

Novelty-Based Evolutionary Design of Morphing Underwater Robots

Francesco Corucci
The BioRobotics Institute
Scuola Superiore Sant'Anna
Pisa, Italy
f.corucci@sssup.it

Marcello Calisti
The BioRobotics Institute
Scuola Superiore Sant'Anna
Pisa, Italy
m.calisti@sssup.it

Helmut Hauser
Department of Engineering
Mathematics
University of Bristol
United Kingdom
helmut.hauser@bristol.ac.uk

Cecilia Laschi
The BioRobotics Institute
Scuola Superiore Sant'Anna
Pisa, Italy
c.laschi@sssup.it

ABSTRACT

Recent developments in robotics demonstrated that bioinspiration and embodiment are powerful tools to achieve robust behavior in presence of little control. In this context morphological design is usually performed by humans, following a set of heuristic principles: in general this can be limiting, both from an engineering and an artificial life perspectives. In this work we thus suggest a different approach, leveraging evolutionary techniques. The case study is the one of improving the locomotion capabilities of an existing bioinspired robot. First, we explore the behavior space of the robot to discover a number of qualitatively different morphology-enabled behaviors, from whose analysis design indications are gained. The suitability of novelty search – a recent open-ended evolutionary algorithm – for this intended purpose is demonstrated. Second, we show how it is possible to condense such behaviors into a reconfigurable robot capable of online morphological adaptation (morphosis, morphing). Examples of successful morphing are demonstrated, in which changing just one morphological parameter entails a dramatic change in the behavior: this is promising for a future robot design. The approach here adopted represents a novel computed-aided, bioinspired, design paradigm, merging human and artificial creativity. This may result in interesting implications also for artificial life, having the potential to contribute in exploring underwater locomotion "as-it-could-be".

CCS Concepts

•Computing methodologies → Evolutionary robotics;
Artificial life; Genetic algorithms;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '15, July 11 - 15, 2015, Madrid, Spain

© 2015 ACM. ISBN 978-1-4503-3472-3/15/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2739480.2754686>

Keywords

Computer aided/automated design; morphogenesis; artificial life; evolutionary robotics; genetic algorithms

1. INTRODUCTION

Legged robotics is notoriously a challenging field. As in other sub-fields of robotics, the bio-inspired approach turns out to be a very effective paradigm [13]. Working at different levels of abstraction and biomimeticism, researchers produced a number of remarkably simple designs for legged robots, that are at the same time efficient and effective [8, 14, 7, 1].

Most of these robots (inspired by or contributing to the field of embodied intelligence [20]) draw from nature in order to facilitate the control: the idea is that, with a proper morphological design, the body can take over (at least partially) the role of the controller. Biological and bio-mechanical studies indeed demonstrated that the robust locomotion of most animals is, to a large extent, the result of self-stabilizing properties of their mechanical structure, see e.g. [9, 12]. As a consequence, most of the aforementioned robots feature a simple, periodic, open-loop control.

Having such a simple control, the morphological design becomes crucial to achieve an effective behavior leveraging the interaction with the environment. Nonetheless this design stage in robotics is usually performed by human experts, that carefully engineer such robots to achieve meaningful behaviors. This entails to take difficult decisions, selecting a single configuration over a whole range of possible solutions, in a design space whose dimension easily explodes also for the simplest robots. Often the path towards the final design is guided by a set of heuristic principles [19], and when experimental tests are performed to "optimize" the design for a given ecological niche, only a small portion of the design space can be explored due to practical limitations. In addition, when highly dynamic behaviors are targeted, the complex intertwining of physics phenomena ensure that the result of such a procedure will be, almost certainly, sub-optimal.

In general this human intervention can limit the possible outcomes of a bioinspired design process. This appears to be true both from an engineering perspective (i.e. in

building the best robot for a given environment or application) and from an artificial life perspective, e.g. in exploring "life-as-it-could-be" [15], producing real robots with surprising lifelike behaviors that deviate from those of biological creatures. Methods from evolutionary robotics [4] can come into play, as they demonstrated an enormous potential to produce complex, functional, lifelike morphologies in simulation (from the seminal work of Karl Sims [24] to more recent, exciting, demonstrations [6, 11, 2, 22]).

In this work our goal is not to evolve from scratch a completely new robot morphology, but rather to improve an existing one [1] by redesigning it, capitalizing elements of bioinspiration that resulted to be effective, but leveraging evolutionary algorithms to explore alternative solutions in a large design space. The emphasis is on morphology evolution in presence of a simple periodic open loop control. We aim at discovering a number of novel, self-stabilized, morphology-enabled, legged gaits with respect to the ones currently exhibited by the robot. Doing so, design indications can be gained by analyzing the evolved morphologies. The focus is thus different with respect to related works [28, 27, 3], in which the goal was mainly to evolve the control for a *fixed*, kinematically complex, robot morphology.

In order to further develop the potential use of morphology to achieve diverse behaviors in presence of a very simple control, we resort to another bioinspired concept: the one of *morphosis* or *morphing*, i.e. the possibility to control morphological adaptation online (a concept that is related to the broader one of morphological computation [10]). This behavior has been observed in biology [26] and preliminary work in this direction has also been done in robotics, mainly for the design of mechanical solutions enabling it [21, 16].

