

Morphological Wobbling Can Help Robots Learn

Fabien C. Y. Benureau

Okinawa Institute of Science and Technology
Okinawa, Japan
0000-0003-4083-4512

Jun Tani

Okinawa Institute of Science and Technology
Okinawa, Japan
0000-0002-9131-9206

Abstract—We propose to make the physical characteristics of a robot oscillate while it learns to improve its behavioral performance. We consider quantities such as mass, actuator strength, and size that are usually fixed in a robot, and show that when those quantities oscillate at the beginning of the learning process on a simulated 2D soft robot, the performance on a locomotion task can be significantly improved. We investigate the dynamics of the phenomenon and conclude that in our case, surprisingly, a high-frequency oscillation with a large amplitude for a large portion of the learning duration leads to the highest performance benefits. Furthermore, we show that *morphological wobbling* significantly increases exploration of the search space.

Index Terms—robots, developmental learning, embodiment

I. INTRODUCTION

Most robots used today have a fixed morphology. The length of their limbs, the strength of their actuators, the mass of their parts are decided during their design and fixed thereafter. In that respect, most robots are similar to adult humans. And since adult humans can acquire skills, can *learn*, it appears reasonable to expect that a fixed-morphology robot, given sophisticated-enough programming, should be perfectly suited to learn as well. Yet, it's difficult to fail to notice that humans acquire the most skills at a time when their morphology is *not* fixed: during their childhood. It is tempting to argue that skill acquisition is at its peak during childhood by necessity: babies, after all, have many things to learn. Morphological development happening at this time would be a coincidence driven by physiological constraints, and in fact would probably impede the pace of skill acquisition. This traditional view is opposed by another theory that has accumulated theoretical and empirical support: that morphological development crucially guides and helps skill acquisition in humans and animals (Jayaraman and Smith, 2020; Thelen et al., 1984). For robots, this is one of the issues at the heart of *developmental robotics* (Lungarella et al., 2003; Cangelosi, 2015), and the one this study explores: can a robot whose morphology changes see its learning performance improved over a fixed-morphology one?

Here we are also deliberately exploring a non-biologically plausible scenario. In particular, we are not going to make a robot's morphology grow from a baby body into an adult one. We explored this idea—in the context of evolutionary robotics—in a previous study (Benureau and Tani, 2022). Other studies have explored this as well (Naya-Varela et al., 2020a,b), and the conclusion is that growing up can make a robot learn better. Designing a robot that grows, however, is

far from trivial. Besides the many technical difficulties this represents, even choosing how it should grow is a tough problem, especially as knowledge on that issue is scarce. It may necessitate essentially designing several different robots for each of the stages (e.g., infant/kid/adult) and ways for the robot to transition from one to the other if development is gradual. All this conspires to make morphological development too hard and costly in many cases to actually use it in robots. Here, we explore a simpler alternative. Rather than growing our robots like animals, we *wobble* the morphology of a robot around its adult shape; we apply, with each passing learning epoch, a sinusoidal perturbation of the morphological values of a robot such as its size, its mass, or actuator strength. By exploring this biologically-removed idea, we can also hope to illuminate the role of morphological growth in animals. Specifically here, we are asking an implicit question: for learning purposes, does an animal morphology needs to *grow* or merely *change*?

The idea of perturbing the target task or environment is not new and one example is Jakobi's minimal simulation approach (Jakobi, 1997, 1998). The method advocates making any non-task-related part of a simulation *noisy*, with the objective to make the discovered behaviors robust enough to experimental dimensions not crucial to the task so the robot can hopefully bridge the reality gap (Koos et al., 2013). In comparison to this approach, morphological wobbling targets dimensions a priori deeply entangled with the task, such as the robot size or motor strength; it modifies the morphology *away* from the target one. It may seem, at first sight, an obstacle rather than a learning help. This is similar to the approach taken by Robust Reinforcement Learning (Morimoto and Doya, 2005; Pinto et al., 2017; Tobin et al., 2017; Peng et al., 2018), where adversarial perturbations are made to a model to make it robust to changes to inaccuracies in the model or to bridge the reality gap. Yet, morphological wobbling assumes that we have access to a *perfect* version of the robot and environment—no uncertainty or reality gap is present—and yet solely focusing training on that target environment leads to subpar performance.

Growing robots or morphologically developing robots (Naya-Varela et al., 2021) is a subject that is still relatively niche, with the majority of the studies focused on the indeterminate growth of plant-like robots (Dottore et al., 2018; Corucci et al., 2017).

A study (Lungarella and Berthouze, 2002) has specifically

looked into how a nonlinear perturbation of a 12 DOFs robot affects performance when present or absent. They concluded, however, that its presence led to lower performance. This is also the conclusion of (Bongard, 2011) when adding morphological development in between trials to a locomotion acquisition task. The solution for that study was to add morphological development during behavior (during a trial), a technique that is used in a number of other studies (Kriegman et al., 2017, 2018a,b; Bongard, 2011; Corucci, 2016). Contrary to those studies, morphological wobbling does not change the morphology during a trial or learning epoch, only in between.

