

Simple or complex bodies? Trade-offs in exploiting body morphology for control

Matej Hoffmann and Vincent C. Müller

Abstract Engineers fine-tune the design of robot bodies for control purposes, however, a methodology or set of tools is largely absent, and optimization of morphology (shape, material properties of robot bodies, etc.) is lagging behind the development of controllers. This has become even more prominent with the advent of compliant, deformable or "soft" bodies. These carry substantial potential regarding their exploitation for control—sometimes referred to as "morphological computation". In this article¹, we briefly review different notions of computation by physical systems and propose the dynamical systems framework as the most useful in the context of describing and eventually designing the interactions of controllers and bodies. Then, we look at the pros and cons of simple vs. complex bodies, critically reviewing the attractive notion of "soft" bodies automatically taking over control tasks. We address another key dimension of the design space—whether model-based control should be used and to what extent it is feasible to develop faithful models for different morphologies.

1 Introduction

It has become increasingly common to explain the intelligent abilities of natural agents through reference to their bodily structure, their morphology, and to make extended use of this morphology for the engineering of intelligent abilities in artificial agents, e.g. robots—thus "offloading" computational processing from a central con-

Matej Hoffmann

iCub Facility, Istituto Italiano di Tecnologia, Via Morego 30, 16163 Genoa, Italy
e-mail: matej.hoffmann@iit.it

Vincent C. Müller

Name, Division of Humanities & Social Sciences, Anatolia College/ACT, Pylaia-Thessaloniki, Greece, e-mail: vmueller@act.edu

¹ This article is a substantially revised version of [20].

troller to the morphology. These uses of morphology for explanation and engineering are sometimes referred to as "morphological computation" (e.g., [15, 16, 39]). However, in our view, only some of the characteristic cases that are embraced by the community as instances of morphological computation have a truly computational flavor. Instead, many of them are concerned with exploiting morphological properties to simplify a control task. This has been labeled "morphological control" in [14]; "mechanical control" could be an alternative label. Developing controllers that exploit a given morphology is only a first step. The space of possible solutions to a task increases dramatically once the mechanical design is included in the design space: imagine having a hand with 10 instead of 5 fingers: there will be completely new ways of grasping things. At the same time, the search space of controllers and mechanical design combinations also becomes enormous.

In this work, we want to take a close look at these issues. First, we will borrow the "trading spaces" landscape from [41] that introduces a number of characteristic examples and distributes them along a metaphorical axis from "informational computation" to "morphological computation". Second, we will analyze under what circumstances can physical bodies be said to compute and then propose the dynamical systems description as the most versatile framework to deal with brain-body-environment interactions. Third, we will critically look at the pros and cons of simple vs. complex (highly dimensional, dynamic, nonlinear, compliant, deformable, "soft") bodies. Fourth, we will address another key dimension of the design space—whether model-based control should be used and to what extent it is feasible to develop faithful models for different morphologies. We will close with an outlook into the future of "soft" robotics.

2 Design "trading spaces"

Pfeifer et al. [41] offer one possible perspective on the problem in Fig. 1. In traditional robots—as represented by industrial robots and Asimo in the figure—, control is essentially confined to the software domain where a model of the robot exists and current state of the robot and the environment is continuously being updated in order to generate appropriate control actions sent to the actuators. In biological organisms, on the other hand, this does not seem to be the case: the separation between "controllers" and "controlled" is much less clear and behavior is orchestrated through a distributed network of interactions of informational (neural) and physical processes. Furthermore, there is no centralized neural control, but a multitude of recurrent loops from the lowest level (e.g., reflexes and pattern generators in the spinal cord) to different subcortical and cortical areas in the brain. At the same time, the bodies themselves tend to be much more complex in terms of geometrical as well as dynamical properties. This has motivated the design of compliant, tendon-driven robots like ECCE [51] or Kenshiro [35], and soft, deformable robots like Octopus (e.g., [29]) (we are moving from left to right in Fig. 1). However, compared to humans or biological octopus, a comparable level of versatility and robustness in

the orchestration of behavior has not yet been achieved in the robotic counterparts. In more restricted settings, the design and subsequent exploitation of morphology is easier, as the jumping and landing robot frog [36], the passive dynamic based walker (Cornell Ranger [4]), or the coffee balloon gripper demonstrate. The predecessors of the Cornell Ranger, the original passive dynamic walkers [33] are a powerful demonstration that appropriate design of morphology can generate behavior in complete absence of software control. Yet, there is only a single behavior and the environmental niche is very narrow. The coffee-balloon gripper [8] employs a similar strategy, but achieves surprising versatility on the types of objects that can be grasped. Body designs that follow this guideline were also labeled "cheap designs" [42].

