Engineered bioinspired natural dynamics and their synergy with control and learning in legged robots

Von der Fakultät Bau- und Umweltingenieurwissenschaften der Universität Stuttgart
zur Erlangung der Würde eines Doktor-Ingenieurs (Dr.-Ing.)
genehmigte Abhandlung

Vorgelegt von

Felix Ruppert

aus Ludwigsburg

Hauptberichter:	**Prof. Dr. Syn Schmitt**
1. Mitberichter:	**Prof. Dr. Renaud Ronsse**
2. Mitberichter:	**Dr. Alexander Badri-Spröwitz**

Tag der mündlichen Prüfung: 5.5.2022

Dynamic Locomotion Group
Max Planck Institute for Intelligent Systems

Institute for Modelling and Simulation of Biomechanical Systems
der Universität Stuttgart

2022

Declaration / Erklärung

I hereby declare that the presented work was accomplished independently and only by the use of the specified resources and literature.

Ich erkläre hiermit, dass die vorgestellte Arbeit unabhängig und nur unter Verwendung der angegebenen Ressourcen und Literatur durchgeführt wurde.

Felix Ruppert

Stuttgart, 13th December 2021

Summary

The performance of legged locomotion relies on the successful mitigation of unstructured, rough terrain in the presence of sparse information and neurosensory delays. Bioinspired walking systems benefit from carefully engineered passive compliant behavior that models the inherent elastic behavior of muscle-tendon structures in animals. To leverage the passive behavior that provides energy efficiency, passive stability as well as simplified control and learning tasks to the system, locomotion control and learning algorithms have to be designed and coordinated with the natural system dynamics in mind to achieve similar locomotion behavior we see in animals.

The major contribution of this thesis is the synergy of a bio-inspired leg design with biarticular muscle-tendon structures, a wearable force and pressure sensor design for closed-loop control in legged locomotion, a biologically inspired closed-loop central pattern generator with reflex-like feedback and a learning approach that enables the locomotion controller to leverage the carefully engineered natural dynamics of the robot to learn convincing locomotion skills and increase energy efficiency.

The first contribution is a biologically inspired leg design focusing on the biarticular lower leg muscle-tendon structure in vertebrate animals. The biarticular elasticity provides two-dimensional passive impedance to the leg and allows the storage of energy orthogonal to the leg axis direction. The leg blueprint is characterized in its capability to store and release elastic energy in the biarticular structure. The stored energy can be recuperated back into the system and increases the energy efficiency of the leg. This leg design achieves the lowest relative cost of transport documented for all dynamically hopping and running robots.

The second contribution introduces the concept of training wheels, temporary mechanical modifications to the system dynamics that shape the learning reward landscape and simplify learning locomotion directly in hardware. Through deliberate changes to the system dynamics, in this case, reduced gravity, the reward landscape can be shaped to simplify the learning process. Learning with this training wheel is safer due to smoother reward landscapes with shallower gradients. Also, the initial guess for initiating the machine learning algorithm is simplified, because the salient gradient set of viable reward signals is bigger. During the learning process, the training wheel influence can be gradually reduced and the learning algorithm converges to the solution of the initial learning problem without training wheels.

The third contribution presents a rugged, lightweight force and pressure sensor for feedback information and biomechanical analysis. The sensor can be mounted on a robotic foot and provides continuous force and pressure feedback during locomotion in unstructured and soft terrain. The sensor is based on a pressure sensor, encapsulated in a polyurethane dome with an air cavity around the pressure sensor. External forces deform the sensor dome and the rising pressure in the air cavity is measured by the pressure sensor. Based on the dome geometry and material, the sensor range can be adjusted for different load cases. The sensor can be used in arrays to measure pressure distributions as well as a wearable force sensor in wet or granular media where classical force plates can not be utilized.

The final contribution synergizes the individual contributions into a neuro-inspired learning approach that matches a bioinspired closed-loop central pattern generator with reflex-like neuroelastic feedback to the natural dynamics of a quadruped robot with biarticular legs. Through sparse contact feedback from the foot sensor, the bioinspired central pattern generator can neuroelastically mitigate short-term perturbations to adapt the robot to its environment. Because the robot dynamics and the control task dynamics initially do not match, the controller uses the neuroelastic feedback to minimize the discrepancy between commanded and measured robot behavior. The amount of feedback activity during level walking can be used as a proxy to estimate the amount of dynamics mismatching. By minimizing the amount of required neuroelastic feedback the robot learns to neuroplastically match its control task dynamics to its natural dynamics through Bayesian optimization.

Through the synergy of mechanics and control the biomechatronic system benefits from both the individual functionality of its components as well as their interplay during locomotion. The designed natural dynamics provide advantageous passive behavior to the robot and the bioinspired controller learns to leverage the natural dynamics to achieve convincing locomotion.

Zusammenfassung

Die Fortbewegung auf Beinen beruht auf der erfolgreichen Bewältigung von unstrukturiertem und unwegsamem Gelände mit der Herausforderung nur spärlich abtastbarer Sensorinformation und neurosensorischen Verzögerungen. Biologisch inspirierte Laufmaschinen profitieren von sorgfältig entworfenem, passiv elastischem Verhalten, welches das inhärente elastische Verhalten des natürlichen Muskel-Sehnen Apparats modelliert. Um das passive Verhalten, welches Energieeffizienz, passive Stabilität sowie ein vereinfachtes Regeln und Lernen ermöglicht, wirksam einzusetzen, müssen Regler und Lernalgorithmen unter Berücksichtigung der Eigendynamik des Systems entworfen und aufeinander abgestimmt werden um zur natürlichen Fortbewegung vergleichbares Verhalten zu erreichen.

Der Hauptbeitrag dieser These ist die Synergie eines biologisch inspirierten Beindesigns mit biartikulären Muskel-Sehnen Strukturen, eines tragbaren Kraft- und Drucksensors zur Regelung, eines biologisch inspirierten, geregelten zentralen Muskelgenerators mit reflexartigem Feedback und eines Lernansatzes, welcher es dem Regler erlaubt, die sorgfältig entworfene Eigendynamik des Roboters wirksam einzusetzen um überzeugende Fortbewegungsmuster zu erlernen und den Wirkungsgrad des Systems zu erhöhen.

Der erste Beitrag ist ein biologisch inspiriertes Beindesign mit Fokus auf die biartikuläre untere Beinmuskulatur von Wirbeltieren. Die biartikuläre Elastizität verleiht dem Bein zweidimensionale mechanische Impedanz und erlaubt die Speicherung von Energie orthogonal zur Längsachse des Beins. Die Blaupause des Beinaufbaus wird in Bezug auf die Speicher- und Rückführfähigkeit von potentieller Energie charakterisiert. Die in den elastischen Elementen gespeicherte Energie kann ins System zurückgeführt werden und erhöht so den Wirkungsgrad des Beins. Dieses Beindesign erzielt den niedrigsten dokumentierten relative Fortbewegungswirkungsgrad (Cost of Transport) für dynamisch hüpfende und rennende Roboter.

Der zweite Beitrag führt das Konzept von Stützrädern, temporären mechanischen Modifikationen der Systemdynamik ein, welche die Lernbelohnungsfunktion formen und das Lernen von Fortbewegung direkt in der Hardware erleichtern. Durch gezielte Anpassungen der Systemdynamik, in diesem Fall durch Änderung der wirkenden Schwerkraft, vergrößert sich die Menge der Regelparameter, in der brauchbare Belohnungssignale abgetastet werden können. Lernen mit Hilfe dieser Stützräder ist durch die ebenere Belohnungslandschaft mit flacheren Gradienten zum Einen sicherer, zum Anderen ist die Wahl der Startbedingungen durch die

größere Menge von Parametern, die brauchbare Erfolgssignale erzeugen, einfacher. Während des Lernens kann der Einfluss der Stützräder langsam reduziert werden und der Lernalgorithmus konvergiert zur Lösung des ursprünglichen Lernproblems ohne Stützräder.

Der dritte Beitrag stellt einen robusten, leichten Kraft- und Drucksensor für Feedback bei der Fortbewegung sowie biomechanische Analysen vor. Der Sensor kann an einem Roboterfuß montiert werden und liefert kontinuierliche Kraft- und Druckmessungen während des Laufens in unwegsamem und weichem Gelände. Der Sensor basiert auf einem Drucksensor, der in einer Polyurethankuppel mit einem Hohlraum um den Drucksensor eingebettet ist. Durch externe Kräfte verformt sich die Sensorkuppel und der Drucksensor registriert den ansteigenden Druck im Hohlraum. Je nach Kuppelgeometrie und -material kann der Sensorbereich für verschiedene Lastfälle angepasst werden. Der Sensor kann sowohl in Sensoranordnungen zur Messung von Druckverteilungen als auch als tragbarer Kraftsensor in nassen oder körnigen Medien verwendet werden, in welchen klassische Kraftmessplatten nicht verwendet werden können.

Der letzte Beitrag führt die einzelnen Beiträge zu einem neurowissenschaftlich inspirierten Lernansatz zusammen, der einen bioinspirierten Closed-Loop-Mustergenerator mit reflexartigem neuroelastischem Feedback an die Eigendynamik eines vierbeinigen Roboters mit biartikulären Beinen anpasst. Durch spärliches Kontaktfeedback des Fußsensors kann der biologisch inspirierte zentrale Mustergenerator kurzfristige Störungen neuroelastisch ausgleichen, um den Roboter an seine Umgebung anzupassen. Da die Eigendynamik des Roboters und die Dynamik der Regelungsaufgabe anfangs nicht übereinstimmen, nutzt der Lernansatz die neuroelastische Rückkopplung, um die Diskrepanz zwischen dem gewünschten und dem gemessenen Roboterverhalten langfristig zu minimieren. Die Stärke der Rückkopplungsaktivität während des Gehens auf ebenem Grund wird hierbei als Maß der Diskrepanz zwischen Regel- und Eigendynamiken verwendet. Durch die Minimierung des erforderlichen neuroelastischen Feedbacks erlernt der Roboter neuroplastisch die Anpassung der Dynamik seiner Steuerungsaufgabe an seine Eigendynamik durch Bayessche Optimierung.

Durch die Synergie der entworfenen, natürlichen Eigendynamik, die dem Roboter ein vorteilhaftes passives Verhalten verleiht, mit einer biologisch inspirierten Regelung, die lernt, die Eigendynamik wirksam zu nutzen, profitiert das biomechatronische System sowohl von der individuellen Funktionalität seiner Komponenten als auch von deren Zusammenspiel während der beinigen Fortbewegung.

Acknowledgment

First of all, I want to thank Alex for the opportunity, support and feedback to build my own legged robot, my initial motivation in high school to study mechatronics.
I would also like to thank the IMPRS-IS for their administrative and scientific support.
I want to thank the DLG team for sharing the last four years of research, lunches and laughs with me. I want to thank the students I got to supervise, who each had a part in the work I present here. And I want to especially thank my lab buddy Alborz for all the chats, discussions, laughs and dances in the lab and for the ideas and solutions we came up with together.
I am thankful for my thesis advisory committee, Syn, Ludo, Jonathan and Daniel who went above and beyond to support me with feedback, ideas and open doors to talk whenever I needed help.
I also want to thank Bernard, Jutta, everyone at robotics ZWE, both mechanical workshops and everyone who had time for a hallway chat that made the last four years worthwhile.
I am incredibly grateful for my awesome family and friends who accompanied me during this whole journey, who supported me during hard times and celebrated successes with me.
And last I want to thank Jule for the endless emotional support, energy, encouragement, distractions, nerves and love she gave me over the last five years.

Contents

Chapter 1

Introduction

> *Imagine standing on top of a flight of stairs in the dark. I tell you it will be 14 steps down. You start walking down the stairs, but it is actually 15 steps. Everyone has experienced this feeling of missing a step and almost falling. And the startled realization: 'That could have gone wrong'. But it didn't. And not only did it not go wrong, you only realize you almost fell afterwards. I want to understand what caused us not to fall even though we couldn't have reacted in time.*

Over the last four years, this has been my one-minute explanation, whenever somebody asked what my Ph.D. topic is. Breaking this story down exposes the current research topics that are relevant in legged locomotion research. A quick calculation of control frequencies in humans reveals that with a nerve speed of $120\,\frac{m}{s}$ [1], [2] in a 2 m tall human the maximum frequency of signals reaching the brain is 60 Hz or a travel time of 17 ms. The frequency is cut in half for a full control cycle foot-brain-foot[1]. On a 0.2 m high step in the staircase, the time between missing the step and the delayed impact is roughly 100 ms, while the human reaction time is ≈ 250 ms [3], [4]. So we really can not actively react neither fast enough nor with a high enough control bandwidth to prevent failure. This means something else has to be responsible for us not falling. In the following chapter I will give an introduction to the topic of legged locomotion and highlight the components of a walking system that contribute to answering the question of what helps us not to fall down the stairs.
From an engineering perspective, legged locomotion is a daunting task. No accurate impact models exist that can reliably and fast estimate the forces exerted onto different substrates for motor control [5]. Muscle properties as well as the leg kinematics are highly nonlinear [6], [7] and can even vary on different limbs of the same animal. Additionally, the terrain conditions in nature are very difficult to handle and are oftentimes unknown before the foot touches the ground. Outside of a laboratory, the ground can be slippery, soft, granular, uneven and cluttered with debris. Yet animals have perfected traversing the most challenging terrains even though, like us humans, they should suffer from the same restrictions and challenges I just mentioned.

[1] This is an underestimation assuming only one bit of data to illustrate the point.

Compared to that, information transfer in robotics is not limited by ion diffusion as it is in neural axons [1]. In electric circuits, information and sensor signals are sampled and shared in kilo- or even gigahertz bandwidth. Additionally, information processing is faster which leads to shorter reaction times [8] and actuators are more torque dense [9] and energy-efficient than muscles [10]. Dealing with the challenges of locomotion is, however, not purely a function of brute computational power and bandwidth. Even with all the technological advantages just mentioned, the most sophisticated legged robots are outperformed by a baby giraffe mere hours after birth.
The ability to locomote gracefully and effortlessly and with lower metabolic energy requirements [11] does, therefore, not come from pure superior technical performance. This hints towards, that the mechanical design, as well as the interplay of mechanics, control and neural processing play a much more important role. In the following chapter I will go into detail about the different aspects that dictate the performance of a walking system from an interdisciplinary biomechatronic design perspective.

1.1 Mechanics

If we break down my introductory story, the underlying effect we see governing the system's behavior is the impedance behavior of the leg [12]. Mechanical impedance describes the resistance of a mechanical structure to motion induced by an external force [13]. From a general equation of motion, the mechanical impedance can be derived in Laplace space as:

$$\begin{aligned} M\ddot{x} + C\dot{x} + Kx &= F \\ (Ms + C + Ks^{-1})\dot{x} &= F(s) \\ Z(s) \cdot \dot{x} &= F(s) \end{aligned}$$

where Z describes the mechanical impedance, M is mass, C is the damping rate, K is the stiffness and F are external forces. While impedance also describes the characteristics of damping, for this thesis I will primarily focus on the stiffness characteristics of legged systems.
In legged locomotion, as in interaction task with an environment in different disciplines like manipulation research [14], the central concept for mechanical interactions is to control the interaction force as well as the motion of the system at the same time [14]–[16]. Controlling the motion and the interaction force provides a compliant[2] manipulator or leg.
A compliant leg is able to deal with uncertainties inherent to interactions with the real world. If for example, the ground during locomotion is not as even or rigid as expected the compliance in the leg will mitigate the unexpected impact forces and stabilize the leg. Because the impact is not hard, but is cushioned through the compliant leg, interactions with the environment are safer and more robust because the negative effects uncertainties would have on a rigid actuator are reduced.
To model the impedance of a mechanical structure it can be described through

[2] Compliance is the inverse of stiffness

a mass-spring-damper system [17]. The stiffness and damping rates define how the system reacts to external forces.
Two general ways of achieving impedance can be distinguished. Either through mechanical components like mechanical springs and dampers or through an actuator and controller that creates a 'virtual' impedance [14].

1.1.1 Passive Mechanics

In locomotion studies, passive elastic elements were discovered in the limbs of animals, that act like mechanical springs. Specifically, the muscle-tendon structures in the leg provide elastic behavior [18]–[21]. This elastic behavior allows the storage and release of energy over a step cycle to increase energy efficiency. Additionally, studies in birds, where running guinea fowls broke through a piece of paper covering a step-down without falling, revealed in detail [22], [23] what I described in my story. The passive mechanical behavior of our muscles and tendons plays a very important role in providing animals with mechanical stability. Energy inserted into the system through a perturbation will lead to a change in momentum in an unelastic system that can destabilize the system. In a system with elasticities the excess energy can be stored in the elastic element. The instantaneous change in energy can then be absorbed and the change in system energy can be mitigated over a longer time during the discharging of the elastic element while the system remains stable (Figure 1.2). In scenarios, where control strategies could not react in time, due to the sensory delay described above, the intrinsic passive behavior of the muscles and tendons takes over the task of mitigating perturbations. The English expression '*a spring in one's step*' as well as the German '*abfedern*' (*engl.* to cushion) also point in the same direction of what Blickhahn et al. [24] found in human walking and running studies. Human running can be described as a point mass bouncing on a spring, the spring-loaded inverted pendulum (SLIP) model [25]. SLIP describes the relationship between the center of pressure and the center of mass (COM) motion of a running human, where the legs are described as massless springs and all mass is concentrated in a point mass (Figure 1.1 top).
The SLIP model is described by its equation of motion that only depends on the physical characteristics of the system [34].

Flight Dynamics

$$\begin{bmatrix} \ddot{x} \\ \ddot{y} \end{bmatrix} = \begin{bmatrix} 0 \\ -g \end{bmatrix}$$

Stance Dynamics

$$\begin{bmatrix} \ddot{x} \\ \ddot{y} \end{bmatrix} = \frac{k}{m}(l - l_0) \begin{bmatrix} \sin(\alpha) \\ \cos(\alpha) \end{bmatrix} - \begin{bmatrix} 0 \\ -g \end{bmatrix}$$

Leg angle

$$\alpha = \arctan2(\frac{y}{x}) - \frac{\pi}{2}$$

Leg length

$$l = \sqrt{(x^2 + y^2)}$$

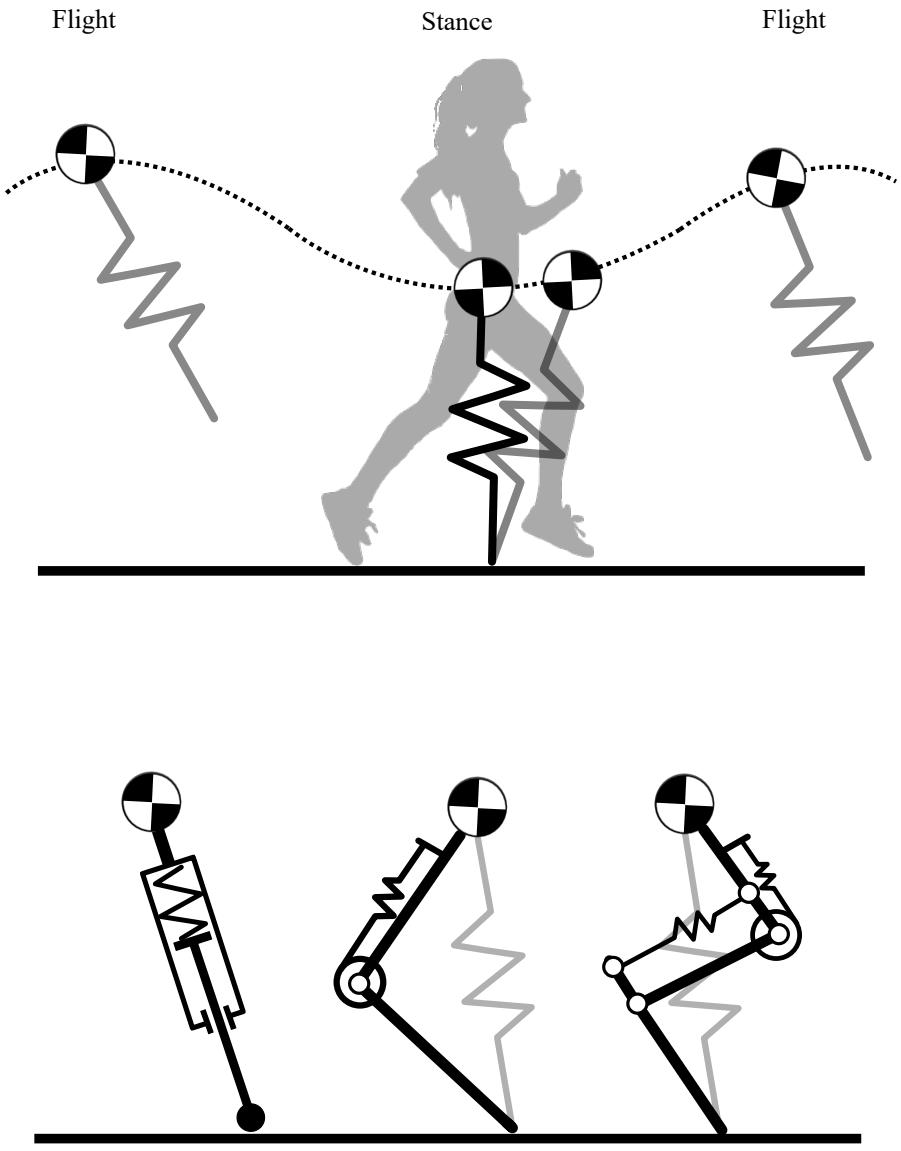

Figure 1.1: The spring loaded inverted pendulum (SLIP) is a template used in biomechanics to describe human running. The model describes the center of mass motion of the human as a point mass, the leg impedance is modeled as a massless spring. The SLIP model can accurately estimate the ground reaction forces produced during human running. In legged robotics, SLIP is used as a template for the design of robotic legs. Based on the concept of SLIP inspired impedance, different leg designs have been proposed. Pogo-stick like leg designs (bottom left) closely follow the SLIP template [26]. Two-segmented legs (bottom middle) are currently used by many research groups [27]–[30]. More biomimetic leg designs (bottom right) have more similarities with their natural role models but still capture the essential features of the SLIP template [31]–[33].

where x and y are the body position, g is the gravitational constant, m is the body mass, l is the leg length, α is the leg angle and k is the leg stiffness. As is visible in the equations above, the dynamic hopping behavior of the SLIP model only depends on the stiffness and mass of the system and does not rely on active control to exhibit hopping. Through the choice of appropriate leg angles during touchdown, the SLIP model entrains on a limit cycle and exhibits periodic hopping behavior. Because the SLIP model exhibits the same spring-mass-damper behavior, previously described in section 1.1, it can walk on uneven ground and mitigate step-down perturbations [25].
Another important aspect of spring-loaded leg designs is their ability to store energy. During the swing phase, the leg falls towards the ground. After impact, the elastic elements capture the kinetic energy from falling and store it until the second half of the stance phase. The stored energy is used to propel the leg forwards for the next step. By storing and recuperating energy that would be lost without elasticity, the legged system requires less energy to locomote and its energy efficiency increases [35]. Additionally, the elastic behavior reduces the impact stress on joints and actuators in both animals and robots [36]. In studies comparing the energy efficiency of animals and robots [11], animals show a higher energy efficiency compared to most state-of-the-art robots [33], [37].
As mentioned above, these springs provide stability and robustness to external perturbations like the missed step in my story. In analytical studies of the SLIP model, Geyer et al. [38] showed, that the SLIP model exhibits self-stabilizing behavior because of the leg's passive elasticity. Daley et al. showed the same stabilizing behavior in animals through their intrinsic passive elasticity provided by the muscle-tendon structures in bird legs [39], [40]. Fischer et al. [41], [42] showed, that the leg morphology[3], specifically the tri-segmentation of femur, shank and foot, and the leg posture, also contribute to the stabilizing performance of legs [43]. The segmentation as well as the scaling of length and mass also influence the resulting leg morphology [44], [45] and performance of walking systems. Even the placement of actuators [46] and the tendon routing [47] inside the leg have an influence on the robustness and efficiency of locomotion.
These findings underline the importance of the mechanical design of the leg and how the different mechanical characteristics have to be carefully designed to enable effective locomotion. These passive characteristics do not depend on control, but only on the physical properties of the leg.

1.1.2 Robotics

Based on the idea of passive impedance captured in the inverted pendulum and SLIP model template [19] it is possible to design robots and simulation models that are able to walk in the absence of control. However, they still capture important aspects of legged locomotion in nature. McGeer's passive walkers [48] are unactuated mechanisms that can walk down a flat slope, converting only the potential energy of the slope into kinetic energy for locomotion.
And the beneficial properties of mechanics can be extended into other parameters that describe a leg's morphology. The segmentation ratio of femur, shank and foot segments has an influence on the performance and energy requirement of a leg [49] due to changing effective lever arms [50]. Where theoretically only one

[3]Study of the form and structure

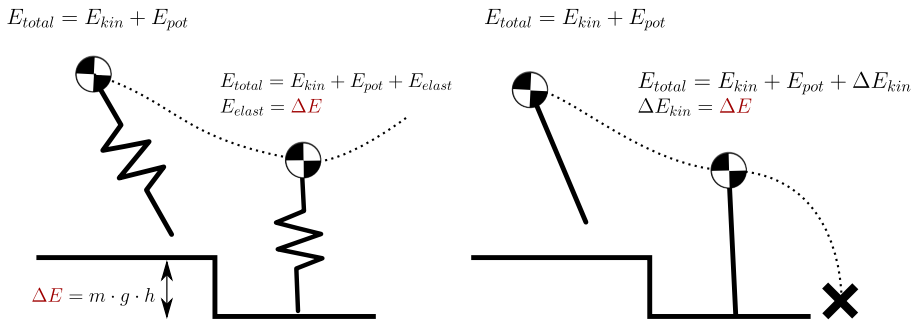

Figure 1.2: In case of a perturbation that injects energy (ΔE) into the system an elastic leg design (left) can store the additional energy in its elastic elements. The energy does not directly accelerate the system which would destabilize the system. An inelastic leg design (right) will become unstable if additional energy is injected into the system because the energy can only be converted into kinetic energy (ΔE_{kin})

degree of freedom along the leg axis is required to lift the leg and create ground clearance during the swing phase, studies have shown that leg morphologies with three segments decrease torque requirements, due to smaller effective lever arms [51], as well as increase stability [38]. The placement of motors influences the coupling through biarticulation[4] [49] as well as antagonistic coupling [52] between actuators. The leg configuration and the placement of elasticities determine the global impedance. Depending on their placement, the elasticities can capture impacts from different angles to extend the advantages shown in reduced-order models [16].

Before high-power computation and controller hardware were as inexpensive and accessible as today, legged locomotion research focused on overcoming these 'neural' shortcomings through mechanical design.

Early research into legged robotics closely followed the SLIP template by the use of pneumatic actuators. The Raibert hopper combines a pneumatic spring with the ability to inject energy back into the system [53] shown in Figure 1.1 bottom left. With a low-complexity and intuitive control design Raibert et al. [26] built machines that were able to walk and resembled human-like gaits both qualitatively as well as in quantitative data.

The ATRIAS robot closely follows the SLIP template in its leg design and control structure and achieves remarkable locomotion behavior without external sensors that resembles the stability of the SLIP model as well as the energy efficiency of human walking [16], [54], [55].

More biomimetic robots like cheetah-cub [31], Oncilla [32], Tekken [56] and SCOUT [57] use mechanical springs in their legs to provide impedance parallel to the respective joint actuator (see Figure 1.1 bottom right). StarlETH [58] and its successor ANYmal [28] use series elastic actuators (SEA) [36] to provide the robot with compliant leg behavior. HyQ uses hydraulic actuators comparable to the Raibert hopper, but in two-segmented legs [30].

Through advances in processing power of microcontrollers and the availability of inexpensive high torque-density motors, it is now possible to achieve leg

[4] muscle-tendon structures spanning more than one joint

impedance without mechanical springs. Contrary to mechanical springs, Seok et al. use motors to mimic the force-deflection behavior of mechanical springs. To achieve the same torque output as a mechanical spring, motors with a high gap radius are needed [15]. These motors, typically used in quadrocopters, also reduce the requirements for high gearing [59] to achieve the torques required for legged locomotion.
With control frequencies in the kilohertz range and communication in the megahertz range, it is possible to provide virtual impedance to a fully actuated leg design without mechanical springs that are able to perform the most versatile locomotion behavior presented to date [55], [60], [61].
Both mechanical and virtual impedance have advantages and disadvantages [62]. A physical spring does not require computation or energy to provide elastic behavior, the mechanical response to excitation is without delay and all the energy stored is released back within the margins of friction losses. The properties of the spring, however, need to be precisely designed to fit into the narrow margins where these properties are advantageous to locomotion [63]. Another disadvantage of physical springs is that once implemented the properties are fixed and can not be changed without additional design effort [64], [65].
This also leads to the problem, that in parts of a stride the motors will work against the spring. When flexing a knee joint with a gravity compensating knee spring, the motor will have to lift the leg and compress the spring which requires additional energy. By designing mechanical impedance and, therefore, altering the passive behavior of the leg, the leg becomes less versatile, because different tasks profit from different passive characteristics which are hardcoded in the mechanical design of the leg.
Virtual spring control, on the other hand, allows to vary the impedance of the virtual spring since the torque generated by the motor can be changed in the control algorithm at runtime. By changing the spring properties, for example, during the swing phase, the motor does not works against the virtual spring. It is, however, disadvantageous that the motor constantly requires energy and active control to provide the gravity compensating elasticity when the leg is on the ground. Also, the virtual spring controller relies on sensors to measure or estimate force and displacement that have inherent delays. This can make reaction to sudden impacts difficult.
Because of advances in the control algorithms (see subsection 1.2.2) the two-segmented leg design currently adopted by many research groups focusing on locomotion control is a simplification of the leg morphologies we see in nature [27], [29], [66], shown in Figure 1.1 bottom middle. These designs focus on linear kinematic chains with two leg segments, high torque-density motors [15] and transparent actuators [59]. By design, these robot morphologies lend themselves to be great tools to study control effects while reducing influences from passive mechanics that limit the controllability and complicate the modeling and control task. The biomimetic similarity, meaning similarity to its natural role models, is traded off for a simplified approach to model-based control algorithms.
Contrary to this, leg designs in nature are defined by highly nonlinear transmissions [6] and actuators [7], are subject to the coupling between mono- and biarticular muscle-tendon structures [49], [67]–[69] as well as underactuated joints [70]. Some of the aspects, primarily the redundancy of actuators as well as the miniature design in nature can not be transferred to simplified blueprints in simulation or walking robots. However, Witte et al. emphasize the importance to

capture the major components of morphology in bio- or robomechanical research [71], [72].

In this thesis, I combine advantages from 'classical' robotic design principles with morphologies extracted from biological role models. The components related to control and actuation are designed to provide high control bandwidth and high torque density by use of high gap-radius motors and low reduction gearing as well as high sensing bandwidth. At the same time, mechanical springs in the leg designs capture the passive elastic behavior of muscle-tendon structures in animals. Additionally, the leg morphology features the same segmentation and closed kinematic loops, namely biarticular muscle-tendon structures [71], [73], as animals to investigate their effects in legged locomotion.

1.2 Control

While carefully designed mechanical properties can make a leg design exhibit walking patterns with little to no control, there is a tradeoff in versatility and maneuverability. A passive dynamic walker can walk down the ramp it was designed for with just a push to start the mechanism. However, if we consider a real-world environment with disturbances in form of uneven ground, obstacles or slippery and soft ground the robot needs to react to these perturbations [74]. Additionally, different tasks like standing, jumping, walking and running require different control strategies to achieve optimal performance [75].
Control strategies for legged locomotion can be differentiated into feed forward and feedback approaches as well as model-based and model-free approaches [76]. In the following section, I will focus on model-free feed forward approaches and model-based feedback approaches.

1.2.1 Feed forward Control

Feed forward control describes control strategies that do not rely on sensory information. The controller dictates a desired behavior without knowledge about errors in the state of the control plant. Instead, feed forward controllers are designed with prior knowledge of the system and its expected behavior [17]. Because feed forward controllers do not rely on sensory information they are inexpensive to implement and can be accomplished with lower control frequencies and computational effort compared to feedback controllers [62]. Feed forward control patterns are therefore implemented to understand natural legged locomotion because, as described in my introduction story, they enable locomotion with the same bandwidth and computational restrictions that animals are subject to.

Central Pattern Generators

A feed forward control method found in nature [78] are central pattern generators (CPG) as shown in Figure 1.3. CPGs are neural networks in the central nervous system that are responsible for creating periodic patterns from non-rhythmic inputs without requiring peripheral sensor feedback [79] as shown in studies on lamprays [80] and frogs [81]. In nature, CPGs are found for tasks like breathing,

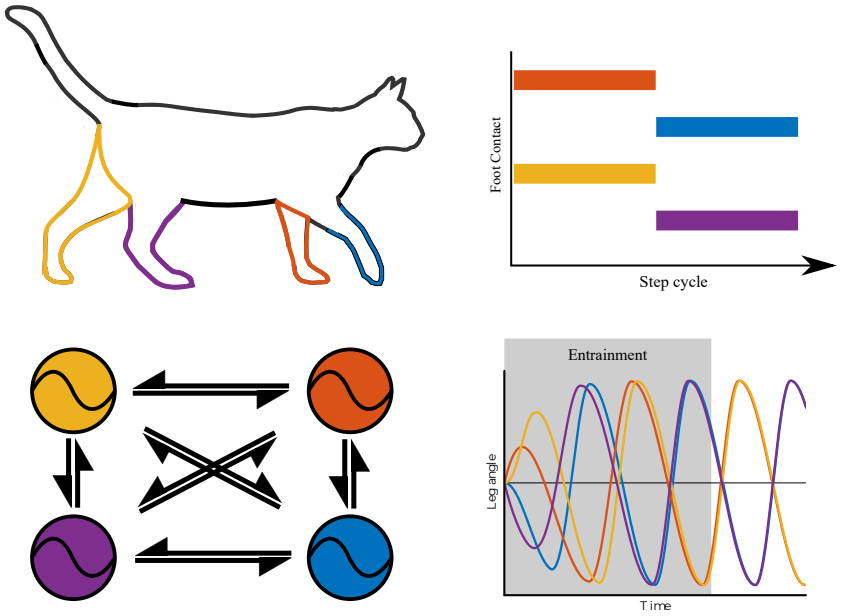

Figure 1.3: Feed forward central pattern generators in animals coordinate periodic patterns. In legged locomotion, each CPG node (circles bottom left) is used to represent the periodic motion of one actuator or leg (top left). The CPG nodes follow a sequence that is predefined by a desired phase difference between the node. Different phase differences result in different gait patterns, in this case, a pace gait (top right). Through coupling in between the nodes (arrows bottom left) the CPG network ensures that the individual nodes follow the desired phase difference. The adaptation of out-of-sync nodes to synchronous motion can be seen in the shaded area in the bottom right (entrainment). The resulting pattern from the CPG can be used as reference trajectories in a legged robot [77].

chewing as well as legged locomotion [82], [83].
These periodic patterns follow a defined sequence of actions. This can be the alternating flexion and extension of the jaw muscles during chewing or the forwards and backwards motion of a leg during locomotion. In a CPG each actuator is represented as a node in the network. The nodes are coupled to one another and exchange their current state in the network. While the external, non-rhythmic input, for example the sequence frequency, only defines the speed at which the sequence is repeated, the coupling in between the nodes defines the timing in between specific actions. If one node is ahead or lags behind in the sequence, the whole network adapts and the speed of all nodes is corrected until the network converges back to its desired behavior. To achieve coupling between nodes different mathematical formulations exist [84]–[87].
In this thesis, I will focus on coupled Hopf oscillators because of their direct representation of joint trajectories, their intuitive design, low dimensional parameter space as well as the possibility to distinguish the stance and flight phases [79], [88]–[90].
During locomotion, the legs of a system swing back and forth periodically in trajectories closely resembling $\cos(\phi)$ functions [91]. In the CPG formulation described in this thesis, the CPG generates the desired joint trajectories for the motors of the robot. The trajectories are described by

$$\Theta_{\text{hip}} = \Theta_{\text{hipOffset}} + \Theta_{\text{hipAmplitude}} \cdot \cos(\phi)$$

where Θ_{hip} is the hip reference trajectory, Θ_{offset} is the hip offset and $\Theta_{\text{hipAmplitude}}$ is the hip amplitude of one CPG node.
Each node of the CPG resembles the oscillatory phase ϕ of one leg. To achieve different gaits, a phase difference between the CPG nodes is introduced. If the phases of both forelegs and both hind legs are symmetric the resulting gait is a bound gait. A phase symmetry between fore-left, hind-right will result in a trot gait and so on. The CPG dynamics and the coupling between nodes are described by:

$$\dot{\phi_j} = 2\pi f + \sum_{k=1}^{N} \alpha_{\text{dyn},j,k} \cdot \mathbf{C}_{jk} \cdot \sin(\phi_k - \phi_j - \mathbf{\Phi}_{jk})$$

where f is the frequency, $\alpha_{\text{dyn},j,k}$ is the conversion constant of the network dynamics between nodes j and k, $\mathbf{C}_{jk}$ is the coupling matrix weight between nodes j and k, $\mathbf{\Phi}_{jk}$ is the desired phase difference matrix value between nodes j and k.
The coupling dynamics act as a proportional controller that ensures, that the desired phase difference in the network is achieved. The angular velocity $\dot{\phi_j}$ of each node is accelerated or decelerated based on the difference between desired and actual phase difference $\phi_j - \mathbf{\Phi}_{jk}$. For a full derivation of the CPG model please refer to the supplementary material of the manuscript of chapter 5.
The generated CPG trajectories are used as angle references for a position controller [31], [32]. The same formulation can also be used for muscle activation patterns [86], [92] that are a more biomimetic resemblance of the CPGs found in nature.
Because feed forward CPGs do not consider the dynamics of the system, their output is identical for each step. This means that feed forward patterns can

be played back on a robot with low computational effort yielding convincing locomotion performance for simple tasks [31]. This performance is the result of the interplay of stability stemming from the passive elastic leg designs and the low computation effort of feed forward control.
But because the patterns are feed forward, they are limited in their usability for real-world applications since they can not react to external perturbations and are limited in the behaviors that can be implemented. Because CPGs are a model-free control approach, their parameters need to be hand-tuned to achieve a desired performance. Since the system does not correct its behavior through sensor information, the CPG trajectories have to be designed to keep the system stable and on a limit cycle. While low dimensional CPGs can be hand-tuned with expert knowledge, the tuning of higher dimensional CPGs requires more sophisticated optimization or learning approaches (see section 1.4).

1.2.2 Feedback Control

Feedback control describes control approaches where the behavior of the system is adapted based on the error between desired and measured behavior [17]. By eliminating this error, the control system follows the desired behavior, can react to external perturbations and actively stabilizes the system. To measure the system's state the controller requires sensory information at a high control and sensing bandwidth to achieve fast reaction time. Feedback control trades off higher computational cost and system integration cost for higher stability and versatility.

Impedance Control

Rather than relying on passive impedance as described in section 1.1, an impedance controller can actively change the impedance behavior of the leg [8], [14]. Through (model-based) inverse dynamics, a set of joint forces and torques is calculated that mitigates perturbations by mimicking passive impedance. Because the impedance can be changed through motor control, the behavior of robots using impedance control is more robust to a wider variety of possible perturbations. The impedance of the leg can also be changed online to adapt to changes in ground properties as well as for different tasks that require different impedance settings.
While impedance controllers increase the versatility of actuators and provide the possibility of online adaptation, the controllers require a high control and sensing bandwidth to mimic a mechanical spring-damper system. Control frequencies are in the range of kilohertz and are, therefore, much higher compared to the aforementioned control frequencies in nature [15], [28]. Because impedance controllers rely on a model, inaccuracies through assumptions and simplifications can have a negative effect on their performance and even production variations between components can lead to problems. Because the impedance behavior is modeled in an actuator, the actuator constantly requires torque during stance for gravity compensation. Physical impedance, in comparison, needs no motor power to hold the robot up.
For an implementation example of a virtual impedance controller please refer to the drop experiments described in the manuscript of chapter 2.