Our methodology attempts to automatically discover candidate morphologies to perform and exploit morphological adaptation online, in order to switch among diverse behaviors. Particularly, we are interested in investigating whether it is possible to trigger a macroscopic change in the behavior by modifying few (ideally one) morphological parameters. Successful examples of this kind of morphing will be presented, and the possibility to build a reconfigurable robot exploiting morphosis to switch among different behaviors is envisioned (e.g. switching from running to swimming to overcome an obstacle). Also, with respect to previous works following similar approaches [25] (focusing on the evolution of quasi-static gaits for terrestrial robots), the setting presented in this paper differs for several reasons. First, robots will evolve in water environment, in which, with respect to the terrestrial one, the importance of morphology is further stressed. As it will be shown, a slight morphological change has the potential to entail dramatic change in the resulting behavior. Second, the reference robot has the potential to perform dynamic multi-modal locomotion (e.g. crawling, running, hopping, swimming), an uncommon feature in legged robotic platforms. Third, we differ in the adopted methodology.

We propose the adoption of genetic algorithms (GA) in combination with a recently introduced evolutionary algorithm called *novelty search* [18] as a suitable tool to explore the behavior space of a robot. Its ability to produce a large and diverse set of solutions is highly desirable here, as it can potentially mitigate the transfer problem (offering more chances of success) when trying to translate some of the insights gained in this work in a future robotic prototype.

To the best of our knowledge, this is the first time novelty search is applied for such a purpose. A method to analyze the results of the evolutionary run is also proposed, that allows the human designer to find candidate morphologies for morphing that are close in morphology space but far apart in behavior space. Eventually this procedure will allow us to design systematically robots that are geared for morphosis, exploiting embodiment to switch among different behaviors, thus being more adaptive and robust.

We propose this design flow as a general approach to mix bioinspired embodied robot design and evolutionary computation, i.e. human and machine skills and creativity. This methodology could potentially bring contributions to the more general field of artificial life, evolving novel artificial life forms that exploit some elements of bioinspiration (provided by humans) and suggesting more that are produced by the evolutionary process, possibly deviating from the starting biological model. As for bioinspired robot design, an advantage of this approach lies in the fact that the system can be improved by the evolutionary process while being exposed (in simulation) to the actual task and environment. This can potentially bring into the final design elements that were overlooked in the bioinspiration process or that (although not present in the reference biological model) could entail benefits in terms of performance or efficiency.

2. METHODS AND MATERIALS

2.1 Reference robot

Our case study will be the PoseiDRONE robot [1] (Fig. 1a), a multi-purpose, soft, underwater robot inspired by the octopus. The robot, composed for most part of soft materials, is made of three main units: a central body (meant for jet-propelled swimming, not relevant to this study), a legged module and a floating module. The legged module comprises four soft silicone limbs, each of which is attached to the rigid part of the robot via a thin, centrally embedded, flexible beam of steel. Each leg is actuated in feed-forward by a rotational motor via a three-bars mechanism, functionally inspired by the push-pull locomotion strategy of the octopus [7]. The mechanism translates a circular rotation of a crank into a pseudo-elliptical loop in the distal part of the leg (Fig. 1b), suitable for locomotion. The floating module can be inflated and deflated to change the robot density and to shift its center of buoyancy. This aspect is peculiar as it can be exploited to change the stance of the robot (e.g. to lift a pair of legs to achieve bipedal locomotion [5]) or to initiate dynamic phenomena (arising from the momentum between buoyancy acting on the center of buoyancy – CoB, and gravity acting at the center of mass – CoM) that can be exploited for locomotion. The robot demonstrates an effective, self-stabilized locomotion in water tank experiments, exhibiting two kinds of legged gaits, i.e. crawling and bipedal walking [1, 5] as configurations were varied. Particularly, mainly for practical reasons, only two morphological parameters were varied during experiments performed with the prototype, i.e. the inclination of the floating module with respect to the central body of the robot, and its inflation state, affecting the resulting robot density. By varying just these two parameters some qualitatively different behaviors were observed: however, by exploring the full morphological space, our work demonstrates that this robot could indeed exhibit

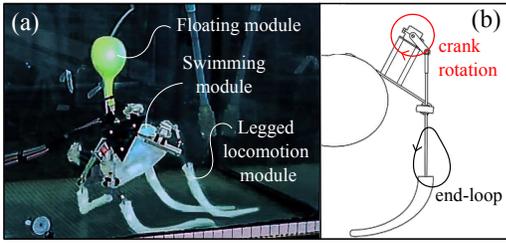


Figure 1: (a) The PoseiDRONE robot in a water tank, (b) detail of the leg mechanism.

a much wider range of morphology-enabled, multi-modal, dynamic behaviors.

2.2 Simulated model

Starting from the robot described in Sect. 2.1 a general model (Fig. 2) was developed and exploited to simulate alternative designs. An exhaustive description of the model is out of the scope of this paper (the interested reader can refer to [5]): in what follows only the most salient features will be reported. Given that the locomotion of the robot is mainly planar, the mathematical model describing its dynamics is planar too, thus featuring three degrees of freedom: the position of the CoM (x and y coordinates) and the pitch angle ϑ . Each leg is actuated by imposing a rotation of the corresponding crank around the center of rotation (CoR), at a given, constant, angular speed $\dot{\vartheta}_r$. The corresponding frequency is thus $f = \dot{\vartheta}_r/2\pi$, and it can be different for front and hind legs (f_{front} , f_{rear}). Moreover, an inter-leg phase shift $\phi_{interleg}$ can exist between the two legs of each pair (front and hind pairs), as well as a front to rear phase shift $\phi_{frontRear}$. From a dynamic point of view the soft legs are described as mass-less spring-damper systems. Each leg exerts to the ground elastic and damping forces proportional to the stiffness k and damping coefficient dr . In addition to gravity, the model describes specific dynamics of the underwater setting, such as buoyancy forces, fluid drag, and fluid dynamics effect induced by sculling movements of the legs [23]. The central body is approximated with a sphere of radius R : this simplifies the computation of relevant quantities such as the isotropic drag, body inertia and added mass. Finally, a second order foot slipping model is also present. The model embeds a considerable number of parameters (geometric, dynamic, related to materials, etc.) affecting the resulting behavior of the robot, whose complex interrelation is not easy to foresee. The ones that have been varied to explore the behavior space of the robot are listed in Tab. 1. The model was implemented and numerically solved in MATLAB[®].