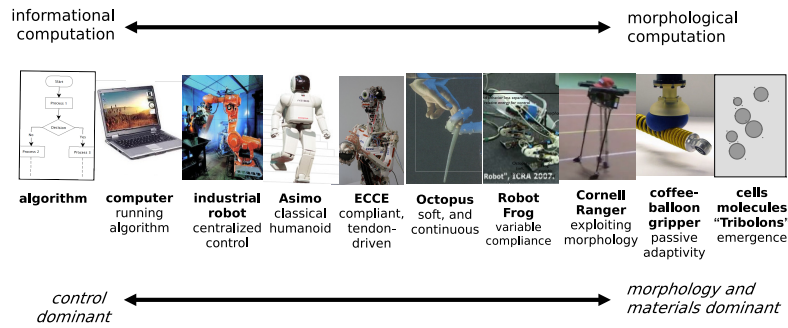


Fig. 1 The design trading space. This figure illustrates the degree to which each system relies on explicit control or self-organization of mechanical dynamics. On the left-hand side of the spectrum, computer algorithms and commercial computers rely on physical self-organization at the minimum level, while towards the right-hand side, more embodied, more soft, and smaller-scale systems require physical interactions as driving forces of behaviors. The design goal then is to find a proper compromise between efficiency and flexibility, taking into account that a certain level of flexibility can also be achieved by changing morphological and material characteristics. (Fig. and caption from [41]).

3 Is the body really computing?

The systems toward the right-hand-side of Fig. 1 rely on physical interactions rather than computer programs to orchestrate their behaviors. This end of the spectrum is labeled "morphological computation". However, in which sense can these systems be said to compute? In the case of the passive dynamic walker and its active descendants (Cornell Ranger) or the jamming-based (coffee-balloon) grippers cited above, the body is ingeniously contributing to its, perhaps primary, function: enabling physical behavior in the real world. This is often interpreted in the "offloading sense": the body design takes over computation from the brain (e.g., [39])—the hypothetical

computation that is needed for walking has been fully off-loaded from a hypothetical controller to the morphology of the walker. However, we argue that this view is hard to defend beyond the level of a metaphor. It is difficult to imagine a real example where one could choose to solve the task “through the brain” or “through the body” and smoothly interchange their contributions.

A word on what we mean by computation is in order. Let’s take the Cornell Ranger example—a robot based on the passive dynamic walker with a simple controller on top. The robot is certainly not performing abstract digital computation (as represented by the Turing model, for example). Borrowing the terminology from [13, p. 5-6], the part of the controller can be said to perform online and embedded computation—such computation is interactive rather than batch, as it relies on a continuous stream of inputs (from sensors in this case) for its execution and produces a continuous stream of outputs (control actions). However, it is the physical interaction, not the controller, that plays the key part in accomplishing the task here. Some authors would subsume this type of interaction under a computational framework as well—e.g., “embodied computation should be understood as a physical process in an ongoing interaction with its environment” [13, p. 6]. Other authors pose much stricter requirements on physical computation: according to Horseman et al. [22], a physical system can be said to compute only if it was designed as such. That is, there needs to be a user that has an abstract computational problem that he wants to solve by a physical machine. This machine (the computer) needs to be designed and its model derived that allows for encoding of abstract inputs into the machine and decoding them again after physical evolution of the machine’s state. In this view, computation cannot be assigned *ex post* and physical systems with interesting computational properties, “intrinsically computing” [11], do not fulfill these requirements.

It is not central to practitioners whether the controlled system is “computing”. However, a unified theory or level of description is desirable. The dynamical systems framework seems to be the most versatile in this context, as it (i) fits the informational and physical processes equally well; (ii) copes with continuous (in time) streams of continuous input and output signals; (iii) is already used by control theory.

