Leveraging Morphological Computation for Controlling Soft Robots

Helmut Hauser, Thrishantha Nanayakkara, Fulvio Forni
POC: H. Hauser (helmut.hauser@bristol.ac.uk)

June 3, 2024

Abstract

Traditional robot designs typically employ rigid body parts and high torque servo motors. This helps to obtain simple, reproducible models and therefore to facilitate control. Soft robotics is a new research field that deliberately expands the design toolbox to a wide range of smart and often soft materials. This approach is generally inspired by the remarkable performance of biological systems, which use soft structures to interact successfully with noisy and hard-to-model environments and, as a result, outperforming state-of-the-art robots in open world scenarios. However, using soft bodies comes with a significant disadvantage. Soft materials often have complex and nonlinear dynamics, which makes them hard to model and therefore difficult to control. To fulfill the potential of soft robotics to achieve performances close to biological systems this control problem has to be solved. A promising solution is another bio-inspired principle called morphological computation, which proposes to outsource functionality directly to the body morphology. From this point of view, the seemingly undesired nonlinear dynamics become suddenly a resource for implementing nonlinear functionalities. This extends the control design problem to the question of how to design the body morphology of the robot. While there exist proof-of-concepts that demonstrate the potential of this approach, the existing work, for the most part, is lacking mathematical rigor and a general framework. We believe that the control community has the right set of tools to support the development of a design framework for morphological computation. The goal of this article is therefore to provide an introduction to the concepts of soft robotics and morphological computation, explain how they can work together, and, with the help of examples, illustrate their potential for control. The hope is to inspire members of the control community to develop novel control frameworks for the next generation of soft robots.

Introduction

The field of robotics is decades old and highly successful. As an industry, even today, it's still one of the fastest growing sectors. Interestingly, the way we build and control robotic systems hasn't changed very much since the early days of robotics in the 1940's. Since then, typically, robots are made of rigid body parts connected with high torque servos. One of the reasons is such systems can be easily modelled and, therefore, easily controlled. Consequently, it allows for the development and use of the wealth of control theoretical tools based on feedback linearization [47], impedance control [29], energy shaping and damping injection [18], and many others.

While feedback control has been incredibly successful for robotics systems in the industrial context, so far, it has had only limited success when transferred to open world applications. One of the main reasons is the underlying assumption that a good enough model of the plant is available, or at least, there is a reliable way to obtain it with sufficient precision, e.g., through observers or system identification. However, for robotic systems, which by definition should physically interact with their surroundings, the plant also includes the environment. While this is not a problem in controlled work spaces such as factory floors, it is a significant challenge in open world scenarios. As a result, a lot of work is put into simplifying the environments of industrial robots. This includes, for example, mounting robot arms on a fixed base, avoiding potential interference from co-workers through fences, placing objects that need to be picked up reliably in the same location and orientation, etc.

Unfortunately, this approach is not transferable to robots that should work in our living spaces or in natural environments. Such environments are dominated by complex dynamics, characterized by continuously changing interactions (humans, pets, etc.), dynamic uncertainties (soft/elastic contacts, unexpected impacts, noise, etc.), and unstructured environmental conditions that are often not fully observable. This means that the corresponding control problem is much more complex than what is typically encountered in an industrial context. As a result, despite the promises of generations of roboticists, the dream of robot butlers and other useful sophisticated robotics helpers have not been realized yet.

Interestingly, biological systems, which have to deal with the exact same challenging environments, seemingly, don't have any problems with this complexity. Through the evolutionary process nature was able to find remarkably effective and efficient solutions to these control problems. As a result, biological systems are outperforming state-of-the-art robots in any task that has to be carried out in unstructured environments.

This observation has inspired roboticists to reconsider the current design approach. The result is a new research field called *soft robotics*, which uses a wide range of materials, often soft, to build intelligent machines (see sidebar). This includes materials like silicone, polymers, hydrogels, as well smart materials that can change physical properties when stimulated. This new wealth of possibilities allows roboticists to mimic biological systems much more closely. However, from a control perspective these materials introduce another challenge, since soft materials often exhibit complex nonlinear dynamics, which are hard to model. As a result, the control of soft robots is challenging and classical control approaches often fail or are not directly applicable. In order to reap the benefits of bio-inspired designs in robots, it's not sufficient to use soft and smart materials, but we also need to develop new control approaches that are able to deal with such complex and highly dynamic bodies. We argue that another bio-inspired principle called *morphological computation* (see sidebar) might be able to provide a good starting point for solutions. Morphological computation proposes to outsource functionality directly to the morphology (i.e., kinematics and dynamics) of the body. This means that complex dynamics can be beneficial, if they implement useful (control) functionality. However, in this context a control framework will need to consider also designing corresponding body morphologies. This raises a number of interesting challenges and requires novel control approaches. This article will introduce the concepts of soft robotics and morphological computation and, with the help of specific examples, provide promising starting points. We hope this paper will not only be a source of information but will also inspire the control community to work on these exciting new challenges in order to help to build the next generation of robotic systems.

Conventional Robots vs. Biological Systems

As a first step, we discuss a few key differences between conventional robotics and biological systems, summarized in Figure 1. In contrast to industrial robots, biological systems (i) can exhibit compliance and nonlinear (uncertain) dynamics, (ii) have a rich coupling within their own body and with the environment, and (iii) show a seamless collaboration between brain and body morphology.

Biological bodies are soft with nonlinear dynamics

Biological systems are surprisingly soft compared to current robotic systems and often exhibit variable nonlinear dynamics. While industrial robots are intentionally built with rigid bodies to simplify the underlying model and control problem, it seems biological systems do not have a need for that. Not only are they soft, but their biological bodies exhibit highly nonlinear dynamics, use noisy sensors, and employ actuation systems (e.g., muscular-skeleton systems) that involve a large number of degrees of freedom (DoF) for even very simple movements. From a control perspective, this is quite counterintuitive as a more complex body typically requires a more complex controller. By taking a closer look, we can see that the nonlinear dynamics in biological bodies are not arbitrary. They play an active role in the emergence of intelligent behavior, actively contributing to the solution of the control problem. For example, nonlinear morphological features in the feet of mountain goats help to stabilise their steps in highly challenging terrains [1]. Similarly, in humans, the complex morphology of the knee helps to control the impact during locomotion [64]. Morphological properties of the cochlea filter out the right frequency band to improve sensing [3] and spiders use their webs' compliance to understand and control their environment [42]. In all these cases, the nonlinear dynamics of the soft body are part of the sensory-control loop and they contribute actively to the control task. It seems this principle of exploiting nonlinear (and often soft) morphologies for control is ubiquitous in nature.

Biological agents are tightly coupled with their environment

A second major difference is that biological systems are tightly coupled with their environment, because their bodies are the result of evolution taking place within an environment. Their morphology is an embodied solution that can solve real-time dynamic problems in a given specific niche. For instance, an octopus performs best in an aquatic environment and it exploits the physical properties of water with its compliant body [58]. Some species of trouts have evolved to exploit passive interaction dynamics between the body and the turbulent water flow to efficiently swim upstream [6], and the insect eye is optimized to efficiently counteract nonlinear visual effects like motion parallax during flying [17]. This coupling is omnipresent in nature; see [50] for more examples. As a result, biological systems are remarkably robust and energy efficient. Conventional robotic design typically neglects this relationship because very often the robot's objective is to maintain a desired temporal behavior rather than working with the context to accomplish a goal. This naturally leads to minimise the influence of the environment. Furthermore, classical robot control design is mostly focused on the agent and does not look actively for synergies with the surroundings. Consequently, control approaches for modern robotic system rely heavily on highly controlled environments and fail in open world scenarios.

The body and the brain are working together in an hierarchical structure

Another distinguishing characteristic of biological systems is the relationship between body and its central nervous system (CNS). In robotics the controller (i.e., the digital implementation of a control algorithm in an electronic device) is conceptually separated from the body. As a consequence, the controller tells the body what to do. The story is, however, more complex in biological systems. On the evolutionary time line, biological systems existed long before the emergence of neural circuits and brains. As a consequence, nature needed to rely on morphological features to successfully interact with the environment. Biological bodies are therefore tuned to solve real-time control problems efficiently without a central controller. As a result, the dynamic interactions between body and environment often give the impression of intelligent behavior. This can be observed at all scales, from DNA replication and protein production [2], to communication and learning in plants [19], to macroscopic behaviors like elegant swimming still observable in dead fish [6].

In addition, evolutionary solutions are known to be not wasteful. Working strategies and morphological designs are not arbitrarily discarded, but rather taken forward. This means the evolutionary developments of dedicated computational systems like neurons are built on top of morphologically intelligent structures. As a consequence, even with humans, who have outstanding cognitive functions, the morphology is still taking over parts of the control task. This can also be observed in the fact that the brain (and other neural circuits) doesn't try to override the body dynamics as is done with high-gain approaches in traditional control, rather it works with the body morphology by exploiting the existing body dynamics. As a result, the body and the brain (neural structures) are tightly intertwined. From a control perspective, the separating line between controller and to-be-controlled is blurred.

Motivated by the outstanding performance of biological systems in unstructured environments, recently, roboticists have taken these three principles more into account. The understanding that nonlinear dynamics of soft bodies are not necessarily bad has led to a newly emerged field of research called soft robotics. In addition, the increased consideration of the body morphology to improve interaction with the environment and to help the central controller in a robot has led to the concept of morphological computation.

The next section provides a short overview of both concepts, i.e., soft robotics and morphological computation. After that, with help of specific examples, we discuss their implications with respect to control and point to a range of interesting research opportunities for the control community.