In recent studies, (Naya-Varela et al., 2020a) and (Naya-Varela et al., 2020b) looked into how growth and gradual increase in range of motion affected learning in quadruped, hexapods, and octopods. In (Benureau and Tani, 2022) we looked into how growth, i.e., increase in size, mass, or muscle strength (or all three combined) can affect learning in 2D and 3D soft robots (the 2D ones are the same as this study). In all three studies, the developmental trajectory starts with a baby morphology that slowly grows, as learning progresses towards the adult morphology. All studies were able to show instances of development outperforming non-development. In this article, rather than starting with a baby morphology that one has to design, morphological wobbling starts with the adult morphology and oscillates its morphological values.

II. METHOD

A. Robot

We consider a simulated 2D soft "starfish" robot, the same as the one in Benureau and Tani (2022). The robots are composed of six tentacles attached to a central body. Every part of the robot is composed of point masses linked together by springs. The springs can adopt a wide range of stiffness, enabling them to simulate rigid and flexible links alike. Some links allow their resting length to be modified by motor commands: those are the muscle of the robot.

Each of the six tentacles is made of eight sections, divided into two motor groups of four sections, as shown in Fig. 1.A. A section has two actuated springs (muscles) on the side that act in an antagonistic manner: when one contracts by a factor $1 - \alpha$ relative to its resting length, the other extends by $1 + \alpha$, bending the section one way or the other. A section is actuated by setting the value of α equal to a fixed sinusoidal signal, parametrized by its period, phase and amplitude (this sinusoidal signal has nothing to do with the one of morphological wobbling). The period value is shared by all the sections of the robot and fixed at 2π seconds (it cannot be modified by learning). Moreover, all sections belonging to the same motor group receive the same sinusoidal signal. With two motor groups per tentacle, and six tentacles, that amounts to 12 motor groups, each needing a value for the phase (in $[-\pi, \pi]$) and amplitude (in $[0, 0.2]$). The controller is therefore fully specified with a 24-scalar vector.

A. ROBOT DESIGN

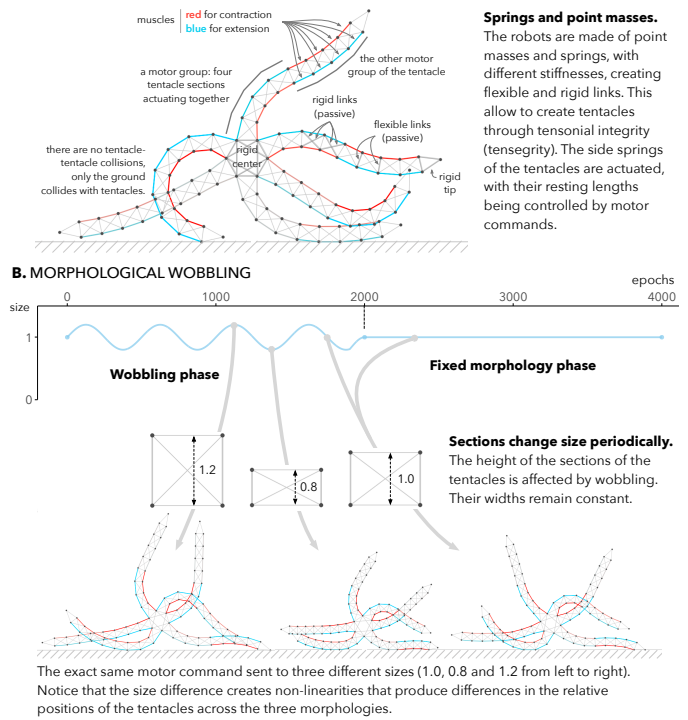


Fig. 1. **A.** Robots are made of springs and point masses. **B.** Morphological wobbling applies a sinus function to chosen morphological characteristics of the robot for the first half of training.

B. Trial & Error Learning

The task of our robot is to learn how to move. The performance is the number of body lengths (of the reference adult size, i.e. fixed, independent of wobbling) the robot is able to do over 60 seconds. The robot is created above ground and dropped at the start of the simulation. To avoid having the robot take advantage of the kinetic energy of the drop to bounce, which experimentally results in chaotic performance, it settles on the ground without actuating for 9.42 seconds ($\approx 3\pi$, 1.5 periods), then actuates for the remaining 50.58 seconds.

We use a simple learning algorithm: trial & error. At each epoch, the robot tries 20 different behaviors, i.e., 20 24-scalar vectors. The five best are kept, and each produces three random perturbations of itself, generating 15 new behaviors that will, with the five kept behaviors, form the 20 behaviors of the next epoch. The robot is trained in this manner for 4000 epochs.

A perturbation is created by selecting two random values in the 24-scalar vector and—assuming those values are normalized in the interval $[0, 1]$ (from the interval $[-\pi, \pi]$ for a phase value or $[0, 0.2]$ for an amplitude value)—adding a random normal perturbation of variance 0.05 to each of those values. If the value exceeds the bounds, the excess is mirrored back into the range.