4 Dynamical systems perspective

Let us look at the concept of self-stabilization, which is often cited in the “morphological computation community”. While maximally exploiting the interaction of the body with the environment can lead to “pure physics walking” like in the passive dynamic walker case [33], what if the agent is perturbed out of this preferred regime? It seems that corrective action needs to be taken. However, it can be the very same mechanical system that can generate this corrective response. This phenomenon is known as self-stabilization and is a result of a mechanical feedback loop. To use the dynamical systems description, certain trajectories (such as

walking with a particular gait) have attracting properties and small perturbations are automatically corrected, without control, or one could say that "control" is inherent in the mechanical system. Examples of this phenomenon are a self-stable bike, driving alone after being pushed and compensating for major disturbances [27] or the contribution of biological muscles to human walking as reviewed by Blickhan et al. [5] in a paper entitled "Intelligence by mechanics" (more examples in this line can be found in [19] or [18] with videos of the bicycle and other material).

A general formulation of a control problem in control theory is making a dynamical system follow a desired trajectory. For our purposes, we will consider the cases where the dynamical system is physical—the body of the agent; in control theory, this is the so-called "plant". There are numerous control schemes and branches of control theory and the reader is referred to abundant literature on the topic (e.g., [3, 12, 25]). The performance of the controller can be evaluated on various grounds: precision of a trajectory with respect to a reference trajectory, or energy expenditure, for example. In addition, performance, stability and robustness guarantees are required by industry. Control theory typically deals with the design of controllers that optimize these criteria. Some control schemes with appropriate cost functions will automatically result in minimal control actions and thus "optimize the contribution of the morphology". For example, Moore et al. [34] used Discrete Mechanics and Optimal Control to steer a satellite while exploiting its dynamics to the maximum. Carbajal [10] developed related methods for reaching, plus offered a formalization of the concept of "natural dynamics". Nevertheless, the plant is treated as fixed in these approaches. Yet, the properties of the physical body obviously have a key influence on the final performance of the whole system (plant + controller), which calls for including them into the design space.

5 Simple or complex bodies?

The spirit of the morphological computation literature that follows the "offloading" or "trade-off" perspective, is that complex (highly dimensional, dynamic, nonlinear, compliant, deformable, "soft") bodies are advantageous for control because they can take over the "computation" that a controller would otherwise have to perform (e.g., [15, 16, 39] or [9] explicitly in Fig. 1). Complex nonlinear bodies certainly give rise to more complex dynamical landscapes where the location of attractors *could*—if properly exploited—facilitate the performance on a given task.

This view is in stark contrast to the views prevalent in control theory. There, linear time-invariant systems are the ideal plants to control. Solutions for nonlinear systems are much more difficult to obtain and they often involve a linearization of the system of some sort. In fact, human-like bodies are a nightmare for control engineers ([43] is an interesting case study) and highly complex models and controllers would be required.

What would be an ideal body then? And, does a complex body imply simple or complex control? Recent attempts at quantifying the amount of morphological

computation shed more light on this issue. Zahedi and Ay [52] propose two concepts for measuring the amount of morphological computation by calculating the conditional dependence of future world states W' (encompassing the body state) on previous world states W and action A taken by the agent. According to Concept 1, the amount of morphological computation is inversely proportional to the contribution of the agent's actions to the overall behavior. That is, if action of the agent's motors (A) has little influence on the future physical state of the agent in the environment (W'), morphological computation is high. Concept 2 calculates the amount of morphological computation by isolating the positive contribution of the world to the overall behavior (effect of W on W'), obtained from the "difference" between conditional probability distributions with and without the action variable, $p(w' | w, a)$ vs. $p(w' | w)$ (see [52] for details). Here, systems with high morphological computation would be those with strong "body dynamics" or "natural dynamics" (see e.g., [23] or [10] for a formal definition). However, optimizing for morphological computation in the above sense, one would arrive at systems with strong internal dynamics (Concept 2), resisting control actions (Concept 1), which seems very impractical for engineers. In fact, Klyubin et al. [26] proposed a different measure relying on information theory, empowerment, which is equivalent to the opposite of morphological computation under Concept 1; maximizing empowerment amounts to maximizing the effect of the agent's actions.

Rückert and Neumann [45] study learning of optimal control policies for a simulated 4-link pendulum which needs to maintain balance in the presence of disturbances. The morphology (link lengths and joint friction and stiffness) is manipulated and controllers are learned for every new morphology. They show that: (1) for a single controller, the complexity of the control (as measured by the "variability" of the controller) varies with the properties of the morphology: certain morphologies can be controlled with simple controllers; (2) optimal morphology depends on the controller used; (3) more complex (time-varying) controllers achieve much higher performance than simple control across morphologies.