New design principles: Soft Robotics and Morphological Computation

From a control theoretic perspective, (i) nonlinear dynamics due to soft materials and complex morphologies, (ii) strong dynamic coupling with the environment, and (iii) the intertwining of body and brain, all pose challenges, as they all increase the complexity of the underlying control problem. This is one of the main reasons why in conventional robot and control designs all three characteristics are actively avoided or suppressed. In contrast, biological systems seem to actively embrace these phenomena often misunderstood as disadvantages and employ them for their benefit. It seems that nature exploits these features for controlling, sensing, and even for computation. This principle is often referred to as embodied intelligence [50] or, when used in the robotics context, as morphological computation [24]. For more information we defer to the sidebar on morphological computation. Partly motivated by this, recently, the new research field of soft robotics has emerged [53]. Soft robotics proposes to use a wider range of material to build intelligent systems. In doing so it can exploit morphological features for control, sensing and computation. For more information, see the sidebar on soft robotics.

The possibilities offered by soft robotics also come with a price. Soft materials notoriously exhibit complex dynamics. They are nonlinear and show hysteresis, they are prone to saturation and drift, and their dynamics are inevitably underactuated. As a result, it is non-trivial to obtain reliable models and it is hard to design effective controllers. This is where morphological computation comes into play. It lets us recruit the body dynamics for control functionalities. This means that the seemingly negative properties of soft materials can provide solutions to the inherent control problem in soft bodies. This is exactly what we can observe in natural systems. While biological solutions were found through the evolutionary process, as designers of intelligent machines, we need to understand and be able to formally describe how control and morphology are related in the context of soft robotics and morphological computation.

Since both concepts are directly inspired by nature, we can use the three previously discussed characteristics of biological systems as starting points. First, we need to understand how the dynamics of soft bodies can be useful for control. Second, we need to ask how the design of body morphology can facilitate and improve the interaction with the environment. Finally, we need to understand how a traditional digital controller can work together with the body morphology to achieve a desired behavior.

These are difficult and still open questions that need to be answered if we want soft robotics to succeed. A body of ideas and prototypes has already started to emerge and it provides promising new control approaches that can effectively take advantage of soft bodies within the context of morphological computation. However, for the most part this body of work still lacks the mathematical rigor and a common general design framework is missing. The control theory community has unique tools to translate these novel ideas into useful and widely applicable approaches for controller design. Below we discuss different aspects of how the principle of morphological computation can help to develop novel control approaches for soft robots. We will use examples to illustrate its potential and, we hope, inspire people to further work on them.

The role of morphology in action and perception

To take full advantage of morphological computation for controlling soft robots, we need to rethink what control design means in this context. With morphological computation, control is not just the algorithm that is implemented in a digital controller, but also includes the design of the body morphology. This leads to a number of questions, including how should we design a body to guarantee the desired morphological properties and therefore control functionality? And, how do we decide which morphological properties are relevant for a specific task?

A second group of natural questions concerns the relationship between morphology and digital controller. How should they work together? How should the control authority be distributed between them? The problem is made even more complex by the dependence of interactions between body morphology and digital controller on the environment (consider Figure 1). The challenge is thus to understand how the environment, the body morphology and the digital controller all can work together within a sensory-motor loop, to obtain a desired behavior. Third, related to this is also the question of morphology and perception, that is, how does morphology affect sensing? This ultimately affects how the robot perceives its environment. Finally, as a fourth point, if functionality is embodied in morphological features of soft robotic bodies, how can and should we adapt and reprogram this functionality? The next sections will explore these four topics in more detail supported by corresponding examples.

The design of the body morphology

Morphological computation suggests that control functionality can be outsourced to the nonlinear dynamics of the soft body. In the context of robotics, this means that the body of a robot does not only have a structural role, but also a computational functionality that can be used for control. Take, for example, a leg of a robot. While it is used for supporting the body weight and to propel it forward during locomotion, we can also incorporate compliant parts (e.g., inspired by the tendons around ankle and knee joints) that can work autonomously as a mechanical stabilizer to deal with unknown roughness in the terrain. Therefore, in a well-designed robotic leg, some of the control functionality is taken over by morphological features, which means that the same performance can now be achieved with a simpler digital controller. Unfortunately, we still do not have general design guidelines that can map directly from a desired control solution to a corresponding physical embodiment. However, there are a number of general insights that might help us to establish such a new control framework. First, the more complex the dynamics of the morphology are, potentially, the more complex controllers can be implemented with it. A morphology with a high-dimensional state space can represent more variables and integrate them over time simultaneously. In addition, nonlinear characteristics like hysteresis and switches can be potentially beneficial, as they introduce exploitable memory effects within the robot body. Hysteresis implicitly carries information about the direction of a movement, which could, e.g., be exploited to control joint trajectories for locomotion depending on the phase of the locomotion cycle. Likewise, an implementation of a binary switch between a limit cycle and an equilibrium point would, for example, allow to control the transition between a periodic gait for locomotion and a resting state / reaching movements using extremely simple control strategies.

Example 1: Morphology as a computational resource for control Hauser et al. proposed to exploit networks of nonlinear mass-spring-damper systems as computational resources [24], [25], [59], [62], [23]. This is interesting because biological tissue and compliant bodies of soft robots can be represented by such a mechanical network. The underlying idea is to exploit the nonlinear dynamics of such body morphologies and combine them with a simple machine learning approach to learn in a supervised fashion interesting dynamic behaviors, e.g., a dynamic controller response. The input could be, for example, a time-varying force (e.g., a disturbance) applied somewhere to the mechanical structure and the output could be a readout from the high-dimensional state space of the network to calculate a corresponding actuation signal to counteract the disturbance. The approach is surprisingly powerful. The combination of such a mechanical network and a machine learning based readout is able to represent any set of smooth nonlinear dynamical differential equations with one equilibrium point. This is even true if the machine learning part (i.e., the readout) is only a simple set of linear, static (i.e., memory-less) output weights that combine linearly (some) of the state variables of the body morphology. Note that in such a setup during the supervised learning process only the linear output weights are adapted (see red arrows in Fig. 2(b)) while the networks structure (i.e., the soft robot body morphology) is kept unchanged. First, this means that the mechanical morphology is only exploited as a computational resource, but not necessarily designed as such. Second, since the readout weights are linear, finding the optimal set of weights can be done by using simple linear regression, which is extremely fast. Recurrent artificial networks, which are typically employed to solve the same class of learning problem (i.e., learn to emulate nonlinear dynamical systems), are using much more complex algorithms like back-propagation through time, which are slow in convergence and get often stuck in local minima.

The discussed approach is based on a more general machine learning approach called *reservoir computing* (see sidebar for more information), where a general high-dimensional nonlinear dynamical system, called the *reservoir*, is exploited as a computational resource. In the case of morphological computation, the reservoir is the soft body of the robot itself (see Fig. 2(b)).

An intuitive understanding of the role of the morphology (i.e., the soft body) in reservoir computing and, therefore, general guidelines on how to design such a body can be given as follows. To emulate a nonlinear dynamical system (e.g., a desired controller), the reservoir computing setup (i.e., the combination of the reservoir and the readout) needs to be able to integrate and nonlinearly combine information over time. Since the readout is linear and static, the nonlinear dynamics of the body morphology have to contribute both parts. Moreover, intuitively, if the state space of the dynamical system is high, potentially more complex functionalities can be implemented. A remarkable implication of that is, loosely speaking, potentially more complex types of computation (controllers) can be realized when there is a higher diversity of dynamics in the morphology. For a more in-depth discussion, see the sidebar on reservoir computing and [26].

Another important property of compliant robot morphologies is their inherent stability. In addition to being beneficial in the context of reservoir computing, a mechanical spring-damper interface, for example, can also reduce the bandwidth of contact forces in interaction control, effectively simplifying the design of the feedback controller.

Example 2: Increasing energy-efficiency through compliance It has also been shown that energy efficiency in locomotion can be significantly increased through compliance. One of the most versatile locomotion models, the so-called Spring-Loaded Inverted Pendulum (SLIP) Model [7], has a linear spring at its core, see Figure 3(a). If, in addition, the compliance can be adapted (e.g., through a variable compliance mechanism) to work with changing ground conditions, the energy efficiency can be significantly improved [61]; see Figure 3(b). In addition to the stiffness, a joint angle dependent damping can also improve locomotion efficiency [41] while at the same time reducing the collision force and its variability [21]. Compliance in locomotion can also help to increase robustness. For example, the so-called *Puppy* robot uses only hip motors and an underactuated, passive DoF (i.e., a mechanical spring) at the knee to deal with unknown rough terrain [30]; see Figure 3(c). $\lrcorner$

Example 3: Better grasping through compliance Compliance can also be beneficial in grasping. Take for example, the so-called coffee-balloon gripper [9], see Figure 3(d), which is simply a balloon that is filled with ground coffee. When the air is sucked out, it hardens (due to the so-called *jamming effect*). It becomes soft again when the air is allowed to enter. When the coffee-balloon gripper is in its soft state it can be pressed onto an object. As a result, it deforms and shapes itself around it. When the air is sucked out it grasps the object. This simple setup has been demonstrated to be an impressively general gripper, able to pick up a wide range of objects in size and shape without the need of a detailed model of the object, nor of a complex controller (a simple on/off for the air pump is sufficient). $\lrcorner$

Another important property of soft bodies is underactuation. The presence of difficult-to-control sub-dynamics is a potential resource for realizing functionality in biological systems and soft robots.

