Morphological Attractors in Darwinian and Lamarckian Evolutionary Robot Systems

Milan Jelisavcic, Karine Miras, and A.E. Eiben
Computational Intelligence Group
Department of Computer Science
Vrije Universiteit Amsterdam, Amsterdam, The Netherlands
Email: m.j.jelisavcic@vu.nl

***Abstract*—Morphological evolution in a robotic system produces novel robot bodies after each reproduction event. This implies the necessity for lifetime learning so that newborn robots can acquire a controller that fits their body. Thus, we obtain a system where evolution and learning are combined. This combination can be Darwinian or Lamarckian and in this paper, we compare the two. In particular, we investigate the evolved morphologies under these regimes for modular robots evolved for good locomotion. Using eight quantifiable morphological descriptors to characterize the physical properties of robots we compare the regions of attraction in the resulting 8-dimensional space. The results show prominent differences in symmetry, size, proportion, and coverage.**

***Index Terms*—Lamarckian evolution, Modular robots, Online learning, Embodied evolution, Artificial life, Evolutionary robotics.**

I. Introduction

Evolutionary Robotics is the field of science that applies evolutionary algorithms to design and optimize the morphologies and/or controllers of simulated or real robots [26]. The approach is a good way to design better robots as well as to test evolutionary hypotheses about biological systems. Through the process of the production of a new robot, potentially novel body designs emerge. The new body calls for a well-adapted controller in order to exploit its full potential. Recent studies show that the choice for the development of controllers has a strong influence on the development of morphologies [8].

It has been shown that making learned knowledge inheritable (i.e. Lamarckian regime) can provide a benefit to a newly-born robot [15]. The same set-up has been tested in another investigation that shows the greater influence of the body structure against the brain throughout the robot's lifetime [16]. In this paper, we investigate how does the inherited knowledge influences the evolutionary development over several generations.

One of the interesting questions that occur is how does it influence the evolution of the morphologies. Namely, we want to answer the following research questions:

- Q1: Could we distinguish regions of attraction in the morphological space after a number of generations?
- Q2: Are these regions of attraction different under Darwinian and Lamarckian regimes?

II. Related Work

The relationship between body, brain, and the environment defines the potential for intelligent behaviour [6]. A robot's behaviour is the result of the interaction between its morphology, controller, and environment [23]. The intricate relation between environment and evolution of morphologies, investigated in 2014, shows that environment greatly determines the complexity of morphologies [4]. As noticed, "for many animals, natural selection may tend to favour structures and patterns of movement that increase maximum speed", and, "in almost every case, legged animals can move faster over land than animals of similar size that lack legs" [2].

The underlying system architecture that fully explores interactions between bodies, brains and environments is called the Triangle of Life and has been put forward in 2013 [11]. This system

978-1-5386-9276-9/18/$31.00

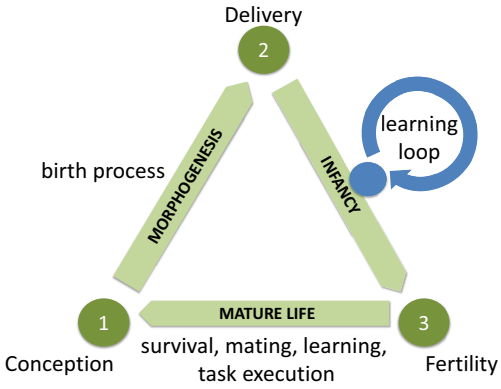

Figure 1. The Triangle of Life. The pivotal moments that span the triangle and separate the three stages are 1) Conception: A new genome is activated, construction of a new robot starts. 2)Delivery: Construction of the new robot is completed. 3) Fertility: The robot becomes an adult, ready to conceive offspring.

captures the pivotal life cycle of an ecosystem of self-reproducing robots as illustrated in Figure 1. This life-cycle does not run from birth to death, but from conception (being conceived) to conception (conceiving one or more children) and it is repeated over and over again, thus creating consecutive generations of robot children. The result is a population of robotic organisms that evolve and thus adapts to the given environment. The Triangle of Life consists of 3 stages, Morphogenesis, Infancy, and Mature Life. The first real-world implementation of the system is presented in 2017 [14].

