

Overcoming Initial Convergence in Multi-Objective Evolution of Robot Control and Morphology Using a Two-Phase Approach

Tønnes F. Nygaard, Eivind Samuelsen, and Kyrre Glette

Department of Informatics, University of Oslo, Norway
{tonnesfn,eivinsam,kyrrehg}@ifi.uio.no

Abstract. Co-evolution of robot morphologies and control systems is a new and interesting approach for robotic design. However, the increased size and ruggedness of the search space becomes a challenge, often leading to early convergence with sub-optimal morphology-controller combinations. Further, mutations in the robot morphologies tend to cause large perturbations in the search, effectively changing the environment, from the controller’s perspective. In this paper, we present a two-stage approach to tackle the early convergence in morphology-controller co-evolution. In the first phase, we allow free evolution of morphologies and controllers simultaneously, while in the second phase we re-evolve the controllers while locking the morphology. The feasibility of the approach is demonstrated in physics simulations, and later verified on three different real-world instances of the robot morphologies. The results demonstrate that by introducing the two-phase approach, the search produces solutions which outperform the single co-evolutionary run by over 10%.

1 Introduction

Evolutionary robotics (ER) is an approach for automatic design and adaptation of robot bodies and control systems, through the use of evolutionary algorithms (EAs). This has the potential for improving future robot design processes, as well as improving the resilience and adaptability of robots [1,2].

While most work in evolutionary robotics concerns the optimization of robotic control systems, there has also been an effort, mainly through the use of software simulations, into simultaneously optimizing robotic bodies and control systems [3]. While examples of such systems are still limited, it is believed that optimizing both aspects simultaneously could lead to interesting and robust behaviors, even when limited control and actuation is available. Most examples of the potential of co-evolution of morphologies and controllers are found within the field of virtual creatures, where a physical counterpart of the simulated system may not be available [4,5]. However, there have also been some examples of evolution in simulation leading to instantiation of working physical robots [3,6,7], and recently even simulationless evolution of robotic morphologies [8].

Optimizing robot morphologies together with control systems is a difficult task, and we have only so far seen relatively simple results, even though significant amounts of computational resources have been allocated [3]. One of the reasons for this difficulty is the increased dimensionality of the search space that comes with the freedom to design the body, in addition to the control system. In effect this requires a much larger amount of exploration to find the desired quality of solutions. However, another aspect of the difficulty is that the co-evolution of morphology and control also leads to a much more difficult search – the search landscape is much more rugged, and small changes in the morphology can easily offset the performance of a previously found good body-controller combination. An interesting view on this, based on the theory of embodied cognition, can be found in [9]. Here, a morphological change is seen as scrambling of the interface between the controller and its physical interface to the environment.

With similar observations in our own work on co-evolving morphologies and control systems [10], in this paper we demonstrate a simple approach to reduce this effect. We demonstrate that it is possible to improve the solution quality after convergence of a morphology-controller evolutionary run by freezing the morphology, and allowing a second phase of continued evolution on the control system alone. This approach leads to significant improvements in the final solutions with a modest number of further evaluations, and indicates that phases of optimizing only the control system can be beneficial. Our method can easily be combined with other approaches, like novelty search [11] or various methods of encouraging behavioral diversity [12]. We demonstrate the approach for a realistic robot design, and validate the method by producing and evaluating two of the evolved robot designs in the real world. The performances of the evolved designs are compared to a hand-designed robot within the same framework.

2 Background

The perhaps most well-known example of co-evolution of morphology and control is Sims’ work on virtual creatures [4]. Here, virtual creatures were evolved in a rigid body physics simulation, displaying diverse and natural behaviors. While these virtual creatures were not targeted for instantiation in the physical world, the setting of virtual creatures has been popular for exploring various aspects of co-evolving morphologies and control. Examples include: evolution of soft bodied creatures [13], evolution of artificial muscle-based creatures [5], diversity-enhanced morphology evolution [14], and investigations of the relationships between environment and body complexity [15].

While results from simulations can give a number of insights, being able to automatically design physical robot morphologies is inarguably even more interesting, and could eventually lead to useful techniques in robot engineering. However, going from the already difficult task of designing virtual creatures to designing physically instantiable robots is challenging. The encoding of the solutions need to take into account physically realizable body parts and actuators, and a sufficiently accurate simulation of these. Further, actually producing the

robots can be a relatively cumbersome and time-consuming task, and usually only a handful of the results from simulation can therefore be evaluated as real-world robots.

