



Kendall Koe^{1,2}

Computer Science,
University of Illinois,
1210 Siebel Center, 201 N. Goodwin,
Urbana, IL 61801
e-mail: kfkoe2@illinois.edu

Samhita Marri²

Electrical and Computer Engineering,
University of Illinois,
306 N Wright St,
Urbana, IL 61801
e-mail: marri2@illinois.edu

Benjamin Walt

Mechanical Science and Engineering,
University of Illinois,
1206 W Green St,
Urbana, IL 61801
e-mail: walt@illinois.edu

Shivani Kamtikar

Computer Science,
University of Illinois,
1210 Siebel Center, 201 N. Goodwin,
Urbana, IL 61801
e-mail: skk7@illinois.edu

Naveen Kumar Uppalapati

National Center for Supercomputing Applications,
University of Illinois,
1205 W Clark St,
Urbana, IL 61801
e-mail: uppalap2@illinois.edu

Girish Krishnan

Industrial and Systems Engineering,
University of Illinois,
104 S Mathews Ave,
Urbana, IL 61801
e-mail: gkrishna@illinois.edu

Girish Chowdhary

Computer Science,
University of Illinois,
1210 Siebel Center, 201 N. Goodwin,
Urbana, IL 61801
e-mail: girishc@illinois.edu

Learning-Based Position and Orientation Control of a Hybrid Rigid-Soft Arm Manipulator

We present a position and orientation controller for a hybrid rigid-soft manipulator arm where the soft arm is extruded from a two degrees-of-freedom rigid link. Our approach involves learning the dynamics of the hybrid arm operating at 4 Hz and leveraging it to generate optimal trajectories that serve as expert data to learn a control policy. We performed an extensive evaluation of the policy on a physical hybrid arm capable of jointly controlling rigid and soft actuation. We show that with a single policy, the arm is capable of reaching arbitrary poses in the workspace with 3.73 cm (<6% overall arm length) and 17.78 deg error within 12.5 s, operating at different control frequencies, and controlling the end effector with different loads. Our results showcase significant improvements in control speed while effectively controlling both the position and orientation of the end effector compared to previous quasistatic controllers for hybrid arms. [DOI: 10.1115/1.4067872]

Keywords: kinematics, dynamics, and control of mechanical systems

¹Corresponding author.

²Kendall Koe and Samhita Marri contributed equally to this work.

Contributed by the Mechanisms and Robotics Committee of ASME for publication in the JOURNAL OF MECHANISMS AND ROBOTICS. Manuscript received February 6, 2024; final manuscript received January 22, 2025; published online February 27, 2025. Assoc. Editor: Vishesh Vikas.

1 Introduction

Robotic manipulation requires high dexterity, adaptability, and safe handling of objects to enable intricate manipulations in agroforestry, telehealth, and other applications in constrained settings [2,3]. Soft continuum arms (SCAs) are inspired by octopus

Table 1 Comparison of sample efficiency with prior methods

Reference	Data points sampled	Actuation space dimension	Length of the manipulator	Workspace	Soft manipulator base is stationary
Bruder et al. [30]	45,103	Three bending actuators	34 cm	Controls laser points on 2D grid	Yes
Bruder et al. [31]	325,733	Three bending actuators	70 cm	Controls 3D position with unknown loads	Yes
Centurelli et al. [22]	10,000	Three bending actuators	44 cm	Controls 3D position	Yes
Ours	41,920	Two servo motors, one stepper motor, one bending actuator, and two rotating actuators	70 cm	Controls 3D position and orientation	No (see θ_1 and θ_2 in Fig. 2)

Note: Bruder et al. [30] have a similar number of data points, but a smaller actuation space, shorter manipulator, smaller workspace, and stationary base making it less sample efficient. Bruder et al. [31] use an order of magnitude more data, but have a smaller actuation space, similar length manipulator, smaller workspace, and a stationary base making it again less sample efficient. No conclusion about sample efficiency can be drawn when comparing to Ref. [22] because it uses less data, but has a smaller actuation space, shorter manipulator, smaller workspace, and stationary manipulator.

tentacles and elephant trunks and do not have rigid joints and actuators [4–8]. These manipulators are predisposed to work well in such cases due to their adaptability [9–12], but they do not lend to precise operations, especially with large payloads. Moreover, the workspace and steerability of these manipulators are strongly affected by gravity [13–16]. Also, controlling both end effector position and orientation is difficult, given the larger workspace and the under-actuated nature of soft manipulators [17].

On the other hand, hybrid rigid-soft manipulators can be used as a conventional rigid arm when precision and load bearing are deemed important, while a distal deployable soft arm can bring in additional dexterity and adaptability. One such platform is the variable length nested soft (VaLeNS) [18] configuration, where an SCA is stored in a hollow robotic link with the ability to extrude out of the link as needed. In the VaLeNS configuration, the rigid link positions itself appropriately for the soft arm to extrude and reach the target position with the required end orientation. Its effectiveness has been demonstrated in agricultural berry harvesting [1,2] and healthcare applications [19] where the soft arm is controlled in a quasistatic manner. While soft arms provide various advantages, a hybrid system offers a broader design space albeit at the cost of increased system complexity. Thus, for practical real-world applications, a fast position and orientation controller must be developed. For instance, in the harvesting scenario, without orientation control, reaching the fruit position does not guarantee that the gripper is facing the fruit. Additionally, if not done in a timely manner, the economic viability of the system degrades. However, due to the rotating base of the hybrid arm, the soft arm is no longer stationary compared to the vertically downward suspended soft arms [20–22]. When coupled with the varying length of the arm and the effects of gravity, more complex shapes and trajectories can be explored that are otherwise unattainable with a purely soft system. Thus, controlling a hybrid rigid-soft manipulator presents more challenges than controlling a purely soft manipulator.

Therefore, we (i) develop a control policy where there is no waiting for the arm to settle unlike in quasistatic controllers; (ii) control our hybrid manipulator in the VaLeNS configuration, which is challenging due to its dexterity and increased range of motion; (iii) control not only the position but also the orientation of the end effector, which is a more difficult learning problem but has more practical applications in any downstream task that requires grasping; (iv) accommodate zero-velocity of the end effector to support downstream manipulation tasks, unlike works like Ref. [20] that have dynamic reaching, and (v) extensively evaluate our physical platform.

2 Related Work

For hybrid arms, a leading approach is to use quasistatic controllers. Dynamic effects are present when operating pneumatic actuators at high speeds. Thus, quasistatic controllers often actuate at low

speeds to dissipate these dynamic effects of the system. In this regime, numerical methods like piece-wise constant curvature and Cosserat rod theory can be used to model and control the dynamics of a soft arm [23–25]. Shi et al. [25] use a Cosserat rod model to control the position and orientation in a quasistatic manner. Using these analytical models or a learned model, Gan et al. [21], Satheshebabu et al. [13,26], Centurelli et al. [22], Wu et al. [27], and Uppalapati et al. [1] learn control policies using a variety of reinforcement learning (RL) methods.

The alternative is dynamic controllers, designed to dampen oscillations, and consider inertial effects and other dynamic effects, which can be incorporated either explicitly or learned through dynamic data collection, to improve control at higher frequencies. In many fields, dynamic controllers have high operating frequencies. However, in Ref. [28], Jacobian-based dynamic pose control is performed at 1–2 Hz, but their system is actuated using magnetic fields. Haggerty et al. [29] also use a magnetically actuated SCA to demonstrate a recurrent neural network (RNN)-based dynamic model to improve upon the real-time prediction of Cosserat rod theory, but this method tends to preserve the structure of the data without considering the underlying physics leading to overfitting and a lack of generalization. Adding a history of previous states, to better model physical phenomena like hysteresis, could improve these results. Using motion capture techniques for state estimation like many previous works [20,22,29–35] is also infeasible for outdoor scenarios due to the unstructured environment and potential occlusions. In contrast, our approach employs a magnetic sensor to track the arm state in real-time offering a more feasible solution for outdoor applications.

