

Safe & Accurate at Speed with Tendons: A Robot Arm for Exploring Dynamic Motion

Simon Guist¹, Jan Schneider¹, Hao Ma¹, Le Chen¹, Vincent Berenz¹, Julian Martus², Heiko Ott¹, Felix Grüninger², Michael Muehlebach¹, Jonathan Fiene², Bernhard Schölkopf¹ and Dieter Büchler¹



Abstract—Operating robots precisely and at high speeds has been a long-standing goal of robotics research. Balancing these competing demands is key to enabling the seamless collaboration of robots and humans and increasing task performance. However, traditional motor-driven systems often fall short in this balancing act. Due to their rigid and often heavy design exacerbated by positioning the motors into the joints, faster motions of such robots transfer high forces at impact. To enable precise and safe dynamic motions, we introduce a four degree-of-freedom (DoF) tendon-driven robot arm. Tendons allow placing the actuation at the base to reduce the robot’s inertia, which we show significantly reduces peak collision forces compared to conventional robots with motors placed near the joints. This configuration, paired with pneumatic muscles, enables high-force, accelerated motions and improves safety through passive compliance. This capacity for rapid motion exploration is particularly valuable for generating extensive datasets of dynamic motions autonomously, without additional direct human oversight. We leverage this feature to generate a proprioceptive dataset of 25 days of diverse robot motions that highlights its robustness and reliability. We also demonstrate its ease of control by quantifying the nonlinearities of the system and the performance on a challenging dynamic table tennis task learned from scratch using reinforcement learning. We open-source the entire hardware design, which can be largely 3D printed, the control software, and the motions dataset can be found at [link removed to comply with anonymity requirements].

I. INTRODUCTION

Tasks such as playing table tennis, harvesting delicate berries, or carrying heavy objects differ inherently in their force, precision, and compliance demands. Humans naturally excel at each of these tasks by encompassing a broad array of various motion characteristics, including slow and precise motions, high-force trajectories, and fast as well as highly accelerated movements. Replicating the full range of these capabilities, along with the inherent safety properties of human arm movements, such as compliance and backdrivability, has been a challenging endeavor in robotics research.

¹Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany. firstname.lastname@tuebingen.mpg.de

²Max Planck Institute for Intelligent Systems, 70569 Stuttgart, Germany. lastname@is.mpg.de

The desired robot capabilities can be roughly divided into (i) the achievable speed and (ii) force, (iii) the closeness to human size, (iv) the ease of control, and (v) the safety properties of the robot and environment at impact. High speed and force are essential to accomplish tasks quickly while potentially handling heavy objects. Executing accurate motions with high safety standards enables handling delicate objects or tasks with low tolerance for imprecision, such as in manufacturing. Inherent safety extends the set of allowed contacts and enables the control algorithm to be less conservative and take more risks to optimize for performance. In that manner, robots could move at higher speeds due to the reduced negative consequences of unintended impacts. The amount of anthropomorphism is crucial when robots should act in human-created environments and tasks, potentially even working alongside humans. The size of the robot affects all categories: while smaller robots can be precise, fast, and safe, executing high forces is rather difficult. On the other hand, a big robot can generate forceful and fast trajectories but sacrifices safety at impact.

Distinct robot designs excel in each of these categories. Industrial motor-driven robots are traditionally heavy and rigid; these properties are necessary to describe the system precisely using the rigid body dynamics equations, which can be used to easily attain high-quality control. Paired with strong motors, industrial robots additionally excel at maximum force and speed. On the downside, industrial robots easily cause damage to themselves and the environment in collisions. Collision avoidance for such robots often necessitates the use of 1) environmental sensors and tracking systems, 2) limiting the robot’s speed, or 3) range of motion. Still, these techniques may not prevent all collisions while constraining the robot’s performance. For this reason, collaborative robots or “cobots” have been invented. These robots are relatively slow and weak but safe for human interaction, compared to industrial robots that are fast and precise but prone to damage upon collision. Alternatively, the negative effects of collisions could be significantly reduced by using entirely soft robots. These systems often rely on stretchable components as well as

compressible fluids. However, continuous scratching along the surfaces of such materials can lead to damage since soft robots are generally less durable than rigid robots. Moreover, accurate control of fully soft robots [43], tends to be problematic.

One way to find a good tradeoff between these desired robot capabilities is to use tendon drives. Tendon drives allow building robot arms with minimal moving masses by transferring the actuation to the robot base. Such robot designs can be safe, even when operating at increased velocities, due to their inherent backdrivability and low inertia. Paired with powerful actuation, such systems can emit high forces and reach high speeds. By positioning the actuation at an arbitrary location in the robot body, the system can more easily take any desired form to approach anthropomorphism. A major drawback of tendon-driven robots is its challenging control. Repeated high-force transmission wears out the tendon guidance, and high amounts of friction typically add nonlinearities and stochasticity.

In this work, we present PAMY2, a durable and fast 4-DoF tendon-driven robot arm roughly of the size of a human arm. PAMY2 features significantly lower friction than previous designs. We pair PAMY2 with powerful pneumatic artificial muscles (PAMs). PAMs offer the advantage of avoiding stiff joints, which results in less severe peak forces at collisions. At the same time, this type of actuator can generate great forces to achieve either fast motions or lift heavy objects. PAMY2 incorporates new low-friction tendon guidances and ball bearings in the joints to improve the ease of control. To that end, we show that our system is more linear than the tendon-driven arm most similar to ours [7, 8]. Moreover, we ensured through various hardware iterations that our design is resilient. We let the robot run for 25 days uninterrupted while quantifying the repeatability of the system throughout. Additionally, we illustrate that our system produces similar impact forces to the Franka Panda and UR5e at $\sim 4\times$ the speed. To illustrate the ease of control of PAMY2, we learn table tennis smashes with reinforcement learning (RL) as in [9] and double the ball’s speed while simultaneously improving precision. The setting is identical down to the hyperparameters, showing that by just using our robot, performance improved substantially. This result is especially interesting since PAMs are nonlinear actuators that change their dynamics with temperature and exhibit hysteresis effects [41] that often require advanced modeling techniques [8, 36, 37]. Figure 1 puts our new robot, PAMY2, into perspective with the other discussed robot designs in terms of the desired robot capabilities.