There are two principal options for evolution to exploit lifetime learning: Darwinian and Lamarckian evolution [27]. Lamarckian evolution, in contrast to Darwinian, does explicitly store the locally learned improvements in the individual genomes, so that lifetime learning can directly accelerate the evolutionary process and vice versa [1]. Up until now, the Lamarckian approach to evolution has seen an initial investigation [9]. While this mechanism has largely not been seen as a correct description of biological evolution, some recent research has reported a Lamarckian type of evolution in nature [10]. The recent researches showed that the implementation of Lamarckian evolution provides benefits at least at the start of a robots life-cycle [15].

Another prominent effort has been put into co-evolving morphologies and environments [7], [18]–[20]. The investigations showed that the robots with more plasticity adapt better to different environments [12], [17]. Recent research explores the influence of fitness functions on the outcome of evolution [21]. The basis for this research is the morphological descriptors defined in [22] along with [16].

III. Robot Description

For the experiments in this paper, the robots are simulated using Revolve [1], a custom simulator based on Gazebo [2]. The robots and their genetic representation are based on RoboGen design [3]. The robot design consists of two parts: a) the *body design* (i.e., morphology), and the *brain design* (i.e., controller)

a) Body Design.: Each robot's genotype describes its layout and consists of a tree structure with the root node representing a core component from which further components branch out. In this study, we use a subset of 3D-printable components: *fixed bricks*, *core component*, and *active hinges* (Figure 2). Components are designated by their type. Each of these components is defined by the two-part model: a detailed mesh suitable for visualisation and 3D-printing and a set of geometric primitives that define the components' mass distribution and a contact surface. Each component also defines the number and placement of possible attachment slots, as well as outputs (motors) contained within 'active hinge' component.

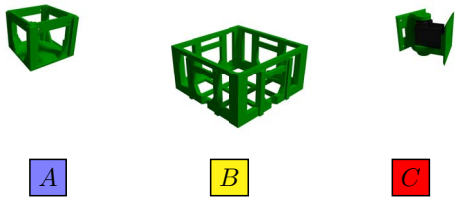

Figure 2. The 3D-printable robot components., **(A)** Fixed brick, **(B)** Core component, and **(C)** Active hinge. These models are used in the simulation, but also could be used for 3D printing and construction of real robots. The blue-, yellow-, and red-coloured blocks bellow components illustrate a 2D representation of robots in Figures 5 and 6.

Robots are genetically encoded by a tree-based representation where each node represents one building block of the robot and edges between nodes represent physical connections between pieces. Every node contains information

[1] https://github.com/ci-group/revolve/
[2] http://gazebosim.org/

about the *type of the component* it represents, its *name*, *orientation*, possible *parametric values*, and its *colour*. Each edge also defines which of the available parent's node *attachment slot* the child will attach to. Construction of a robot from this representation begins with the root node, defined to always represent the requisite core component. The robot body is then constructed by traversing the tree edges and attaching the components represented by child nodes to the current component at the specified slot positions and orientations.

b) Brain Design.: The brain design for the robot locomotion consists of two main components – a *CPG controller structure* derived from a robot's body structure and *weights of CPG connections* derived as outputs of a CPPN network. Figure 3 depicts the resulting architecture. The CPG is strongly grounded in the morphology of a given robot (explained below). The part that can be transferred between different robots is the CPPN. This is very important as it enables us to transfer controllers between different bodies.

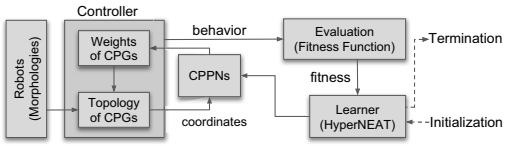

Figure 3. The overall architecture of the learning system. The learning method is implemented by an evolutionary algorithm (HyperNEAT). It evolves the CPPN that defines the connection weights of the CPG-based controller whose topology is based on the morphology of the given robot.

The main components of the CPG controllers are differential oscillators. One oscillator is defined for each active hinge. The consequence of assigning the nodes in a differential CPG structure a location in an n-dimensional hypercube is the inclusion of HyperNEAT as a learning mechanism. The assigned relative positions should in some way reflect a relationship between the nodes allowing the algorithm to exploit the geometry of the problem. The CPPN evolves using the HyperNEAT learning method [25] so that the CPG structure's performance is optimised.