Li et al. [36] use an adaptive Kalman filter to enable an optimal controller of a two-section pneumatic arm. However, Kalman filter-based approaches are limited by linear models and Gaussian noise. A popular approach for modeling continuum arms is Cosserat rod theory, but it is computationally expensive, requires parameter estimation, and is restricted to simulation studies [37–39]. Finite element methods [32] have also been studied for dynamic position control, but struggle with high computational loads slowing real-time deployment. Several Koopman-operator theory-based methods [29–31] seek to overcome this limitation by constructing a linear model in a higher dimensional space, but require large amounts of data leading to wear on the robot, especially when the size of state space increases. Compared with these methods, ours is more sample efficient. Bruder et al. [30] use a similar number of data points as our method to model and control the projected position of a vertically suspended manipulator with three actuators. To extend their work, Bruder et al. [31] control a longer arm with variable loading, but the number of samples grows by an order of magnitude. Our method controls the position and orientation of a manipulator of comparable length to that in Ref. [31] with six actuators with the same magnitude of samples as in Ref. [30]. Our system additionally has a mobile soft arm base and a hybrid bending/rotational actuator combination as opposed to purely

bending actuators. Thus, our method finds a mapping between more complex actuation and task spaces and is more sample efficient. A comparable method on a smaller workspace and an actuation space can be seen in Ref. [22]. These comparisons are summarized in Table 1. Sample efficient models combined with RL have been explored [1,21,22,26,27]. Methods for overcoming the simulation-to-reality gap in RL have also been explored [40–42]. Although these methods are very precise, they require a new policy for each learned trajectory, limiting the scalability of their application. Toward overcoming model degradation issues in order to adapt as the system wears, Piqué et al. [43] explore continual learning methods in the soft robot control field under continuously changing loads. However, these approaches involve trade-offs, particularly between avoiding catastrophic forgetting and managing computational and memory overheads.

Many works [31,44–49] use a model predictive control (MPC) framework for the dynamic control of soft systems but have several limitations. Approaches that model the system as linear [45,46,49] offer real-time control but fail to capture the complex nonlinearities. Preiss et al. [44] incorporate nonlinearities using an RNN-based model followed by an extended Kalman filter into the optimization framework but it is limited by the filter’s assumptions of noise and modeling. Tang et al. [48] model the dynamics using the Gaussian process, incorporating uncertainty estimation and propagation, but can be limited computationally as the training data increase. Within the same MPC framework, Jensen et al. [47] explore using a genetic evolutionary optimization framework with integral control but can be slow due to the number of iterations to converge particularly as the number of states increases.

In this work, we present a closed-loop controller that controls both the position and line-of-sight orientation of a hybrid rigid-soft continuum manipulator’s end effector aimed at real-time deployment in applications where low computational power is available and medium to high operating speeds are necessary. Similar to Ref. [20], the forward model is first learned from dynamic data to model our hybrid arm. Zero-velocity tip data are incorporated

to enable static reaching and to better support downstream manipulation tasks. Second, the forward model is used to generate optimal trajectories offline to serve as expert data. Third, the generated trajectories are run on the physical robot to avoid the simulation-to-reality gap, and a control policy is learned offline that can be leveraged in a real-time closed-loop controller as shown in Fig. 1. This policy is evaluated and analyzed on the physical platform showing a single control policy reaching arbitrary poses in the workspace, being robust to external loads, following trajectories, and operating at different control frequencies. Note that soft arms in previous works are vertically suspended with a fixed orientation of the base and solely have bending actuations [20–22,26–31,33,36,40]. In contrast, our system varies the orientation of the soft arm base, the soft arm varies in length, and a torsional actuator is considered. As a result, the workspace is not completely symmetric and the effect of gravity is quite severe on the soft portion presenting an additional challenge. Because of the little compute needed at inference time, this policy is ideal in settings like agricultural environments where computing power is limited.

3 Proposed Approach

3.1 Robotic Platform. The robotic platform used is an extrudable BR² SCA nested within a 3D-printed rigid arm, also called the VaLeNS arm [18] where the rigid portion can rotate in two degrees-of-freedom as seen in Fig. 2(d). Custom camera and gripper combinations can be mounted on the end of this hybrid arm to support various manipulation tasks. To measure the end effector pose, a Polhemus Patriot Motion Tracker³ is used. It consists of a pea-sized magnetic sensor mounted at the end effector whose pose is measured relative to a magnetic source placed at the base of the arm as seen in Fig. 2(a). The soft arm is made from a combination of one bending and two rotational fiber-reinforced elastomeric enclosures [14,50], hence the name BR² [18]. These are pneumatically pressurized tubes wrapped with an inextensible fiber woven in a helical shape. The angle of the wrapping, the material of the fiber, and the number of fibers affect the behavior of the BR² when pressurized [14]. The stepper motor extrudes the soft arm by length, L , using a lead screw enclosed within the rigid arm. The rigid arm can adjust its pitch θ_1 and roll θ_2 .

Each actuation (Fig. 2(d)) is limited to a set range to support a sufficiently large workspace while ensuring operational safety. The pitch, θ_1 , is limited to $[-45 \text{ deg}, 45 \text{ deg}]$ with 0 being the level with the ground. The roll, θ_2 , is limited to $[-90 \text{ deg}, 90 \text{ deg}]$ with 0 being the home position. The extrusion, L , is limited to $[10, 16] \text{ cm}$. The bending pressure, B , is limited to $[48.26, 241.32] \text{ kPa}$, and the rotating pressures, R_1 and R_2 , are limited to $[0, 241.32] \text{ kPa}$. Both rotational actuators are pressurized to a base pressure of 241.32 kPa and then differentially pressurized to rotate in the desired direction. The differential pressurization serves to minimize oscillations, thereby enhancing the system’s load-bearing ability. A single actuation variable can be used to control the pressures in both R_1 and R_2 with differential actuation.

3.2 Control Algorithm. This section elaborates on Fig. 3, which describes a pipeline for developing a fast, real-time position and orientation controller. Controlling the orientation is especially important for vision-based tasks where target objects must be in the view of a tip camera as seen in Fig. 4. Therefore, the line of sight vector tangent with respect to the arm length, as shown in Fig. 2(b), was selected to represent the arm’s orientation. This decreases the number of parameters to control and allows freedom of rotation axially along the soft arm as long as the line of sight is maintained. The control policy is developed in three steps offline. First, data are collected on the robot and a forward dynamic model, f , is learned (see Fig. 3(a)). Next, optimal

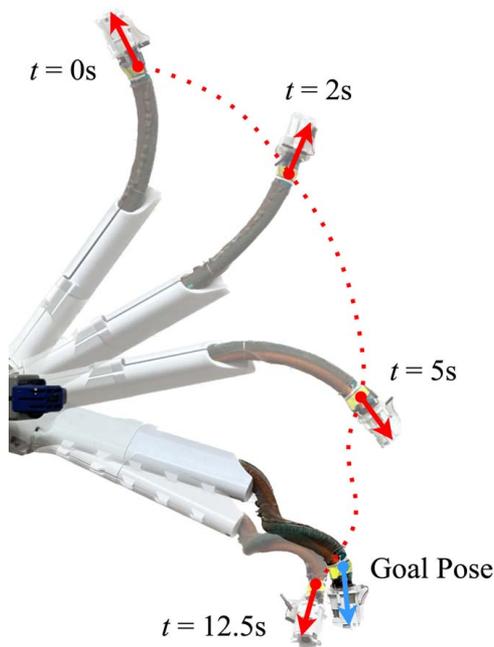


Fig. 1 We developed a non-quasistatic pose controller for a hybrid continuum arm. Our approach enables quick and accurate pose control anywhere in the workspace. The manipulator’s inherent dexterity and adaptability make it promising for practical applications [1,2]. This figure shows a sample evaluation of our controller where it reaches the desired pose within 12.5 s at a control frequency of 4 Hz.