To accelerate progress on learning for dynamic tasks, we *open-source the design* as well as the *entire software infrastructure* required to run our system, including a C++ and Python interface. Since our system uses mainly off-the-shelf, commercially available components, 3D-printed parts, and only a small number of custom-machined parts, practitioners are welcome to modify the design and customize our system to their needs. PAMY2 can be equipped with any electrically driven end-effector (such as the custom-fabricated articulated 3D-printed hand on the front page). We also open-source an

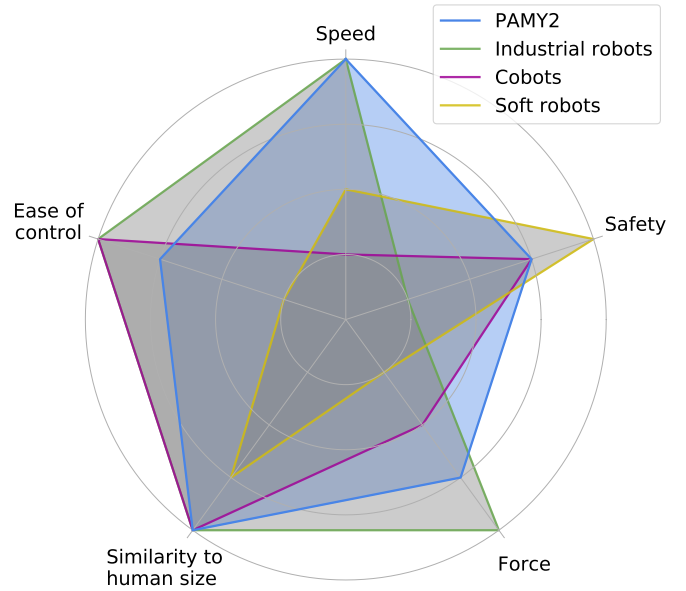


Fig. 1: Visualization of the capabilities of different robot designs. Industrial robots excel in speed, generated forces, and ease of control but are not safe to operate in the proximity of humans. Cobots are easy to control and safe but sacrifice speed and force. Soft robots are generally superior in terms of safety but are hard to control and are often unable to generate high forces. Our robot, PAMY2, is capable of generating high-force and high-velocity trajectories while being significantly safer than most robots. Furthermore, the reduced friction makes our robot easier to control than typical tendon-driven systems.

electrically driven two-DoF wrist that can be combined with, for instance, a racket for table tennis (see Figure 3).

The main contribution of this work is the design of a robot arm (i) that is less prone to damage upon collision due to its lightweight construction, passively compliant actuation, and tendon drives, (ii) with enhanced ease of control, achieved by minimizing nonlinearities, primarily caused by high friction, (iii) that allows for repeatable dynamic motions facilitating the collection of large amounts of data for long-term training, (iv) designed for replicability and adaptability, which allows researchers to build upon and customize our robot for their specific research question.

In Section III, we present the key design decisions that accomplish these objectives. These include a description of the tendon-driven design and the choices that reduce friction in the tendons and joints and increase the system’s robustness. In Section IV, we conduct experiments to demonstrate our system’s effectiveness. We perform long-term dynamic motions to verify robustness and conduct measurements to quantify impact safety. To showcase enhanced ease of control, we demonstrate the increased linearity of our robot. Finally, we apply our system to a challenging dynamic table tennis task, illustrating its capability for rapid yet precise movements.

II. RELATED WORK

a) *Safety through collision avoidance*: Safety in robotics is typically tackled by instrumenting the environment with sensors to detect and track humans and obstacles in the workspace. Sensors employed for this purpose range from distance sensors mounted to the robot [1] to depth cameras [34] and marker-based motion capture systems [24]. If the distance between the robot and a human or obstacle is below a threshold, a safety controller adapts the robot motion to avoid a collision [34] or slows down and stops the motion [24, 32]. These methods generally come with additional costs for the sensors and the need for sensor calibration. Due to the potential for occlusions, they typically require multiple sensors that capture the scene from different angles. Furthermore, collision avoidance strategies tend to be very conservative because they aim to avoid collisions at all cost, resulting in a robot that is heavily constrained in its motions.

Robot safety is particularly important when training policies via RL. During training, the RL agent typically explores random actions. Due to the uncontrolled nature of these exploration strategies, the resulting motions can be dangerous for both the robot and its environment. Safe RL aims to mitigate these safety concerns by discouraging the agent from visiting unsafe states. To that end, these methods modify, e.g., the optimization objective [19, 6, 16, 2], the exploration behavior [13, 17, 4, 10], or the action selected by the policy [11, 31, 38]. Unless provided with additional domain knowledge, safe RL methods need to explore dangerous states at least once during training to learn that these states are unsafe [14]. Domain knowledge, e.g., in the form of a dynamics model or an expert policy, might not always be available, and visiting an unsafe state even once can already cause severe damage to the robot or its environment.

b) *Safety through compliance*: Inherently compliant robots are a viable option to alleviate some safety requirements and avoid the extensive use of sensors for collision avoidance and the dependence on domain knowledge for safe RL. Soft robot components, like passively compliant joints [30] or links [29], can significantly reduce contact forces upon collision. Gealy et al. [15] built a 7-DoF robot arm that achieves passive compliance through backdrivable transmissions. GummiArm [39, 40] and BioRob [26] are notable examples that employ elastic tendons for mechanical compliance. Both systems leverage tendon-driven architectures to enhance compliance and safety.

GummiArm is characterized by its agonist-antagonist actuation system, employing pairs of opposing elastic tendons that control the movement and stiffness of each joint via variable co-contraction levels. Such a design allows the arm to absorb impacts through its joints' natural flexing, reducing the risk during accidental human contact. The elastic materials used for the tendons contribute both to the safety and the bio-fidelity of the arm, allowing it to execute smooth movements.

BioRob focuses on incorporating a series elastic actuation concept within a highly lightweight structure. These

actuators introduce significant compliance at each joint, serving both to mitigate impact forces and to enhance energy efficiency through the storage and release of kinetic energy during tasks. BioRob's design emphasizes minimal moving mass and compact actuator integration.

An alternative is offered by PAMs, which are inherently compliant actuators that can achieve high forces. These actuators are widely used in human-inspired robot arms [42, 5, 21, 20, 18]. Passively compliant robots can also be combined with active collision avoidance strategies, such as in [33]. In such a combination, the robot's compliance enhances safety in the event of undetected or unavoidable collisions.

c) *Generating dynamic motions*: Safety through compliance is essential in dynamic tasks because collisions are more likely at high velocities, and impacts are more severe due to the high momentum of the robot. However, fast motions apply additional strain to the system and, therefore, few robot designs in the literature are capable of generating dynamic motions. Ikemoto et al. [20] evaluate their design on a dynamic throwing task. Mori et al. [28] designed a high-speed robot for a highly dynamic badminton task, which achieves racket speeds of 21 m/s.

