Effects of Selection Preferences on Evolved Robot Morphologies and Behaviors

Karine Miras¹, Evert Haasdijk¹, Kyrre Glette², A.E. Eiben¹

¹Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

²University of Oslo, Oslo, Norway
k.s.m.a.dasilvamirasdearaujo@vu.nl

Abstract

This paper investigates the evolution of modular robots using different selection preferences (i.e., fitness functions), aiming at novelty, speed of locomotion, number of limbs, and combinations of these. The outcomes are analyzed from different perspectives: sampling of the search space, evolved morphologies, and evolved behaviors. This results in a wealth of findings, including a surprise about the number of sampled regions of the search space and the effect of different fitness functions on the evolved morphologies.

Keywords: evolutionary robotics, morphology evolution, generative encoding, novelty search, search space

Introduction

In this paper we consider evolutionary robot systems where both the “bodies” and the “brains”, i.e., the morphologies and the controllers, of the robots are evolvable. As opposed to the majority of evolutionary robotics studies, we focus on the evolution of morphologies. In an earlier paper we investigated the evolution of several morphological features, e.g., Size, Symmetry, Number of Limbs, Proportion, in a system where selection ignored the behavior, and the fitness was based on the novelty of the newly generated robot morphology, cf. (Miras et al., 2018). The main goal of this paper is to investigate how the evolved morphological features change when using different criteria for selection.

To this end, we define a task (locomotion) that induces a behavior-based fitness measure (speed) and evolve robots under various fitness functions. In particular, we compare four options, 1) based purely on a *relative morphological* measure (novelty) as in our previous paper, 2) based purely on a *behavioral* measure (speed), 3) based on the combination of novelty and speed, and 4) based on novelty, speed, and an *absolute morphological* measure (length of limbs). Our specific research questions can be grouped around three subjects:

- Search space sampling: What proportions of the search space are being explored under these different fitness functions?

- Morphology: How do the dominant morphologies depend on the selection preferences? Which morphological properties are influenced most by varying the fitness function?
- Behavior: How is the behavior of the evolved robots affected if we combine a behavior-oriented fitness with morphological preferences?

Related Work

The task of locomotion is commonly applied in Evolutionary Robotics, as for instance Sims (1994), that in his seminal work, evolved morphologies and controllers of robots using a directed-graph-based generative encoding. Likewise, Hornby and Pollack (2001) utilized an L-System as a representation to conjointly evolve “body” and “brain”, despite the difficulty of body-brain co-evolution pointed out in Cheney et al. (2016). Moreover, Wampler and Popović (2009) investigated methods for generating efficient morphological shapes and motion styles. Samuelsen and Glette (2014) studied measures for morphological diversification when evolving generative robot morphologies. And, another study (D’Angelo et al., 2013) examines the influence of the size of the robots in the success of learning a gait, while Aoi et al. (2017) analyzed the coordination of multiple limbs for legged robots. Regarding the relationship between the environment and the morphological traits of the robots, Auerbach and Bongard (2014) verified an interesting influence of the environmental complexity on the complexity of the morphologies. These results, however, do not regard modular robots.

Design Space

Robot Morphology

The phenotypes of the morphologies (“bodies”) are composed of modules (Auerbach et al., 2014) as shown in Fig. 1. The morphologies (Fig. 2) are flat, i.e., the modules do not allow attachment on the top or bottom slots, but only the lateral ones. And, any module can be attached to any other through its attachable slots, except for the sensors, which

Figure 1: Robot modules: Core-component with controller board (C); Structural brick (B); Active hinges with servo motor joints in the vertical (A1) and horizontal (A2) axes; Touch sensor (T). C and B have attachment slots on their four lateral faces, and A1 and A2 have slots on their two opposite lateral faces; T has a single slot which can be attached to any slot of C or B. The sequence of letters (T or n) in C and B indicate if there is a sensor on the laterals left, front, right and back (for C only), in this order.