Example 4: Compliance for swimming It has been speculated that the undulatory movement of octopus tentacles in the biological animal during swimming is not the result of the full control of the entire arm. Rather, for the most part it is actuated only at the shoulder. The compliant, underactuated tipped arm is simply responding to this stimulation with this movement. The result is a traveling wave that propels the animal forward. The validity of this approach was demonstrated with silicone-based tentacles in [55]. $\lrcorner$

All of these properties, i.e., (i) nonlinear dynamics, (ii) large state space, (iii) compliance, and (iv) underactuation, are able to contribute to control functionality. However, they are all actively suppressed in conventional robotics design while, at the same time, these properties are intrinsic to the materials used to build soft robots. This means that soft morphologies are better suited to implement complex functionality than the rigid bodies of classical industrial robots.

While the examples above show that clever morphology can endow complex functionalities in a (soft) robot, it's not obvious how to design such systems from scratch. Some success in finding useful morphologies has been made using evolutionary algorithms. However, a system-theoretic framework is needed to reduce the fragility of these approaches and to prune the search space to find optimal solutions. We need to understand the design of the body in the wider context of feedback control. Rethinking the morphology through the lens of passivity theory, for example, would allow to develop a modular approach based on the interconnection of simple, thus predictable, components. Likewise, the extension of robust and adaptive control techniques (i.e., adapting morphological instead of control parameters) to soft robotics would ensure resilience of the complex behaviors achieved through the morphology.

The examples of this section illustrate the potential of considering the design of a robot body as part of the overall control design. From a system-theoretic perspective, the controller is not just the feedback law implemented in a digital device but also extends to the robot body. This view also calls for novel feedback principles for soft robot control, to enable a robust and tractable design of the morphology.

Control authority

This section looks at the interplay between digital controller and complex morphology, focusing on examples where the control action is shared among the two.

The more control functionality we can outsource to the body, the simpler the digital controller can be. An extreme case of this approach are so-called *passive walkers*. Their mechanical structure converts the gravitational pull into a stable locomotion behavior on a slope, without the need of any digital controller [11], see Figure 4(a) for an example. While this is an intriguing example, its working range is rather limited since the mechanical structure is optimized for a very specific environment. A slight change in the angle of the slope and the walker doesn't locomote anymore. Also, passive walkers can only counteract very small perturbations and are not very robust in general. This means, getting rid of the digital controller entirely may not be a practical approach.

Conventional robot manipulators represent the other side of this spectrum. Their rigid structure shows that no contribution from the morphology is expected, with the control effort fully provided by the digital controller. However, as argued before, this is one of the main limiting factors for bringing robots into more challenging environments.

The emerging picture from these two contrasting approaches is that sharing control authority between morphology and digital controller can be beneficial. The question is then, how can the workload be best distributed between the two? How do we design them to *work together*? This also means that the dynamics of the body shouldn't be overruled by the controller, but rather exploited. A clear example of the advantage of combining morphology and digital controller is provided by digitally controlled passive walkers, like the Cornell Ranger [11].

Example 5: Localized and timed control to exploit dynamical features The mechanics of the Cornell Ranger (Figure 4(b)) guarantees the existence of a limit cycle (walking motion) within the robot dynamics, like in passive walkers. However, the Cornell Ranger also works on flat ground. It does not harvest energy from the gravitational field but uses a controller, which supplies energy to balance mechanical losses. The controller is not continuously driving the system (i.e., it does not provide a control signal at every control time step as typically done in conventional robot control), but instead it introduces a kick of energy once a locomotion cycle. In doing so, it does not interfere too much with the actual passive dynamics of the system. The resulting robot is remarkably energy efficient. The trade-off between morphology and digital control in the Cornell Ranger is not far from the way humans walk. We let our leg swing and thereby convert potential into kinematic energy without controlling every single degree of freedom (i.e., muscles) during this swing phase. This is one of the main contributing factors why humans outperform humanoid robots in energy efficient walking by far. Both the Cornell Walker and the human body show how energy efficiency can be significantly improved when morphology and digital/brain control work in synergy. ⌟

Another interesting way to distribute the control authority has been described in a series of papers by Corucci et al. [13], [14], [12]. The main idea was that a simple digital control action can be translated by the compliant morphology into useful and stable locomotion behaviors.

Example 6: Switching behavior by switching morphology Corucci et al.'s bio-inspired robot loosely mimics an octopus (Figure 4(c)). The robot has four passive silicone tentacles, driven through a crank-slider mechanism by four rotational motors operating at the same constant speed. The control signal is thus very simple. To optimize the overall locomotion of the robot, Corucci et al. proposed to optimize only the morphological features without changing the existing simple digital control signal. They used genetic algorithms (GA). Specifically, they employed the so-called Novelty Search algorithm [38] which optimizes for novelty (i.e., a gene is ranked higher the more it differs from all the previous designs). In the octopus example, the GA found a large number of combinations of all 24 morphological parameters (including stiffness, center of mass, different body lengths, spring constants, etc. (see Figure 4(d)). The result was a large database of morphologically different robots. Interestingly, different morphologies translated the same (unchanged and not optimized) digital control signal, i.e., the constant rotational movement of the crank-slider mechanism, into different behaviors, e.g., walking, swimming, hopping, etc. It is not difficult to see how the role of the morphology is central in Corucci et al.'s approach. It enables and facilitates the emergence of different locomotion behaviors. Notably, the morphology and the digital controller are acting at different levels of competency. The digital controller still balances energy losses (as in the Cornell Ranger) but also enforces a sort of digital clock to set the rhythm/frequency of the locomotion cycle. The morphology shapes this simple signal into an actual behavior, producing swimming, walking, hopping, etc. One could say, the digital control signal is defining a base limit cycle and the morphology is determining the shape of the limit cycle. Notably, this is similar to the learning framework of Dynamic Movement Primitives proposed by Ijspeert et al. [31] which is used to learn repeating (rhythmic) trajectories. Their dynamic movement primitives are based on synthetic central pattern generators, producing rhythmic patterns (limit cycles), which are then superimposed by a learned nonlinear function that shapes the final rhythmic behavior.

Corucci et al. took their work even further by asking the question, how can we switch between the different behaviors that have been found? One naive way would be to take two morphologies associated with two different behaviors and adapt their parameters from one configuration to another. However, this might require the simultaneous adaptation of a significant number of morphological parameters, which is non-trivial and practically hard to implement. Instead, Corucci et al. looked at morphologies that were close (in parameter space) but still exhibited sharp behavioral differences. Interestingly, they found a number of such morphologies, separated only by a single parameter. They showed that the change of this critical parameter was enough to stably modulate the robot behavior without changing the digital control signal. From a feedback control perspective, this is a question of sensitivity in morphology space. It shows that there are directions in the parameter space associated to the morphology whose variation has a maximum impact on the robot behavior (singularities).

Example 7: Multiple behaviors encoded in morphology The octopus example could also be generalized. A robot morphology could be designed to implement multiple limit cycles, for example, to represent different gaits. This has been demonstrated in simulation with the previously discussed networks of mass-spring-damper systems from Example 1 that uses the concept of reservoir computing [25]. In that case, the input was given by constant (external) forces squeezing the soft structure. Depending on how strong the force was (i.e., three different amplitudes were explored) the system produced three different corresponding limit cycles. While the force in [25] was assumed to be coming from the environment, one could implement a mechanical structure, e.g., a buckling mechanism, to easily switch between the limit cycles. Similarly, one could conceive a morphological design that switches from a limit cycle to an exponentially stable equilibrium point, therefore, switching from locomoting to reaching.

The discussion in this section makes it clear that the co-design of morphology and digital control is central for the realization of controllable complex behaviors in soft robots. The control authority is shared between digital and morphological control and they play different roles. In the particular examples given, the digital controller restores energy, defines the rhythm, and contributes to the overall system behavior. This is mediated by the morphology, which shapes the final locomotion behavior to obtain the desired function, with different degrees of sensitivity.

Perception through morphology

So far, we have only discussed the morphological layer in the context of control tasks by showing how the dynamics of a soft body can be exploited to implement a broad set of computations and nonlinear behaviors. However, another crucial part of the feedback loop is sensing, and morphological features can also be beneficial in this context. The morphology, being a dynamical system itself, can be designed as a physical implementation of a dynamical filter enhancing or suppressing in real-time certain features of physical signals. It can thus improve the sensing performance by carrying out real-time computation on the inflowing data stream and therefore simplifying the interaction with the world.

Example 8: Linearization through nonlinear morphology The eyes of flying insects provide an excellent illustration of this point. The morphological arrangement of the ommatidia is non-homogeneous, with higher density on the side than in the front and back. This nonlinear spatial arrangement is optimized to counteract the nonlinear visual effects of motion parallax during flying [17]; see Figure 5(a). From a functional perspective, this means that the morphology implements a static nonlinear mapping to linearize the stream of data. As a result, in the flying insect, the computation needed to navigate is simpler. ⌟

The mapping implemented in the morphology, however, doesn't have to be static. For example, compliant morphologies can be designed to realize a dynamic filter (including an exponentially fading memory) that transforms the external stimulation into useful information. In biological systems we can observe that the dynamic properties of morphologies are tuned towards specific sensing modalities. However, the same principle can also be used in the design of artificial sensors. The role of the morphology as a dynamical filter is illustrated in the following examples.