³<https://polhemus.com/motion-tracking/all-trackers/patriot/>

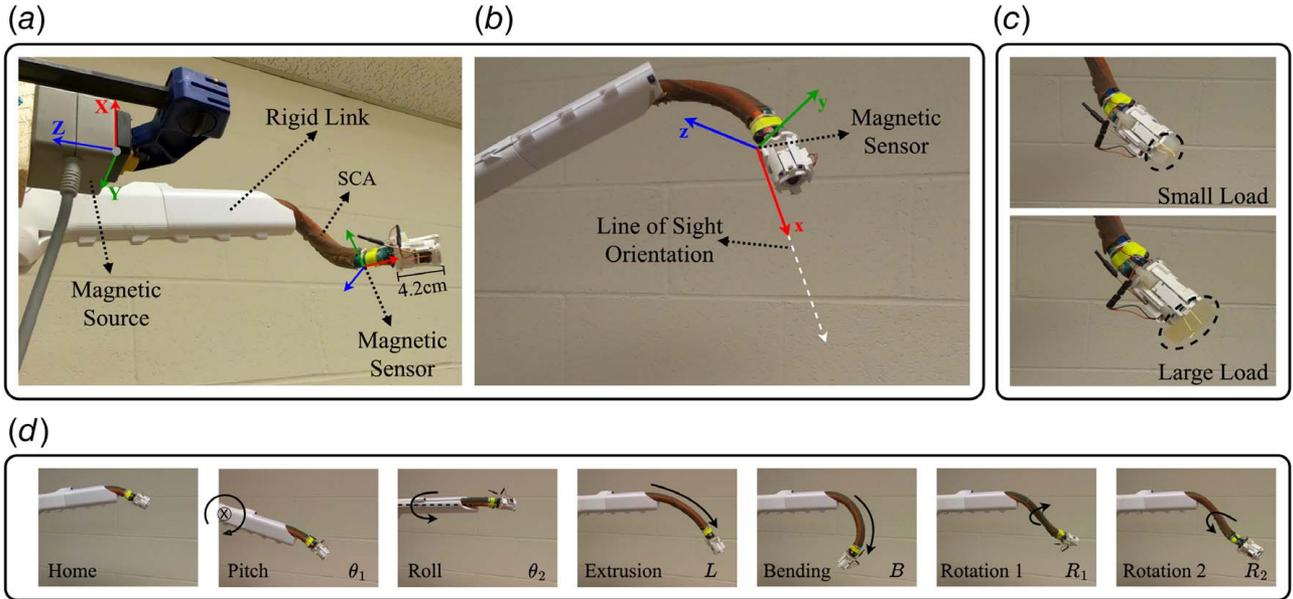


Fig. 2 (a) Overall hybrid manipulator system. The magnetic sensor measures the pose of the end effector relative to the magnetic source. (b) The controller enables control of the end effector's position and line of sight orientation. (c) A small load (3.8 g, 10% of the SCA's weight) and a large load (10.8 g, 30% of the SCA's weight) are added at the tip to test the controller's robustness. (d) Each individual actuation of the hybrid arm relative to the home position is shown.

trajectories are generated using the forward model (see Fig. 3(b)). Finally, a control policy, g , is learned from the generated optimal trajectories (see Fig. 3(c)). Here, f and g are neural network-based. We also verify that the LMU-based model outperforms other architectures (see Tables 2 and 3).

In this article, the following notations are used. State at time t , $s(t)$, consists of position, $p(t)$, and line of sight orientation, $r(t)$. The line of sight is the x -axis of the sensor's frame (Fig. 2(b)). Actuation is $u = [\theta_1, \theta_2, L, B, R_1, R_2]$ whereas change in actuation is $\Delta u = [\Delta\theta_1, \Delta\theta_2, \Delta L, \Delta B, \Delta R]$. D is the size of the dataset under discussion. The subscript gt will designate the ground truth label. *Forward Model*. The forward model data are collected on the robot by first setting the bounds on actuations ($\theta_1, \theta_2, L, B, R_1, R_2$) as explained in Sec. 4.1. The bounds on

changes in actuations ($\Delta\theta_1, \Delta\theta_2, \Delta L, \Delta B, \Delta R$) are also set so that the system has sufficient time to respond to the control signals. The changes in the rigid arm joint angles, $\Delta\theta_1$ and $\Delta\theta_2$, were bound between $[-5 \text{ deg}, +5 \text{ deg}]$ and $[-10 \text{ deg}, +10 \text{ deg}]$, respectively. The changes in the pressures, ΔB and ΔR , were bound between $[-20.68, 20.68] \text{ kPa}$, and changes in the extrusion, ΔL , between $[0, 0.7] \text{ cm}$. After random control inputs are actuated for a predetermined time of 5 s, the actuations are held constant for 3 s to let the system settle as the oscillations dampen and data are collected throughout 8 s at 4 Hz. When operating faster than 4 Hz, the set range of actuation changes for the stepper motor would result in high speeds that risk damaging the magnetic sensor cable that goes along the length of the soft arm. This being the most expensive part of the system, a 4 Hz operating frequency was selected due to the

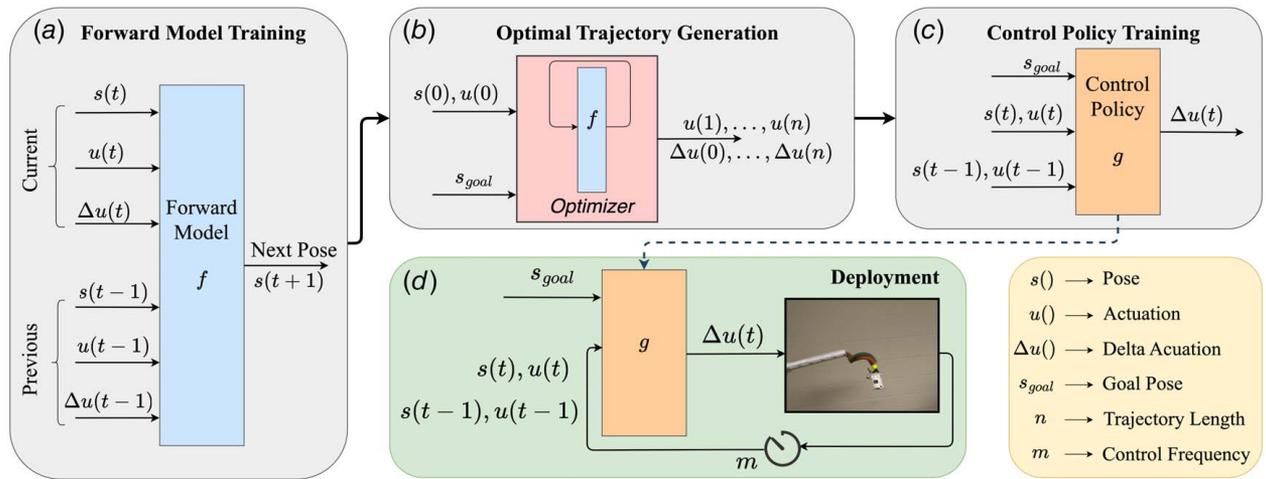


Fig. 3 Pipeline for obtaining and deploying control policy of the hybrid arm. (a) Step one trains the neural-network-based forward model f using data collected on the robot where inputs are poses, actuations, and changes in actuations at current and previous time-steps and the output is the pose at the next time-step. (b) Step two generates optimal trajectories where the trained forward model f is used iteratively to obtain optimal actuations and optimal changes in actuation in reaching a given initial pose, initial actuation, and goal pose. (c) Step three runs generated trajectories on the robot and trains a control policy g that learns optimal changes in actuation. (d) Finally, the learned offline policy g is deployed on the robot iteratively at m control frequency to reach a goal pose given current and previous poses and actuations.