The oscillators of neighbouring hinges (i.e., hinges separated by a single component) are interconnected by means of weighted connections between their x neurons. This results in a chain-like a neural network of differential oscillators that extends across the robot body, as illustrated on the left side of Figure 4.

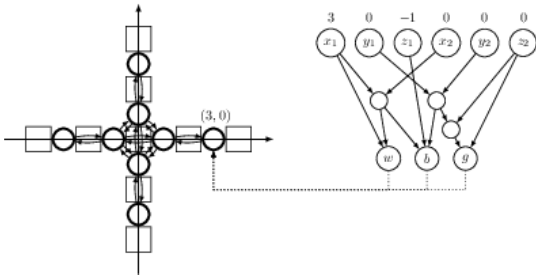

Figure 4. Example of the lifetime process of applying the proper weights from a CPPN network onto CPG connections. The arrow is pointing to a neuron within a differential oscillator with coordinates $(x, y) = (3, 0)$ and $z = -1$ for one of the nodes within the oscillator (the other one is designated with $z = 1$).

Like a neural network, a CPPN is a network of mathematical functions with weighted connections. The CPPNs have six inputs denoting the coordinates of a source and a target node when querying connection weights or just the position of one node when obtaining node parameters with the other three inputs being initialised as zero. The CPPNs have three outputs: the weight of the connection from source to target as well as the bias and gain values when calculating parameters for a node. To determine the weight of a connection in the CPG network that controls the robot (the substrate), the coordinates of the two substrate nodes are fed into the CPPN which then returns the connection weight [24]. In order to obtain the parameters of a node, the coordinates of that node and the all-zero vector (instead of a coordinate of the other node) are used as inputs. This way enables us to select either a connection between two nodes, or a specific node itself.

Example for the process of applying parameters to a specific neuron in a CPG network is illustrated in Figure 4. On the CPG structure, the coordinates of each active hinge are illustrated. In order to define the values for the y node on the coordinate $(3, 0)$, we designate $z_1 = -1$ and $(x_2, y_2, z_2) = (0, 0, 0)$. Based on values feed into a CPG network, the different output pattern is produced for every actuator resulting in a different locomotion behaviour. The HyperNEAT then evolves the CPPN in order to optimise the connection weights, the node biases, and the gain levels of the output nodes

produced by it.

IV. Experimental Set-up

The main logic behind the experiments is to separately run and compare two systems:

- Darwinian evolution of morphologies and Darwinian evolution of controllers,
- Darwinian evolution of morphologies and Lamarckian evolution of controllers.

With lifetime learning by means of an on-line evolutionary algorithm as in this research, each robot carries an internal population of controllers that evolve during the robot's lifetime. It is important to note once more that, in this experimental set-up, two evolutionary processes are ongoing: **(1)** evolution of *morphologies*, and **(2)** evolution of *controllers* in an individual robot. The process of lifetime-learning of gait controllers does not necessarily have to include an evolutionary algorithm, but since we are using HyperNEAT-CPPN pair to develop controllers, it can be viewed as evolutionary.

We have been using versatile robot morphologies and a unique controller architecture in combination with HyperNEAT learning algorithm. We should emphasize that these versatile morphologies are a product of a nature of evolutionary systems that we must count on. In such a system, a simple but effective implementation of Lamarckian evolution is to seed an individual's population from that of its parents.

The process of adapting CPPNs through the recombination and mutation and further on applying and testing with them a robot's locomotory performance represents the learning process in our system (Figure 3). Recombination and mutation of genomes are implemented through the standard operators defined in RoboGen. As illustrated in Section III, the morphologies of the robots can be represented as tree structures where every node represents one component. Therefore, conveniently, we can use the recombination and mutation operators that are well-established in genetic programming practice [5].

In both tested system, the evolution of morphologies goes through the same conditions, meaning that we apply recombination and mutation on directly-encoded body genome. The main difference is contained within the within the evolutionary process of controllers. When considering the Darwinian evolution of controllers, the lifetime-learning process does not have an influence on the evolution of controllers – the controllers that robot inherited at his birth will be used in the recombination and mutation for its offspring. Quite the opposite, in the Lamarckian evolution of controllers, instead of the initial controllers the system will use, for the production of offspring, the best controllers developed throughout the robot's lifetime.