Model-Based Full-Body Control

The control approach that is currently able to achieve the most versatile behaviors is called full-body control. The general idea is that the body position and orientation are described by a centroid, Bledt et al. [93] call it the 'potato model'. A generic shape with the mass and inertia of the robot's body dictates the COM movement, as well as the forces that need to be exerted on this body to reach a goal location. Based on the calculated forces, the legs and joints are coordinated accordingly to achieve the resultant ground reaction forces and leg impedance through inverse dynamics. Through trajectory optimization, the controller finds the optimal way to achieve a desired behavior. A variation of this approach is divergent component of motion (DCM). DCM [94] aims to eliminate force and acceleration components that would diverge the centroid from its desired trajectory.
All full-body controllers rely on a computational model that can be used for trajectory optimization and the inverse dynamics calculations that run in the background of these control algorithms. These algorithms directly observe the system state and coordinate interactions with the environment to steer the system state along a desired state space trajectory.
Most robots that are controlled by full-body control algorithms do not rely on passive elasticities as described above but rather on virtual impedance control. The passive elasticities limit the impedance range since passive stiffness cannot be changed arbitrarily while the robot is running.
While full-body controllers, at least theoretically, have unlimited versatility, they require an accurate model of both the robot and their surroundings to calculate the correct forces to control the dynamics. This comes with a high demand for model accuracy and sensor information. Additionally, the trajectory optimization and inverse dynamics calculations are computationally expensive and require expensive, high-end hardware as well as increased energy [28], [93]. This becomes obvious in an exaggerated example. A full-body controller for a robot walking on even ground will repeatedly do the same calculations for every step of the robot, always reaching the same result and consequently motor trajectories since the motion is periodic. The cost of computation for this repetitive task with an identical outcome is high. Compared to this example, a feed forward controller uses less energy to compute its desired trajectories as it is periodic by design.
In uneven terrain with unexpected perturbations, the full-body controller requires high bandwidth sensor signals that have small sensory delays. Because impacts happen in short time windows [22] even small delays in the millisecond range can be disadvantageous. In comparison, a passively elastic system will react instantaneously and does not suffer from such delays.

1.2.3 Hybrid Feedback Control

To combine the advantages of feed forward and feedback control two approaches are used. Feed forward controllers can be augmented with feedback to make them more robust. Alternatively, feedback signals can be used to trigger feed forward responses.

Feedback CPGs

As described in subsubsection 1.2.1, CPGs are limited in their versatility because they can not react to external perturbations. To be able to react to perturbations, it is necessary to add feedback loops and sensor information to the CPGs to provide a measure of the system state and their environment. The challenge lies in finding appropriate sensory information and reactive measures. Different sensors and feedback pathways like contact information for timing discrepancies [89], body orientation for trajectory adaptation [95], load sensing for gait transitions [96] or exhibitory or inhibitory feedback for gait emergence [56] were described to extend the capabilities of CPGs.
Because CPGs are model-free control approaches, it is important to also feed the system state back into the CPG to match the desired behavior dictated by the CPG to the capabilities and physical properties of the mechanical system. Because CPGs are often formulated in joint angle space it is challenging to find appropriate feedback reactions to external perturbations that exert forces onto the robot. Errors in timing [89] and body pitch corrections [95] are easier to correct because the CPG formulation allows direct control on these state parameters. The feedback loops can also be used to dictate the coupling between CPG nodes for spontaneous gait emergence [56], [92], [96].

Reflexes

Feedback mechanisms that combine both sensory information as well as feed forward actions are reflexes. Reflexes are a feed forward reaction to an external stimulus. Well-known examples are eye blinking or the patellar reflex. These patterns are triggered through a sensor signal that measures the presence of a stimulus. While the sensing is identical to feedback control, the reaction to said stimulus is a feed forward pattern that does not take the quantity of stimulus into account. A simple example: if an eye blinking reflex is triggered, the eye lid closes at the same speed independent of what triggered the reflex, how fast and how close the object is moving towards the eye.
This way the simple presence of the stimulus triggers the fastest possible reaction. The reaction time mentioned in my introductory story can be shortened through reflexes. Because the stimulus is not signaled to the brain and, therefore not consciously processed, the time between stimulus and reaction can be decreased to 20 ms [97]. While reflexes can prevent imminent failure, they are not suited for feedback pathways where the amount of stimulus is important in scaling the required reaction.
In robotics, reflexes are used for trip recovery [56], [98]–[100] and reflex-based gait generation [101], [102] where sensory information is used to trigger feed forward adaptation maneuvers to external perturbations. As mentioned in the first section, these reflexes are faster than continuous error-based feedback loops because their reaction is predefined and requires less computational effort and time.

1.3 Sensing

Feedback is a vital part of the control of legged robots. To provide the robot with a sense of its environment, appropriate information needs to be measured. While

many different sensor signals can be measured on a robot from an engineering perspective like motor currents, joint angles or body orientation and position, I will focus on signals relevant to biomechanical analysis, specifically force sensing for ground reaction force measurements [103], [104].
All force sensors depend on a deformable medium as well as a sensor modality that transduces measured deflections back into forces. Since force sensors are generally heavy and delicate, mounting both mechanical or electrical force sensors on a robot's foot is an engineering challenge.
Leveraging the force-torque relationship and the transparent gearbox design that is currently state-of-the-art in legged robotics, it is possible to implement a sense of proprioception [15], [59] without the use of external force sensors. Under the assumption of a pointed foot as the only point of contact with the ground, ground reaction forces can be calculated from motor torques, given an accurate robot model. While this assumption works in controlled environments where the assumption is met, proprioceptive force sensing is prone to failure if the legs hit obstacles not with the foot but other parts of the leg or if the assumption of flat ground is not met.
In unstructured terrain, contact information is also important to determine the switch between the flight phase and the stance phase in control [105]. While assumptions about the timing of the contact phases can be made when the terrain is flat, these assumptions also break in the presence of obstacles. This problem is pronounced when gain scheduling for distinct stance and flight control algorithms are used that rely on precise timing for when to switch the control approach [106], [107]. When using virtual impedance control, the controllers also rely on good timing information for when and how to change the leg's elastic behavior. Because the handling of impact is crucial in legged locomotion these sensor modalities need to be fast and robust to ensure good performance of the control structure.
For research purposes force and pressure sensing is an important tool to understand the dynamics of locomotion in animals where no direct proprioceptive force information can be measured, for example in human or animal studies. The forces and torques exerted in animals as well as in their bioinspired robotic copies provide information about the performance and comparability to reduced-order models like the SLIP model. For this reason, force and pressure plates are the gold standard sensor modality in legged locomotion research [23]. To build mobile robots that can locomote autonomously, it is advantageous to implement sensors that are mounted on the feet of the robot to capture continuous force and pressure data [5], [103], [108].

1.4 Optimization and Learning

To push the boundaries of what is possible in engineering today, optimization is used to increase performance, minimize energy requirements and reduce the cost in terms of material, time and possible failure. The goal of optimization and learning is to increase the performance of a system. The performance can be measured by a cost to reach a goal, or a reward for achieving a certain goal [5], [109], [110]. In a nutshell, the goal is to achieve the best performance while spending as little cost as possible. In an iterative process, the optimizer can follow the gradient (slope) of the reward function to find a parameter set that

optimizes the cost reward ratio.
If an analytical or numerical model of the cost/performance function exists, the optimization is called model-based. If the cost/performance function can only be sampled in discrete data points from a model or from hardware, and the optimization relies on those samples to approximate the optimal point it is called model-free.
In legged locomotion research, the formulation of a cost or reward function is a challenge because the goal is not clearly defined. When designing a robot, no adequate guess for the maximum speed the robot can reach exists prior to implementation. Additionally, mathematical models that can be optimized in a model-based approach are hard to formulate because of the lack of models for contact, impact and the uncertainty related to real-world interactions as described in subsection 1.1.2. By treating not well-understood features of legged locomotion as a black box, these problems can be optimized using model-free optimization approaches.
The challenge in model-free optimization is to find an optimization approach, that is data efficient, meaning how many data samples are required to approximate the reward function, that is robust to noise in the sample data and that does not get stuck in local maxima or minima.
The optimization problems in legged locomotion are often high-dimensional [111]–[113]. This means that the reward functions are not intuitive and can have local extreme points that are costly to sample in experiments. One optimization approach that can deal with these challenges is Bayesian optimization (BO).
In BO a Gaussian kernel is used to approximate the reward function as well as the uncertainty of the prediction made by the model [112], [114]. An algorithm calculates the next sample point based on it's estimation where the model can reduce its prediction uncertainty or where it expects a high objective/reward.
Because for this work I am only interested in optimization and learning as a tool, I will not make a clear distinction between the used optimization and learning approaches but rather focus on the characteristics that led the decision process of which tool to use.
When looking at optimization in the context of biology, the topic of data-driven optimization or learning is an exciting subject. Understanding how animals learn a task, through imitation or by conceiving a novel idea promises to provide insight into how brains and nervous systems work as well as how we can improve the acquisition of new skill sets both in animals and robotics.
In recent years the broad field of machine learning has gained popularity. Due to inexpensive and readily available computation, artificial neural networks (ANN) are considered an alternative to classical algorithm approaches. Comparable to biological neural networks, ANNs consist of layers of nodes (neurons) with weighted connections (synapses) between nodes [115]. Through training, the weights between the layers are adapted based on the current estimation error to approximate an arbitrary policy function. Deep neural nets (DNNs) and convolutional neural nets (CNN) excel in image and pattern recognition and are used for classification and computer vision [115].
Because robot manipulation and locomotion are tasks that require continuous control, classical DNNs are not applicable for this task. In this thesis, I will focus on the machine learning tools that are commonly used for legged locomotion.
The common approaches to decide on the action of an agent, in this case, the legged robot, are reinforcement learning (RL) [111], [113], [116] and recurrent

neural nets (RNNs) [61], [110].
In RL the agent is 'reinforced' by a reward to improve its choice of actions, called the policy. Like a dog learns new tricks with the help of food as a treat, the same concept can be used to maximize the performance of a robot. Based on the robot's performance a reward is calculated and is fed back to the learning algorithm. The learning algorithm determines how to adjust the learnable parameters to achieve a higher reward in the next steps.
The way the algorithm determines how to change the learnable parameters depends on the different approaches. Simpler approaches like 'gradient descent' estimate the gradient of the reward function and follow this gradient to a minimum or maximum point. This is comparable to finding a valley in the mountains by always walking downhill to eventually find the bottom of the valley.
Reinforcement learning algorithms primarily differ in their data efficiency, meaning how much data, and consequently how many learning experiments, are needed to progress the learning process. Going back to the 'find the valley' example from the gradient descent method, we would have to take four orthogonal steps in a two-dimensional world to determine where the steepest slope on the mountain is. This methods scale to 2^n for n learnable parameters and quickly becomes inefficient for large parameter sets.
More complex gradient methods like simultaneous perturbation stochastic approximation (SPSA) [117] limit the number of required samples by trading them off against the learning speed, also called the learning rate. SPSA only requires $2 \cdot n$ samples for n parameters but does not follow the steepest gradient. There exist many more ways to estimate reward gradients, all trading off required samples and learning rate in different ratios.

RNNs are deep neural networks with a 'feedback loop' that retains the state of the network to give the network a memory of its past state [110]. Like a DNN, it consists of layers of neutrons with weighted connections. Additionally, the internal states and the output state are fed back into the input layer of the RNN [110], [115]. A RNN can theoretically approximate the complete control structure of a robotic system with all feedback loops and only take a desired speed as the input to the network. Because deep neural nets are complex and have a large parameter space, training RNNs is very data inefficient. Current findings in robotics report millions of trials for learning a successful walking controller [61], [110].
Both RL and RNNs are especially attractive in a field like manipulation or locomotion robotics. Many aspects of these fields are still not understood well enough to formulate accurate mathematical models. The modeling of contact with the environment in the presence of uncertainty is an issue in robotics that has not yet been solved. Dealing with impact, switching dynamics between the stance and the flight phase, highly nonlinear dynamics, mechanical coupling, antagonistic actuation as well as inhomogeneous ground conditions make machine learning a viable approach in legged locomotion research. Since machine learning in this context is a black box, that automatically approximates the dynamics of the robot numerically, without requiring a mathematical formulation for the above-mentioned issues, it can excel in areas where model-based control still lacks understanding.
While the prospect of the robot learning its own behavior from a user-crafted reward function sounds promising, machine learning has inherent drawbacks that

still require an expert user with intuition and experience. Because the shape of the reward function is unknown prior to learning, it is hard to determine the hyperparameters of the approach, like the learning rate. If the learning rate is too low, the learning will take a long time. If it is too high the algorithm might miss features in the reward function by stepping over them. Another big issue is the choice of the initial guess. If the initial guess is wrong, the robot might either not produce any reward at all, or it might misbehave in a way that will damage the robot. A way to limit the hardware costs in terms of wear, experimentation time and failure is to learn in a simulation. This way the cost of time and failure are reduced, can be parallelized and the only cost is that of computation power and time. Since simulations are only an approximation of the real system, a learned policy will also only map to this approximation. Most policies learned in simulation need to be adjusted to be transferable to hardware. This discrepancy between simulation and hardware is called the sim2real gap [61], [111], [113]. The sim2real gap can be closed by either identifying the simulation model inaccuracies that prevent the transfer [113] or by using the same learning algorithm in hardware and continued learning from the optimal point found in simulation [61].
Alternatively, the learning algorithm can be implemented directly in hardware [116]. For low-dimensional parameter spaces, learning in hardware eliminates the sim2real gap as the data is directly sampled in the real system. Because exploration in hardware is prone to failure, meticulous and robust design are vital to prevent the robot from breaking during failure.

An important distinction between RNNs and RL is the scope of outcome. RL relies on an existing control structure and learns a parameter set that uses the full potential of said architecture. Compared to that, RNNs learn the control architectures themselves. RL is comparable to animals learning skills after birth with preexisting neural pathways. RNNs are more comparable to a combination of evolution, as in forming new neural pathways or control structures, as well as then learning to use them to complete a task successfully.

1.5 Objective

In my introduction story, I raised the question of why we did not fall when we could not have reacted in time. The topics described in the previous sections give a partial answer to why animals can locomote so gracefully and robustly. Here I would like to raise four, more specific, questions that are the central subject of this thesis.

Is there more to the passive impedance of legs than what the SLIP model captures? Compared to the SLIP model, animals have distal muscle-tendon structures that do not contribute to impedance in leg length direction. What is the functional morphology of these elastic structures?
In chapter 2 I investigate the question, which role the lower leg elasticities play in the passive impedance behavior of three-segmented legs. The role of elasticities in leg length direction, inspired by the SLIP model is evident in many robotic designs [12], [16], [30], [54]. Biarticular muscle-tendon structures in vertebrate animals possess elastic components that do not contribute to the

elasticity described by the SLIP model. In this chapter I investigate the functional morphology of these elastic components in a robotic model.

Learning to walk is difficult and prone to failure. Is it possible to provide a learning system with training wheels comparable to learning to ride a bike? And what is the effect of these temporary mechanical modifications to the system dynamics?
In chapter 3 I present a training wheel concept that can aid reinforcement learning in hardware, described in section 1.4 through temporary mechanical changes that increase the set of learnable parameters that yield a viable reward.

Can we provide legged robots with a sense of touch on their feet that can be used for both feedback control and biomechanical analyses?
In chapter 4 I describe a rugged wearable force and pressure sensor design that can be used as a feedback sensor as well as a biomechanical sensor for legged robots as described in section 1.3.

Passive elasticities provide stability and robustness to walking systems. How can we measure, how well the controller leverages the passive mechanics? And how can animals and robots learn to effectively leverage the advantages of their passive dynamics?
In chapter 5 I combine the findings of the previous chapters to investigate how a robot with bioinspired passive elastic legs can learn to match its neuroinspired CPG controller to its natural dynamics by minimizing the feedback activity required to mitigate the effects of mismatched control task and natural dynamics.

In the following chapters I will present my published work that is the subject of this thesis. Based on the introduction, I will give a more detailed motivation for each project and classify them in the field of locomotion research. I will shortly describe the publishing venue and state my contributions to each project.

Publications

F. Ruppert and A. Badri-Spröwitz, "Series elastic behavior of biarticular muscle-tendon structure in a robotic leg", *Frontiers in Neurorobotics*, vol. 13, Aug. 2019. DOI: 10.3389/fnbot.2019.00064.

S. Heim, F. Ruppert, A. A. Sarvestani, *et al.*, "Shaping in practice: Training wheels to learn fast hopping directly in hardware", in *©2018 IEEE. Reprinted, with permission 2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2018. DOI: 10.1109/icra.2018.8460984.

F. Ruppert and A. Badri-Sprowitz, "FootTile: A rugged foot sensor for force and center of pressure sensing in soft terrain", in *©2018 IEEE. Reprinted, with permission 2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2020. DOI: 10.1109/icra40945.2020.9197466.

F. Ruppert and A. Badri-Spröwitz, "Learning plastic matching of robot dynamics in closed-loop central pattern generators", *Nature Machine Intelligence*, Jul. 2022. DOI: 10.1038/s42256-022-00505-4.

Chapter 2

Series Elastic Behavior of Biarticular Muscle-Tendon Structure in a Robotic Leg

2.1 Topic

This work is placed at the intersection of robotic leg design and functional morphology in quadrupedal animals and their biomimetic blueprints in robotics. The paper investigates the function and advantages lower leg biarticular muscle-tendon structures have in a simplified robotic system concerning energy efficiency.

2.1.1 Motivation

As discussed in section 1.1 elastic behavior plays an important role in legged locomotion. The SLIP model as an established template [24], [25], [35], [38] is a paradigm used in the design of many robotic systems [26], [54], [55]. Because SLIP only represents the elasticity in leg length direction this is also reflected in the resulting mechanical designs. Most robotic leg designs with passive elasticities use their springs in leg length direction, i.e. for gravity compensation [26]. Because the SLIP model does not capture leg dynamics and torques, it does not capture the internal leg dynamics and the concept of two-dimensional impedance. Forces that are orthogonal to the leg axis are, therefore, not captured in the SLIP model or the elastic elements of SLIP-inspired robotic legs. However, all natural leg morphologies that are not erect like the human leg, have elastic elements that can potentially act as elasticities orthogonal to the leg length axis [42]. While some biomimetic leg designs exist that have elements capturing these forces [31], [32], there is little data describing the advantages of these leg designs. In this paper, I investigate the functional morphology of muscle-tendon structures in a robotic framework, that provide two-dimensional impedance, specifically, the biarticular lower leg muscle-tendon structures found in legged vertebrate animals.

2.1.2 Abstract

We investigate the role of lower leg muscle-tendon structures in providing serial elastic behavior to the hip actuator. We present a leg design with physical elastic elements in leg angle and virtual leg axis direction, and its impact on energy-efficient legged locomotion. By testing and comparing two robotic lower leg spring configurations, we can provide potential explanations of the functionality of similar animal leg morphologies with lower leg muscle-tendon network structures. We investigate the effects of leg angle compliance during locomotion. In a proof of concept, we show that a leg with a gastrocnemius inspired elasticity possesses elastic components that deflect in leg angle directions. The leg design with elastic elements in leg angle direction can store hip actuator energy in the series elastic element. We then show the leg's advantages in mechanical design in a vertical drop experiment. In the drop experiments, the biarticular leg requires 46% less power. During drop loading, the leg adapts its posture and stores the energy in its springs. The increased energy storing capacity in leg angle direction reduces energy requirements and cost of transport by 31% during dynamic hopping to a cost of transport of 1.2 at 0.9 kg body weight.
The biarticular robot leg design has major advantages, especially compared to more traditional robot designs. Despite its high degree of under-actuation, it is easy to converge into and maintain dynamic hopping locomotion. The presented control is based on a simple-to-implement, feed forward pattern generator. The biarticular leg's lightweight design can be rapidly assembled and is largely made from elements created by rapid prototyping. At the same time it is robust, and passively withstands drops from 200% body height. The biarticular leg shows, to the best of the authors' knowledge, the lowest achieved relative cost of transport documented for all dynamically hopping and running robots of 64% of a comparable natural runner's COT.

Takeaway Message

Impedance in leg angle direction, which is not captured in the SLIP model, plays an important role in energy-efficient locomotion and can be achieved through biarticular lower leg muscle-tendon structures that provide series elastic behavior to the hip actuator.

2.2 Classification

2.2.1 Venue

The journal Frontiers in Neurorobotics is a young journal at the intersection of neuroscience, artificial intelligence and robotics. The journal has a good impact factor and all publications are open access. This study was published in a special issue on 'Bioinspired Design and Control of Robots with Intrinsic Compliance'.

2.2.2 Contribution

My contribution to this work was to form the initial idea and conceive the experimental protocol. I designed and implemented the robot, the experimental setups, the motor controller and the sensor instrumentation. As the first author

I conducted all experiments, analyzed, interpreted and visualized the data, wrote the manuscript and supplied figures and video material. Throughout the project, I participated in regular meetings with my supervisor to discuss project management as well as content feedback.

Chapter 3

Shaping in Practice: Training Wheels to Learn Fast Hopping Directly in Hardware

3.1 Topic

This work is classified at the intersection of machine learning, specifically, reinforcement learning, and the interplay of mechanical passive mechanics. The paper provides insight into how temporary changes to the natural dynamics of a system can simplify control and learning tasks.

3.1.1 Motivation

Learning in the context of locomotion is a very attractive topic as described in section 1.4. Because we do not have a good understanding of impact and contact modeling as well as dealing with the nonlinear dynamics of legged systems (see subsection 1.1.2), machine learning in hardware is a good alternative to approaches that require mathematical formulations (section 1.4). Additionally, hand-tuning parameters in high-dimensional spaces becomes impossible, especially in coupled systems. For the work presented in chapter 2 even the parameter space of 'only' three dimensions showed how time-consuming and costly in terms of failure the tuning of robotic systems is. The CPG described later in chapter 5 consists of 26 tunable parameters, not even considering the internal coupling parameters. Because the parameters are coupled, meaning changing one parameter might influence other parameters, it is neither possible nor feasible in a timely manner to tune these CPGs by hand.
To speed up the tuning process, machine learning can be an alternative. Since learning in hardware, like in nature, is almost guaranteed to fail in the first few trials, external help is necessary. Because robotic systems rarely have the safety net of additional reflexes, arms and the associated controllers to catch a fall and the ability to self heal in case of failure, an external concept that provides

assistance in learning is a viable alternative. While the concept of training wheels is applied in everyday life, the idea has not been transferred to robotics yet. In this paper, we investigate the concept of training wheels for learning control parameters in hardware and investigate their effect relating to the simplification of choosing initial conditions and consequently the reduction of critical failures.

3.1.2 Abstract

Learning instead of designing robot controllers can greatly reduce engineering effort required, while also emphasizing robustness. Despite considerable progress in simulation, applying learning directly in hardware is still challenging, in part due to the necessity to explore potentially unstable parameters. We explore the concept of shaping the reward landscape with *training wheels*; temporary modifications of the physical hardware that facilitate learning. We demonstrate the concept with a robot leg mounted on a boom learning to hop fast. This proof of concept embodies typical challenges such as instability and contact while being simple enough to empirically map out and visualize the reward landscape. Based on our results we propose three criteria for designing effective training wheels for learning in robotics.

Takeaway Message

Learning in hardware is prone to failure and can be aided through temporary mechanical changes that increase the size of the gradient set where viable reward signals can be sampled.

3.2 Classification

3.2.1 Venue

The IEEE Conference for Robotics and Automation (ICRA) is the largest and one of the most important conferences in robotics. The venue has a high impact and is considered a top publishing venue in robotics research.

3.2.2 Contribution

I designed and implemented the boom setup with sensor instrumentation. I conducted a large part of the experiments and helped analyze and interpret the data. I wrote the hardware section of the manuscript, helped write the rest of the manuscript and supplied several figures and the video material. Throughout the project, I participated in regular team meetings to help plan and adjust milestones.

Chapter 4

FootTile: a Rugged Foot Sensor for Force and Center of Pressure Sensing in Soft Terrain

4.1 Topic

This work is classified in the field of soft sensor design. The sensor design was designed with applications in robotics and biomechanical research in mind.

4.1.1 Motivation

Proprioceptive force sensing in legged locomotion is a concept that is established in the scientific literature (see section 1.3) and was also used in the work presented in chapter 2 to estimate torques and forces from motor currents. As described in section 1.3, proprioceptive sensing reaches its limits in unstructured terrain. The limitations of proprioception in combination with experiences in my own research showing the necessity for contact sensing, lead to the development of this sensor concept. In the work presented in chapter 2 the proprioceptive data from joint angles and motor torques did not allow a precise estimation of ground contact on the leg. Because the springs in the leg design show a smooth transition from resting state to deflected state, contact can only be visually evaluated and the accuracy is limited to several control cycles. The same issue was previously observed in the work for chapter 3 where it would have been helpful to have contact data to quantify the performance of the learned controller. Especially during early testing the development of a contact sensor was discussed in the team but was not feasible in the given time frame.
Because legged locomotion is primarily defined by the interaction between a leg and the ground, ground reaction forces are an important part of a dataset to discuss the dynamics of legged locomotion. This work was motivated by the necessity to find a sensor solution that can be worn on the foot of either a robot or animal and provides continuous contact and possibly force information

without negatively influencing the leg mass and walking behavior.

4.1.2 Abstract

In this paper, we present FootTile, a foot sensor for reaction force and center of pressure sensing in challenging terrain. We compare our sensor design to standard biomechanical devices, force plates and pressure plates. We show that FootTile can accurately estimate force and pressure distribution during legged locomotion. FootTile weighs 0.9 g, has a sampling rate of 330 Hz, a footprint of 10×10 mm and can easily be adapted in sensor range to the required load case. In three experiments, we validate: first, the performance of the individual sensor, second an array of FootTiles for center of pressure sensing and third the ground reaction force estimation during locomotion in granular substrate. We then go on to show the accurate sensing capabilities of the waterproof sensor in liquid mud, as a showcase for real-world rough terrain use.

Takeaway Message

FootTile is a rugged, lightweight, waterproof force and pressure sensor that can be mounted on a robot's foot for both biomechanical analyses as well as for feedback information.

4.2 Classification

4.2.1 Venue

The IEEE Conference for Robotics and Automation (ICRA) is the largest and one of the most important conferences in robotics. The venue has a high impact and is considered a top publishing venue in robotics research.

4.2.2 Contribution

I helped form the initial idea and conceived the experimental protocol. As the first author, I designed and implemented the sensors, the test stands, measuring electronics and the robotic leg. I conducted all experiments, analyzed, interpreted and visualized the data, wrote the manuscript, supplied figures and video material and managed the project. Throughout the project, I participated in regular meetings with my supervisor to discuss project management as well as content feedback.

Chapter 5

Learning Plastic Matching of Robot Dynamics in Closed-loop CPGs

5.1 Topic

This work classifies at the intersection of legged robotics, neuromechanically inspired control for legged locomotion and optimization/machine learning for legged locomotion.

5.1.1 Motivation

In chapter 2 we developed and characterized a leg design that relies on its engineered passive elastic behavior to improve its energy efficiency. This passive behavior strongly influences the behavior of the legged robot [48], [54].
In combination with the mechanical coupling in the leg design and the high-dimensional CPGs similar to the one learned in chapter 3, the task of tuning CPG parameters is a challenging task. The CPG needs to actuate the robot while taking the passive behavior of the leg morphology as well as its underactuation into account. At the same time, the CPG requires feedback to mitigate perturbations and increase the robustness of the system. Because the CPG is a model-free control approach (see subsection 1.2.1), it is important to find a sensor signal that can be incorporated into the CPG as feedback.
The sensor design presented in chapter 4 provides contact information that can be used both as a feedback signal to correct timing errors in the CPG [89] as well as a proxy for how well the controller takes the mechanics into account. Rather than using a RL algorithm in hardware, it is safer to learn in simulation and transfer the findings to hardware. To make sure the reward landscape is accurately sampled, a global Bayesian optimizer is implemented to prevent not finding a salient gradient set (compare chapter 3). By optimizing the amount, by which the controller takes the robot mechanics into account, I aim to improve the locomotion performance of the robot in terms of energy efficiency and learn CPG parameters that enable the robot to walk in fewer trials compared to state-of-the-art RL algorithms.

Abstract

Animals achieve agile locomotion performance with reduced control effort and energy efficiency by leveraging compliance in their muscles and tendons. However, it is not known how biological locomotion controllers learn to leverage the intelligence embodied in their leg mechanics. Here we present a framework to match control patterns and mechanics based on the concept of short-term elasticity and long-term plasticity. Inspired by animals, we design a robot, Morti, with passive elastic legs. The quadruped robot Morti is controlled by a bioinspired closed-loop central pattern generator that is designed to elastically mitigate short-term perturbations using sparse contact feedback. By minimizing the amount of corrective feedback on the long term, Morti learns to match the controller to its mechanics and learns to walk within 1 h. By leveraging the advantages of its mechanics, Morti improves its energy efficiency by 42% without explicit minimization in the cost function.

Takeaway Message

Matching controller and natural dynamics is important for efficient legged locomotion. Matching can be achieved by minimizing the amount of short-term feedback activity in a closed-loop CPG as a proxy for mismatching dynamics.

5.2 Classification

5.2.1 Venue

The manuscript is published at Nature Machine Intelligence.
Nature Machine Intelligence is a high impact journal from the Nature family. The journal sits at the intersection of machine learning and embodied intelligence. The journal publishes high quality contributions in the field of intelligent systems that arouse interest in a variety of research communities related to machine learning, robotics and artificial intelligence.

5.2.2 Contribution

I formed the initial idea and conceived the experimental setup. As the first author, I designed and implemented the robot hardware, electronics, the experimental setup, the robot controller and the robot firmware. I wrote the code for the multi-body simulation and the optimization code. I conducted all experiments in simulation and hardware, analyzed, interpreted and visualized the data, wrote the manuscript and supplied figures and video material. Throughout the project, I participated in regular meetings with my supervisor to discuss project management as well as content feedback.

Chapter 6

Discussion

In this thesis, I provide insight into the functional morphology of biarticular muscle-tendon structures in chapter 2. In chapter 3 I present a concept to employ temporary mechanical changes to a legged system to simplify learning a control task. To augment the feed forward control task used in chapter 2 and chapter 3 I design and characterize a foot sensor design in chapter 4 that provides force and pressure information for feedback and biomechanical analysis. In chapter 5 I combine a robot based on the leg design investigated in chapter 2 with the controller from chapter 3, augmented with feedback loops using the sensor design in chapter 4 to enhance the performance of the robot. By matching the control task dynamics to the natural dynamics of the robot, the energy efficiency of the system can be improved and the robot learns to walk in one hour.
For a detailed discussion of the findings of each work please refer to the discussion section in each manuscript.

In section 1.5 I raised four research questions that lay the foundation for the studies in this thesis. In the following I want to provide the (partial) answers that my work provides to these questions.

6.1 Mechanics

Is there more to the passive impedance of legs than what the SLIP model captures? Compared to the SLIP model, animals have distal muscle-tendon structures that do not contribute to impedance in leg length direction. What is the functional morphology of these elastic structures?

In section 1.1 I describe, how the concept of passive impedance, specifically, the SLIP template [24], were developed based on biomechanical observations in animal locomotion. Because SLIP, as a reduced-order model, does not require torque and leg dynamics for its formulation it also can not capture energy efficiency and the dynamics of different leg morphologies. While SLIP is an excellent tool to investigate the behavior of the center of mass [25], [38] and the relationship of the emergence of gait based on stiffness and damping characteristics [118] it should not be used as the sole inspiration when it comes to mechanical leg

design.
When comparing state-of-the-art robotic leg designs [28]–[30], [110] with animal anatomy [41], [42], [71], [72] it becomes clear that the functional morphology of many major anatomical features in animals are not yet captured. The work in chapter 2 provides evidence for the functionality and advantages of lower-leg biarticular elastic structures that are currently not implemented in the majority of state-of-the-art robots.
While it is consensus that leg length elasticity aids legged systems in both energy storage and passive robustness [22], [24], [54], the effects in leg angle direction are mostly neglected. While lower leg elasticities were previously implemented in robotic structures [31], [32], their effects on the dynamics and performance of the system are not yet investigated and understood. In my work, I provide insights, that energy storage and impact mitigation capabilities in leg angle direction are equally beneficial. The addition of leg angle elasticity, that acts as a series elasticity to the hip actuator, increases the energy efficiency of the system and provides passive stability which became apparent during the experiments for chapter 2 where the pantograph leg would fail for most initial conditions when dropped on the floor. The biarticular leg would rarely fail and could even recover after stumbling and re-entrain itself into a hopping gait. I can only provide anecdotal evidence for the stability increase by implementing the biarticular elasticity into the leg. A general issue in legged locomotion remains, that no rigurous formulations for stability exist in the context of locomotion. While many groups use random obstacles to show stability and robustness no benchmarks or mathematical formulations exist to date that allow comparisons or optimization of robust locomotion. While approaches like viability theory [34] exist, they have so far not been shown to scale to real-world robotic systems and require further research.
The addition of the biarticular structure is purely mechanical. The proposed design is simple and robust and can be incorporated into most leg designs without much added complexity. Because the biarticular structure is purely passive there is no additional control effort.
In the work presented here, the same controller was used for both design variants and only slight adjustments to the gait parameters were necessary due to the changes to the natural dynamics.
The general concept of elasticity orthogonal to the leg axis scales to different leg postures, leg designs and the number of legs as long as a component of the biarticular elasticity remains orthogonal to the leg axis during the stance phase.
In my proof of concept study, I investigate the charging behavior of the biarticular elastic structures under torque to prove the series elastic behavior. By optimizing the ratio of knee and biarticular elasticity, the resting posture of the shank segment as well as the amount of knee flexion during the design of successor designs, the advantageous effects can be further improved with respect to energy efficiency, but also with respect to increased robustness of the mechanical design.

Learning to walk is difficult and prone to failure. Is it possible to provide a learning system with training wheels comparable to learning to ride a bike? And what is the effect of these temporary mechanical modifications to the system dynamics?

In chapter 3 we implemented mechanical changes to a legged system consisting of a boom structure and a hopper leg. The simple 'training wheel' of changes to the counterweight mass, simulating reduced gravity, enabled the system to learn safely in hardware from initial conditions, where learning without the training wheel would have failed. This shows, that the natural dynamics of a system can have beneficial effects, not only on the energy efficiency of the mechanical system but can also simplify the control task.
Our toy example of reduced gravity opens the door for extended training wheel ideas. Training wheels are used in the recovery of stroke patients learning to walk in pools, as well as kids learning to walk with carts or doorway jumpers or literal training wheels for bicycles. In robotics, the concept of training wheels has so far not been implemented in aiding the tuning or learning of control task parameters to the best of my knowledge.
While the general idea of training wheels is intuitive, an appropriate training wheel has to be chosen by an expert user since not every robotic system will benefit from the same training wheel based on morphology, function and desired behavior. While I believe that rigorous mathematical formulation for the training wheel idea is challenging, our work, as well as the everyday examples mentioned above provide evidence that training wheels are a non-negligible aspect of machine learning for hardware applications.

6.2 Control

The control schemes used in this thesis revolve around the implementation of bioinspired central pattern generators described in subsubsection 1.2.1. In chapter 2 I implemented a basic CPG to actuate the hopper legs. In this work, the CPG resembles a simple controller that benefits from the passive mechanical behavior of the leg [31], [86]. The series elastic behavior [36] of the leg converts the force control problem of contact and locomotion into a position control problem. The simplicity and ease of implementation make CPGs a great tool for locomotion research that focuses on mechanical performance in controlled environments and limits the design cost for sophisticated control algorithms. In chapter 3 we implemented a CPG for the same reason to enable locomotion behavior in a reduced and intuitive parameter space, where the focus is not on the actual control algorithm but rather on the effects mechanical changes have on learning.
In chapter 5 I augmented the feed forward CPG with both feedback and reflexes to enable robust behavior when mitigating perturbations. While the robots in chapter 2 and chapter 3 relied on mechanical leg impedance to mitigate perturbations, the CPG in chapter 5 can actively react to external perturbations. By adapting the timing characteristics [88] of the desired trajectories to the actual measured behavior of the robot, the CPG can react to unforeseen terrain changes.
The CPGs that are used in this thesis are designed as joint trajectory generators. As perturbations create forces and torques that change the system dynamics, it is hard to correlate changes in the dynamics to appropriate joint trajectory changes in the CPG. The feedback mechanisms are, therefore, hand-crafted and require hand-tuning. Because of their mathematical formulation, feedback mechanisms are hard to implement because they have to fit the rest of the CPG dynamics

and cannot arbitrarily change parameters without breaking the inter-limb coordination due to the constant internal coupling described in subsubsection 1.2.1. A common critique when designing feedback mechanisms and reflexes is their heuristic, handcrafted nature. While we know about reflexes from biological observations, we do not have enough understanding of how these feedback mechanisms work internally, to formulate design methodologies that could simplify the creation of feedback pathways. To understand how feedback is triggered and what system state is corrected requires more research, both in neuroscience as well as in biomechanics. With a better understanding of biological feedback pathways, we can then create feedback blueprints in robotics that do not rely on heuristic expert knowledge.
In comparison to full-body controllers [28], [29], [110], CPGs lack the versatility and robustness that a full state space model provides. On the other hand, CPGs are computationally far less expensive. The CPG implemented on the Morti robot runs on a cheap Raspberry Pi computer consuming 5 W, whereas modern full-body controllers and trajectory optimizers for optimal control struggle to run on state-of-the-art workstations with many times more computation power consuming upwards of 100 W [28], [37], [110]. While the tradeoff between versatility and efficiency will have to be made in every design process for legged machines, energy will always remain a bottleneck for both legged robots in research that are evaluated based on their energy efficiency in the robotics community as well as robots deployed in the real world that rely on finite battery power to fulfill their tasks.

6.3 Sensing

Can we provide legged robots with a sense of touch on their feet that can be used for both feedback control and biomechanical analyses?

In biomechanical studies (chapter 2) as well as for feedback information during control (chapter 5) sensor information is an important component in a legged robot. In chapter 4 I presented FootTile, a foot sensor design that can be used for both use-cases. While cheap and simple to produce, FootTile provides a rugged design and can be used for online force and pressure sensing. The sensor is lightweight and small and can be placed on the foot of a robot without negative effects on leg inertia.
I demonstrate the force and pressure sensing capabilities of the sensor design and provide an exemplary experiment in a real-world substrate. In chapter 5 I successfully implement the FootTile design as a contact sensor for closed-loop CPG control in the Morti robot.
In its current form, the sensor provides force information at 330 Hz which can be increased by implementing a higher bandwidth pressure sensor if required. With additional effort spent on reworking the housing of the off-the-shelf pressure sensor, the footprint could also be reduced further, or a sensor array could be implemented in one polyurethane mold to enable directional force sensing. The directional sensing capabilities and the higher bandwidth can further improve the sensor design.
For experiments, where the use of pressure plates and force sensors is not possible, the FootTile sensor can provide biomechanical data outside of a laboratory in

human or animal experiments if integrated into a wearable shoe.
In a control scenario, FootTile can be used as a binary contact sensor, as a pressure distribution sensor or as a force sensor. Because the sensor is easy to adapt and fast to manufacture it lends itself as a great tool for research, where rapid prototyping is an important factor.

6.4 Optimization and Learning

Passive elasticities provide stability and robustness to walking systems. How can we measure, how well the controller leverages the passive mechanics? And how can animals and robots learn to effectively leverage the advantages of their passive dynamics?