Figure 2: Example evolved robots, resembling a sea turtle on land that locomotes by powerstroke (left), and a salamander that locomotes by side-to-side (right).

can not be attached to joints. Each module type is represented by a distinct letter in the genotype and by a different color in the phenotype.

For analyzing the morphological properties of the robots, we used a framework proposed in our previous work (Miras et al., 2018), containing eight morphological descriptors. For the current study the framework was incremented with a new descriptor, and additionally, had one of them reformulated. The descriptors, for which a more detailed explanation can be found in our previous work, range from 0 to 1 and are named: Branching, Number of Limbs, Length of Limbs, Coverage, Joints, Proportion, Symmetry and Size. The descriptor Joints was reformulated with the only difference that the concept of effective joint is no longer being attached by both slots to the core-component or a brick, but to any module type.¹ The new descriptor is called Sensors (Fig. 3) and is defined with the Eq. (1):

$$C = \begin{cases} \frac{c}{c_{max}}, & \text{if } c_{max} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where c is the number of sensors and c_{max} is the number of free slots in the morphology.

¹We decided to reformulate it after observing that very often the robots develop limbs formed purely by a sequence of joints.

Figure 3: Example morphology containing a single sensor (green arrow), and two more free slots (orange arrows).

Robot Controller

The phenotype of the controller (“brain”) is a multilayer Artificial Neural Network, not necessarily fully connected (Fig. 4, right). For every joint in the morphology, there exists an equivalent Oscillator neuron in the network, and every sensor is reflected as an input of the network. The intermediate topology may vary for each robot and is composed of neurons which might have Linear or Sigmoid transfer functions. The weights of the connections range from -1 to 1 . In this paper, the neurons are also referred to as nodes and the connections as edges.

Evolution

Generative Encoding

Our generative encoding to represent the genotypes of the robots is a grammatical parallel rewriting system called Lindenmayer-System (Jacob, 1994), and conjointly includes elements relative to both morphology and controller, as in (Hornby and Pollack, 2001). The grammar of an L-System is defined as a tuple $G = (V, w, P)$, where

- V , the alphabet, is a set of symbols containing replaceable and non-replaceable elements.
- w , the axiom, is a symbol from which the system starts.
- P is a set of production-rules for the replaceable symbols.

The following didactic example depicts the process of iterative-rewriting of an L-System. For a determined number of iterations each replaceable symbol is simultaneously replaced by the elements of its production-rule. Given $w = C$, $V = \{C, B, T\}$ and $P = \{C : \{C, B\}, B : \{T\}, T : \{C, T\}\}$, the rewriting follows as:

Iteration 0: C
Iteration 1: C B
Iteration 2: C B T
Iteration 3: C B T C T

Each genotype is a distinct grammar, making use of the same alphabet (Tab. 1), and the alphabet is formed by types of morphology modules and commands to attach them together, as well as commands for defining the structure of the controller. To construct a robot, firstly (early development), the axiom of the grammar is rewritten into a more complex string of elements, according to the production-rules of the grammar (the parameter of number of iterations was set 3). Secondly (late development), this string is decoded into a

Figure 4: Process of decoding an early-developed phenotype into a late-developed phenotype with morphology and controller.

phenotype. The decoding process of the phenotype of morphology and controller is illustrated in Fig. 4. During this construction phase, two references are always maintained in the phenotype, one for the morphology (pointing to the current module) and one for the controller (pointing to the current edge). The application of the commands to the phenotype happens in the current module for the morphology and in the current edge for the controller.