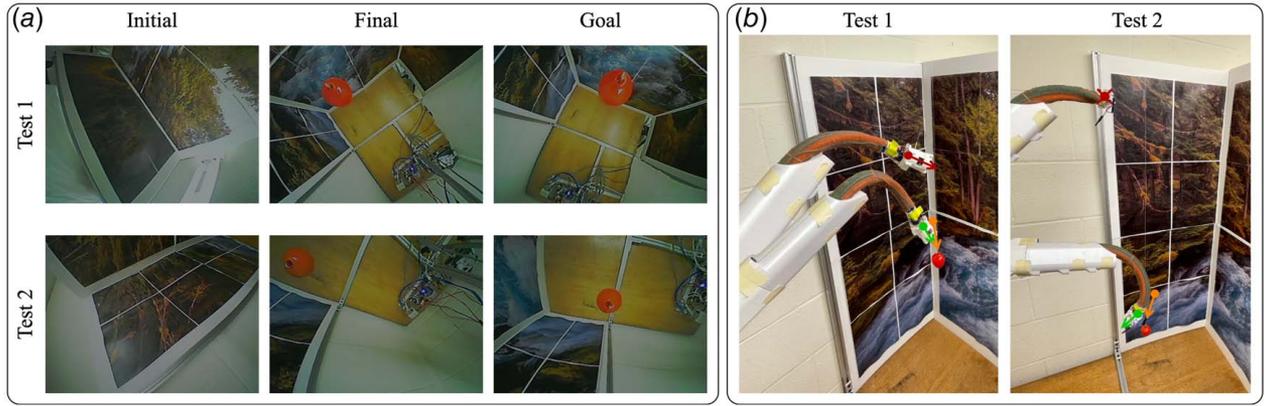


Fig. 4 Bringing the target in view of the tip camera using our policy that controls both position and line-of-sight orientation, as seen in (a) tip camera view and (b) third-person view, where the start, final, and goal poses are indicated by arrows

Table 2 Forward model test errors on 480 unseen poses

Architecture	Position error (cm) ↓	Orientation error (deg) ↓
MLP	1.74	9.46
RNN	1.42	10.91
LSTM	1.14	10.13
LMU	1.13	10.49

Note: The LMU has the best position error while the MLP has the best orientation error. see Fig. 3(a). Best errors are bolded.

limitation of the extrusion component. In total, 9600 training samples were collected in this manner, and the training and testing data were split in a 0.8–0.2 ratio randomly.

The history of inputs must be considered when modeling the forward model due to hysteresis in the system, a phenomenon in soft robotics that typically occurs in dynamic regimes [17]. A system has hysteresis if the behavior of a system depends on the input and the history of previous states. Given that our approach involves using both current and past states as inputs to the forward model and control policy neural networks, RNN-based architectures [51] are better suited for capturing these temporal dependencies compared to MLPs (Multilayer Perceptrons) [52], which lack temporal memory. RNNs [51] are one of the earliest architectures that deal with such time-dependent problems. However, they struggle to capture long-term dependencies and

suffer from issues like vanishing and exploding gradients during training. To address these limitations, more advanced architectures such as long short term memory (LSTM) [53] and LMUs (Legendre Memory Units) [54] have been introduced, which are capable of capturing long-term dependencies. With the data collected, MLP, RNN, LSTM, and LMU-based architectures were trained and tested to evaluate their effectiveness in capturing temporal relations.

As shown in Fig. 3(a), the inputs to the network are the poses, $s(t)$, $s(t-1)$, actuations, $u(t)$, $u(t-1)$, and changes in actuations, $\Delta u(t)$, $\Delta u(t-1)$ at current time-step, t , and previous time-step, $t-1$, respectively. The output of the forward model is the pose at the next time-step, $s(t+1)$. Hyperparameters of these models include memory size, batch size, learning rate, number of layers, hidden layer size, and window length. The overall function of the forward models is represented mathematically in Eq. (1).

$$s(t+1) = f(s(t), u(t), \Delta u(t), s(t-1), u(t-1), \Delta u(t-1)) \quad (1)$$

$$\frac{1}{\|D\|} \sum_{i=1}^D \|\hat{p}^i - p_{gt}^i\|^2 - \frac{1}{\|D\|} \sum_{i=1}^D \frac{(\hat{r}^i \cdot r_{gt}^i)}{(\|\hat{r}^i\| \|r_{gt}^i\|)} \quad (2)$$

The loss function for learning the forward model is composed of the mean square error for learning position $p = (x, y, z)$ and cosine similarity for learning orientation, r , which is shown in Eq. (2). Here, \hat{p}^i and \hat{r}^i denote the (x, y, z) and line of sight orientation prediction from the forward model of i th data in D , whereas, p_{gt}^i and r_{gt}^i denote the corresponding ground truth. Ideally, the first term goes to 0 while the second term goes to 1.

$$\min_{u(1), \dots, u(n), \Delta u(0), \dots, \Delta u(n)} \alpha \|p_{\text{goal}} - p(n)\|^2 + \lambda \sum_{t=0}^{n-1} \|p_{\text{goal}} - p(t)\|^2 + \beta \left(1 - \frac{r_{\text{goal}} \cdot r(n)}{\|r_{\text{goal}}\| \|r(n)\|}\right)^2 + \gamma \sum_{t=0}^n \|\Delta u(t)\|^2$$

where $u_{\min} \leq u(t) \leq u_{\max}$, $\Delta u_{\min} \leq \Delta u(t) \leq \Delta u_{\max} \quad \forall t \in [0, 1, \dots, n]$
 $s(t+1) = f(s(t), u(t), \Delta u(t), s(t-1), u(t-1), \Delta u(t-1)) \quad \forall t \in [1, \dots, n]$
 $s(t+1) = f(s(t), u(t), \Delta u(t), s(t), u(t), \Delta u(t)) \quad \text{if } t = 0$ (3)

Trajectory Optimization. Once the forward model is trained, an optimizer is used to find optimal trajectories for a given start pose, $s(0)$, its corresponding start actuation, $u(0)$, and desired goal pose, s_{goal} , in the workspace as shown in Fig. 3(b). These optimal trajectories serve as expert data that the control policy will use to learn the inverse kinematics of the system.

To optimize the trajectory, the Adam optimizer [55] was used. To solve this constrained optimization problem more quickly, a “warm-start” was used where the solution to the unconstrained optimization was used as the initial guess for the constrained

optimization [56]. A subset of 10 trajectories were tested to compare optimization with a “cold-start” and “warm-start.” Empirically, using a warm-start optimization increased speeds by about threefold. With this setting, the Adam optimizer produced more accurate results in fewer iterations. The objective function is shown in Eq. (3) where minimization for a given start and goal states along with the constraints is performed.