2.3 Optimization setup

The model described in Sect. 2.2 is used to explore the behavior space of the robot in simulation, by adopting genetic algorithms (GAs). To conduct such an exploration *novelty search* is suggested and adopted, operating on top of a standard genetic algorithm. To demonstrate the advantages of this approach for the present study, a conventional objective-based GA is first used to achieve a comparison baseline. In the next sections we will describe the two optimization approaches in more detail.

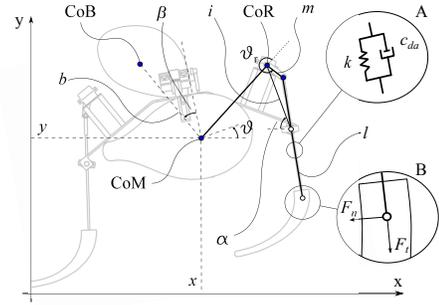


Figure 2: Schematic view of the robot showing the kinematic structure of the leg mechanism as well as some of the geometric and dynamic parameters. Legs are depicted in phase for simplicity. Detail A highlights parameters governing contact forces, while detail B highlights sculling forces that can be exploited for swimming.

2.3.1 Encoding

The structure of the artificial genome is the same for both settings (objective-based and novelty-based), embedding the geometric, actuation and materials parameters described in Tab. 1. Direct encoding is adopted, the resulting genome length being 24. Some parameters are expressed as multipliers of other quantities in such a way to embed necessary constraints into the encoding and impose only boundary constraints.

2.3.2 Experimental details

The model was allowed to run for $t_e = 25s$ for each genome. All the robots are dropped from a constant distance to the floor. If a robot falls (i.e. touches the ground with parts of the body other than feet) the event is detected and the simulation stops. Both the objective-based and the novelty-based settings rely on a standard GA implementation from MATLAB's Global Optimization Toolbox (real-coded linear constrained solver with bound constraints). Relevant GA parameters are listed in Tab. 2. All the processing was performed in MATLAB 2013b[®]. Each execution of the GA ran in parallel over 12 cores of a Workstation (Intel Xeon 3.07 GHz, 48GB RAM), for a total execution time of ~ 48 h.

2.3.3 Objective-based GA

The objective-based setting serves as a baseline to highlight the benefits of novelty search in exploring the behavior space of a robot. The adopted fitness is based on the space traveled by each robot, normalized by a characteristic body length (BL). The problem was formulated as a minimum, and the fitness was thus defined as: $fitness = -|s_x/BL|$, where s_x is the abscissa of the CoM at the end of the simulation, and $BL = l_{front} + l_{rear}$. Falling robots receive a fixed positive penalty of 100, plus an additional penalty (the longer the time a robot managed to keep the balance, the lower the additional penalty): $P_{fall} = \lceil 100 + 5 \cdot (t_e - t_{fall}) \rceil$, where t_e is the total execution time and t_{fall} is the falling time.

2.3.4 Novelty-based GA

Novelty search [18] is a recently introduced open-ended approach to evolutionary computation. Instead of maximiz-

Table 1: Model parameters composing the genome and corresponding bounds. The indication "x2" means that two parameters of the described type are present, one for hind legs, one for front legs.

Parameter	Description	Bounds
Geometry		
R	Radius of the central spherical body	$[0.05, 0.30] m$
β	Inclination of the CoB with respect to the medial plane	$[-90, +90]^\circ$
bR_{mul}	Multiplier determining the distance b between CoB and CoM ($b = bR_{mul} \cdot R$)	$[0, 3]$
$lMul$	x2 - Multiplier determining the length of the legs l $l = lMul \cdot (m + i)$	$[1, 2]$
$mMulI$	x2 - Multiplier determining the length of the cranks $m = \min(mMulI \cdot i, mMulR \cdot R)$	$[0.01, 0.95]$
$mMulR$	x2 - Multiplier determining the length of the cranks (see above row)	$[0.01, 0.95]$
i	x2 - Distance of crank's CoR from the spherical joint	$[0.01, 0.40] m$
α	x2 - Inclination of leg's mechanisms	$[0, 120]^\circ$
Actuation		
f	x2 - Rotation frequency of the cranks (the sign encodes the direction)	$[-3, 3] Hz$
$\phi_{interleg}$	Phase shift between legs of the same pair	$[0, 360]^\circ$
$\phi_{frontRear}$	Phase shift between frontal and rear legs	$[0, 360]^\circ$
Materials		
λ_t	Tangential drag coefficient of the legs	$[0, 0.04] kg/m$
λ_n	Normal drag coefficient of the legs	$[0, 0.15] kg/m$
k	Stiffness of the legs	$[50, 400] N/m$
dr	Damping reduction factor $c_{da} = dr \cdot 2\sqrt{k} M$	$[0.1, 1.5]$
c_{sf}	Static friction coefficient	$[0.1, 1]$
c_{dfmul}	Dynamic friction coefficient multiplier ($c_{df} = c_{dfmul} \cdot c_{sf}$)	$[0.1, 0.9]$
ρ_r	Robot density	$[1010, 1400] \frac{kg}{m^3}$

ing an explicit, task-based objective, it maximizes a novelty metric, that quantifies how novel an individual is with respect to others that have been previously observed. Instead of rewarding individuals performing *well*, the algorithm rewards individuals performing *differently*. Even in problems in which a clear objective can be defined, novelty search completely ignores the objective. Nevertheless, it was shown that this strategy is effective in several scenarios, and it is indeed capable of producing far better solutions with respect to settings in which the objective is explicitly defined. Proposers of this algorithm suggest [18] that this could be due to the fact that novelty search is most likely to reward counter-

Table 2: MATLAB GA settings.

