begin{cases} C_p[i_p, j_p], & \text{if } w_i, w_j \text{ correspond to parent} \\ & \text{weights} \\ \delta_{ij}, & \text{otherwise} \end{cases} \quad (4)$$

where δ_{ij} is the Kronecker delta, initializing new parameters with unit variance and zero covariance. This preserves learned dependencies in the search distribution while ensuring $\mathbf{C}_o$ remains symmetric and positive definite.

2) *Step Size and Mean Adaptation*: The mean vector μ represents the current center of the search distribution, while the step size σ controls its overall scale. For inherited parameters, the offspring initializes its mean from the parent's mean; newly introduced parameters are initialized at zero. To account for morphological change, the step size is blended with the initial step size based on the relative change in controller dimensionality:

$$\sigma_o = (1 - \alpha)\sigma_p + \alpha\sigma_{\text{init}}, \quad \alpha = \frac{|d_o - d_p|}{d_o}, \quad (5)$$

$$\mu_o[i] = \begin{cases} \mu_p[i_p], & \text{if } w_i \text{ corresponds to a parent} \\ & \text{parameter} \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

3) *Parent Selection*: When there are multiple candidate parents, inheritance is taken from the parent with the most similar morphology, measured using tree edit distance. This maximizes the relevance of the inherited search distribution, particularly the covariance structure that captures parameter correlations.

IV. EXPERIMENTAL SETUP

A. Methodology and Procedure

We employ a 2×2 factorial experimental design to systematically investigate the effects of inheritance mechanism and fitness criterion on evolutionary dynamics and outcomes. The inheritance mechanism factor comprises two methods: Darwinian inheritance, wherein offspring controllers undergo random reinitialization, and Lamarckian inheritance, wherein offspring inherit parental learned controllers. These mechanisms are operationalized as described in Section 3. The fitness criterion factor similarly comprises two levels:

pure locomotion fitness, measured as translational distance, and **combined locomotion-novelty fitness**, incorporating both locomotion performance and morphological diversity.

Morphological novelty is quantified through vectors composed of morphological descriptors, following the methodology established in our prior work [19]. Specifically, the morphological descriptors we use are: branching, limbs, length of limbs, coverage, joints, proportion, symmetry, size, all normalized between 0 and 1. For each individual i , novelty N_i is computed as the Euclidean distance to its nearest neighbor in morphological descriptor space, thereby incentivizing exploration of diverse phenotypic configurations. The combined fitness function is formulated as:

$$f_{\text{combined}} = f_{\text{locomotion}} \cdot f_{\text{novelty}} \quad (7)$$

Each experimental condition executes for 50 generations with a population size of 30 offspring per generation. Individual fitness evaluation requires 1,000 learning episodes, yielding 50 generations $\times$ 30 offspring $\times$ 1,000 evaluations = 1,500,000 simulation episodes, plus 30 initial population evaluations, totaling 1,530,000 time-steps per experimental run or 1,530 evaluations. With an estimated computational cost of approximately 1 minute per evaluation, each run requires approximately 48 compute-hours. We repeated each experiment with different seeds for 20 runs, generating data for 80 runs in total.

B. Performance Metrics and Statistical Analysis

Experimental outcomes are evaluated using a set of performance metrics. Population fitness progression is tracked via the mean and maximum fitness across generations to characterize evolutionary progress. To assess structural inheritance, parent—offspring morphological similarity is quantified as the Euclidean distance between parental and offspring descriptors, indicating the degree to which morphology is preserved through reproduction. Initial locomotion performance is defined as the first fitness evaluation of the offspring, conducted immediately after inheriting the weight of the parents. Comparisons of locomotion fitness Between-conditions of the final-generation (mean per run, 20 runs each) were performed using Welch’s t-test (two-sided, unequal variances assumed), while differences in initial fitness (Figures 3–4) were assessed using the Mann–Whitney U test (two-sided).