α , λ , β , and γ are all manually tuned hyperparameters that weigh each term in the objective function. The time horizon, n , is another hyperparameter. The first term in the objective function drives the

Table 3 Comparison of policies with different architectures of our learned control policy with respect to position and orientation errors over 11 random start-goal pairs each run 5 times

Controller	Average final pose error		Average best position in $x(t)$			Average best orientation in $x(t)$		
	Pos (cm) ↓	Ori (deg) ↓	Pos (cm) ↓	t (s) ↓	Corr Ori (deg) ↓	Ori (deg) ↓	t (s) ↓	Corr Pos (cm) ↓
Ours (MLP)	7.31	25.65	5.18	8.82	22.44	18.37	8.48	6.93
Ours (RNN)	5.44	21.34	3.85	9.01	22.09	16.5	6.97	7.97
Ours (LSTM)	4.18	21.47	3.01	8.06	23.11	15.47	7.42	5.53
Ours (LMU)	3.73	17.78	2.75	9.15	20.52	13.13	8.09	5.92
Ours (LMU)+3.8 g	4.33	16.85	3.04	8.97	17.7	12.74	8.92	5.44
Ours (LMU)+10.8 g	5.22	18.34	3.61	8.69	16.76	12.74	6.46	7.05

Note: Additionally, the best position and corresponding time and orientation errors are reported. Similarly, the best orientation error and corresponding time and position error are reported. The LMU-based has the best performance overall. Finally, the LMU control policy performance is evaluated with different loads on the end effector to show robustness to external loads unseen during training. Best errors are bolded.

final position toward the desired goal position. The second term encourages each step of the trajectory to proceed toward the goal position. The third term drives the final orientation to the desired goal orientation. The final term encourages minimal actuation effort across the trajectory. The forward model and the physical limitations of the system act as the constraints. The sequence of actuations, $u(0), \dots, u(n)$, and the sequence of change in actuations, $\Delta u(0), \dots, \Delta u(n)$, are the optimization variables. The bounds on the changes in actuation are the same as in the forward model. This prevents the optimizer from exploiting unexplored regions with larger actuations at the cost of limiting the system's ability to dampen oscillatory behavior. The optimizer was run until convergence, and if the solution was not within 1 cm and 25 deg of the goal, then the optimizer was initialized with a different seed. Optimization would continue as such until convergence of the tolerance or until a 1 h time limit was reached. To speed up the generation of trajectories, optimization was done in parallel. To summarize, this objective function seeks to minimize the end effector's position error, orientation error, and minimal actuation effort for reaching

tasks subject to physical constraints. *Control Policy Network.* Solving nonlinear optimization may not always lead to convergence, so trajectories with final position errors larger than 10 cm and orientation errors larger than 36 deg were discarded. About 4% of the 2926 trajectories were discarded and visualized in Fig. 5. For clarity, only the desired end pose of the trajectories is plotted. Additionally, the optimizer may have exploited inaccuracies in the learned forward model. So, prior to training the control policy, the generated optimal trajectories are collected on the hybrid arm to eliminate the simulation-to-reality gap, since directly training on simulated optimal trajectories resulted in poor performance: 15.8 cm position error and 43.56 deg orientation error when evaluated on 11 test cases of each goal pose each run 5 times.

With the filtered data, various models are trained to learn the inverse kinematics of the system shown in Eq. (4). As seen in Fig. 3(c), this policy takes the poses and actuations at current and previous time-steps along with the goal pose as input. The output is the next optimal change in actuation to progress toward the goal pose. More specifically, RNN, LMU, LSTM, and a 5-layer MLP-based models are trained and compared. To train these policies, the mean square error between the predicted change in actuations, $\Delta \hat{u}$, and corresponding ground truth, Δu_{gt} , is used as the loss function as shown in Eq. (5). The optimal trajectories provide Δu_{gt} at each time-step to reach toward the goal pose, s_{goal} . At inference time, these policies can be used iteratively at an operating control frequency, m Hz, to navigate to a goal position as illustrated in Fig. 3(d).

$$\Delta u(t) = g(s_{goal}, s(t), u(t), s(t-1), u(t-1)) \quad (4)$$

$$\frac{1}{\|D\|} \sum_{i=1}^D \|\Delta \hat{u}^i - \Delta u_{gt}^i\|^2 \quad (5)$$

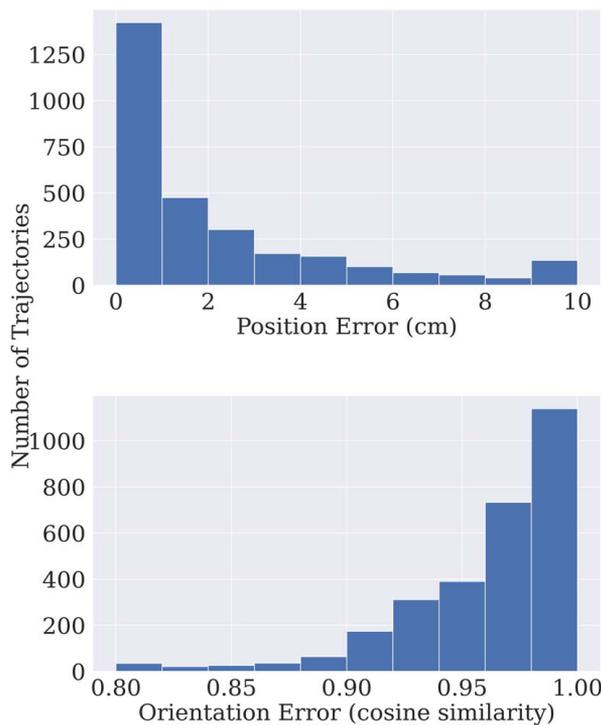


Fig. 5 Of all the generated trajectories, 115 trajectories (4%) were filtered due to poor optimization. A histogram of all their final convergence errors is presented.

4 Experimental Results

This section describes the suite of tests performed to evaluate the method described in Sec. 3.2. First, the predicted forward model poses were evaluated with the measured pose data. Then, the performance of control policies with different architectures was evaluated on 11 different random poses in the workspace. These tests are repeated with loads attached to the end effector as shown in Fig. 2(c) to test the policy's robustness. Additionally, line-tracking tests were evaluated with varying control frequencies.

4.1 Forward Model Evaluation. The forward models are trained using various architectures and evaluated on 480 unseen test poses as shown in Table 2. The visualization of the training and testing data distribution, which includes the entire trajectory along with the end poses, in the form of relative dexterity measure is shown in Fig. 6. To understand the visualization, take, for example, the xyz-grid areas highlighted in yellow which have