Option	Value
<i>GenomeSize</i>	24
<i>PopulationSize</i>	800
<i>CrossoverFraction</i>	0.7
<i>EliteCount</i>	4
<i>TimeLimit</i>	<i>Inf</i>
<i>Generations</i>	70
<i>CreationOperator</i>	<i>gacreationuniform</i>
<i>SelectionOperator</i>	<i>gaselectionstochunif</i>
<i>CrossoverOperator</i>	<i>crossoverscattered</i>
<i>MutationOperator</i>	<i>mutationgaussian</i>

intuitive intermediate solutions that are not appealing for an objective-based algorithm, but that can, eventually, lead to a far better final solution.

Here we adopt novelty search to explore the space of morphology enabled behaviors of a robot. This idea is appealing for a number of reasons. First, it has the potential to encourage the production of diverse locomotion strategies with respect to the ones already implemented in the reference robotic prototype and, possibly, also with respect to the ones observed in nature. Second, it is capable of producing *several* alternative designs: this is highly attractive with respect to the possibility to cross the reality-gap. Having a number of various designs to test into the real world, potentially exploiting very different locomotion strategies, the chances to achieve a successful transfer from simulation to reality increase significantly. Third, it produces a potentially large set of alternative solutions in a single run, and without the need to specify each time hand-crafted fitness functions.

In order to apply novelty search to a problem, a *novelty score* must be defined for each individual, being the scalar quantity actually optimized by the algorithm. A possible procedure to compute such a score is to first map each individual into a features space, and then apply a novelty metric acting on such a space to quantify how novel an individual is with respect to previously observed ones. As for the novelty metric (essentially a sparseness metric), the average Euclidean distance to the k nearest neighbors is adopted as in related works [18]. It has to be noted that, for each individual, the k nearest neighbors are searched in a set that contains all the current population plus the so called *novelty archive* (a repository that keeps track of the most novel individuals discovered during evolution).

Features are domain-specific, and should be defined in such a way that by exploring novelty in the features space interesting behaviors can be instantiated. In our setting the behavior of each robot design is thus characterized by a vector \mathbf{x} of locomotion-related behavioral features defined as $\mathbf{x} = (\overline{s_x}, \overline{s_y}, \overline{\delta_{front}}, \overline{\delta_{rear}}) \in \mathbb{R}^4$, where $\overline{s_x}$ and $\overline{s_y}$ are, respectively, the normalized space traveled in the x and y directions, and $\overline{\delta_{front}}$ and $\overline{\delta_{rear}}$ are the mean duty factors of front and rear legs (i.e. the percentage of time legs stay in contact with the ground).

Resorting to the Euclidean distance, novelty search may be biased towards exploring features with a bigger magnitude, while overlooking those with a smaller one. In general this is not desirable, as the algorithm may completely

ignore some of the features, ending up exploring just a portion of the behavior space. Therefore, before computing the novelty score (at the end of each epoch), the following normalization is applied to the behavior space: $\bar{\mathbf{x}} = (\mathbf{x} - \mathbf{x}_{min}) / (\mathbf{x}_{max} - \mathbf{x}_{min})$, where $\bar{\mathbf{x}}$ is the normalized behavior vector, \mathbf{x}_{min} , \mathbf{x}_{max} are vectors containing, respectively, the minimum and maximum values for each feature, and the division is intended as an element-wise vector operation. Novelty search was implemented on top of MATLAB’s GAs following the C++ reference implementation by Joel Lehman [17], the only difference being that in our setting the N_{add} most novel individuals of each generation are added to the novelty archive ($N_{add} = 0.025 \cdot PopulationSize$).

2.4 Morphing

2.4.1 Morphing among close morphologies

The output of the evolutionary run is expected to be a large set of different designs exhibiting diverse locomotion strategies, that we would like to embed into a reconfigurable morphing robot. The more morphological parameters must be changed to entail a relevant behavior switch, the more complex must be the robot to allow such a complete reconfiguration (e.g. in terms of additional actuators). It is thus important to understand whether it is possible to achieve a macroscopic variation in the behavior by changing the minimum number of parameters (ideally just one). In what follows, we propose a search procedure to identify robot designs that are close in the morphology space, but far apart in behavior space.

2.4.2 Searching for candidate morphologies

The output of novelty search comprised the final population and the novelty archive, the size of such a set being, in our experiments, of 2200 individuals. The question arises of how could we search for candidate morphologies to attempt morphing in such a big set. Here we propose the following procedure:

1. Normalize the genomes space (comprising the last population and the novelty archive) as explained in Sect. 2.3.4;
2. Perform a clustering procedure in morphology space (Sect. 2.4.3);
3. Select a number of promising clusters to analyze (e.g. in which the maximum speed is above a certain threshold);
4. Inside each of the promising clusters, rank individuals for their speed, and select one candidate for morphing;
5. For each of the selected candidates search, within the same cluster, for individuals that although being close in morphology space (e.g. differing in just one parameter) maximally differ in their behavior, i.e. are far apart in behavior space. Euclidean distance is employed to estimate proximity in such spaces. Among them, select the candidate(s) for morphing;
6. Start simulating one of the selected morphologies and trigger morphing at some point of the execution (multiple times, if multiple candidates were selected).