V. RESULTS AND DISCUSSION

We start by analysing the performance of the four searches on the task of locomotion (Fig. 3). All searches succeed in discovering robots that locomote effectively, with the rate of improvement decreasing substantially after approximately generation 30. The Lamarckian evolution outperforms Darwinian evolution in the search for locomotion performance ($p=0.039$), consistent with prior work showing that Lamarckian inheritance accelerates learning. However, when morphological novelty is introduced as part of the fitness, the advantage of Lamarckian evolution in comparison to Darwinian evolution




Fig. 3: Average locomotion performance ($f_{\text{locomotion}}$) per generation. Higher is better.




Fig. 4: Average difference in locomotion performance. Lower is better. Performance differences between pure locomotion only and combined locomotion-novelty fitness, comparing Darwinian and Lamarckian setups. **The introduction of novelty has a more detrimental effect on Lamarckian evolution.**

for locomotion performance is no longer significant ($p=0.136$). While the Lamarckian system presents a significant reduction in locomotion performance under the pressure for morphological novelty ($p=0.016$), this reduction in the Darwinian system under the same pressure is insignificant ($p=0.071$).

Figure 4 makes this effect explicit by showing the difference in performance between the searches that maximize locomotion and the searches that maximize for locomotion and novelty. The results directly support our hypothesis: *the Lamarckian system suffers a substantially larger performance drop when novelty becomes part of the fitness, whereas the Darwinian system is almost unaffected.*

Now that the hypothesis has been confirmed, we conduct a deeper analysis to evaluate whether the underlying intuition is valid: the Lamarckian system is more affected by novelty because it depends more heavily on similarity – a dependence that follows from assuming no external pressure toward morphological diversity. If this intuition is correct, then three assumptions must hold, connected as follows:

- (1) The Lamarckian system ‘prefers’ similar bodies more




Fig. 5: Mean parent-child morphological dissimilarity across generations (lower means more similar). Measured as the Euclidean distance between parental and offspring descriptors. **Lamarckian evolution exhibits a bias toward parent-offspring similarity, but this breaks down when novelty is introduced into the fitness function.**

than the Darwinian system – meaning that it works better with similar bodies (Fig. 6, comparisons between boxplots of the same color); therefore, it produces higher parent-offspring similarity when selecting only for performance (Fig. 5, see red and dark blue curves). (2) However, parent-offspring similarity is reduced by novelty (Fig. 5, see orange and light blue curves). (3) This reduction results in decreased locomotion performance of the offspring in the Lamarckian system immediately after controller inheritance, and this effect is more pronounced when parent-offspring similarity is higher (Fig. 6, see comparisons among red and orange box-plots).

We examine the data supporting each of these assumptions in detail and find that all are upheld. First, we inspect the morphological similarity between parents and children and their evolution over time. Figure 5 shows the mean morphological similarity between parents and children for all conditions. In all experiments, similarity increases sharply during the initial generations. When novelty is incorporated, this initial increase is more pronounced, indicating stronger divergence of offspring from their parents; after this early phase, offspring become progressively more dissimilar. In the locomotion-only condition, however, offspring continue to become increasingly similar over time, but the Darwinian and Lamarckian systems differ in how convergence unfolds: the Lamarckian system produces greater parent-offspring similarity (assumption 1). This distinction disappears when novelty is introduced: similarity is reduced under both Darwinian and Lamarckian evolution (assumption 2), substantially narrowing the gap between the two approaches.

Subsequently, we assess whether parent-offspring similarity modulates performance transference. Figure 6 shows the initial locomotion performance of the offspring immediately after transference – this quantifies the benefit of the parent’s learned controller for the child’s initial performance. Under Darwinian




Fig. 6: Initial locomotion performance within subgroups of parent-offspring morphological similarity. The two groups, more similar and less similar, are split using the median parent-offspring similarity of all individuals. **Under Lamarckian inheritance, greater parent-offspring similarity translates to higher initial locomotion performance.**




Fig. 7: Comparison of initial locomotion performance. In this analysis, any individuals with performance above the maximum performance achieved in the combined experiments were removed from the sample; this aimed at preventing a bias that could be introduced by the high performances values of the pure experiments. **Lamarckian evolution exhibits greater performance degradation when novelty is combined with a locomotion task, and this is due to decreased parent-offspring similarity.**

evolution, the initial performance is largely insensitive to morphological similarity, while Lamarckian evolution shows a strong dependence. In the latter, offspring with greater parent-offspring similarity achieve substantially higher initial performance than less similar ones – this reinforces the evidence that the Lamarckian system works better under greater parent-offspring similarity (assumption 1). Additionally, the initial performance is significantly decreased in the novelty experiments, and this decrease is only significant for the subgroup with high parent-offspring similarity (assumption 3).