BioRob [26] and GummiArm [39] are also both capable of dynamic motions. BioRob arm is able to achieve end effector velocities of 7.4 m/s. Analysis of Figure 10 in [39], showing ballistic movements, suggests GummiArm is slower than BioRob. This compares to PAMY2's end effector velocities of 12 m/s during the table tennis experiments shown in Section IV-C2. GummiArm also has its actuators distributed along the arm, potentially increasing moving mass and impacting safety at high speeds or necessitating the use of lighter, less powerful actuators. Furthermore, neither BioRob nor GummiArm extensively focus on reducing internal friction within their systems. This is particularly relevant when considering their dynamic motion capabilities and overall ease of control. To our knowledge, long-term experimental evaluations, which are critical for assessing durability and reliability, have not been reported for these systems, so it remains unclear whether these motions can be executed robustly over long periods.

III. REALIZATION OF PAMY2

In this section, we present the design of our robot, PAMY2, which substantially improves upon the tendon-driven robot introduced by Büchler et al. [7, 8], referred to as PAMY1 throughout the paper. We detail the improvements to the mechanical design, Bowden tubes, bearings, and pneumatics. Furthermore, we discuss our design choices in light of the goals of impact safety, robustness, and ease of control. The mechanical design of the arm is depicted in Figure 2.

A. Design Choices to Improve Impact Safety

Because collisions cannot always be avoided without limiting the performance of the robot, one of the primary design objectives is to ensure that the impact of such collisions is limited. We achieve this goal by incorporating a tendon-driven design and leveraging passively compliant actuators. As shown



Fig. 2: Design of the tendon-driven robot arm (a). The arm has a rotational and a swivel DoF within the first (e), (f), and second joint (c), (d). It features ball bearings, which are low in friction. Many parts are self-designed and 3D-printed, which are shown colored in black. The four angle encoders are shown with a small green circuit board. The bowden tubes (b) guide the tendons from the muscles to the joints. They feature an inner tube and outer support elements that help maintain constant tendon length.

in Figure 2, the actuators of our robot are not on the robot arm but at the base. Therefore, the moving masses (about 1.3 kg) and inertia are small compared to traditional robot designs where actuators are typically located at the joints.

Active and passive compliance are two distinct approaches to achieving compliance. Active compliance utilizes sensor data and feedback control to adapt and respond to external forces. One example of active compliance are collision reaction schemes. However, the reaction times are often too long to prevent damage. In contrast, passive compliance is achieved through the inherent mechanical properties of the robot, such as elastic joints or soft materials. Our robot achieves passive

compliance through the use of PAM actuators. These actuators are inherently compliant, allowing the robot to absorb and dissipate external forces without the need for complex control schemes.

B. Design Choices to Increase Ease of Control and Extend Durability

Reducing friction is essential for improving the robot’s ease of control, as it reduces uncertainties and nonlinearities in the dynamics. Furthermore, friction leads to wear, which limits the system’s longevity. This section highlights the design choices that help minimize friction in our robot.

1) *Bowden Tubes and Tendons*: Improving the Bowden tubes is key to addressing friction, durability, and maintenance challenges commonly encountered in tendon-driven systems. Our design incorporates continuous Polytetrafluoroethylene (PTFE) Bowden tubes, which have very low friction. These tubes also exhibit resistance to kinking and separation minimizing the risk of tendon entanglement. To further reduce friction within the PTFE tube, we use Ballistol universal oil to lubricate the tendon strings. In addition, our system features a design consisting of an inner tube and custom outer support elements, as depicted in Figure 2b. These outer support elements were specifically engineered for easy 3D printing and manufactured out of Onyx, a carbon-fiber reinforced polyamide filament that withstands exceptionally high forces without breaking. At the same time, this new design fulfills the usual task of a Bowden tube: ensuring a constant tendon length during arm movements by providing external support. Consequently, the movement of one joint influences others only through the rigid body dynamics rather than the tendon drives, improving overall ease of control.

The tendons themselves are made from a Dyneema string that has high strength-to-weight ratio and durability. Dyneema strings are chosen due to their maximum load capacity, and smooth surface, which reduces friction inside the guiding tubes. The tendon strings used in our robot have a diameter of 1.8 mm, but can withstand a load capacity of 500 daN. To improve Dyneema’s creep resistance under high load, we use a pre-tensioned version of the string that is heat-treated to reduce elongation to less than 1%, enhancing its load capacity compared to similar strings of the same diameter and ensuring minimal length change under tension. The lower temperature resistance of Dyneema is mitigated by our design focusing on reducing friction. The effectiveness of this design in reducing heat is shown in experiment IV-B1. For the tendon connections, we utilize a knot-based method, which circumvents the disadvantages of adhesive-based attachments, such as long drying times and potential weakening over time. To address the decreased tear-strength of strings because of the tightening of the knot, our knotting method involves guiding the tendon along a rounded curve before knotting. More details about the parts and materials used can also be found on our project website.

2) *Bearings*: The previous design of B uchler et al. [7, 8] utilizes gliding bearings, which offer the benefit of being highly compact. However, these bearings exhibit considerable friction and stiction. In contrast, PAMY2 includes industry-standard ball bearings at the shoulder and elbow joints to significantly reduce the friction and stiction, while also providing increased off-axis rigidity and improved longevity. The primary tradeoff is a slight increase in the mass of the joints and packaging complexity.

C. Improved Pneumatics

In the pneumatic system of our tendon-driven robot arm, we have implemented several optimizations to enhance performance and reliability. First, we employed optimized tube routing

to improve airflow and reduce pressure losses, particularly by avoiding sharp 90-degree angles in the pneumatic lines. Second, we incorporated a buffer reservoir to stabilize the air pressure supply in front of the valves, ensuring consistent and efficient actuation. Lastly, we designed a ring circuit for the pneumatic system, further improving the air pressure distribution among the valves.

D. Open-Source Hardware and Software

To facilitate further research and development in dynamic robotic tasks, we have made both the hardware and software components of our robot open-source.

1) *Hardware*: Our approach aims at enabling others to build upon our work and adapt the robot for their specific applications. Therefore, the design primarily employs off-the-shelf, commercially available components, 3D-printed parts, and only a limited number of custom-machined parts, thus making it cheaper than many other industrial or research robots. The total material costs amount to approximately €14,185, broken down as follows: Base at €2,652, Arm at €1,448, Electronics at €4,053, and Pneumatics at €6,032. It is important to note, however, that the 3D-printed components require specialized printers, capable of reinforcing parts with continuous fibers.