Example 9: Signal-processing through nonlinear morphologies Spiderwebs are highly complex dynamical structures. The spider uses a number of different types of silk with different dynamic properties, building highly sophisticated structures with distinct morphological features, as illustrated in Figure 5(b). It has been speculated that the spider uses the web as an external filter [42] to recognise and distinguish the mechanical vibrations introduced by prey, predators and mating partners. Recent results show that even very simplified 3D-printed versions of spiderweb-inspired structures exhibit a rich set of nonlinear dynamics [20]. ⌟

Example 10: Frequency separation through morphology The cochlea, a spiral-shaped membrane of the inner ear, provides another biological example of how morphology can help the sensing process. Due to its shape, the cochlea acts as a passive frequency filter with different locations along it spiral sensitive to different frequencies. Consequently, neurons distributed along the cochlea will be exposed to different frequency bands and their spikes will correspond only to certain frequencies [3]. In a similar fashion, a tapered whisker was used for terrain classification on a mobile robot, where different parts of a sensing whisker (which was a tapered spring) resonated at different frequencies [63]. ⌟

Example 11: Morphological dynamics as full body sensor Judd et al. [33] showed another clever use of the morphology for sensing in robots. The authors of the study moved a silicone-based octopus arm in the vicinity of objects to learn from the associated deformation of the soft body *if* (classification) and *where* (regression) the object was without ever touching it. See Figure 5(c) for an illustration of the setup. The flow of water induced by the movement of the arm is affected by the presence of an object in the vicinity. This generates an *echo*, a feedback through water, which is enough to introduce a measurable deformation of the soft body, picked up by bending sensors along the arm. ⌟

This last example also shows that perception through the morphology doesn't have to be necessarily passive. Exploration through motion and adaptation of morphological parameters (i.e., active perception) can both be used to improve the sensing performance.

Example 12: Adaptive morphology to increase sensing performance The estimation of the depth of a hard nodule by doctors in soft tissue palpation shows the role of motor action and morphology in perception through two key phenomena. The first phenomenon is a process of stiffness control via muscle co-contraction with statistical patterns unique to the depth of the nodule. This was also verified by robotic experiments [57]. Independent verification using a controllable stiffness robotic finger shows that such stiffness control helps to tune the sensor model or the likelihood function in a recursive Bayesian estimation framework [28]. Through this adaptation the morphology optimizes the sensing performance for the depth of the nodule. The second phenomenon is the regulation of palpation behaviors in terms of force and speed of the fingers [35]. This is an extrinsic process to maximise rewards in the task. The efficiency depends on the internal dynamics due to the presence of the nodule. This interplay is usually called behavioral lensing, because it leads to an augmented perception of task relevant states with lower uncertainty, leading to higher rewards. ⌟

Sensing through morphology shows potential to open new directions in perception and interaction control [51], [49], [48]. But this requires a methodological effort to move from prototypes based on clever intuition to design principles that can guide the design of the morphology, to take advantage of active exploration and morphology adaptation.

Adaptive morphology

The previous section illustrated how morphology shapes the interaction between robot and external world. Typically, the design of morphology is performed offline and its features remain consistent throughout the life of the robot (with possible undesired drifts due to material stress/aging and faults). However, as the palpation and the octopus examples have shown, adapting the morphology to a specific task/environment can improve performance [39]. The idea is that control and perception functionalities embodied as morphologies can be made more flexible and resilient to deal with ever-changing environmental and task requirements by endowing the morphology with adaptable features, like variable stiffness, variable directional rigidity, re-shaping abilities, etc. These features find relate to the classic and still fundamental approach of adaptive control. The difference between the use of soft robotics / morphological computation and the traditional approach of adaptive control is that the adaptation is not limited to the digital control side, but extends to the robot morphology [4]. A broad discussion on shape-changing robots to optimize interaction dynamics can be found in [56].

Although the design of the morphology is an established practice in soft robotics, the development of adaptive morphology is still in its infancy. The shared control authority and feedback between morphology and digital controller mean that any adaptation (be it morphological or digital) affects both. Thus, particular attention is needed to preserve key properties like stability or passivity during adaptation while taking into account large uncertainties due to compliant materials/structures and due to the interaction with the environment. The adaptation of the morphology is also subject to practical constraints. A robot cannot be built with the capability to change everything. Therefore, the identification of morphological parameters that have a major effect on the robot behavior is required—similar to the ones found through Genetic Algorithms (GA) in the octopus example. However, instead of using blind nonlinear optimization techniques like GAs, general design frameworks need to be developed that can map specific performance goals to morphological designs. Furthermore, even if a clever adaptive morphology has been devised, we need to find a way to optimize/learn the best adaptation strategy, given the specific environment/task. In this context, machine learning will play a crucial role in improving and optimizing performance. Two examples of clever collaboration between adaptive morphology and control/learning strategies are provided below.

Example 13: Control learning space through adaptive morphology When we learn a new movement, like the tennis swing, we typically stiffen up at the beginning of our training naturally constraining the movement space. Then, when our motor skills have improved, we relax to move more dynamically and refine our skills. This strategy effectively reduces the search space at the beginning of our learning, by constraining it to a subspace of all possible movements [27], [60]. A similar approach could be explored with soft robots, by using smart materials that are capable of changing stiffness in response to a controlled stimulation, e.g., electric current. The initial learning could operate within a reduced motion subspace, characterized by high rigidity/stiffness. Then, the robot could gradually become softer to explore more dynamic motions. ⌟

Example 14: Progressive control learning via growing morphology Similar strategies are used in the context of growing morphologies: growing robots that evolve from simple (low numbers of degrees of freedom) to complex systems by going through growing stages. Learning in the earlier stages is easier due to the constrained parameter space, but hard in the full state-space of a complex robot. However, control policies learned from simpler morphologies can be progressively transferred to more complex growth stages if the morphological differences are not too large. Zhu et al. [66] demonstrated that such an approach can lead to accelerated learning, using a simulated tadpole robot that grew into a frog robot in seven stages. Similar results were obtained by Bongard [8] on a worm-like robot that grew into a quadruped. Their results showed that the controllers found by genetic algorithms and that have gone through the progression of the four growing stages were more robust than those found directly on the last quadruped stage.

In general, there is a large set of interesting learning and design problems at the intersection of adaptive morphology and learning, and we believe that control theory could play a crucial role in connecting both worlds.

Discussion and Future Directions

The field of soft robotics is growing rapidly. In little more than a decade it transformed from a niche exploratory field into one of the biggest trends in robotics. Despite this success, the field is still struggling to demonstrate its full potential. One of the biggest challenges is the inherent control problem. As laid out in this article, the seemingly unfavorable characteristics of soft bodies, i.e., their nonlinear dynamics, can be beneficial when properly exploited as proposed by morphological computation.

Unfortunately, so far, morphological computation is mostly based on a combination of bio-inspiration, intuition and ingenuity. The field lacks a unifying theory, a framework that can produce fundamental design guidelines. Such a theory will also have to include a level of design beyond mechanical systems. The choice to focus on mechanical bodies in this article was deliberate to facilitate the communication of the underlying concepts and ideas. However, our discussion holds true also for general robotic systems, taking advantage of chemical reactions, electromagnetic effects, biological tissues, etc. In that sense, the framework of nonlinear systems offers a very broad perspective and can serve as a basis for a general theory of morphological control. We believe that the control community is uniquely equipped to contribute to such a framework. Control theorists are familiar with systems thinking and feedback, and are trained to think of physical devices as operators for processing, sensing, and control. Our hope is that this article will inspire researchers in the control community to engage with the wide range of interesting control challenges in soft robotics, to help to build the next generation of intelligent robots.

Sidebar: Soft Robotics

Soft robotics is a recently emerged field of research [34], [53]. Partially inspired by nature, soft robotics explores a wide range of new materials to build intelligent systems . While conventional robot designs typically use rigid body parts and high-torque servo motors, soft robotics actively embraces soft materials for building robots. Soft robotics does not intend to replace conventional robotics but proposes to expand the set of building blocks of modern robotics, by considering materials such as silicone, polymers, hydrogels, and numerous smart materials [37]. Soft robotics also considers partially-rigid structures like tensegrity designs, foldable robots, origami-based mechanisms and many others [36]. The result is a wealth of new types of robots with a great potential to solve problems that rigid robots are struggling with. The inherent compliance in soft robots translates into the ability to increase safety in human-robot interaction, improve the handling of delicate produce, e.g., in the context of harvesting fruit or manipulating bakery goods, facilitate non-invasive surgery [10], boost energy efficiency in locomotion, enrich haptic interfaces, and many other aspects [52].

Sidebar: Morphological Computation

Morphological computation is a principle in robotics inspired by natural observations that show how morphological features of biological systems play a crucial role in the emergence of intelligent behavior. It suggests that biological systems use their bodies and their corresponding morphological features to implement useful functionality. Loosely speaking, biological systems devolve functionality to their physical bodies. Consequently, the seemingly complex task to control such morphologically challenging bodies interacting in noisy and partially unknown environments can be solved with rather simple controllers, since part of the needed computation/control takes place in the morphology.

Researchers have used this principle to control robots in a wide range of contexts including sensing [5], [57], [28], energy-efficient locomotion [11], [61], pure computational tasks [44], and many others. For a good overview of examples we refer the reader to [50].

Sidebar: Reservoir Computing

Reservoir computing is a machine learning framework that allows to learn nonlinear and dynamic computational mappings in a supervised fashion. This means, it can be used to learn to emulate nonlinear (smooth) dynamical systems. Traditionally, recurrent neural networks (RNN) are used for this class of tasks. However, learning algorithms for RNNs, like backpropagation through time, are prone to get stuck in local minima and are known to converge slowly. Reservoir computing is a powerful alternative.