In summary, the performance on a task will always depend on a complex interplay of the controller, body, and environment: taking out the controller is just as big a mistake as taking out the body was. The tasks that can be completely solved by appropriate tuning of the body, like passive dynamic walking, are the exception rather than the rule. A controller will thus be needed too. A complex body may have the potential to partially solve certain tasks on its own; yet, it may present itself as difficult to control, model (if the controller is relying on models), design, and manufacture. An optimal balance thus needs to be found. For that, however, new design methodologies that would encompass complex cost functions (performance on a task, versatility, robustness, costs associated with hardware whose parameters can be manipulated etc.) are needed. Hermans et al. [17] very recently proposed such a method that uses machine learning to optimize physical systems; an approximate parametric model of the system's dynamics and sufficient examples of the desired dynamical behavior need to be available though—which leads us to the next section.

6 With or without a model?

Including the parameters of the body into the design considerations may give rise to better performance of the whole system; these may be solutions involving a simpler controller, but also solutions that were previously unattainable when the body was fixed. Following the dynamical system's perspective, [14] provide an illustration of the possible goals of the design process: (1) To design the physical dynamical system such that desired regions of the state space have attracting properties. Then it is sufficient to use a simple control signal that will bring the system to the basins of attraction of individual stable points that correspond to target behaviors. (2) More complicated behavior can be achieved if the attractor landscape can be manipulated by the control signal.

If a mathematical formulation of the controller and the plant is available, this design methodology can be directly applied. The first part is demonstrated by McGeer [33] on the passive dynamic walker: The influence of scale, foot radius, leg inertia, height of center of mass, hip mass and damping, mass offset, and leg mismatch is evaluated. In addition, the stability of the walker is calculated. Recently, Jerrold Marsden and his coworkers presented a method that allows for co-optimization of the controller and plant by combining an inner loop (with discrete mechanics and optimal control) and an outer loop (multiscale trend optimization). They applied it to a model of a walker and obtained the best position of the knee joints ([38] – Ch. 5).

However, typical real-world agents are more complex than simple walkers. Holmes et al. [21] provide an excellent dynamical systems analysis of the locomotion of rapidly running insects and derive implications for the design of the RHex robot. Yet, they conclude that "a gulf remains between the performance we can elicit empirically and what mathematical analyses or numerical simulations can explain. Modeling is still too crude to offer detailed design insights for dynamically stable autonomous machines in physically interesting settings." Hermans et al. [17] similarly note that applying their method to robotics, which is known to suffer from lack of accurate models, is a challenge. The modeling and optimization of more complicated morphologies—like compliant structures—is nevertheless an active research topic (e.g., Wang [50] and other work by the author). The second point of Fuchslin et al. [14]—achieving "morphological programmability" by constructing a dynamical system with a parametrized attractor landscape—remains even more challenging though.

One of the merits of exploiting the contributions of body morphology should be that the physical processes do not need to be modeled, but can be used directly. However, without a model of the body at hand, several body designs need to be produced and—together with the controller—tested in the respective task setting. The design space of the joint controller-body system blows up and we may be facing a curse of dimensionality. This is presumably the strategy adopted by the evolution of biological organisms that could cope with the enormous dimensionality of the space. In robotics, this has been taken up by evolutionary robotics [37]. The simulated agents of Karl Sims [47] demonstrate that co-evolving brains and bodies

together can give rise to unexpected solution to problems. More recently, Bongard [6] showed that morphological change indeed accelerates the evolution of robust behavior in such a brain-body co-evolution setting. With the advent of rapid prototyping technologies, physics-based simulation could be complemented by testing in real hardware [31], but this reintroduces the modeling through the back door: the phenotypes in the simulator now become models and they need to sufficiently match their real counterparts. Yet, a "reality gap" [24, 28] always remains between simulated and real physics. The only alternative is to optimize in hardware directly, which is in general slow and costly. Brodbeck et al. [7] provide an interesting illustration how locomoting cube-like creatures can be evolved in a model-free fashion through automated manufacturing and testing. However, in summary, the design decisions—which parameters to optimize—are based on heuristics and a clear methodology is still missing. Furthermore, with the absence of an analytical model of the controller and plant, no guarantees on the system's performance can be given.