2) *End-effectors*: PAMY2 can be combined with any electrically driven end-effector, since we feed cables seamlessly through the inside of the arm. In this manner, hands could be mounted, such as the custom-fabricated articulated 3D-printed hand on the front page. We also open-source a two-DoF wrist that features position and torque control and can be combined with, for instance, a racket for table tennis Figure 3.

3) *Software*: We provide an open-source software framework with a versatile API in Python and C++ for controlling and monitoring the robot, based on the o80 framework [3]. The o80 software framework interfaces with the robot’s Programmable Logic Controller (PLC). Communication between the PLC and the o80 software running on the PC is facilitated through UDP, transmitting data such as the robot state (joint angles and velocities, muscle pressures, and valve positions), actions (target pressures or target joint positions, depending on the control mode), and error information.




Fig. 3: Electrically driven wrist for PAMY2 that can be combined with an end-effector

IV. EXPERIMENTS & EVALUATIONS

In this section, we present a series of experiments designed to assess the efficacy of our new robot arm. Our experiments focus on evaluating the following characteristics: Impact




Fig. 4: Experimental setup for the collision force measurements. A Pilz PRMS is mounted to a table onto which the end effector of the robot is colliding.

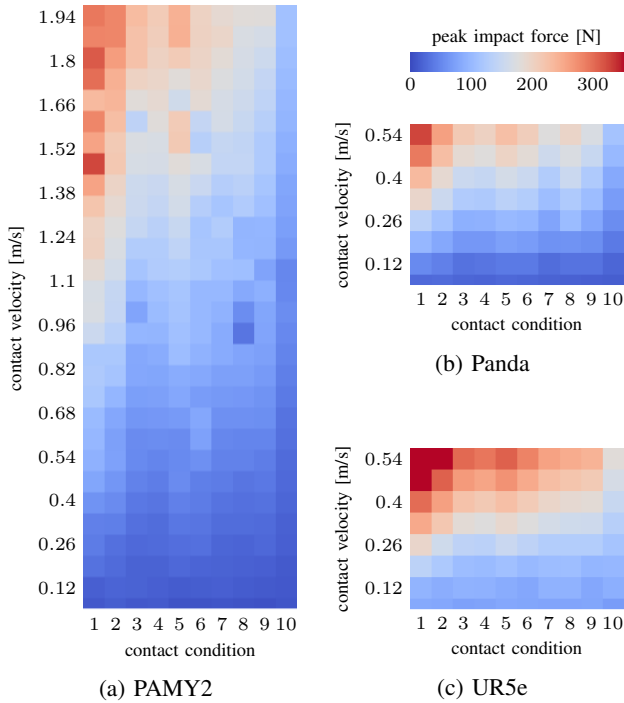


Fig. 5: Collision force map depicting peak impact forces resulting from varying impact velocities and contact scenarios for our robot, alongside the Franka Emika Panda and the Universal Robot UR5e for comparison. Our findings reveal that our robot, when operating at high velocities, generates impact forces akin to those exhibited by the other two robots at considerably lower velocities. We express our gratitude to Kirschner et al. [22] for generously sharing the impact data for the Panda and UR5e robot arms.

safety, robustness, ease of control, and the ability to perform rapid and precise movements.

A. Evaluating Impact Safety

One of the main goals of our robot arm is to ensure superior impact safety compared to traditional motor-driven systems.




Fig. 6: Experiment comparing the friction of our Bowden tubes (top) with those utilized by [7] (bottom). The experimental setup is illustrated in (a): Both types of Bowden tubes are actuated by rapidly switching muscles of the same type. The thermal camera’s heat map (b) and the temperature evolution over time (c) both show a significantly lower temperature increase for our Bowden tubes. As the temperature increase is caused by friction, this finding implies that our Bowden tubes exhibit significantly lower friction.

We achieve this objective primarily through the tendon-driven design, which relocates the heavy actuators to the robot base. Although the use of compliant actuators contributes to the overall safety of the robot, the system’s inertia primarily determines the peak force at impact.

To evaluate the impact safety of our robot, we examine the peak force occurring during potential collisions. For our experiments, we employ the Pilz Robot Measurement System (PRMS) to measure forces during collisions. This device comprises a one-dimensional load cell, a spring, and a rubber cover. Various springs and covers are available to adjust the stiffness and hardness according to different human body parts as specified by ISO/TS 15066:2016 [12]. This technical specification introduces a model of the human body, covering

TABLE I: Definition of contact conditions based on ISO/TS 15066 [12].

No.	Body part	Stiffness	Hardness	Pain Threshold
1	Skull	150 N/mm	70 ShA	130 N
2	Face/hand	75 N/mm	70 ShA	65 N
3	Lower legs	60 N/mm	30 ShA	260 N
4	Thighs	50 N/mm	30 ShA	300 N
5	Neck	50 N/mm	70 ShA	440 N
6	Lower arms	40 N/mm	70 ShA	320 N
7	Back	35 N/mm	30 ShA	420 N
8	Upper arms	30 N/mm	30 ShA	300 N
9	Chest	25 N/mm	70 ShA	280 N
10	Abdomen	10 N/mm	10 ShA	220 N




Fig. 7: Target pressures during an episode of the long-term experiment. All actions are executed open-loop. An episode consists of multisine signals with a frequency of up to 10Hz which are randomly sampled before each episode to explore different areas of the state space. A reset motion sequence is aimed at minimizing the robot’s final position dependence on its previous state. At the end of this reset sequence, the repeatability measurement is taken to assess PAMY2’s consistency and repeatability across the duration of the experiment. Finally, there are movements executed from sets of fixed target pressures at lower and higher speed.

21 body regions. For each body region, it provides contact conditions and a pain tolerance, which we display in Table I.

We compare our measurements to the results obtained in previous studies by Kirschner et al. [22, 23]. To ensure that our measurements accurately reflect the impact safety of our robot, we measure the generated forces at a position close to the robot center. This position ensures that most of the robot’s mass contributes to the peak force experienced during a collision. Figure 4 illustrates the robot’s position during the measurements. To collect the data, we execute trajectories with linear changes in muscle pressure. Upon detecting a sudden change in velocity, indicating a collision, we halt the movement by keeping the target pressure fixed. By modifying the rate of change of the target pressure, we evaluate the impact safety of our robot at various velocities.

Figure 5 shows peak forces during collisions for our tendon-driven arm and for the Franka Emika Panda and Universal Robots UR5e, two conventional motor-driven systems, investigated in prior research. The results clearly show the superior impact safety of our robot, as it achieves similar contact forces while moving at almost four times the speed.