At the core of reservoir computing is a high dimensional nonlinear system, i.e., the *reservoir* (compare Figure 2(a)), which is exploited as a computational resource. Traditionally, the reservoir is a randomly initialised network of nodes. The connection are simple linear weights (randomly chosen from a given distribution) and the nodes are one-dimensional exponentially stable differential equations with a nonlinear output (with randomly chosen time constants from a given distribution). The network forms a high-dimensional nonlinear dynamical system. A low-dimensional input (e.g., a time series $\mathbf{u}(t)$) will be propagated through this network producing a dynamic response of the reservoir in the form of state trajectories. These states are then linearly combined with a set of learned output weights to produce a desired target response (see Figure 2(a)).

Learning in reservoir computing Learning happens in a supervised fashion. A target system (i.e., a nonlinear dynamical system with a desired behavior) that represents a target behavior (mapping) is provided. The learning data is given in form of a time series of inputs $\mathbf{u}(t)$ and a corresponding time series of outputs $\mathbf{y}(t)$. The goal of the learning process is to find the optimal set of output weights $\mathbf{w}^*$ (red arrows) that map the states of the reservoir onto the desired target output. This means that when these optimal weights are found, the reservoir computing should emulate the target system (e.g., the controller) as closely as possible. Practically, the learning is done as follows: The reservoir is fed with the input $\mathbf{u}(t)$ and the states trajectories of the reservoir are collected in a large matrix $\mathbf{M}$. The learning is then to find the optimal weights $\mathbf{w}^*$ that linearly combine all states at all time steps to produce the desired target $\mathbf{y}(t)$ output as close as possible (reducing the quadratic error). The optimal weights $\mathbf{w}^*$ can be calculated through linear regression, i.e., $\mathbf{w}^* = \mathbf{M}^+\mathbf{y}$, with $\mathbf{M}^+$ being the Moore-Penrose pseudoinverse.

Physical reservoir computing Originally, the concept of reservoir computing was introduced independently by two groups under different names. Maass et al. [40] called it Liquid State Machines (LSM) and Jaeger et al. [32] introduced the concept under Echo State Networks (ESN). LSM was inspired by the computational power of the brain and used therefore (artificial) networks of spiking neurons to construct a reservoir. ESN, on the other hand, was driven by machine learning frameworks and, hence, used linear connections of simple, nonlinear integration nodes. Later both approaches were systematically combined under the name of reservoir computing as they both shared the same underlying principles [54].

As the origins of reservoir computing demonstrate the concept of a reservoir is quite general and there are different ways to implement a reservoir. It has been shown that only a few, quite broad characteristics are needed for a dynamical system to be computationally useful as a reservoir. It needs a high-dimensional state space, exponentially stable, nonlinear dynamics, and some form of integration with fading memory. For a detailed discussion on how these abstract properties (state space, nonlinearity and memory) map onto concrete physical properties in soft robots and their implication on the design we refer to [26].

Interestingly, many real-world physical systems do have exactly these kinds of properties. This implies that reservoir computing can be implemented with the help of real physical systems and not just as abstract computational models in computers. The first work demonstrating that used a bucket of water as a reservoir [16]. In robotics Hauser et al. showed that networks of mass-spring damper systems (which is a common way to describe soft bodies) can be used as a reservoir as well [24], [25]. Since then the field of *physical reservoir computing* has emerged [43] and a wide range of physical systems have been proposed and used for reservoirs. This includes the exploitation of nonlinear effects in lasers, memristors, quantum effects and many others. Inspired by the mass-spring-damper work by Hauser et al. physical reservoir computing has also been employed in numerous robots (see [22] for a review) ranging from artificial octopus arms that can compute [44], [46], [45], to controlling the locomotion of quadruped robots [65], and to follow trajectories with a pneumatically driven robot arm [15].

References

[1] S.A. Abad, N. Sornkarn, and T. Nanayakkara. The role of morphological computation of the goat hoof in slip reduction. In *IEEE International Conference on Intelligent Robots and Systems*, pages 5599–5605, 2016.

[2] L. M. Adleman. Molecular computation of solutions to combinatorial problems. *Science*, 266(5187):1021–1024, November 1994. Publisher: American Association for the Advancement of Science.

[3] H. Ammari and B. Davies. Mimicking the active cochlea with a fluid-coupled array of subwavelength Hopf resonators. *Proceedings of the Royal Society A*, 476(2234), 2020.

[4] R. Baines, S. Freeman, F. Fish, and R. Kramer-Bottiglio. Variable stiffness morphing limb for amphibious legged robots inspired by chelonian environmental adaptations. *Bioinspiration & Biomimetics*, 15(2):025002, 2020.

[5] Y. Bar-Cohen. Biological senses as inspiring model for biomimetic sensors. *IEEE Sensors Journal*, 11(12):3194–3201, 2011.

[6] D.N. Beal, F.S. Hover, M.S. Triantafyllou, J.C. Liao, and G.V. Lauder. Passive propulsion in vortex wakes. *Journal of Fluid Mechanics*, 549:385, 2006.

[7] R. Blickhan. The spring-mass model for running and hopping. *Journal of Biomechanics*, 22(11–12):1217–1227, 1989.

[8] J. Bongard. Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences*, 108(4):1234–1239, 2011.

[9] E. Brown, N. Rodenberg, J. Amend, A. Mozeika, E. Steltz, M. R. Zakin, H. Lipson, and H. M. Jaeger. Universal robotic gripper based on the jamming of granular material.

Proceedings of the National Academy of Sciences, 107(44):18809–18814, 2010.

[10] M. Cianchetti, C. Laschi, A. Menciassi, and P. Dario. Biomedical applications of soft robotics. *Nature Reviews Materials*, 3(6):143–153, June 2018. Number: 6 Publisher: Nature Publishing Group.

[11] S. Collins, A. Ruina, R. Tedrake, and M. Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085, 2005.

[12] F. Corucci, M. Calisti, H. Hauser, and C. Laschi. Evolutionary discovery of self-stabilized dynamic gaits for a soft underwater legged robot. In *Advanced Robotics (ICAR), 2015 International Conference on*, pages 337–344, 2015.

[13] F. Corucci, M. Calisti, H. Hauser, and C. Laschi. Novelty-based evolutionary design of morphing underwater robots. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, GECCO '15, pages 145–152, New York, NY, USA, 2015. ACM.

[14] F. Corucci, M. Calisti, H. Hauser, and C. Laschi. Shaping the body to shape the behavior: a more active role of the morphology in the brain-body trade-off. *13th European Conference on Artificial Life (ECAL2015), Late Breaking Proceedings*, pages 7–8, 2015.

[15] M. Eder, F. Hisch, and H. Hauser. Morphological computation-based control of a modular, pneumatically driven, soft robotic arm. *Advanced Robotics*, 32(7):375–385, 2017.

[16] C. Fernando and S. Sojakka. Pattern recognition in a bucket. In W. Banzhaf, J. Ziegler, T. Christaller, P. Dittrich, and J. T. Kim, editors, *Advances in Artificial Life SE - 63*, volume 2801 of *Lecture Notes in Computer Science*, pages 588–597. Springer Berlin Heidelberg, 2003.

[17] N. Franceschini, J.M. Pichon, C. Blanes, and J.M. Brady. From insect vision to robot vision. *Philosophical Transactions Biological Sciences*, 337(1281):283–294, 1992.

[18] E. Franco, A.G. Casanovas, and A. Donaire. Energy shaping control with integral action for soft continuum manipulators. *Mechanism and Machine Theory*, 158:104250, 2021.

[19] M. Gagliano, V.V. Vyazovskiy, A.A. Borbély, M. Grimonprez, and M. Depczynski. Learning by association in plants. *Scientific Reports*, 2016.

[20] M.S. Garrad, L. Renson, and H. Hauser. Nonlinear dynamics of a synthetic spider web. In *NODYCON 2021 Proceedings*, Switzerland, September 2020. Springer, Cham.

[21] E. Hamid, N. Herzig, S. Abad, and T. Nanayakkara. A state-dependent damping method to reduce collision force and its variability. *IEEE Robotics and Automation Letters*, 6(2):3025–3032, 2021.

[22] H. Hauser. Physical reservoir computing in robotics. *Natural Computing Series*, pages 169–190, 2021.

[23] H. Hauser and G. Griesbacher. Moving a robot arm by exploiting its complex compliant morphology. In R. Pfeifer, S. Hidenobu, R.M Füchslin, H. Hauser, K. Nakajima, and

S. Miyashita, editors, *Proceedings of the 2nd International Conference on Morphological Computation*, 2011.

[24] H. Hauser, A.J. Ijspeert, R.M. Füchslin, R. Pfeifer, and W. Maass. Towards a theoretical foundation for morphological computation with compliant bodies. *Biological Cybernetics*, 105(5-6):355–370, 2011.

[25] H. Hauser, A.J. Ijspeert, R.M. Füchslin, R. Pfeifer, W. Maass, R.M. Füchslin, R. Pfeifer, and W. Maass. The role of feedback in morphological computation with compliant bodies. *Biological Cybernetics*, 106(10):595–613, 2012.

[26] H. Hauser, K. Nakajima, and R.M. Füchslin. Morphological computation — the body as a computational resource. In *E-book on Opinions and Outlooks on Morphological Computation*, chapter 20, pages 226–244. 2014.

[27] J.B. Heald, D.W. Franklin, and D.M. Wolpert. Increasing muscle co-contraction speeds up internal model acquisition during dynamic motor learning. *Scientific Reports*, 8(1):16355, 2018.