Modules

C	core-component
B	brick
A1	vertical joint
A2	horizontal joint
T	touch sensor

Morphology-mounting commands

addr	add the next module to the right
addf	add the next module to the front
addl	add the next module to the left

Morphology-moving commands

moveb	move reference to the module at the back
mover	move reference to the module to the right
movef	move reference to the module to the front
movel	move reference to the module to the left

Controller-change commands

bnode	add a new node to the neural network
bedge	add a new edge to the neural network
bperturb	perturb the weight of an edge
bloop	add a self-connection edge to the network

Controller-moving commands

bmvFTC	move current edge origin-reference to a child
bmvFTP	move cur. edge origin-reference to a parent
bmvFTS	move cur. edge origin-reference to a sibling
bmvTTC	move cur. edge dest.-reference to a child
bmvTTP	move cur. edge dest.-reference to a parent
bmvTTS	move cur. edge dest.-reference to a sibling

Table 1: Alphabet of the grammars. (cur.=current and dest.=destination)

Evolutionary Operators

For all the experiments, the same evolutionary operators and parameters were applied. The population size was $\mu = 100$, being evolved for 100 generations. In each generation: an offspring $\lambda = 50$ was created by producing 1 individual from each of 50 binary tournaments for parent selection and mutating the new individual. From the resulting population of μ parents and λ offspring, 100 individuals were selected for survival, also using binary tournaments. For each fitness function, the experiments were repeated 10 times. The genotypes were initialized by adding four elements to each production-rule, one of each of the categories, Controller-moving, Morphology-mounting, Morphology-moving and Modules, and, each element of the categories was chosen randomly. The maximum amount m of modules allowed in a morphology was 100. So, during the phase of decoding the genotype into the phenotype, after reaching the maximum, the succeeding modules were ignored. Additionally, morphologies without at least one joint or with intersecting parts were considered invalid, and though having been kept in the population, were not evaluated, receiving a value of zero to the measurement of speed. The crossovers were performed by taking the production-rules randomly from the parents, and all individuals underwent mutation by adding/deleting/swapping one random element from a random production-rule/position. The crossover probability was 100%, as it is possible that during the rewriting of the L-System, only the rules of one single parent end up being expressed. And, as it is not rare that one mutation happens for non-expressed genes, thus, to minimize this effect the mutation rate was 100%.

Fitness Functions

The fitness function for the diversity-oriented search (Novelty Search) is defined as $N = n$, where n is a measure of novelty which is calculated as the average distance to the k -nearest neighbors of an individual, for which $k = 1$ and the distance is the Euclidean distance (Lehman and Stanley, 2008) regarding the nine morphological descriptors. The set of neighbors for the comparison is formed by the current population, plus an archive, to which every new individual has a 10% probability of being added, with the individuals added to the archive remaining in it until the end. This fitness has a purely relative morphological property.

Concerning the definition of the fitness functions for the behavior-oriented searches, some of them combine different properties. The first fitness (Speed 1) is calculated as $S_1 = s$, where s is a measure of the speed (m/s) of the displacement of the robot's head from its initial position to its final position during the evaluation time, having a purely behavioral character. The second fitness (Speed 2) is calculated as $S_2 = s * n$, combining behavioral and relative morphological characters. And, the third fitness (Speed 3) is calculated as $S_3 = s * n * \max(0.1, 1 - E)$, where E is the

descriptor Length of Limbs defined by Eq. (2). This fitness combines behavioral, relative morphological, and absolute morphological characters.

$$E = \begin{cases} \frac{e}{e_{max}}, & \text{if } m \geq 3 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where m is the total number of modules of the body, e is the number of modules which have two of their faces attached to other modules (except for the core-component), and $e_{max} = m - 2$ – the maximum amount of modules that a body with m modules could have with two of their faces attached to other modules, if containing the same amount of modules arranged in a different way. The types of modules would not have to be necessarily the same, as long as the body had the same amount of modules.

Experimental Setup

This study² conducted four distinct experiments³ evolving robots for the behavior of undirected locomotion in a plain terrain, using a different fitness function in each. One of the fitness functions is a diversity-oriented search, i.e., Novelty Search (Lehman and Stanley, 2008), while the other three are behavior-oriented searches. The rationale for having three different behavior-oriented fitness functions is explained hereafter.

A search space can be seen as different layers which build-up over each other, being them the design space, the representation and the reproduction operator. These layers result in a set of phenotypes, which in interaction with the environment can produce a series of behaviors. In our previous work we analyzed the morphological search space of the encoding that we use to manipulate our design space (Miras et al., 2018). The results showed the possibility of reaching great levels of morphological diversity, there being though, a tendency to very commonly discovering morphologies with few and long limbs (note that “long” is relative to the size of the body).