However, it is important to emphasize that although our tendon-driven arm is significantly safer compared to other robotic arms, it could still cause serious injuries if a human is struck with full force at high speeds. Future work should focus


on further enhancing the safety of our system by combining its inherently safer hardware design with algorithmic approaches, additional sensors, or safety mechanisms.

B. Evaluating Robustness


Reinforcement learning enables robots to achieve high performance in complex tasks. However, for these approaches, training for long durations is crucial. We aim to produce a system that lasts and is reliable and, therefore, minimize the main contributor to failure: friction.

1) *Friction Quantification*: Friction plays a significant role in tendon-driven robots as it converts kinetic energy into thermal energy. Because it is widely acknowledged that using ball bearings instead of sliding contact bearings is an effective way to reduce friction, we focus our experiments on evaluating the other major source of friction in our tendon-driven robot: the Bowden tubes.

Higher friction implies that more heat is generated in the Bowden tubes. Therefore, we compare the friction in our Bowden tubes with the original system from [7] by capturing the heat generated during operation with a thermal camera. We set up an experiment where two muscles are connected directly using a Bowden tube, as shown in Figure 6a. The antagonistic muscle pair is contracted in an alternating manner at a frequency of 3.33 Hz for 30 minutes. Figure 6b displays the thermal image captured at the end of the experiment, which



(a) Mean of the final position



(b) Standard deviation of the final position

Fig. 8: Moving mean position relative to the initial mean position $\bar{q}_{\text{moving}} - \bar{q}_0$ and the standard deviation $(\bar{\sigma}_q)_{\text{moving}}$ of the robot’s final position after the open-loop reset motion during the approximately 25 days of the long-term experiment. It uses a moving average window of 400 episodes. A stable mean position (a) alongside a small and slightly decreasing standard deviation (b) over time indicates consistent performance throughout the long-term experiment.

indicates that our Bowden tube remained substantially cooler than the original. A more detailed analysis of the temperature change in the two Bowden tubes can be seen in Figure 6c. The data points in the plot represent the median temperature of the five hottest pixels in the corresponding Bowden tube’s image region. The graph illustrates that the temperature of the original tubes rises more rapidly and ultimately converges to a higher value. These measurements confirm that the design of our Bowden tubes offers significantly lower friction than the original system.

2) *Long-term Dynamic Motions*: Learning complex skills with real robots often necessitates numerous interactions with the environment. To showcase our robot’s capability in collecting the data required for mastering dynamic tasks, we conducted a long-term experiment involving various dynamic movements. This experiment evaluates the system’s reliability and robustness over extended periods.

We designed an array of movement patterns, encompassing random multisine signals, fixed target pressure movements with varying time intervals, and reset motions. The reset motion, executed in an open-loop fashion without feedback from joint position measurements, involves moving to medium pressures, then to minimum pressures, and finally to medium pressures again. This particular reset motion was chosen to make the final position more independent of the preceding

position or motion, thus making differences in this position more indicative of changes in hardware. Figure 7 displays the actuation signal for one episode of the long-term experiment.

During the experiment, the robot operated continuously for approximately 25 days, amassing a comprehensive dataset of diverse robotic motions. The data was recorded at a high sampling rate of 500 Hz. The dataset includes the observed and desired pressure for each muscle, the position, and velocity of each joint, as well as timestamps.

Throughout the long-term experiments, the robot exhibited a high level of robustness and reliability, with no significant signs of wear or damage. This outcome underscores the effectiveness of our design in ensuring the robot’s durability during dynamic tasks.

To quantitatively assess the system’s repeatability, we analyzed the robot’s final position after executing the reset motion multiple times. The small deviation in the final positions, illustrated in Figure 8, indicates that the system maintains consistent performance even after prolonged usage. This repeatability is crucial for successful reinforcement learning, as it ensures the robot can reliably execute learned policies.

C. Evaluating Ease of Control

1) *Increased Linearity*: We employ the method proposed by Ma et al. [27] for system identification in the frequency domain to quantify the nonlinearity in PAMY1 and PAMY2. Each degree of freedom is treated as a single-input and single-output (SISO) system. We design ten different excitation signals with the same frequency spectrum and ten randomly chosen phase spectra, exciting the frequency lines $\Omega = \{0.1\text{Hz}, 0.2\text{Hz}, \dots, 10\text{Hz}\}$.

We excite each degree of freedom individually. Each excitation signal is applied for ten periods continuously, and the response signals of the first two periods are discarded to avoid the effect of transients. Let $U^i(j\omega_k)$ and $Y^i(j\omega_k)$ with $i = 1, \dots, p$ and $\omega_k \in \Omega$ denote the discrete Fourier transformation (DFT) of the input (difference in target pressure for antagonistic muscle pairs) and output signals (joint angles) of the i -th period, respectively, where p represents the total number of periods after discarding (here $p = 8$), and $j = \sqrt{-1}$ denotes the imaginary number. First of all, we calculate the average value of the input and output signals in the frequency domain over all the different periods.

$$\hat{Y}(j\omega_k) = \frac{1}{p} \sum_{i=1}^p Y^i(j\omega_k) \quad (1)$$

$$\hat{U}(j\omega_k) = \frac{1}{p} \sum_{i=1}^p U^i(j\omega_k) \quad (2)$$

The average frequency response function (FRF) \hat{G} is given by

$$\hat{G}(j\omega_k) = \frac{\hat{Y}(j\omega_k)}{\hat{U}(j\omega_k)}. \quad (3)$$

Since the identified system is nonlinear, the discrepancy between the measured average FRFs that arises when having