Since the pure locomotion fitness and the locomotion-novelty fitness yield distinct locomotion-performance dis-

tributions as a result of their respective selection pressures, one might question whether the performance differences between the novelty and pure settings are attributable solely to these distributions. If this were the case, the same type of difference would be observed consistently in both systems; however, it happens only in the Lamarckian system. To strengthen this evidence, we include an additional comparison in support of assumption 3. Since the pure setups achieve higher fitness than the combined setups, an effect that is more pronounced in the Lamarckian case, we cap individual fitness values in the pure experiments at the maximum fitness achieved in the corresponding combined experiments. Specifically, individuals in the Darwinian pure condition are retained only if their fitness does not exceed the maximum fitness observed in the Darwinian combined condition, and the same procedure is applied to the Lamarckian setups. This enables a direct comparison of initial fitness while avoiding bias introduced by differing performance ceilings. Figure 7 shows that in the Darwinian system, there is no observable difference, whereas in the Lamarckian system, a significant difference in initial fitness emerges.

In summary, these results show that the advantages of the Lamarckian system were hindered under the pressure towards morphological novelty, which in the present study was implemented as an explicit component of the fitness function. In non-open-ended evolutionary systems, where fitness is a user-defined function utilized for selection a priori rather than measured a posteriori as an outcome of natural selection, the fitness function serves as an abstraction of environmental pressures. Morphological novelty here, can be viewed as a controlled proxy for environmental forces that introduce transgenerational variability. In natural or dynamic environments, such forces may arise through frequency-dependent selection, niche differentiation, or environmental heterogeneity. What ultimately matters for controller inheritance is the effect on parent-offspring morphological continuity, regardless of whether variability originates from external environmental change or from diversity-promoting selection, as both disrupt the similarity assumption underlying Lamarckian inheritance.

VI. CONCLUSION

This paper investigated the benefits and limitations of Lamarckian evolution in comparison with Darwinian evolution in morphologically evolving robots that learn during their lifetime. We showed that while Lamarckian inheritance consistently improves performance when evolution favors morphological similarity, its advantage is reduced when novelty is explicitly rewarded. Darwinian evolution is largely unaffected. The key challenge arises when the pressure for morphological novelty reduces the similarity between parents and offspring, thereby affecting the conditions under which Lamarckian inheritance is effective. These findings reveal an important limitation of Lamarckian evolution by exposing a fundamental trade-off between inheritance-based exploitation and diversity-driven exploration. Although Lamarckian evolution remains a powerful approach, its suitability depends on the conditions

under which it is applied, and this work helps to delineate those conditions.

These observations open several avenues for future research aimed at identifying the regimes in which Lamarckian evolution is most effective. For example, it remains to be determined at what rate and magnitude of environmental change or variability introduction Lamarckian strategies cease to be advantageous. Future work could also explore adaptive inheritance mechanisms, such as similarity-aware transfer, partial controller reuse, or hybrid evolutionary schemes, to better harvest the benefits of Lamarckian learning with the demands of morphological diversity.

VII. ACKNOWLEDGMENTS

The authors would like to acknowledge that this work builds upon the ARIEL simulator, which is currently in the early stages of development. We sincerely thank the developers Jacopo Di Matteo, Áron Richárd Ferencz, Giannis Grigoriadis, and Lilly Schwarzenbach for their contributions to the development of this tool.

The research reported in this paper was supported by the European Commission Horizon project SPEAR (EU grant number 101119774, <https://www.spear-robotics.com/>).