In this procedure the human designer actively re-enters into the design process, evaluating which behaviors are more interesting, either for morphing, or because they embed promising design configurations that are most likely to be transferable to the real world at a later stage.

2.4.3 Details on the clustering procedure

For the experiments reported in this paper, *agglomerative hierarchical clustering* was applied. This approach groups data at different scales, creating a cluster tree (also called *dendrogram*, Fig. 5). The procedure does not produce a single set of clusters, but instead a multi-level hierarchy that the designer can “cut” to achieve clustering at different scales. Standardized Euclidean distance was used to compute similarity, while agglomeration is based on unweighted average distance. Clustering functions from MATLAB[®] Statistics Toolbox were adopted.

2.5 Human-machine collaborative design

In the proposed methodology the human designer works together with the evolutionary algorithm and the other computational tools during the design process, in a sort of collaborative design process. The goal is to preserve the best from both worlds. The human designer first performs the conceptual work needed to apply bioinspiration in a proper manner. Clever bioinspired solutions devised by human designers in previous work (such as the leg mechanism and the presence of soft limbs) are not overlooked, but instead used to constrain the search space and simplify the job of the evolutionary algorithm in finding simple, functional, and effective robots. Also, constraining the algorithm to adopt technical solutions that have been already implemented in several prototypes ensures that evolved morphologies can be actually built. The novelty-based evolutionary algorithm then leverage these building blocks to explore different solutions in a “creative” way. It can play with the basic morphology of the robot in any conceivable way, possibly deviating (as we will show in Sect. 3) from the original biological model. In this sense this methodology can provide hints for alternative solutions with respect to the one observed in biology. The human designer then reenters the design process with the important task of searching for interesting solutions within the results of the evolutionary run.

3. RESULTS

3.1 Novelty-based vs. objective based GA

An extensive comparison between the objective-based and the novelty-based settings is beyond the scope of this work. Here we report some experimental results to support the idea that novelty search is suitable for the intended study, comparing typical runs performed with the two approaches.

First, our experiments confirm previous observations [18] suggesting that novelty search is indeed capable of outperforming objective-based algorithms even though it completely ignores the objective. Fig. 3 shows the best horizontal speed achieved at each generation by the two algorithms. Although for novelty search this quantity is only one of the four explored features (the others being vertical speed and mean duty factors of front and hind legs) and there is no pressure applied toward discovering faster robots, the algorithm outperforms the objective-based GA.

More relevant for our intended study is the comparison reported in Fig. 4, in which the four features are plotted for individuals produced by the two evolutionary runs. The objective-based GA produces individuals that are concentrated in a narrow region of the behavior space. It can be observed how the objective-based GA converges toward a bipedal robot running backward (duty factors of front legs

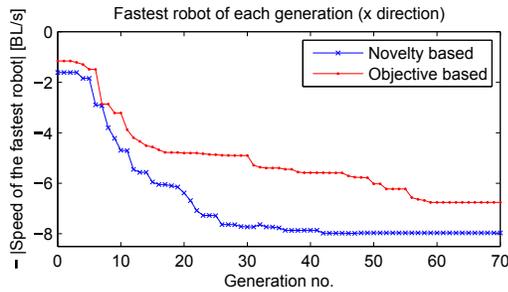


Figure 3: Speed of the fastest individual of each generation: comparison of the objective-based and novelty-based genetic algorithm. Novelty search finds faster individuals (compare absolute values).

are null for almost all individuals). Basically the algorithm returns *one* design, and individuals in the last run are mostly slight variations of the same design, exploiting the same locomotion principle. In contrast, novelty search seems to explore the four dimensional behavior space as intended, producing a heterogeneous population. Qualitatively different solutions can be observed. From the plot of space traveled, it can be observed how novelty search produces robots that travel not only forward and backward in the x direction, but also swim in all possible directions (in contrast with the objective-based setting, in which the space traveled in y never exceeds $1BL$). The production of qualitatively different solutions is confirmed by the plot of duty factors, in which robots can be observed that exhibit a wide range of behaviors (bipedal – on both pairs of legs, quadrupedal, etc.).

Results thus suggest that novelty search can be effective in exploring the behavior space of a robot and, as a consequence, its morphology space (given the marginal role of the control in our setting). The final population contains a rich repertoire of possible designs to search in. Hereafter the focus will thus be on analyzing results of the novelty-based evolutionary run.

3.2 Evolved morphologies and morphing experiments

We now analyze some of the designs extracted from the final population that, in addition to exhibiting interesting locomotion strategies, provided examples of successful morphing. Those designs are the result of the search procedure described in Sect. 2.4.2. Fig. 5 shows a typical dendrogram, resulting from clustering in morphology space. The cophenetic correlation coefficient – an index that measures how faithfully the clustering preserves the structure in the data – is 0.9043 (the closer to 1 the better). Fig. 6 details some of the most interesting designs. Configurations discovered by novelty search suggest different, multi-modal, locomotion strategies, each embedding possible design choices. For example, many designs suggest that, although being originally designed for walking and running, the leg mechanism is indeed particularly effective to achieve swimming by sculling. By keeping $m \approx i$ (as a reminder, m is the length of the crank, i is the distance between the crank’s CoR and the spherical joint), the distal parts of the legs reach high speeds, thus generating high swimming forces. According to experimental data this appears as a particularly effective strategy, confirmed by the fact that the fastest