In chapter 3 we applied reinforcement learning to learn CPG parameters while investigating the effects of temporary mechanical changes to the system. While the RL approach was successful in finding viable gaits with the help of the training wheel, we also showed, how hard the initial guess is in machine learning as well as how prone to failure learning in hardware is. While learning in hardware worked for this reduced-order example, the number of mechanical parts, as well as learned parameters do not scale well for more complex systems. The tuning of hyperparameters like the learning rate or the choice of when to switch the training wheel settings (changing gravity) still rely mostly on experienced users and testing.
In chapter 5 I implemented a Bayesian optimizer to minimize the mismatching between the control task dynamics and the robot's natural dynamics. In contrast to the RL approach, the BO approach does not rely on a 'good' initial guess. Additionally, it is more data-efficient and more robust to noise compared to gradient descent methods [112], [114]. Because the BO can also sample parameter sets that result in violent failures and broken robot legs, I chose to optimize the parameters for the work in chapter 5 in simulation. Through optimization in software, the rollouts (that were comparable in duration to the rollouts in chapter 3) were less time-consuming both in terms of experimentation time as well as in repairing time after inevitable failures.
Through the implemented Bayesian optimizer, the robot is able to learn convincing locomotion skills in less than one hour from scratch. Compared to learning approaches using either RL or RNNs that report thousands or even millions of learning rollouts [110], [111], [113], it seems beneficial in terms of time to learn specific aspects of a locomotion controller compared to learning the complete control structure. Since the parameterized CPG structure exists prior to the optimization, learning is essentially bootstrapped. The learning approach does not need to approximate the control structure but instead only has to optimize the preexisting structure. This way, the required number of learning rollouts is orders of magnitude lower compared to deep learning approaches [110], [111], [113] and is comparable to the results of Tedrake et al. [116] because of a similar approach but a more complex robotic system.
The primary challenge in optimization and learning in my opinion, depends less on the choice of tools but rather on the crafting of a cost function. Locomotion is not a one-dimensional task. Where the performance of a simple system like a sensor can be evaluated by its accuracy, legged locomotion is described by many

performance factors like speed, energy efficiency, robustness, versatility but also computational efficiency. The desired cost or reward is therefore hard to quantify and hard to weigh. The performance of the robot in chapter 5 relies heavily on the weights and reward terms of the cost function. In my opinion, legged locomotion is an especially complex field for optimization because the goals are often hard to formulate in a way that is both quantifiable and comparable.
While energy efficiency, as mentioned above, is an important bottleneck in robotics, other characteristics like robust locomotion, metabolic energy conservation or 'natural looking' gait characteristics are difficult to formulate in mathematical equations. Where a system should, in layman's terms, 'learn from failure', failure in robotics can be hard to quantify, and the same is true for success. While a simple task like walking in a straight line towards a target is easy to evaluate, more complex tasks that can not be evaluated by speed or efficiency are difficult to characterize. Because we do not have insights into the cost function that drives natural evolution, optimization still requires hand-crafted and hand-tuned cost functions from expert users. The question remains what makes a good cost function for locomotion. For robotic engineers task-specific rewards for speed and energy efficiency seem intuitive [113], [119]. For simple tasks like walking in a straight line reaching the end position fast and efficient is sufficient [113], [116]. However in complex tasks like turning, jumping and climbing hand-crafted cost functions reach their limitations because the tasks themselves are hard to describe mathematically and therefore also hard to validate by a performance measure. Learning algorithms often resort to imitation learning approaches [110], [120] for more complex task formulations. While some examples for more complex tasks like fall recovery exist [121] they also require hand-tuning of weights and cost functions. The tuning of parameters that influence the optimization process is a task that is still not solved. This task will however become easier with cost functions that rely on intrinsic motivations like learning from failure, or reducing control effort to increase the synergy of the legged robot's subsystems, compared to measures that only focus on output performance like speed.
The results in chapter 5 highlight one possibility to learn locomotion based on a intrinsic measure. Rather than strive towards a task based performance like speed, the learning system is encouraged to leverage its natural mechanics as best as possible.
For more sophisticated learning approaches like curriculum learning [121] where the cost function changes with the learning success, these intrinsic measure can be especially helpful in the exploration phase where the robot needs to learn a basic synergy between its controller and its mechanics. Later on in the learning process, measures like efficiency can then be weighted higher, to fine tune the performance after basic locomotion skills were established.
These problems in crafting cost functions also limit the scalability of learning approaches. While the general machine learning or optimization approaches can be transferred to other systems, the tuning of hyperparameters and cost function weights still have to be adjusted for different system dynamics. The constraint of reducing energy efficiency [11], motor torque [119], body pitch [60] or high speed can also limit the possible outcome of an optimization. Where limiting torque can lead to conservative behaviors that do no exhaust the full potential of the system [121], maximizing speed can lead to behavior that is infeasible in hardware [113]. In my opinion, more general rewards like the mismatching measure in

chapter 5 are better measures for the performance of walking machines than energy efficiency.
As I show in chapter 5, energy efficiency is coupled to the successful synergy of control and mechanics and will passively increase without the need to constrain the system by actively minimizing energy requirements. As we constantly learn more about the inner workings of locomotion, more generalizable performance measures are required to further improve the performance of legged machines.
Another big topic in machine learning is data efficiency. Because most learning approaches attempt to learn a policy from scratch [111], [113] rollouts in the order of millions are required. Because the policies are initialized with random parameters and also sample the action space randomly these approaches cost a lot of computation and time. Even when parallelized in simulation these approaches are still costly in terms of computation power and electrical energy. I believe that, basic neural circuits like CPGs can be incorporated here, to establish a baseline functionality that does not need to be learned. If the CPG is parameterized correctly, it provides 'prior knowledge' in the form of a general system behavior. During learning, the CPG trajectories can then be optimized to increase performance and to match control and dynamics. Additionally, the policy that was bootstrapped with a CPG can be augmented with learned feedback mechanisms that increase the robustness of the system and leverage the generalizability of learning approaches. Learned feedback loops can also provide insights into the development of natural feedback pathways that are still not understood well enough to create design principles.

Additionally, temporary measures to increase the salient gradient set of reward signals during the initial exploration can increase the success rate of learning approaches when little to no prior knowledge about the reward landscape is available. Admittedly, the implementation of training wheels requires some trial and error as well as expert knowledge in the choice of an appropriate training wheel due to the early stage of this research idea. Additionally, the switch between gravity landscapes as described in chapter 3 also requires further research to achieve a rigorous approach that can eventually scale across hardware platforms and control approaches.

6.5 Biomechatronic Synergy

Throughout this thesis, I present different aspects of legged systems, that contribute to successful locomotion performance. Mechanical structures provide passive impedance, controllers provide active perturbation mitigation, learning algorithms optimize locomotion performance. More important than these individual contributions is the synergy of these classical engineering disciplines.
Coming back to my introduction story, the individual components of mechanics, control, sensing and learning can not prevent us from falling after missing that last step. However, the interplay of embodied intelligence in the mechanics, controllers that can exploit the natural dynamics and enhance the system performance through active adaptation with feedback from sparse sensor information enable us to learn behaviors that prevent falls when missing a step.
In chapter 5 the same synergy of action, perception and learning is applied to achieve convincing locomotion behavior. The Morti hardware, consisting of four

legs described in chapter 2 provides passive elastic behavior to the robot. In this case, the 'blind' CPG controller benefits from the passive impedance provided by the leg design. Because the leg impedance assumes the role of impact mitigation, a simpler control algorithm can be used to reduce the computational complexity and consequently metabolic energy requirements. The controller, therefore, does not have to actively react to perturbations, but their mitigation is handled in the leg mechanics. To increase the synergy, the controller needs to leverage the mechanics of the robot and the control task dynamics and natural dynamics have to match. In chapter 3 the CPG parameters were optimized for speed, the matching of dynamics was achieved indirectly as a black box. To actively quantify and minimize the mismatching, a measure for the amount of mismatching is required. I implement the FootTile sensor from chapter 4 as a contact sensor that provides feedback to the CPG. Because the mismatching dynamics manifest as feedback activity through the discrepancy between desired and measured robot behavior I can quantify and minimize the amount of mismatching.
By learning control parameters that enable the CPG to leverage the passive mechanics rather than overpowering them through actuation, control and mechanics synergize to increase the robot's performance in terms of energy efficiency that would not be possible with only one separate component described in previous chapters.
In a broader view, the idea of synergy should not only be limited to the components of one robotic system but should also concern the incorporation of different existing approaches for the design of legged systems. Most contemporary robot designs focus on control since the field of model-based full-body control became attractive and accessible due to inexpensive and powerful computation and hardware [27], [30], [55], [60], [61]. The robotic designs are used as controller demonstrators where passive elastic behavior would complicate the control task and hinder controllability. While these robots show convincing behavior that is currently not matched by other control schemes, I believe that the synergy of advantages from carefully engineered passive behavior in combination with the superior performance and versatility of model-based full-body controllers [93] can further improve the capabilities of walking machines.
The two approaches of feed forward and feedback control, while seemingly very different, in my opinion, describe the same view from different perspectives. Because baseline data in life sciences come primarily from biological or biomechanical observations, these observations can be recreated with different models and approaches. In the legged locomotion community, the two approaches seem to diverge further from each other. In my opinion, it would be beneficial to merge the approaches and combine their respective advantages. As mentioned in the introduction story, animals lack the bandwidth and reaction time to facilitate the same control as a full-body controller on a cerebral level. Furthermore, CPGs have been shown to exist in the locomotion controllers of animals like cats on a spinal level. For the control of legged robots, it therefore seems intuitive to rely on feed forward patterns for locomotion on flat ground and only adapt the feed forward patterns through feedback from a full-body controller in the presence of perturbations. In a sense, the feed forward pattern could act as a bootstrap for the trajectory optimizer in the full-body controller. When no feedback is required the feed forward patterns can be directly used as a reference trajectory. In the presence of perturbations, the feed forward pattern provides an initial guess for the trajectory optimizer to minimize the computational effort and optimization

time that currently limits the performance of full-body controllers. With this synergetic approach, the versatility and robustness of a full-body controller with active perturbation mitigation could be combined with the low computational effort of a feed forward pattern when no feedback intervention is necessary.
For the further understanding of animal locomotion as well as to leverage the previously published results in the biomechanical analysis of leg morphologies I also believe that effort in the synergy of bioinspired leg designs that combine the versatility of virtual impedance with the instantaneous impact mitigation capabilities of passive elastic legs will further increase the performance of legged machines. As the performance leap provided by full-body control by far outperforms the previous control algorithms and passive elastic leg designs, the advantages that were previously achieved in the absence of powerful computation seem to get more and more neglected. I believe, that the low computational passive behavior of legs in combination with more powerful control algorithms that could eventually be approximated by neural networks to reduce computational effort, are all valuable key points in our understanding of legged locomotion. Further system integration and synergetic approaches like the approach presented in this thesis will push the boundaries of legged locomotion further.
Because of the close resemblance of their biological role models, robots like the Morti platform are a great tool to investigate both the functional morphology of different leg designs as well as the interplay of bioinspired control, learning techniques and engineered natural dynamics in the context of legged locomotion research.
Cycling back to my introduction story at the beginning of this thesis, the result I presented here provide insights into how animals can mitigate unforeseen perturbations. How their passive elastic leg designs provide passive impedance where control frequencies are not sufficiently high. How sparse feedback signals are sufficient for neuroelastic feedback to actively mitigate perturbations. And how the synergy of mechanical design, neuroinspired control and learning approaches aid in the short and long term adaptations of walking systems to their environment to achieve the locomotion performance roboticists strive to understand and replicate in legged robots.
But the synergy does not stop at robotics. Because the legged robots I built are inspired by nature, they also allow inference in the field of life sciences. In chapter 2 I provide evidence, that there are phenomena in the leg morphology of quadrupedal vertebrates, that are not captured in one of the major computational models of legged locomotion, the SLIP model. And while my findings do not transfer directly into natural runners, since the biarticular leg is only a simplified blueprint of animal legs, it still provides insight into the functional morphology of animal leg designs. The charging and stabilizing behavior I saw in my experiments can lay the foundation for further research in animal locomotion to verify the hypothesis that the same behavior exists in animals. Which can in turn further our basic understanding of the mechanic blueprints that we believe define the embodied intelligence of legged animals.
Concerning the learning of locomotion tasks, my results in chapter 5 provide a new perspective on the cost function that drives learning in animals as well as in robots. Where robotic design focuses on task specific performance measures like speed and energy efficiency, I show that a more intrinsic motivation that allows learning from failure over a long time horizon is a viable addition to the complex cost functions that enable agile and elegant locomotion as seen in nature.

From my perspective, the findings in this thesis act as part of a feedback loop, that allows biologists and neuroscientists to build upon the blueprints that I adopted from their research and extended with more details presented in this thesis. Through this continuous cycle of synergizing biological findings with in-depth studies in robotic surrogates, my research contributes to our understanding of locomotion.

6.6 Outlook

Last, I want to give a perspective of where my research fits into the current state-of-the-art in legged research and where I believe my research can be incorporated for further studies.
Building on top of research that founded our understanding of elastic structures in legs [11], [18]–[21] and their portation into robotics [26], [29], [37], [55] and biomechanics [24], [38] my work extends the currently accepted models and methods by an additional focus on lower leg elastic structures and their advantages for legged locomotion. Where the SLIP model has been a template for legged robots [54], this template can now be extended by another layer of design choices that extend the SLIP idea more towards bioinspired real-world robots. Currently, mechanically simpler two-segmented legs are popular in the legged robotics community, that resemble the slip model in their leg design [29], [30], [37], [55], [111]. With the results presented in chapter 2 I hope to inspire leg designs, that combine the advantages of fully actuated legs with the advantages that passive elastic structures provide to legged locomotion into a hybrid design that combines the advantages of both worlds. While I provide evidence for a previously undescribed function in passive lower leg elasticities, in-depth studies are required to understand and optimize design choices like insertion locations, stiffness ratios and stability studies. To date, no design principles exist, that go much further than reduced inertia [37] and transparent transmissions [59]. It could be helpful to consider measures like morphological computation [122] and control effort [123] when discussing leg morphology. These measures allow comparisons between different design approaches like fully actuated robots [29], [30], [74] and robots with passive elastic structures [31]–[33]. Measures that go beyond the purely mechanical system and also consider control can emphasize the importance of both actively controlled and passive mechanic approaches that I also mentioned in my introduction. While reduced control effort [54], [55] is a commonly cited heuristic design goal for legged robots, a quantifiable measure would allow an evidence-based trade-off and an optimization of active and passive components. With these measures a more informed choice can be made in the design of leg morphologies to design legs as a part of a synergetic mechatronic system rather than designing legs that do not interfere with the controller that forms the centerdata basedpiece in these legged systems.
Many findings in leg designs are based on mechanical solutions to overcome computational deficits, primarily in the early days of legged locomotion research when computation was expensive and inaccessible [26][31], [48], [53]–[55]. Ever since computational performance increased, these findings are neglected in current leg designs [29], [37], [93], [124], rather than augmented with current findings in full-body control. In the long term, I see passive structures being reintroduced into the currently used leg designs for tasks that can not be achieved with active

control due to uncertainty and sensory delays [62]. Passive mechanisms and structures should therefore remain an important topic in legged locomotion research even if other fields currently overshadow their contribution to successful legged locomotion.
In locomotion control feed forward approaches are losing importance because of the impressive locomotion capabilities of full-body controllers with trajectory optimization and kinodynamic planning [93], [124]. Compared to these controllers feed forward approaches lack robustness and versatility. However, I do not believe feed forward approaches should be compared to full-body controllers, and neither can they. Position control based feed forward approaches like CPGs will never reach the same performance as full-body controllers when used by themselves. In my opinion CPGs should not be seen as a standalone control approach but rather as a reference generator for other control approaches. CPGs generate inexpensive trajectories that, as I showed in several examples, can be tuned to actuate a legged robot for basic locomotion skills. If formulated for an impedance control problem rather than a position control problem, CPGs are able to provide the periodic, computationally inexpensive behavior that full-body control approaches lack and compensate with constant computation. In turn, the CPG's lack of robustness will be compensated by the superior feedback mechanisms of full-body controllers.
When combined in a (deep) learning approach, the full-body controller can be approximated by a neural net [120] to reduce its computational cost.

While many machine learning studies try to approximate the control structure in an end-to-end approach [110], [111], [113], I believe it is a better approach to augment existing control approaches with the ability of machine learning approaches to generalize [111], [125] rather than replace them entirely. By leveraging preexisting knowledge in the control of legged systems and combining them with learned features to tackle the uncertainty inherent to legged locomotion in unstructured terrain in the presence of sensory delays, better locomotion performance can be achieved. And if only subsystems of the locomotion controller are learned, the system can be adapted without retraining the whole network. By treating the learning approach as a mechatronic subsystem of the walking system rather than a replacement for locomotion control, we can also prevent a 'reinvention of the wheel' while learning and approximating the underlying control approaches and instead leverage the basic principles of locomotion that are captured in the state-of-the-art control approaches described above. By bootstrapping the learning approaches it might also be possible to drastically reduce the number of learning rollouts that are required during the exploration of the learning landscape.
By synergizing the results of this thesis with the state-of-the-art as I just described, I believe we can expand what is currently possible in legged robotics, and in return use these systems to further our understanding of animal locomotion by using legged robotics as proxies into nature as I did in this thesis.

Bibliography

[1] A. Siegel, *Essential neuroscience*. Philadelphia: Lippincott Williams & Wilkins, 2006, ISBN: 9780781750776.

[2] H. L. More and J. M. Donelan, "Scaling of sensorimotor delays in terrestrial mammals", *Proceedings of the Royal Society B: Biological Sciences*, vol. 285, no. 1885, p. 20 180 613, Aug. 2018.

[3] A. Jain, R. Bansal, A. Kumar, *et al.*, "A comparative study of visual and auditory reaction times on the basis of gender and physical activity levels of medical first year students", *International Journal of Applied and Basic Medical Research*, vol. 5, no. 2, p. 124, 2015.

[4] H. Geyer and H. Herr, "A muscle-reflex model that encodes principles of legged mechanics produces human walking dynamics and muscle activities", *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 18, no. 3, pp. 263–273, Jun. 2010.

[5] M. Y. Chuah and S. Kim, "Enabling force sensing during ground locomotion: A bio-inspired, multi-axis, composite force sensor using discrete pressure mapping", *IEEE Sensors Journal*, vol. 14, no. 5, pp. 1693–1703, May 2014.

[6] S. Regnault, V. R. Allen, K. P. Chadwick, *et al.*, "Analysis of the moment arms and kinematics of ostrich (struthio camelus) double patellar sesamoids", *Journal of Experimental Zoology Part A: Ecological and Integrative Physiology*, vol. 327, no. 4, pp. 163–171, Apr. 2017.

[7] T. Siebert, C. Rode, W. Herzog, *et al.*, "Nonlinearities make a difference: Comparison of two common hill-type models with real muscle", *Biological Cybernetics*, vol. 98, no. 2, pp. 133–143, Nov. 2007.

[8] H.-W. Park, M. Y. Chuah, and S. Kim, "Quadruped bounding control with variable duty cycle via vertical impulse scaling", in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, Sep. 2014.

[9] C. J. PENNYCUICK and M. A. REZENDE, "The specific power output of aerobic muscle, related to the power density of mitochondria", *Journal of Experimental Biology*, vol. 108, no. 1, pp. 377–392, Jan. 1984.

[10] D. Böning, N. Maassen, and M. Steinach, "The efficiency of muscular exercise", *Deutsche Zeitschrift für Sportmedizin*, vol. 2017, no. 09, pp. 203–214, Sep. 2017.

[11] V. A. Tucker, "The energetic cost of moving about.", *American scientist*, vol. 63, pp. 413–419, 4 1975, ISSN: 0003-0996.

[12] J. H. Park, "Compliance/impedance control strategy for humanoids", in *Humanoid Robotics: A Reference*, Springer Netherlands, Oct. 2018, pp. 1009–1028.

[13] A. Šabanović and K. Ohnishi, *Motion Control Systems*. John Wiley & Sons (Asia) Pte Ltd, Jan. 2011.

[14] N. Hogan, "Impedance control: An approach to manipulation", in *1984 American Control Conference*, IEEE, Jul. 1984.

[15] S. Seok, A. Wang, D. Otten, *et al.*, "Actuator design for high force proprioceptive control in fast legged locomotion", in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems October 7-12, 2012. Vilamoura, Algarve, Portugal*, IEEE, 2012.

[16] A. Abate, R. L. Hatton, and J. Hurst, "Passive-dynamic leg design for agile robots", in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2015.

[17] B. Siciliano, *Springer handbook of robotics*. Berlin: Springer, 2008, ISBN: 9783540303015.

[18] R. ALEXANDER, "Elastic energy stores in running vertebrates", *American Zoologist*, vol. 24, no. 1, pp. 85–94, Feb. 1984.

[19] R. M. Alexander and H. C. Bennet-Clark, "Storage of elastic strain energy in muscle and other tissues.", *Nature*, vol. 265, pp. 114–117, 5590 Jan. 1977, ISSN: 0028-0836.

[20] R. M. Alexander, "Tendon elasticity and muscle function", *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, vol. 133, no. 4, pp. 1001–1011, Dec. 2002.

[21] Biewener and Baudinette, "In vivo muscle force and elastic energy storage during steady-speed hopping of tammar wallabies (macropus eugenii)", *The Journal of experimental biology*, vol. 198, pp. 1829–1841, Pt 9 1995, ISSN: 1477-9145.

[22] M. A. Daley, "Running over rough terrain: Guinea fowl maintain dynamic stability despite a large unexpected change in substrate height", *Journal of Experimental Biology*, vol. 209, no. 1, pp. 171–187, Jan. 2006.

[23] M. A. Daley, "Biomechanics: Running over uneven terrain is a no-brainer", *Current Biology*, vol. 18, no. 22, R1064–R1066, Nov. 2008.

[24] R. Blickhan, "The spring-mass model for running and hopping", *Journal of biomechanics*, vol. 22, no. 11-12, pp. 1217–1227, 1989.

[25] H. Geyer, A. Seyfarth, and R. Blickhan, "Compliant leg behaviour explains basic dynamics of walking and running", *Proceedings of the Royal Society B: Biological Sciences*, vol. 273, no. 1603, pp. 2861–2867, Nov. 2006.

[26] M. H. Raibert, H. B. Brown, and S. S. Murthy, "Machines that walk", in *Robotics and Artificial Intelligence*, Springer Berlin Heidelberg, 1984, pp. 345–364.

[27] F. Grimminger, A. Meduri, M. Khadiv, *et al.*, "An open torque-controlled modular robot architecture for legged locomotion research", *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3650–3657, Apr. 2020.

[28] M. Hutter, C. Gehring, D. Jud, *et al.*, "Anymal - a highly mobile and dynamic quadrupedal robot", Tech. Rep. interaction, 2016.

[29] B. Katz, J. D. Carlo, and S. Kim, "Mini Cheetah: A platform for pushing the limits of dynamic quadruped control", in *2019 International Conference on Robotics and Automation (ICRA)*, IEEE, May 2019.

[30] C. Semini, N. G. Tsagarakis, B. Vanderborght, *et al.*, "HyQ - hydraulically actuated quadruped robot: Hopping leg prototype", in *2008 2nd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics*, IEEE, Oct. 2008.

[31] A. Spröwitz, A. Tuleu, M. Vespignani, *et al.*, "Towards dynamic trot gait locomotion: Design, control, and experiments with cheetah-cub, a compliant quadruped robot", *The International Journal of Robotics Research*, vol. 32, no. 8, pp. 932–950, Jun. 2013.

[32] A. T. Spröwitz, A. Tuleu, M. Ajallooeian, *et al.*, "Oncilla robot: A versatile open-source quadruped research robot with compliant pantograph legs", *Frontiers in Robotics and AI*, vol. 5, Jun. 2018.

[33] F. Ruppert and A. Badri-Spröwitz, "Series elastic behavior of biarticular muscle-tendon structure in a robotic leg", *Frontiers in Neurorobotics*, vol. 13, Aug. 2019.

[34] S. Heim and A. Sprowitz, "Beyond basins of attraction: Quantifying robustness of natural dynamics", *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 939–952, Aug. 2019.

[35] R. Blickhan, A. Seyfarth, H. Geyer, *et al.*, "Intelligence by mechanics", *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 365, no. 1850, pp. 199–220, 2007.

[36] G. A. Pratt and M. M. Williamson, "Series elastic actuators", in *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, IEEE Comput. Soc. Press, 1995.

[37] H.-W. Park and S. Kim, "The MIT Cheetah, an electrically-powered quadrupedal robot for high-speed running", *Journal of the Robotics Society of Japan*, vol. 32, no. 4, pp. 323–328, 2014.

[38] H. Geyer, A. Seyfarth, and R. Blickhan, "Spring-mass running: Simple approximate solution and application to gait stability", *Journal of Theoretical Biology*, vol. 232, no. 3, pp. 315–328, Feb. 2005.

[39] M. A. Daley and A. A. Biewener, "Muscle force-length dynamics during levelversusincline locomotion: A comparison of in vivo performance of two guinea fowl ankle extensors", *Journal of Experimental Biology*, vol. 206, no. 17, pp. 2941–2958, Sep. 2003.

[40] M. A. Daley, A. Voloshina, and A. A. Biewener, "The role of intrinsic muscle mechanics in the neuromuscular control of stable running in the guinea fowl", *The Journal of Physiology*, vol. 587, no. 11, pp. 2693–2707, Jun. 2009.

[41] M. S. Fischer, N. Schilling, M. Schmidt, *et al.*, "Basic limb kinematics of small therian mammals", *Journal of Experimental Biology*, vol. 205, no. 9, pp. 1315–1338, May 2002.

[42] M. S. Fischer and R. Blickhan, "The tri-segmented limbs of therian mammals: Kinematics, dynamics, and self-stabilization—a review", *Journal of Experimental Zoology Part A: Comparative Experimental Biology*, vol. 305A, no. 11, pp. 935–952, 2006.

[43] R. C. Payne, J. R. Hutchinson, J. J. Robilliard, *et al.*, "Functional specialisation of pelvic limb anatomy in horses (equus caballus)", *Journal of Anatomy*, vol. 206, no. 6, pp. 557–574, Jun. 2005.

[44] A. Biewener, "Scaling body support in mammals: Limb posture and muscle mechanics", *Science*, vol. 245, no. 4913, pp. 45–48, 1989.

[45] S. Gatesy and A. Biewener, "Bipedal locomotion: Effects of speed, size and limb posture in birds and humans", *Journal of Zoology*, vol. 224, no. 1, pp. 127–147, 1991.

[46] A. Seyfarth, R. Blickhan, and J. V. Leeuwen, "Optimum take-off techniques and muscle design for long jump", *Journal of Experimental Biology*, vol. 203, no. 4, pp. 741–750, Feb. 2000.

[47] J. R. Hutchinson, "The evolution of hindlimb tendons and muscles on the line to crown-group birds", *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, vol. 133, no. 4, pp. 1051–1086, Dec. 2002.

[48] T. McGeer, "Passive dynamic walking", *The International Journal of Robotics Research*, vol. 9, no. 2, pp. 62–82, Apr. 1990.

[49] A. Seyfarth, M. Günther, and R. Blickhan, "Stable operation of an elastic three-segment leg", *Biological Cybernetics*, vol. 84, no. 5, pp. 365–382, Apr. 2001.

[50] A. A. Biewener, C. T. Farley, T. J. Roberts, *et al.*, "Muscle mechanical advantage of human walking and running: Implications for energy cost", *Journal of Applied Physiology*, vol. 97, no. 6, pp. 2266–2274, Dec. 2004.

[51] H. S. Im, O. Goltzer, and F. T. Sheehan, "The effective quadriceps and patellar tendon moment arms relative to the tibiofemoral finite helical axis", *Journal of Biomechanics*, vol. 48, no. 14, pp. 3737–3742, Nov. 2015.

[52] A. Abate, J. W. Hurst, and R. L. Hatton, "Mechanical antagonism in legged robots", in *Robotics: Science and Systems XII*, Robotics: Science and Systems Foundation, 2016.

[53] M. H. Raibert, "Running with symmetry", in *Autonomous Robot Vehicles*, Springer New York, 1986, pp. 45–61.

[54] D. Renjewski, A. Sprowitz, A. Peekema, *et al.*, "Exciting engineered passive dynamics in a bipedal robot", *IEEE Transactions on Robotics*, vol. 31, no. 5, pp. 1244–1251, Oct. 2015.

[55] C. Hubicki, J. Grimes, M. Jones, *et al.*, "ATRIAS: Design and validation of a tether-free 3d-capable spring-mass bipedal robot", *The International Journal of Robotics Research*, vol. 35, no. 12, pp. 1497–1521, Jul. 2016.

[56] Y. Fukuoka, H. Kimura, Y. Hada, *et al.*, "Adaptive dynamic walking of a quadruped robot 'tekken' on irregular terrain using a neural system model", in *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, IEEE.

[57] M. Buehler, R. Battaglia, A. Cocosco, *et al.*, "Scout: A simple quadruped that walks, climbs, and runs", in *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146)*, IEEE, 1998.

[58] M. HUTTER, C. GEHRING, M. BLOESCH, *et al.*, "Starleth: A compliant quadrupedal robot for fast, efficient, and versatile locomotion", in *Adaptive Mobile Robotics*, WORLD SCIENTIFIC, Jul. 2012, pp. 483–490.

[59] P. M. Wensing, A. Wang, S. Seok, *et al.*, "Proprioceptive actuator design in the MIT Cheetah: Impact mitigation and high-bandwidth physical interaction for dynamic legged robots", *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 509–522, Jun. 2017.

[60] H.-W. Park, P. M. Wensing, and S. Kim, "High-speed bounding with the MIT Cheetah 2: Control design and experiments", *The International Journal of Robotics Research*, vol. 36, no. 2, pp. 167–192, Jan. 2017.

[61] T. Apgar, P. Clary, K. Green, *et al.*, "Fast online trajectory optimization for the bipedal robot cassie", in *Robotics: Science and Systems XIV*, Robotics: Science and Systems Foundation, Jun. 2018.

[62] M. S. Ashtiani, A. A. Sarvestani, and A. Badri-Spröwitz, "Hybrid parallel compliance allows robots to operate with sensorimotor delays and low control frequencies", *Frontiers in Robotics and AI*, vol. 8, Jun. 2021.

[63] D. Lakatos, K. Ploeger, F. Loeffl, *et al.*, "Dynamic locomotion gaits of a compliantly actuated quadruped with SLIP-like articulated legs embodied in the mechanical design", *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3908–3915, Oct. 2018.

[64] S. Wolf, G. Grioli, O. Eiberger, *et al.*, "Variable stiffness actuators: Review on design and components", *IEEE/ASME Transactions on Mechatronics*, vol. 21, no. 5, pp. 2418–2430, Oct. 2016.

[65] T. Verstraten, P. Beckerle, R. Furnémont, *et al.*, "Series and parallel elastic actuation: Impact of natural dynamics on power and energy consumption", *Mechanism and Machine Theory*, vol. 102, pp. 232–246, Aug. 2016.

[66] C. Gehring, P. Fankhauser, L. Isler, *et al.*, "ANYmal in the field: Solving industrial inspection of an offshore HVDC platform with a quadrupedal robot", in *Field and Service Robotics*, Springer Singapore, 2021, pp. 247–260.

[67] G. J. V. I. Schenau, "From rotation to translation: Constraints on multi-joint movements and the unique action of bi-articular muscles", *Human Movement Science*, vol. 8, no. 4, pp. 301–337, Aug. 1989.

[68] G. J. van Ingen Schenau, "Proposed actions of bi-articular muscles and the design of hindlimbs of bi- and quadrupeds", *Human Movement Science*, vol. 13, no. 5, pp. 665–681, Oct. 1994.

[69] L. Wang, F. Meng, H. Liu, *et al.*, "Design and implementation of jumping robot with multi-springs based on the coupling of polyarticular", in *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, IEEE, Dec. 2018.

[70] A. Shkolnik and R. Tedrake, "High-dimensional underactuated motion planning via task space control", in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, Sep. 2008.

[71] H. Witte, H. Hoffmann, R. Hackert, *et al.*, "Biomimetic robotics should be based on functional morphology", *Journal of Anatomy*, vol. 204, no. 5, pp. 331–342, May 2004.

[72] H. Witte, R. Hackert, K. E. Lilje, *et al.*, "Transfer of biological principles into the construction of quadruped walking machines", in *Proceedings of the Second International Workshop on Robot Motion and Control. RoMoCo'01 (IEEE Cat. No.01EX535)*, Poznan Univ. Technol, 2001.

[73] W. P. Lombard, "The tendon action and leverage of two-joint muscles of the hind leg of the frog, with special reference of the spring movement.", *Contributions to Medical Research*, vol. 34, no. 4, pp. 280–301, 1903.

[74] M. Hutter, C. Gehring, M. A. Hopflinger, *et al.*, "Toward combining speed, efficiency, versatility, and robustness in an autonomous quadruped", *IEEE Transactions on Robotics*, vol. 30, no. 6, pp. 1427–1440, Dec. 2014.

[75] G. Bledt and S. Kim, "Extracting legged locomotion heuristics with regularized predictive control", in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2020.

[76] W. S. Levine, *The control handbook*. Boca Raton, Fla: CRC Press, 2011, ISBN: 9781420073669.

[77] F. Ruppert and A. Badri-Spröwitz, "Learning plastic matching of robot dynamics in closed-loop central pattern generators", *Nature Machine Intelligence*, Jul. 2022.

[78] K. Baev, V. Esipenko, and Y. Shimansky, "Afferent control of central pattern generators: Experimental analysis of locomotion in the decerebrate cat", *Neuroscience*, vol. 43, no. 1, pp. 237–247, Jan. 1991.

[79] A. J. Ijspeert, "Central pattern generators for locomotion control in animals and robots: A review", *Neural networks*, vol. 21, no. 4, pp. 642–653, 2008.

[80] A. Cohen and P. Wallen, "The neuronal correlate of locomotion in fish", *Experimental Brain Research*, vol. 41, no. 1, Dec. 1980.

[81] I. Delvolvé, P. Branchereau, R. Dubuc, *et al.*, "Fictive rhythmic motor patterns induced by NMDA in an in vitro brain stem–spinal cord preparation from an adult urodele", *Journal of Neurophysiology*, vol. 82, no. 2, pp. 1074–1077, Aug. 1999.

[82] M. H. Dickinson, "How animals move: An integrative view", *Science*, vol. 288, no. 5463, pp. 100–106, Apr. 2000.

[83] A. Frigon and S. Rossignol, "Experiments and models of sensorimotor interactions during locomotion", *Biological Cybernetics*, vol. 95, no. 6, pp. 607–627, Nov. 2006.

[84] P. Arena, "The central pattern generator: A paradigm for artificial locomotion", *Soft Computing*, vol. 4, no. 4, pp. 251–266, Dec. 2000.

[85] M. Okada, K. Tatani, and Y. Nakamura, "Polynomial design of the nonlinear dynamics for the brain-like information processing of whole body motion", in *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, IEEE.

[86] A. J. Ijspeert, A. Crespi, D. Ryczko, *et al.*, "From swimming to walking with a salamander robot driven by a spinal cord model", *Science*, vol. 315, no. 5817, pp. 1416–1420, Mar. 2007.

[87] H. Kimura, S. Akiyama, and K. Sakurama, *Autonomous Robots*, vol. 7, no. 3, pp. 247–258, 1999.

[88] L. Righetti and A. Ijspeert, "Design methodologies for central pattern generators: An application to crawling humanoids", in *Robotics: Science and Systems II*, Robotics: Science and Systems Foundation, Aug. 2006.

[89] L. Righetti and A. J. Ijspeert, "Pattern generators with sensory feedback for the control of quadruped locomotion", in *2008 IEEE International Conference on Robotics and Automation*, IEEE, May 2008.

[90] N. Kopell and G. Ermentrout, "Coupled oscillators and the design of central pattern generators", *Mathematical Biosciences*, vol. 90, no. 1-2, pp. 87–109, Jul. 1988.

[91] R. M. Alexander, "The gaits of bipedal and quadrupedal animals", *The International Journal of Robotics Research*, vol. 3, no. 2, pp. 49–59, Jun. 1984.

[92] R. Thandiackal, K. Melo, L. Paez, *et al.*, "Emergence of robust self-organized undulatory swimming based on local hydrodynamic force sensing", *Science Robotics*, vol. 6, no. 57, Aug. 2021.

[93] G. Bledt, M. J. Powell, B. Katz, *et al.*, "MIT Cheetah 3: Design and control of a robust, dynamic quadruped robot", in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2018.

[94] J. Englsberger, C. Ott, and A. Albu-Schaffer, "Three-dimensional bipedal walking control using divergent component of motion", in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, Nov. 2013.

[95] G. Sartoretti, S. Shaw, K. Lam, *et al.*, "Central pattern generator with inertial feedback for stable locomotion and climbing in unstructured terrain", in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2018.

[96] D. Owaki and A. Ishiguro, "A quadruped robot exhibiting spontaneous gait transitions from walking to trotting to galloping", *Scientific Reports*, vol. 7, no. 1, Mar. 2017.

[97] P. Hoffmann, "Die reflexzeit", in *Untersuchungen über die Eigenreflexe (Sehnenreflexe) Menschlicher Muskeln*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1922, pp. 45–59, ISBN: 978-3-662-24789-1.

[98] G. Urbain, V. Barasuol, C. Semini, *et al.*, "Stance control inspired by cerebellum stabilizes reflex-based locomotion on HyQ robot", in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2020.

[99] A. Murai and K. Yamane, "A neuromuscular locomotion controller that realizes human-like responses to unexpected disturbances", in *2011 IEEE International Conference on Robotics and Automation*, IEEE, May 2011.

[100] N. Rathod, A. Bratta, M. Focchi, *et al.*, "Model predictive control with environment adaptation for legged locomotion", *IEEE Access*, vol. 9, pp. 145 710–145 727, 2021.

[101] T. Tanikawa, Y. Masuda, and M. Ishikawa, "A reciprocal excitatory reflex between extensors reproduces the prolongation of stance phase in walking cats: Analysis on a robotic platform", *Frontiers in Neurorobotics*, vol. 15, Apr. 2021.

[102] C. Ferreira, V. Matos, C. P. Santos, *et al.*, "Quadrupedal locomotion based in a purely reflex controller", in *Proceedings of the 11th International Conference on Informatics in Control, Automation and Robotics*, SCITEPRESS - Science, 2014.

[103] R. Kaslin, H. Kolvenbach, L. Paez, *et al.*, "Towards a passive adaptive planar foot with ground orientation and contact force sensing for legged robots", in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2018.

[104] X. A. Wu, T. M. Huh, R. Mukherjee, *et al.*, "Integrated ground reaction force sensing and terrain classification for small legged robots", *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 1125–1132, Jul. 2016.

[105] Y. Takahashi, K. Nishiwaki, S. Kagami, *et al.*, "High-speed pressure sensor grid for humanoid robot foot", in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2005.

[106] E. Guglielmino, C. Semini, H. Kogler, *et al.*, "Power hydraulics - switched mode control of hydraulic actuation", in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, Oct. 2010.

[107] J.-Y. Kim, C. G. Atkeson, J. K. Hodgins, *et al.*, "Online gain switching algorithm for joint position control of a hydraulic humanoid robot", in *2007 7th IEEE-RAS International Conference on Humanoid Robots*, IEEE, Nov. 2007.

[108] J. Hwangbo, C. D. Bellicoso, P. Fankhauser, *et al.*, "Probabilistic foot contact estimation by fusing information from dynamics and differential/forward kinematics", in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2016.

[109] I. Griva, S. G. Nash, and A. Sofer, *Linear and Nonlinear Optimization*. CAMBRIDGE, Mar. 31, 2009, 764 pp., ISBN: 0898716616.

[110] X. B. Peng, E. Coumans, T. Zhang, *et al.*, "Learning agile robotic locomotion skills by imitating animals", in *Robotics: Science and Systems XVI*, Robotics: Science and Systems Foundation, Jul. 2020.

[111] J. Siekmann, S. Valluri, J. Dao, *et al.*, "Learning memory-based control for human-scale bipedal locomotion", Jun. 3, 2020. arXiv: 2006.02402 [cs.RO].

[112] R. Calandra, A. Seyfarth, J. Peters, *et al.*, "Bayesian optimization for learning gaits under uncertainty", *Annals of Mathematics and Artificial Intelligence*, vol. 76, no. 1-2, pp. 5–23, Jun. 2015.

[113] J. Tan, T. Zhang, E. Coumans, *et al.*, "Sim-to-real: Learning agile locomotion for quadruped robots", in *Robotics: Science and Systems XIV*, Robotics: Science and Systems Foundation, Jun. 2018.

[114] J. Mockus, *Bayesian Approach to Global Optimization*. Springer Netherlands, Dec. 6, 2012, 270 pp., ISBN: 9789400909090.

[115] P. Winston, *Artificial intelligence*. Reading, Mass: Addison-Wesley Pub. Co, 1993, ISBN: 0201533774.

[116] R. Tedrake, T. Weirui, Z. H. Sebastian, *et al.*, "Learning to walk in 20 minutes", 224.

[117] J. C. Spall, "An overview of the simultaneous perturbation method for efficient optimization", *Johns Hopkins Apl Technical Digest*, vol. 19, pp. 482–492, 1998.

[118] Z. Gan, Y. Yesilevskiy, P. Zaytsev, *et al.*, "All common bipedal gaits emerge from a single passive model", *Journal of The Royal Society Interface*, vol. 15, no. 146, p. 20 180 455, Oct. 2018.