For the comparison of the two systems, we randomly generated five populations of robots, each containing 20 individuals. For the reasons of computational costs in the first system with the Lamarckian evolution of controllers, the number of generations timespan is limited to 10 generations. We test the variant of seeding an offspring's population that initialises the HyperNEAT population with the best five CPPNs from each parent. The first generation of robots does not have a parental seed to start from, so their initial HyperNEAT population consists of randomly initialised networks only containing the input and output neurons and connections from every input to every output neuron with randomly initialised weights and neuron parameters. In the second system, with the Darwinian evolution of controllers, the population of 10 CPPNs is generated by recombining the best performing CPPN from the first parent to the best five CPPNs from the second parent and vice versa, producing in total nine CPPNs. The additional 10th CPPN is provided by random selection from the best five of parental CPPN pool.

As the system of choice, Revolve [13] simulator was used, which is specifically designed for managing Triangle of Life-based experiments.

V. Analysis

For the analysis of the morphological properties of the evolved populations we measure and compare a set of morphological descriptors [22]. The morphological descriptors are a tool for quantifying the properties of each robot's morphology. In short, there are eight defined descriptors:

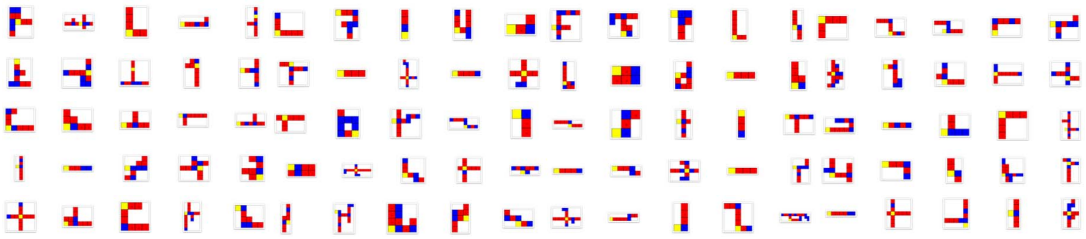

Figure 5. The first generation in five lineages used in both scenarios. Each row represents one lineage with 20 robots.

- **Branching** quantifies how the attachments of the components are grouped together in a body;
- **Number of Limbs** quantifies the number of extremities of a body;
- **Length of Limbs** quantifies the extensiveness of extremities in a body;
- **Coverage** quantifies the fulfilment of the rectangular space created by a body;
- **Joints** quantify degrees of freedom of a body;
- **Proportion** quantifies the two-directional proportion of a body;
- **Symmetry** quantifies two-directional reflexive symmetry of a body;
- **Size** quantifies the extent of a body in terms of the number of components;

Figure 6 presents the morphologies evolved by Lamarckian and Darwinian scenarios. Both scenarios had the same initial generations, as shown in Figure 5, but the final populations present distinct predominant morphological properties. For the Lamarckian scenario, the body outlines are predominantly X- and T-shaped (multiple-limbs robots). While for the Darwinian scenario, the body outlines are predominantly I- and L-shaped (snake-like robots). This is understandable, once by having the chance to learn coordination, robots could benefit from having multiple limbs through maintaining a constant speed, and thus producing faster and more stable locomotive patterns.

Figure 7 presents confidence intervals for all of the morphological descriptors in the final populations We can notice the clear differences in the confidence intervals for the 'proportion', 'coverage', and 'size'. The most interesting is the symmetry descriptor. What attracts the most of the attention, over the course of evolution, robots tend to evolve more symmetrical in the Lamarckian regime. Apart from the robot's symmetry, the size and proportion also tend to increase in the Lamarckian setup, whilst the coverage decrease compared to the Darwinian setup. This is very important considering that none of the described aspects is implemented in the system as a requirement.

Figure 8 illustrates emerged body features using evolutionary learning for bodies and lifetime learning for minds. We can conclude that the different morphological niches are covered in two different scenarios. While under the Darwinian regime the descriptor space tends to be more covered, under the Lamarckian 'coverage' and 'proportion' tend to cluster.