Having these observations in mind, for the current work we tested one behavior-oriented fitness, aiming high speed on the behavior of undirected locomotion, and two more with the same goal but, including penalties regarding preferences for morphological traits. The motivation to apply the penalties was attempting to drive the search away from the previously observed tendencies for the common morphological trait of having few limbs. For the behavior-oriented fitness functions, each robot was evaluated during a period of 30 seconds.

²The code for reproducing all experiments can be found at the link <https://tinyurl.com/y8469tt2>.

³The resulting data of the experiments can be found at <https://tinyurl.com/y9xsyh6q>.

Considering the natural trade-off⁴ among the descriptors Number of Limbs, Length of Limbs and Branching, we may suppose that any of them would play the same role of tackling the tendency for few limbs. Nevertheless, rewarding for Number of Limbs or Branching could lead evolution to simply explore several short limbs or several branching areas with short limbs, respectively. Thus, the descriptor Length of Limbs was chosen for the penalty.

Results and Discussion

Search Space Sampling and Diversity Levels

To examine the sampling of the morphological phenotypic search space, we use a measure we call Number of Sampled Cubes (NSC). The NSC is accumulated along the full evolutionary run, and accounts for the number of cubes in the morphological multidimensional space discovered, that is, the amount of distinct morphologies (within a determined granularity for our morphological descriptors) ever found along the full evolution. In these experiments, the grid of cubes has its dimensions composed of the nine morphological descriptors, with each one divided into 100 bins of size 0.01. Every new morphology is attributed to a proper cube given its descriptors, and if the cube does not contain any morphology yet, one more unit is summed to the NSC.

When we observe the diversity levels, naturally, all the fitness functions show a decline tendency (Mann-Kendall $p < 0.001$) for the mean of the measure of novelty n along the generations (Fig. 5c). Furthermore, as expected, the means of n for N , S_1 , S_2 , and S_3 in the final populations are significantly (Wilcoxon $p < 0.05$) different. This diversity differentiation corroborates a disparity for the speed s , shown in Figure 5a, so that the non-penalized fitness functions are more greedy and better performers, being thus less diverse.

For the NSC measure in the final population (Fig. 5b), the diversity-oriented fitness N naturally presents a much higher value than the behavior-oriented fitness S_1 . Nevertheless, the fitnesses S_2 and S_3 , which combine searches for behavior and diversity, present even higher values (all differences are significant, Wilcoxon $p < 0.05$). The average distance among the points is, though, higher for N than the other functions. Thus, S_2 and S_3 sample more cubes, while N samples cubes further from each other.

In Fig. 6 we see the concentration of morphologies in the cubes of the morphological multidimensional space, explored by each fitness function along all the evolutionary runs (grouped together). These plots show the existence of cubes which are clearly much more often revisited than the majority of the cubes, with N and S_1 having much more concentration of revisited cubes than S_2 and S_3 . This observation helps to explain why the last two have a significant

⁴The higher the Branching, the higher the Number of Limbs (Pearson 0.3); The higher the Length of Limbs the lower the Limbs (−0.4) and the Branching (−0.1), $p < 0.05$.

Figure 5: Progression of the mean of the population averaged over all runs for (a) The speed measure s for the behavior-oriented fitness functions; (b) The NSC measure; (c) The novelty measure n .

Figure 6: Concentration of morphologies in each sampled cube of the morphological multidimensional space considering all runs. Including all morphologies explored, each bar represents the frequency of morphologies in one specific cube, and the bars are sorted in descending order. All charts are scaled on the y-axis limited to 200 and in the x-axis limited to 500.

greater value for the measure NSC, once by revisiting discovered cubes less often they explore more the space.

When we correlate the frequencies for cubes of the individual runs for each function, we see (Fig. 7, left) that the runs of N are highly positively correlated, while the runs of S_1 , S_2 and S_3 have much weaker correlations.