7 Conclusion and outlook

"Morphological computation" or "morphological control" are very attractive concepts, receiving significant attention and carrying great potential. The rich properties of "soft" bodies (highly dimensional, dynamic, nonlinear, compliant and deformable) have been largely overlooked or deliberately suppressed by classical mechatronic designs, as they are largely incompatible with traditional control frameworks, where linear plants are preferred. This is definitely a missed opportunity. On the other hand, while complex bodies carry a lot of "auto-control" potential, this property does not come for free. In this article, we provided a critical review of the design "trading spaces", an imaginary landscape from "control dominant" systems whose natural dynamics is suppressed to designs that capitalize on self-organization of the physical system interacting with the environment. We conclude that the contributions of the body to the task are not computational in any substantial sense and proposed the dynamical systems descriptions as the most versatile in order to facilitate description, understanding, control and design of brain-body-environment systems. The pros and cons of simple vs. complex bodies were illustrated on examples. It has to be said that the exploitation of truly complex bodies to accomplish tasks is still mostly at a "proof-of-concept" stage. A closely connected issue is the one of modeling of these systems—soft bodies are notoriously difficult to model. The model may not be necessary for the system to perform the task, however, without a model, the understanding and design is more complicated and performance guarantees are limited. The field, which has been dominated by heuristics so far, needs to embrace more systematic approaches that allow to navigate in this complex landscape.

In terms of applications, the most relevant area where exploitation of morphology is and will be the key is probably robotics, and in particular soft robotics (see [2, 40, 41, 49] and the first issue of the journal *Soft Robotics* [48]). "Soft" robots, with

the robot Octopus (e.g., [29]) serving as a good representative, break the traditional separation of control and mechanics and exploit the morphology of the body and properties of materials to assist control as well as perceptual tasks. Pfeifer et al. [40] even discuss a new industrial revolution. Appropriate, "cheap", designs lead to simpler control structures, and eventually can lead to technology that is cheap in a monetary sense and thus more likely to impact on practical applications. Yet, a lot of research in design, simulation and fabrication is needed (see [30, 32, 46] for reviews).

The area of soft robotics and morphological computation seems to be rife with different trading spaces [41]. As we move from the traditional engineering framework with a central controller that commands a "dumb" body toward delegating more functionality to the physical morphology, some convenient properties will be lost. In particular, the solutions may not be portable to other platforms anymore, as they will become dependent on the particular morphology and environment (the passive dynamic walker is the extreme case). The versatility of the solutions is likely to drop as well. To some extent, the morphology itself can be used to alleviate these issues—if it becomes adaptive. Online changes of morphology (like changes of stiffness or shape) thus constitute another tough technological challenge (see also project LOCOMORPH [1]). Completely new, distributed control algorithms that rely on self-organizing properties of complex bodies and local distributed control units will need to be developed [32, 44].

Acknowledgements M.H. was supported by the Marie Curie Intra European Fellowship iCub Body Schema (625727) within the 7th European Community Framework Programme. M.H. also thanks Juan Pablo Carbajal for fruitful discussions and pointers to literature. Both authors thank the EUCogIII project (FP7-ICT 269981) for making us talk to each other.

References

1. Project LOCOMORPH. FP7-ICT-231688.
2. A. Albu-Schaffer, O. Eiberger, M. Grebenstein, S. Haddadin, C. Ott, T. Wimbock, S. Wolf, and G. Hirzinger. Soft robotics. *Robotics Automation Magazine, IEEE*, 15(3):20–30, september 2008.
3. K.J. Aström and R. M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 2008.
4. Pranav A Bhounsule, Jason Cortell, Anoop Grewal, Bram Hendriksen, JG Daniël Karssen, Chandana Paul, and Andy Ruina. Low-bandwidth reflex-based control for lower power walking: 65 km on a single battery charge. *The International Journal of Robotics Research*, 33(10):1305–1321, 2014.
5. R. Blickhan, A. Seyfarth, H. Geyer, S. Grimmer, H. Wagner, and M. Guenther. Intelligence by mechanics. *Phil. Trans. R. Soc. Lond. A*, 365:199–220, 2007.
6. Josh Bongard. Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences*, 108(4):1234–1239, 2011. It is in the "Legged locomotion" folder.
7. Luzius Brodbeck, Simon Hauser, and Fumiya Iida. Morphological evolution of physical robots through model-free phenotype development. *PloS one*, 10(6):e0128444, 2015.