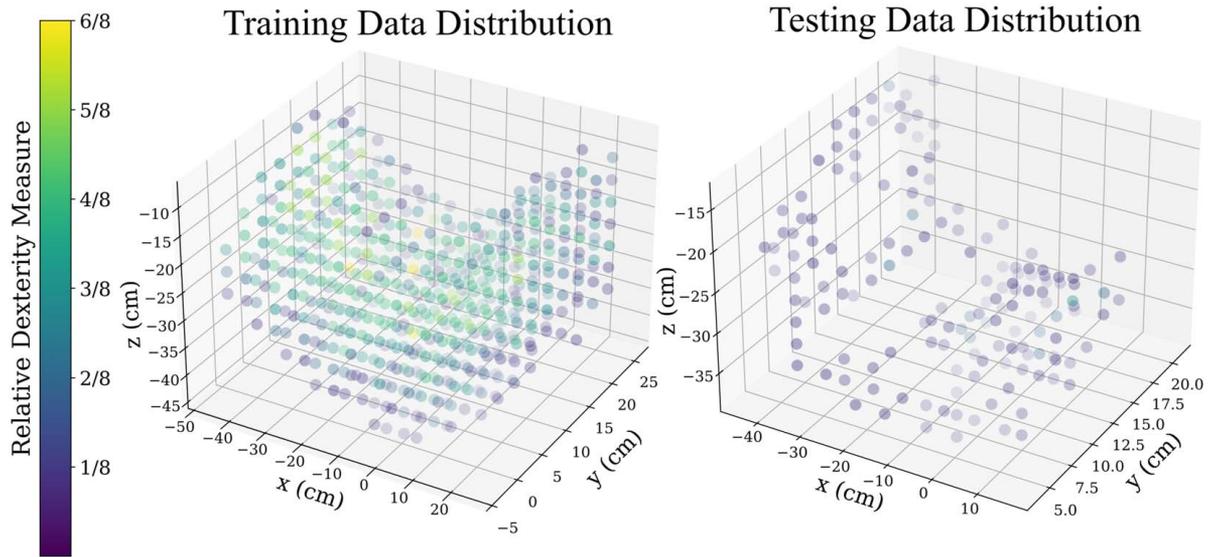


Fig. 6 Three-dimensional visualization of training and testing data in the workspace. Each grid cell is shaded based on the number of octants covered by the orientation vectors at that cell.

a 6/8 dexterity measure in the training data distribution. If we draw a sphere in any of these grid cells, the line-of-sight of our arm-tip can point in six directions out of the eight possible octants in the sphere. The evaluation of forward models on test data in Table 2 indicates the overall learning of dynamics in the workspace with varying dexterity due to the random split of training and testing data as mentioned in Sec. 3.2. Therefore, the test data which were randomly withheld during training help evaluate our learned models on unseen data. Note that the sparsity on the right side of Fig. 6 results from the 20% random sampling of the entire dataset unseen during training.

To evaluate the learned forward model, for each trajectory, the mean squared error and cosine similarity for position and orientation errors are computed over the entire trajectory and averages are calculated. Then, the combined average error is computed over the entire test data and reported in Table 2 for different architectures.

Each architecture was trained for 200 epochs with a batch size of 100 and a learning rate of 5×10^{-4} for a fair comparison. The recurrent architectures had a hidden size of 128, and the LMU had a window length of 2. These settings were found using a manual hyperparameter search. Looking at the position and orientation

errors together, LSTM and LMU perform better than other models and the LMU model was selected for the trajectory generation as its position error was the lowest. A sample trajectory predicted by the forward model based on LMU is shown in Fig. 7.

4.2 System Stochasticity. To baseline the stochasticity of the system, 10 random actuations were sent 5 times each to the system. The range and standard deviation in the measured poses were averaged across the 10 actuations. The results are shown in Table 4, which shows the inherent variability in the hardware.

4.3 Control Policy Evaluation. To construct the control policies, the architectures are trained using 2926 optimal trajectories. Each architecture was trained for 1000 epochs with a batch size of 64 and a learning rate of 1×10^{-4} . The recurrent architectures had a hidden size of 256, and θ was 700 for the LMU. These settings were found manually using a hyperparameter search. A suite of tests is run to evaluate these policies. First, policies are evaluated on 11 different random start-goal pairs each tested 5 times. The LMU is found to have the best performance and is further evaluated by adding external load, tracking trajectories, and determining control frequency robustness. *Reaching Arbitrary Pose.* Control policies trained using various architectures are evaluated on 11 random poses. A summary of the results is given in Table 3. The steady-state position and orientation errors are reported in the first two columns of Table 3. The best position errors, corresponding orientation errors, and their respective times are also provided, as are the best orientation errors, corresponding position errors, and their corresponding times. For each start-goal pair, five tests were run to capture the variance. These results baseline the system. In our previous works

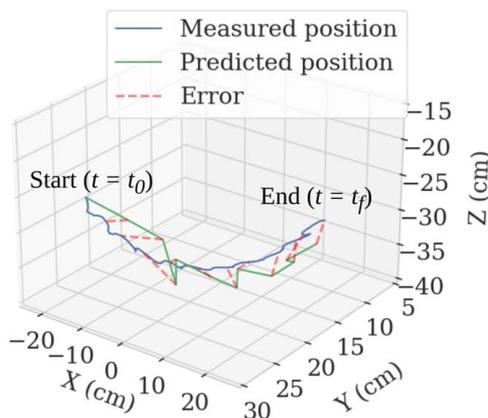


Fig. 7 Three-dimensional visualization of the LMU's predicted position of the end effector versus its measured position over time where the forward model predicts the next state based on current and previous states, and control inputs at each time-step

Table 4 Variability of the pose given an arbitrary actuation

Actuation	Position (cm) ↓	Orientation (deg) ↓
θ_1	0.63 ± 0.66	3.74 ± 3.94
θ_2	0.74 ± 0.78	4.44 ± 1.94
L	1.09 ± 1.23	5.76 ± 4.85
B	1.66 ± 1.80	9.87 ± 9.03
R	0.42 ± 0.45	2.20 ± 1.47

Note: Bending has the highest variance in both position and orientation compared to other actuations.

Table 5 Baseline comparisons

Controller	Dataset	Control frequency (Hz)	Average final position error (cm)	Average final orientation error (deg)	Average ITAE position (min cm)	Average ITAE orientation (min deg)
Learned (LMU)	Random	4	18.74	54.28	19.68	58.37
Learned (LMU)	Dynamic	0.25	6.48	26.34	7.56	28.00
Learned (LMU)	Quasistatic	4	9.49	37.13	9.67	43.84
Learned (LMU)	Optimal	4	3.73	17.78	6.05	28.57
MPC	–	4	27.54	62.35	28.28	62.56
MPPI (1)	–	4	16.28	17.19	18.16	21.48
MPPI (2)	–	4	8.94	23.12	11.26	23.32
Learned (LMU)	Optimal	4	3.73	17.78	6.05	28.57
Learned (LMU)	Dynamic	0.25	6.48	26.34	7.56	28.00

Note: (Top row) Learned policy trained on random trajectories of forward model data, (rows 2 and 3) learned policies trained on optimal trajectories in quasistatic settings tested at quasistatic and dynamic rates, (rows 4, 5, and 6) MPC, MPPI (1), and MPPI (2) where optimization is performed online, and (bottom row) learned policy on optimal trajectories whose combined position and orientation error is better than the baselines. ITAE: Integral of Time-weighted Absolute Error, which is a typical metric used to evaluate control systems. Best errors are bolded.

[26,57], quasistatic methods obtained results of 1–2 cm errors in the order of minutes. Our approach sacrifices a small degree of positional accuracy in exchange for significant improvements in control speed and the ability to effectively control end effector orientation. On average, the poses are reached with 3.73 cm and 17.78 deg error, which is 1.13× the SCA diameter (3.30 cm).

We further evaluate two test cases where the goal pose of a berry-like target is provided such that it is visible from the camera mounted at the tip of the arm. Our control policy is able to make the targets appear in view of the tip camera as seen in Fig. 4, using which image-based control methods like visual servoing can be explored in the future to close the gap. *Baseline Comparison.* To validate the performance of our proposed approach, we perform several baseline tests. First, we evaluate using random trajectories of the forward model data to directly train the control policy. The results in Table 5 show that the policy trained on the forward model random data has less generalization compared to the policy trained on optimal trajectories. But by leveraging the forward model, generating optimal trajectories helps explore the workspace more efficiently resulting in a better policy.