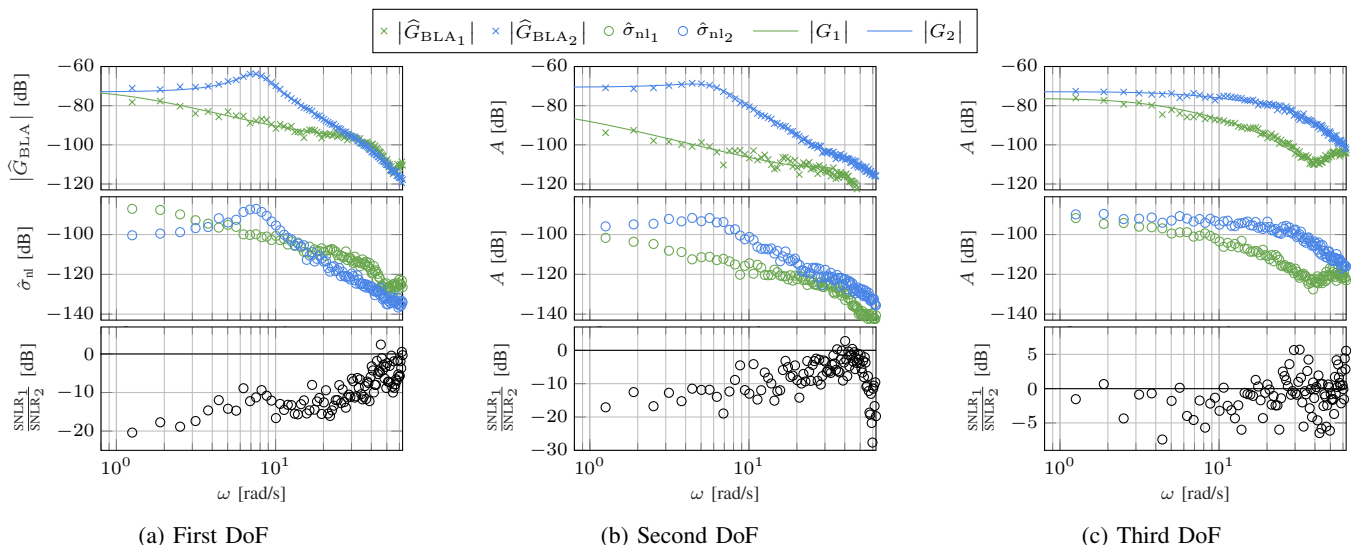


Fig. 9: Comparison of the linearity between PAMY1 (index 1) and PAMY2 (index 2). The new system demonstrates increased amplitude (A) and a higher bandwidth across all DoFs (first row). The second row shows the absolute level of nonlinearity of the two systems. To make this nonlinearity between the systems comparable, we examine the SNLR (third row), which is significantly lower for the first two DoFs at most frequencies, and marginally lower for the third DoF at lower frequencies. These findings underscore the enhanced tracking performance and reduced nonlinearity of the new system.

excitation signals with different phase spectra provides a means to characterize the nonlinearities. If the system were linear, then applying excitation signals with different phase realizations would not affect the average FRF. First, we calculate the corresponding average FRF $\hat{G}^l, l = 1, \dots, m$ for each input signal according to Equation (3), where the superscript l refers to the different excitation signals (here $m = 10$). Then, the average FRF over all excitation signals is given by

$$\hat{G}_{\text{BLA}}(j\omega_k) = \frac{1}{m} \sum_{l=1}^m \hat{G}^l(j\omega_k), \quad (4)$$

where the subscript BLA refers to “Best Linear Approximation”. Lastly, an estimate of the system’s nonlinearity is given by

$$\hat{\sigma}_{\text{nl}}^2(j\omega_k) = \frac{1}{m(m-1)} \sum_{l=1}^m \left| \hat{G}^l(j\omega_k) - \hat{G}_{\text{BLA}}(j\omega_k) \right|^2, \quad (5)$$

where the subscript nl refers to “nonlinearity”.

We note that the absolute value of $\hat{\sigma}_{\text{nl}}^2$ is not a good indicator since it would be affected by a simple re-scaling of the output variable. To better compare the degree of nonlinearity of the two systems, we define the signal-to-nonlinearity ratio (SNLR) as

$$\text{SNLR}(j\omega_k) = \left(\frac{|\hat{G}_{\text{BLA}}(j\omega_k)|}{\hat{\sigma}_{\text{nl}}} \right)^2. \quad (6)$$

Figure 9 displays the system identification results. We observe that the new system exhibits higher amplitudes compared to the old system, indicating improved tracking of dynamic

motions. At the same time, each system’s degree of nonlinearity and amplitude show the same trend. The third row shows the ratio of the SNLR for the two systems. We notice that in a large portion of the frequency spectrum, the SNLR value of the old system is significantly smaller than that of the new system, especially for the first and second DoF. When adjusted for amplitude, this result shows that the new system has significantly lower nonlinearity than the old system. This indicates that the new system is easier to control, which we will demonstrate in the next experiment.

2) *Learning a Dynamic Task*: To demonstrate our system’s capabilities in a highly dynamic task, we repeat the table tennis smashing experiment from [9]. This task employs a reward function, detailed in [9], that includes both the ball’s speed and accuracy concerning a target location on the other side of the table. It is a good demonstration of the robot’s potential, as it requires rapid and precise movements, as well as high forces to accelerate the ball to high velocities. Similar to Büchler et al. [9], we train the robot in a Hybrid-Sim-and-Real setup (HySR), with the real robot and a simulated ball as shown in Figure 10. After training, the robot can return real balls with high speeds, albeit training only with simulated balls.

We use a stochastic policy, with actions being changes in target pressures, and use Proximal Policy Optimization (PPO) [35] as the backbone RL algorithm. Although there were changes in the software, we aimed to keep the setup as similar as possible to [9] to allow for a direct performance comparison. It is worth noting that the learning hyperparameters were optimized for the old system, and we decided not to adapt or optimize them further for our new system to




Fig. 10: Table tennis smashing experiment with PAMY2. In the HySR setup, we learn with a real robot and a simulated ball. The robot’s initial position at the beginning of an episode is shown in (a). During the training, the robot learns a motion in which it first draws back to generate momentum before striking the ball (b). The racket reaches speeds of up to 12 m/s during this motion.

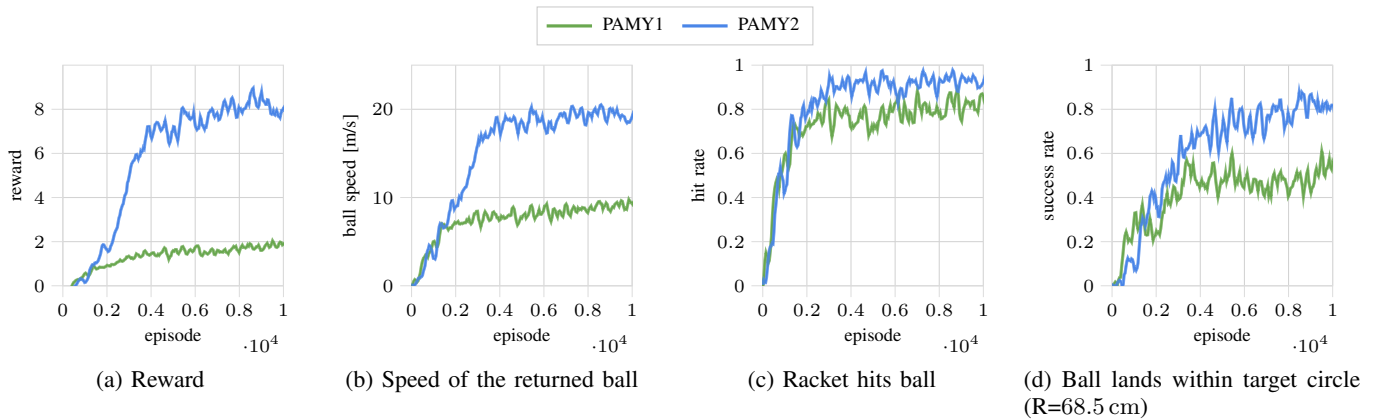


Fig. 11: Results of the table tennis smashing experiment. Compared to the design of Büchler et al. [7, 8], PAMY2 reaches a significantly higher final reward (a). Analyzing this result in detail, we find that it returns balls with higher average speed (b), but at the same time also hits the ball more often (c), and more often the ball lands close to the target position (d). This experiment demonstrates the benefits of the new design for motions that are highly dynamic but also precise.