[119] M. H. Yeganegi, M. Khadiv, S. A. A. Moosavian, *et al.*, "Robust humanoid locomotion using trajectory optimization and sample-efficient learning", in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, IEEE, Oct. 2019.

[120] J. Viereck and L. Righetti, "Learning a centroidal motion planner for legged locomotion", in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2021.

[121] J. Hwangbo, J. Lee, A. Dosovitskiy, *et al.*, "Learning agile and dynamic motor skills for legged robots", *Science Robotics*, vol. 4, no. 26, Jan. 2019.

[122] R. Pfeifer and J. Bongard, *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.

[123] D. F. B. Haeufle, M. Günther, G. Wunner, *et al.*, "Quantifying control effort of biological and technical movements: An information-entropy-based approach", *Physical Review E*, vol. 89, no. 1, Jan. 2014.

[124] B. Ponton, A. Herzog, S. Schaal, *et al.*, "A convex model of humanoid momentum dynamics for multi-contact motion generation", in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, IEEE, Nov. 2016.

[125] C. Yang, K. Yuan, Q. Zhu, *et al.*, "Multi-expert learning of adaptive legged locomotion", *Science Robotics*, vol. 5, no. 49, Dec. 2020.

Appendices

frontiers in Neurorobotics

ORIGINAL RESEARCH
published: 13 August 2019
doi: 10.3389/fnbot.2019.00064

Series Elastic Behavior of Biarticular Muscle-Tendon Structure in a Robotic Leg

Felix Ruppert and Alexander Badri-Spröwitz*

Dynamic Locomotion Group, Max Planck Institute for Intelligent Systems, Stuttgart, Germany

We investigate the role of lower leg muscle-tendon structures in providing serial elastic behavior to the hip actuator. We present a leg design with physical elastic elements in leg angle and virtual leg axis direction, and its impact onto energy efficient legged locomotion. By testing and comparing two robotic lower leg spring configurations, we can provide potential explanations of the functionality of similar animal leg morphologies with lower leg muscle-tendon network structures. We investigate the effects of leg angle compliance during locomotion. In a proof of concept, we show that a leg with a gastrocnemius inspired elasticity possesses elastic components that deflect in leg angle directions. The leg design with elastic elements in leg angle direction can store hip actuator energy in the series elastic element. We then show the leg's advantages in mechanical design in a vertical drop experiment. In the drop experiments the biarticular leg requires 46% less power. During drop loading, the leg adapts its posture and stores the energy in its springs. The increased energy storing capacity in leg angle direction reduces energy requirements and cost of transport by 31% during dynamic hopping to a cost of transport of 1.2 at 0.9 kg body weight. The biarticular robot leg design has major advantages, especially compared to more traditional robot designs. Despite its high degree of under-actuation, it is easy to converge into and maintain dynamic hopping locomotion. The presented control is based on a simple-to-implement, feed-forward pattern generator. The biarticular legs lightweight design can be rapidly assembled and is largely made from elements created by rapid prototyping. At the same time it is robust, and passively withstands drops from 200% body height. The biarticular leg shows, to the best of the authors' knowledge, the lowest achieved relative cost of transport documented for all dynamically hopping and running robots of 64% of a comparable natural runner's COT.

Keywords: robotics, bioinspired robotics, intrinsic compliance, locomotion, energy efficiency, biarticular, leg design, series elastics

OPEN ACCESS

Edited by:
Zhao Guo,
Wuhan University, China

Reviewed by:
Muye Pang,
Wuhan University of Technology, China
Ning Tan,
Sun Yat-sen University, China

****Correspondence:***
Felix Ruppert
ruppert@is.mpg.de

Received: *18 April 2019*
Accepted: *22 July 2019*
Published: *13 August 2019*

Citation:
Ruppert F and Badri-Spröwitz A (2019) Series Elastic Behavior of Biarticular Muscle-Tendon Structure in a Robotic Leg. Front. Neurorobot. 13:64. doi: 10.3389/fnbot.2019.00064

1. INTRODUCTION

A persistent question in legged locomotion relates to the functional morphology of compliant elements in segmented leg structures. Elastic elements in legs enhance locomotion performance in terms of stability, robustness to perturbations and impact mitigation in legged walking systems (Hurst, 2008; Rummel et al., 2010). Leg elasticity can simplify the control task (Verstraten et al., 2016; Beckerle et al., 2017) by giving the system favorable natural dynamics. Biological observations

show that muscles and tendons act like elastic elements (Biewener, 1998; Alexander, 2002) that enable rich locomotion skills with high energy efficiency at low control effort (Daley, 2008; Lakatos et al., 2018).

In bioinspired robotics, the concept of elasticity was first introduced in series elastic actuators (SEA) (Pratt and Williamson, 1995) and prismatic actuators (Raibert et al., 1984). Many robotic designs use a minimal order template, the spring-loaded inverted pendulum (SLIP) model (Blickhan, 1989; Seyfarth et al., 2001; Geyer et al., 2006), as a design baseline for walking systems. Based on SLIP models much effort has gone into designing compliance in virtual leg axis direction for robots. Compliance is implemented as either motor controlled compliance (Ding and Park, 2017; Park et al., 2017) or physical springs (Fukuoka et al., 2003; Renjewski et al., 2015; Semini et al., 2015). Like the SLIP model, these robots have elastic elements in their joints to help them achieve the same energy efficient and robust behavior as their biological role models (Alexander and Bennet-Clark, 1977).

The primary focus in designing compliance in robots using physical springs has been on *virtual leg axis direction compliance*. We pronounce the influence of additional physical elastic elements acting in *leg angle direction*. In a real world locomotion scenario we show compliance in leg angle direction to be important as well as compliance in virtual leg axis direction which has been shown in SLIP model and SLIP-inspired robots. To achieve intrinsic compliance we implement a mechanism inspired by a biological blueprint.

In studies of quadrupedal leg morphology, a four-bar-like mechanism has been observed by Lombard (1903). This simplified mechanism describes the functional morphology of lower leg muscle-tendon groups. It was extended by Witte et al. (2001, 2004) to a pantograph structure, including muscle-tendon structures. Because of the distal elastic tendon structures (Roberts, 2016), the simplified pantograph structure is spring-loaded. The concept implemented in a robot (Spröwitz et al., 2013), briefly suggested a potential function as an effective elastic element in leg angle direction (Spröwitz et al., 2014). The element is oriented so that its elastic elements possess deflection components orthogonal to virtual leg axis direction. Unlike in SLIP model, these components do not primarily contribute to deflection in virtual leg axis direction. However, they deflect under the presence of hip torque and perturbations that reflect as a torque to the hip actuator.

This leg morphology has been applied in robots before, empirically showing its advantages concerning the simplification in creating stable gaits. However, the general morphology has not yet been characterized, and the differences and advantages are not yet documented.

We investigate the effects, leg angle compliance in combination with virtual leg axis compliance has on spring behavior and resulting energy efficiency in the leg.

In this paper, we characterize one leg design with virtual leg axis compliance and one with virtual leg axis *and* leg angle compliance. We show the differences in leg morphology first on a simple kinematic model. To decompose virtual leg axis and leg angle effects we conduct static experiments to examine isolated virtual leg axis and torque influence on the elastic elements. In a drop test experiment, we investigate the mechanical behavior under dynamic loading without considering control design. At last, we compare both legs in a monoped hopping experiment and analyze the differences in dynamic behavior and energy stored and recuperated in the springs under a realistic load case.

1.1. Related Work

The functional morphology of multiple degrees of compliance in multi-segmented legs in animals and robotics has not been understood yet by either biologists nor roboticists. While two-segmented legs with one degree of compliance have been studied thoroughly (Raibert et al., 1984; Hutter et al., 2012; Semini et al., 2015; Park et al., 2017), the placement and interplay between multiple compliant elements is still an unsolved research topic.

Because of observations in biological examples, implementations of multi-segmented legs with several compliant elements have been tested in robotic hardware as well as in simulations to understand their behavior. Spröwitz et al. (2013, 2018) implemented a leg with a biarticular spring to investigate self stabilizing behavior on a quadruped during dynamic locomotion. They showed, that a simple sensorless central pattern generator with a position controller can allow dynamic feed forward locomotion. Iida et al. (2007) investigated the possibility to create both walking and running gaits in a humanoid biped with biarticular springs as well as the ability to create more human-like gaits. Sato et al. (2015) implemented a robot with only one biarticular spring but no intrinsic compliant knee. There, the biarticular spring provided elastic behavior to the leg for jumping and landing motions.

An aspect that has not been in the research focus yet is the interplay between both an intrinsically compliant knee and a biarticular spring in a multi-segmented leg. No systematic and comparative research exists so far comparing multiple compliant elements in highly under-actuated segmented legs, specifically for the combination of leg-angle and virtual leg axis compliance. As energy fluctuates in both directions in animal legs (Alexander, 1984) one can expect that compliant passive mechanisms evolved benefiting from these resources, i.e. energetically. We focus our research on the torque influence onto a series elastic biarticular spring and the increase in energy efficiency the additional stored energy provides.

In this paper we present a leg design with compliance in virtual leg axis direction as well as in leg angle direction. We show that the element in leg angle direction charges under torque influence, providing series elastic behavior for the hip. We show how the implementation of this element can drastically increase the amount of elastic energy stored in the leg.

2. MATERIALS AND METHODS

2.1. Leg Design and Implementation

The bio-inspired leg designs under investigation (**Figure 1**) consist of three segments. A femur segment, a shank segment with a four-bar structure and a foot segment. The arrangement of segments and elastic elements (**Figure 2B**), is inspired by the

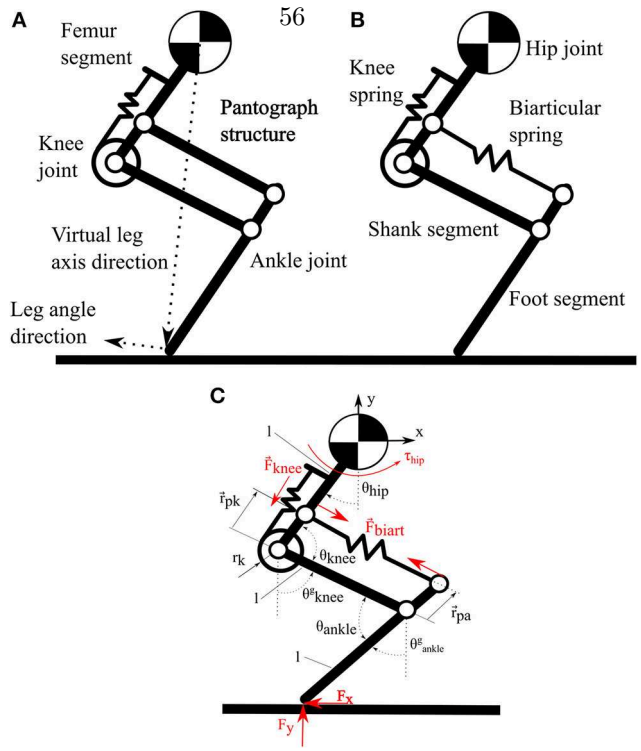

FIGURE 1 | (A) Pantograph leg with one spring around the knee joint and a rigid pantograph segment. **(B)** Biarticular leg with one spring around the knee joint and segment with a biarticular spring. **(C)** Schematic of three segment leg with angle definitions and symbol annotations.

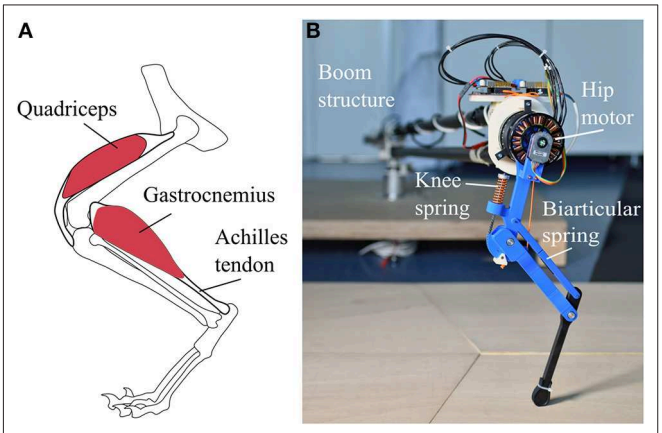

FIGURE 2 | (A) Schematic representation of cat leg anatomy for comparison. Quadriceps originates on the upper femur and inserts into the shank segment via the patella. Gastrocnemius originates on the lower femur and inserts into the upper foot segment spanning the knee and ankle joint. **(B)** Photo of the biarticular leg mounted on the boom structure. The biarticular spring is hidden under the two parts of the biarticular segment; the point of contact is visible as a slit.

leg anatomy of mammalian quadrupeds. The hip joint connects the femur segment to the trunk; the knee joint connects the hip and shank segment; the ankle joint connects the shank and foot segment. Leg segments on these legs represent the major bone groups of vertebrate animals, namely femur, tibia and fibula and the bones forming the foot segment. For simplicity, all segments in both designs have the same length.

Elastic elements in the robot are placed to mimic the functionality of big muscle-tendon groups in animal legs. In placement and functionality, the knee spring represents the quadriceps femoris muscle and patella tendon of the biological example. On the pantograph leg the segment parallel to the shank segment of the biarticular leg consists of a rigid element, forming a pantograph structure (**Figure 1A**). The segment parallel to the shank segment consists of a second spring connecting the hip and foot segment (**Figure 1B**). This biarticular segment spans two joints. The biarticular spring models the lower leg muscle-tendon apparatus of gastrocnemius muscle and Achilles tendon in a quadrupedal animal (**Figure 2A**). We refer to the leg with the biarticular spring as the *biarticular leg*, to the leg with the pantograph structure as the *pantograph leg*. Unloaded, the biarticular segment has the same length as the shank segment. The femur and foot segment are parallel when the biarticular spring is not deflected. Both ankle and knee joint have a hard stop to prevent over-extension.

The knee joint stiffness is realized by a spring that wraps around the knee joint on a cam mechanism, inspired by a knee cap (patella) (Allen et al., 2017; Heim et al., 2018). The knee cam mechanism linearizes the knee spring deflection over knee angle. Knee stiffness is designed to provide sufficient torque to hold the leg during running, exerting ground reaction forces three times the body weight of the robot at 10% virtual leg length deflection. The cam radius on the knee is designed to enable 35° knee angle deflection or about 70 mm leg length change. Empirically, we choose the biarticular spring stiffness similar to the knee spring stiffness, so the biarticular spring does not saturate, and the knee spring deflects similar to the pantograph leg. Through the biarticular spring the ankle joint can deflect by 60°, equivalent to 160 mm leg length change. The hip joint is articulated with a brushless motor. In combination with a 5:1 planetary gear box the nominal output hip torque is 6.2 Nm. To measure the joint deflections, all joints on the leg are instrumented with rotary absolute encoders.

The leg design consists of a hip joint rigidly connected to an actuator and two passive joints on the knee and ankle. The leg design builds on previous research on the Cheetah cub and Oncilla robots. Instead of the servo motors used in the previous robots we implemented high torque density brushless motors. To increase the backdriveability of the gear train a low ratio gearbox was used. This way the actuator can potentially be used as a proprioceptive actuator (Seok et al., 2012).The new knee spring placement in our design largely reduces the nonlinearity of the spring force to joint angle relationship of the knee joint, compared to the design used in Cheetah-cub and Oncilla. This simplifies modeling, and reduces the complexity of the mechanical design. The general mechanical leg design was improved to be more dureable and robust while at the same time reducing the complexity of the design to enable faster prototyping, as well as simplified manufacturing and assembly. The biarticular leg's lightweight

design can be rapidly assembled, and is largely made from elements created by rapid prototyping. At the same time it is however robust, and passively withstands drops from 200% body height.

Here we reduce the investigation to a single leg hopping in the saggital plane. This is common practice (Semini et al., 2008; Hutter et al., 2011; Ding and Park, 2017; Liu et al., 2018). It also reduces the effects of body inertia, multiple legs and the system complexity.

2.2. Kinematic Model

In this section, we investigate the governing equations describing the difference in behavior for both legs. All future assertions talk about joint angles implying resultant deflection of the associated elastic elements.

By formulating the kinematic equations for the pantograph leg, we show the basics of our hypothesis. Writing the forward kinematics to obtain the foot position with the reference at the hip joint, shows that the system rank $r = 2$ with 2 parameters (θ_{hip} and θ_{knee}), since $\theta_{hip} = \theta_{ankle}$ because of the pantograph structure. The equation system is fully defined. In comparison, the pantograph segment in the biarticular leg is replaced by a biarticular spring. The rank of the system matrix is also $r = 2$ but because $\theta_{hip} \neq \theta_{ankle}$, an additional parameter or Degree Of Freedom (DoF) is added to the leg. Annotations are depicted in **Figure 1** and **Table 1**.

Forward kinematics for pantograph leg:

$$\begin{aligned} x_{foot} &= -2 \cdot l \cdot \sin(\theta_{hip}) + l \cdot \sin(\theta^g_{knee}) \\ y_{foot} &= -2 \cdot l \cdot \cos(\theta_{hip}) - l \cdot cos(\theta^g_{knee}) \\ \text{for } \theta^g_{hip} &= \theta^g_{ankle} \end{aligned} \tag{1}$$

TABLE 1 | Leg parameters and robot implementation components.

Leg Parameters		
Segment length	l	150 mm
Knee and ankle resting angle		127°
Resting leg length		408 mm
Knee cam radius	r_k	30 mm
Knee - pantograph insertion distance	r_{pk}	30 mm
Mass biarticular leg		0.91 kg
Mass pantograph leg		0.88 kg
Knee spring stiffness	k_k	10.89 $\frac{N}{mm}$
Biarticular spring stiffness	k_{biart}	9.8 $\frac{N}{mm}$
Implementation		
Motor	TMotors MN7005 KV115	$m = 188\,g$, $\tau_{max} = 1.3\,Nm$
Motor driver	TI TMS320x2806x	24 V/15 A max.
Computer	pre-empt Ubuntu 14.04	1 kHz control frequency
Joint encoders	Broadcom AEAT8800-Q24	12-bit
Planetary gearbox	Matex RS3505S	gear ratio = 1:5

*Schematic in **Figure 1C**.*

Forward kinematics for biarticular leg:

$$\begin{aligned} x_{foot} &= -l \cdot \sin(\theta_{hip}) + l \cdot sin(\theta^g_{knee}) - l \cdot sin(\theta^g_{ankle}) \\ y_{foot} &= -l \cdot cos(\theta_{hip}) - l \cdot cos(\theta^g_{knee}) - l \cdot cos(\theta^g_{ankle}) \end{aligned} \tag{2}$$

with,

$$\begin{aligned} \theta^g_{knee} &= \pi - \theta_{hip} - \theta_{knee} \\ &\text{and} \\ \theta^g_{ankle} &= \theta^g_{knee} - \theta_{ankle} \end{aligned} \tag{3}$$

To describe the joint positions of the biarticular leg, an additional kinetic constraint is necessary to describe the coupling of the two springs:

$$\begin{aligned} &\frac{r_k \cdot F_{knee} + \overrightarrow{r_{pk}} \times \overrightarrow{F_{biart}} + \tau_{hip}}{cos(\theta^g_{hip})} = \\ &\frac{F_x \cdot l \cdot cos(\theta^g_{ankle}) + \overrightarrow{r_{pa}} \times \overrightarrow{F_{biart}}}{cos(\theta^g_{ankle})} \end{aligned} \tag{4}$$

$$\begin{aligned} F_{knee} &= k_k \cdot r_k \cdot \Delta\theta_{knee} \\ \overrightarrow{F_{biart}} &= k_p \cdot \left(\overrightarrow{P} - l \cdot \frac{\overrightarrow{P}}{||\overrightarrow{P}||} \right) \end{aligned} \tag{5}$$

where, F_{knee} is the force of the knee spring and F_{biart} is the force of the biarticular spring.

$$\overrightarrow{r_{pa}} = \begin{bmatrix} -||r_{pk}|| \cdot sin(\theta^g_{ankle}) \\ ||r_{pk}|| \cdot cos(\theta^g_{ankle}) \end{bmatrix}$$

and

$$\overrightarrow{P} = \begin{bmatrix} ||r_{pk}|| \cdot sin(\theta_{hip}) - l \cdot sin(\theta^g_{knee}) - ||r_{pa}|| \cdot sin(\theta^g_{ankle}) \\ -||r_{pk}|| \cdot cos(\theta_{hip}) - l \cdot cos(\theta^g_{knee}) + ||r_{pa}|| \cdot cos(\theta^g_{ankle}) \end{bmatrix} \tag{6}$$

where, $\overrightarrow{P}$ is the position of the biarticular spring insertion into the foot segment. If we assume the hip and foot fixed with rotary joints to a global frame (**Figure 3**), the pantograph leg cannot change its joint angles. Because of the increased DOF, the biarticular leg has an infinite number of joint orientations with a fixed hip and foot point.

By changing the torques and forces acting on the biarticular leg, the joint orientation can be changed based on the ratio of chosen stiffnesses. Under hip torque, the pantograph leg increases the forces on the hip and foot bearings but does not change joint angles. The biarticular leg orients its joints to satisfy the kinetic constraint described above. By changing its posture, the biarticular leg deflects the springs attached to each joint. When torque is exerted on the hip, the leg can store the energy from hip actuation in the biarticular spring. The energy storage potentially enables the biarticular leg to recuperate the energy stored in the springs.

Spring energies for leg comparison are calculated as:

$$E_{knee} = \frac{(F_{knee})^2}{2 \cdot k_k} \quad \text{and} \quad E_{biart} = \frac{(F_{biart})^2}{2 \cdot k_{biart}} \tag{7}$$

where E_i is the energy stored in the corresponding spring, k_k is the knee spring stiffness, and k_{biart} is the biarticular spring stiffness.

2.3. Experiments

During locomotion, legs are subject to dynamic forces in leg length, as well as leg angle direction. In this section, we investigate the behavior of both legs under loads in both directions. To show the basic functionality of the leg we reduce the experiment complexity compared to a hopping experiment. In a reduced order experiment, we investigate the effects of virtual leg axis and leg angle forces separately. Then we investigate the mechanical leg behavior in a vertical drop test without control influence. Last we show that the leg shows series elastic behavior in the biarticular spring under combined loads during dynamic hopping to provide a realistic locomotion load case.

2.3.1. Static Virtual Leg Axis Forces

First, we implemented a simplified setup neglecting weight and inertia effects to show the virtual leg axis related behavior clearly. Both the foot and hip joint of each leg were fixed to a ground frame by a rotational pin joint. The joint restricts both hinge points to one rotational DOF (**Figure 3A**). The ball bearings used to implement the rotational joints only allow forces to be transmitted, but no torques. This experiment investigates the change in joint angles purely based on change in virtual leg length. We fixed the hip joint to the frame at different virtual leg lengths in steps of 5 mm from resting leg length to 65 mm deflection, and measure the joint angles with rotary encoders.

2.3.2. Static Leg Angle Torque

In the next step, we investigated the effects of hip torques on the legs in the static test stand and observed the joint angles. The legs were fixed to the same static test setup as before. Both legs were deflected by 10 mm initial leg length. We applied hip torque from 0 to 2.5 Nm in steps of 0.1 Nm every 2 s to exclude acceleration effects. We measured the resulting joint deflections as well as the forces exerted onto the foot fixture with a force sensor (K3D60 me-systeme) to verify the applied hip torque.

2.3.3. Vertical Drop Experiment

After investigating the static behavior of the leg, we focus on behavior under dynamic loading. We separate the effects of virtual leg axis forces and leg angle torque for vaulting the leg during forward hopping with a vertical drop experiment (**Figure 4**). Holding the legs at a defined position requires a motor with a position controller. We want to investigate only the mechanical response to dynamic virtual leg axis forces without effects induced by the controller. Using the same controller could give advantages to one leg, and different controllers for both legs would be hard to compare in mechanical performance. To eliminate this potential bias, we implemented a virtual spring on the hip actuator. The motor mimicked a torsion spring between the hip joint and a global frame. The virtual spring had its set point at 17.5°. At this hip angle, the virtual leg was vertical. By using a direct drive motor as a proprioceptive actuator, we avoided measuring inaccuracies through backlash, friction and reflected gearbox inertia. We used the proprioceptive actuator as a sensor to directly measure the hip angle deflection as well as the resulting forces. The virtual spring stiffness was chosen at 5.8 $\frac{Nm}{rad}$ to match the position controller gains used for the hopping experiments. The leg with joint encoders was connected to a boom structure, to restrict the motion to the sagittal plane (**Figure 2B**). The boom was instrumented with rotational encoders (AMT 102-V) to measure the horizontal and vertical angle of the boom representing the center of mass

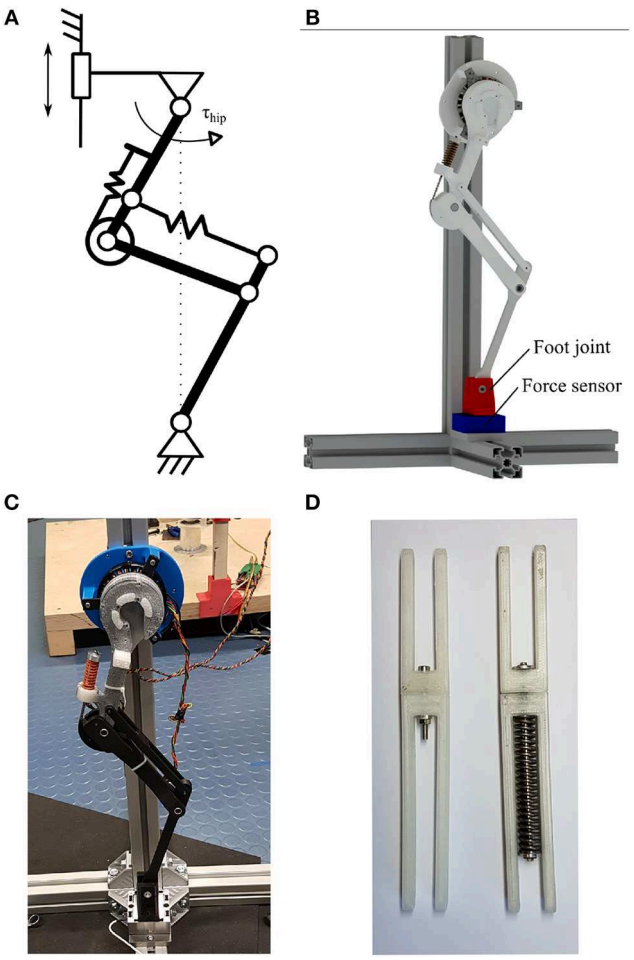

FIGURE 3 | (A) Schematic of the static setup with the leg fixed to a rotary joint on the foot and hip joint. The fixture on the hip joint can be moved to change the virtual leg length. **(B)** Rendered leg test stand with leg fixed into a rotary joint on hip and foot (red) on force sensor (blue). **(C)** Photo of the static test setup with pantograph leg. **(D)** Photo of the two four-bar segments. Fixed pantograph segment left. Spring loaded biarticular segment right.

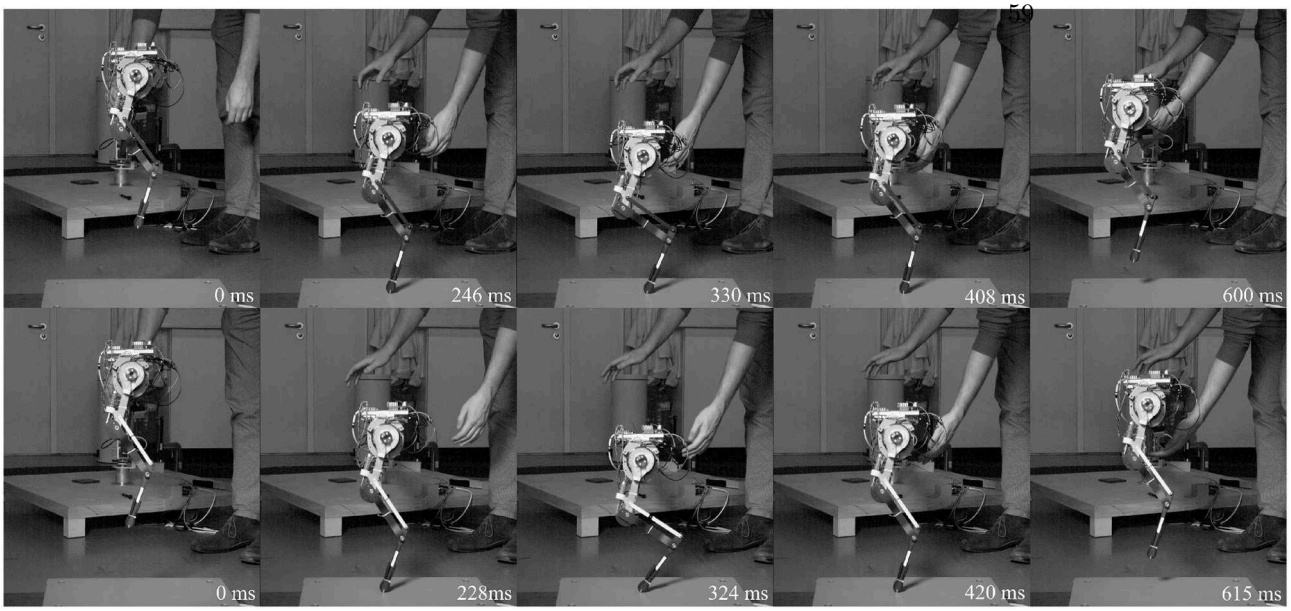

FIGURE 4 | High-speed snapshots of drop experiment starting from release to next hip apex. Pantograph leg top row, biarticular leg bottom row. Depicted are left to right: drop, touchdown, maximum deflection, lift-off, apex. Due to the difference in leg stiffness the biarticular leg is on the ground longer. The biarticular spring deflection is visible in the gap in the biarticular segment at maximum deflection. The biarticular leg adapts its posture, visible in the difference of hip and foot segment angles at maximum deflection.

position of the robot. The biarticular and pantograph leg were dropped from 590 mm hip height. We also measured the input power consumption of the motor driver with a current sensor (ACS713). Data was normalized from hip dropping height to first apex (drop cycle) and averaged over 30 drops and displayed with a 95% confidence interval. Touchdown and liftoff are determined by when the spring-loaded ankle joint starts to deflect. Hip torque is calculated from armature motor current as:

$$\tau_{hip} = i_{armature} \cdot k_t \tag{8}$$

where $i_{armature}$, is the armature motor current measured on the motor driver and k_t is the torque constant of the motor. Electrical system input power is calculated as:

$$P_{el} = U \cdot I \tag{9}$$

where U, is the constant power supply voltage of 24 V, and I is the input current measured from the power supply.

2.3.4. Hopping Experiment

In this section, we provide a realistic locomotion showcase to investigate the behavior of a single hopping leg under a combination of virtual leg axis forces and hip torque. The leg is again constrained to movement in the sagittal plane without trunk rotation by a boom structure.

Both legs use the same gearbox in this experiment to provide enough hip torque for forward locomotion. We implemented a sine wave position controller on the hip actuation resulting in a hopping gait,

$$\theta_{hip_{desired}} = \theta_0 + \theta_1 \cdot sin(2 \cdot \pi \cdot f \cdot t) \tag{10}$$

where $\theta_{hip_{desired}}$ is the desired hip angle, θ_0 is the hip angle offset, θ_1 is the hip angle amplitude, f is the hopping frequency, and t is time.

A PD position controller calculates the desired current for the low level current controller on the motor driver according to:

$$\begin{aligned} i_{motor_{desired}}(t) = kp \cdot (\theta_{hip}(t)_{desired} - \theta_{hip_{encoder}}(t)) \\ + kd \cdot \frac{d(\theta_{hip}(t) - \theta_{hip_{encoder}}(t))}{dt} \end{aligned} \tag{11}$$

where kp and kd are the proportional and derivative control gains, $\theta_{hip_{encoder}}$ is the measured hip angle and $\theta_{hip_{desired}}$ is the desired hip angle calculated above.

The PD position controller, schematic in **Figure 5**, of the motor driver was tuned to the same gains for both legs. Gait parameters θ_0, θ_1 and the gait frequency were hand-tuned for both legs to get stable hopping at 0.99 $\frac{m}{s}$ with $f = 2\,Hz$ for the pantograph leg and 0.95 $\frac{m}{s}$ with $f = 2.2\,Hz$ on the biarticular leg (**Figure 6**). We defined stable hopping, if the robot hopped for trials longer than 2 min, equivalent to $\cong$ 240 steps. Data was collected 1 min after the robot achieved a stable gait. All data sets were normalized over time from hip apex to apex (step cycle) and averaged over 30 consecutive steps. Average data was displayed with 95% confidence intervals. Touchdown and liftoff were determined by when the springs started to deflect. All future discussions are conducted using the averaged data set to get a representative picture.

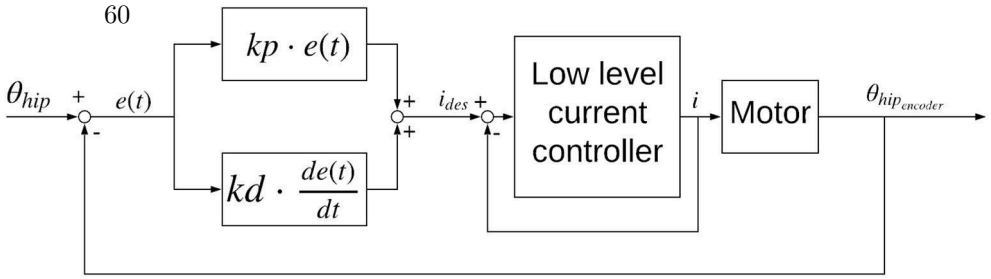

FIGURE 5 | Controller diagram for the sine wave position controller used for the forward hopping experiments.

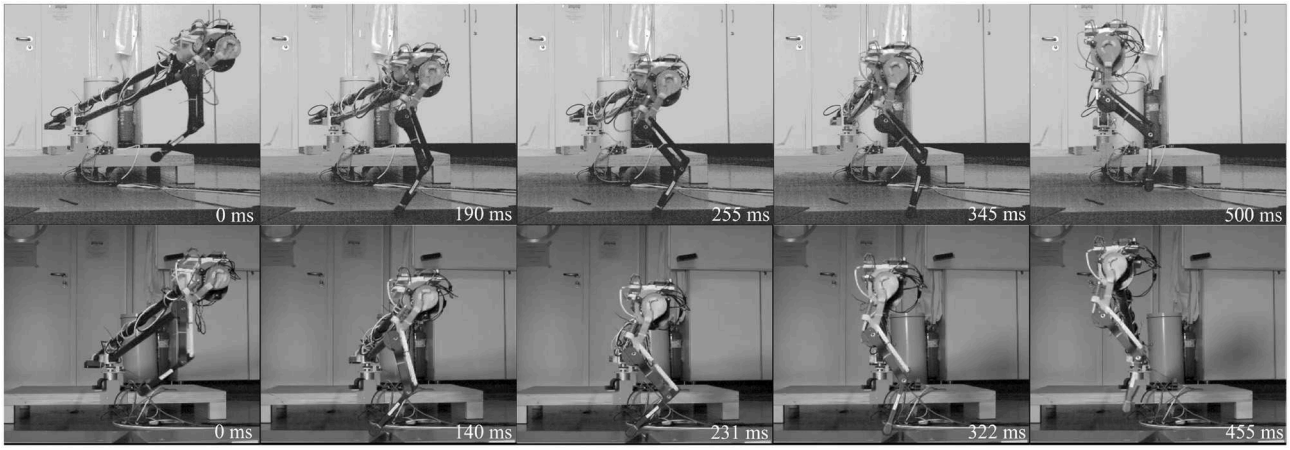

FIGURE 6 | High-speed snapshots of both legs hopping forward for one step cycle from release to hip apex. Pantograph leg top row, biarticular leg bottom row. Depicted are left to right: hip apex, touchdown, maximum deflection, lift-off, second hip apex. The biarticular spring deflection is visible in the gap in the biarticular segment at maximum deflection. Timing differences stem from different duty factor and difference in gait frequency, 2 Hz for the pantograph leg and 2.2 Hz for the biarticular leg. The biarticular leg adapts its posture, visible in the difference of hip and foot segment angles at maximum deflection.

3. RESULTS

In this section, we present data and results from the static leg force experiment, the static leg angle torque experiment, the vertical drop test, and the hopping experiment.

3.1. Static Virtual Leg Axis Forces

In the pantograph leg knee and ankle angles change equally (**Figure 7**). Because of the parallelogram geometry in the leg's four-bar mechanism, knee and ankle angles are kinematically coupled to be equal. Play in the joints causes the small deviation between the pantograph knee and ankle angle curve. At maximum leg deflection, the pantographs knee and ankle angle deflect by 32°. The model prediction fits the data, neglecting the small deviation of $\leq 2°$.

In the biarticular leg, the change in knee and ankle angles are not equal. Because of the biarticular spring, the ankle deflects more than the knee. At maximum deflection the ankle in the biarticular leg deflects by 43°, the knee deflects to 17°. This first experiment shows, that knee and ankle are not kinematically coupled in the biarticular leg.

3.2. Static Leg Angle Torque

Because of the kinetic coupling, both knee and ankle angles in the biarticular leg change under torque (**Figure 8**). The knee on the biarticular leg deflects by 0.85°, the ankle deflects by 2.1° at 1 Nm. The model shows a reasonable prediction for the angles. The deviation and flat line stem from the hard stop at the knee preventing the knee from over-extension over its resting angle. The hard stop effect is less pronounced in the experimental data, due to material elasticity. It can be seen as a change in slope. The knee model is only valid before hitting the hard stop.

This experiment shows that the biarticular leg can store energy from hip actuation in the biarticular spring. Under the influence of hip torque, the pantograph leg does not change its joint angles.

We argue that the distal elastic element, mimicking the lower leg muscle-tendon structures, acts to the hip actuation like a serial spring. Different from a classical SEA hip actuator, the biarticular spring has components acting in both virtual leg axis and leg angle direction. The ratio of components that act in virtual leg axis or leg angle direction depends on the virtual leg deflection as well as the resting joint angles, and the chosen spring stiffnesses.

These experiments are abstracted from the behavior of a hopping robot. Under only hip torque the data shows that the

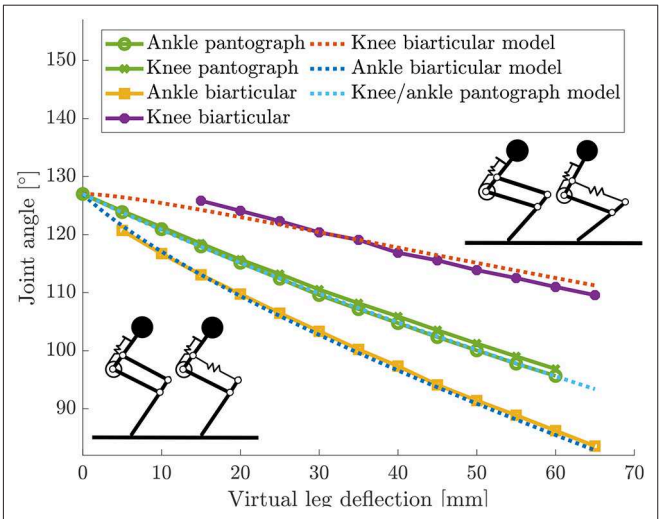

FIGURE 7 | Knee and ankle angle changes from resting angles for the static virtual leg axis experiment with both legs. Knee and ankle angles change equally because of the pantograph structure. Knee and ankle change not equally in the biarticular leg because of the additional degree of freedom.