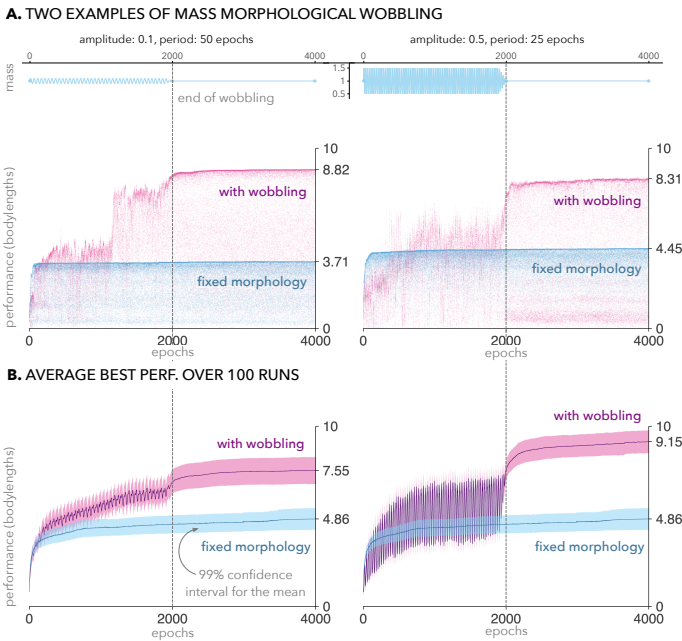


Fig. 2. Morphological wobbling improves performance. **A.** Scatter plots of the trials of a run for the fixed-morphology and morphological wobbling conditions. Each dot is a trial. Furthermore, two oscillation conditions are shown and correspond to the left and right column, and extend to **B.**, which presents the average of the best trial of each epoch, over 100 runs. The shaded area is the 99% confidence interval.

C. Morphological Wobbling

Our morphological wobbling consists in changing the morphological characteristics around the values of the adult morphology (see Fig. 1.B). The morphological characteristics we consider are the mass of the nodes of the robot, the stiffness of the muscle springs, and the size (the height) of the tentacle sections. Typically, we apply a sinusoid perturbation of one of those values during the first 1900 epochs, with a given amplitude and perturbation. From epoch 1900 to 2000, the amplitude, whatever its value, is linearly decreased to zero. From epoch 2000 to 4000, the robot’s morphology is fixed at the adult morphology values. For simplicity, we normalize all values so that the adult morphology values of mass, muscle stiffness, and tentacle section size is 1.0.

It is important to note that the morphology of the robot does not change during a learning trial or even during a whole epoch. It remains fixed and is same for all trials. The morphology changes in-between epochs, during the first 2000 epochs. Another point is that our wobbling trajectory is not a function of learning performance, and does not change, for instance, when a performance threshold is reached.

III. RESULTS

Fig. 2 shows the impact of mass morphological wobbling for two conditions over the fixed-morphology condition. In Fig. 2.A, two selected runs are shown. We can observe that the performance during wobbling has a high variance that correlates with the wobbling oscillation. That is not surprising.

Given similar behaviors, increasing or decreasing the mass of a robot will usually increase or decrease his locomotion performance. During wobbling, the performance oscillates at levels of performance inferior, similar, or superior to the fixed-morphology performance, with sudden changes to the mean. After wobbling ends, however, the performance of the wobbled robot settles to a value much higher than the fixed morphology.

Those results are validated when we look at the average performance of the best trial of each epoch across 100 runs in Fig. 2.B. We can see that the performance indeed oscillates precisely with the wobbling (else, it would get smoothed out by the average). And the final performance is significantly higher than the fixed morphology one, up to twice as much for an amplitude of 0.5 and a 25-epoch period.

Those last parameters represent quite a high amplitude and a fast period. It is perhaps the most unexpected result of this study: we assumed that a moderate period and amplitude would reap the benefit of wobbling while giving the time to the learning algorithm to adapt. Yet—and this is the point of the next two sections—performance is generally increased the higher the amplitude and the shorter the period.

A. Period and Wobbling Pauses

To apply morphological wobbling, we must choose a few variables for the oscillations: the period, the amplitude, and how long they last. Let’s start with the period, i.e., how fast the morphology wobbles. As stated above, we could expect that a slow wobbling, perhaps even interlaced with some fixed-morphology “rest” periods, would be best for learning. The results of Fig. 3.A show it is not the case for our robot. We consider three conditions across different period values. The first condition is the regular wobbling around the 1.0 mass value, with an amplitude of 0.1. The second is the same wobbling, with an upper bound at a value of 1.0; this creates a fixed value at 1.0 for half of the period. Finally, the third condition wobbles around a mean of 0.9, with an amplitude of 0.95 (the fixed value during the second half of learning is still 1.0). This has the same range as condition 2 without any fixed-morphology portions. What we observe is that the period offering the most performance is 25 epochs, rather than the slower period we expected. A shorter period than 25 does lead to some modest loss of performance, suggesting that the relative proximity of successive morphologies along the oscillation range is of some importance. Stopping to wobble periodically as in condition 2 does not offer any advantage over condition 3 and has a consistently inferior mean performance over condition 1. Comparing conditions 1 and 3 further suggests that having the mean of the oscillation be the value of the target morphology might be important, but this conclusion is muddled by the difference in amplitude, especially in light of the results of the next section.