ensure that any performance gains are caused by hardware improvements, not by hyperparameter changes.

As a result of the increased activation bandwidth and also due to the robot discovering an effective strategy of moving along the joint limit of the second DoF before striking the ball, we found that the learning algorithm frequently pushed the robot towards its joint limits, which is detrimental to the robot’s longevity. To address this issue and minimize wear and tear on the robot, we introduced a minor term to the reward function that discourages reaching the joint limits. Crucially, aside from this adjustment, no further safety measures are necessary, underscoring the inherent robustness of the PAMY2 design.

Figure 11 shows that the new design, PAMY2, achieves significantly higher ball speeds than the design of Büchler et al. [7, 8]. Despite the higher ball speeds, it is also more precise in terms of more frequent ball contacts and lower distance to the ball’s target location. When evaluating the trained policy with real balls instead of simulated balls, we

achieve similar ball speeds of 20 m/s. To our knowledge, this is the fastest robot table tennis play to date and comparable to professional human players [25]. Overall, PAMY2 reaches performance far superior to PAMY1, demonstrating the benefits of the improved design for highly dynamic and precise motions.

V. DISCUSSION AND CONCLUSION

In this paper, we have presented a novel 4-DoF tendon-driven robot arm actuated by PAMs. Our design focuses on reducing friction, passive compliance, and inherent impact safety, allowing the robot to operate efficiently and safely during dynamic tasks. Through various experiments, we have demonstrated the effectiveness of our robotic arm in terms of these design goals.

Our work contributes to the growing field of soft robotics, which aims to create more adaptable and safer robots, particularly for human-robot interaction scenarios. Although our robot arm showcases several advantages over traditional

motor-driven systems, there are limitations to our design. Despite the improvements in terms of ease of control, PAM-driven systems still face challenges in achieving the repeatability and precision offered by their motor-driven counterparts. Identifying the optimal set of tasks for our robot arm, where the benefits of safety and dynamic performance outweigh the limitations, is an important avenue for future research. Recent advances in machine learning for robotics hold great opportunities for enhancing the capabilities of robots like ours. In the table tennis task, we showed that with our improved design, and by leveraging data-driven approaches, it is possible to develop advanced control strategies that address the inherent challenges of PAM-driven systems.

Our work represents a step forward in the development of robotic systems that can achieve high performance while maintaining safety in shared human environments. By making our design and resources open-source, we hope to inspire future research efforts that build upon and refine our work, fostering a new generation of collaborative and versatile robots.

REFERENCES

- [1] Giovanni Buizza Avanzini, Nicola Maria Ceriani, Andrea Maria Zanchettin, Paolo Rocco, and Luca Bascetta. Safety control of industrial robots based on a distributed distance sensor. *Transactions on Control Systems Technology (TCST)*, 22(6):2127–2140, 2014.
- [2] Arnab Basu, Tirthankar Bhattacharyya, and Vivek S Borkar. A learning algorithm for risk-sensitive cost. *Mathematics of operations research (MOR)*, 33(4):880–898, 2008.
- [3] Vincent Berenz, Maximilien Naveau, Felix Widmaier, Manuel Wüthrich, Jean-Claude Passy, Simon Guist, and Dieter Büchler. The o80 c++ templated toolbox: Designing customized python apis for synchronizing realtime processes. *Journal of Open Source Software (JOSS)*, 6(66):2752, 2021.
- [4] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [5] Ivo Boblan and Andreas Schulz. A humanoid muscle robot torso with biologically inspired construction. In *International Symposium on Robotics (ISR) and German Conference on Robotics (ROBOTIK)*, 2010.
- [6] Vivek S Borkar. Q-learning for risk-sensitive control. *Mathematics of operations research (MOR)*, 27(2):294–311, 2002.
- [7] Dieter Büchler, Heiko Ott, and Jan Peters. A lightweight robotic arm with pneumatic muscles for robot learning. In *International Conference on Robotics and Automation (ICRA)*, 2016.
- [8] Dieter Büchler, Roberto Calandra, Bernhard Schölkopf, and Jan Peters. Control of musculoskeletal systems using learned dynamics models. *Robotics and Automation Letters (RA-L)*, 3(4):3161–3168, 2018.
- [9] Dieter Büchler, Simon Guist, Roberto Calandra, Vincent Berenz, Bernhard Schölkopf, and Jan Peters. Learning to play table tennis from scratch using muscular robots. *Transactions on Robotics (T-RO)*, 38(6):3850–3860, 2022.
- [10] Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A lyapunov-based approach to safe reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [11] Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757*, 2018.
- [12] International Organization for Standardization. *ISO-TS 15066: Robots and Robotic Devices: Collaborative Robots*. Geneva, CH, 2016.
- [13] Javier Garcia and Fernando Fernández. Safe exploration of state and action spaces in reinforcement learning. *Journal of Artificial Intelligence Research (JAIR)*, 45: 515–564, 2012.
- [14] Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research (JMLR)*, 16(1):1437–1480, 2015.
- [15] David V Gealy, Stephen McKinley, Brent Yi, Philipp Wu, Phillip R Downey, Greg Balke, Allan Zhao, Menglong Guo, Rachel Thomasson, Anthony Sinclair, et al. Quasi-direct drive for low-cost compliant robotic manipulation. In *International Conference on Robotics and Automation (ICRA)*, 2019.
- [16] Peter Geibel and Fritz Wysotzki. Risk-sensitive reinforcement learning applied to control under constraints. *Journal of Artificial Intelligence Research (JAIR)*, 24:81–108, 2005.
- [17] Alborz Geramifard, Joshua Redding, and Jonathan P How. Intelligent cooperative control architecture: a framework for performance improvement using safe learning. *Journal of Intelligent & Robotic Systems (JINT)*, 72:83–103, 2013.
- [18] Daoxiong Gong, Rui He, Yu Wang, and Jianjun Yu. Bionic design of a 7-dof human-arm-like manipulator actuated by antagonized pneumatic artificial muscles. In *International Conference on CYBER Technology in Automation, Control, and Intelligent Systems (CYBER)*, 2019.
- [19] Matthias Heger. Consideration of risk in reinforcement learning. In *Machine Learning Proceedings*, pages 105–111. Elsevier, 1994.
- [20] Shuhei Ikemoto, Fumiya Kannou, and Koh Hosoda. Humanlike shoulder complex for musculoskeletal robot arms. In *International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [21] Shuhei Ikemoto, Yoichi Nishigori, and Koh Hosoda. Direct teaching method for musculoskeletal robots driven by pneumatic artificial muscles. In *International Conference on Robotics and Automation (ICRA)*, 2012.