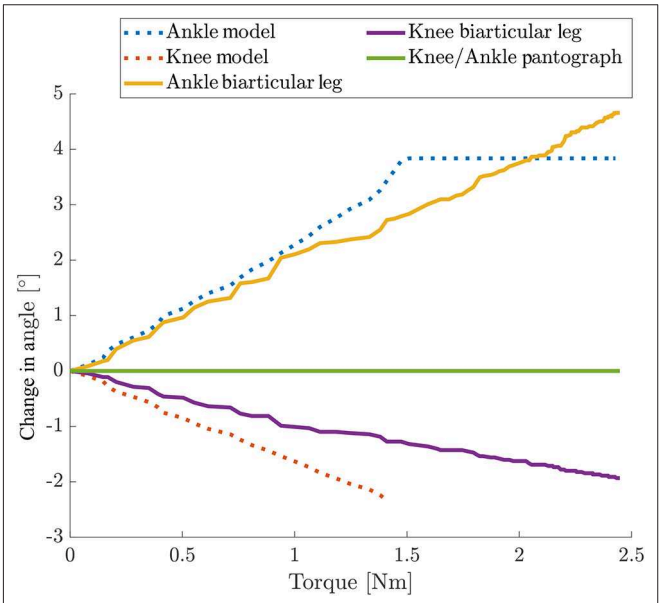

FIGURE 8 | Change in joint angles over hip torque for static torque experiments for both legs. Hip position was fixed to 10 mm leg length deflection. Cut-off in the model are due to the hard stop to prevent over-extension. The cut in the experimental data is only visible as a change in slope due to material elasticity. The knee model is only valid until hitting the hard stop. Only the biarticular leg can deflect its joints under torque. Because of the kinematic coupling in the biarticular leg the knee joint deflects as well when the ankle joint deflects. In the pantograph leg the hip and foot position are fixed and the joints do not deflect under hip torque.

biarticular leg has the ability to store hip actuator energy in its springs. For any given initial posture the leg can adapt its posture and store energy in the springs that can potentially be recuperated.

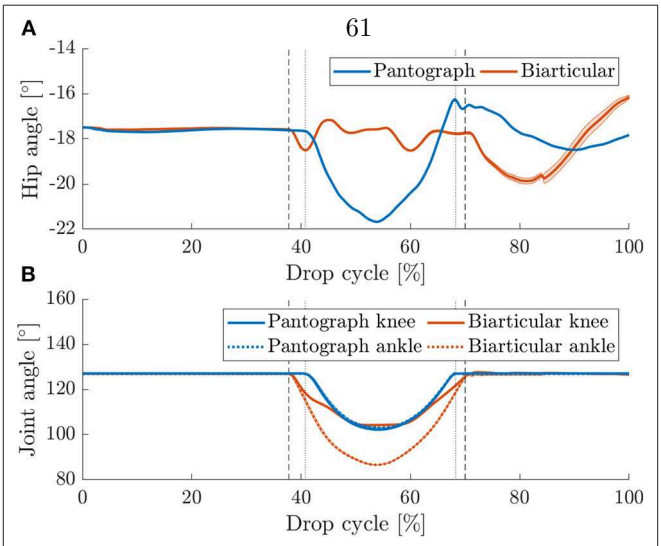

FIGURE 9 | (A) Mean hip angles for vertical drop tests with 95% confidence levels for 30 drops. Vertical touchdown (left markers) and lift off (right markers) are plotted for pantograph (dotted) and biarticular leg (dashed). The hip angle in the pantograph leg deflects more since the leg only has one DOF. In the biarticular leg, the leg can adapt its internal posture to mitigate the dynamic forces without reflecting them to the femur segment. **(B)** Knee and ankle angles for vertical drop tests with 95% confidence levels for 30 consecutive steps. Touchdown (left markers) and lift off (right markers) are plotted for pantograph (dotted) and biarticular leg (dashed). In the pantograph leg, the joints deflect symmetrically because of the kinematic coupling. In the biarticular leg, the ankle joint deflects nearly twice as much as the knee joint.

3.3. Vertical Drop Test

During stance phase (**Figure 4**), the hip angle in the pantograph leg deflects by 4° (**Figure 9A**). Because of the kinematic coupling in the leg, any force reflects into the femur segment and change its angle. Because the biarticular leg has one more DOF, it adepts its posture (**Figure 9B**). By changing the ankle and knee angle the energy is stored in the springs, and the data shows that the hip angle does not change. As the virtual spring induces a torque when the hip angle deflects from its resting angle, the hip torque during stance phase is much higher in the pantograph leg (**Figure 10**). Hip torque is calculated from measured armature current on the motor. The hip torque on the pantograph leg peaks at 2.2 Nm while the biarticular leg peaks at 1.2 Nm. Mean hip torque during stance phase is 0.22 and 0.02 Nm for the pantograph leg and biarticular leg, respectively.

After liftoff, the torque requirement is higher in the biarticular leg. Since the biarticular leg has elastic components in leg angle direction, a force resulting from unloading the joint rapidly reflects into leg angle direction. The virtual leg shoots forward when the two parts of the biarticular spring mount collide due to the hard stop. The collision can also be seen in the provided high-speed video in the **Supplementary Material**. We ignore this reflection effect and do not compensate or utilize it here.

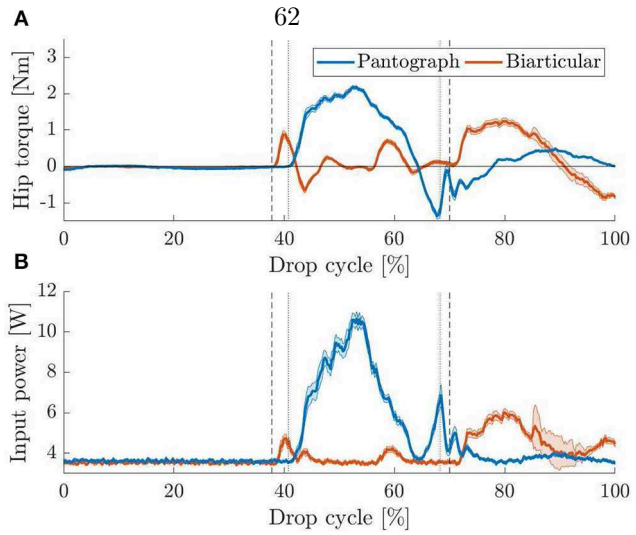

FIGURE 10 | (A) Mean hip torque for vertical drop tests with 95% confidence levels. Vertical touchdown (left markers) and lift off (right markers) are plotted for pantograph (dotted) and biarticular leg (dashed). Because the biarticular leg adapts its posture forces are reflected less into the femur segment. This results in a lower hip angle change. **(B)** Mean input power consumption for vertical drop tests with 95% confidence levels. Touchdown (left markers) and lift off (right markers) are plotted for pantograph (dotted) and biarticular leg (dashed). In the biarticular leg, the motor does not need additional power to counteract the force. The biarticular leg requires 81% less motor power to hold the virtual leg vertical.

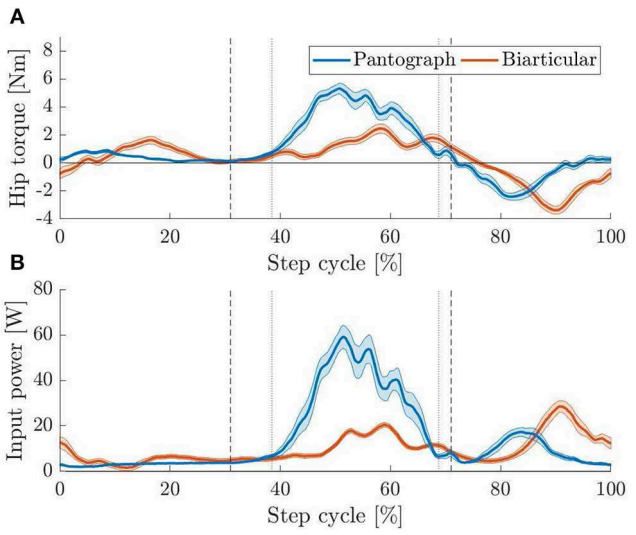

FIGURE 11 | (A) Apex to apex normalized hip torque for both hopping legs with 95% confidence levels. Vertical touchdown (left markers) and lift off (right markers) are plotted for pantograph (dotted) and biarticular leg (dashed). During stance phase, the biarticular leg has lower peak torque requirements than the pantograph leg. After toe-off, the hip actuator suppresses the reflection effect shooting the leg forwards. The active suppression increases the torque requirements on the biarticular leg. **(B)** Apex to apex normalized input power for both hopping legs with 95% confidence levels. Touchdown (left markers) and lift off (right markers) are plotted for pantograph (dotted) and biarticular leg (dashed). Peak input power for the biarticular leg is 50% lower and mean power is 31% lower than on the pantograph leg.

The duration of stance phase varies between 27% on the pantograph leg and 32% on the biarticular leg. We suspect the difference is due to the higher mass of 29 g as well as the lower global leg stiffness of the biarticular leg.

Because the drop experiment is not a periodic motion, the beginning and end points of the graphs do not match as the leg moves differently for the subsequent lower hops.

As a result of the higher torque requirement, input power shows that the biarticular leg needs less power during stance phase to keep the desired leg posture (**Figure 10B**). After liftoff, the same rise in power that was explained in the hip torque curve is visible. Since oppressing the reflection effect requires high torque at high speed, a drastic rise in power consumption during swing phase is visible. Over the full step cycle, mean power consumption for the pantograph leg is 4.6 and 3.9 W for the biarticular leg. Mean power consumption during stance phase for the pantograph leg is 6.8 W for the pantograph leg and 3.7 W for the biarticular leg. The biarticular leg shows a 46% lower power requirement during stance phase and 15% lower power consumption over the whole drop cycle.

3.4. Hopping Experiment

During forward hopping (**Figure 6**) the pantograph and biarticular leg show a similar trend in torque requirements as during the drop experiments before. Hip torque is calculated from measured armature current on the motor. The peak hip torque is 5.2 Nm for the pantograph leg and 2.4 Nm for the biarticular leg (**Figure 11A**). The mean torque for the pantograph leg is 0.19 and 0.05 Nm for the biarticular leg over the whole step cycle. Mean torque during stance phase is 0.64 and 0.23 Nm for pantograph and biarticular leg respectively. The biarticular leg requires 53% less peak torque and 74% less mean torque.

The previously observed difference in knee and ankle angle between the two legs is also visible during hopping (**Figure 12**). While the pantograph leg deflects both joints the same way due to the four-bar geometry, the ankle deflects more in the biarticular leg. We observe that the knee joints deflect similar in both legs. The biarticular leg's ankle, however, deflects by 45° compared to 19° on the pantograph leg.

Additionally, the duty factor, the fraction of stance phase over one step cycle period:

$$d_{duty} = \frac{t_{Stance}}{T_{Stepcycle}} \tag{12}$$

is much smaller at 31% of step cycle on the pantograph leg than on the biarticular leg where the duty factor is 40%. We assume the duty factor to be higher due to the lower global leg stiffness resulting in an extended stance phase duration as the leg deflects more.

Power requirements during hopping are higher in the pantograph leg than in the biarticular leg (**Figure 11B**). The pantograph leg power peaks at 60 W where the biarticular leg

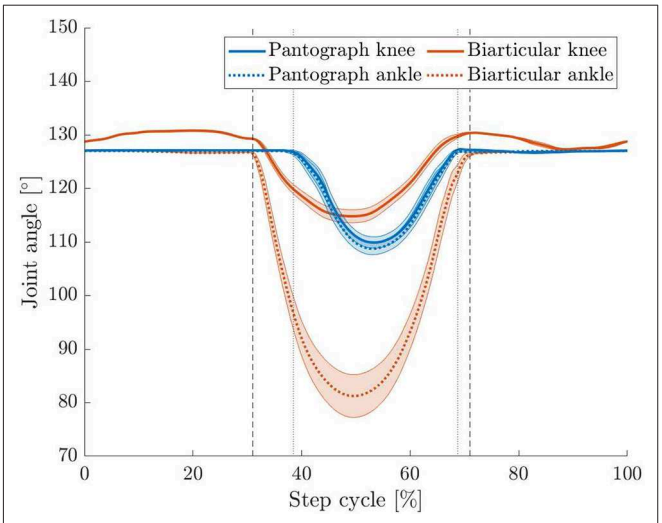

FIGURE 12 | Apex to apex normalized knee and ankle angle change for both hopping legs with 95% confidence levels. Vertical touchdown (left markers) and lift off (right markers) are plotted for pantograph (dotted) and biarticular leg (dashed). The ankle angle in the biarticular leg deflects by 45° compared to 19° on the pantograph leg. Knee angles for both legs deflect similarly.

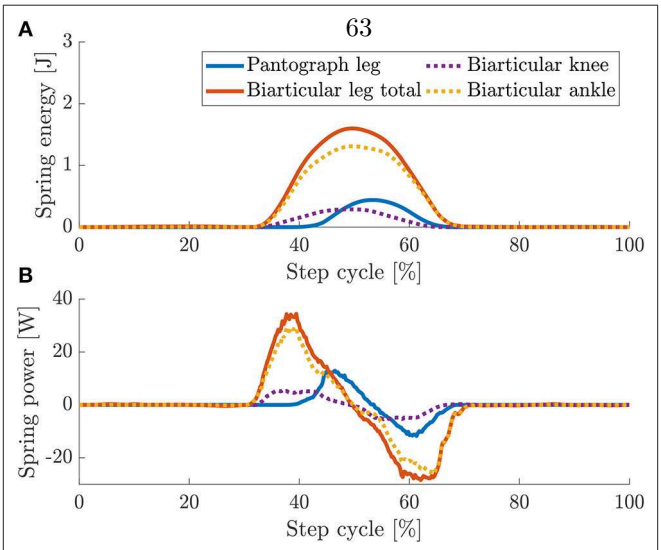

FIGURE 13 | Overlaid spring energy **(A)** and power **(B)** for pantograph (orange) and biarticular leg (blue). Due to the biarticular spring, the energy stored in the biarticular leg is 2.4 times higher. In dashed lines are shown the energy and power for the individual springs in the biarticular leg. The energy in the knee springs are similar. The energy in the biarticular spring is ≅ 80% higher.

peaks at 20 W during stance and 30 W during swing because of the reflection effect. Mean input power for the pantograph leg is 14.1 and 9.7 W for the biarticular leg. Mean power requirement on the biarticular leg is 31% lower and peak power requirement is 50% lower. The difference in input power requirement is evident in the cost of transport (COT) (Tucker, 1975),

$$COT = \frac{P_{in}}{m \cdot g \cdot v}, \tag{13}$$

where P_{in} is electrical input power to the motor driver, m is the robot mass, g is the gravitational acceleration, and v is the forward speed of the robot.

Total COT is calculated using overall input power. The total COT for the pantograph leg is 1.7 compared to the biarticular leg at 1.2. COT when substracting 3 W idle power consumption of the system is 1.3 for the pantograph leg and 0.8 for the biarticular leg. To investigate this further, we calculate the energy stored in the springs and compare the two leg designs.

In the biarticular leg, the overall stored energy is considerably higher than the energy stored in the pantograph leg (**Figure 13A**). The maximum total spring energy in the pantograph leg is 0.45 J vs. 1.56 J for the biarticular leg. Mean spring energy is 0.06 and 0.34 J for the biarticular leg. Total spring energy for the pantograph leg is 82% lower than for the biarticular leg. As we show a higher energy efficiency in the biarticular leg, we conclude that the leg design has a higher recuperation rate. Higher recuperation means the biarticular leg can use the energy stored in its spring more effectively for locomotion.

To get a clearer picture on the biarticular leg, we also plot the individual spring energy contribution to the total energy. The knee spring on the biarticular leg stores roughly the same amount of energy as the single knee spring in the pantograph leg. Peak knee spring energy for the biarticular leg is 0.28 J and peak biarticular spring energy is 1.3 J. Mean energy stored in the biarticular knee is 0.05 J and mean biarticular spring energy is 0.28 J. The mean energy stored in the biarticular spring is 82% higher.

The maximum power in the springs in the pantograph leg at 12.9 W is 62% lower than in the biarticular leg at 34.4 W (**Figure 13B**). The maximum released power is −11.7 W in the pantograph leg and −28.4 W in the biarticular leg.

The biarticular leg stores more energy in its springs due to the elasticity in both virtual leg axis and leg angle direction. In the vertical drop and the dynamic hopping experiment, the leg recuperates more energy from its springs which reduces the overall required torque and input power.

4. DISCUSSION

The placement and functional morphology of elastic elements in legs is an important research question in legged locomotion. In this paper, we show that the biarticular spring, which mimics the elasticity of lower leg muscle tendon structures, has elastic components that can provide series elastic behavior to hip actuation. In a model as well as a static experiment, we show *how* the biarticular spring enables the leg to deflect its joints at a fixed leg length without changing the hip and foot position. We then show that the additional degree of freedom allows the leg to store energy provided by hip actuation in this elastic element. In a vertical drop test with a virtual spring on the hip, we show that the favorable lower peak torque and power consumption of series elastic behavior do not depend on the motor controller but result from leg mechanics. In the drop experiment, we show that the leg changes its internal posture to adapt to external forces

instead of reflecting these forces into the hip actuation. As the hip actuation does not need to compensate for the dynamic loading, no additional torque and power is required, which increases energy efficiency.

Last we show that in a combined load case of torque and virtual leg axis forces, the peak torque and power requirements are lower for the leg with distal series elastic components. By reducing the overall leg stiffness, the leg has a smaller leg length which acts as the lever arm for hip torque to produce a ground reaction force. The higher leg length deflection of 64 mm on the biarticular leg vs. 37 mm on the pantograph leg reduces the mean torque requirement for the leg.

Compared to the vertical drop the biarticular ankle joint during hopping deflects more by 6°, even though the hopping height, 490 mm for the pantograph leg and 470 mm for the biarticular leg, is lower than the drop height for the vertical drop. The difference in hip angle stems from the deflection under the additional hip torque to move the leg forward. The higher joint deflection is the result of the combined load case of virtual leg axis forces through dynamic loads and the torque required to vault the leg forward. As expected the biarticular leg stores energy provided by hip actuation in the biarticular spring even under a combined load case of virtual leg axis forces and leg angle torques.

Through the implementation of this elasticity, it is possible to reduce the peak power requirement by 26%, the mean power requirement by 31%, the peak torque requirement by 53% and mean torque by 71% in the hopping experiment.

We show that the biarticular leg with elastic components in leg angle direction possesses the same effects as a series elastic element, namely reduced torque and power requirements. We can, therefore, conclude, that the biarticular leg adds series elastic behavior to the leg. Because the biarticular spring stores 82% more energy we can further conclude that the biarticular spring also reduces the mean power and torque requirements of the leg. The reduced energy requirement shows that in robotic legs leg, compliance in leg angle direction is an equally important design parameter to virtual leg axis compliance.

To put the COT of our design into perspective, **Figure 14** shows the COT values for a selection of robots as well as the regression from Tucker (1975) for animal data over their respective masses. Both the pantograph leg as well as the biarticular leg are below the line for comparable natural runners. We include SPEAR (Liu et al., 2018) as a direct comparison to our monoped hopper. Comparing the COT without base consumption of our biarticular design to SPEAR, the COT of our design is lower at 0.8 than SPEAR at 0.86.

Since power, speed and mass do not scale linearly, as shown by Tucker, we believe that a better comparison than absolute COT numbers, is the comparison to a natural runner of comparable weight, the relative COT. We calculate the relative distance of the biarticular leg's COT to the COT of a model animal of the same weight from the Tucker linear regression.

Comparing the biarticular leg's total COT (including base consumption) to natural runners, the biarticular leg is still below the natural runners line and roughly on the same level as the pantograph leg without base consumption. The relative COT for the biarticular leg is 64% of a natural runner's COT. The relative

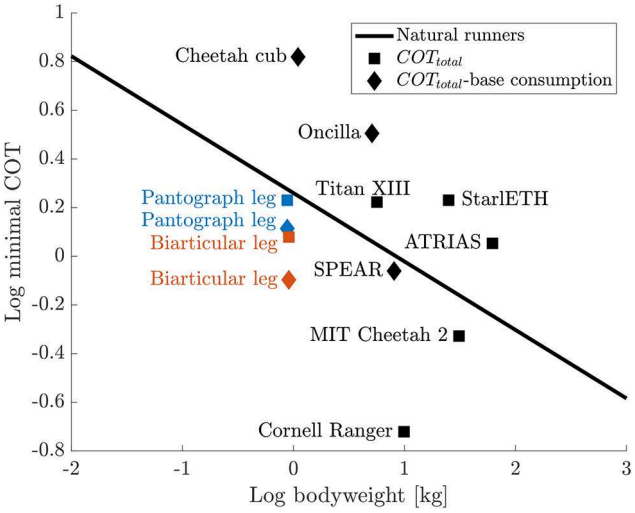

FIGURE 14 | COT comparison for a selection of legged robots (Spröwitz et al., 2013, 2018; Bhounsule et al., 2014; Hutter et al., 2014; Renjewski et al., 2015; Kitano et al., 2016; Park et al., 2017; Liu et al., 2018) compared to the regression for animal runners from Tucker (1975) that shows a linear regression for the minimal COT from various running animals. Total COT values are shown as squares, COT values where base consumption (communication, electronics, etc.) is substracted, are shown as diamonds. The COT for the biarticular leg is 64% of the COT of a natural runner with the same weight. All COT values are listed in the **Supplementary Material**.

COT without base consumption is at 43% of the comparable natural runner's COT. To the best of the authors' knowledge, the biarticular leg shows the lowest achieved relative cost of transport documented for all dynamically hopping and running robots including MIT Cheetah at 68% relative COT.

The only legged robots with a lower relative COT are Cornell Ranger (Bhounsule et al., 2014) and Cargo (Guenther and Iida, 2017). Cornell Rangers COT of 0.19 is 20% of the COT of a comparable natural runner. Cornell Ranger was optimized for COT efficient walking, unlike the here shown dynamic hopping locomotion of the biarticular leg. Cargos COT of 0.1 is 21% of a comparable natural runner. Cargo was designed to run at its natural frequency to increase COT. We exclude Cargo because of its non-practical ground clearance of (Guenther and Iida, 2017, **Figure 12**).

With the results of this paper we create a novel, robotic perspective on the placement and functional morphology of elastic elements in legs. Our research raises the question whether a transfer from the insights from this abstracted model back to biology is possible which has not been shown or discussed in previous research in biology. By showing the same joint deflection behavior under similar load cases, it might be possible, to verify the behavior of the biarticular leg in its natural role models. By finding similar behavior we could then conclude that the anatomy of vertebrate animals is in parts due to the functional morphology shown in this paper. During experimentation, we show a reflection effect that shoots the leg forward at the end of stance phase. While not considering the effect in this study, we will focus our future research on implementing controllers that

utilize the effect to further reduce power requirements during the swing phase.

Additionally, we will investigate whether the distal series elastic element increases robustness to perturbations. To follow up the findings in this paper we want to optimize the energy recuperation through an investigation into the effects of posture, segmentation and spring stiffness ratio on the elastic behavior of the leg.

5. CONCLUSION

In this paper, we investigate the effects of a distal biarticular elastic element. We show that a bio-inspired distal elastic element has components that deflect in leg angle direction. To characterize the leg we provide a mathematical model, to show the underlying behavior. We then investigate the leg behavior first under virtual leg axis forces. We show that the distal elastic element provides an additional degree of freedom to the leg. In a second step we investigate the leg behavior under only leg angle torque. The second experiment shows that the elastic components in leg angle direction deflect under hip torque and store hip actuator energy. Then we show that the leg can reconfigure its internal posture during a vertical drop experiment. The leg adapts its posture to the loading force, leading to a lower femur deflection. This decreases the power requirement during drop experiments by 46% compared to the leg with only virtual leg axis compliance. The leg angle actuator will therefore require less torque and power to hold the leg during stance. Last we show that the effects investigated in reduced complexity experiments are visible in a realistic monoped hopping experiment with combined leg angle torques and virtual leg axis forces. In the hopping experiment we show that the distal elastic element reduces the power requirements by 31% and the peak torque requirements by 71%. We record a 31% reduced COT of 1.2 for our leg design of 0.9 kg at 1 $\frac{m}{s}$. The relative COT of our biarticular leg design is 64% of a comparable natural runner's COT.

DATA AVAILABILITY

The datasets generated for this study are available on request to the corresponding author.

AUTHOR CONTRIBUTIONS

FR contributed to the concept, robot design, experimental setup, experimentation, data discussion, and writing. AB-S contributed to the concept, data discussion, and writing.

ACKNOWLEDGMENTS

We thank the Robotics ZWE under Dr. Fiene as well as the general workshop, the fine-mechanics workshop and the carpentry shop at MPI for providing components for the robots. We thank Felix Widmaier for programming the motor driver interface. We thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting the academic development of FR.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fnbot.2019.00064/full#supplementary-material

REFERENCES

Alexander, R. (1984). Elastic energy stores in running vertebrates. *Am. Zool.* 24, 85–94.

Alexander, R. M. (2002). Tendon elasticity and muscle function. *Compar. Biochem. Physiol. Part A Mol. Integrat. Physiol.* 133, 1001–1011. doi: 10.1016/S1095-6433(02)00143-5

Alexander, R. M. and Bennet-Clark, H. C. (1977). Storage of elastic strain energy in muscle and other tissues. *Nature* 265, 114–117. doi: 10.1038/265114a0

Allen, V. R., Kambic, R. E., Gatesy, S. M., and Hutchinson, J. R. (2017). Gearing effects of the patella (knee extensor muscle sesamoid) of the helmeted guineafowl during terrestrial locomotion. *J. Zool.* 303, 178–187. doi: 10.1111/jzo.12485

Beckerle, P., Verstraten, T., Mathijssen, G., Furnemont, R., Vanderborght, B., and Lefeber, D. (2017). Series and parallel elastic actuation: Influence of operating positions on design and control. *IEEE/ASME Trans. Mech.* 22, 521–529. doi: 10.1109/TMECH.2016.2621062

Bhounsule, P. A., Cortell, J., Grewal, A., Hendriksen, B., Karssen, J. G. D., Paul, C., et al. (2014). Low-bandwidth reflex-based control for lower power walking: 65 km on a single battery charge. *Int. J. Robot. Res.* 33, 1305–1321. doi: 10.1177/0278364914527485

Biewener, A. A. (1998). Muscle function *in vivo*: a comparison of muscles used for elastic energy Savings versus Muscles used to generate mechanical power1. *Am. Zool.* 38, 703–717.

Blickhan, R. (1989). The spring-mass model for running and hopping. *J. Biomechan.* 22, 1217–1227.

Daley, M. A. (2008). Biomechanics: running over uneven terrain is a no-brainer. *Curr. Biol.* 18, R1064–R1066. doi: 10.1016/j.cub.2008.09.050

Ding, Y., and Park, H.-W. (2017). "Design and experimental implementation of a quasi-direct-drive leg for optimized jumping," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Vancouver, BC: IEEE).

Fukuoka, Y., Kimura, H., and Cohen, A. H. (2003). Adaptive dynamic walking of a quadruped robot on irregular terrain based on biological concepts. *Int. J. Robot. Res.* 22, 187–202. doi: 10.1177/0278364903022003004

Geyer, H., Seyfarth, A., and Blickhan, R. (2006). Compliant leg behaviour explains basic dynamics of walking and running. *Proc. R. Soc. B Biol. Sci.* 273, 2861–2867. doi: 10.1098/rspb.2006.3637

Guenther, F., and Iida, F. (2017). Energy-efficient monopod running with a large payload based on open-loop parallel elastic actuation. *IEEE Trans. Robot.* 33, 102–113. doi: 10.1109/TRO.2016.2623342

Heim, S., Ruppert, F., Sarvestani, A. A., and Sprowitz, A. (2018). "Shaping in practice: training wheels to learn fast hopping directly in hardware," in *2018 IEEE International Conference on Robotics and Automation (ICRA)* (Brisbane, QLD: IEEE). doi: 10.1109/icra.2018.8460984

Hurst, J. W. (2008). *The Role and Implementation of Compliance in Legged Locomotion*. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA.

Hutter, M., Gehring, C., Bloesch, M., Hoepflinger, M. A., Remy, C. D., and Siegwart, R. (2012). "Starleth: a compliant quadrupedal robot for fast, efficient, and versatile locomotion," in *Adaptive Mobile Robotics*, eds A. K. M. Azad, N. J. Cowan, M. O. Tokhi, G. S. Virk, and R. D. Eastman (Baltiomore, MD: World Scientific), 483–490. doi: 10.1142/9789814415958_0062

Hutter, M., Gehring, C., Hopflinger, M. A., Blosch, M., and Siegwart, R. (2014). Toward combining speed, efficiency, versatility, and robustness in an autonomous quadruped. *IEEE Trans. Robot.* 30, 1427–1440. doi: 10.1109/TRO.2014.2360493

Hutter, M., Remy, C. D., Hoepflinger, M. A., and Siegwart, R. (2011). "High compliant series elastic actuation for the robotic leg SCARLETH," in *Field Robotics* eds P. Bidaud, M. O. Tokhi and G. S. Virk (World Scientific), 507–514. doi: 10.1142/9789814374286_0059

Iida, F., Rummel, J., and Seyfarth, A. (2007). "Bipedal walking and running with compliant legs," in *Proceedings 2007 IEEE International Conference on Robotics and Automation* (Rome: IEEE). doi: 10.1109/robot.2007.364088

Kitano, S., Hirose, S., Horigome, A., and Endo, G. (2016). TITAN-XIII: sprawling-type quadruped robot with ability of fast and energy-efficient walking. *ROBOMECH J.* 3:8. doi: 10.1186/s40648-016-0047-1

Lakatos, D., Ploeger, K., Loeffl, F., Seidel, D., Schmidt, F., Gumpert, T., et al. (2018). Dynamic locomotion gaits of a compliantly actuated quadruped with SLIP-like articulated legs embodied in the mechanical design. *IEEE Robot. Automat. Lett.* 3, 3908–3915. doi: 10.1109/LRA.2018.2857511

Liu, X., Rossi, A., and Poulakakis, I. (2018). A switchable parallel elastic actuator and its application to leg design for running robots. *IEEE/ASME Trans. Mech.* 23, 2681–2692. doi: 10.1109/TMECH.2018.2871670

Lombard, W. P. (1903). The tendon action and leverage of two-joint muscles of the hind leg of the frog, with special reference of the spring movement. *Contrib. Med. Res.* 34, 280–301.

Park, H.-W., Wensing, P. M., and Kim, S. (2017). High-speed bounding with the mit cheetah 2: control design and experiments. *Int. J. Robot. Res.* 36, 167–192. doi: 10.1177/0278364917694244

Pratt, G. A., and Williamson, M. M. (1995). "Series elastic actuators," in *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots* (Pittsburgh, PA: IEEE). doi: 10.1109/iros.1995.525827

Raibert, M. H., Brown, H. B., and Murthy, S. S. (1984). "Machines that walk," in *Robotics and Artificial Intelligence* eds M. Brady, L. Gerhardt, and H. F. Davidson (Berlin; Heidelberg: Springer), 345–364. doi: 10.1007/978-3-642-82153-0_18

Renjewski, D., Spröwitz, A., Peekema, A., Jones, M., and Hurst, J. (2015). Exciting engineered passive dynamics in a bipedal robot. *IEEE Trans. Robot.* 31, 1244–1251. doi: 10.1109/TRO.2015.2473456

Roberts, T. J. (2016). Contribution of elastic tissues to the mechanics and energetics of muscle function during movement. *J. Exp. Biol.* 219, 266–275. doi: 10.1242/jeb.124446

Rummel, J., Blum, Y., Maus, H. M., Rode, C., and Seyfarth, A. (2010). "Stable and robust walking with compliant legs," in *2010 IEEE International Conference on Robotics and Automation* (Anchorage, AK: IEEE). doi: 10.1109/robot.2010.5509500

Sato, R., Miyamoto, I., Sato, K., Ming, A., and Shimojo, M. (2015). "Development of robot legs inspired by bi-articular muscle-tendon complex of cats," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Hamburg: IEEE). doi: 10.1109/iros.2015.7353574

Semini, C., Barasuol, V., Boaventura, T., Frigerio, M., Focchi, M., Caldwell, D. G., et al. (2015). Towards versatile legged robots through active impedance control. *Int. J. Robot. Res.* 34, 1003–1020. doi: 10.1177/0278364915578839

Semini, C., Tsagarakis, N. G., Vanderborght, B., Yang, Y., and Caldwell, D. G. (2008). "HyQ - hydraulically actuated quadruped robot: Hopping leg prototype," in *2008 2nd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics* (IEEE). doi: 10.1109/biorob.2008.4762913

Seok, S., Wang, A., Otten, D., and Kim, S. (2012). "Actuator design for high force proprioceptive control in fast legged locomotion," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Vilamoura: IEEE).

Seyfarth, A., Günther, M., and Blickhan, R. (2001). Stable operation of an elastic three-segment leg. *Biol. Cybernet.* 84, 365–382. doi: 10.1007/PL00007982

Spröwitz, A., Ajallooeian, M., Tuleu, A., and Ijspeert, A. J. (2014). Kinematic primitives for walking and trotting gaits of a quadruped robot with compliant legs. *Front. Comput. Neurosci.* 8:27. doi: 10.3389/fncom.2014.00027

Spröwitz, A., Tuleu, A., Vespignani, M., Ajallooeian, M., Badri, E., and Ijspeert, A. J. (2013). Towards dynamic trot gait locomotion: design, control, and experiments with cheetah-cub, a compliant quadruped robot. *Int. J. Robot. Res.* 32, 932–950. doi: 10.1177/0278364913489205

Spröwitz, A. T., Tuleu, A., Ajallooeian, M., Vespignani, M., Möckel, R., Eckert, P., et al. (2018). Oncilla robot: a versatile open-source quadruped research robot with compliant pantograph legs. *Front. Robot. AI* 5:67. doi: 10.3389/frobt.2018.00067

Tucker, V. A. (1975). The energetic cost of moving about. *Am. Sci.* 63, 413–419.

Verstraten, T., Beckerle, P., Furnémont, R., Mathijssen, G., Vanderborght, B., and Lefeber, D. (2016). Series and parallel elastic actuation: Impact of natural dynamics on power and energy consumption. *Mech. Mach. Theory* 102, 232–246. doi: 10.1016/j.mechmachtheory.2016.04.004

Witte, H., Hackert, R., Lilje, K. E., Schilling, N., Voges, D., Klauer, G., et al. (2001). "Transfer of biological principles into the construction of quadruped walking machines," in *Proceedings of the Second International Workshop on Robot Motion and Control* (Poznań: Poznan University of Technology).

Witte, H., Hoffmann, H., Hackert, R., Schilling, C., Fischer, M. S., and Preuschoft, H. (2004). Biomimetic robotics should be based on functional morphology. *J. Anatomy* 204, 331–342. doi: 10.1111/j.0021-8782.2004.00297.x

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Shaping in Practice: Training Wheels to Learn Fast Hopping Directly in Hardware

Steve Heim, Felix Ruppert, Alborz A. Sarvestani, Alex Spröwitz
Dynamic Locomotion Group
Max Planck Institute for Intelligent Systems, Germany

***Abstract*— Learning instead of designing robot controllers can greatly reduce engineering effort required, while also emphasizing robustness. Despite considerable progress in simulation, applying learning directly in hardware is still challenging, in part due to the necessity to explore potentially unstable parameters. We explore the of concept shaping the reward landscape with *training wheels*; temporary modifications of the physical hardware that facilitate learning. We demonstrate the concept with a robot leg mounted on a boom learning to hop fast. This proof of concept embodies typical challenges such as instability and contact, while being simple enough to empirically map out and visualize the reward landscape. Based on our results we propose three criteria for designing effective training wheels for learning in robotics.**

I. INTRODUCTION

In nature, animals learn to move with a grace and agility that is the envy of robotics engineers. One major challenge is that most algorithms rely on accurate models, which in turn also take a lot of engineering effort. Alternatively, reinforcement learning (RL) is a powerful paradigm that can work both model-based or *model-free*. In addition, reinforcement learning is often able to learn from generic and even highly delayed reward signals: for example a legged robot might receive a reward for reaching a specific target location within a set time limit, and no reward for getting progressively closer. This allows for easy and intuitive assignment of rewards without constraining the behavior for achieving the goal. Despite these attractive features and promising achievements in simulation [1][2], applying RL directly in hardware has proven challenging [3][4][5] with only a handful of successes that actually run model-free [6][7].

One major challenge in hardware comes from the necessity to *explore* the reward landscape. This landscape is usually non-convex, and often only subsets represent behaviors that actually accumulate reward: the rest of the landscape often looks flat, representing different behaviors that all receive the same or even no reward. Sampling from these regions provides no gradient information for the robot to learn from. This is particularly true when the reward is generic and delayed such as in the previous example: a policy that causes the robot to fall over immediately would get the same reward of 0 as a policy that hops in place, even though the second policy arguably solves part the locomotion problem [8].

contact: lastname@is.mpg.de
Dynamic Locomotion Group, Max Planck Institute for Intelligent Systems, Heisenbergstr. 3, 70569 Stuttgart, Germany

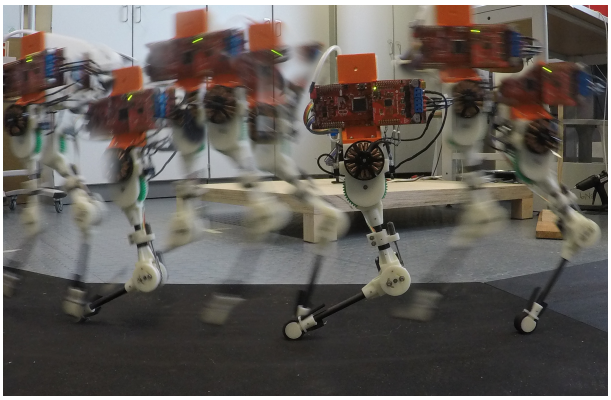

Fig. 1: A multiple-exposure image of our robotic leg mounted on a boom hopping, showing the distinct stance and flight phases.

This problem is even more accentuated in robots that are unstable, since instabilities often quickly lead to direct failure states. These failures generally lead to no reward, and can also damage the hardware. In practice this means exploration is executed cautiously, usually locally. This combination means that large parts of the reward landscape are flat, and there is no salient gradient to lead the learning agent in the correct direction.

This exploration challenge can be solved by choosing a more appropriate policy parameterization. This can however be difficult to find and does not eliminate the flat regions.

Another approach is to shape the reward landscape. A common method of shaping is to encode more information of the task in the reward [9]. The drawback is that it requires more prior knowledge of the task, and goes against the attractiveness of being able to choose rewards based on achieving a task rather than specifying a behavior[1]. It is also possible to shape the landscape by proper mechanical design. For example, walking robots designed after passive-dynamic walkers [10] have good stability properties for a wide range of policy parameters, allowing quick and reliable learning from even poor initializations [6]. The drawback is that designing the system around one specific behavior can be limiting in terms of versatility and design options.

We build on these ideas with the concept of *training wheels*: shaping the landscape with *temporary mechanical modifications of the robot that allow for easier learning*.

[1]Whether the robot crawls, walks or runs should depend on the context of the situation and not on the goal.

To the best of our knowledge, this concept has only been briefly explored in simulation despite initial results showing promise [11]. This paper presents a proof of concept directly in hardware, applied to learning fast hopping of a monoped robot with a rolling foot: an underactuated, unstable system featuring hybrid dynamics. A video of the robot hopping can be found in the supplementary material.

We would like to note that we largely use the terminology of the RL community. In particular, the term *environment* signifies everything that is beyond the learning agent. Take for example an agent whose policy outputs a desired joint position; then the environment includes not only the physical world the robot moves in, but also the robot itself and the PD motor controller used to track the desired joint position. For a more thorough treatment of RL see [5][12].

II. SETUP: MECHANICS, POLICY AND LEARNING SCHEME

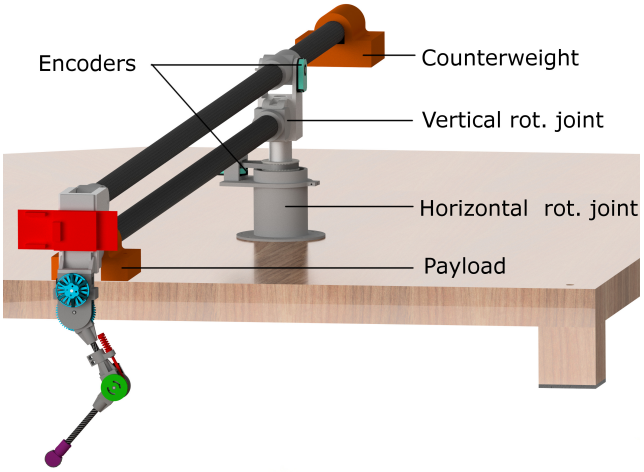

Fig. 2: The entire robot consists of a leg mounted on a boom, with a total of four degrees of freedom. The counterweight balances out the mass of the boom without the leg or payload. The payload represents the mass of the batteries and additional electronics, which are offloaded via a tether.