B. Amplitude and Morphological Characteristics

The other crucial parameter of wobbling is its amplitude, i.e.: how much to wobble? Fig. 3.B shows both the effect of amplitude and of wobbling different morphological

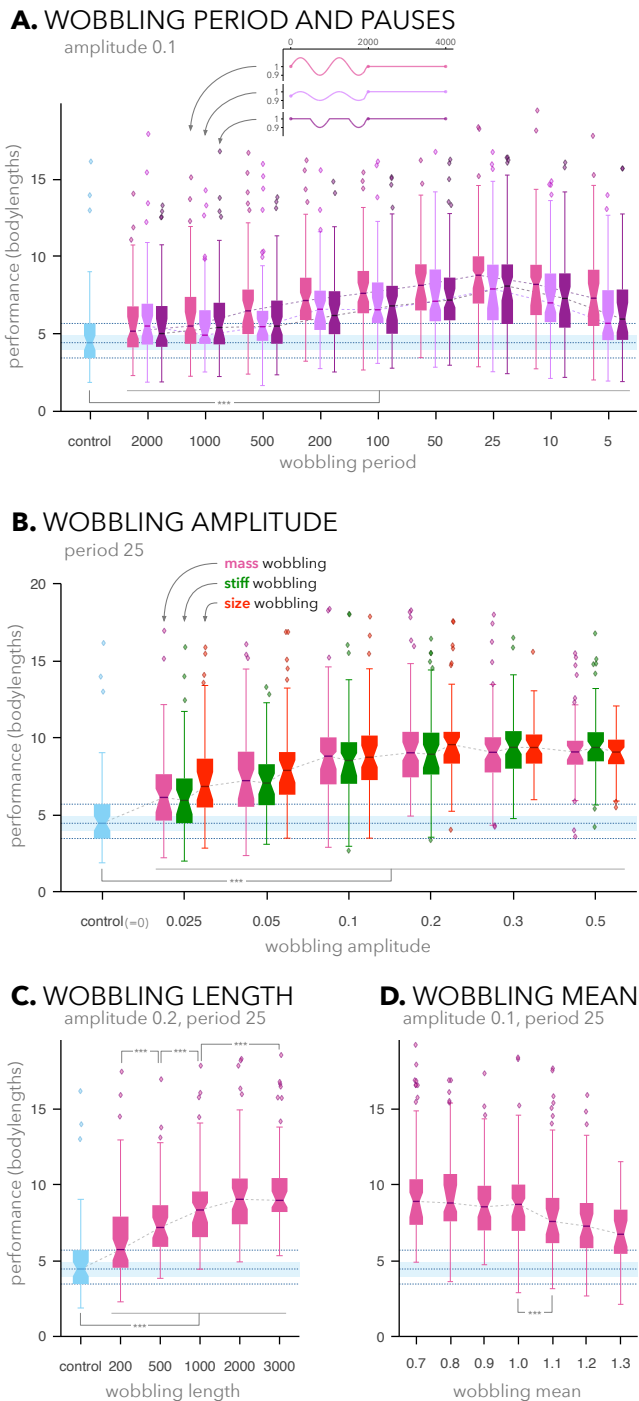


Fig. 3. Robots should wobble for a long time (C.), at high frequency (A.) high amplitude, regardless of the morphological dimension of wobbling (B.), and at a wobbling mean inferior or equal to the target morphology when wobbling mass (D.). Each box-and-whisker represents 100 runs. All boxplots use the same 100 starting populations (100 same random seeds), creating paired experiments. This allows to compute significance stars using a Wilcoxon signed-rank test (paired performances regularly deviate significantly from normality under Shapiro-Wilk testing). The significance threshold is 0.01. Boxplots show the first and third quantile, with notches for the 99% confidence interval; whiskers show the minimum and maximum values, clipped to $1.5 \times \text{IQR}$. Beyond that, outliers are shown as diamonds.

characteristics on performance for a period of 25 epochs. The performance increases as the amplitude of wobbling increases, until 0.3. After, further increases in amplitude do not increase performance significantly. This is again unexpected. The higher the amplitude, the further away the robot is from the target, adult morphology. Behaviors that are effective on a morphology are rarely effective on a highly different one, with robotic search spaces often dense with non-linearity and discontinuities. This could easily lead to a situation where a learning algorithm cannot keep up with the fast-oscillating environmental changes caused by a changing morphology, and regularly discard good solutions because half of the learning time is spent on morphology significantly different from the target one. Yet, we do not observe this here: a high amplitude of oscillation is beneficial, and though we observe diminishing returns, we do not observe any adverse effects of high amplitudes.

Interestingly, even for qualitatively different forms of morphological wobbling—node mass, muscle stiffness, or size—the performance is impacted the same. This non-specificity of the wobbling is surprising. We would expect that some morphological characteristics among the ones we considered would be coupled more tightly to the learning performance than others.