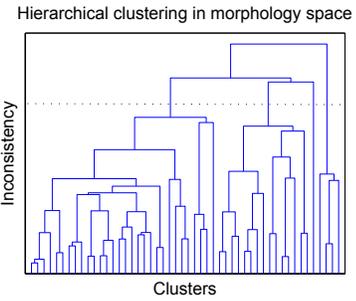


Figure 5: The dendrogram resulting from the clustering in morphology space, and a possible cut.

robot is indeed a swimming one. As for morphing, the diverse final population offers multiple alternatives to perform it, switching among qualitatively different behaviors. Examples of successful morphing are reported in Fig. 6. Note that in those examples varying just one parameter resulted to be enough to entail a macroscopic variation in the behavior (e.g. switching from running to swimming). Although here omitted for space constrains, an in-deep analysis of evolved behaviors reveals very elaborated solutions, regularities, and embodiment properties (e.g. effective locomotion leveraging passive dynamics). Finally, although evolved behaviors appear natural and life-like in many cases, lot of the evolved designs deviate from any conceivable biological model, therefore suggesting novel ways to locomote underwater.

4. CONCLUSIONS

Results confirm our hypothesis that novelty search is indeed a suitable tool to explore the behavior and morphology space of a robot. Also, the proposed methodology for human-machine collaborative design proved to be effective in order to design self-stabilized legged robots, mixing the strengths of both worlds. The human designer enters into the design process when conceptual work is needed to apply bioinspiration in a proper manner, as well as when results of an evolutionary process must be analyzed. On the other hand, the evolutionary algorithm is able to explore a huge behavior space, producing a wide range of effective alternative solutions. Morphing was also demonstrated, encouraging to build a reconfigurable robot exploiting self-stabilization to switch among different behaviors. In this regard it was also demonstrated that in water environment a single parameter is enough to determine macroscopic variations in the resulting behavior. From a design perspective, some of the most interesting parameters to exploit for morphing have been identified. Design indications were gained (as for example the suitability of the three-bar leg mechanism for performing sculling-based swimming) and more can still be gained by further analyzing evolved designs. From an artificial life perspective novel underwater “creatures” have been produced in simulation that exhibit surprising behaviors, in some cases resembling biological creatures, in some other completely deviating from biological examples.

Acknowledgments

This work is partly supported by RoboSoft - A Coordination Action for Soft Robotics (FP7-ICT-2013-C # 619319)

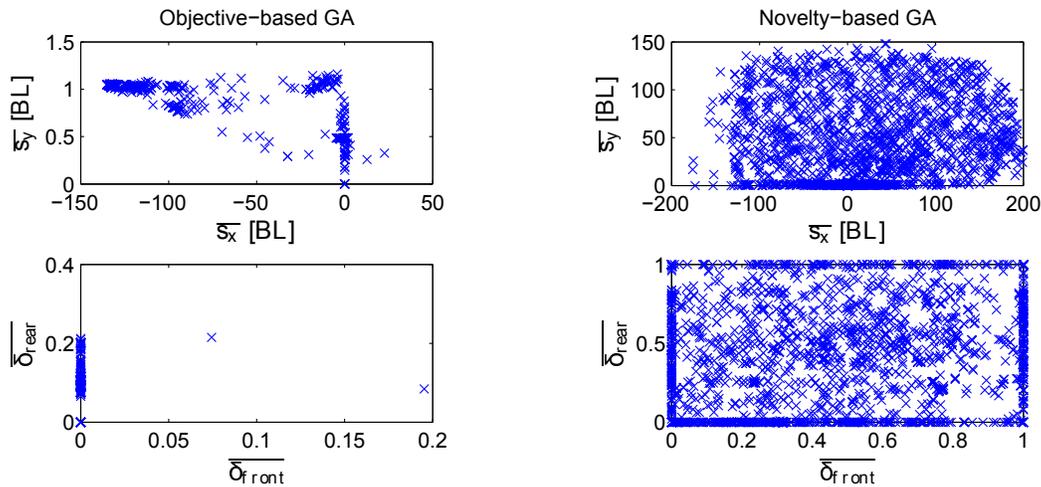


Figure 4: Behavioral space exploration: objective-based GA vs. novelty-based GA.

and by the Fondazione Livorno within the framework of the PoseiDRONE project Grant.