Second, we have shown that the forward model data collected in a dynamic manner account for dynamic effects by running our pipeline shown in Fig. 3 but learned on quasistatic data and then testing in a dynamic regime. If no dynamic effects were learned by the policy, then this method would have similar results to a policy trained on dynamic data. However, the results from Table 5 show that the policy trained on data collected in a quasistatic manner does not perform at non-quasistatic operating frequencies. On the other hand, the performance of policy learned on dynamic data does not deviate significantly with increasing operating frequency as shown in the control frequency robustness test in Sec. 4.3. This shows that our policy learns how to control the system to achieve the desired poses quickly although oscillations are not dampened.

Third baseline comparison is done with an MPC [58] and a model predictive path integral (MPPI) controller [59] using the learned LMU forward model. Our findings are summarized as follows:

(i) Since MPC operates in real-time, it requires conducting the same nonlinear optimizations as described in Eq. (3) at every time-step. To enable the 4 Hz operating frequency on a laptop with 16 CPUs running at 2.3 GHz and a GeForce RTX 3060 Mobile Graphics card running at 1.4 GHz, we chose a time horizon, n , of 2 and a maximum of 20 iterations for optimization. Because of the nonconvexity of the optimization as well as the limited computation time, the MPC has poor real-time performance as seen in Table 3, though recent studies have explored methods to improve its real-time efficiency [60–62].

(ii) The MPC controller was also implemented in simulation to evaluate its accuracy in a dynamic regime given more compute time. With a time horizon of 20 steps and max iterations of 100 k for optimization, a single trajectory required about 3 h of computation time to reach position accuracy of 1.1 cm and orientation accuracy of 5 deg with 33 actuations applied in a closed-loop fashion. Simulation offers superior accuracy, but computation time and simulation-to-reality errors limit realizing these accuracies on the physical system at the operating frequency of 4 Hz.

(iii) Two MPPI controllers with different weights are also tested as baselines. Each computes 10,000 rollouts each with a time horizon of 15 steps in parallel. MPPI (1) uses the same optimization objective (Eq. (3)), and MPPI (2) weights the position error 4× more. The next actuation is then computed as a weighted average as shown in Eq. (6) where S_k is the objective function value of rollout k .

$$\Delta u(t+1) = \frac{\sum_{k=0}^K w_k u_k}{\sum_{k=0}^K w_k}, \quad w_k = e^{-(S_k/\lambda)} \quad (6)$$

While their performance is better than MPC, they could not achieve lower position and orientation errors compared to our learned controller because our method learns this difficult nonconvex optimization offline.

Our learned controller outperforms these baselines mainly due to the lack of simulation-to-reality error. The baselines are optimized over a learned forward model, where low data regimes may be

Table 6 Evaluation of position control and position and orientation control of the soft component of our hybrid system, where their respective policies can take the arm to the desired goal pose within 4–7 s

Workspace	Average time to converge (s)	Convergence percentage (%)	Average final position error (cm)	Average final orientation error (deg)
$(L, B, R_1, R_2) + (x, y, z)$	4.4	100	2.29 (tol = 2.5)	–
$(L, B, R_1, R_2) + (x, y, z, r_x)$	4.4	90	2.56 (tol = 3)	0.86 (tol = 0.8)
$(L, B, R_1, R_2) + (x, y, z, r_x)$	7.6	60	2.19 (tol = 2.5)	0.90 (tol = 0.9)

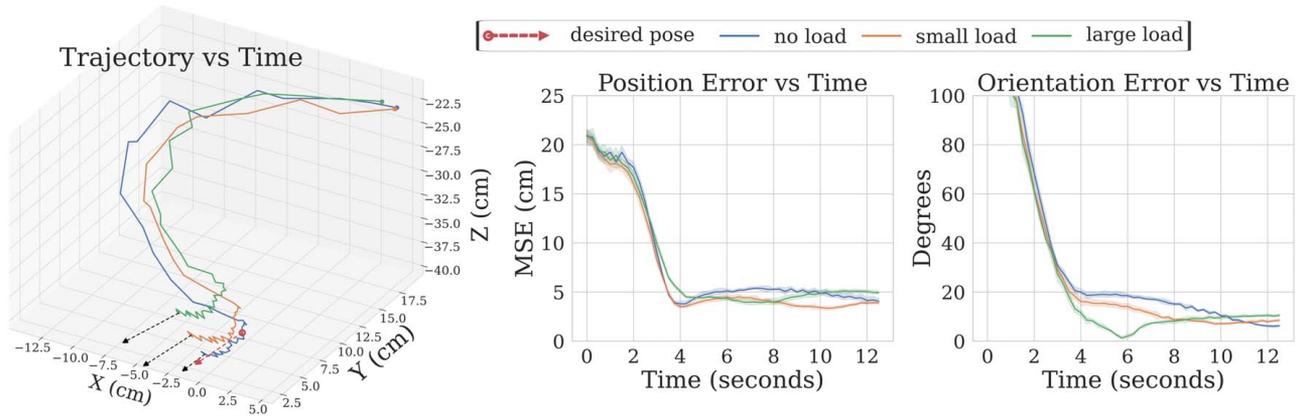


Fig. 8 Comparison of trajectory, position error, and orientation error across different loads. The starting positions change due to different loads, but the position and orientation converge to the desired pose over time.

Table 7 Control frequency and pose tolerance analysis of the line test

Control frequency (Hz)	Best tolerance with 100% convergence	Average trajectory time (s)	Average final position error (cm)	Average final orientation error (deg)
1	2 cm, 5 deg	15.0 ± 1.1	1.73 ± 0.24	1.89 ± 0.69
2	3 cm, 5 deg	5.5 ± 0.3	2.50 ± 0.53	1.77 ± 0.88
4	3 cm, 5 deg	2.8 ± 0.1	2.58 ± 0.20	1.55 ± 1.29
8	3 cm, 5 deg	1.6 ± 0.0	2.16 ± 0.44	2.49 ± 0.72

Note: The higher the frequency, the shorter the trajectory times. Lowering the control frequency results in better performance. Errors and standard deviation are reported for the final position and orientation.

exploited that do not translate to the real system. On the contrary, our proposed approach uses optimal trajectories run on the real system to train the learned policy and avoid the simulation-to-reality gap. *Scalability.* To evaluate the scalability of learning policies with the number of states, a comparison with two smaller workspaces involving just the soft components (extrusion, bending, and rotation) is performed. The first workspace involves just position control and the second one also involves orientation control, where the pipeline shown in Fig. 3 is implemented with appropriate state modifications. Table 6 shows the performance of the two policies. With the full workspace of our system, we require 2926 optimal trajectories to learn a policy that can reach desired poses with an average position error of 3.73 cm and orientation error of 17.78 deg error within 12.5 s, compared to 1115 trajectories in both versions of smaller workspace in reaching less than 3 cm position and 25–36 deg orientation within 4–7 s.

Robustness to External Load. The policy was evaluated with different loading conditions. A small load of 3.8 g ($\sim 10\%$ of the SCA's weight) and a large load of 10.8 g ($\sim 30\%$ of the SCA's weight) were attached to the tip of the arm as seen in Fig. 2(c). The same 11 trajectories were tested 5 times in these different loading scenarios to capture the variance of the trajectories. As seen in Table 3, the loads have minimal effect on the errors achieved by the policy. Averaging over 55 tests, the final position (column 1) and orientation error (column 2) degrade slightly, but the loads even allow the best positions (column 5) and orientations (column 8) to be achieved more quickly. We attribute this outcome to the damped oscillations due to the larger end effector load. Errors of a test case at different load conditions are shown in Fig. 8 showing the robustness of our policy. *Control Frequency Robustness.* To further evaluate the policy's robustness, multiple control frequencies were also tested for the line-following trajectory. The line was still segmented into 10 evenly spaced waypoints, and control frequencies of 1, 2, 4, and 8 Hz were evaluated. Frequencies faster than 8 Hz were not attempted due to hardware limitations. The results (see Table 7)

show that if a pose can be reached within a certain tolerance, it is reachable with any slower control frequency. For example, if 3 cm and 5 deg can be converged upon at 8 Hz, then it can also be converged at 1 Hz. This is an interesting finding because the control policy was learned with data collected at 4 Hz.