[28] N. Herzig, P. Maiolino, F. Iida, and T. Nanayakkara. A variable stiffness robotic probe for soft tissue palpation. *IEEE Robotics and Automation Letters*, 3(2):1168–1175, 2018.

[29] N. Hogan. Impedance control: An approach to manipulation. In *1984 American Control Conference*, pages 304–313, 1984.

[30] F. Iida and R. Pfeifer. Sensing through body dynamics. *Robotics and Autonomous Systems*, 54(8):631–640, 2006.

[31] A.J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural Computation*, 25(2):328–73, feb 2013.

[32] H. Jaeger. Adaptive nonlinear system identification with echo state networks. in S. T. S. Becker and K. Obermayer, eds, ”Advances in Neural Information Processing Systems 15”, MIT Press, Cambridge, MA, pp. 593–600., 2003.

[33] E. Judd, G. Soter, J. Rossiter, and H. Hauser. Sensing through the body – non-contact object localisation using morphological computation. In *RoboSoft 2019 - 2019 IEEE International Conference on Soft Robotics*, 2019.

[34] S. Kim, C. Laschi, and B. Trimmer. Soft robotics: A bioinspired evolution in robotics. *Trends in Biotechnology*, 31(5):287–294, 2013. ISBN: 0167-7799.

[35] J. Konstantinova, M. Li, G. Mehra, P. Dasgupta, K. Althoefer, and T. Nanayakkara. Behavioral characteristics of manual palpation to localize hard nodules in soft tissues. *IEEE Transactions on Biomedical Engineering*, 61(6):1651–1659, 2014.

[36] C. Laschi, B. Mazzolai, and M. Cianchetti. Soft robotics: Technologies and systems pushing the boundaries of robot abilities. *Science Robotics*, 1(1):eaah3690, 2016. ISBN: 2470-9476.

[37] C. Laschi and R.J. Wood. Smarter materials for smarter robots. *Science Robotics*, 6(53):eabh4443, April 2021. Publisher: Science Robotics.

[38] J. Lehman and K.O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.

[39] L. Lichtensteiger. The need to adapt and its implications for embodiment. In *Embodied Artificial Intelligence*, pages 98–106. Springer, 2004.

[40] W. Maass, T. Natschlaeger, and H. Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560, 2002.

[41] M.F.E. Morland, K. Althoefer, and T. Nanayakkara. Novel method to form adaptive internal impedance profiles in walkers. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 7764–7767. IEEE, 2015.

[42] B. Mortimer, A. Soler, L. Wilkins, and F. Vollrath. Decoding the locational information in the orb web vibrations of araneus diadematus and zygiella x-notata. *Journal of the Royal Society Interface*, 2019.

[43] K. Nakajima. Physical reservoir computing—an introductory perspective. *Japanese Journal of Applied Physics*, 59(6):060501, 2020.

[44] K. Nakajima, H. Hauser, T. Li, and R. Pfeifer. Information processing via physical soft body. *Scientific Reports*, 5:10487, 2015.

[45] K. Nakajima, H. Hauser, T. Li, and R. Pfeifer. Exploiting the dynamics of soft materials for machine learning. *Soft Robotics*, 2018.

[46] K. Nakajima, T. Li, H. Hauser, and R. Pfeifer. Exploiting short-term memory in soft body dynamics as a computational resource. *Journal of The Royal Society Interface*, 11(100):20140437, 2014.

[47] G. Palli, C. Melchiorri, and A. De Luca. On the feedback linearization of robots with variable joint stiffness. In *2008 IEEE International Conference on Robotics and Automation*, pages 1753–1759. IEEE, 2008.

[48] N. Pestell, T. Griffith, and N. F. Lepora. Artificial SA-I and RA-I afferents for tactile sensing of ridges and gratings. *Journal of the Royal Society Interface*, 19(189), 2022. arXiv: 2107.02084 Publisher: The Royal Society.

[49] N. Pestell and N. F. Lepora. Artificial SA-I, RA-I and RA-II/vibrotactile afferents for tactile sensing of texture. *Journal of the Royal Society Interface*, 19(189), 2022. arXiv: 2204.02475 Publisher: The Royal Society.

[50] R. Pfeifer and J.C Bongard. *How the body shapes the way we think*. The MIT Press, 2006.

[51] Q. Qi and V.A Ho. Wrinkled soft sensor with variable afferent morphology: Case of bending actuation. *IEEE Robotics and Automation Letters*, 5(3):4102–4109, July 2020.

[52] J. Rossiter and H. Hauser. Soft robotics - The next industrial revolution? *IEEE Robotics and Automation Magazine*, 23(3), 2016.

[53] D. Rus and M.T. Tolley. Design, fabrication and control of soft robots. *Nature*, 521(7553):467–475, 2015.

[54] B. Schrauwen, D. Verstraeten, and J. Van Campenhout. An overview of reservoir computing: theory, applications and implementations. In *Proceedings of the 15th European Symposium on Artificial Neural Networks*, pages 471–482, 2007.

[55] M Sfakiotakis, A Kazakidi, and DP Tsakiris. Octopus-inspired multi-arm robotic swimming. *Bioinspiration & Biomimetics*, 10(3):035005, 2015.

[56] D. Shah, B. Yang, S. Kriegman, M. Levin, J. Bongard, and R. Kramer-Bottiglio. Shape changing robots: bioinspiration, simulation, and physical realization. *Advanced Materials*, 33(19):2002882, 2021.

[57] N. Sornkarn, P. Dasgupta, and T. Nanayakkara. Morphological computation of haptic perception of a controllable stiffness probe. *PloS one*, 11(6):e0156982, 2016.

[58] G. Sumbre, G. Fiorito, T. Flash, and B. Hochner. Motor control of flexible octopus arms. *Nature*, 433(7026):595–596, 2005.

[59] H. Sumioka, H. Hauser, and R. Pfeifer. Computation with mechanically coupled springs for compliant robots. In *IEEE International Conference on Intelligent Robots and Systems*, pages 4168–4173, 2011.

[60] M.H. Vincken and J.J. Denier van der Gon. Stiffness as a control variable in motor performance. *Human Movement Science*, 4(4):307–319, 1985.

[61] H. Vu Quy, H. Hauser, D. Leach, and R. Pfeifer. A variable stiffness mechanism for improving energy efficiency of a planar single-legged hopping robot. In *Advanced Robotics (ICAR), 2013 16th International Conference on*, pages 1–7, nov 2013.

[62] Y. Yamanaka, T. Yaguchi, K. Nakajima, and H. Hauser. Mass-spring damper array as a mechanical medium for computation. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 11141 LNCS, pages 781–794, 2018.

[63] Z. Yu, S. Perera, H. Hauser, P.R.N. Childs, and T. Nanayakkara. A tapered whisker-based physical reservoir computing system for mobile robot terrain identification in unstructured environments. *IEEE Robotics and Automation Letters*, 7(2):3608–3615, 2022.

[64] L.Q. Zhang, G. Nuber, J. Butler, M. Bowen, and W.Z. Rymer. In vivo human knee joint dynamic properties as functions of muscle contraction and joint position. *Journal of Biomechanics*, 31(1):71–76, November 1997.

[65] Q. Zhao, K. Nakajima, H. Sumioka, H. Hauser, and R. Pfeifer. Spine dynamics as a computational resource in spine-driven quadruped locomotion. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013.

[66] Y. Zhu, J. Rossiter, and H. Hauser. Learning in growing robots: Knowledge transfer from tadpole to frog robot. In *Proceedings of Biomimetic and Biohybrid Systems*, 2019.

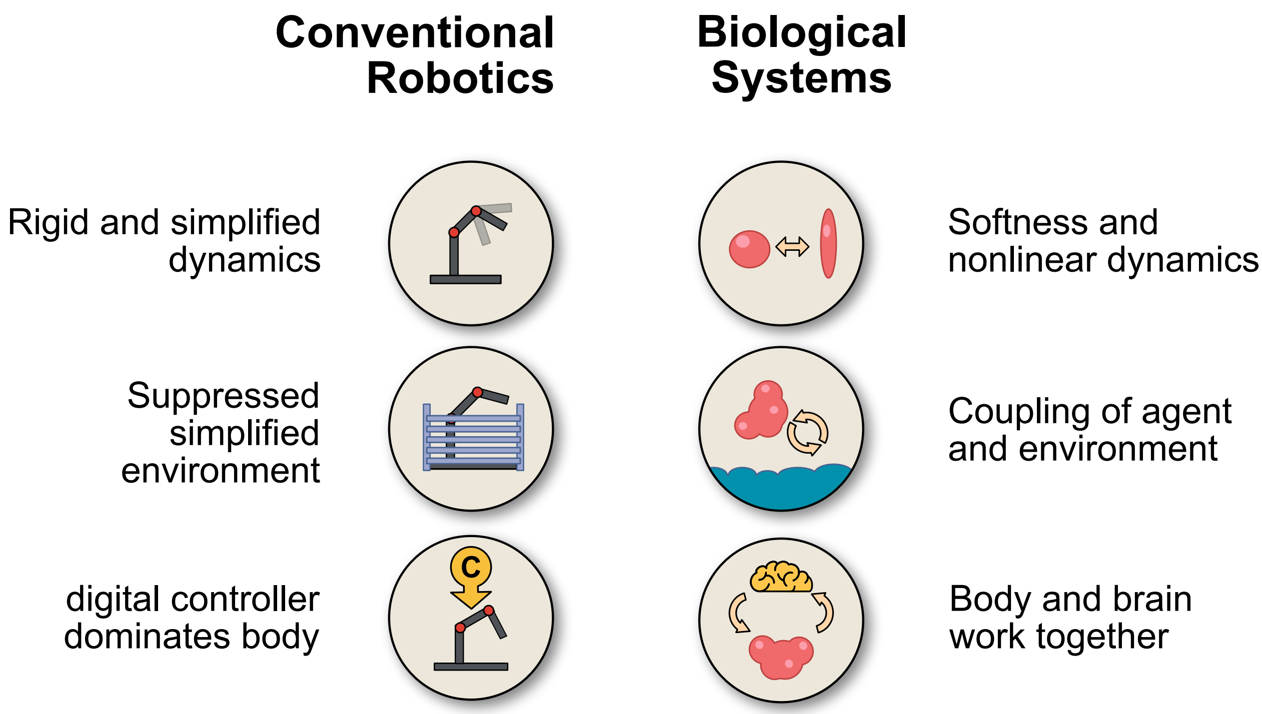

Figure 1: Three main characteristics of biological systems compare to industrial robots. Biological bodies, (i) exhibit nonlinear dynamics and are compliant, (ii) are tightly coupled with their environment and, (iii) show a close synergistic collaboration between brain and body morphology.