Finally, the combined morphology-controller robots are prone to reality gap issues, potentially even more so than when control systems only are evolved. In particular, it seems to often be the case that some morphology-controller combinations work relatively as expected from simulation, while other combinations suffer from a large reality gap [3,6,7].

Experiments on evolving morphologies and controllers employ a variety of different building blocks, from convenient physics simulation primitives in the case of virtual creatures [4,13], to building blocks which are convenient to prototype and integrate with actuators in the case of physical instantiation [3,6,16,7]. In these cases, focus has been more on the evolutionary design process, and the resulting physical robots are mostly meant to validate the approach.

However, there have also been some attempts at evolutionary morphology design with an engineering perspective in mind, range from rather free-form structures, which are given large design freedom, to structures inspired by more "realistic" robot shapes, which are intended to solve a task or carry a payload. These include the design of a variety of shapes for robotic manipulators [17], or parametric legged robots [18]. It should be noted that while these works worked on realistic robot shapes, they have so far only been validated in simulation. There are also some good examples of parametric design optimizations for bio-inspired robot morphologies, such as the optimization of the caudal fin dimensions for a robotic fish [19], or the optimization of leg parameters for an octopus-inspired robot [20].

It has been pointed out that co-evolution of morphologies and controllers is difficult [9]. Methods which have been explored to tackle this challenge includes generative encodings [6,13,10], morphological diversity-enhanced algorithms [14,10,20], and more complex environments [15].

3 Implementation

3.1 The Robot

Six legged robots with three degrees of freedom for each leg were used for all experiments. The legs are arranged in a spider configuration, and vacant slots for tools or two more legs were added to the front of the robot for versatility and future experiments.

A parametric blueprint was defined from which each robot was defined. The choice of parameters was limited by the computational complexity of the increasing solution space, while the ambition of a more general robot system provides a contrast with the demand for more dimensions. The final compromise assumes symmetry along the movement direction of the robot to allow a sufficiently complex robot with fewer parameters. The lengths of the two outer leg segments can be set, with minimum lengths, given by the size of the servos used, and the

maximum given by the manufacturing equipment available. Six legs with two configurable segment lengths each yields 12 parameters, but assuming symmetry reduces this to 6. The placement of each servo on the base plate was also made parameterizable, as seen in Fig. 1. A total of 11 parameters defines the complete morphology of the robot, with details given in table 1.




Fig. 1: Drawing of the base, showing all parameters. Table 1: The parameters used in morphology.

The physical instantiation of each robot is built using plastic parts printed in ABS by a Fortus 250mc 3d printer. The parts were designed for high strength and low weight, and FEM simulation was used extensively in the design process. Rapid prototyping principles were used to reduce risk of design errors and encourage good project progression, and the first manually designed robot can be seen in Fig. 2. We chose Dynamixel AX-18A smart servos as actuators for their ease of use and relatively high power, and implemented the whole control scheme on a windows based computer.




Fig. 2: Image of the printed and assembled hand designed robot. The face was added for recruitment events, to make the robot seem less frightening to children.

3.2 Control System

We implemented a simple control system that serves as an extension of a popular wave controller [21]. The traditional controller contains a collection of amplitude and a phase parameters that together with the source wave signal generates commands for all actuators. We also needed an offset value for each servo, but through initial tests found the resulting controller to be hard to limit and evolve. We therefore introduced two new equations that uses v as minimum angle and w as maximum angle, as seen in equation 1. An offset ϕ is added to provide the full controller equation, given in equation 2. The time t is given in seconds. Symmetry was also used in the control system, and each pair of servos was thus defined by its minimum movement, maximum movement, and phase offset. This results in a total of 36 parameters to define the movement of the robot, as seen in table 2.

$$\begin{aligned}\alpha &= \frac{(\nu - \omega)}{2} \\ \beta &= \frac{(\nu + \omega)}{2}\end{aligned}\tag{1}$$

$$\chi(t, \nu, \omega, \phi) = \frac{(\nu - \omega)}{2} \times \tanh(4 \times \sin(2 \times \pi \times (t + \phi))) + \frac{(\nu + \omega)}{2}\tag{2}$$

Coxa movement ₁	4	-0.81	1.64
Coxa movement ₂	2	-1.64	1.64
Femur movement	6	-2.49	2.49
Tibia movement	6	-2.49	2.49
Phases	18	$-\pi$	π

Table 2: Table showing search space for morphology at the top and control at the bottom. Coxa movement₁ is moving along the side of the robot, while Coxa movement₂ is moving behind the robot.

3.3 Evaluation

The main goal of each individual is a fast walking gait, but the weight of individuals was also used in the evolutionary algorithm to ensure a higher degree of variation in the resulting population, and to aid in the evolutionary search by providing two partially conflicting objectives [22].

A custom in-house developed simulation framework was used to rapidly test a range of different individuals. The framework acts as an interface to Nvidia PhysX library, and generates a simplified robot morphology based on the morphological parameters. It uses the same controller code as used in hardware tests.

The evolutionary framework ParadisEO [ref] is used for running the evolutionary algorithm.

Real world evaluations were done in the University of Oslo’s motion capture lab. It features an OptiTrack motion capture system to efficiently estimate the position and pose of the robot, using infrared cameras and the light from reflective spheres mounted on the robot. Each gait is evaluated for four seconds. A wiring harness was built to ensure that the robot did not trip in its own power and command wires, and turning gaits were evolved and used to ensure the robot stays within the area covered by the cameras at all times. We achieved a fully autonomous testing system that needs no human intervention during testing or evolutionary experiments, which allowed for prolonged tests in hardware.

3.4 Evolutionary Setup

We chose NSGA-II [23] as the evolutionary algorithm for our experiments. 32 bit floating point numbers were used to represent all parameters in the genome.

We used non-uniform mutation using the normal distribution with standard deviations 0.025 and 0.01 for control and morphology mutation respectively. Random reset mutation was used with a per-gene probability of 0.02. Tests with several different crossover operators did not yield significant advantages, so these were not used.

4 Experiments and Results

4.1 Experiments

The main co-evolutionary runs in simulation were done using 1024 generations of 256 individuals. A total of 20 runs were done, to reduce statistical anomalies. Two morphologies were then selected from the results, and were, together with the hand designed robot, selected for further evolution. Control systems for these three robots were then evolved with three new groups of runs, using the same evolutionary parameters as before. Morphology was locked, and the controllers were all reset to random values at the start of the new runs. A smaller number of evolutionary runs were then conducted on each morphology to generate turning gaits needed for continuous automated tests in hardware without human intervention.

The best performing individual from each of the three control-exclusive evolutionary runs were 3d printed and built according to their evolved morphology. The performance was then verified in hardware by running the gait on the physical platform in the motion capture lab, and movement was recorded and compared to performance in simulation. Each individual gait speed was recorded over 4 seconds, and two sets of 128 evaluations were done for each individual. The two sets were done at different times with other tests in between, to reduce the effect of noise in the test setup.

4.2 Results

Fig. 3 shows a typical development of fitness over the run time for an EA. We can see from the relatively low increase during the last generations that the algorithm is close to convergence. We also see a fairly large difference between the global best distance and the average best distance of the 20 evolutionary runs.




Fig. 3: The distance travelled for morphology runs. Highest fitness across all runs is shown in green, while the blue line shows the average of the maximum fitness for all runs. Average distance of all last generation individuals is seen in red.

Two robots were selected from the runs, one smaller and one larger than the hand designed robot. The smaller evolved robot had a 6% lower body weight (not including servos), while the larger robot had a 2% higher weight. The robots can be seen in Fig. 4

Table 3 shows the improvement of the additional control-exclusive phase, compared to the single co-evolutionary run. We see that the two robots with evolved morphologies gets a considerable improvement at 15.5% and 11.7%. The manually designed robot only achieves a 3.2% increase over the best performing individual from the co-evolutionary run.

Fig. 5 shows the results of the co-evolutionary runs, with the Pareto front marked in red. The results from the evolution of control is seen in green, which outperforms the co-evolutionary runs significantly. As in table 3, we also here see that the manually designed robot, in the middle green line, does not show the same improvement over the Pareto front as the other two.




Fig. 4: The three robots used for real world evaluations. The small evolved robot is to the left, the hand designed robot is in the middle, and the large evolved robot is to the right.




Fig. 5: The Pareto front of the final morphology runs, compared to the three groups of evolutionary runs for the robot control systems.

Robot	1st phase result	2nd phase result	Improvement
Manually designed	-	0.299m/s	3.2%*
Small evolved robot	0.203m/s	0.234m/s	15.5%
Large evolved robot	0.304m/s	0.339m/s	11.7%

Table 3: The performance of the 1st phase (co-evolutionary run), 2nd phase (control-exclusive evolutionary run) and the percentage increase when compared to the Pareto front. A 1st phase result is missing from the manually designed robot since this morphology is not picked from the Pareto front of the co-evolutionary run. The improvement is compared to the highest performing robot of comparable weight.