5 Conclusion

Developing fast, robust, and accurate soft robotic controllers remains an open challenge. In this work, the difficulty of this challenge is increased by tackling orientation as well as position control. Furthermore, the hybrid soft-rigid system considered in this article is challenging due to variations in the orientation of the soft manipulator base, varying soft manipulator length, and the torsional pneumatic actuator. A dynamic model is first learned and evaluated using neural networks. Then, optimal trajectories were generated to serve as expert data for training a control policy. This optimal control policy was trained and then evaluated on a physical platform. Its capability of reaching arbitrary poses in the workspace, tracking pose trajectories, and operating at different control frequencies and with unknown loads was analyzed. Arbitrary poses in the workspace were reached within 3.73 cm and 17.78 deg in 12.5 s. The policy was also robust to operating control frequency and minimally affected by external loads without further training or fine-tuning.

While this work demonstrates that accuracy in both position and orientation can be achieved quickly in a complex 3D workspace, there are limitations. Without online learning, our method requires proper data collection to ensure good generalization in all parts of the workspace for both the forward model and the policy. Similarly, if the system changes over time due to issues like wear and tear, online learning must be incorporated to adapt to these changes. Additionally, the policy's performance is limited by the quality of the expert trajectories. Ensuring high-quality nonconvex

optimizations is difficult with current methods lacking global convergence guarantees. Despite these limitations, our method is a step toward fast and accurate dynamic pose control of soft manipulators. With these promising results, future work will seek to outperform this baseline. Visual feedback can be incorporated if target images or objects are provided. More application-focused extensions like obstacle avoidance and controlling a serial soft manipulator configuration can also be explored.

Acknowledgment

This paper benefits from a joint NSF-USDA COALESCE Grant 2021-67021-34418, the NSF Grant No. DBI-2019674 for CROPPS, and is supported by the AIFARMS National AI Institute in agriculture.

Conflict of Interest

There are no conflicts of interest.

Data Availability Statement

The datasets generated and supporting the findings of this article are obtainable from the corresponding author upon reasonable request.

References

- [1] Uppalapati, N. K., Walt, B., Havens, A. J., Mahdian, A., Chowdhary, G., and Krishnan, G., 2020, "A Berry Picking Robot With a Hybrid Soft-Rigid Arm: Design and Task Space Control," *Robotics Science and Systems*, Virtual, July, p. 95.
- [2] Chowdhary, G., Gazzola, M., Krishnan, G., Soman, C., and Lovell, S., 2019, "Soft Robotics as an Enabling Technology for Agroforestry Practice and Research," *Sustainability*, **11**(23), p. 6751.
- [3] Wen, T., Hu, J., Zhang, J., Li, X., Kang, S., and Zhang, N., 2024, "Design, Performance Analysis, and Experiments of a Soft Robot for Rescue," *ASME J. Mech. Rob.*, **16**(7), p. 071011.
- [4] Kadylak, T., Uppalapati, N., Huq, A., Krishnan, G., and Rogers, W. A., 2023, "Engaging Healthcare Providers to Design a Robot for Telehealth," *Ergon. Des.*, p. 10648046231193287.
- [5] Liu, Y., Ge, Z., Yang, S., Walker, I. D., and Ju, Z., 2019, "Elephant's Trunk Robot: An Extremely Versatile Under-Actuated Continuum Robot Driven by a Single Motor," *ASME J. Mech. Rob.*, **11**(5), p. 051008.
- [6] Konda, R., Bombara, D., Chow, X., and Zhang, J., 2024, "Kinematic Modeling and Open-loop Control of a Twisted String Actuator-Driven Soft Robotic Manipulator," *ASME J. Mech. Rob.*, **16**(4), p. 041007.
- [7] Banerjee, H., Pusalkar, N., and Ren, H., 2018, "Single-Motor Controlled Tendon-Driven Peristaltic Soft Origami Robot," *ASME J. Mech. Rob.*, **10**(6), p. 064501.
- [8] Xu, Y., Yan, D., Zhang, K., Li, X., Xing, Y., and Shao, L.-H., 2022, "Soft Robot Based on Hyperelastic Buckling Controlled by Discontinuous Magnetic Field," *ASME J. Mech. Rob.*, **14**(1), p. 011008.
- [9] Yip, M. C., and Camarillo, D. B., 2016, "Model-less Hybrid Position/Force Control: A Minimalist Approach for Continuum Manipulators in Unknown, Constrained Environments," *IEEE Rob. Autom. Lett.*, **1**(2), pp. 844–851.
- [10] Ma, K., Chen, X., Zhang, J., Xie, Z., Wu, J., and Zhang, J., 2023, "Inspired by Physical Intelligence of an Elephant Trunk: Biomimetic Soft Robot With Pre-programmable Localized Stiffness," *IEEE Rob. Autom. Lett.*, **8**(5), pp. 2898–2905.
- [11] Trivedi, D., Rahn, C. D., Kier, W. M., and Walker, I. D., 2008, "Soft Robotics: Biological Inspiration, State of the Art, and Future Research," *Appl. Bionics Biomech.*, **5**(3), pp. 99–117.
- [12] Dou, W., Zhong, G., Cao, J., Shi, Z., Peng, B., and Jiang, L., 2021, "Soft Robotic Manipulators: Designs, Actuation, Stiffness Tuning, and Sensing," *Adv. Mater. Technol.*, **6**(9), p. 2100018.
- [13] Satheeshbabu, S., Uppalapati, N. K., Chowdhary, G., and Krishnan, G., 2019, "Open Loop Position Control of Soft Continuum Arm Using Deep Reinforcement Learning," International Conference on Robotics and Automation (ICRA), Montreal, Canada, May, pp. 5133–5139.
- [14] Uppalapati, N. K., and Krishnan, G., 2021, "Design and Modeling of Soft Continuum Manipulators Using Parallel Asymmetric Combination of Fiber-Reinforced Elastomers," *ASME J. Mech. Rob.*, **13**(1), p. 011010.
- [15] George Thuruthel, T., Falotico, E., Manti, M., Pratesi, A., Cianchetti, M., and Laschi, C., 2017, "Learning Closed Loop Kinematic Controllers for Continuum Manipulators in Unstructured Environments," *Soft Rob.*, **4**(3), pp. 285–296.

- [16] Rolf, M., and Steil, J. J., 2013, "Efficient Exploratory Learning of Inverse Kinematics on a Bionic Elephant Trunk," *IEEE Trans. Neural Netw. Learn. Syst.*, **25**(6), pp. 1147–1160.
- [17] George Thuruthel, T., Ansari, Y., Falotico, E., and Laschi, C., 2018, "Control Strategies for Soft Robotic Manipulators: A Survey," *Soft Rob.*, **5**(2), pp. 149–163.
- [18] Uppalapati, N. K., and Krishnan, G., 2020, "Valens: Design of a Novel Variable Length Nested Soft Arm," *IEEE Rob. Autom. Lett.*, **5**(2), pp. 1135–1142.
- [19] Uppalapati, N. K., Kadylak, T., Rogers, W., and Krishnan, G., 2020, "Morphological Switching Robots to Support Independent Living for Older Adults," International Conference on Intelligent Robotics and Systems, