Towards Growing Robots: A Piecewise Morphology-Controller Co-adaptation Strategy for Legged Locomotion

David Hardman, Thomas George Thuruthel(✉), and Fumiya Iida

Bio-Inspired Robotics Lab, Department of Engineering,
University of Cambridge, Cambridge, UK
{dsh46,tg444,fi224}@cam.ac.uk

Abstract. Control of robots has largely been based on the assumption of a fixed morphology. Accordingly, robot designs have been stationary in time, except for the case of modular robots. Any drastic change in morphology, hence, requires a remodelling of the controller. This work takes inspiration from developmental robotics to present a piecewise morphology-controller growth/adaptation strategy that facilitates fast and reliable control adaptation to growing robots. We demonstrate our methodology on a simple 3 degree of freedom walking robot with adjustable foot lengths and with varying inertial conditions. Our results show not only the effectiveness and reliability of the piecewise morphology controller co-adaptation (PMCCA) strategy, but also highlight the need for morphological adaptation as a robot design strategy.

Keywords: Morphological adaptation · Growing robots · Control optimization

1 Introduction

Adaptation in an embodied agent can be found along two time scales: evolutionary and developmental. Evolutionary adaptations occur along larger time-scales at a slower rate. Over several generations, through a selection process, high functioning morphologies (body) and control parameters are evolved. Developmental adaptations are faster and occur within a generation [1,17]. They are more specialized to the working environment and are not transferable within generations [3,5]. There are numerous technological and algorithmic difficulties involved with the co-optimization of the body and control in evolved systems [11]. Likewise, there are many technological challenges in creating morphologically developing systems.

The body morphology has an important role in shaping the behavior of an embodied system [14,18], the influence of which can even extend through generations [9]. The role of morphological development, on the other hand, is not well

Supported by Mathworks Inc. Media available at: https://youtu.be/Xd6axFgqYyg.

A. Mohammad et al. (Eds.): TAROS 2020, LNAI 12228, pp. 357–368, 2020.
https://doi.org/10.1007/978-3-030-63486-5_37

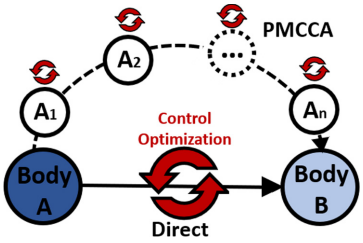

Fig. 1. The morphology, control and behavior of a dynamical system are interconnected. The PMCCA strategy undergoes piecewise morphological changes and hence local control adaptations, leveraging information from its ascendants. A single step morphological change, as shown above, would most likely require a complete recalibration of the control parameters.

understood [4]. Kreigman et al. showed that artificial organisms which developed environment-mediated morphological growth exhibited better robustness to slight abberations [10]. Similarly, another study showed that incorporating the evolution-of-development along with the evolution-of-body led to faster discovery of desirable behavior and higher robustness [1], a conclusion backed even by experimental studies [19].

One of the key differences between morphological changes induced by evolution and development is the continuity of the body plan over time. Evolved body plans can be drastically different from their parents. Developed bodies, by virtue of physical constraints, maintain continuity in their *morphology space*. This constraint causes smooth variations in the body dynamics which is expected to lead to continuous changes in the behavioral dynamics (before the bifurcation point) [7]. The control adaptation required to retain desired behavior, hence, can also be expected to be continuous. This way the control adaptation problem can be reduced to a simple local search problem for each piecewise change in morphology (referred to jointly as the piecewise morphology controller co-adaptation strategy).

This work, unlike other related works on design optimization of morphology [16] and co-optimization of morphology and controller [15] for locomotion, is not a global optimization strategy. The main objective is to obtain locally optimum controllers or morphologies using data from the real-world quickly and without failures (falling or self-collisions). Remarkably, as our results indicate, the proposed strategy actually performs better than a global optimization strategy.

The relevance of this study is two-fold. First, this work presents an efficient reliable controller-morphology co-adaptation strategy for robots that need to undergo morphological changes (shown by the body transformation from A to B in Fig. 1) [8]. Such cases arise for modular robots, when considerable load is added, or when new functional components are attached. Second and more importantly, this work proposes morphological adaptation as a design strategy for tuning robots to their working environment (In Fig. 1, a higher performing body B can be obtained by searching through the intermediate morphologies

A1, A2... An). We use a simple 3 degree of freedom walking robot for our study. The morphological parameters we explore are the length of the robot feet and the inertial parameters of the robot. The behavior in study is the locomotion speed. Even with a simple robot design and a common robotic task, development of dynamic controllers for the system is not straightforward. Using extensive experimental tests we show that the proposed localized controller-morphology co-adaptation strategy is highly desirable for adapting controllers for morphologically adapting robots. Our results not only indicate fast control adaptation with a low risk of failure and damage, but were also able to find better performing solutions than a global optimization algorithm like simulated annealing. A PMCCA search around the morphological space also showed that superior morphologies better suited for the task and environment can be easily found and tuned.

2 Experimental Setup

For this study we use a planar robot, comprising of four rigid medium density fibreboard (MDF) links and three rotary joints. At each end of the four link mechanism, a MDF foot of variable length f is perpendicularly joined (see Fig. 2), which can be varied between 5 cm and 25 cm in 5 cm intervals. For our study, we consider the feet length to be our morphological parameter to be tuned. Since, it required more complex mechanisms to adjust the feet length continuously, we rely on a discretized version of the same. Rubber tape on the underside of the feet controls the slip between the robot and the horizontal surface. Servo motors actuate the joints between links, each with a range of 120°. All are set at their lower bound when the four links are parallel. They are signalled by an on-board microcontroller, which is tethered to a constant voltage power supply and a PC running *MATLAB*. Six 33 g M10 bolts may optionally be arranged around the leading leg to add mass to the robot, as in Fig. 2. Hence, this robot can be parameterized in two directions in the morphological space, although quite discreetly.

The movement of the robot within the plane is monitored using a six camera *OptiTrack* motion capture system. Four markers are attached at each of the outermost joints, and the position of each cluster's geometric centre is followed. Of particular interest is the average velocity of the joints in the x direction (Fig. 2), used to prescribe the behavioral scores.

3 Procedure

Like the morphological space, the action space of the robot has to be parameterized to make the search problem tractable. Knowing that the behavior of concern is arising from locomotion, we can constrain our control actions to periodic signals. Hence, the configuration space of all possible motions was parameterized into ten dimensions by controlling each of the three servo motors with a sine wave. Though the time period of the wave was consistent between the motors,

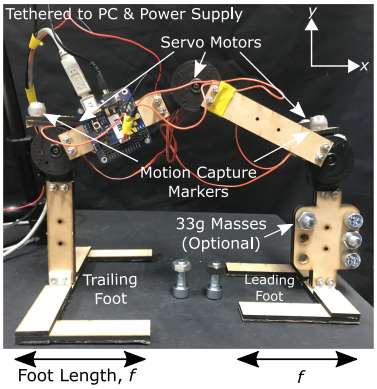

Fig. 2. Experimental Setup: the four link tethered robot. Six masses may optionally be added around the leading leg.

each of the other parameters - amplitude, phase, and mean - could be independently varied, subject to physical limits. The control objective is then to find the optimal set of control parameters that maximizes the desired behavior: locomotion speed, in our case. More specifically, the optimization process aimed to maximise the average velocity of the robot in the x direction over the final three sine periods of a five period run. This was done in order to remove effects of any transient dynamics and ensure that the locomotion was stable, at least for a short duration.

3.1 Simulated Annealing

The initial controller for the 'base' morphology can be obtained by any global optimization algorithm. We use simulated annealing for this study. The adaptation of the controller for subsequent *local* morphologies can then be performed again by simulated annealing (naive approach) or by our proposed PMCCA search algorithm that leverages the continuity in the action-morphology space.

Table 1. Simulated annealing parameters

No. of cycles	15	Starting Phases	0 rad
Iterations per cycle	15	Starting Means	$\pi/15$ rad
Starting Temp. T_0	1	σ_t	8T s
T_{i+1}/T_i	0.85	σ_M	$2T\pi/3$ rad
Starting Time Period	5 s	σ_A	$2T\pi/3$ rad
Starting Amplitudes	$\pi/15$ rad	σ_P	$4T\pi/3$ rad

For comparison purposes, each of the five foot sizes independently underwent a 225 iteration simulated annealing process, the parameters of which are given in

Table 1. During every iteration, each new parameter was proposed by sampling from a one dimensional Gaussian curve centred at the current parameter and truncated at its physical limits. The time period was limited between 1 & 15 s. Standard deviation values σ_t (time period), σ_M (mean), σ_A (amplitude) & σ_P (phase) depended on the current temperature, T_i, which was decreased by 15% at the end of each 15 step cycle. Any score, S, above the currently accepted score was immediately accepted, otherwise acceptance had a probability of $e^{\frac{-1}{T_i S}}$. After each iteration, the data was manually flagged with a descriptor of the robot's behaviour. This later allowed the determination of number/location of catastrophic (C) failures (in which the robot fell over or collided with its own body) and harmless (H) failures (in which the robot's movement had no effect on its position or caused it to move backwards) to be identified. Both scenarios were assigned a minimum score of 0.01 cm/s.

3.2 Proposed Methodology: PMCCA Search from Peak

The proposed PMCCA search algorithm for morphologically adapting robots is based on the use of priors from the previous morphology to greatly reduce the search space region. First for the 'base' morphology ($f = 15$ cm), the highest three scores (henceforth referred to as Peaks A, B and C) from the simulated annealing process were identified and a refining search was performed around each. This search comprised of 30 additional iterations using parameters drawn randomly from a space near the peak, each limited to one-fifth of its range during the simulated annealing. Note that the time period was permitted to extend below 1 s since the robot was deemed less likely to collide with itself; a high speed collision could have easily broken joints or prompted a dangerous current spike from the power supply. Note that reducing the search space to one-fifth of a single parameter cumulatively amounts to a significant reduction in the overall search volume. This PMCCA search (Algorithm 1) is then used to adapt the controller to the robot's changing morphology. After every search, the new best scores and their control parameters are used as the next starting point for the new morphology.

For example, after the refined search through the peaks A, B, and C of the 'base' morphology ($f = 15$ cm), the best parameters are selected and the morphology is modified in a piecewise manner to a setting of $f = 20$ cm or $f = 10$ cm. Before moving to the next morphology ($f = 25$ cm or $f = 5$ cm, respectively), a refined search is performed around the current peaks and the next best parameters identified. Next, f was set to 10 cm, and the same co-adaptation rule was applied for sequential addition of six 33 g masses to the leading leg. The analysis of the experimental results are presented next.

```
while not final morphology do
    S_current = 0;
    repeat 30 times
        propose parameters near the current set;
        run 5 sine periods using proposed parameters;
        if no failure then
            calculate S from OptiTrack data;
            if S > S_current then
                S_current = S;
            end
        end
    end
    accept S_current parameters;
    increment morphology;
end
```

Algorithm 1: PMCCA search algorithm, beginning from the base morphology

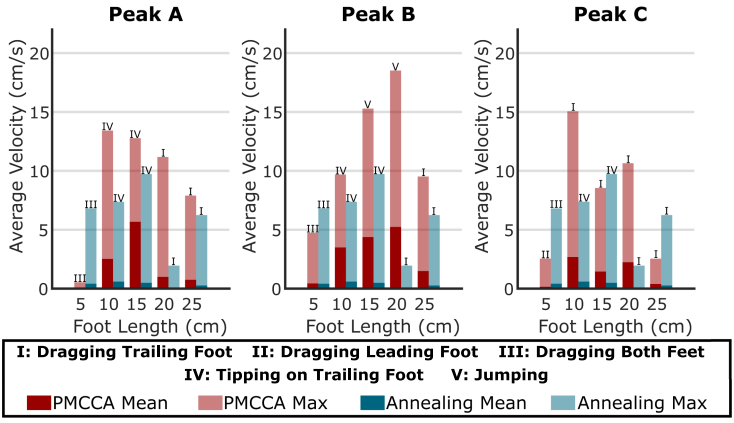

Fig. 3. PMCCA searches branching outwards from the 15 cm base morphology. The average velocities and gait pattern is shown here. (Color figure online)

4 Experimental Results

4.1 Morphological Adaptation by Varying Foot Length, f

Figure refpeaksearchbarfig shows the results of the 30-iteration PMCCA searches for each of the three starting peaks. These are compared with the exhaustive simulated annealing results, displayed in blue beside each foot length. With the exception of $f = 5$ cm, we see that the PMCCA method outperforms the global simulated annealing process for all new morphologies. We believe this is because the high performing behaviors lie on a very small region in the action space, which is unlikely to be visited by global search algorithms in practical time-limits. Given infinite searching time, both algorithms will have the same maximum peaks, while PMCCA will have a better mean value as it searches locally near the peaks. This hypothesis is valid in the coming sections.

Now, the 900-iteration process of independently optimizing all four new morphologies has been replaced by a 450-iteration process - a reduction of 50%. The

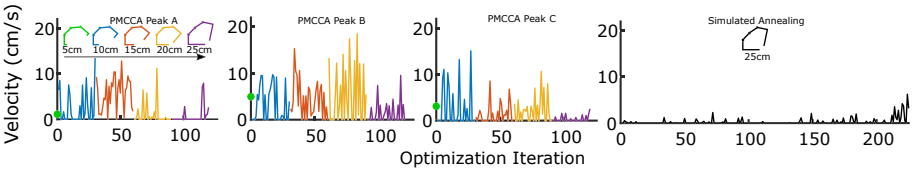

Fig. 4. Comparing a simulated annealing process starting from scratch on $f = 25\,\text{cm}$ morphology with a PMCCA search algorithm that propagates from $f = 5\,\text{cm}$ morphology to $f = 25\,\text{cm}$.

process could be sped up by reducing the number of searches and stopping the refined search once good control candidates have been found. Additionally, the average catastrophic failure rate of the entire process falls from 27.9% to 17.9%, even though it is skewed upwards by the comparatively high failure rate of $f = 25\,\text{cm}$ during the PMCCA search (see Table 2).

Table 2. Comparing the simulated annealing and the PMCCA method catastrophic failure rates

(%)	5 cm	10 cm	15 cm	20 cm	25 cm
Annealing Failures (C)	22.7	15.6	30.7	37.3	33.3
PMCCA Failures A (C)	0	6.7	0	3.3	70.0
PMCCA Failures B (C)	0	0	0	13.3	43.3
PMCCA Failures C (C)	0	0	3.3	13.3	20.0

It is clear from Table 2 that transitioning downwards in foot size is significantly less risky. Of the 180 iterations doing so, only 2 failures occurred: a failure rate of 1.1%. However, Fig. 3 seems to suggest that this direction is associated with a reduction in performance of the 5 cm ft. Though the highest $f = 5\,\text{cm}$ score found during the PMCCA search is only 69.6% of the simulated annealing maximum, the significant increase in safety may prove to justify this in applications where failure of trials carries a high risk and is highly undesirable.

The PMCCA search algorithm is very powerful because of its efficiency. An example is illustrated with scenario shown in Fig. 4, where the method of searching directly on a new morphology ($f = 25\,\text{cm}$) is compared with the method of morphology-controller co-adaptation. The PMCCA search algorithm starts with the $f = 5\,\text{cm}$ morphology and iteratively proceeds to a $f = 25\,\text{cm}$ morphology. This search noticeably led to the development of a range of gaits suited to each morphology, and had high scores averaging 292% that of the simulated annealing maxima, compared to the 204% & 183% of peaks A & C respectively. The PMCCA method not only found superior gaits quicker when compared to the global simulated annealing case, but it also provides valuable information about other morphologies along the 'way'. For instance, the best score achieved

among all the morphologies was found on the $f = 20\,\text{cm}$ morphology, through the PMCCA search. Surprisingly, the simulated annealing process performed the worst on the $f = 20\,\text{cm}$ morphology. The top scoring gaits emerging from the simulated anneal were observed to be static, with the sine waves of the two refined peaks having time periods greater than 6 s. In contrast, the higher scoring $f = 20\,\text{cm}$ gaits developed whilst incrementally extending the robot's morphology were observed to be dynamic, relying on the effect of gravity and foot elasticity to achieve faster locomotion.

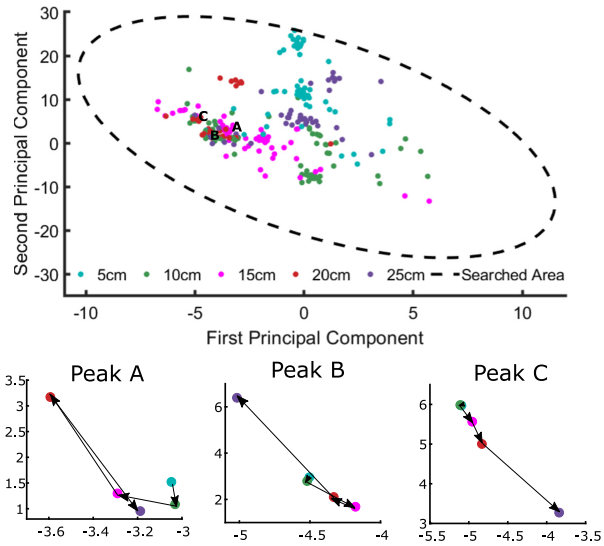

Fig. 5. Clustering of behavioral scores above 2 cm/s compared to the complete search space. Inset shows the shifting of the three peaks in a low dimensional action space representation.

The reason PMCCA performed so well can be understood from Fig. 5. Parameters producing a score greater than 2 cm/s are plotted in the reduced dimensionality control space, in which the axes represent the first two principal components of the normalized 10 dimensional control space. A dashed line marks the boundary of the valid configuration space in the reduced dimensions. On the coarsest of scales, an improvement in search efficiency is quickly deduced from the behavior distribution in the low dimensional representation of the control space. The highest scoring parameters cluster into a central region of the searched area, rarely with a first principal component greater than 5 or a second below −10. Upon further inspection this excluded region largely contains points flagged as catastrophic failures or harmless failures. Efficiency may thus have been improved by eliminating parameter ranges in the 10 dimensional configuration space corresponding to this region of the graph before optimization

began. If thoroughness of the search were not a priority, only the densely populated region to the upper left of $(1, -10)$ could be searched. How the peaks shifted with piecewise morphological changes is also quite localized.

The PMCCA search method being evaluated does not prioritise any direction of search over the others, and is equally likely to search any of the physically valid directions. However, it can be clearly observed that peaks have directional correlation. This is most apparent around peak C, suggesting a consistency in the direction of shift of a peak for a directional change in morphology. This indicates that the PMCCA search method can be further extended to include priors from preceding morphologies to also guide the search direction. Further experimentation would be necessary to confirm this hypothesis.

Distinct clusters have emerged for each value of f in the two principal component directions, confirming that optimum control parameters are dependent on the body morphology of the robot. These clusters are often manifested in the form of different gait types, as is seen in Fig. 6. The top half shows four patterns developed during the simulated annealing process. The first and third were somewhat dynamic in nature, whereas the second and fourth were static. Conversely, the final three rows of Fig. 6 contain the gaits developed during the search from peak B. Interestingly, novel gaits, highly dynamic in nature emerged from the refined search even in the 'base' morphology.

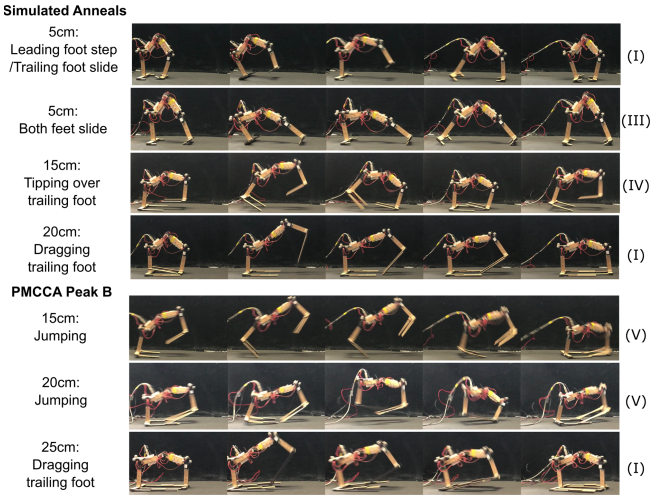

Fig. 6. A range of gaits developed whilst f was varied, from both the annealing and PMCCA searches