C. Wobbling Mean

In Fig. 3.D, we quickly explore the effect of different values of the mean of the wobbling, for mass wobbling. In accordance with the results of [Benureau and Tani \(2022\)](#), a heavier average mass during wobbling leads to lower performance, yet still significantly higher than fixed-morphology learning. A lighter mass during wobbling does not significantly increase performance.

D. Wobbling Length

A final major question is: how long should the wobbling last? Fig. 3.C tells us that even a short wobbling phase of 200 epochs brings a significant learning performance benefit. And the longer the wobbling phase is, the higher the performance benefit. This is not a trivial result. Where is that dynamic coming from? To answer this, we need to start understanding why morphological wobbling works.

E. Morphological Wobbling Fosters Exploration

Given a task, a robot learns by exploring a behavioral space and finding behaviors that solve the task as best as possible. One danger of learning is being trapped in a local extremum of the behavioral space and spending time improving a specific solution when better ones exist in other areas of the space. Such premature convergence is an inevitable risk of any empirical learning happening in a behavioral space that cannot be exhaustively explored. Many algorithmic techniques exist in machine learning to minimize that risk.

Among other things, morphological wobbling can be viewed as such a technique, specifically an *embodied* one. Rather than finding efficient ways to explore a behavioral space,

morphological wobbling *modifies the behavioral space*. It is difficult to stay trapped in a valley of the behavioral space when the valley disappears under you.

One way to see the relationship between a learning algorithm and a dynamic environment is to see the learning algorithm as having to constantly play catch up with the environmental changes. Another is to consider that the environment pushes the learning process around, especially out of stable attractors.

Fig. 5 gives us an indication that this is indeed what is happening here. We initialized a robot with a single behavior. The robot learns to modify this behavior to improve its locomotion. In one condition, the robot has a fixed morphology, and in the other, the robot undergoes morphological wobbling. Since at each epoch the robot tries new behaviors by modifying the existing ones, we can measure the number of modifications that we explored from the initial behavior to the best behavior of the last epoch. We obtain a *search distance* as the sum of the euclidean norm of all these modifications. For a given behavior b , we define $predecessor(b)$ as the behavior b' in the previous epoch that b was created from, either through replication (the behavior is one of the five that was kept from the previous epoch) or random perturbation. We can then build the sequence of behaviors that led to behavior b_q , at epoch q , from epoch p ($p < q$): b_p, b_{p+1}, \dots, b_q with $b_i = predecessor(b_{i+1})$ for $p \leq i \leq q - 1$.

Given two epochs p and q , we define the search distance between epochs p and q as the sum of the euclidean norm between the sequence of behaviors above, starting at the best performer of epoch q :

$$\sum_{i=p}^{q-1} |b_i - b_{i+1}|$$

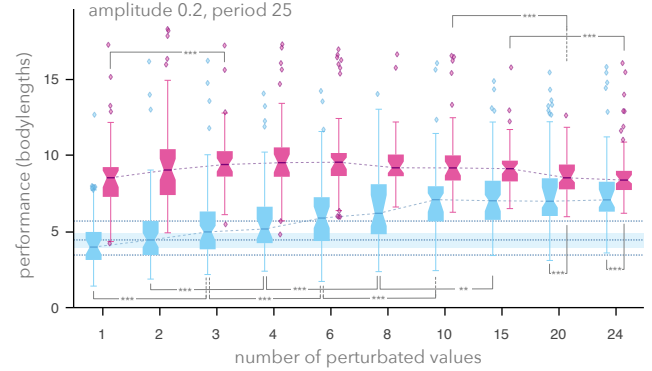
In the run of Fig. 5.A, the fixed morphology covered a distance of 6.79 between epochs 1 and 2000 (6.66 ± 0.84 (99% CI) over 100 runs); the wobbling robot covered a distance of 60.7 (58.23 ± 1.49 over 100 runs) in the same timeframe. As soon as the wobbling stops, the search distance drops sharply, to cover a distance of 1.56 (3.33 ± 0.65) for the wobbling morphology between epoch 2000 and 4000 versus 0.74 (1.81 ± 0.45) for the fixed one. Wobbling drives exploration, and we confirm this by computing the average distance between the best behaviors of each two consecutive epochs in Fig. 5.B.

An informal observation that can be made from Fig. 5.A is that the interaction between a morphology that dislodges learning from the local extremum and learning that gravitates toward better behaviors does not seem to lead to a stationary dynamic in the search space. The search does not seem to visit again and again the same place of the search space, as could be expected of a repetitive, sinusoidal, perturbation: the search *moves around* in the behavioral space, as shown through the PCA.

F. Comparison with Perturbation Exploration

Morphological wobbling here seems to generate good, useful exploration. But wouldn't any other methods that increase exploration fare as well? One straightforward way to achieve that is to change how the random perturbations of the controllers are generated. So far, we only changed two values in the controller 24-length vector, with a normal perturbation of variance 0.05. In Fig. 4, we explore what happens for other values of those two parameters, for fixed morphology and morphological wobbling conditions.

A. NUMBER OF PERTURBATIONS



B. PERTURBATION VARIANCE

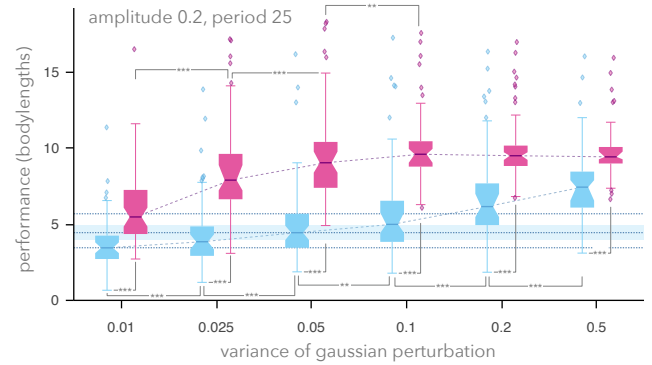
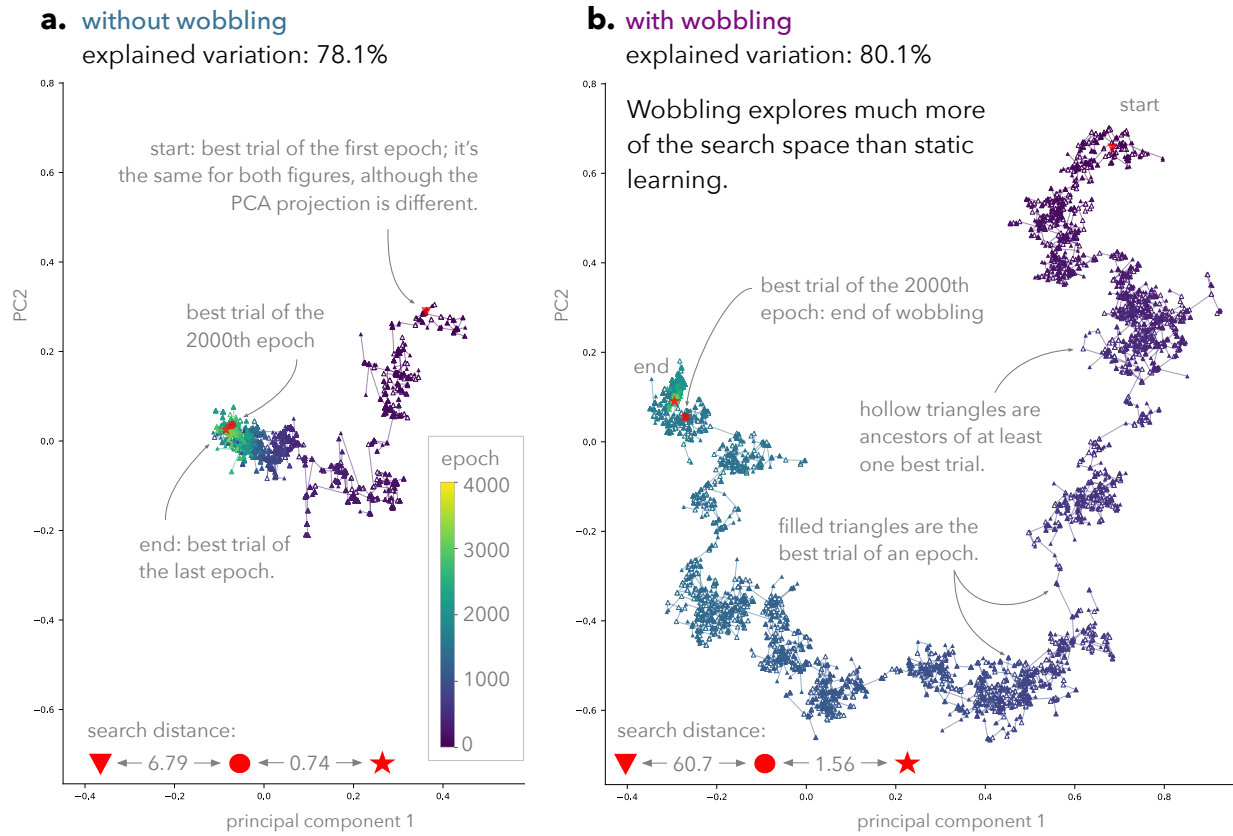


Fig. 4. Increasing perturbations does not improve performance as much as development. **A.** Performance for different values of the number of perturbations per epoch (variance 0.05), for the fixed morphology (blue) and morphological wobbling (pink, amplitude 0.2, period 25). The same conventions as Fig. 3 are used for boxplots. **B.** Same as **A.**, but varying the variance of the perturbations (two perturbations per epoch).

While an increased number of perturbations or increased perturbation variance does improve the performance of the fixed morphology, in all cases, the performance of corresponding development remains significantly higher ($p < 0.001$). The base development condition (amplitude 0.2, period 25, two perturbations of variance 0.05) also has a significantly higher performance than *all* the fixed morphology conditions across Fig. 4.A and B ($p < 0.001$). Interestingly the performance of development dips when the number of perturbations per epoch is high. Further analysis is needed to understand why.