This way, although also unclear why, it seems that by searching purely for diversity, evolution has a stronger bias when sampling the search space in independent runs than by searching for behavior. We should clarify here that this tendency regards only a regularity in the most commonly explored morphologies. For instance, although S_1 has a much lower value of NSC than N when sampling the search space,

Figure 7: (left) Median Pearson correlations of the frequencies of the cubes between the different runs for each fitness function. For each function, a matrix comparing all pairs of runs was performed, and the significant cases ($p < 0.05$) were used to calculate medians of Pearson correlations. (right) Pearson correlations for the frequencies of the cubes between all pairs of fitness functions.

the most common cubes explored by the different runs of N present a much greater intersection than the ones explored by S_1 . That is, N samples more and is more consistent when the search is repeated.

Additionally, comparing the cubes⁵ explored by all runs grouped together between each pair of fitness functions, we verify also that all pairs are positively correlated (Fig. 7, right). Considering that it includes the pair of functions that searches purely for diversity and purely for locomotion (N vs S_1), this could perhaps be evidence of the influence of the search space in the morphological traits of the population.

⁵The frequencies presented in Fig. 6 were correlated between each pair of fitness function. As each function might find different sets of cubes, the cubes that exist for one function of the pair and not for the other are valued as zero for this last one. Pearson correlations with $p < 0.001$.

Figure 8: The most commonly generated morphologies in 3 sample runs. From left to right, the most to the least common.

Morphological Traits

To illustrate the tendency of the reproduction operator, Figure 8 shows the 10 most common morphologies of three sample runs for each fitness function. These morphologies usually possess one or two limbs only, using the other modules to make these limbs a little longer. Although S_3 mainly follows the same trend, there are also cases of tree or four limbs, demonstrating the influence of the selection towards the traits related to the penalty.

On the other hand, when observing the best robots in Fig. 9, we see that S_3 and N produce champions with morphologies very different from the common ones, possessing actually multiple limbs. The best individuals in S_2 also deviate from the common shapes, while S_1 presents largely similar robots having a single limb.⁶

So far the presented results might lead us to suspect that the predominant morphological traits of the champion robots for S_1 and also largely S_2 are due to the observed tendency of the search space for few limbs. However, assessing the distributions of the morphological descriptors in the final population, we can have additional insights about the influence of the selection on the morphologies. Figure 10a shows the difference in the means of the morphological descriptors among the fitness functions, and the first relevant fact is that the predominant traits when realizing the diversity-oriented search are usually different from when realizing any of the behavior-oriented searches.⁷ And more importantly,

⁶A video showing the locomotion of some of these robots is available at <https://tinyurl.com/ydbs7nt8>.

⁷All pairs of comparisons between N and the behavior-oriented

Figure 9: Best individuals of the last generation in one sample run. From left to right, the best to the worst. The locomotion styles are: Speed 1: Rolling, for all morphologies; Speed 2: 1-Sidewinding, 2-Undulating Worm, 3-Rowing, 4-Rolling, 5-Walking, 6-Rowing, 7-Rowing, 8-Undulating Worm, 9-Rowing, and 10-Rowing; Speed 3: 1-Rowing, 2-Rowing, 3-Walking, 4-Walking, 5-Sidewinding, 6-Walking, 7-Walking, 8-Crawling, 9-Walking, and 10-Walking.

these differences seem characterized in different directions, according to the nature of each fitness function. For instance, S_1 and S_2 present lower values of Branching, Number of Limbs, and Symmetry, with a greater value of Length of Limbs and Joints than N , while for S_3 , which includes a penalty for long limbs, it is exactly the opposite. This example demonstrates the morphological penalties leading to more multi-limb, branching, and symmetrical morphologies. When observing the mean of the descriptors along the generations in Fig. 11, we see the tendency (Mann-Kendall $p < 0.001$) of Branching and Number of Limbs to increase and of Length of Limbs to decrease for S_3 , still according to the penalty. Thus, when we evaluate the result of the full search with S_3 (Fig. 8) morphologies with few limbs seem dominant, while studying the progress over generations (Fig. 9) this tendency is decreased. And with S_2 , it is also decreased, though much less.