5. REFERENCES

- [1] A. Arienti, M. Calisti, F. G. Serchi, and C. Laschi. Poseidrone: design of a soft-bodied ROV with crawling, swimming and manipulation ability. In *Proceedings of the MTS/IEEE OCEANS. MTS/IEEE*, 2013.
- [2] J. E. Auerbach and J. C. Bongard. On the relationship between environmental and morphological complexity in evolved robots. In *Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation*, pages 521–528. ACM, 2012.
- [3] J. Bongard, V. Zykov, and H. Lipson. Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121, 2006.
- [4] J. C. Bongard. Evolutionary robotics. *Commun. ACM*, 56(8):74–83, Aug. 2013.
- [5] M. Calisti, F. Corucci, A. Arienti, and C. Laschi. Bipedal walking of an octopus-inspired robot. In *Biomimetic and Biohybrid Systems*, pages 35–46. Springer, 2014.
- [6] N. Cheney, R. MacCurdy, J. Clune, and H. Lipson. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceeding of the Fifteenth Annual Conference on Genetic and Evolutionary Computation*, pages 167–174. ACM, 2013.
- [7] M. Cianchetti, M. Calisti, L. Margheri, M. Kuba, and C. Laschi. Bioinspired locomotion and grasping in water: the soft eight-arm octopus robot. *Bioinsp. Biomim.*, to appear, 2014.
- [8] K. C. Galloway, G. C. Haynes, B. D. Ilhan, A. M. Johnson, R. Knopf, G. A. Lynch, B. N. Plotnick, M. White, and D. E. Koditschek. X-RHex: A highly mobile hexapedal robot for sensorimotor tasks. 2010.
- [9] S. Grimmer, M. Ernst, M. Günther, and R. Blickhan. Running on uneven ground: leg adjustment to vertical steps and self-stability. *Journal of Experimental Biology*, 211(18):2989–3000, 2008.
- [10] H. Hauser, R. M. Füchslin, and R. Pfeifer. *Opinions and Outlooks on Morphological Computation*. Zürich, 2014 - ISBN 978-3-033-04515-6.
- [11] J. D. Hiller and H. Lipson. Evolving amorphous robots. In *ALIFE*, pages 717–724, 2010.
- [12] F. Iida, J. Rummel, and A. Seyfarth. Bipedal walking and running with spring-like biarticular muscles. *Journal of Biomechanics*, 41(3):656–667, 2008.
- [13] A. J. Ijspeert. Biorobotics: Using robots to emulate and investigate agile locomotion. *Science*, 346(6206):196–203, 2014.
- [14] S. Kim, J. E. Clark, and M. R. Cutkosky. iSprawl: Design and tuning for high-speed autonomous open-loop running. *The International Journal of Robotics Research*, 25(9):903–912, 2006.
- [15] C. Langton. *Artificial Life: An Overview*. A Bradford book. BRADFORD BOOK, 1997.
- [16] J. C. Larsen, D. Brandt, and K. Stoy. Locokit: A robot construction kit for studying and developing functional morphologies. In *From Animals to Animats 12*, pages 12–22. Springer, 2012.
- [17] J. Lehman. Novelty search C++ 1.0 - <http://eplex.cs.ucf.edu/uncategorised/software>, 2008.
- [18] J. Lehman and K. O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.
- [19] R. Pfeifer, F. Iida, and J. Bongard. New robotics: Design principles for intelligent systems. *Artificial Life*, 11(1-2):99–120, 2005.
- [20] R. Pfeifer, M. Lungarella, and F. Iida. Self-organization, embodiment, and biologically inspired robotics. *Science*, 318(5853):1088–1093, 2007.
- [21] H. V. Qu, G. Ramstein, F. Casanova, L. Aryananda, M. Hoffmann, F. I. Sheikh, and H. Hauser. Gait versatility through morphological changes in a new quadruped robot. In *5th International Symposium on Adaptive Motion of Animals and Machines*, 2011.
- [22] J. Rieffel and S. Smith. A face-encoding grammar for the generation of tetrahedral-mesh soft bodies. In *ALIFE*, pages 414–420, 2010.
- [23] M. Sfakiotakis, A. Kazakidi, N. Pateromicelakis, J. A. Ekaterinaris, and D. P. Tsakiris. Robotic underwater propulsion inspired by the octopus multi-arm swimming. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3833–3839. IEEE, 2012.
- [24] K. Sims. Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer Graphics and Interactive Techniques*, pages 15–22. ACM, 1994.
- [25] B. G. Smith, C. M. Saaj, and E. Allouis. Evolving legged robots using biologically inspired optimization strategies. In *Robotics and Biomimetics (ROBIO), 2010 IEEE International Conference on*, pages 1335–1340. IEEE, 2010.
- [26] J. T. Watson, R. E. Ritzmann, and A. J. Pollack. Control of climbing behavior in the cockroach, *blaberus discoidalis*. ii. motor activities associated with joint movement. *Journal of Comparative Physiology A*, 188(1):55–69, 2002.
- [27] J. Yosinski, J. Clune, D. Hidalgo, S. Nguyen, J. Zagal, and H. Lipson. Evolving robot gaits in hardware: the hyperneat generative encoding vs. parameter optimization. In *Proceedings of the 20th European Conference on Artificial Life*, pages 890–897, 2011.
- [28] V. Zykov, J. Bongard, and H. Lipson. Evolving dynamic gaits on a physical robot. In *Proceedings of Genetic and Evolutionary Computation Conference, Late Breaking Paper, GECCO*, volume 4, 2004.