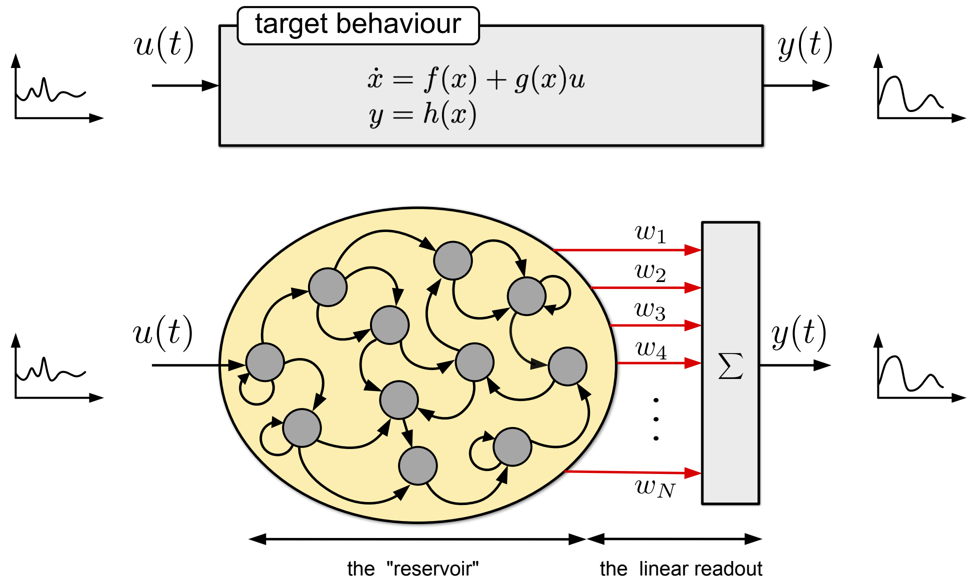

(a)

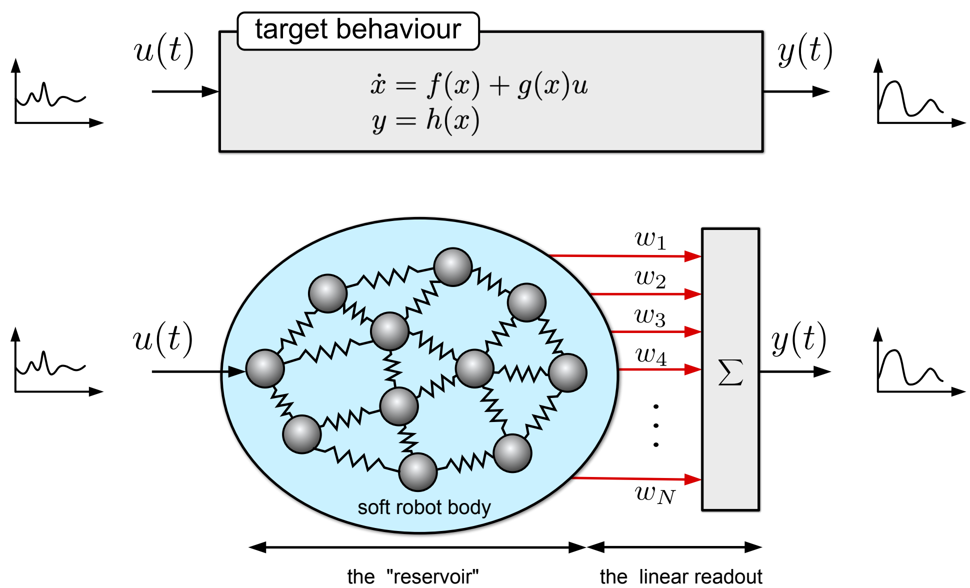

(b)

Figure 2: Principles of reservoir computing. The goal is to emulate a desired target behavior that maps an input $u(t)$ onto a desired target output $y(t)$. This means if the reservoir computing setup receives the same input $u(t)$ it should produce the same output $y(t)$ (or as close as possible). Such a setup can learn to emulate behaviors that can be described by smooth, nonlinear dynamical systems as described in the target behavior box. (a) Standard reservoir computing setup. Simple computational nodes are randomly connected to build a high-dimensional nonlinear dynamical system, i.e., the *reservoir*. The output is a weighted sum of the N-dimensional state space of the reservoir. Only the readout weights (red arrows) are adapted during learning and are found with simple linear regression. (b) Example of a physical system to implement a reservoir. In this case, a randomly connected network of mass-spring-damper systems (which is a good representation of the bodies of soft robots) is used. Again, only the readout from the state space (red arrows) is adapted during learning to emulate a desired behavior.

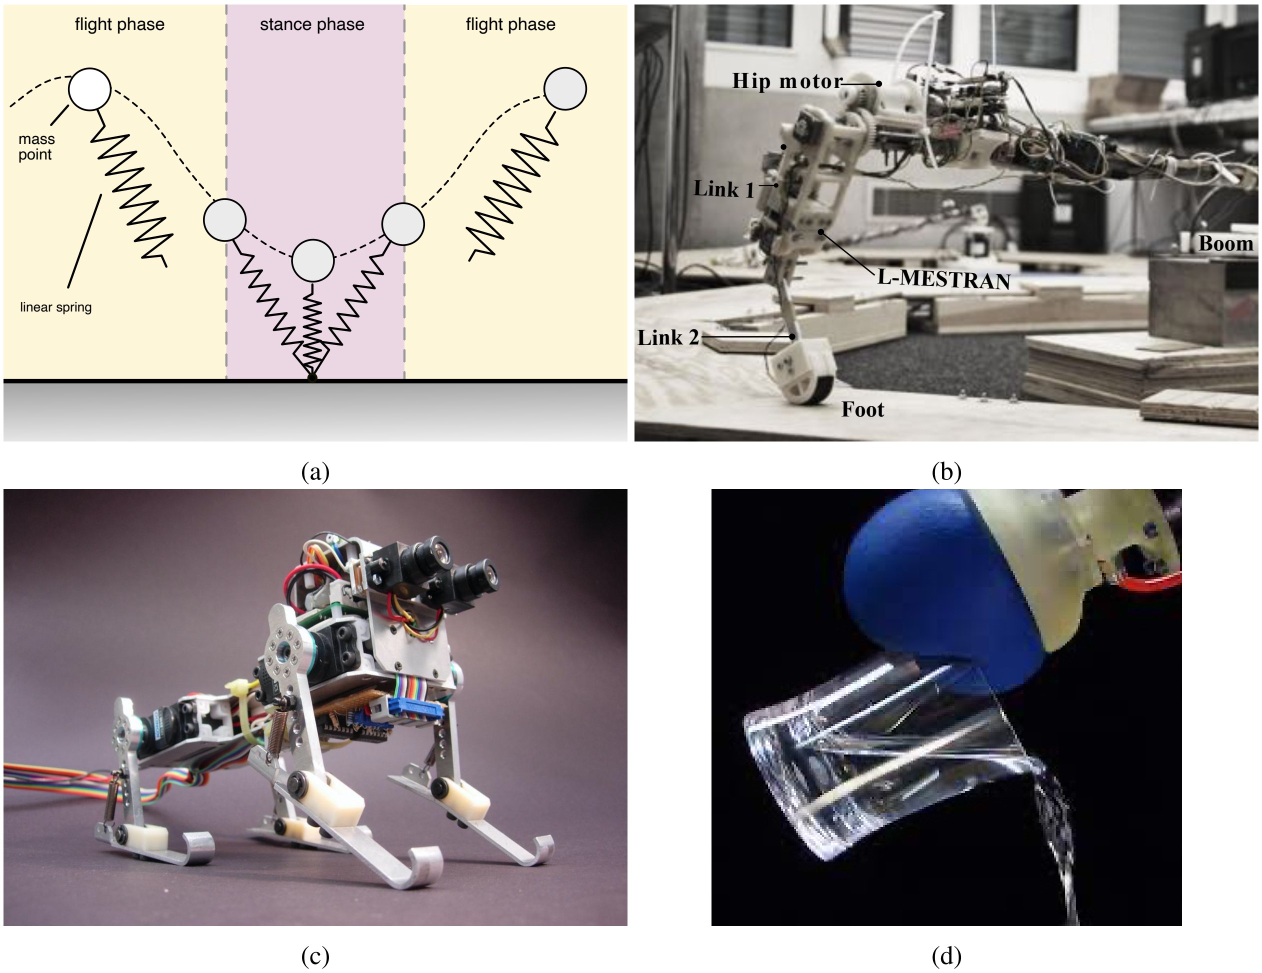

Figure 3: Examples for designing morphologies for control. (a) The Spring Loaded Inverted Pendulum (SLIP) model is a fundamental model describing a wide range of biological locomotion [7]. (b) Hopper with changeable stiffness. Picture taken from [61]. Adapting the leg stiffness enabled the leg to hop energy-efficiently over a range of different surface stiffnesses. (c) The Puppy robot (picture taken from [30]) has only 1 active DoF per leg and a passive spring at the knee joint. The robot's locomotion is highly robust against rough terrain. (d) The coffee-balloon gripper is a highly versatile gripper [9]. A balloon filled with ground coffee can be transitioned between soft and stiff by sucking out the air or letting it in. In the soft mode the balloon conforms to the unknown object (e.g., the glass in the picture). When the air is sucked out, it can pick it up. Back to the soft mode (air in) it releases the object.