Our robot platform (Fig. 2) consists of a two-segment leg with a passive-compliant ankle joint and an actuated hip joint, mounted on a boom which constrains the body to motion on a 2D surface. The robot thus has four degrees of freedom (DoFs) and a single actuator. The passive compliance at the ankle joint (Fig. 3) results in favorable natural dynamics [13], though the system is still passively unstable. The learning task is to achieve fast hopping, and the reward for each rollout is the average speed with one additional condition: potentially damaging behavior, such as landing on the ankle instead of the foot, is tagged as a failure and receives no reward. The training wheels for this proof of concept are a simple change of the total mass of the robot body: essentially we allow learning in a reduced gravity environment. We choose these training wheels for two reasons: in addition to somehow making learning easier for the learning agent, they should also be easy to apply

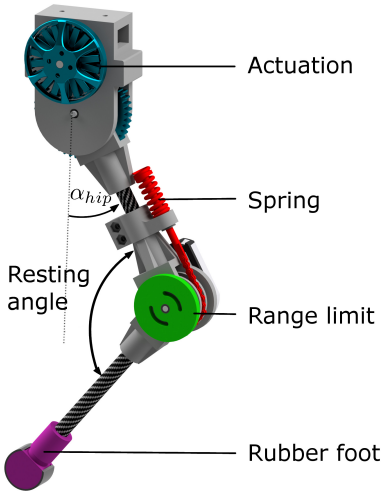

Fig. 3: The two-segment leg has a torque-controlled actuator at the hip, and a passive compliant ankle joint. The spring is mounted to a cam mechanism, and the joint itself is limited in extension range.

such that less engineering effort is required. Based on this we introduce our first criterion in designing training wheels: *how easy are training wheels to apply to a generic set of robots*. For example, the weight of batteries, computation or other payload can easily be offloaded during an initial training phase for most robot designs.

We choose a simple policy with a 2D parameter space: the hip motor tracks an open-loop sinusoidal position trajectory as follows

$$\alpha_{hip} = \theta_0 + \theta_1 \sin(\omega t) \tag{1}$$

where α_{hip} is the angular position of the hip, ω is a hardcoded angular frequency while θ_0 and θ_1, offset and amplitude parameters respectively, form the parameter space of the policy. This simple policy parameterization serves two purposes: first, a low-dimensional deterministic policy is amenable to the simplest of learning schemes, thus eliminating the ambiguity of whether the training wheels or the algorithm implementation are responsible for the change in performance. In the results presented we choose $\omega = 9\ [\frac{rad}{s}]$, based on experience. Higher values could achieve higher performance, but failures were also more violent and prone to damaging hardware. Since we had to also sample failure parameters to map out the landscape, we chose a compromise between safety and performance.

We use stochastic gradient descent based on simple finite-difference methods [14]. More importantly, the low dimensionality allows us to empirically map out and inspect the landscape of the learning problem as a 3D surface as seen in Fig. 4. This allows us to compare the landscapes with and without training wheels in detail, and show the change in learning performance across each landscape.

A. Hardware Details

Each DoF of the boom and leg is measured with a rotary encoder (CUI ATM102-V). The boom arm has a length of $1.5[m]$ from pivot to the leg, and is counterweighted to

completely offset its own mass without the leg. The ankle joint of the leg (Fig. 3) is mechanically limited to 130° in one direction, and has a spring with a stiffness of 6 $[\frac{N}{mm}]$ attached to a cam mechanism with a radius of 15 $[mm]$. This spring is slightly preloaded such that it always returns to the resting angle of 130°. The upper and lower leg segments measure 110 $[mm]$ and 136 $[mm]$ respectively, and the virtual leg length from hip to foot is 223 $[mm]$ at rest. The hip is actuated with a brushless outrunner motor (T-motor MN-4006) with a 1:5 gearbox. The motor control board (Texas Instruments TMS320F28069M with DRV8305EVM booster packs) uses field-oriented control for direct torque control of the motor. A Xenomai real-time linux operating system handles all the high-level control. Electric power and computational power are both off-loaded via tether. A representative mass is directly attached on the boom just behind the leg. With the entire payload, the robot has a body weight of 600 $[g]$. For our two training wheel environments the representative mass is replaced with an intermediate mass or completely removed. This results in a body weight of 505 $[g]$ (0.84 g_0) and 415 $[g]$ (0.69 g_0) respectively.

III. RESULTS

We test three environments: the robot with full payload and two environments with training wheels which reduce the weight to 0.69 g_0 and 0.84 g_0, where g_0 is the total weight of the robot in the original environment. We will refer to these two environments with training wheels as the *beginner environment* and the *intermediate environment* respectively. We map out the entire learning landscape for each environment by sampling and then interpolating the parameter space (Fig. 4). The parameter space is limited to $\theta_0 = [0\ 40]^\circ$ and $\theta_1 = [10\ 45]^\circ$. Parameters outside this range are either unreachable due to mechanical hard-stops, or in the 0-reward region for all environments and cropped out for clarity.

All three landscapes have a mountain-like shape emerging out of a flat surface. While not quite convex, the landscapes each have a prominent peak. This makes the problem amenable to stochastic gradient descent. Also present in all three landscapes is a cliff: a sudden sharp drop from high to 0 reward. This is found in the upper right quadrant of the parameter space and can be recognized in Fig. 4 by the steep contour lines. This cliff represents the border between parameters which exhibit stable high performance and unstable parameters. In practice it is both difficult as well as dangerous to learn from beyond this cliff: policies with high-amplitude tend to crash violently and damage the hardware. It is interesting to note that the orientation of the cliff does change in each environment, though its proximity to the peak does not.

A. Salient Gradient Sets

We are interested in the region that achieves non-zero reward which we will refer to as a *salient gradient set* (SCS), delimited in the figures by the thicker, outermost contour. This is the set of parameters that a learning agent needs to sample from in order to learn. The second criterion for choosing effective training wheels is *how much the training wheels increase the size of the SCS*. Indeed the SCS of the beginner environment covers 46% of the total parameter space, compared to 25% of the intermediate environment and 20% in the original environment. This increase in size is important both for gradient-based and gradient-free methods. With stochastic gradient descent, for example, the gradient is estimated by local sampling. This means the agent must start inside, or at least within local sampling distance of the SCS in order to estimate a gradient. Increasing the size of the SCS directly increases the basin of attraction of the learning system. Other exploration approaches such as eps-greedy can sample from the SCS despite being initialized well away of the SCS. In this case, increasing the size of the SCS increases the probability of sampling from it, thus improving rate of convergence.

B. Funneling Sets

In Randløv's simulated work [11], the training-wheels environment converges gradually to the original environment. In practice, it is often difficult to implement a gradual mechanical change: in many cases it is desirable to have training wheels that are either on or off, or at least that require only a few stages. This brings up an important requirement for training wheels: *successive environments must funnel into each other*. In Fig. 5, the peak of the beginner environment, located at $[22.5\ 32.6]^\circ$, lies only barely within the SCS of the original environment. In general, there are no guarantees that the peak of the training-wheels will be contained in the SCS of the original problem. If it isn't, or in our case if the policy has not fully converged, there is a good chance that when removing the training wheels the policy is still too far from the next SCS to effectively sample from it. This issue is solved by having an intermediate environment, whose SCS contains a large area around the peak of the beginner environment. In other words, each training-wheel environment should easily *funnel into* the SCS of the next environment to be effective. This is particularly important when using a local exploration strategy, and is conceptually the same as designing controller funnels [15][16].

There is a second consideration that should be kept in mind: while the peak of an earlier environment *must*[2] be contained in the SCS of the successive environment, the reverse is not true. This means that switching *back* to an earlier training environment must be done only cautiously, especially in hardware.

C. LEARNING ACROSS LANDSCAPES

As a proof of concept we use offline stochastic gradient descent with finite-differences, using parameter perturbations ranging between 0.5° and 2°. Larger perturbations tend to be more robust to noise in estimating the gradient, and where the gradient is very shallow such as at the edges of the SCS. On the other hand, they become unstable closer to the peak

[2]This condition is necessary when exploration is strictly local, and can be relaxed otherwise.

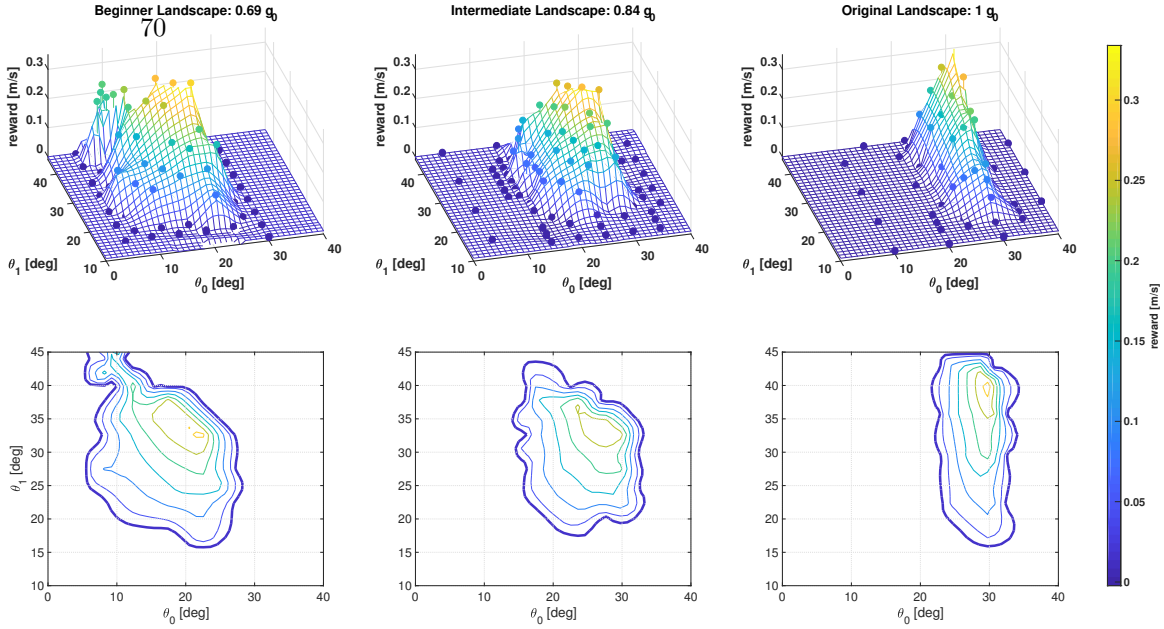

Fig. 4: The landscapes of the beginner, intermediate and original environments are visualized here. The upper row shows the sampled points (circles) and the resulting interpolated mesh, slightly offset for visibility. The more gradual climb in the lower gravity environments is visible. The contour maps in the second column more clearly show the change in shape of the 'reward mountain', the shape of of the cliff and most importantly, the size of the basin of attraction for the learning system. The outer contour showing the set of parameters which can provide a gradient is outlined with a thicker line. If the learning agent only samples outside this set, it will not be able to accurately estimate a gradient.

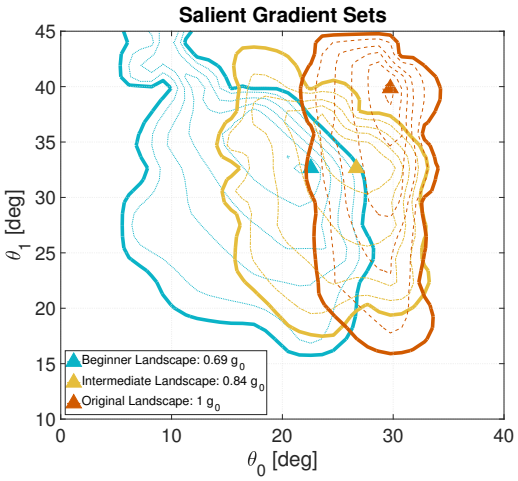

Fig. 5: The salient gradient set (SCS) for each environment is mapped out with contour lines and the peak of each set marked by a triangle. The location of the peak of one training environment with respect to the SCS of the successive environment is very important. To be effective, the training wheels must guide the current policy towards parameters that will sample from the salient gradient set of the next landscape with higher probability.

and especially when close to the cliff. We also choose a constant, relatively large learning rate of 2.5. Again, larger steps have the risk of overshooting and stepping over the cliff, but otherwise perform well. In both cases a cleverer, variable choice of these parameters would help the learning process, but is not relevant for showing the training-wheels concept and thus they are kept constant.

Several typical learning sessions (Fig. 6a), with the most successful reaching a velocity of 0.35 $[\frac{m}{s}]$. We also purposefully initialize several trials outside the SCS of the original environment, and as expected observe meandering paths (Fig. 6b). In Fig. 6b, agents initialized with parameters outside of their salient gradient sets are depicted. As expected, without sufficient gradient information the agents will simply take steps in random direction. While this random exploration has a non zero probability of entering the salient gradient set and therefore converging, it can take a large amount of iterations, especially when starting at some distance from the set and with smaller learning-step sizes. Especially when learning directly in robot hardware, reducing the number of trials necessary has a high priority since running trials can be time and resource costly.

IV. CONCLUSIONS AND OUTLOOK

We build on the concept of training wheels, temporary mechanical modifications of the system, to shape the learning landscape[11]. We apply it to learning open-loop legged locomotion in a constrained test stand, as a simple, low-

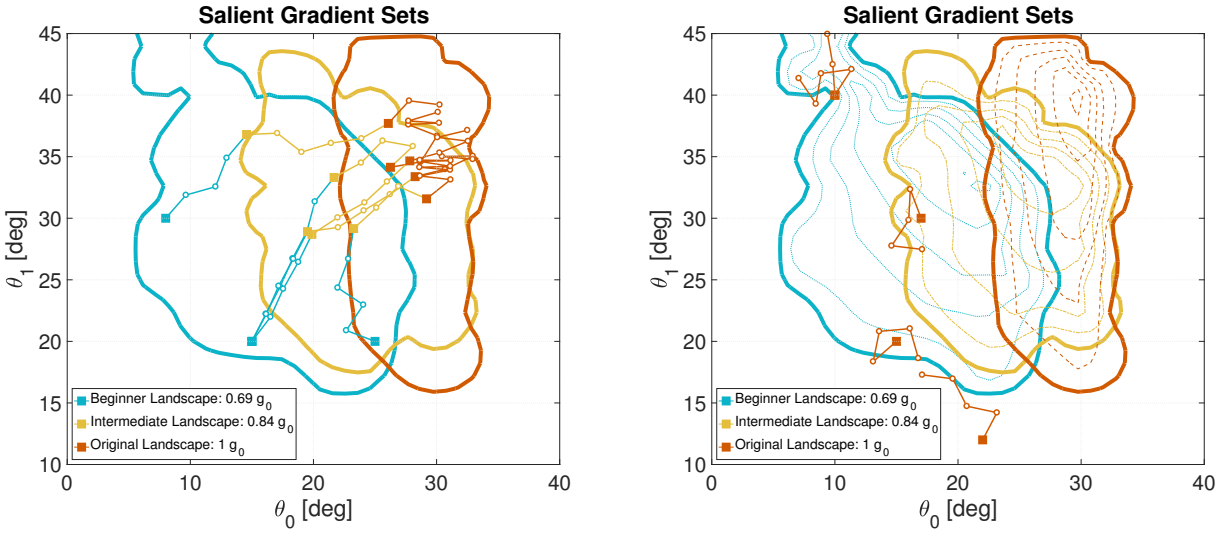

(a) Typical learning paths

(b) Learning attempted with the original landscape, marked in red.

Fig. 6: Several typical learning paths are shown here: above are successful learning trials progressing from the beginner to intermediate and finally original environment. Contour lines are not shown to clearly show the learning paths. One of these learning trials typically took around 10 minutes. Below are trials initialized directly with the original environment but outside its salient gradient set. These take learning steps in random directions without improving.

dimensional problem that is unstable, underactuated and features impacts. We've identified three key criteria to designing effective training wheels in practice.

1) Ease of application to a generic set of robots
2) Increase in probability of sampling from the salient gradient set
3) Ease of funneling from a training environment to the successive environment

Since reducing the engineering effort a main attraction for applying learning to robotics, it is important that training wheels are easy to implement and apply. As an example, it is intuitively predictable that adding damping to the joints of a robot would help stabilize the system and greatly help learning, especially for force control [17]. Implementing mechanical damping on small joints is however much more difficult than simply temporarily offloading some of the payload, and would require a custom design for each new robot. This partially defeats the purpose of reducing the engineering effort. While we are planning on exploring solutions to this in future work, there is a lot of merit in solutions as effective yet simple as reducing the payload.

The second criterion is the main qualifier for the effectiveness of the training wheels in shaping the learning landscape. To be more precise, increasing the probability of sampling from the salient gradient set is what makes a training wheel environment easier to learn in. The actual size of the set in relation to the sample space is a good proxy; it is more generalizable to arbitrary exploration strategies, and makes is more intuitive to predict when designing training wheels. Since in practice it will rarely be possible to quantify the change in landscape mathematically or empirically, as we've done here, it is important to be able to intuitively predict the qualitative effect of the training wheels.

The final criterion is particularly relevant when using local exploration strategies. As this strategy is common in robotics, we feel it is an important criterion to include. Training wheels that can be continuously tuned out, until the dynamics converge back to the original environment, would be guaranteed to satisfy this criteria [11]. However the implementation of such training wheels generally goes against the first criteria, and a trade-off will have to be made. For future work, it will be important to develop a systematic way of making this trade-off.

In this proof of concept, the timing for switching between environments was chosen heuristically. Having the actual landscape maps as reference, not to mention a lot of experience with this system, we were very confident that the funnel overlap between environments was very big and we did not have to completely converge on one environment before switching to the next. For future work, it will be very interesting to find a more general rule for switching environments. Since the number of trials needed to converge is particularly important when learning in hardware, finding an optimal policy to learn with the fewest numbers of iterations would be particularly useful.

Although we have presented these landscape shaping results

in the context of reinforcement learning, the challenge of traversing a landscape in $\mathbb{R}^2$ parameter-space is inherent to optimization problems as a whole. In particular, the concepts we develop should be useful for applying derivative-free optimization in hardware [18].

ACKNOWLEDGMENTS

We would like to thank Özge Drama, Felix Grimminger and Julian Viereck for fruitful discussions, advice and help along the way. We also thank Felix Widmaier who wrote the code to run and interface with the motor control boards. We thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting the academic development of Felix Ruppert. This work was partially supported by a MPI Grassroots grant provided to Ludovic Righetti, Felix Grimminger and Alexander Spröwitz in 2017.

REFERENCES

[1] X. B. Peng, G. Berseth, and M. Van de Panne, "Terrain-adaptive locomotion skills using deep reinforcement learning," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, p. 81, 2016.

[2] T. P. Lillicrap, J. J. Hunt, A. Pritzel, *et al.*, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.

[3] P. Abbeel, A. Coates, M. Quigley, *et al.*, "An application of reinforcement learning to aerobatic helicopter flight," in *Advances in neural information processing systems*, 2007, pp. 1–8.

[4] J. Peters and S. Schaal, "Reinforcement learning of motor skills with policy gradients," *Neural networks*, vol. 21, no. 4, pp. 682–697, 2008.

[5] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.

[6] R. Tedrake, T. W. Zhang, and H. S. Seung, "Learning to walk in 20 minutes," in *Proceedings of the Fourteenth Yale Workshop on Adaptive and Learning Systems*, Yale University New Haven (CT), vol. 95585, 2005, pp. 1939–1412.

[7] N. Kohl and P. Stone, "Policy gradient reinforcement learning for fast quadrupedal locomotion," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, IEEE, vol. 3, 2004, pp. 2619–2624.

[8] M. H. Raibert, *Legged robots that balance*. MIT press, 1986.

[9] V. Gullapalli and A. G. Barto, "Shaping as a method for accelerating reinforcement learning," in *Intelligent Control, 1992., Proceedings of the 1992 IEEE International Symposium on*, IEEE, 1992, pp. 554–559.

[10] T. McGeer, "Passive dynamic walking," *The International Journal of Robotic Research*, vol. 9, no. 2, pp. 62–82, 1990.

[11] J. Randløv, "Shaping in reinforcement learning by changing the physics of the problem.," in *ICML*, 2000, pp. 767–774.

[12] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*, 1. MIT press Cambridge, 1998, vol. 1.

[13] J. Rummel and A. Seyfarth, "Stable running with segmented legs," *The International Journal of Robotics Research*, vol. 27, no. 8, pp. 919–934, 2008.

[14] J. Peters and S. Schaal, "Policy gradient methods for robotics," in *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, IEEE, 2006, pp. 2219–2225.

[15] Q. Cao, A. T. Van Rijn, and I. Poulakakis, "On the control of gait transitions in quadrupedal running," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, IEEE, 2015, pp. 5136–5141.

[16] R. Tedrake, I. R. Manchester, M. Tobenkin, *et al.*, "Lqr-trees: Feedback motion planning via sums-of-squares verification," *The International Journal of Robotics Research*, vol. 29, no. 8, pp. 1038–1052, 2010.

[17] J. E. Colgate and J. M. Brown, "Factors affecting the z-width of a haptic display," in *Robotics and Automation, 1994. Proceedings., 1994 IEEE International Conference on*, IEEE, 1994, pp. 3205–3210.

[18] A. Spröwitz, R. Moeckel, J. Maye, *et al.*, "Learning to move in modular robots using central pattern generators and online optimization," *The International Journal of Robotics Research*, vol. 27, no. 3-4, pp. 423–443, 2008.

FootTile: a Rugged Foot Sensor for Force and Center of Pressure Sensing in Soft Terrain

Felix Ruppert and Alexander Badri-Spröwitz

***Abstract*— In this paper, we present FootTile, a foot sensor for reaction force and center of pressure sensing in challenging terrain. We compare our sensor design to standard biomechanical devices, force plates and pressure plates. We show that FootTile can accurately estimate force and pressure distribution during legged locomotion. FootTile weighs 0.9 g, has a sampling rate of 330 Hz, a footprint of 10×10 mm and can easily be adapted in sensor range to the required load case. In three experiments, we validate: first, the performance of the individual sensor, second an array of FootTiles for center of pressure sensing and third the ground reaction force estimation during locomotion in granular substrate. We then go on to show the accurate sensing capabilities of the waterproof sensor in liquid mud, as a showcase for real world rough terrain use.**

I. INTRODUCTION

In walking robotics and in biomechanics, sensors are a fundamental tool to understand the mechanics and control of a system. Biomechanic data is especially important in outdoor environments to collect data during natural locomotion. Standard tools in biomechanics, are ground reaction force plates and pressure plates to measure the force and pressure distribution on a foot. Force plates and pressure plates are easy to use but are heavy, delicate, immobile and expensive. It is also not possible to measure pressure distribution on soft and granular substrate. The biggest drawback of these systems, however, is that they are not wearable. Both systems can only capture a small number of strides because of their size. Therefore it becomes hard to collect average data or data where the conditions change between steps.

To have a pressure or force reading for a long duration at every step of a system, wearable devices are required. The most direct sensors are strain gauge based force/torque sensors [1]. Attached to the foot, these sensors can read forces and torques with high range, accuracy, and frequency. Stiff sensors, however, tend to be heavy and increase the rotational inertia when placed at the end of a leg. They are also expensive and very sensitive to impact forces. Soft sensors solve the problems of weight and impact sensitivity by implementing soft, deformable elements. Under load, the soft element deforms, and various sensing principles can measure the deformation. The Optoforce sensor measures the changes in light reflection on the inside of a soft dome due to deformation. Alternatively, the deformation of a soft sensor body can be detected by Hall effect sensors [2], resistive [3–6] or capacitive [7] sensors. While these sensors are easy to use and cheap to implement, they are not sufficiently robust against the high impact forces during running and can not easily be scaled down and miniaturized. Tenzer and colleagues [9] developed a tactile sensor called 'Takktile', based on a barometric pressure sensor. The micro-electro-mechanical (MEMS) sensor's cavity gets filled with rubber that transmits external forces to the pressure sensor's membrane. The miniature sensor can measure forces up to 5 N and can be overloaded with 400 % of its nominal force. To extend the sensor range, Chuah and colleagues [10] implemented a higher range pressure sensor in the same

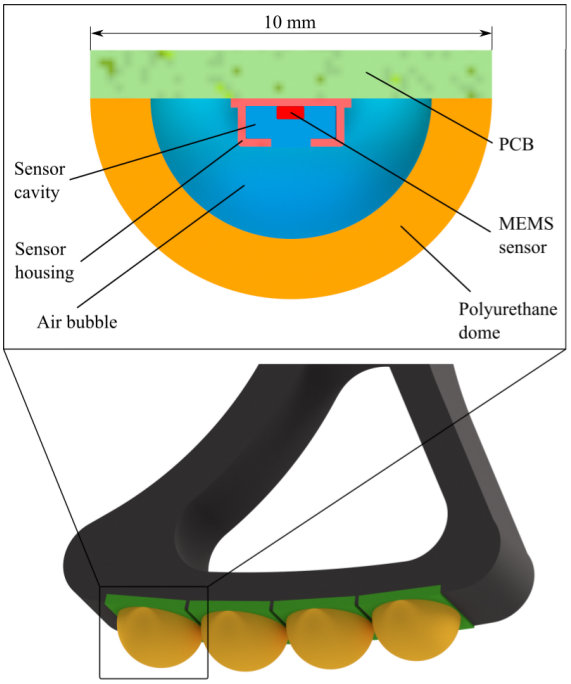

Fig. 1: a) FootTile sensor design, inspired from [8], with significant changes to bubble shape, size, and simplified dome morphology. The size of the air bubble, and the shape of the dome influence the sensor's characteristics. The blue areas show the air bubble volume. External forces will compress the air, and increase the air pressure. The FootTile in this image weighs 0.8 g, has a size of 10 x 10 x 5.5 mm, a sample rate of 330 Hz and a resolution of 10 bit in 30 N. b) FootTile array on a robot foot to estimate the center of pressure and ground reaction forces.

Authors are with the Dynamic Locomotion Group, Max Planck Institute for Intelligent Systems, Stuttgart, Germany

Correspondance: ruppert@is.mpg.de

FR contributed to concept, design, experiments, data analysis and writing. ABS contributed to concept and writing.

We thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting the academic development of Felix Ruppert. This work was made possible thanks to a Max Planck Group Leader grant awarded to ABS by the Max Planck Society.

design to circumvent the range limits. Lèon and colleagues [8] enhanced the sensor range by including an engineered bubble into the rubber dome. This bubble acts as a pressure transducer. Based on the dome deformation, the bubble volume reduces, and the pressure in the bubble rises. The range can be designed based on the dome design and can be adjusted for any load case. The sensor was used in the paper to measure the forces on a rotorcraft robotic landing gear.
In this project, we adapt the sensor design from [8], for use in legged locomotion, with our requirements for force range, sensitivity, sensor size and sampling frequency. Specifically, we implement a sensor array and showcase measurements of forces as well as the center of pressure on solid and soft terrain. We present a foot sensor array that is *lightweight*, *rugged*, *wearable*, *waterproof* and allows *force and pressure sensing* in flat and rough terrain and granular media as well as mud. As shown in previous research, the sensor can be used on robotic systems for control purposes. We illustrate the advantages of this miniature sensor for use in biomechanical legged locomotion experiments for animal or robotic subjects. We compare the sensor performance to force plates and pressure plates, the two standard devices for biomechanical data acquisition. We show how the sensor can estimate ground reaction forces and the center of pressure. We present experimental data on how the sensor can be used on solid ground, as well as granular medium and real world rough terrain, like mud.

II. METHODS

A. Sensor design

The sensor in Figure 1 consists of a barometric pressure sensor (MPL115A2, *NXP*). To date, it is the miniature barometric pressure sensor with the highest sample rate on the market. In this paper, we implement simplified polyurethane half-spheres with a spherical bubble inside to reduce the mold complexity. By using different PU materials, three design parameters can be used to adapt the sensor range to the application. The spheres are produced in a two-part 3D printed mold (Onyx, *Markforged*) and glued airtight to the sensor PCB with instant glue. To make the FootTile as small as possible, we chose a bubble diameter of 6 mm to fit around the sensor diagonal of 5.4 mm. The total sensor design has production costs of less than 10 € per sensor.
The pressure sensor is internally temperature-compensated with an individual calibration. Using temperature-compensated pressure for calculations, we can assume an isotherm process ($T = \text{const}$), the general gas law results in:

$$p_1 \cdot V_1 = p_2 \cdot V_2\big|_{T=\text{compensated}}$$

which leads to

$$p_2 = p_1 \cdot \frac{V_1}{V_2} \quad (1)$$

where p_i is the pressure in the bubble, V_i is the bubble volume, and T is the gas temperature in the bubble, before and after deformation, respectively. The MEMS sensor has a range of 50 kPa to 115 kPa. The initial pressure in the bubble during manufacturing is one atmospheric pressure, around p_1 =1 bar≡ 100 kPa. The sensor saturates at $1.15 \cdot p_1$ =115 kPa. To maximize the range of the MEMS sensor, the bubble should compress by 13% of its initial volume at the maximum external load. By changing the bubble volume and the dome geometry, the sensor range can be designed for many load cases.
To show the range dependency on the dome and bubble volume as well as the material, we test the parameter space shown (Table I).

TABLE I: Experimental parameters for sensor design

Parameter	Values
Dome diameter d_{dome} [mm]	10, 11, 12
Bubble radius r_{bubble} [mm]	3
Material (E_{Young})	Vytaflex40 (E=0.69 MPa), Vytaflex60 (E=2.17 MPa)

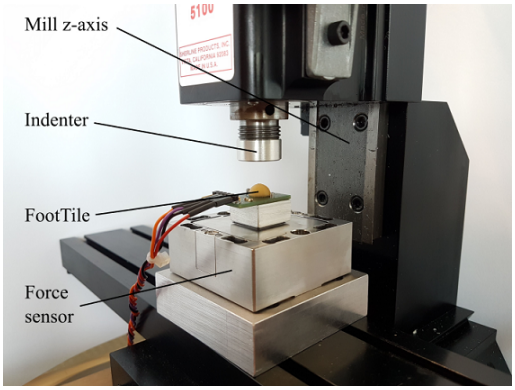

Fig. 2: Sensor validation setup. FootTile mounted to a force sensor for calibration. The guide of the manual mill guarantees purely vertical deflection through the indenter, fixed to the milling head. The FootTile is compressed until its barometric pressure sensor saturates.

B. Sensor validation

To calibrate the sensor reading to the applied force, we use a 3-axis force sensor (K3D60, *ME-Systeme*) as ground truth. The FootTile is fixed to the force sensor mounted on the z-axis of a manual mill (Figure 2). By moving the mill in the vertical direction, the indenter on the mill head deforms the FootTile, and the force sensor data can be compared to the FootTile data. The FootTile is connected via I^2C to a Raspberry Pi 3B+. The force sensor is connected to an amplifier (BA9236, *Burster*) and an analog-digital converter (MCP3208, *Microchip*). The different sensor configurations from Table I are indented until the sensor saturates at 115 kPa to determine the maximum sensor range. The sensor has a conversion time for one data sample of 1.6 ms. This results in a maximum theoretical reading frequency of 625 Hz. Due to speed restrictions in the Raspberry I^2C driver, data for both sensors is sampled at 330 Hz. The indentation experiments are repeated four times, and the data is averaged and shown with 95% confidence intervals.

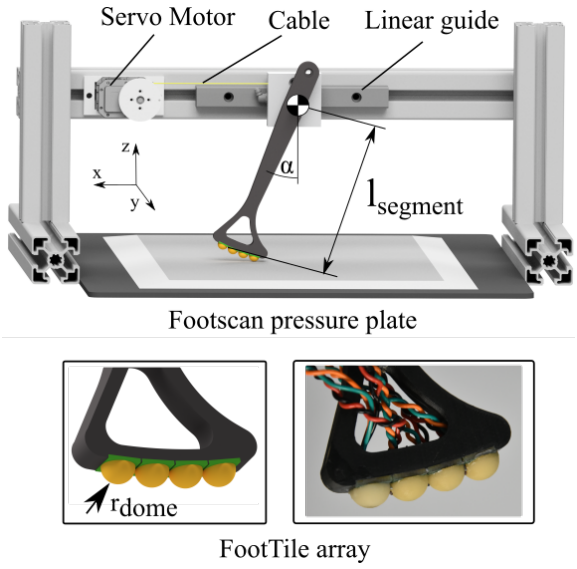

Fig. 3: Pressure plate comparison. Robotic leg equipped with an array of four FootTile sensors to measure the pressure distribution during locomotion. The foot rolls on a footscan pressure plate as ground truth. The linear guide is pulled by a servo motor to ensure reproducible experiments.

C. Pressure plate comparison

In order to measure the pressure distribution and center of pressure (COP) along a footpad, we implement an array of FootTiles. The center of pressure describes the point where the resultant force vector of a pressure field acts on a body [11]. To simplify the experiment, we restrict the foot to only roll on the ground in the sagittal plane. Therefore we only use a one-dimensional array of sensors. The foot of our robot [12] is redesigned to have a constant radius of 150 mm measured from the ankle joint over the footpad. The footpad arc spans the angle the robot leg sweeps during actuated hopping. We place four sensors along the arc of the foot segment. As ground-truth, the foot rolls onto a pressure plate (Advanced footscan pressure plate, *rsscan*). The FootTiles are connected to an I^2C multiplexer (TCA9548A, *Texas Instruments*), to use several sensors on one I^2C bus and reduce the delays caused by the sensor conversion time. To measure the foot position and segment angle, we place two markers along the leg axis. The experiments are recorded with a camera at 50 fps. From the supplementary video, we can extract the marker position to calculate angle and position data of the foot segment. The pressure plate data is recorded with 330 Hz as well as the FootTile array data. The leg's joint is connected to a linear rail with two ball bearings (Figure 3). The linear rail guarantees that the arc on the footpad of the leg segment experiences pure rolling without height changes. To make the experiment reproducible, we implement a servo motor (MX-28AR *Dynamixel*) that pulls the linear guide sled with a constant velocity by a Teflon cable.

D. Granular and muddy substrate

Here, we compare the ground reaction force (GRF) estimation capabilities of the FootTile array. We simulate locomotion in granular medium by using poppy seeds as substrate [13]. The leg with a total weight of 909 g is dropped into a box filled with poppy seeds. As ground-truth we place a force plate (9260AA, *Kistler*) under the box (Figure 4). We record the experiment with a high-speed camera from the sagittal plane at 1000 Hz. To simulate locomotion, a servo motor (MX-64, *Dynamixel*) is connected to the hip of the robot to achieve realistic leg swinging behavior. The center of mass (COM) motion of the robot is constrained to the sagittal plane by a boom structure. The leg is dropped from a foot height of 10 cm above the substrate. Data from the FootTile array is recorded at a maximum achievable frequency of 330 Hz. We further test the sensor under realistic conditions in rough terrain by replacing the poppy seeds with mud from a nearby forest, to showcase the rugged and waterproof design. The leg is dropped into mud while the same data as before is recorded. After use, the sensors are cleaned with water while remaining fully functional. To make the sensor waterproof, only the back of each PCB had to be covered with a layer of waterproof protective urethane resin (Urethane *Cramolin*). The resin protects the PCB from moisture, chemicals, and abrasion. The domes seal the sensors water- and airtight by design.

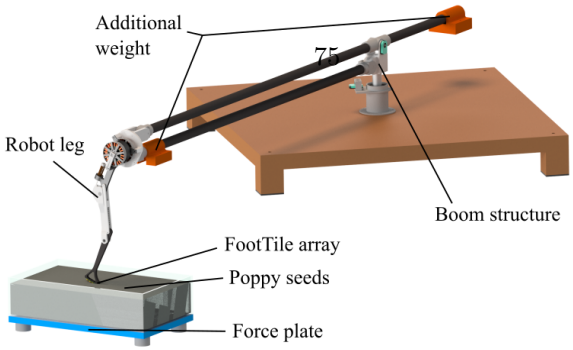

Fig. 4: Granular medium setup. Render of the robot leg with FootTile array. We use poppy seeds as a substitute for a granular medium. The robot is constrained to the sagittal plane by a boom structure. As ground truth for the overall force applied by the leg the poppy seed box is placed on a force plate.

III. RESULTS

A. Sensor validation

We compare the data from the FootTile to the force data recorded from the force sensor (Figure 5). The FootTile data is normalized to the maximum value when the barometric sensor saturates. As expected, the FootTiles saturate at different forces depending on the material and dome geometry. The sensors molded with Vytaflex 40 saturate at lower forces than the senors with Vytaflex 60 domes. Sensors saturate at higher forces for bigger dome diameters. All data displayed is averaged over four experiments and displayed with 95% confidence intervals. As shown in previous studies [9], the sensor measurements are repeatable with a standard deviation of 0.03 for the normalized sensor output. Because of the deforming bubble geometry, the sensor output is nonlinear with respect to the applied force. The sensor fitting best for the

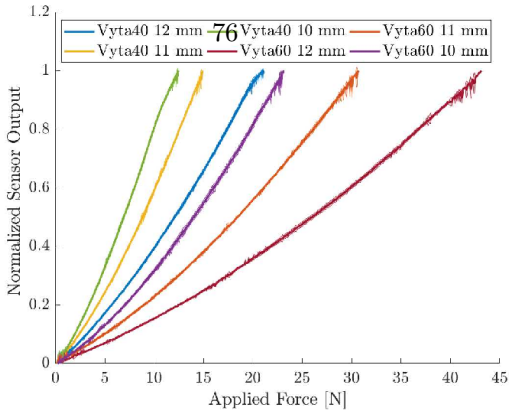

Fig. 5: Normalized FootTile pressure data over reference input force from an external force sensor for sensor validation. Increasing dome material stiffness and dome diameter increase the sensor range. With dome material and diameter, our FootTile can be adapted for any load case. Displayed is averaged data over four experiments with 95 % confidence intervals.

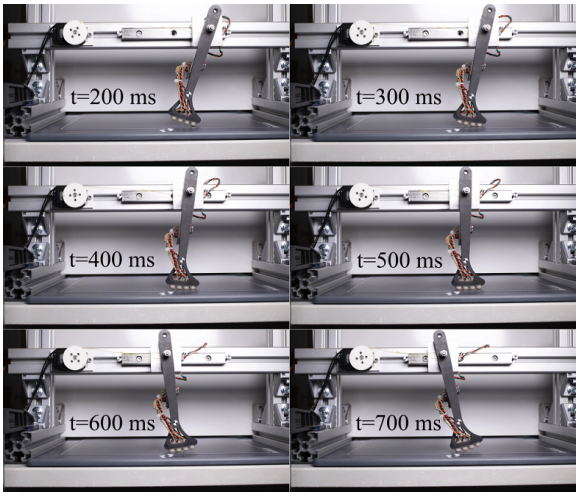

Fig. 6: Video snapshots of the pressure plate comparison. The array is rolling on the pressure plate for different moments in time.

required sensor range of 25-35 N from previous experiments with the robot is the sensor with the 11 mm Vytaflex 60 dome. The sensor with PCB and dome has a size of 10 x 10 x 5.5 mm and weighs 0.86 g. All subsequent experiments use this sensor.

Since the measurements have such a small standard deviation, we can approximate the sensor output with a third-order polynomial to correct sensor data in coming experiments. The pressure-force equation for the selected sensor (Figure 5 orange) can be approximated with

$$F(p) = 0.13 \cdot p^3 + 0.02354 \cdot p^2 + 0.5702 \cdot p + 0.1309 \quad (2)$$

with $R_{square} = 99.99\%$, where F is the applied force, and p is the FootTile output pressure.

B. Pressure plate comparison

To validate how well the FootTiles can estimate the pressure distribution, we plot the pressure plate data versus the data of a 4×1 array of FootTiles (Figure 7). The pressure plate data shows an interpolation of all the active sensor cells. The FootTile data shows the raw data from the individual FootTiles. We calculate the active area of the FootTile array by using the video data (Figure 6) to calculate the leg segment angle and find the contact points for all sensors with

$$\begin{aligned} y_{contact} &= l_{segment} \cdot \sin(\alpha) - r_{dome} \cdot \cos(\alpha) \\ z_{contact} &= -l_{segment} \cdot \cos(\alpha) - r_{dome} \cdot \sin(\alpha) \end{aligned} \quad (3)$$

where $y_{contact}$ and $z_{contact}$ are the y and z contact point coordinates with respect to the leg segment joint, $l_{segment}$ is the segment length, α is the segment angle to the vertical, and r_{dome} is the dome radius as shown in Figure 3. The origin is in the rotary joint on the linear guide. To calculate the COP position in the sensor grid, we use a weight function in x and y direction,

$$\begin{aligned} x_{COP} &= \frac{1}{p_{total}} \cdot \sum_{i=1}^{m} (p(i) \cdot i) \\ y_{COP} &= \frac{1}{p_{total}} \cdot \sum_{j=1}^{n} (p(i) \cdot j) \end{aligned} \quad (4)$$

where x_{COP} and y_{COP} are the Cartesian COP coordinates (Figure 3) in the $m \times n$ grid, p_{total} is the normalized sum of all grid pressure, and index i and j of the grid are used as the linear weight functions in x and y directions. We construct a 3×50 matrix with the raw FootTile data in the second column to be able to compare the pressure distributions visually (Figure 7). Since the FootTile array in this paper is one-dimensional, we only compare the x-component of the COP position (Figure 8). The COP estimation from the FootTile array fits the ground truth from the pressure plate. The error is less than $4\,\mathrm{mm}$ throughout the whole experiment. Only at the end of the experiment, the error gets bigger due to the FootTile lifting from the pressure plate.