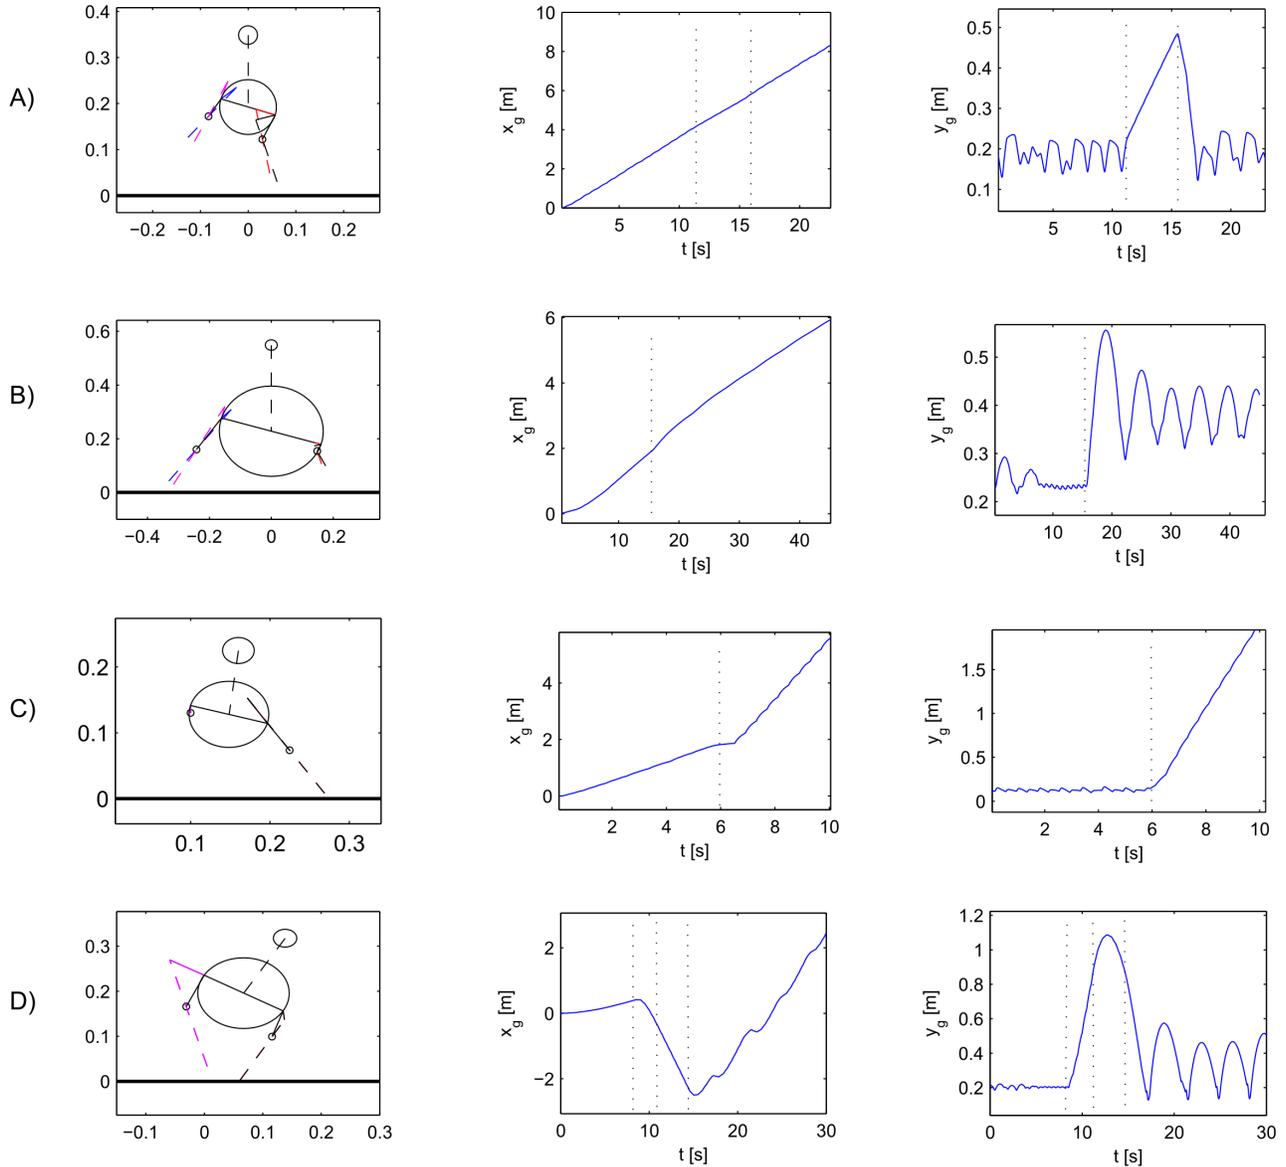


Figure 6: *Simulation results: some of the evolved morphologies and examples of successful morphing.* Morphing from one morphology m_1 to another morphology m_2 is implemented as a gradual linear transformation in morphology space, occurring in a predefined time (*morphingTime*) A) A quadruped design resembling some kind of fast terrestrial quadruped (e.g. antelopes). The robot shows functional differentiation of front and hind legs. Hind legs provide powerful swimming strokes during the flight phase, modifying the ballistic trajectory of the robot and accelerating its descent towards the ground. Frontal legs absorb the subsequent impact, then releasing energy and making the robot bounce. Around $t = 11s$ the β parameter (Tab. 1) is decreased, with the effect of tilting the robot backward: the robot starts swimming. Shortly after $t = 15s$, β is brought back to the original value: the robot absorbs the impact with the ground and converges again toward a terrestrial gait. B) A bipedal robot that uses hind legs to push himself forward (front legs disappeared during evolution). Fluid resistance helps the robot to balance. At $t = 15s$ the length of the rear cranks is increased, entailing a change in the trajectory of rear feet. As a consequence the gait changes considerably. C) Another bipedal design, this time exploiting front legs to run (hind legs disappeared during evolution). Around $t = 6s$ the β parameter is decreased (similarly to A), tilting the robot backward: the robot starts to swim. D) This robot is tilted toward its front side, and exploits frontal legs to run. However, it has a functional pair of hind legs that, although being deactivated, appear to be optimized for swimming ($m \approx i$). Around $t = 8s$ rear legs are activated with a forward rotation: the robot inverts its motion and starts swimming backward. Around $t = 11s$ the β parameter is brought close to zero: this redirects the swimming propulsive action of rear legs, bringing the robot closer to ground. Around $t = 14.57s$ the rotation applied to hind legs is inverted: when the robot touches the ground, it converges to a hopping terrestrial gait. Other configurations are here omitted for space constraints.