REFERENCES

- [1] A. E. Eiben, S. Kernbach, and E. Haasdijk, “Grand challenges for evolutionary robotics,” *Frontiers in Robotics and AI*, vol. 2, p. 4, 2015.
- [2] M. Jelisavcic, K. Glette, E. Haasdijk, and A. E. Eiben, “Lamarckian evolution of simulated modular robots,” *Frontiers in Robotics and AI*, vol. 6, p. 9, 2019. DOI: 10.3389/frobt.2019.00009.
- [3] M. Jelisavcic, R. Kiesel, K. Glette, E. Haasdijk, and A. E. Eiben, “Analysis of lamarckian evolution in morphologically evolving robots,” in *Artificial Life Conference Proceedings*, MIT Press, 2017, pp. 214–221.
- [4] M. Jelisavcic, K. Miras, and A. E. Eiben, “Morphological attractors in darwinian and lamarckian evolutionary robot systems,” in *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, IEEE, 2018, pp. 859–866.
- [5] M. Jelisavcic, R. Kiesel, K. Glette, E. Haasdijk, and A. E. Eiben, “Benefits of lamarckian evolution for morphologically evolving robots,” in *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, ACM, 2017, pp. 65–66.
- [6] J. Luo, K. Miras, J. M. Tomczak, and A. E. Eiben, “Enhancing robot evolution through lamarckian principles,” *Scientific Reports*, vol. 13, no. 1, p. 21 109, 2023.
- [7] J. Luo, C. Longhi, and A. E. Eiben, “A comparative study of brain reproduction methods for morphologically evolving robots,” in *Artificial Life Conference Proceedings*, MIT Press, vol. 2023, 2023, p. 44.

- [8] J. Luo, J. M. Tomczak, K. Miras, and A. E. Eiben, "A comparison of controller architectures and learning mechanisms for arbitrary robot morphologies," in *2023 IEEE Symposium Series on Computational Intelligence (SSCI)*, IEEE, 2023, pp. 1518–1525.
- [9] J. Luo, K. Miras, C. Longhi, O. Weiszl, and A. E. Eiben, "Lamarckian inheritance improves robot evolution in dynamic environments," *IEEE Transactions on Evolutionary Computation*, 2025.
- [10] J. M. Mingo, R. Aler, D. Maravall, and J. de Lope, "Investigations into lamarckism, baldwinism and local search in grammatical evolution guided by reinforcement," *Computing and Informatics*, vol. 32, no. 3, pp. 595–627, 2013.
- [11] T. Sasaki and M. Tokoro, "Comparison between lamarckian and darwinian evolution on a model using neural networks and genetic algorithms," *Knowledge and Information Systems*, vol. 2, no. 2, pp. 201–222, 2000.
- [12] Z. Sharifi, K. Soltanian, and A. Amiri, "Developing convolutional neural networks using a novel lamarckian co-evolutionary algorithm," in *13th International Conference on Computer and Knowledge Engineering (ICCKE)*, Ferdowsi University of Mashhad, Mashhad, Iran, Nov. 2023.
- [13] N. A. M. Tack, Y. Yamazaki, and F. Iida, "Lamarckian co-design of soft robots via transfer learning," in *Proceedings of the Genetic and Evolutionary Computation Conference*, ACM, 2024, pp. 129–137.
- [14] 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.
- [15] J. K. Pugh, L. B. Soros, and K. O. Stanley, "Quality diversity: A new frontier for evolutionary computation," *Frontiers in Robotics and AI*, vol. 3, p. 40, 2016.
- [16] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, 2015.
- [17] J. E. Auerbach and J. C. Bongard, "Environmental influence on the evolution of morphological complexity in machines," *PLOS Computational Biology*, vol. 10, no. 1, e1003399, 2014.
- [18] N. Cheney, R. MacCurdy, J. Clune, and H. Lipson, "Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding," *ACM SIGEVOlution*, vol. 7, no. 1, pp. 11–23, 2013.
- [19] K. Miras, E. Haasdijk, K. Glette, and A. E. Eiben, "Search space analysis of evolvable robot morphologies," in *Applications of Evolutionary Computation: 21st International Conference, EvoApplications 2018, Parma, Italy, April 4-6, 2018, Proceedings*, K. Sim and P. Kaufmann, Eds., ser. Lecture Notes in Computer Science, vol. 10784, Cham: Springer, 2018, pp. 703–718. DOI: 10.1007/978-3-319-77538-8_47.