If we consider morphological wobbling as producing noise during the learning process, then one thing to remark is that

A. PRINCIPAL COMPONENT ANALYSIS OF THE LEARNING SEARCH



B. MOVING AVERAGE OF THE PERTURBATION DISTANCE OVER 100 RUNS




Fig. 5. Wobbling makes the learning process explore the search space much more. **A.** Starting from the same initial behavior, we perform a 2D principal component analysis of the search through the behavioral space. Each behavior except the initial one is derived from another and we can therefore draw a graph of the behavior. Here we represent only the best behavior of each epoch and its ancestors. The starting point is a large inverted red triangle, the halfway point (epoch 2000) is a red disk, and the endpoint is a red star. The search distance, as the sum of the euclidean distance between each consecutive behavior (encoded as a 24-scalar vector), is given between those points for the fixed-morphology and morphological wobbling conditions. The two PCAs are made independently for each condition (the projection axes are different). **B.** We confirm that the exploration dynamics are shared over the 100 runs. We show the moving average of the *epoch perturbation distance* (the euclidean distance from the best behavior genotype vector from one epoch to the next) over 100 runs and a 101-epoch smoothing window.

this noise is intrinsically specific to the task and the robot. The interaction of the morphology with the world creates the behavioral landscape, and thus changing the morphology creates a task-specific, environment-computed noise signal. Increasing perturbations, on the other hand, creates a task-agnostic source of noise, divorced from the environment and task. This hints at a reason for the difference in performance observed here. Yet, a comprehensive explanation of why morphological wobbling generates useful exploration pattern remains, at this point, elusive.

IV. DISCUSSION

There are many limitations to the current study. We have only one type of robot, the robots are in 2D, in simulation, the task is locomotion on a perfectly flat surface, there are no sensors, no neural networks. They all warrant further study and expanding the work to figure out how much, if at all, this result generalizes to different robots, tasks, and environments.

There's one aspect that warrants an extended discussion: the fixed morphology converges *quickly*, as is evidenced in Fig. 2.A. This fast convergence is probably a reason why a short period and therefore fast oscillations produce good results in this study. It could lead to the conclusion that morphological wobbling is effective here because it prevents the premature convergence of a crude trial-and-error learning algorithm; in contrast, more sophisticated learning algorithms, not susceptible to such a premature convergence, would not get any benefit from morphological wobbling.

It's a valid point. Except that premature convergence can befall even sophisticated learning algorithms, and that the additional help of morphological wobbling, although not useful in many situations and tasks, might be crucial in some others. But a better justification here is to argue that this quick convergence of the task may actually capture some of the dynamics at play in the biological world.

Indeed, there is evidence that a simple trial and error algorithm may suffice to acquire new sensorimotor skills in biology, simply because the search space has a lot of good solutions, and we may be able to converge on one of them quickly, wherever our starting point is (Raphael et al., 2010; Loeb, 2012). Quickly finding a good-enough solution to a problem is useful for survival in the animal kingdom. Yet it could prevent finding a better solution that might represent a negative selection pressure over time, especially if other members of the population discover it. Here, morphological change, i.e. physical growth, may help improve over time a behavior that converged quickly. We would have both advantages: a fast convergence toward a useful behavior in early childhood and a high-performing behavior by the time the animal reaches adulthood.

In this work, we chose to apply a sinusoidal oscillation to the morphology value. The main reason behind that choice among other possible oscillation types was its simplicity. There's an interesting parallel to make between our work and the field of adaptive control, where convergence can be guaranteed if a *persistently excitatory* signal is present, with

sinus being often used as such a signal (Boyd and Sastry, 1986; Lee et al., 2015). Exploring in future studies which theoretical tools and insights from that field could be adapted and applied to our case might be a fruitful avenue of research.

Another point to discuss is morphological wobbling practical implementation difficulties. Designing a robot that changes its morphology is complicated and runs contrary to many of the other engineering constraints of robot design. Furthermore, designing developmental paths for growing robots from baby to adults, as we have done in Benureau and Tani (2022), is far from trivial, and may sometime present a difficulty equal or superior to the design of the robot itself. Morphological wobbling removes some of those difficulties, since the morphology deviates around the adult values in a straightforwardly controlled manner. Many of our experiments have dealt with changing the mass of our robots. In simulation, this can be achieved directly. Even more simply, a similar effect can be achieved by wobbling the gravity of the simulation, which amounts to changing the value of a single vector, and removes the need to recompute inertia matrices. With real robots, changing gravity is not possible and modifying the mass—or the size—can be challenging or impossible. A simple intervention, however, is to change the motor strength by modifying the max torque for instance or by using specialized actuators (Vu et al., 2013). In all those cases, setting the amplitude appropriately provides a direct way to constrain the wobbling within the physical limits of the robot. Thus, with only a handful of parameters to consider—period, amplitude, duration—wobbling the motors' torque is a straightforward way to implement morphological wobbling in a real robot.

CONCLUSION

We presented morphological wobbling and showed that deliberately applying a sinusoidal perturbation to the morphology of a robot can increase its learning performance. We showed that, surprisingly, in our case study, a high frequency and high amplitude perturbation applied for a high number of epochs provided the best results. We analyzed how exploration was affected by morphological wobbling, showing that the distance covered in the search space increased significantly during the oscillations.