C. Granular and muddy substrate

We compare the reading of the FootTile array with the force plate recorded ground reaction forces during actuated stepping into the poppy seed box (Figure 9). The sensor values are recalculated into forces using Equation 2. The sum of all forces is plotted together with the vertical ground reaction force of the force plate (Figure 10).

The leg hops three times (visible in supplementary video). During the first hop (Figure 11), a single FootTile measures most of the force. During the second hop, the force is distributed among all sensors. During the third hop, the leg comes to rest in the poppy seed box. During all three hops, the GRF estimation is reasonable and follows the ground truth data of the force plate. Around the force peaks, there is a small deviation between the FootTile and the ground truth. We believe this to be caused by the hyper-elastic material

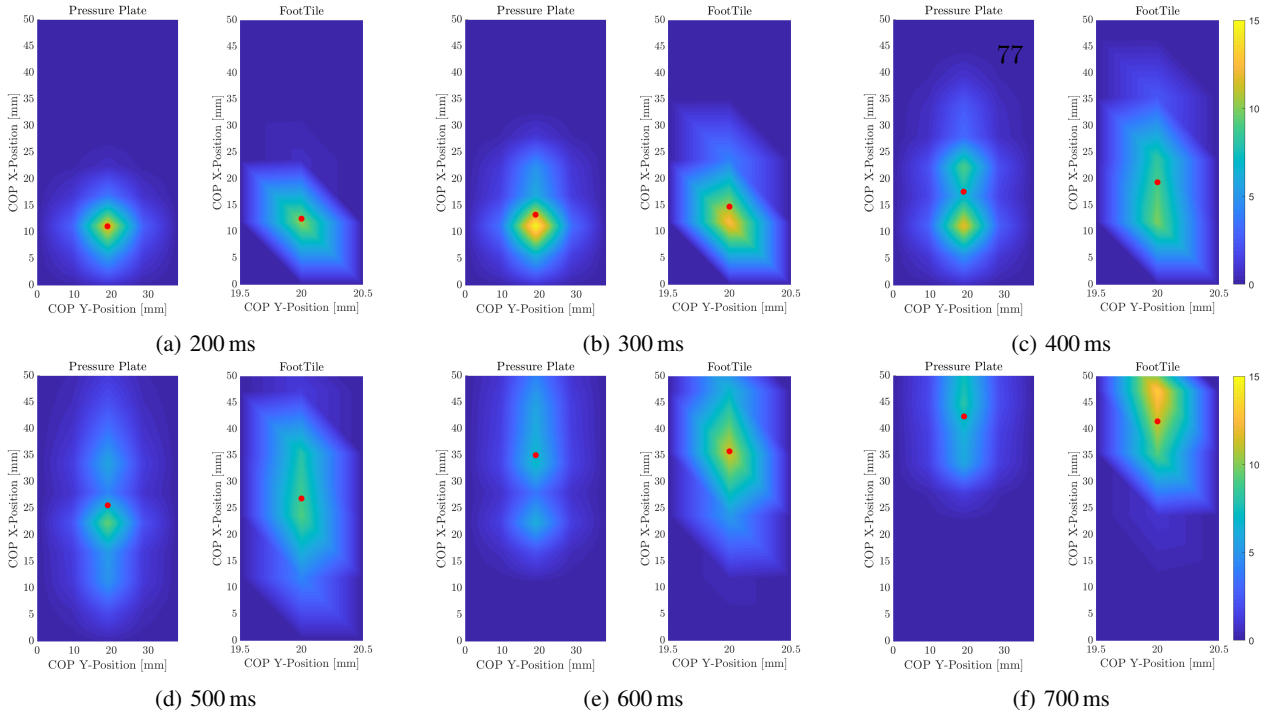

(a) 200 ms (b) 300 ms (c) 400 ms

(d) 500 ms (e) 600 ms (f) 700 ms

Fig. 7: Pressure plate comparison spatial and temporal pressure value readout. Pressure plate data, as ground-truth, shows interpolated data of the active 4x50 sensors. FootTile data shows one-dimensional raw sensor data integrated into a 3x50 grid for comparable visuals. The COP position for each dataset is indicated by the red dot. The error is less than the size of a single FootTile. The small error between ground-truth and FootTile COP is quantified in Figure 8.

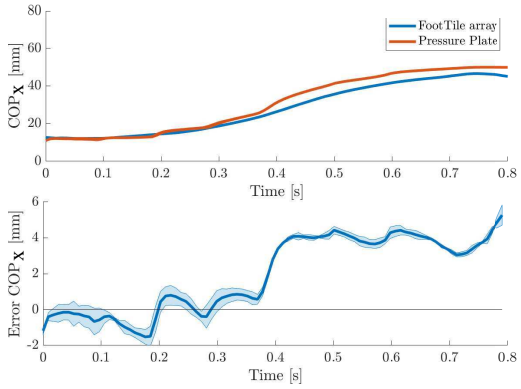

Fig. 8: (A) Pressure plate comparison COP x-coordinate for both the FootTile array and the pressure plate as ground truth. The displayed data is the average of four experiments.
(B) Pressure plate comparison difference in COP x-coordinate. The displayed data is the average of four experiments with 95% confidence interval. The error during stance is smaller than 4 mm. At the end of the experiment, the sensor array loses contact to the pressure plate and the COP estimation becomes less accurate.

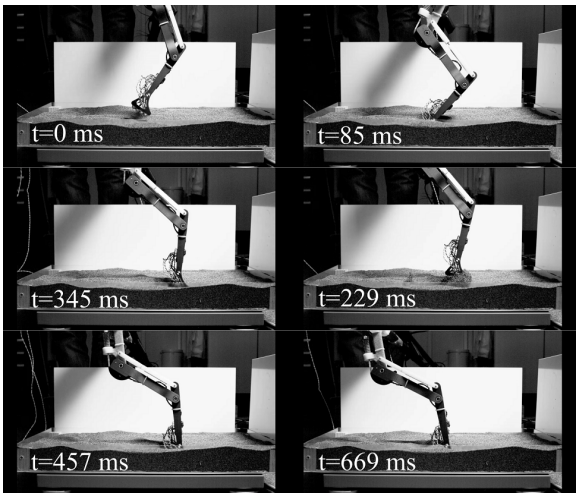

Fig. 9: High-speed snapshots of the FootTile leg hopping in poppy seeds as a granular medium. The leg hops three times with a different number of FootTiles engaging with the ground. Timestamps are displayed in the figure.

properties of the polyurethane dome, as well as the torque influences that we neglect here.
The FootTile array is able to estimate the vertical ground reaction force with a mean error of less than 1 N. Discrepancies between the ground truth and the FootTile estimation stem from the one-dimensional sensor array. Since the sensor is modular, it is possible to use multiple lines of sensors shifted by half the sensor length to increase the sensing resolution in one direction. This way, more sensors are engaged at the same time and potential error during the transition from one sensor to the next could be prevented.

To showcase the ruggedness, the FootTile array is dropped into mud, where the sensor array is fully submerged in liquid mud (Figure 12). The array stays fully functional during the experiment (Figure 13). After use, the FootTile array can be washed in water without compromising the functionality

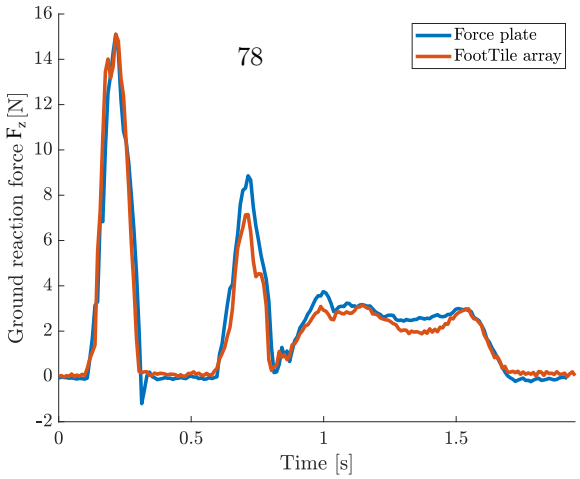

Fig. 10: Comparison of the raw force plate vertical ground reaction force as ground truth (blue) and the estimated FootTile forces (orange) from raw pressure data during hopping in poppy seeds. Mean error is smaller than 1 N.

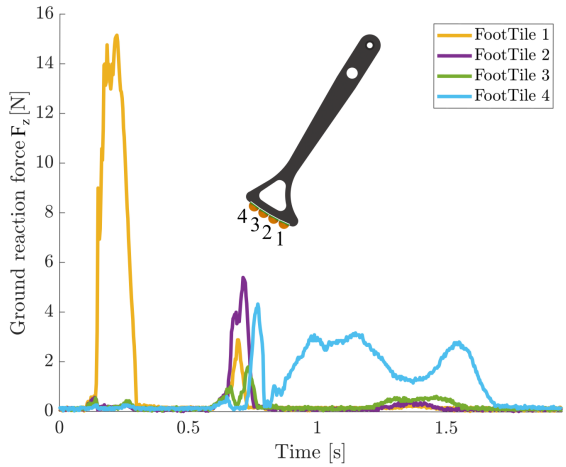

Fig. 11: Individual FootTile contribution from Figure 10 to the ground reaction force estimation hopping in poppy seeds. During the first hop, mostly the first sensor is excited. During the second hop, all sensors are excited, during the third hop, mostly the fourth sensor is excited.

of the sensor (supplementary video). Again, the summation of FootTile GRF estimation follows the ground-truth from the force plate. Small deviations could stem from the highly anisotropic material that includes small stones and roots. This shows the capability of the FootTile to be used outside of a laboratory environment in soft and moist real world terrains.

IV. DISCUSSION

In this paper, we present FootTile, a force and pressure sensor that is lightweight, small, portable, rugged, modular, and low cost. The sensor can easily be adapted to any required load case. We adapt the previously presented sensor design in size, measuring range, sampling frequency and reduced complexity. We present experimental data and show that FootTile can be used as an alternative to standard biomechanical tools like force plates and pressure plates.

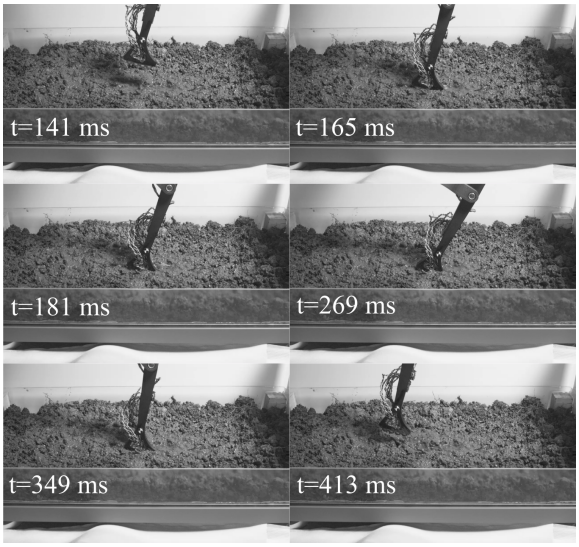

Fig. 12: High-speed screenshots of the FootTile array hopping in mud as a real world substrate. The waterproof FootTile array is fully submerged in water and mud and remains functional.

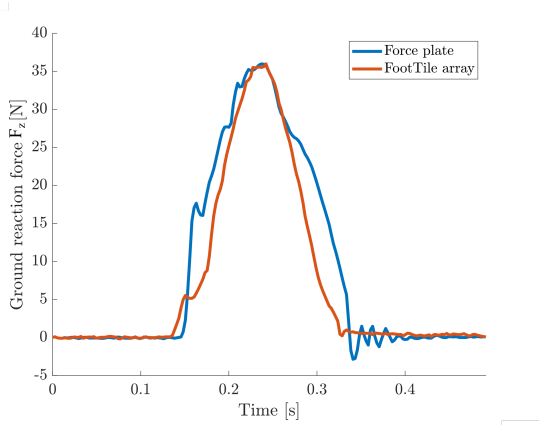

Fig. 13: Ground reaction force estimation of the FootTile array hopping in mud compared to a standard force plate.

FootTile can accurately estimate ground reaction forces as well as the center of pressure position in one device.
If required, the reading frequency could be improved with a dedicated microcontroller (in the boundaries of material responsiveness). With the current development of smaller MEMS pressure sensors for smartphones the sensor could be miniaturized further. In the future, we plan to mount the small-sized FootTile sensor units on robotic feet as well as animal feet to investigate force and pressure distribution in rough terrain and granular media in the dynamic locomotion of robots and running animals.

ACKNOWLEDGMENT

We thank Alborz Aghamaleki Sarvestani for the many fruitful discussions in the lab, Annalena Daniels for the material characterization and initial experiments and the Robotics support group (ZWE) for the 3D printed parts.

REFERENCES

[1] R. Kaslin, H. Kolvenbach, L. Paez, *et al.*, "Towards a passive adaptive planar foot with ground orientation and contact force sensing for legged robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2018.

[2] A. Ananthanarayanan, S. Foong, and S. Kim, "A compact two DOF magneto-elastomeric force sensor for a running quadruped," in *2012 IEEE International Conference on Robotics and Automation*, IEEE, May 2012.

[3] M. Hutter, C. Gehring, M. Bloesch, *et al.*, "Starleth: A compliant quadruped robot for fast, efficient, and versatile locomotion," in *Adaptive Mobile Robotics*, WORLD SCIENTIFIC, Jul. 2012, pp. 483–490.

[4] Y. Takahashi, K. Nishiwaki, S. Kagami, *et al.*, "High-speed pressure sensor grid for humanoid robot foot," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2005.

[5] K. Fondahl, D. Kuehn, F. Beinersdorf, *et al.*, "An adaptive sensor foot for a bipedal and quadrupedal robot," in *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, IEEE, Jun. 2012.

[6] M. Lungarella and L. Berthouze, "Robot bouncing: On the synergy between neural and body-environment dynamics," in *Embodied Artificial Intelligence*, Springer Berlin Heidelberg, 2004, pp. 86–97.

[7] X. A. Wu, T. M. Huh, R. Mukherjee, *et al.*, "Integrated ground reaction force sensing and terrain classification for small legged robots," *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 1125–1132, Jul. 2016.

[8] B. Leon, J. J. Rimoli, and C. V. D. Leo, "Elastomer encapsulated pressure sensor with engineered air cavity for force sensing," *IEEE Sensors Journal*, vol. 19, no. 16, pp. 6628–6643, Aug. 2019.

[9] Y. Tenzer, L. P. Jentoft, and R. D. Howe, "The feel of MEMS barometers: Inexpensive and easily customized tactile array sensors," *IEEE Robotics & Automation Magazine*, vol. 21, no. 3, pp. 89–95, Sep. 2014.

[10] M. Y. Chuah and S. Kim, "Enabling force sensing during ground locomotion: A bio-inspired, multi-axis, composite force sensor using discrete pressure mapping," *IEEE Sensors Journal*, vol. 14, no. 5, pp. 1693–1703, May 2014.

[11] P. R. Cavanagh, "A technique for averaging center of pressure paths from a force platform," *Journal of Biomechanics*, vol. 11, no. 10-12, pp. 487–491, Jan. 1978.

[12] F. Ruppert and A. Badri-Spröwitz, "Series elastic behavior of biarticular muscle-tendon structure in a robotic leg," *Frontiers in Neurorobotics*, vol. 13, Aug. 2019.

[13] C. Li, A. M. Hoover, P. Birkmeyer, *et al.*, "Systematic study of the performance of small robots on controlled laboratory substrates," in *Micro- and Nanotechnology Sensors, Systems, and Applications II*, T. George, M. S. Islam, and A. K. Dutta, Eds., SPIE, Apr. 2010.

nature machine intelligence

ARTICLES

https://doi.org/10.1038/s42256-022-00505-4

81

Check for updates

OPEN

Learning plastic matching of robot dynamics in closed-loop central pattern generators

Felix Ruppert ✉ and Alexander Badri-Spröwitz

Animals achieve agile locomotion performance with reduced control effort and energy efficiency by leveraging compliance in their muscles and tendons. However, it is not known how biological locomotion controllers learn to leverage the intelligence embodied in their leg mechanics. Here we present a framework to match control patterns and mechanics based on the concept of short-term elasticity and long-term plasticity. Inspired by animals, we design a robot, Morti, with passive elastic legs. The quadruped robot Morti is controlled by a bioinspired closed-loop central pattern generator that is designed to elastically mitigate short-term perturbations using sparse contact feedback. By minimizing the amount of corrective feedback on the long term, Morti learns to match the controller to its mechanics and learns to walk within 1h. By leveraging the advantages of its mechanics, Morti improves its energy efficiency by 42% without explicit minimization in the cost function.

Animals can locomote with grace and efficiency due to intelligence embodied in their leg designs[1]. Owing to compliant mechanisms in their leg designs, animals can safely traverse rough and unstructured terrain[2,3] in the presence of neural delays and limited actuator power and bandwidth[4,5]. These compliant mechanisms are important components of the natural dynamics of a system. Natural, or passive, dynamics[6] describes the system's passive dynamic behaviour governed by its mechanical characteristics, such as impedance or inertia. More specifically, it describes the dynamics of the unactuated plant transfer function[7].

Compliant mechanisms help to mitigate the interaction forces between walking systems and the environment that are hard to model and are defined by a high degree of uncertainty[8].

To gain a better understanding of the underlying mechanics, bioinspired robots with passive compliant structures that provide the same advantages to robots and simplify the control task have been investigated[2,9–11]. By designing mechanical properties such as impedance[12–14] and spring-loaded inverted pendulum behaviour[3,15,16], the natural dynamics can be designed to achieve viable behaviour with no or reduced control effort, improved energy efficiency and robustness[17–19] comparable to nature.

In a system with strong natural dynamics, the mechanical elements produce forces comparable to the actuators. The challenge of how a controller learns to leverage those natural dynamics then arises. How can animals and bioinspired robots learn to match the control patterns (meaning the desired muscle or motor activation patterns) they produce to their natural dynamics to leverage advantageous passive characteristics?

If the control patterns do not match the natural dynamics, the controller requires additional energy to enforce a desired behaviour (see Supplementary Section 5) as it has to overcome the forces and torques produced by the passive mechanical elements. There is a lack of model-free learning formulations for the matching of control patterns to a given robot's dynamics, especially for robots with strong engineered passive compliant elements. Previous work focused on designing specific aspects of natural dynamics to fit a given control scheme[9,20,21]. In this work we focus on quantifying the match between control patterns and natural dynamics, and how to improve and learn matching in a bioinspired quadruped robot (Fig. 1a).

The neural structure and neuromuscular pathways of animals evolved over many generations and are inherent to each individual at birth[22]. In robotics, the control approach and electrical connections are hardcoded in the design phase before deployment. The timing and intensity of muscle activity patterns in animals have to be matched to the system's natural dynamics as a lifelong learning task[23–25], whereas in robotics the controller has to be learned or tuned for optimal performance during the testing phase or during the robot's lifetime[26–28].

To learn matching the control patterns to the natural dynamics, we separated system perturbations by their time horizon. A one-time stochastic perturbation, like stumbling, should not trigger a long-term adaptation. However, if stumbling occurs frequently, the system should adapt to this systematic discrepancy between the desired control patterns and the system's behaviour governed by the natural dynamics.

To implement this approach we took inspiration from the concept of neuroelasticity and long-term neuroplasticity from neuroscience[29], as well as the concept of elasticity and plasticity in mechanics[30] that describes the reaction to environmental stimuli based on its intensity. A one-time stimulus with low intensity will be mitigated and the control pattern elastically returns to its initial state (Fig. 2a, top). Permanent or frequent stimuli will plastically adapt the control pattern to remove the discrepancy between desired control pattern and natural dynamics behaviour (Fig. 2a, bottom).

In this study we implemented a quadruped robot, Morti, with engineered natural dynamics that is controlled by a central pattern generator (CPG). CPGs are neural networks found in animals that produce rhythmic output signals from non-rhythmic inputs[31,32] for tasks such as chewing, breathing and legged locomotion[33,34]. In robotics CPGs are used as joint trajectory generators[13,35,36] or bioinspired muscle activation pattern generators[37,38]. Feedforward CPGs dictate control and coordination of motor or muscle activation without knowledge of the system's dynamics. These model-free feedforward patterns work well in combination with passively compliant leg designs that provide passive stability and robustness[13,37,38]. By closing feedback loops in CPGs, the system can actively react to unforeseen influences from its environment and mitigate perturbations[32,35,39] such as unstructured terrain.

Dynamic Locomotion Group, Max Planck Institute for Intelligent Systems, Stuttgart, Germany. ✉e-mail: ruppert@is.mpg.de

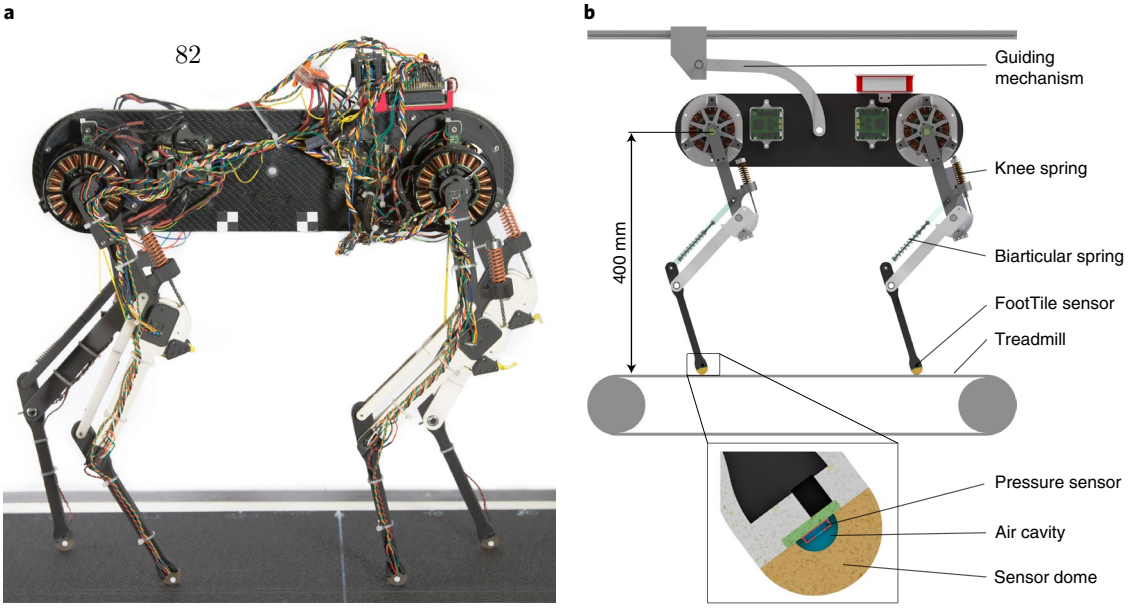

Fig. 1 | Quadruped robot Morti. a, Photograph of Morti. **b**, Render of Morti on top of the treadmill. Morti was constrained to the sagittal plane by a linear rail and lever guiding mechanism that allowed body pitch around its centre of mass. Contact by Morti was measured using four FootTile contact sensors. Inset: close-up cross-section of the FootTile sensor on the foot segment of Morti. The polyurethane sensor dome deforms under loading and the pressure sensor measures the increasing pressure in the air cavity to detect foot contact.

In our quadruped robot Morti, we implemented feedback loops based on continuous sensor data and feedforward reflexes triggered by discrete perturbation events. We then observed Morti's behaviour as measured using sparse feedback from contact sensors on its feet. This short-term feedback acts as a mechanism for mitigating elastic short-term perturbations.

To quantify how well the control pattern matched Morti's natural dynamics, we used the elastic feedback activity as a proxy. If the dynamics did not match, the feedback mechanisms constantly had to intervene to correct for the discrepancy between the commanded and measured behaviour of Morti. The matching of the control patterns needed to be increased plastically.

To improve the matching plastically, we optimized the CPG parameters that generate the control patterns by minimizing the amount of elastic feedback activity (Fig. 2b,c).

Different methods have been used to optimize and tune control patterns, such as optimization[40,41], self-modelling[42], adaptive CPGs[43–45] and machine learning techniques[45–51]. For this study, we applied Bayesian optimization[52,53] to minimize the amount of elastic feedback activity to plastically adapt the CPG parameters.

In previous work, Owaki and Ishiguro[54] presented a control approach that showed spontaneous gait transition based on mechanical coupling ('physical communication'). The CPG was coupled through mechanical coupling. Buchli et al.[45] presented an adaptive oscillator that adapted its frequency to the natural frequency of a spring-loaded inverted pendulum-like simulation model. In their adaptive frequency oscillator approach, the matching of control frequency and natural frequency led to performance improvements and a reduction in energy requirements. Fukuoka et al.[43] implemented short-term reflexes that adapted the robot's controller to the motion of the robot induced by external perturbations. Through a closed-loop CPG that incorporated the 'rolling body motion' the robot could actively adapt to its surrounding. Thandiackal et al.[38] showed that feedback from hydrodynamic pressure in CPGs can lead to self-organized undulatory swimming. However, there are no approaches for long-term plastic matching of control patterns and the natural dynamics of complex walking systems in passive elastic robots at present.

As learning and exploration in hardware are prone to critical failure, the control patterns were first optimized in simulation (Fig. 3a), as is common practice in robotics[47,50,51]. After successful optimization in simulation, the acquired optimal parameter set was applied in hardware (Fig. 3b). We transferred the optimized CPG parameter set into hardware to measure the performance of the real robot and validate the effectiveness of our approach by evaluating a performance measure.

Although optimization and learning in simulation are efficient and cheap, the transfer of control policies can be difficult due to the sim2real gap[50,51,55]. We examined the transferability of our approach by quantifying the sim2real gap by comparing simulation and hardware experiments.

Here we implemented elastic CPG feedback pathways triggered by foot contact. We utilized this elastic feedback activity to mitigate short-term perturbations. Over the long term, we used the feedback activity as a proxy for the mismatching between Morti's natural dynamics and the control pattern. We plastically minimized the required elastic feedback activity through model-free Bayesian optimization. Our approach enabled Morti to learn a trot gait at 0.3 m s^{-1} within 1 h. Matching improved energy efficiency without explicit formulation in the cost function. The improved energy efficiency is evidence of increased matching.

Results

We first examined the performance of the feedback mechanisms in simulation (Fig. 4). The feedback mechanism for late touchdown ($\mathbf{r}_{LTD}$), shown in red, decelerated the phase of the front left leg to wait until ground contact was established. The deceleration was visible in the flatter gradient of the oscillator phase when the mechanism was active.

The early touchdown mechanism (r_{ETD}) triggered a knee pull-up reflex (purple line) to shorten the leg to prevent further impact. In the event of early toeoff (r_{ETO}), shown in yellow, the knee flexion started earlier than instructed by the feedforward CPG. The late toeoff mechanism (r_{LTO}) measured the mismatching of control task and natural dynamics but did not trigger a feedback mechanism.

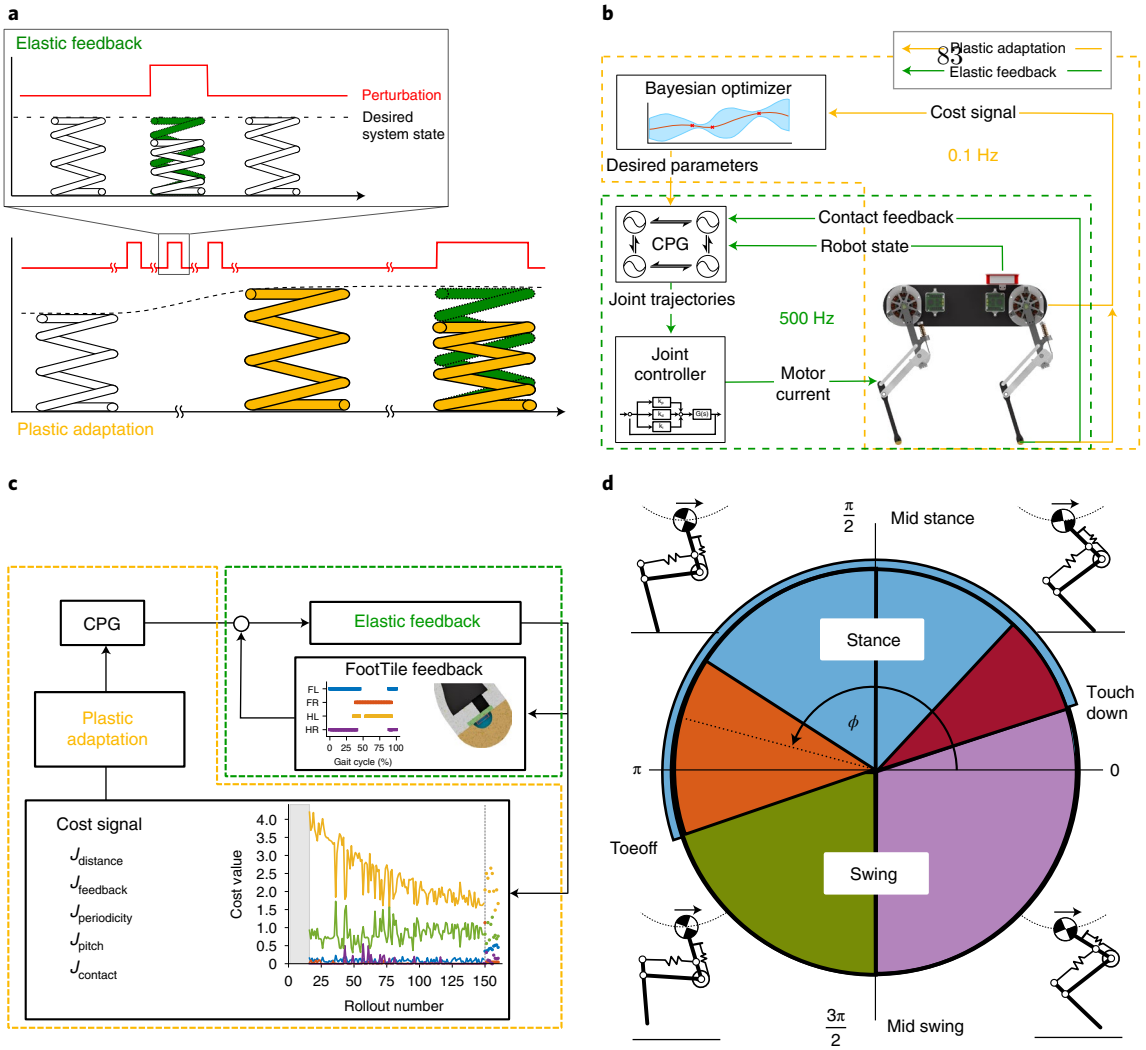

Fig. 2 | Schematic depiction of the elasticity and plasticity framework. a, Short-term elastic feedback (top) and long-term plasticity (bottom). Elastic feedback (green) mitigates stochastic short-term perturbations (red), such as pot holes, that disturb the system (spring) from its desired state (dashed line). Elastic activity is reversible and only active when a perturbation is present, just as a spring only deflects as long as an external force is active and then returns to its initial state. Plasticity (yellow) changes the system behaviour permanently to adapt to long-term active stimuli from the environment. If the same perturbation is frequently present, the system adapts to the perturbation. In our example the spring adapts its set point (spring length, dashed line) and stiffness (spring thickness). In this way, an initial desired system state that might be encoded in the initial control design can be adapted to better deal with perturbations throughout its life span, as well as changing environments. After plastic adaptation the spring deflects less (green, bottom right). **b**, Control structure of Morti. k_p, k_d and k_i are the joint controller gains; $G(s)$ is the plant transfer function in Laplace space s. **c**, Flowchart of the matching approach. The elastic feedback activity mitigates short-term perturbations through sparse contact feedback from the FootTile contact sensors. We measured the amount of elastic feedback activity as a proxy for the mismatching of dynamics. Over a longer time window, the optimizer minimizes the elastic feedback activity to plastically match the control pattern of the CPG to Morti's natural dynamics. F, front; L, left; H, hind; R, right. Colour representations are similar to those in Fig. 5b. **d**, Diagram of a step cycle in phase space (ϕ). The segments are colour-coded by feedback mechanism: late touchdown (red) later than the desired touchdown time ($\delta_{overSwing}$), late toeoff (yellow) later than the desired toeoff time ($\delta_{\phi,knee}$), early toeoff (green) and early touchdown (purple). The stance phase from touchdown to toeoff is shaded blue.

The feedback mechanisms helped Morti to mitigate perturbations stemming from dynamics mismatching. This mitigation effect was especially important in the first rollouts of the optimization, where good dynamics matching was not yet achieved.

In this rollout, the late touchdown mechanism was active for 7% of the step cycle, the early touchdown mechanism was active for 5% of the step cycle, the late toeoff mechanism was active for 8% of the step cycle and the early toeoff mechanism was active for 9% of the step cycle.

We found that 150 rollouts in simulation (Fig. 5a) were sufficient to learn a gait at a speed of 0.3 m s^{-1}. Each rollout in simulation took an average of 23 s for 20 s of simulation runtime on an Intel i7 CPU, making the whole optimization duration roughly 1 h. The hardware rollouts were roughly 1 min long to ensure stable locomotion. When Morti reached stable behaviour, 10 s were evaluated, as in the simulation rollouts.

After initialization of the Gaussian kernel with 15 rollouts with random CPG parameters, the optimizer started to approximate the cost function and performance converged toward the optimum point.

During the whole optimization, Morti fell 16 times or 11% of rollouts. Nine of the failed rollouts occurred during the first 15 rollouts with random CPG parameters.

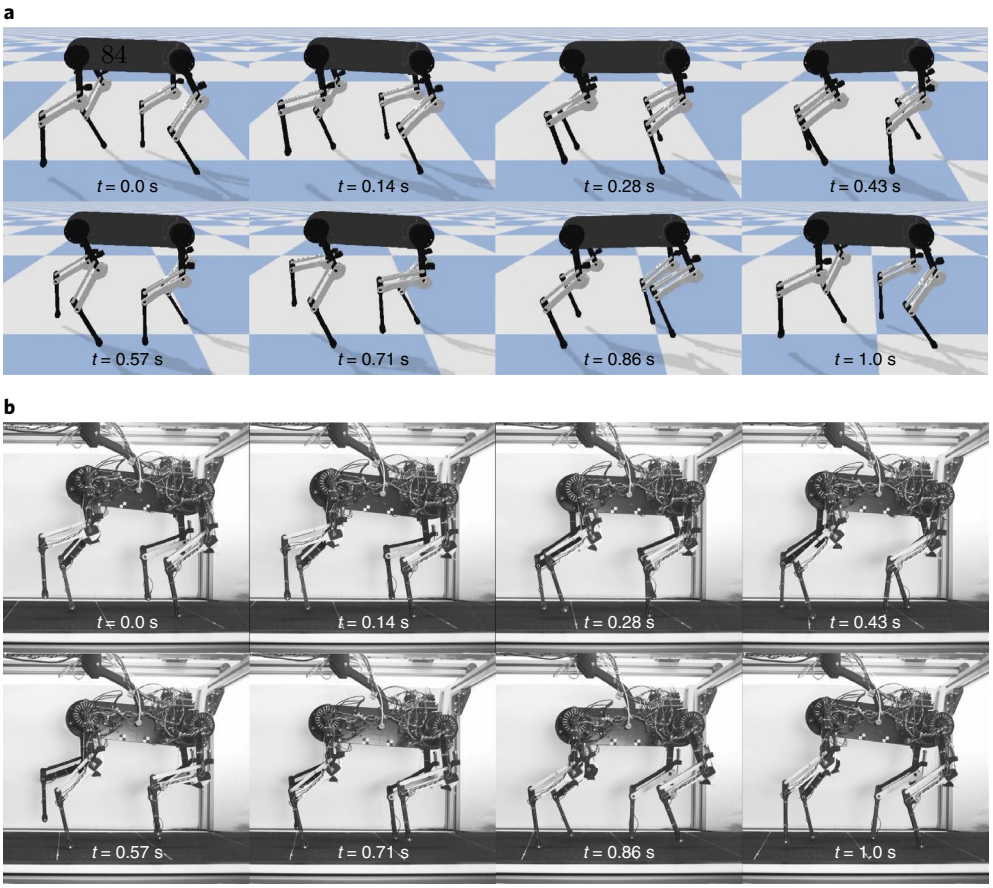

Fig. 3 | Snapshots of the simulation and hardware rollouts. a, Snapshots of the simulated Morti walking during one of the optimization rollouts in the PyBullet multibody simulation at the times (*t*) indicated. **b**, Snapshots of Morti walking on a treadmill during one of the hardware rollouts.

Through optimization, the simulated robot increased its performance from the least-performing rollout (rollout 107, cost 5.62) to the optimal rollout (rollout 109, cost 2.59) by 215%. In comparison, the simulation results transferred to hardware scored a cost between 5.65 and 4.41. The mean simulation cost was 3.49 ± 0.66 and the median simulation cost was 3.34. The mean hardware cost was 4.96 ± 0.38. The best simulation result was 41% lower than the lowest hardware result.

To validate the performance, as well as the differences between simulation and hardware rollouts in detail, we investigated the individual cost factors (Supplementary Table 3) for both simulation and hardware rollouts (Fig. 5a). We found that no single cost factor was responsible for the higher returned cost. Instead, all cost factors were slightly higher and their summation led to the higher cost returned for the hardware results. The distance cost term (J_{distance}) and the feedback cost term (J_{feedback}) contributed the highest difference between the simulation and hardware cost values: J_{distance} had a mean hardware cost of 2.13 ± 0.36 compared with a simulation cost of 1.67 and J_{feedback} had a mean hardware cost of 0.43 ± 0.06 compared with a simulation cost of 0.13. We assumed that the difference was due to modelling assumptions that were made in the simulation. The hardware robot showed a lower speed due to contact losses, gearbox backlash, friction and elasticity in the FootTile sensors. During touchdown, imperfect contact of the feet led to higher feedback activity, which imposed a penalty via the feedback cost term. The body pitch cost term J_{pitch} was in the range of the simulated cost; the mean hardware cost was 0.85 ± 0.22 compared to 0.80 in simulation. Morti showed more body pitch both during the optimization shown here and initial tests for untuned CPG parameters, and it flipped over during several rollouts. This did not happen in hardware—even in early experiments the hardware robot never pitched more than 30°. The periodicity cost term ($J_{\text{periodicity}}$) (hardware: 0.15 ± 0.33; simulation: 0.0) and the contact cost term (J_{contact}) (hardware: 0.12 ± 0.09, simulation: 0.03) behaved similarly in simulation and hardware rollouts. This similarity was expected as both simulation and hardware gaits converged to the desired gait, and the latter three cost terms were introduced to guide the optimizer to find gaits similar to the desired CPG patterns, mostly during the first rollouts.

At the core of our approach, we hypothesized that matching dynamics improves energy efficiency. We therefore explicitly did not incorporate energy efficiency into the cost function. To quantify how matching dynamics improved energy efficiency we calculated a normalized torque as a measure of performance. We chose a normalized mean torque as we showed in previous work[12] that the torque signal has no major oscillations (Supplementary Fig. 5). It is therefore a sufficient representation of the system's energy requirement and is simple to measure both in simulation and in the hardware robot.

$$\tau_{\text{normal}} = \sum_{n=1}^{4} \frac{(\overline{\tau}_{\text{hip},n} + \overline{\tau}_{\text{knee},n})}{\overline{v}_{\text{body}}} \tag{1}$$

where n is the leg index, $\overline{\tau}_{\text{knee}}$ and $\overline{\tau}_{\text{hip}}$ are the mean knee and hip torque per rollout per leg and $\overline{v}_{\text{body}}$ is the mean body velocity of the respective rollout.