Fig. 6: The evolution of control for the small evolved robot. The original top fitness is marked in purple, and is surpassed already at generation 126.

The performance of the robots tested in the real world can be seen in Fig. 7. We see that the evolved gait for the manually designed robot does not outperform the manually designed gait, though the variance in measured speed is greater. We see that the speed of the smaller evolved robot is slightly lower than the manually designed robot, and that the large evolved robot outperforms both. This is consistent with the simulation results. The original fitness values in the simulation is seen in green, and shows a fairly large reality gap between simulation and real world.




Fig. 7: Box plot of the speeds from the motion capture of evolved gaits on all robots. The green dashed line indicates the speed of the robot in simulation.

5 Discussion

We made the following observations from the experimental results:

- The selected robots display different optimization trade-offs in terms of weight reduction and speed increase, and we confirmed that these optimizations were still valid on the real world robots. We see from Fig. 7 that the smaller evolved morphology has comparable performance to the much heavier hand designed robot, giving further proof of the effectiveness of evolving the morphology in addition to control. We also see that the reality gap is very different between the different robots, though this is most likely caused by inaccuracies in the modeling of the motors in simulation, as it seems proportional to the weight of the robot. The reality gap is relatively large for all morphologies, but preliminary experiments on real-world adaptation show promising results towards reducing the performance loss.
- We see from Table 3 and Fig. 5 that the hand designed robot performs worse than the two evolved robots when compared to the Pareto front of the co-evolutionary run. This is due to the evolutionary algorithm finding better morphologies than the hand designed version, and the comparison of the hand designed robot with an evolved controller to an arbitrary co-evolved robot with similar weight only yields a slight improvement.
- We see from Fig. 3 that the evolutionary algorithm has most likely converged on local optima after 1024 generations, and we do not expect considerable jumps in fitness if we were to run this algorithm for another 1024 generations. The two stage evolution we propose adds a second run of controller only optimization, and there are mainly two features of this second run that could explain the increase in the quality of the solutions found: A reduced search space due to the frozen morphology genes, and a difference in fitness landscape ruggedness. Smaller search spaces may typically require fewer evaluations before convergence, but do not necessarily affect the ability to escape local optima. We therefore believe that the changed shape of the fitness landscape may be an important factor here, in line with the thoughts in [9] of locking the controller’s interface to the environment. It would however be necessary to do further experiments to fully confirm this.
- The second evolutionary run works in a smaller search space due to the locked morphology, but another aspect of this approach is that it also starts with new, random control systems. We originally did not expect this step to make a large difference, since we are using random reset mutation throughout the co-evolutionary run, but early experiments showed significant improvements. This most likely shows that the ability to surpass the original convergence is both due to the changes in landscape, and partial reset of genes. The action of resetting the genomes could be viewed as spreading out the initial solutions over the entirety of the now morphologically frozen search landscape, and may thus lead to better exploration towards optimal solutions, as compared to a less diversified population of already converged control genes.

- While it may not be required to do a full evolutionary run in the second phase, the experiments demonstrate that there is room for adaptation and performance increase after the morphology has been frozen, which warrants further explorations into different types of controller optimizations combined with the co-evolutionary search for morphology and control.

6 Conclusion and Future Work

In this paper, we investigated using a two-phase evolution method to exceed the initial convergence found in the single co-evolutionary case. We saw that locking morphology and doing a re-evolution of control on two different morphologies showed continued improvement of between 10% and 15%, which was unlikely to be found by extending the original co-evolutionary run. Tests in hardware showed that the gaits and morphologies found worked also in the real world.

It would be interesting to further investigate the inner workings of why we are able to exceed the initial fitness found, and if this could be implemented as an operator within a single evolutionary run. Perhaps using either some sort of random mutation operator to emulate the reset of controllers between runs, or an operator that locks the morphology for some number of generations, might show some improvement over traditional co-evolutionary techniques. Expanding on this two-phase evolutionary method by using many alternating phases might show an even higher ability to escape local minima than we found here, and might allow continued evolution on not only the controller, but the morphology as well. It would also be natural to consider lifetime learning schemes, i.e. allowing the controller to adapt to the morphology before the fitness value is recorded. We already have some promising results from using such schemes for fixed morphologies [24]. In any case, these approaches could also be combined with other methods for preventing premature convergence, such as diversity-enhancing search methods [14].