To verify these tendencies for the descriptors over generations, in the two tested scenarios, we applied a Mann-Kendall trend test. Results of the test are presented in Table I. For the Lamarckian scenario the results are statistically significant for all eight descriptors. Thus, there is a trend to growth or decay for all morphological properties. The most significant positive trend is noted for 'size' and 'length of limbs', but tendency also exists for 'symmetry', 'proportion', and 'branching'. Table II shows results that corroborate with this, by comparing the differences in an average of the descriptors from the initial to the final population. In almost all cases, except for 'joints', the same descriptors that present trend, present also an average in the final population that is different from the initial one.

In the Darwinian scenario, the significant trends were only 'number of limbs', 'length of limbs', and 'joints'. However, the p values of 'proportion' and 'symmetry' are not so high. As the number of generations is reasonably low, perhaps there was not time enough to see a clearly significant trend in the second scenario. Thus, we assessed the differences for the descriptors when comparing

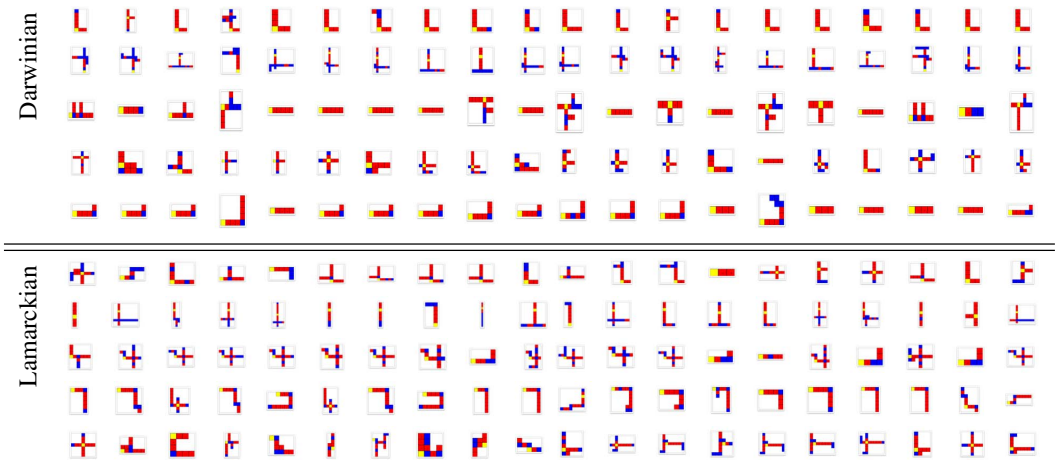

Figure 6. Morphologies of final (10th) generations in both scenarios. Darwinian setup (upper group) develops predominantly I- and L-shaped robots. Lamarckian setup (lower group) develops predominantly X- and T-shaped robots.

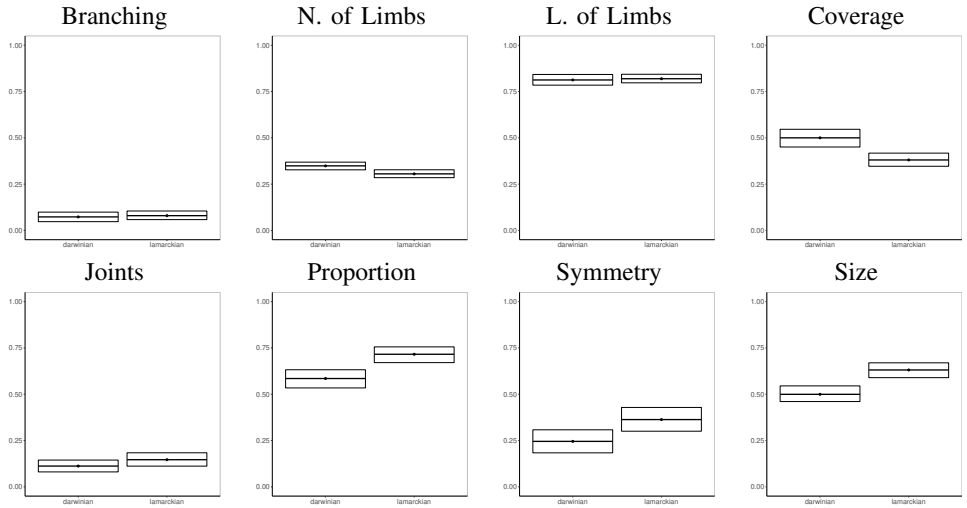

Figure 7. Confidence intervals for the morphological descriptors in the 10th generation Confidence intervals in morphological descriptors over 10 generations in Lamarckian and Darwinian evolution. The x-axis represents two investigated scenarios, Darwinian and Lamarckian, and the y-axis is a coefficient for every measure. Note the positive trend for 'symmetry', 'size', coverage', and 'proportion' for the Lamarckian setup.