All behavior-oriented searches seem to explore Size, which is not a surprise, since a large body more easily can produce a large displacement of the head and thus achieve a high speed s . Proportion, which is kept at a stable level for most functions, drops drastically for S_1 , as this case was dominated by single-limb, and thus disproportional, robots. Coverage decreases for most functions but S_1 , as its predominant final morphologies are similar to snakes (Fig. 9), covering the whole body area.

searches are significantly (Wilcoxon $p < 0.05$) different, except for N vs S_2 of Branching, Size and Sensors, and for N vs S_3 of Proportion.

Figure 10: Comparison of the mean of the runs for the morphological descriptors. (a) shows the measures for all fitness functions in the final populations (S and I refer whether the difference in the averages between a determined fitness function and N is significant or insignificant, respectively), (b) compares the High and Low Rankers during the full evolutionary process; only significant differences (Wilcoxon $p < 0.05$) were plotted. M1=Branching, M2=Number of Limbs, M3=Length of Limbs, M4=Coverage, M6=Proportion, M7=Symmetry, M8=Size, and M9=Sensors.

Figure 11: Mean of the population averaged over all runs for the morphological descriptors along the generations.

Moreover, we try to understand the traits of the high and low rankers, seeing clearly the influence of the selection in the morphological traits of the population. The Low and High Rankers regard two groups of robots with significantly different average speed. To divide the populations according to this concept, for each fitness function, the median of the speed was calculated to separate the population into two groups between which the averages of the runs were tested statistically significant (Wilcoxon $p < 0.001$). The group with the highest average speed was considered the group of High Rankers, and the one with lowest speed was considered the group of Low Rankers. Fig. 10b shows the difference in the means of the morphological descriptors for the Low and High Rankers. The most interesting cases are Number of Limbs and Symmetry, which present higher means for the Low Rankers with all fitness functions, that is, the more limbs, the slower. This is in accordance with the previous observation that the more penalized the fitness function,

the higher the mean of Number of Limbs/Symmetry and the lower the mean of s .

Behavior

Figure 5a displays the average speed for the three behavior-oriented fitness functions. Comparing the mean speeds between each function in the final populations, we verify that they are significantly (Wilcoxon $p < 0.05$) different. The more penalized the fitness, the lower the speed, thus, in the order S_1 , S_2 , S_3 , and finally N , which is purely novelty.

When observing the prevalent locomotion style obtained with each fitness function, we see that for S_3 , the style is Walking (Fig. 9), characterized by a synchronized alternation of limb movement, while for S_2 it is Rowing, characterized by simultaneous limb stroking, and for S_1 , it is Rolling.⁸ Additionally, in Fig. 2 we see examples of man-

⁸Videos shown at <https://tinyurl.com/ydbs7nt8>.

ually selected robots which are not included in the best-rank, but which present an interesting life-like locomotion style. Thus, in our experiment, walking and other complex locomotion strategies proved relatively slow compared to the simple, rolling ones. One possible explanation for this could be that one of the most efficient ways to achieve a high speed according to our fitness measure and environment, is to grow large and through this achieve a large absolute displacement of the robot's head during the fitness evaluation. Also, rolling can be an efficient locomotion style as long as the environment permits it, and similar behavior has been observed in other studies (Samuelson and Glette, 2014).

However, “for many animals, natural selection may tend to favor 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” (Alexander, 2003). It is not clear why the trait of few long limbs showed to be predominant, and the above remark makes us wonder if it might be due, not to some advantage of having fewer limbs, but to the challenge of having multiple limbs. For example, having one limb that permits locomotion is a challenge in itself, while having multiple limbs not only multiplies this challenge, but also carries an additional challenge of synchronization, to avoid limbs pulling in different directions and impairing displacement. Perhaps adding a life-time learning ability to the robots would allow them to learn how to use their limbs better and obtain higher speed. Moreover, given that the complexity of the environment may influence the complexity of the resulting robot morphologies (Auerbach and Bongard, 2014), maybe an increase in the complexity of our environment would have an interesting effect on the morphological traits of the population.