8. Eric Brown, Nicholas Rodenberg, John Amend, Annan Mozeika, Erik Steltz, Mitchell R Zakin, Hod Lipson, and Heinrich M Jaeger. From the cover: Universal robotic gripper based on the jamming of granular material. *Proc Natl Acad Sci U.S.A.*, 107(44):18809–18814, Nov 2010.
9. K. Caluwaerts, M. D’Haene, D. Verstraeten, and B. Schrauwen. Locomotion without a brain: physical reservoir computing in tensegrity structures. *Artificial Life*, 19(1):35–66, 2013.
10. J. P. Carbajal. *Harnessing Nonlinearities: Behavior Generation from Natural Dynamics*. PhD thesis, University of Zurich, 2012.
11. J.P. Crutchfield, W.L. Ditto, and S. Sinha. Introduction to focus issue: intrinsic and designed computation: information processing in dynamical systems-beyond the digital hegemony. *Chaos*, 20(3):037101.1–037101.6, 2010.
12. A. Emami-Naeini, G. F. Franklin, and J. D. Powell. *Feedback control of dynamic systems*. Prentice Hall, 2002.
13. Nir Fresco. *Physical computation and cognitive science*. Springer, 2014.
14. R.M. Fuchslin, A. Dzyakanchuk, D. Flumini, H. Hauser, K.J. Hunt, R.H. Luchsinger, B. Reller, S. Scheidegger, and R. Walker. Morphological computation and morphological control: steps towards a formal theory and applications. *Artificial Life*, 19(1):9–34, 2013.
15. H. Hauser, A.J. Ijspeert, R.M. Fuchslin, R. Pfeifer, and W. Maass. Towards a theoretical foundation for morphological computation with compliant bodies. *Biological Cybernetics*, 105:355–370, 2011.
16. H. Hauser, A.J. Ijspeert, R.M. Fuchslin, R. Pfeifer, and W. Maass. The role of feedback in morphological computation with compliant bodies. *Biological Cybernetics*, 106:595–613, 2012.
17. M. Hermans, B. Schrauwen, P. Bienstman, and J. Dambre. Automated design of complex dynamic systems. *PLOS ONE*, 9(1):e86696, 2014.
18. M. Hoffmann, D. Assaf, and R. Pfeifer. A tutorial on embodiment, 2011. <http://www.eucognition.org/index.php?page=tutorial-on-embodiment>.
19. M. Hoffmann and R. Pfeifer. *The Implications of Embodiment: Cognition and Communication*, chapter The implications of embodiment for behavior and cognition: animal and robotic case studies, pages 31–58. Exeter: Imprint Academic, 2011.
20. Matej Hoffmann and Vincent C. Müller. Trade-offs in exploiting body morphology for control: from simple bodies and model-based control to complex bodies with model-free distributed control schemes. In Helmut Hauser, Rudolf M. Fuchslin, and Rolf Pfeifer, editors, *E-book on Opinions and Outlooks on Morphological Computation*, chapter 17, pages 185–194. 2014.
21. P. Holmes, R. J. Full, D. Koditschek, and J. Guckenheimer. The dynamics of legged locomotion: Models, analyses and challenges. *SIAM Review*, 48(2):207–304, 2006.
22. C. Horsman, S. Stepney, R.C. Wagner, and V. Kendon. When does a physical system compute? *Proc. R. Soc. A*, 470(2169):20140182, 2014.
23. F. Iida, G. Gómez, and R. Pfeifer. Exploiting body dynamics for controlling a running quadruped robot. In *Proceedings of the 12th Int. Conf. on Advanced Robotics (ICAR05)*, pages 229–235, Seattle, U.S.A., 2005.
24. Nick Jakobi, Phil Husbands, and Inman Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In *Advances in artificial life*, pages 704–720. Springer, 1995.
25. D. Kirk. *Optimal control theory: an introduction*. Dover Publications, 2004.
26. Alexander S Klyubin, Daniel Polani, and Chrystopher L Nehaniv. All else being equal be empowered. In *Advances in Artificial Life*, pages 744–753. Springer, 2005.
27. JDG Kooijman, JP Meijaard, Jim M Papadopoulos, Andy Ruina, and AL Schwab. A bicycle can be self-stable without gyroscopic or caster effects. *Science*, 332(6027):339–342, 2011.
28. Sylvain Koos, Jean-Baptiste Mouret, and Stéphane Doncieux. The transferability approach: Crossing the reality gap in evolutionary robotics. *Evolutionary Computation, IEEE Transactions on*, 17(1):122–145, 2013.
29. C. Laschi, M. Cianchetti, B. Mazzolai, L. Margheri, M. Follador, and P. Dario. Soft robot arm inspired by the octopus. *Advanced Robotics*, 26(7):709–727, 2012.
30. H. Lipson. Challenges and opportunities for design, simulation, and fabrication of soft robots. *Soft Robotics*, 1:21–27, 2013.