only the last generation of each scenario. Figure III shows the significances for differences in an average of all morphological descriptors between the two scenarios. It is interesting to see that from this perspective, 'proportion' and 'symmetry' are significantly different. Also, the 'coverage' is higher for the non-learning scenario, which makes sense, once an I-shape covers the whole body area, while an X-shape creates sparseness among the body parts. In summary, in the Lamarckian scenario (learning) the population is predominantly symmetrical, proportional and with [multiple limbs], while in the purely Darwinian scenario it is the opposite, with disproportional, asymmetrical robots.

VI. Discussion

In this study, we examined the regions of attraction in morphological space with the Darwinian and Lamarckian evolution of controllers. Our results showed differences in evolved morphologies. In the Darwinian system, robots tend to develop to more simple forms. In the Lamarckian system, robots evolve more symmetrical and larger structures.

In the Darwinian regime, we observed that robots tend to develop snake-like shapes after a number of generations. The evolution of morphologies with the Lamarckian evolution of controllers tend to produce more complex body structures with three and four limbs. A plausible explanation of this

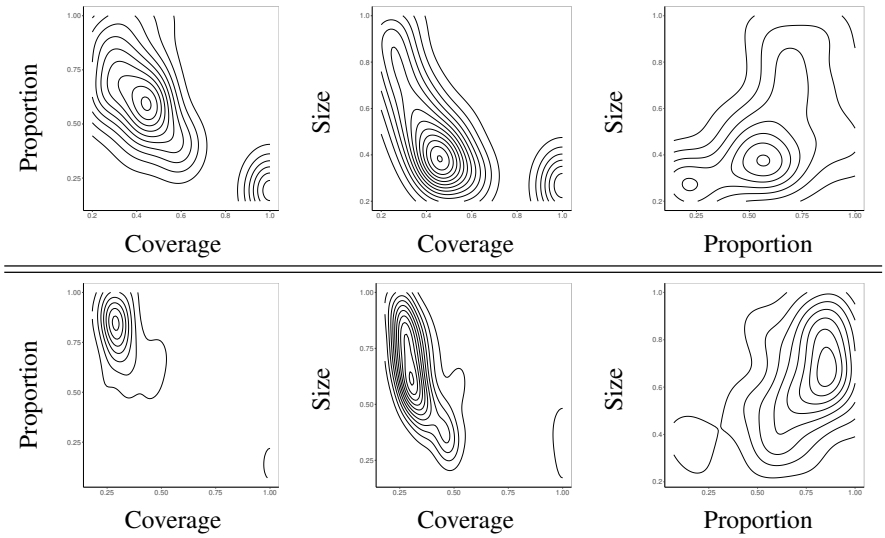

Figure 8. Density areas for the three most prominent morphological descriptors in the 10th generation of both Darwinian (top row) and Lamarckian (bottom row) regimes. Every plot represents the density correlation between two descriptors.

	Lamarckian		Darwinian	
	coef.	p-value	coef.	p-value
Branching	0.07	8×10^{-3}	0.01	0.72
N. of Limbs	–0.12	2.6×10^{-7}	–0.1	4.6×10^{-4}
L. of Limbs	0.07	8.7×10^{-3}	0.1	3.2×10^{-5}
Coverage	–0.17	2.5×10^{-14}	–0.03	0.17
Joints	–0.05	2.8×10^{-2}	–0.06	0.02
Proportion	0.05	1.3×10^{-2}	–0.03	0.09
Symmetry	0.05	2.4×10^{-2}	–0.04	0.07
Size	0.16	2.2×10^{-16}	0.006	0.76

Table I
COEFFICIENT VALUES OF MANN-KENDALL STATISTICAL TREND TEST FOR MORPHOLOGICAL DESCRIPTORS OVER GENERATIONS. THE GREY-COLOURED CELLS HIGHLIGHT THE NOT ENOUGH SIGNIFICANT CORRELATION.