Conclusion and Further Work

We investigated the evolution of modular robots from different perspectives: coverage of the search space, evolved morphologies, and evolved behaviors. We compared several selection preferences (i.e., fitness functions) and observed that novelty selection sampled fewer cubes in the space of possible morphologies than a combination of speed and novelty, cf. Fig. 5 (b). This is a surprising result and currently we are researching this ‘anomaly’ to understand and clarify it. Furthermore, we saw a clear effect of the fitness function on the evolved morphologies. The final populations in Fig. 9 exhibit visible patterns internally and clear differences from each other. Looking at the morphological descriptors we noted that all of them are influenced by the selection preferences. Considering evolved behaviors, the outcomes are not surprising in terms of pure speed values, cf. Fig 5 (a), but the locomotion patterns are interesting to see as they often resemble gaits of existing animals, <https://tinyurl.com/ydbs7nt8>. Overall, the penalized fitness functions resulted in slower robots, pointing to a trade-off between maintaining diversity and increasing

quality. For further work we intend to 1) deepen the investigation of the search space coverage; 2) evolve in a non-plain terrain to test the effect of the environment in the morphologies and controllers; 3) add life-time learning to assess its effect on the morphologies and the structure of the controllers.

Acknowledgements

This work has received funding from the European Union's PHOENIX project under grant agreement No 665347.

References

- Alexander, R. M. (2003). *Principles of animal locomotion*. Princeton University Press.
- Aoi, S., Manoonpong, P., Ambe, Y., Matsuno, F., and Wörgötter, F. (2017). Adaptive control strategies for interlimb coordination in legged robots: A review. *Frontiers in neurorobotics*, 11:39.
- Auerbach, J., Aydin, D., Maesani, A., Kornatowski, P., Cieslewski, T., Heitz, G., Fernando, P., Loshchilov, I., Daler, L., and Floreano, D. (2014). Robogen: Robot generation through artificial evolution. In *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 136–137. The MIT Press.
- Auerbach, J. and Bongard, J. (2014). Environmental influence on the evolution of morphological complexity in machines. *PLOS Computational Biology*, 10(1):e1003399.
- Cheney, N., Bongard, J., Sunspiral, V., and Lipson, H. (2016). On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures. In *Proceedings of the Artificial Life Conference (ALIFE XV)*, pages 226–234. MIT Press.
- D'Angelo, M., Weel, B., and Eiben, A. (2013). Online gait learning for modular robots with arbitrary shapes and sizes. In *International Conference on Theory and Practice of Natural Computing*, pages 45–56. Springer.
- Hornby, G. S. and Pollack, J. B. (2001). Body-brain co-evolution using l-systems as a generative encoding. In *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation*, pages 868–875. Morgan Kaufmann Publishers.
- Jacob, C. (1994). Genetic l-system programming. *Parallel Problem Solving from Nature PPSN III*, pages 333–343.
- Lehman, J. and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pages 329–336.
- Miras, K., Haasdijk, E., Glette, K., and Eiben, A. E. (2018). Search space analysis of evolvable robot morphologies. In *Applications of Evolutionary Computation - 21st International Conference, EvoApplications 2018*, volume 10784 of *Lecture Notes in Computer Science*, pages 703–718. Springer.
- Samuelson, E. and Glette, K. (2014). Some distance measures for morphological diversification in generative evolutionary robotics. In *GECCO 2014 - Proceedings of the 2014 Genetic and Evolutionary Computation Conference*, pages 721–728.
- Sims, K. (1994). Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372.
- Wampler, K. and Popović, Z. (2009). Optimal gait and form for animal locomotion. In *ACM Transactions on Graphics (TOG)*, volume 28, page 60. ACM.