A lot of work remains to investigate how successful morphological wobbling can be for different learning robots and across environments and tasks. Whether it proves to be widely applicable or only applicable to a handful of cases, such as our 2D robots, even a failure would be interesting as it illuminates the relationship between development and learning. If morphological wobbling improves the performance of some robots but not others, understanding why may yield precious insights into how development affects us.

ACKNOWLEDGMENT

We are grateful for the help and computational resources provided by the Scientific Computing and Data Analysis section of Research Support Division at OIST.

SUPPLEMENTARY DATA AND SOURCE CODE

The supplementary data containing necessary simulation details, along with the source code and full simulation records to reproduce the results, regenerate the plots, and recompute the numerical figures (distances, confidence interval, significance, explained variation of PCA) is available at: <https://doi.org/10.5281/zenodo.6513360>.

REFERENCES

- Benureau, F. C. Y. and Tani, J. (2022). Morphological development at the evolutionary timescale: Robotic developmental evolution. *Artificial Life*, pages 1–19.
- Bongard, J. (2011). Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences*, 108(4):1234–1239.
- Boyd, S. and Sastry, S. S. (1986). Necessary and sufficient conditions for parameter convergence in adaptive control. *Automatica*, 22(6):629–639.
- Cangelosi, A. (2015). *Developmental robotics : from babies to robots*. The MIT Press, Cambridge, Massachusetts.
- Corucci, F. (2016). Evolutionary developmental soft robotics: Towards adaptive and intelligent soft machines following nature’s approach to design. In *Soft Robotics: Trends, Applications and Challenges*, pages 111–116. Springer International Publishing.
- Corucci, F., Cheney, N., Kriegman, S., Bongard, J., and Laschi, C. (2017). Evolutionary developmental soft robotics as a framework to study intelligence and adaptive behavior in animals and plants. *Frontiers in Robotics and AI*, 4.
- Dottore, E. D., Sadeghi, A., Mondini, A., Mattoli, V., and Mazzolai, B. (2018). Toward growing robots: A historical evolution from cellular to plant-inspired robotics. *Frontiers in Robotics and AI*, 5.
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive Behavior*, 6(2):325–368.
- Jakobi, N. (1998). *Minimal simulations for evolutionary robotics*. PhD thesis, University of Sussex.
- Jayaraman, S. and Smith, L. B. (2020). The infant’s visual world. In *The Cambridge Handbook of Infant Development*, pages 549–576. Cambridge University Press.
- Koos, S., Mouret, J.-B., and Doncieux, S. (2013). The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation*, 17(1):122–145.
- Kriegman, S., Cheney, N., and Bongard, J. (2018a). How morphological development can guide evolution. *Scientific Reports*, 8(1).
- Kriegman, S., Cheney, N., Corucci, F., and Bongard, J. C. (2017). A minimal developmental model can increase evolvability in soft robots. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM.
- Kriegman, S., Cheney, N., Corucci, F., and Bongard, J. C. (2018b). Interoceptive robustness through environment-mediated morphological development. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM.
- Lee, T.-C., Tan, Y., and Netic, D. (2015). Stability and persistent excitation in signal sets. *IEEE Transactions on Automatic Control*, 60(5):1188–1203.
- Loeb, G. E. (2012). Optimal isn’t good enough. *Biol Cybern*, 106(11-12):757–765.
- Lungarella, M. and Berthouze, L. (2002). On the interplay between morphological, neural, and environmental dynamics: A robotic case study. *Adaptive Behavior*, 10(3):223–241.
- Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.
- Morimoto, J. and Doya, K. (2005). Robust reinforcement learning. *Neural Computation*, 17(2):335–359.
- Naya-Varela, M., Faina, A., and Duro, R. J. (2020a). An experiment in morphological development for learning ann based controllers. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE.
- Naya-Varela, M., Faina, A., and Duro, R. J. (2020b). Some experiments on the influence of problem hardness in morphological development based learning of neural controllers. In *Lecture Notes in Computer Science*, pages 362–373. Springer International Publishing.
- Naya-Varela, M., Faina, A., and Duro, R. J. (2021). Morphological development in robotic learning: A survey. *IEEE Transactions on Cognitive and Developmental Systems*, 13(4):750–768.
- Peng, X. B., Andrychowicz, M., Zaremba, W., and Abbeel, P. (2018). Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. (2017). Robust adversarial reinforcement learning. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2817–2826. PMLR.
- Raphael, G., Tsianos, G. A., and Loeb, G. E. (2010). Spinal-like regulator facilitates control of a two-degree-of-freedom wrist. *Journal of Neuroscience*, 30(28):9431–9444.
- Thelen, E., Fisher, D. M., and Ridley-Johnson, R. (1984). The relationship between physical growth and a newborn reflex. *Infant Behavior and Development*, 7(4):479–493.
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.
- Vu, H. Q., Hauser, H., Leach, D., and Pfeifer, R. (2013). A variable stiffness mechanism for improving energy efficiency of a planar single-legged hopping robot. In *2013 16th International Conference on Advanced Robotics (ICAR)*. IEEE.