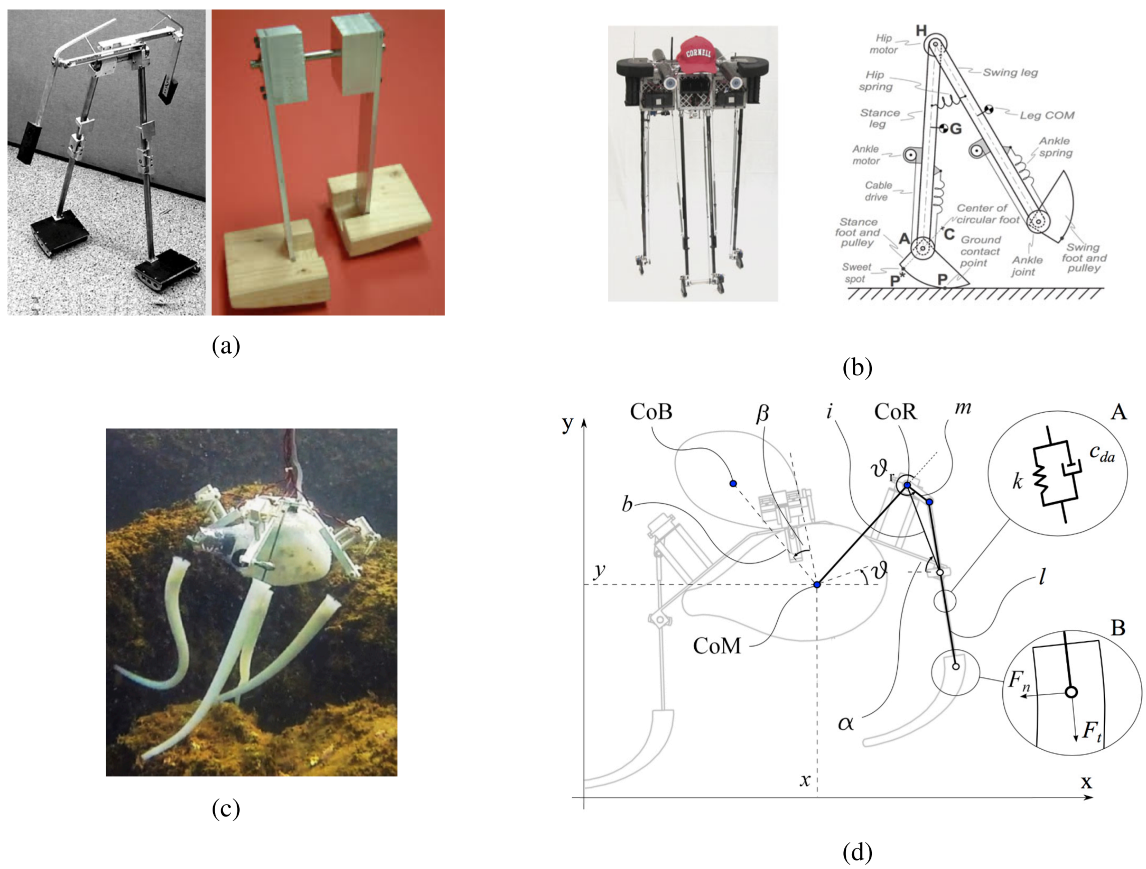

Figure 4: Examples to demonstrate shared control authority between morphological and digital control. (a) Example of passive walkers [6]. The mechanical design (morphology) implements a stable walking cycle on a slope. There is no actuation (digital control), only morphological control, (b) The Cornell Ranger [11] is an extension of the passive walker concept, which can walk on flat ground. The controller only provides input once per locomotion cycle resulting in highly efficient locomotion. (c) Octopus robot used as a starting point for Corucci et al.'s work [13]. (d) Morphological parameters used to optimize the body of the octopus [12].

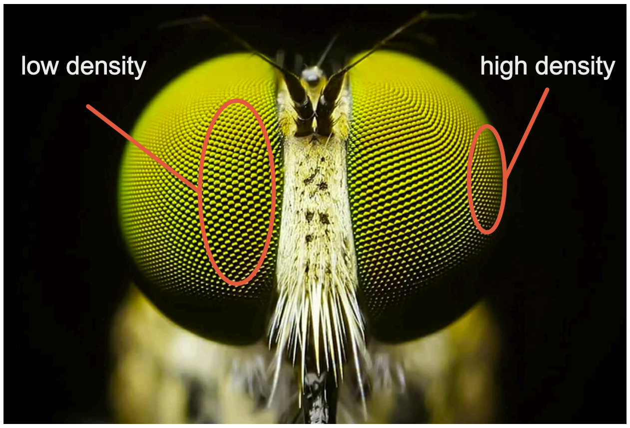

(a)

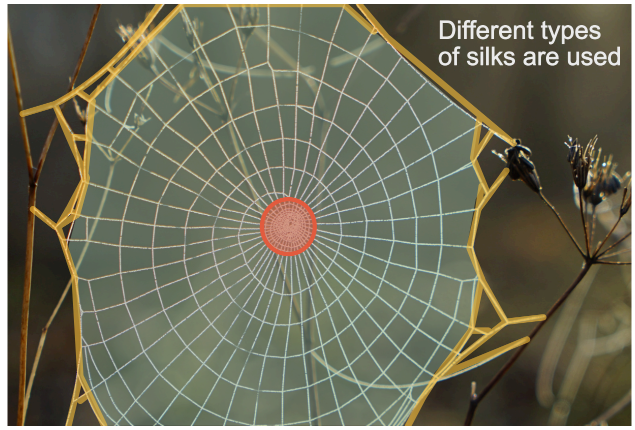

(b)

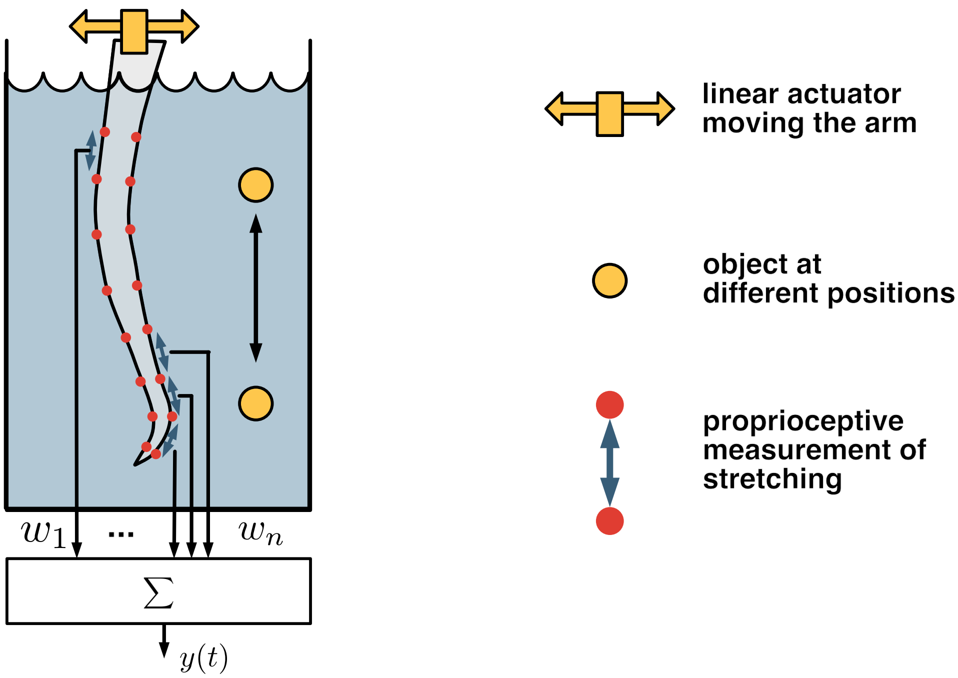

(c)

Figure 5: Examples for morphologies for perception. (a) non-homogeneous morphological arrangement of ommatidia in an insect eye compensates for the nonlinear motion parallax effect. The density at the front is higher than at the side, (b) Spider webs are highly complex and dynamic structures that are exploited by spiders as (vibration) signal processing devices [42]. (c) Implementation example of physical reservoir computing for perception. The moving soft (silicone-based) octopus arm was able to detect objects (yellow) in its vicinity without touching them by only observing its propioreceptive state (i.e., level of bending along the body by measuring distance between red points). This is a physical reservoir computing setup as depicted in Figure 2(b).