Gaits emerging from the PMCCA searches also developed with the morphology: for example, Fig. 6 sees a dynamic jumping gait become a slower, static, gait in which the trailing foot is dragged along the surface. Very few variations of gait pattern observed during the simulated annealing process did not subsequently

emerge during the PMCCA search, though some gaits - such as the jumping displayed - only appeared during this search. This occurred most frequently when $f = 20\,\text{cm}$, for which the simulated annealing had the least effect (Fig. 3). However, $f = 20\,\text{cm}$ produced the highest overall score during the PMCCA search - 18.5 cm/s - by developing a jumping gait from the $f = 15\,\text{cm}$ ft. Whilst this morphology would rely on momentum of the fast-swinging links to propel it into the air, the $f = 20\,\text{cm}$ setup used the elasticity of the trailing foot, which would noticeably bend with each step, to spring from the table. Even when searching a small region around the peaks, such noticeably different gaits emerge: if variations in the gait pattern were deemed important to the thoroughness of the search, extending the range and fineness of the refined search can be done.

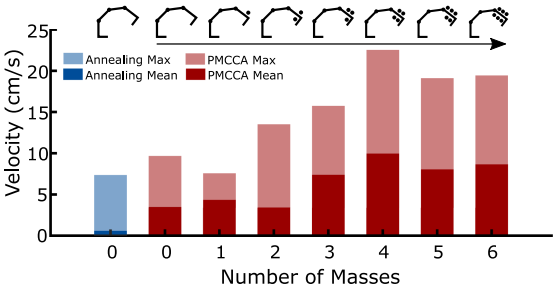

Fig. 7. Results of PMCCA search with incrementally added mass. The simulated annealing results are shown for the 'base' morphology only. (Color figure online)

4.2 Morphological Adaptation by Varying Mass Parameters

In practice, the foot length of a robot may stay consistent whilst another change in morphology occurs. For example, the addition of an arm and manipulator to the top of the robot would not change the size of the action space searched in Fig. 5, but would affect the dynamics of the robot. Such an effect was simulated using the six 33 g masses around the leading leg. The whole robot weighs around 510 g. The results of the incremental addition of masses are presented in Table 3 & Fig. 7. The blue simulated annealing bar corresponds to the optimization

Table 3. Comparing the simulated annealing and the PMCCA performance for additional weight

	Simulated annealing ($M0$)	PMCCA ($M0 \rightarrow M6$)
No. Iterations	225	180
Failures (C) (%)	15.6	6.7
Failures (H) (%)	40.0	7.8

process for the case $f = 10\,\text{cm}$. The simulated annealing process was not done for the added masses due to the higher risk of damage. Figure 7 shows us that the end of the 180-iteration search using PMCCA had produced a score not only higher than the $f = 10\,\text{cm}$ peak at which it began, but higher than any score seen without the added mass for any morphology, with almost one third of the failure rate of the full $f = 10\,\text{cm}$ optimization. Indeed, the 3 & 4 mass searches were clearly benefiting from the added mass, using the additional momentum to propel the leading foot forwards during stepping. This contrasted with the behaviour when all six masses were initially added at once: the sudden increase in weight tended to pin the leading leg to the surface, achieving very little locomotion. This shows that morphological growth can also be a desirable adaptation strategy for improving performance of the robot, tuned to the real-world conditions.

5 Conclusions

The results presented in Sect. 4 suggest that a naive global optimization, like simulated annealing, of the entire control space is rarely the safest or most efficient way for a controller to achieve high performing behavior when prior knowledge can be extracted from 'neighboring' morphologies and their sub-optimal controllers. Utilising the knowledge of a neighbouring peak along an axis of shifting body morphology enables our proposed PMCCA search to quickly attain, and often exceed the performance of the optimization in considerably fewer iterations. High performing unique gaits were also observed using the PMCCA search possibly because of its finer search process in the regions of high performance. It must however be noted that the stability of the gaits were not analyzed in this study but can also be easily be incorporated in the behavioral score.

This method of piecewise controller-morphology adaptation, in fact, parallels the process of growth seen throughout in nature; developing the controller starting from a stable morphology towards a higher performing but less stable morphology [12]. The proposed methodology is not just limited to shape and mass changes. Other morphological parameters like stiffness, damping and actuator distribution can be be similarly investigated. With advances in soft robotics, morphological adaptations are becoming more prevalent with tunable properties [2,6,13] and hence new and elegant control adaptation algorithms are required.