31. H. Lipson and J.B. Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974–978, 2000.
32. MA McEvoy and N Correll. Materials that couple sensing, actuation, computation, and communication. *Science*, 347(6228):1261689, 2015.
33. T. McGeer. Passive dynamic walking. *The International Journal of Robotics Research*, 9(2):62–82, 1990.
34. A. Moore, S. Ober-Blöbaum, and J.E. Marsden. Trajectory design combining invariant manifolds with discrete mechanics and optimal control. *Journal of Guidance, Control, and Dynamics*, 35(5):1507–1525, 2012.
35. Yoichiro Nakanishi, Yuji Asano, Toyotaka Kozuki, Hiroshi Mizoguchi, Yotaro Motegi, Masahiko Osada, Tokimasa Shirai, Junichi Urata, Kenichi Okada, and Masayuki Inaba. Design concept of detail musculoskeletal humanoid kenshiro – toward a real human body musculoskeletal simulator. In *Humanoid Robots (Humanoids), 2012 12th IEEE-RAS International Conference on*, pages 1–6. IEEE, 2012.
36. R. Niiyama, A. Nagakubo, and Y. Kuniyoshi. Mowgli: A bipedal jumping and landing robot with an artificial musculoskeletal system. In *IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 2546–2551. IEEE, 2007.
37. Stefano Nolfi and Dario Floreano. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*, volume 26. MIT press Cambridge, 2000.
38. D.N. Pekarek. *Variational methods for control and design of bipedal robot models*. PhD thesis, California Institute of Technology, 2010.
39. R. Pfeifer and J.C. Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT Press, Cambridge, MA, 2007.
40. R. Pfeifer, M. Lungarella, and F. Iida. The challenges ahead for bio-inspired ‘soft’ robotics. *Communications of the ACM*, 55(11):76–87, 2012.
41. R. Pfeifer, H.G. Marques, and F. Iida. Soft robotics: the next generation of intelligent machines. In *Proc. 23rd Int. Joint Conf. on Artificial Intelligence*, pages 5–11. AAAI Press, 2013.
42. R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, 1999.
43. V. Potkonjak, B. Svetozarevic, K. Jovanovic, and O. Holland. The puller-follower control of compliant and noncompliant antagonistic tendon drives in robotic systems. *Int. Journal Advanced Robotic Systems*, 8(5):143–155, 2011.
44. J.A. Rieffel, F.J. Valero-Cuevas, and H. Lipson. Morphological communication: exploiting coupled dynamics in a complex mechanical structure to achieve locomotion. *Journal of the Royal Society Interface*, 7(45):613–621, 2010.
45. E. Rückert and G. Neumann. Stochastic optimal control methods for investigating the power of morphological computation. *Artificial Life*, 19:115–131, 2013.
46. Daniela Rus and Michael T Tolley. Design, fabrication and control of soft robots. *Nature*, 521(7553):467–475, 2015.
47. K. Sims. Evolving 3D morphology and behavior by competition. *Artificial Life*, 1(4):353–372, 1994.
48. B. Trimmer. A journal of soft robotics: Why now? *Soft Robotics*, 1:1–4, 2013.
49. D. Trivedi, C.D. Rahn, W.M. Kier, and I.D. Walker. Soft robotics: Biological inspiration, state of the art, and future research. *Applied Bionics and Biomechanics*, 5(3):99–117, 2008.
50. Michael Yu Wang. A kinetoelastic formulation of compliant mechanism optimization. *Journal of Mechanisms and Robotics*, 1(2):021011, 2009.
51. S. Wittmeier, C. Alessandro, N. Bascarevic, K. Dalamagkidis, D. Devereux, A. Diamond, M. Jäntschi, K. Jovanovic, R. Knight, H.G. Marques, et al. Toward anthropomorphic robotics: Development, simulation, and control of a musculoskeletal torso. *Artificial life*, 19(1):171–193, 2013.
52. K. Zahedi and N. Ay. Quantifying morphological computation. *Entropy*, 15(5):1887–1915, 2013.