References

1. Bongard, J.C.: Evolutionary robotics. *Commun. ACM* **56**(8) (2013) 74–83
2. Doncieux, S., Bredeche, N., Mouret, J.B., Eiben, A.E.G.G.: Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI* **2**(March) (2015) 4
3. Lipson, H., Pollack, J.B.: Automatic design and manufacture of robotic lifeforms. *Nature* **406**(6799) (2000) 974–978
4. Sims, K.: Evolving virtual creatures. In: *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, ACM (1994) 15–22
5. Lessin, D., Risi, S.: Soft-Body Muscles for Evolved Virtual Creatures: The Next Step on a Bio-Mimetic Path to Meaningful Morphological Complexity. In: *European Conference on Artificial Life*. (2015) 761–762
6. Hornby, G.S., Lipson, H., Pollack, J.B.: Generative representations for the automated design of modular physical robots. *IEEE Transactions on Robotics and Automation* **19**(4) (2003) 703–719

7. Samuelsen, E., Glette, K.: Real-world reproduction of evolved robot morphologies: Automated categorization and evaluation. In: *Lecture Notes in Computer Science*. Volume 9028. (2015) 771–782
8. Brodbeck, L., Hauser, S., Iida, F.: Morphological evolution of physical robots through model-free phenotype development. *PLoS ONE* **10**(6) (2015) 1–17
9. Cheney, N., Bongard, J., Sunspir, V., Lipson, H.: On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures. In: *Proceedings of the Artificial Life Conference 2016 (ALIFE XV)*, MIT Press (2016) 226–234
10. Samuelsen, E., Glette, K.: Some distance measures for morphological diversification in generative evolutionary robotics. In: *GECCO 2014 - Proceedings of the 2014 Genetic and Evolutionary Computation Conference*. (2014) 721–728
11. Lehman, J., Stanley, K.O.: Abandoning objectives: Evolution through the search for novelty alone. *Evol. Comput.* **19**(2) (June 2011) 189–223
12. Mouret, J.B., Doncieux, S.: Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary computation* **20**(1) (2012) 91–133
13. Cheney, N., MacCurdy, R., Clune, J., Lipson, H.: Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding. In: *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*, ACM (2013) 167–174
14. Lehman, J., Stanley, K.O.: Evolving a diversity of virtual creatures through novelty search and local competition. In: *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, ACM (2011) 211–218
15. Auerbach, J.E., Bongard, J.C.: Environmental influence on the evolution of morphological complexity in machines. *PLoS computational biology* **10**(1) (2014)
16. Lund, H.H. In: *Co-evolving Control and Morphology with LEGO Robots*. Springer (2003) 59–79
17. Leger, C.: *Automated synthesis and optimization of robot configurations: an evolutionary approach*. PhD thesis, Carnegie Mellon University (1999)
18. Passault, G., Rouxel, Q., Fabre, R., N’Guyen, S., Ly, O.: Optimizing Morphology and Locomotion on a Corpus of Parametric Legged Robots. In: *Biomimetic and Biohybrid Systems: 5th International Conference, Living Machines 2016*, Edinburgh, UK, July 19-22, 2016. *Proceedings*. Springer (2016) 227–238
19. Clark, A.J., Moore, J.M., Wang, J., Tan, X., McKinley, P.K.: Evolutionary Design and Experimental Validation of a Flexible Caudal Fin for Robotic Fish. In: *Artificial Life 13*. (2012) 325–332
20. Corucci, F., Calisti, M., Laschi, C.: Novelty-Based Evolutionary Design of Morphing Underwater Robots. In: *Proceedings of Genetic and Evolutionary Computation Conference*. (2015) 145–152
21. Koos, S., Cully, A., Mouret, J.B.: Fast damage recovery in robotics with the t-resilience algorithm. *International Journal of Robotics Research* **32**(14) (December 2013) 1700–1723
22. Deb, K., Srinivasan, A.: Innovization: Discovery of innovative design principles through multiobjective evolutionary optimization. In Knowles, J., Corne, D., Deb, K., Chair, D., eds.: *Multiobjective Problem Solving from Nature*. *Natural Computing Series*. Springer Berlin Heidelberg (2008) 243–262
23. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: Nsga-ii. *Trans. Evol. Comp* **6**(2) (April 2002) 182–197
24. Ruud, E.L., Samuelsen, E., Glette, K.: Memetic Robot Control Evolution and Adaption to Reality. In: *ICES: 2016 IEEE International Conference on Evolvable Systems, Proceedings*. (2016)