Although we have restricted our piecewise morphology search to a single dimension, with appropriate automated morphological adaptation mechanisms, the process can be extended to a multidimensional search space. Another interesting observation which was not investigated in this paper is the directional dependencies found in the piecewise morphological changes. This indicates that the control adaptation can be further polished from the current local search to include directional information. The definition and parameterization of the robot behavior is another aspect to be addressed if the method is to be extended to other applications. To extend this procedure to other tasks such as manipulation, appropriate behavioral scores have to be defined and estimated.

References

1. Bongard, J.: Morphological change in machines accelerates the evolution of robust behavior. Proc. Natl. Acad. Sci. **108**(4), 1234–1239 (2011)
2. Booth, J.W., et al.: OmniSkins: robotic skins that turn inanimate objects into multifunctional robots. Sci. Robot. **3**(22), eaat1853 (2018)
3. Brodbeck, L., Hauser, S., Iida, F.: Morphological evolution of physical robots through model-free phenotype development. PLoS ONE **10**(6), e0128444 (2015)
4. Doursat, R., Sayama, H., Michel, O.: A review of morphogenetic engineering. Nat. Comput. **12**(4), 517–535 (2013). https://doi.org/10.1007/s11047-013-9398-1
5. Eggenberger, P.: Evolving morphologies of simulated 3D organisms based on differential gene expression. In: Proceedings of the Fourth European Conference on Artificial Life, pp. 205–213 (1997)
6. Hawkes, E.W., Blumenschein, L.H., Greer, J.D., Okamura, A.M.: A soft robot that navigates its environment through growth. Sci. Robot. **2**(8), eaan3028 (2017)
7. Iqbal, S., Zang, X., Zhu, Y., Zhao, J.: Bifurcations and chaos in passive dynamic walking: a review. Robot. Auton. Syst. **62**(6), 889–909 (2014)
8. Jin, Y., Meng, Y.: Morphogenetic robotics: an emerging new field in developmental robotics. IEEE Trans. Syst. Man Cybern. Part C (Appl. Rev.) **41**(2), 145–160 (2010)
9. Kriegman, S., Cheney, N., Bongard, J.: How morphological development can guide evolution. Sci. Rep. **8**(1), 13934 (2018)
10. Kriegman, S., Cheney, N., Corucci, F., Bongard, J.C.: Interoceptive robustness through environment-mediated morphological development. In: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 109–116. ACM (2018)
11. Lipson, H., Sunspiral, V., Bongard, J., Cheney, N.: On the difficulty of co-optimizing morphology and control in evolved virtual creatures. In: Artificial Life Conference Proceedings 13, pp. 226–233. MIT Press (2016)
12. Malina, R.M.: Motor development during infancy and early childhood: overview and suggested directions for research. Int. J. Sport Health Sci. **2**, 50–66 (2004)
13. Manti, M., Cacucciolo, V., Cianchetti, M.: Stiffening in soft robotics: a review of the state of the art. IEEE Robot. Autom. Mag. **23**(3), 93–106 (2016)
14. Pfeifer, R., Bongard, J.C.: How the body shapes the way we think: a new view of intelligence (bradford books) (2006)
15. Rosendo, A., von Atzigen, M., Iida, F.: The trade-off between morphology and control in the co-optimized design of robots. PLoS ONE **12**(10), e0186107 (2017)
16. Saar, K.A., Giardina, F., Iida, F.: Model-free design optimization of a hopping robot and its comparison with a human designer. IEEE Robot. Autom. Lett. **3**(2), 1245–1251 (2018)
17. Sadeghi, A., Tonazzini, A., Popova, L., Mazzolai, B.: A novel growing device inspired by plant root soil penetration behaviors. PLoS ONE **9**(2), e90139 (2014)
18. Thuruthel, T.G., Falotico, E., Renda, F., Flash, T., Laschi, C.: Emergence of behavior through morphology: a case study on an octopus inspired manipulator. Bioinspiration Biomim. **14**(3), 034001 (2019)
19. Vujovic, V., Rosendo, A., Brodbeck, L., Iida, F.: Evolutionary developmental robotics: improving morphology and control of physical robots. Artificial Life **23**(2), 169–185 (2017)