Descriptor	L. 1-10	D. 1-10
Branching	6×10^{-3}	0.67
N. of Limbs	8×10^{-6}	1×10^{-3}
L. of Limbs	2×10^{-3}	2×10^{-3}
Coverage	9×10^{-8}	0.1
Joints	0.17	4×10^{-3}
Proportion	5×10^{-3}	0.4
Symmetry	6×10^{-3}	0.67
Size	2×10^{-8}	0.2

Table II
SIGNIFICANCES FOR DIFFERENCES IN THE AVERAGES OF THE DESCRIPTORS BETWEEN THE INITIAL AND FINAL GENERATION OF EACH SYSTEM. THE TEST USED WAS WILCOXON. AVERAGES ARE THE MEAN OF ALL LINEAGES.

Descriptor	Lamarckian vs. Darwinian
Branching	0.43
N. of Limbs	0.02
L. of Limbs	0.96
Coverage	1×10^{-5}
Joints	0.11
Proportion	4×10^{-5}
Symmetry	7×10^{-3}
Size	5×10^{-6}

Table III
SIGNIFICANCES FOR DIFFERENCES IN THE AVERAGES OF THE DESCRIPTORS BETWEEN THE TWO SYSTEMS IN THE FINAL GENERATION. THE TEST USED WAS WILCOXON. AVERAGES ARE THE MEAN OF ALL LINEAGES.

effect is that the larger number of actuators in larger bodies increases the amount of interference between moving limbs. The body complexity makes the learning task more difficult and starting the learning process from the adapted controllers of the parents (Lamarckian) obviously has an advantage over starting the learning process from the random situation (Darwinian). In case of a 'weak'

learning method (random start) the more complex morphologies cannot acquire suitable controllers, hence the simple shapes become dominant.

Future work will be devoted to research the scope of this effect and investigate how it depends on the environment.

References

[1] Ackley, D., Littman, M.: A case for distributed lamarckian evolution. Artificial Life III (1994)

[2] Alexander, R.M.: Principles of animal locomotion. Princeton University Press (2003)

[3] Auerbach, J., Aydin, D., Maesani, A., Kornatowski, P., Cieslewski, T., Heitz, G., Fernando, P., Loshchilov, I., Daler, L., Floreano, D.: RoboGen: Robot Generation through Artificial Evolution. In: Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems. pp. 136–137. The MIT Press, New York, New York, USA (July 2014)

[4] Auerbach, J., Bongard, J.: Environmental Influence on the Evolution of Morphological Complexity in Machines. PLoS Computational Biology **10**(1), e1003399 (jan 2014). https://doi.org/10.1371/journal.pcbi.1003399, http://dx.plos.org/10.1371/journal.pcbi.1003399

[5] Banzhaf, W., Nordin, P., Keller, R.E., Francone, F.D.: Genetic programming: an introduction, vol. 1. Morgan Kaufmann San Francisco (1998)

[6] Beer, R.D.: The Dynamics of Brain—Body–Environment Systems. In: Handbook of Cognitive Science, pp. 99–120. Elsevier, Amsterdam (2008). https://doi.org/10.1016/B978-0-08-046616-3.00006-2, http://linkinghub.elsevier.com/retrieve/pii/B9780080466163000062

[7] Buason, G., Bergfeldt, N., Ziemke, T.: Brains, bodies, and beyond: Competitive co-evolution of robot controllers, morphologies and environments. Genetic Programming and Evolvable Machines **6**(1), 25–51 (Mar 2005). https://doi.org/10.1007/s10710-005-7618-x, https://doi.org/10.1007/s10710-005-7618-x

[8] Buresch, T., Eiben, A.E., Nitschke, G., Schut, M.: Effects of Evolutionary and Lifetime Learning on Minds and Bodies in an Artifical Society. In: 2005 IEEE Congress on Evolutionary Computation. vol. 2, pp. 1448–1454. IEEE (2005). https://doi.org/10.1109/CEC.2005.1554860, http://ieeexplore.ieee.org/document/1554860/

[9] Cortez, P., Rocha, M., Neves, J.: A lamarckian approach for neural network training. Neural Processing Letters **15**(2), 105–116 (Apr 2002). https://doi.org/10.1023/A:1015259001150, https://doi.org/10.1023/A:1015259001150

[10] Dias, B.G., Ressler, K.J.: Parental olfactory experience influences behavior and neural structure in subsequent generations. Nature neuroscience **17**(1), 89 (2014)