- [22] Robin Jeanne Kirschner, Alexander Kurdas, Kübra Karacan, Philipp Junge, Seyed Ali Baradaran Birjandi, Nico Mansfeld, Saeed Abdolshah, and Sami Haddadin. Towards a reference framework for tactile robot performance and safety benchmarking. In *International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [23] Robin Jeanne Kirschner, Nico Mansfeld, Saeed Abdolshah, and Sami Haddadin. Experimental analysis of impact forces in constrained collisions according to iso/ts 15066. In *International Conference on Intelligence and Safety for Robotics (ISR)*, 2021.
- [24] Przemyslaw A Lasota, Gregory F Rossano, and Julie A Shah. Toward safe close-proximity human-robot interaction with standard industrial robots. In *International conference on Automation science and engineering (CASE)*, 2014.
- [25] Marcus JC Lee, Hiroki Ozaki, Wan Xiu Goh, et al. Speed and spin differences between the old celluloid versus new plastic table tennis balls and the effect on the kinematic responses of elite versus sub-elite players. *International Journal of Racket Sports Science*, 1(1), 2019.
- [26] Thomas Lens, Jürgen Kunz, Oskar Von Stryk, Christian Trommer, and Andreas Karguth. Biorob-arm: A quickly deployable and intrinsically safe, light-weight robot arm for service robotics applications. In *International Symposium on Robotics (ISR) and German Conference on Robotics (ROBOTIK)*, 2010.
- [27] Hao Ma, Dieter Büchler, Bernhard Schölkopf, and Michael Muehlebach. A learning-based iterative control framework for controlling a robot arm with pneumatic artificial muscles. In *Robotics: Science and Systems (RSS)*, 2022.
- [28] Shotaro Mori, Kazutoshi Tanaka, Satoshi Nishikawa, Ryuma Niiyama, and Yasuo Kuniyoshi. High-speed and lightweight humanoid robot arm for a skillful badminton robot. *Robotics and Automation Letters (RA-L)*, 3(3): 1727–1734, 2018.
- [29] Jung-Jun Park, Byeong-Sang Kim, Jae-Bok Song, and Hong-Seok Kim. Safe link mechanism based on nonlinear stiffness for collision safety. *Mechanism and Machine Theory*, 43(10):1332–1348, 2008.
- [30] Jung-Jun Park, Jae-Bok Song, and Hong-Seok Kim. Safe joint mechanism based on passive compliance for collision safety. In *Recent Progress in Robotics: Viable Robotic Service to Human: An Edition of the Selected Papers from the 13th International Conference on Advanced Robotics*, 2008.
- [31] Tu-Hoa Pham, Giovanni De Magistris, and Ryuki Tachibana. Optlayer-practical constrained optimization for deep reinforcement learning in the real world. In *International Conference on Robotics and Automation (ICRA)*, 2018.
- [32] Paul Rybski, Peter Anderson-Sprecher, Daniel Huber, Chris Niessl, and Reid Simmons. Sensor fusion for human safety in industrial workcells. In *International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [33] Riccardo Schiavi, Antonio Bicchi, and Fabrizio Flacco. Integration of active and passive compliance control for safe human-robot coexistence. In *International Conference on Robotics and Automation (ICRA)*, 2009.
- [34] Bernard Schmidt and Lihui Wang. Depth camera based collision avoidance via active robot control. *Journal of Manufacturing Systems*, 33(4):711–718, 2014.
- [35] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [36] Vaisakh Shaj, Philipp Becker, Dieter Büchler, Harit Pandya, Niels van Duijkeren, C James Taylor, Marc Hanheide, and Gerhard Neumann. Action-conditional recurrent kalman networks for forward and inverse dynamics learning. In *Conference on Robot Learning (CoRL)*, 2021.
- [37] Vaisakh Shaj, Dieter Buchler, Rohit Sonker, Philipp Becker, and Gerhard Neumann. Hidden parameter recurrent state space models for changing dynamics scenarios. *arXiv preprint arXiv:2206.14697*, 2022.
- [38] Yifei Simon Shao, Chao Chen, Shreyas Kousik, and Ram Vasudevan. Reachability-based trajectory safeguard (rts): A safe and fast reinforcement learning safety layer for continuous control. *Robotics and Automation Letters (RA-L)*, 6(2):3663–3670, 2021.
- [39] Martin F Stoelen, Fabio Bonsignorio, and Angelo Cangelosi. Co-exploring actuator antagonism and bio-inspired control in a printable robot arm. In *From Animals to Animats 14: 14th International Conference on Simulation of Adaptive Behavior (SAB)*, 2016.
- [40] Martin F Stoelen, Ricardo de Azambuja, Beatriz López Rodríguez, Fabio Bonsignorio, and Angelo Cangelosi. The gummiarm project: A replicable and variable-stiffness robot arm for experiments on embodied ai. *Frontiers in Neurorobotics*, 16:836772, 2022.
- [41] Bertrand Tondu. Modelling of the mckibben artificial muscle: A review. *Journal of Intelligent Material Systems and Structures*, 23(3):225–253, 2012.
- [42] Bertrand Tondu, Serge Ippolito, Jérémie Guiochet, and Alain Daidie. A seven-degrees-of-freedom robot-arm driven by pneumatic artificial muscles for humanoid robots. *The International Journal of Robotics Research (IJRR)*, 24(4):257–274, 2005.
- [43] Michael Wehner, Ryan L Truby, Daniel J Fitzgerald, Bobak Mosadegh, George M Whitesides, Jennifer A Lewis, and Robert J Wood. An integrated design and fabrication strategy for entirely soft, autonomous robots. *Nature*, 536(7617):451–455, 2016.