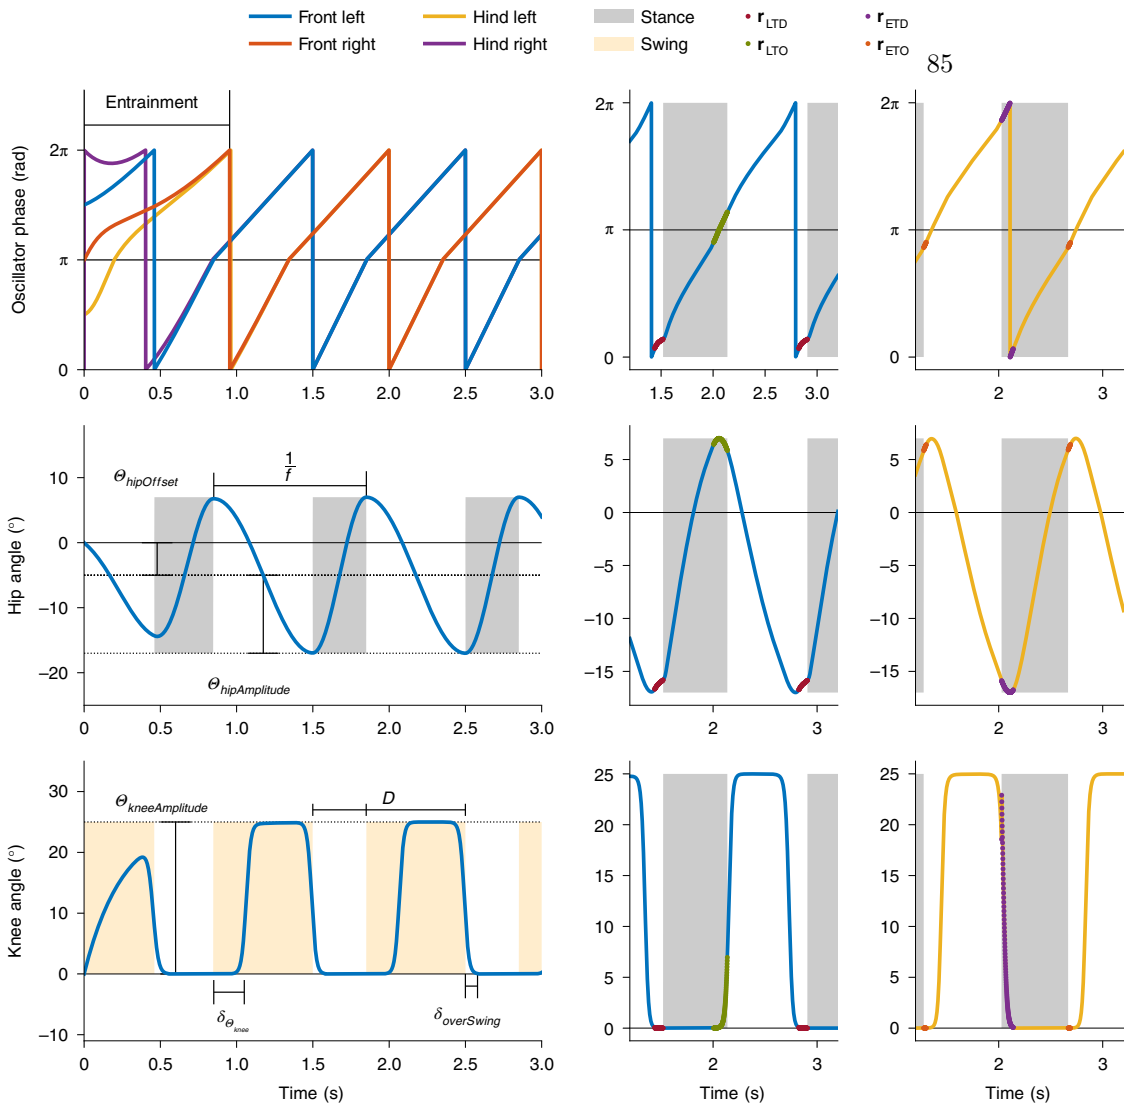

Fig. 4 | CPG parameters and elastic feedback activity. Left: examples of the CPG output for four coupled oscillators and the generated trajectories. The coupled phases are shown at the top, and the hip (middle) and knee (bottom) joint trajectories for one oscillator with their respective CPG parameters are shown below (Supplementary Table 2). Parameters here are $D = 0.35$, $\delta_{\phi,knee} = 0.3$, $\delta_{overSwing} = 0.2$, $f = 1$. Right: simulation results showing the four feedback mechanisms (same colour coding as Supplementary Fig. 2). Data are shown for the front left and front right legs. Late touchdown (red) on the front left leg phase shows the phase delay to wait for touchdown. Early touchdown (purple) on the right leg shows the knee pull-up reflex. Late toeoff (yellow) is shown on the left leg. Early toeoff (green) is shown on the right leg. The stance phase is shaded grey. $\Theta_{hipAmplitude}$, hip amplitude; $\Theta_{kneeAmplitude}$, knee amplitude; $\Theta_{hipOffset}$, hip offset; f, frequency; δ_{knee}, knee phase shift; $\delta_{overSwing}$, knee overswing; D duty factor as described in Supplementary Section 2.

The initial normalized torque was 2.52, and the final value 1.02. The mean normalized torque was 1.7 ± 0.5 and the median normalized torque was 1.55 (Fig. 6). As expected, the normalized torque reduced over the optimization by 42% from plastically unmatched initial conditions (compare with Supplementary Section 5). The reduction in normalized torque as an efficiency measure confirmed our hypothesis, that matching the control pattern to the system's natural dynamics has beneficial effects on energy requirements.

Discussion

We suggested that enabling a locomotion controller to leverage the passively compliant leg structures could increase the energy efficiency indirectly. By minimizing the required elastic feedback activity, the controller learns to increase the matching between its control pattern and the natural dynamics. We showed that 150 optimization rollouts sufficed to learn a stable trot gait on flat ground at a speed of $0.3\,m\,s^{-1}$ from random initial conditions with an optimization duration of 1 h. In our experiment, we showed that matching dynamics is indeed beneficial for energy-efficient locomotion. We calculated a normalized performance measure that showed a decrease in power requirements.

In the normalized torque measure (τ_{normal}), Morti benefited from the increase in distance cost and a reduction in the required torque. Even though Morti increased its speed more than two-fold, the required normalized torque did not increase. Instead, the normalized torque decreased with a trend comparable to the cost function. The improved control pattern matching enabled the controller to leverage the natural dynamics to achieve better performance (Fig. 6).

The designed passive behaviour of Morti enabled a simple matched CPG control structure to leverage the natural dynamics of the leg design. Through sparse binary feedback from touch sensors, the controller was able to elastically mitigate the perturbations stemming from initial mismatching. Through synergy of the natural dynamics and matched CPG, Morti learned to walk on inexpensive

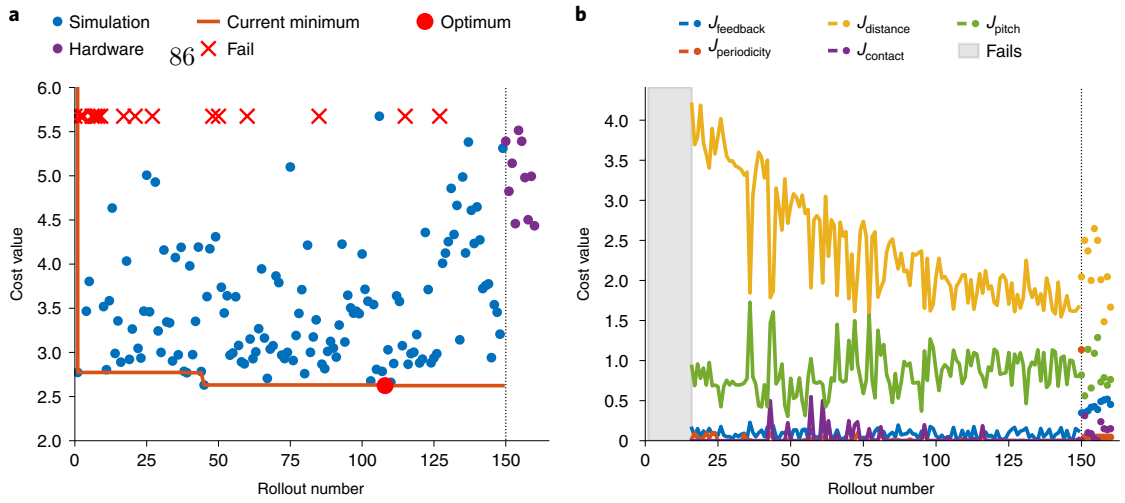

Fig. 5 | Results from plastic adaptation. a, Cost function for the Bayesian optimization. The Bayesian optimization did not monotonically minimize the cost function, so the current minimum is shown. **b**, Individual cost values from the different cost function terms (Supplementary Table 3). The individual cost terms show similar results between the simulation results (lines) and the hardware samples (dots). The mean hardware cost values are similar to the optimal costs from the simulation.

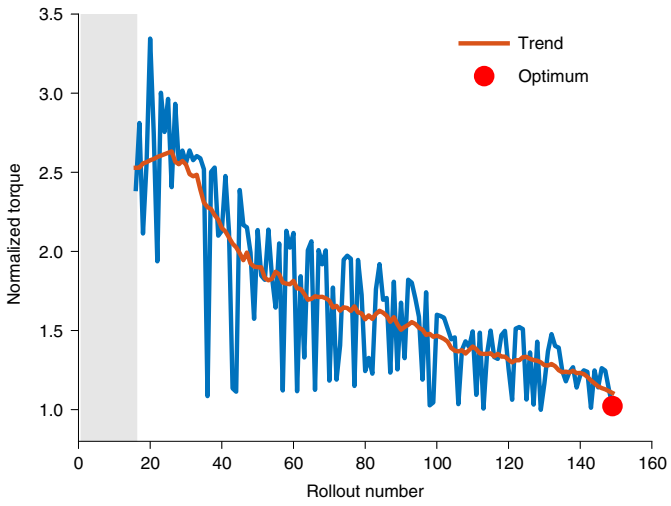

Fig. 6 | Normalized torque performance measure. Normalized torque (torque/speed) reduces over the optimization. Data is sorted by reward. Shaded area marks failed rollouts. By matching the control pattern to Morti's natural dynamics, the required energy for locomotion reduces. Because the controller learns to exploit Morti's natural dynamics, less energy is used to achieve the desired behaviour. The normalized torque was not part of the cost function but minimizes because the matching increases (compare with Supplementary Section 5).

hardware (<€4,000) with low computational power (5 W Raspberry power) and with lower control (500 Hz control loop) and sensor (250 Hz binary sensor signal) frequencies than state-of-the-art model-based locomotion controllers that require high-bandwidth computation and high control frequencies (>2 kHz control frequency, >17 W processor power)[50,56].

Closely examining the cost (Fig. 5b) showed that through dynamics matching and the minimization of $J_{feedback}$, Morti travelled longer distances in the given time, as shown by the improved $J_{distance}$. There was little change in J_{pitch} over the optimization, which was expected because the CPG cannot actively control body pitch. The $J_{periodicity}$ and $J_{contact}$ terms were used as penalty terms for undesired gait characteristics. They are an order of magnitude lower than the remaining cost terms, and only peaked for less performant rollouts.

Although the gait learned in this work was simpler than state-of-the-art full-body control approaches, we provide evidence that minimizing feedback activity stemming from systematic mismatching between the control pattern and the system's natural dynamics provides an alternative learning approach.

Compared with end-to-end learning approaches[50,51], our method requires fewer rollouts. As the underlying control structure (the CPG network) was predefined, it required no approximation in the learning approach first. On the other hand, the versatility of CPG-based locomotion hinges on the complexity of the chosen CPG model. In this proof of concept, we chose a simplified CPG to limit the complexity of the underlying model and thus the technical implementation of our approach. However, we believe that our approach of minimizing elastic feedback activity to increase the matching between control pattern and natural dynamics is not limited by the choice of pattern generator, and could be transferred to other locomotion controllers that possess a metric for the amount of required elastic feedback.

The discrete feedback events (Fig. 2d) allowed the amount of feedback activity to be measured easily in our example. In systems with continuous feedback such as whole-body control[16,56,57], we believe control effort[58] could be an alternative measure. Our approach is also not limited to Bayesian optimization approaches, and we believe it could be used as part of the cost/reward function for different optimization or learning approaches. In this way, our model-free matching approach could scale to different compliant robots. More generally, our approach could be adapted for more versatile control approaches such as model-based full-body control[56,57] or CPGs with more versatile behaviours[43].

Although other studies reported problems with transferring simulation results to hardware (the sim2real gap)[59–61], we successfully transferred our simulation results to the Morti hardware without post-transfer modifications. The hardware performance was comparable both quantitatively (Fig. 5b) and in qualitative observation of the resultant gaits (Supplementary Videos 1 and 2). We believe that because the joint torques of Morti were not calculated from potentially inaccurate model parameters, the sim2real gap is not as evident here. Learned controllers that directly influence joint torques and leg forces might suffer more from the sim2real gap because smaller

inaccuracies between model and hardware behaviour can have a direct effect on the forces exerted on Morti. However, more research will be required to understand the transferability of results for underactuated robots with strong natural dynamics.

In future work, we intend to extend the CPG, taking body pitch into account when generating the hip trajectories as done by ref. [39]. With an inertial measurement unit the body pitch could be fed back into the CPG. In the current formulation, the CPG assumes no body pitch and relies on the robustness the passive elasticity adds to the system to compensate the existing body pitch. Abduction/adduction degrees of freedom with their respective feedback loops[39,62] could also be added to Morti to enable 3D locomotion without a guiding mechanism. The optimization loop could be implemented to run online on the hardware robot's computer. With online optimization and 3D locomotion, it would become possible to investigate the lifelong adaptation of the CPG control patterns to changing ground conditions and surface properties over extended time windows, as well as adaptations to wear and tear throughout Morti's lifetime.

In this Article, we examined how a walking system with limited control and sensor bandwidth could learn to leverage the intelligence embodied in its leg mechanics. Energy efficiency and speed are often used as criteria to evaluate performance in robotic systems. Here we proposed an additional measure that focuses on the synergy of passive mechanical structures and neural control. By separating feedback by its time horizon, we achieved perturbation mitigation in the short term and at the same time quantified the mismatching of control patterns and natural dynamics. We optimized the long-term performance of the system and adapted the controller to its mechanical system. Although investigated in a robotic surrogate, our findings could provide a new perspective on how learning in biological systems might happen in the presence of neural limitations and sparse feedback. Matching is probably not the sole driving factor in animal learning, but our study suggests that a quantitative measure for 'long-term learning from failure' could in part be influenced by the goal of maximizing the synergy between locomotion control and the robot's or animal's mechanical walking system. In contrast to task-specific cost functions such as speed or energy efficiency, our matching approach provides an intrinsic motivation to leverage the embodied intelligence in the natural dynamics as much as possible.

Methods

For both the experimentation and simulation, we designed and implemented quadruped robot Morti. Morti has a monoarticular knee spring and a biarticular spring between hip and foot that provide series elastic behaviour[12]. It was controlled by a closed-loop CPG. Through reflex-like feedback mechanisms, Morti could elastically mitigate short-term perturbations. To minimize the elastic activity, we implemented a Bayesian optimizer that plastically matched the control pattern to Morti's natural dynamics.

Robot mechanics. Morti consists of four 'biarticular legs' (Fig. 1b; ref. [12]) mounted on a carbon fibre body. Each leg has three segments: femur, shank and foot. The femur and foot segments are connected by a spring-loaded parallel mechanism that mimics the biarticular muscle–tendon structure formed by the gastrocnemius muscle–tendon group in quadruped animals[63]. A knee spring inspired by the patellar tendon in animals provided passive elasticity of the knee joint.

Morti walked on a treadmill and was constrained to the sagittal plane by a linear rail that allowed body pitch (Fig. 1b). It was instrumented with joint angle sensors, position sensors and the treadmill speed sensor. To measure ground contact, four FootTile sensors[64] were mounted on Morti's feet. Using a threshold, these analogue pressure sensors could be used to determine whether it established ground contact. Detailed descriptions of the experimental set-up can be found in Supplementary Section 1.

Simulation. We implemented the simulation in PyBullet[59], a multibody simulator based on the bullet physics engine (Fig. 3a). The robot mechanics were derived from the mechanical robot and its computer-aided design model (Supplementary Table 1). To increase the matching between the robot hardware and simulation, we imposed motor limits and set the motor controller to resemble the real actuator limits[55]. The simulation ran at 1 kHz, the CPG control loop ran at 500 Hz and ground contacts are polled at 250 Hz to resemble the hardware implementation. The control frequency was chosen for technical reasons to guarantee stable position control and fast data acquisition. It could be lower, as shown in similar systems[5,13], to more closely resemble the neural delays and low technical complexity in animals.

CPG. The CPG used in this work was a modified Hopf oscillator[13,35,62] that was modelled in phase space. More biologically accurate and biomimetic CPG models do exist[37,43,54]; we chose this representation because of its reduced parameter space while retaining the functionality required to generate joint trajectories for locomotion. Similar to their biological counterparts, CPGs can be entrained through feedback from external sensory input[38] or from internal coupling to neighbouring nodes[54]. The CPG in this work consisted of four coupled nodes, representing the four legs (see Supplementary Fig. 1). The hip and knee of each leg were coupled through a variable phase shift. Depending on the desired phase shifts in between oscillator nodes (legs), a variety of gaits can be implemented by adapting the phase difference matrix while keeping the network dynamics identical (Supplementary Section 6). The joint trajectories generated by the CPG are described by eight parameters (Supplementary Table 6): the hip offset ($\Theta_{\text{hipOffset}}$) and hip amplitude ($\Theta_{\text{hipAmplitude}}$) describe the hip trajectory, the knee offset amplitude ($\Theta_{\text{kneeOffset}}$) describes the knee flexion, the frequency f determines the robot's overall speed, duty factors (D) describe the ratio of stance phase to flight phase, the knee phase shift ($\delta_{\phi,\text{knee}}$) describes the phase shift between hip protraction and knee flexion and overswing ($\delta_{\text{overSwing}}$) describes the amount of swing leg retraction[65]. The mathematical description of the CPG dynamics can be found in Supplementary Section 2.

Elasticity. As the CPG implemented here was written as a model-free feedforward network, it could be difficult to find parameters that lead to viable gaits with given robot dynamics. Essentially, the CPG commands desired trajectories without knowledge of the robot's natural dynamics. In the worst-case scenario the CPG would command behaviour that the robot cannot fulfil because of its own natural dynamics and mechanical limitations such as inertia, motor speed or torque limitations. To address this shortcoming, feedback can be used to mitigate the differences between desired and measured behaviour.

The feedback implemented here was an adaptation from Righetti et al.[35] that has been shown to aid in entrainment and can mitigate perturbations in foot contact information. This contact information can be integrated into the CPG to measure timing differences between the desired and measured trajectories.

The trajectories created by the CPG can be influenced through feedback by changing the CPG dynamics (meaning accelerating or decelerating the CPG's phases). Alternatively, feedback can influence the generated joint angle trajectories. During a step cycle (Fig. 2d), contact signals were used for several feedback mechanisms (Supplementary Fig. 2). The feedback mechanisms reacted to timing discrepancies in the touchdown and toeoff events and corrected the CPG trajectories if Morti established or lost ground contact earlier or later than instructed by the CPG. Righetti et al.[35] showed how these feedback mechanisms can actively stabilize a CPG controlled robot. We adapted these mechanisms to a phase-space CPG formulation and robot hardware to achieve similar performance in a different class of robot.

Using the feedback mechanisms implemented in the CPG, we corrected the timing of discrete events. If touchdown and toeoff happened earlier or later than commanded by the feedforward control pattern, the individual phases of each leg (CPG node) could be accelerated or decelerated to correct Morti's behaviour. We implemented a phase deceleration when touchdown was delayed (Fig. 2d, red). We accelerated knee flexion when a foot lost ground contact too early (Fig. 2d, orange), in addition to a phase deceleration when toeoff occurred later than commanded (Fig. 2d, green). We combined these feedback mechanisms with a knee pull-up reflex when a leg hit the ground too early to mimic a patellar reflex as adapted from ref. [66]. If a leg hit the ground during the swing phase (Fig. 2d, purple), the knee flexed in a predefined trajectory to generate more ground clearance. In addition to the mechanism in ref. [66], we disabled the hip motor from interfering in the passive impact mitigation of the mechanical leg springs. An in-depth description of the feedback mechanisms can be found in Supplementary Section 3 and the CPG output with active feedback can be seen in Fig. 4.

Plasticity. To match the CPG to the robot dynamics, we wanted to tune the CPG parameters p_m to achieve optimal performance. To do so, we evaluated the performance of Morti for a number of steps. The timescale of the optimization was designed to be much bigger than the frequency of the elastic feedback activity mechanisms ($\leq$0.1 Hz versus $\geq$100 Hz). Consequently, the effects of the elastic feedback activity were minimized and small perturbations within one step were not captured in the plastic optimization that will only improve long-term performance.

To achieve long-term (close to) optimal behaviour we used Bayesian optimization for its global optimization capabilities, data efficiency and robustness to noise[52,53].

Bayesian optimization. Bayesian optimization is a black-box optimization approach that uses Gaussian kernels for function approximation. It is model-free, derivative-free and has been used successfully in many robotic optimization approaches[67–69]. Bayesian optimization is favoured over other data-driven

optimization and learning approaches because of its data efficiency for ten or more parameters.

We implemented a Bayesian optimizer that was based on skopt gp_minimize (ref. [70]). The optimizer evaluated the PyBullet simulation for 10 s (approximately ten step cycles) of each rollout with a cost function. Morti walked for 10 s to entrain the CPG from its initial condition (standing still; see Fig. 4, top) to achieve steady-state behaviour before the evaluation began. One complete rollout therefore took 20 s. We sampled 15 rollouts with random CPG parameters before approximating the cost function. Then we optimized for 135 rollouts with the gp_hedge acquisition function, which is a probabilistic choice of the lower confidence bound, negative expected improvement and negative probability of improvement.

To reduce complexity we limited the parameter space to six parameters. The six parameters are $\Theta_{\text{hipOffset}}$, $\Theta_{\text{hipAmplitude}}$, D_{front} and D_{hind}, $\delta_{\phi,\text{knee}}$ and $\delta_{\text{overSwing}}$ (Fig. 4). More parameters would probably improve performance more, but would also lead to more corner cases where the selected cost function could be exploited by the optimizer and result in undesired gait characteristics (such as skipping gaits) or gaits where the feet drag on the ground. For this proof of concept, we chose independent duty factors D_{front} and D_{hind} to allow some front–hind asymmetry that could help the optimizer find gaits that reduce body pitch. Where only one CPG parameter was selected, the parameter was used for all legs. For simplicity, we also fixed the frequency to $f = 1$ Hz to reduce experimental cost in terms of hardware wear from violent motions at high speed. The hip amplitude $\Theta_{\text{kneeAmplitude}}$ is set to 30° to ensure adequate ground clearance.

Cost function. We evaluated Morti using a cost function comprising three major components. The first component evaluated J_{feedback}, specifically the percentage of a step cycle that one of the elastic feedback mechanisms was active:

$$J_{\text{feedback}} = \frac{1}{fnT}\sum_{n=0}^{4}\sum_{t=0}^{T}\mathbf{r}_{\text{ETO}} + \mathbf{r}_{\text{ETD}} + \mathbf{r}_{\text{LTO}} + \mathbf{r}_{\text{LTD}} \tag{2}$$

where J_{feedback} is the average percentage of active feedback per step and leg, T is the evaluation time and $\mathbf{r}_{\text{ETO}}$, $\mathbf{r}_{\text{ETD}}$, $\mathbf{r}_{\text{LTO}}$ and $\mathbf{r}_{\text{LTD}}$ are the time vectors when the specific feedback was active for each leg.

The second component evaluates the distance travelled (J_{distance} $(1/m)$) to encourage forward locomotion.

$$J_{\text{distance}} = \frac{1}{x_{\text{body}}} \tag{3}$$

where x_{body} is the centre-of-mass position in the walking direction.

The third component penalized deviations from the commanded gait characteristics. It ensured that Morti moved with a low mean J_{pitch}:

$$J_{\text{pitch}} = \| \max(\alpha_{\text{pitch}}) - \min(\alpha_{\text{pitch}}) \| \tag{4}$$

and was calculated as the difference between the mean minimum and the mean maximum of body pitch angle (α_{pitch}) of all strides during one rollout.

It also imposed a penalty if more than one ground contact phase per foot and step (J_{contact} [% of step cycle]) occurred, as would take place during stumbling or dragging of the feet.

$$J_{\text{contact}} = \frac{1}{fnT}\sum\nolimits_{n=0}^{4}\sum\nolimits_{t=0}^{T}\left(\| \frac{\mathrm{d}}{\mathrm{dt}}\ \text{contact}\ \| > 0\right) \tag{5}$$

where J_{contact} is the mean number of flight-stance changes per step, t is time, T is the evaluation time and contact is the contact sensor data matrix for all four legs.

The third component penalized differences between the desired gait frequency and the measured gait frequency to prevent non-periodic gaits or multi-step gaits ($J_{\text{periodicity}}$ (Hz)).

$$\begin{aligned} S_{\text{pitch}} &= \alpha_{\text{pitch}} \times \alpha_{\text{pitch}} \\ f_{\text{bodyPitch}} &= \| \max S_{\text{pitch}} \| \end{aligned} \tag{6}$$

$$J_{\text{periodicity}} = \sqrt{\frac{1}{N}\left(f_{\text{bodyPitch}} - f_{\text{cpg}}\right)^2} \tag{7}$$

where S_{pitch} is the frequency spectrum of α_{pitch}, $f_{\text{bodyPitch}}$ is the frequency of the body pitch measurement, $J_{\text{periodicity}}$ is the standard deviation of the periodicity measure and f_{cpg} is the commanded CPG frequency.

Detailed descriptions can be found in Supplementary Section 4.

Hardware rollouts. To validate the optimal set of CPG parameters from simulation, we tested the same parameters on the hardware robot. The hardware controller had the same elastic mechanisms described in the 'Elasticity' section. We tested ten parameter sets and randomly varied the CPG parameters obtained from simulation by ≤10% to validate the hardware cost function around the optimal point found in simulation. We then evaluated Morti's performance with the same cost function used for the simulation. As in the simulation, Morti ran for 10 s to entrain itself. The performance was measured for 10 s after Morti converged on a stable gait.

Videos of Morti walking are available as Supplementary Videos 1 and 2.

Data availability

The experimental data are available at https://doi.org/10.17617/3.XDOQNW (ref. [71]). The robot model and CAD design are available for non-commercial use at the same link.

Code availability

The code and a demo are available at https://doi.org/10.17617/3.XDOQNW (ref. [71]).

Received: 25 October 2021; Accepted: 19 May 2022;
Published online: 18 July 2022

References

1. Iida, F. *Embodied Artificial Intelligence* (Springer, 2004).
2. Alexander, R. Elastic energy stores in running vertebrates. *Am. Zool.* **24**, 85–94 (1984).
3. Blickhan, R. The spring-mass model for running and hopping. *J. Biomech.* **22**, 1217–1227 (1989).
4. More, H. L. & Donelan, J. M. Scaling of sensorimotor delays in terrestrial mammals. *Proc. R. Soc. B* **285**, 20180613 (2018).
5. Ashtiani, M. S., Sarvestani, A. A. & Badri-Spröwitz, A. T. Hybrid parallel compliance allows robots to operate with sensorimotor delays and low control frequencies. *Front. Robot. AI* **8**, 645748 (2021).
6. Collins, S. Efficient bipedal robots based on passive-dynamic walkers. *Science* **307**, 1082–1085 (2005).
7. Franklin, G. *Feedback Control of Dynamic Systems* (Prentice Hall, 2002).
8. Daley, M. A. Running over rough terrain: guinea fowl maintain dynamic stability despite a large unexpected change in substrate height. *J. Exp. Biol.* **209**, 171–187 (2006).
9. Renjewski, D., Spröwitz, A., Peekema, A., Jones, M. & Hurst, J. Exciting engineered passive dynamics in a bipedal robot. *IEEE Trans. Robot.* **31**, 1244–1251 (2015).
10. Luo, X. & Xu, W. Planning and control for passive dynamics based walking of 3D biped robots. *J. Bionic Eng.* **9**, 143–155 (2012).
11. Ruina, A. Passive dynamics is a good basis for robot design and control, not! *Princeton University MAE* https://mae.princeton.edu/about-mae/events/passive-dynamics-good-basis-robot-design-and-control-not (2017).
12. Ruppert, F. & Badri-Spröwitz, A. Series elastic behavior of biarticular muscle-tendon structure in a robotic leg. *Front. Neurorobotics* **13**, 8 (2019).
13. Spröwitz, A. et al. Towards dynamic trot gait locomotion: design, control, and experiments with cheetah-cub, a compliant quadruped robot. *Int. J. Robot. Res.* **32**, 932–950 (2013).
14. Lee, H. & Hogan, N. Time-varying ankle mechanical impedance during human locomotion. *IEEE Trans. Neur. Syst. Rehab. Eng.* **23**, 755–764 (2015).
15. Tedrake, R. Zhang, T. W. & Seung, H. S. Learning to walk in 20 minutes. In *IEEE International Conference on Robotics and Automation* 4656–4661 (IEEE, 2004).
16. Bhounsule, P. A. et al. Low-bandwidth reflex-based control for lower power walking: 65 km on a single battery charge. *Int. J. Robot. Res.* **33**, 1305–1321 (2014).
17. Geyer, H., Seyfarth, A. & Blickhan, R. Spring-mass running: simple approximate solution and application to gait stability. *J. Theor. Biol.* **232**, 315–328 (2005).
18. Geyer, H., Seyfarth, A. & Blickhan, R. Compliant leg behaviour explains basic dynamics of walking and running. *Proc. R. Soc. B* **273**, 2861–2867 (2006).
19. Rummel, J. & Seyfarth, A. Stable running with segmented legs. *Int. J. Robot. Res.* **27**, 919–934 (2008).
20. Rummel, J., Blum, Y. & Seyfarth, A. Robust and efficient walking with spring-like legs. *Bioinspir. Biomim.* **5**, 046004 (2010).
21. Kenneally, G., De, A. & Koditschek, D. E. Design principles for a family of direct-drive legged robots. *IEEE Robot. Autom. Lett.* **1**, 900–907 (2016).
22. Bicanski, A. et al. Decoding the mechanisms of gait generation in salamanders by combining neurobiology, modeling and robotics. *Biol. Cybernet.* **107**, 545–564 (2013).
23. Dominici, N. et al. Locomotor primitives in newborn babies and their development. *Science* **334**, 997–999 (2011).
24. Grasso, R. et al. Distributed plasticity of locomotor pattern generators in spinal cord injured patients. *Brain* **127**, 1019–1034 (2004).
25. Kudithipudi, D. et al. Biological underpinnings for lifelong learning machines. *Nat. Mach. Intell.* **4**, 196–210 (2022).
26. Marjaninejad, A., Urbina-Meléndez, D., Cohn, B. A. & Valero-Cuevas, F. J. Autonomous functional movements in a tendon-driven limb via limited experience. *Nat. Mach. Intell.* **1**, 144–154 (2019).
27. Mastalli, C. et al. Trajectory and foothold optimization using low-dimensional models for rough terrain locomotion. In *2017 IEEE International Conference on Robotics and Automation* 1096–1103 (IEEE, 2017).

28. Kwiatkowski, R. & Lipson, H. Task-agnostic self-modeling machines. *Sci. Robot.* **4**, 26 (2019).
29. Mitteroecker, P. & Stansfield, E. A model of developmental canalization, applied to human cranial form. *PLoS Comput. Biol.* **17**, e1008381 (2021).
30. Sadd, M. *Elasticity: Theory, Applications, and Numerics* (Elsevier/Academic Press, 2009).
31. Marder, E. & Bucher, D. Central pattern generators and the control of rhythmic movements. *Curr. Biol.* **11**, R986–R996 (2001).
32. Matsuoka, K. Mechanisms of frequency and pattern control in the neural rhythm generators. *Biol. Cybernet.* **56**, 345–353 (1987).
33. Bizzi, E., Tresch, M. C., Saltiel, P. & d'Avella, A. New perspectives on spinal motor systems. *Nat. Rev. Neurosci.* **1**, 101–108 (2000).
34. Dickinson, M. H. How animals move: an integrative view. *Science* **288**, 100–106 (2000).
35. Righetti, L. & Ijspeert, A. J. Pattern generators with sensory feedback for the control of quadruped locomotion. In *2008 IEEE International Conference on Robotics and Automation* 819–824 (IEEE, 2008).
36. Xie, F., Zhong, Y., Du, R. & Li, Z. Central pattern generator (CPG) control of a biomimetic robot fish for multimodal swimming. *J. Bion. Eng.* **16**, 222–234 (2019).
37. Ijspeert, A. J., Crespi, A., Ryczko, D. & Cabelguen, J.-M. From swimming to walking with a salamander robot driven by a spinal cord model. *Science* **315**, 1416–1420 (2007).
38. Thandiackal, R. et al. Emergence of robust self-organized undulatory swimming based on local hydrodynamic force sensing. *Sci. Robot.* **6**, eabf6354 (2021).
39. Sartoretti, G. et al. Central pattern generator with inertial feedback for stable locomotion and climbing in unstructured terrain. In *2018 IEEE International Conference on Robotics and Automation* 5769–5775 (IEEE, 2018).
40. Oliveira, M., Matos, V., Santos, C. P. & Costa, L. Multi-objective parameter CPG optimization for gait generation of a biped robot. In *2013 IEEE International Conference on Robotics and Automation* 3130–3135 (IEEE, 2013).
41. Yeganegi, M. H. et al. Robust humanoid locomotion using trajectory optimization and sample-efficient learning. In *International Conference on Humanoid Robots* 170–177 (IEEE, 2019).
42. Bongard, J., Zykov, V. & Lipson, H. Resilient machines through continuous self-modeling. *Science* **314**, 1118–1121 (2006).
43. Fukuoka, Y., Kimura, H., Hada, Y. & Takase, K. Adaptive dynamic walking of a quadruped robot 'Tekken' on irregular terrain using a neural system model. In *2003 IEEE International Conference on Robotics and Automation* IEEE Cat. No.03CH37422 (IEEE, 2003).
44. Buchli, J. & Ijspeert, A. J. Self-organized adaptive legged locomotion in a compliant quadruped robot. *Auton. Robots* **25**, 331–347 (2008).
45. Buchli, J., Righetti, L. & Ijspeert, A. J. in *Advances in Artificial Life* 210–220 (Springer, 2005).
46. Pearlmutter, B. A. Learning state space trajectories in recurrent neural networks. *Neur. Comput.* **1**, 263–269 (1989).
47. Heim, S., Ruppert, F., Sarvestani, A. A. & and Spröwitz, A. Shaping in practice: training wheels to learn fast hopping directly in hardware. In *2018 IEEE International Conference on Robotics and Automation* 5076–5081 (IEEE, 2018).
48. Matsubara, T., Morimoto, J., Nakanishi, J., Sato, M. & Doya, K. Learning CPG-based biped locomotion with a policy gradient method. In *5th IEEE-RAS International Conference on Humanoid Robots* 208–213 (IEEE, 2005).
49. Nakamura, Y., Mori, T. & Ishii, S. Natural policy gradient reinforcement learning for a CPG control of a biped robot. In *Parallel Problem Solving from Nature VIII* (eds Yao, X. et al.) (Springer, 2004).
50. Siekmann, J. et al. Learning memory-based control for human-scale bipedal locomotion. In *Robotics: Science and Systems Conference* (Robotics: Science and Systems Foundation, 2020).
51. Peng, X. B. et al. Learning agile robotic locomotion skills by imitating animals. In *Robotics: Science and Systems XVI* (Robotics: Science and Systems Foundation, 2020).
52. Calandra, R., Seyfarth, A., Peters, J. & Deisenroth, M. P. Bayesian optimization for learning gaits under uncertainty. *Ann. Math. Artif. Intell.* **76**, 5–23 (2015).
53. Mockus, J. *Bayesian Approach to Global Optimization* (Springer, 2012).
54. Owaki, D. & Ishiguro, A. A quadruped robot exhibiting spontaneous gait transitions from walking to trotting to galloping. *Sci. Rep.* **277**, 3 (2017).
55. Tan, J. et al. Sim-to-real: learning agile locomotion for quadruped robots. In *Robotics: Science and Systems XIV* (Robotics: Science and Systems Foundation, 2018).
56. Park, H.-W. & Kim, S. The MIT cheetah, an electrically-powered quadrupedal robot for high-speed running. *J. Robot. Soc. Jpn* **32**, 323–328 (2014).
57. Hutter, M. et al. ANYmal: a highly mobile and dynamic quadrupedal robot. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2016).
58. Haeufle, D. F. B., Günther, M., Wunner, G. & Schmitt, S. Quantifying control effort of biological and technical movements: an information-entropy-based approach. *Phys. Rev. E* **89**, 012716 (2014).
59. Coumans, E. & Bai, Y. pybullet version (3.0.7) (2016); http://pybullet.org
60. Rosser, K., Kok, J., Chahl, J. & Bongard, J. Sim2real gap is non-monotonic with robot complexity for morphology-in-the-loop flapping wing design. In *2020 IEEE International Conference on Robotics and Automation* 7001–7007 (IEEE, 2020).
61. Heiden, E., Millard, D., Coumans, E., Sheng, Y. & Sukhatme, G. S. NeuralSim: augmenting differentiable simulators with neural networks. In *IEEE International Conference on Robotics and Automation* 9474–9481 (IEEE, 2021).
62. Spröwitz, A. T. et al. Oncilla robot: a versatile open-source quadruped research robot with compliant pantograph legs. *Front. Robot. AI* **5**, (2018).
63. Witte, H. et al. Transfer of biological principles into the construction of quadruped walking machines. In *Second International Workshop on Robot Motion and Control* IEEE Cat. No.01EX535 245–249 (Poznan University Technology, 2001).
64. Ruppert, F. & Badri-Spröwitz, A. FootTile: a rugged foot sensor for force and center of pressure sensing in soft terrain. In *2020 IEEE International Conference on Robotics and Automation* 4810–4816 (IEEE, 2020).
65. Seyfarth, A., Geyer, H. & Herr, H. Swing-leg retraction: a simple control model for stable running. *J. Exp. Biol.* **206**, 2547–2555 (2003).
66. Focchi, M. et al. in *Nature-Inspired Mobile Robotics* (eds Waldron, K. J. et al.) 443–450 (World Scientific, 2013).
67. Gianni, M., Garcia, M. A. R. & Pirri, F. Learning the dynamics of articulated tracked vehicles. *Zenodo* https://doi.org/10.5281/zenodo.1124704 (2016).
68. Marco, A., Hennig, P., Bohg, J., Schaal, S. & Trimpe, S. Automatic LQR tuning based on Gaussian process global optimization. In *2016 IEEE International Conference on Robotics and Automation* 270–277 (IEEE, 2016).
69. Seyde, T., Carius, J., Grandia, R., Farshidian, F. & Hutter, M. Locomotion planning through a hybrid Bayesian trajectory optimization. In *2019 International Conference on Robotics and Automation* 5544–5550 (IEEE, 2019).
70. Head, T., Kumar, M., Nahrstaedt, H., Louppe, G. & Shcherbatyi, I. scikit-optimize version 0.9.0 (Python Software Foundation, 2020); https://pypi.org/project/scikit-optimize/
71. Ruppert, F. & Badri-Spröwitz, A. Learning plastic matching of robot dynamics in closed-loop central pattern generators: data. *Edmond* https://doi.org/10.17617/3.XDOQNW (2022).

Acknowledgements

F.R. thanks the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. This work was made possible thanks to a Max Planck Group Leader grant awarded to A.B.S. by the Max Planck Society. We thank ZWE robotics for support for 3D printing, the MPI machine shop for support for metal manufacturing, M. Khadiv and L. Righetti for discussions concerning Bayesian optimization, A. Sarvestani for the many fruitful discussions and R. Petereit for his input concerning Morti's firmware.

Author contributions

F.R. contributed to the concept, design, experiments, data analysis and writing. A.B.S. contributed to the concept and provided feedback and supervision.

Funding

Open access funding provided by Max Planck Society.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s42256-022-00505-4.

Correspondence and requests for materials should be addressed to Felix Ruppert.

Peer review information *Nature Machine Intelligence* thanks Joao Ramos, Fernando Perez-Peña, Ren Qinyuan and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.