[11] Eiben, A., Bredeche, N., Hoogendoorn, M., Stradner, J., Timmis, J., Tyrrell, A., Winfield, A.: The Triangle of Life: Evolving Robots in Real-Time and Real-Space. In: Liò, P., Miglino, O., Nicosia, G., Nolfi, S., Pavone, M. (eds.) Advances In Artificial Life, ECAL 2013. pp. 1056–1063. MIT Press (2013)

[12] Hornby, G.S., Pollack, J.B.: Creating high-level components with a generative representation for body-brain evolution. Artificial Life **8**(3), 223–246 (2002). https://doi.org/10.1162/106454602320991837

[13] Hupkes, E., Jelisavcic, M., Eiben, A.E.: Revolve: A versatile simulator for online robot evolution. In: Sim, K., Kaufmann, P. (eds.) Applications of Evolutionary Computation. pp. 687–702. Springer International Publishing, Cham (2018)

[14] Jelisavcic, M., De Carlo, M., Hupkes, E., Eustratiadis, P., Orlowski, J., Haasdijk, E., E. Auerbach, J., Eiben, A.: Real-World Evolution of Robot Morphologies: A Proof of Concept. Artificial life **23**(2), 206–235 (June 2017), pMID: 28513201

[15] Jelisavcic, M., Kiesel, R., Glette, K., Haasdijk, E., Eiben, A.: Analysis of Lamarckian Evolution in Morphologically Evolving Robots. In: Proceedings of the European Conference on Artificial Life 2017, ECAL 2017. pp. 214–221. MIT Press (September 2017)

[16] Jelisavcic, M., Roijers, D.M., Eiben, A.E.: Analysing the Relative Importance of Robot Brains and Bodies. In: Ikegami, T., Virgo, N., Witkowski, O., Suzuki, R., Oka, M., Iizuka, H. (eds.) Proceedings of the Artificial Life Conference 2018 (ALIFE 2018). MIT Press, Tokyo (2018)

[17] Krcah, P.: Effects of morphological plasticity on evolution of virtual robots. Adaptive Behavior **25**(2), 44–59 (2017)

[18] Larpin, K.: Co-evolution of morphology, control and behavior (2011)

[19] Linder, C.R.: Embodiment in two dimensions. In: Climbing and Walking Robots, pp. 313–320. Springer (2005)

[20] Lund, H.H.: Co-Evolving Control and Morphology with Lego Robots. In: Morpho-functional machines: the new species, pp. 59–79. Springer (2003)

[21] Miras, K., Haasdijk, E., Glette, K., Eiben, A.E.: Effects of Selection Preferences on Evolved Robot Morphologies and Behaviors. In: Ikegami, T., Virgo, N., Witkowski, O., Suzuki, R., Oka, M., Iizuka, H. (eds.) Proceedings of the Artificial Life Conference 2018 (ALIFE 2018). MIT Press, Tokyo (2018)

[22] Miras, K., Haasdijk, E., Glette, K., Eiben, A.E.: Search space analysis of evolvable robot morphologies. In: Sim, K., Kaufmann, P. (eds.) Applications of Evolutionary Computation. pp. 703–718. Springer International Publishing, Cham (2018)

[23] Pfeifer, R., Iida, F.: Embodied artificial intelligence: Trends and challenges. In: Embodied artificial intelligence, pp. 1–26. Springer (2004)

[24] Stanley, K.O.: Compositional pattern producing networks: A novel abstraction of development. Genetic programming and evolvable machines **8**(2), 131–162 (2007)

[25] Stanley, K.O., D'Ambrosio, D.B., Gauci, J.: A Hypercube-Based Encoding for Evolving Large-Scale Neural Networks. Artificial Life **15**(2), 185–212 (apr 2009). https://doi.org/10.1162/artl.2009.15.2.15202, http://www.mitpressjournals.org/doi/abs/10.1162/artl.2009.15.2.15202

[26] Vargas, P.A., Di Paolo, E.A., Harvey, I., Husbands, P.: The horizons of evolutionary robotics. MIT Press (2014)

[27] Whitley, D., Gordon, V.S., Mathias, K.: Lamarckian evolution, the baldwin effect and function optimization. In: International Conference on Parallel Problem Solving from Nature. pp. 5–15. Springer (1994)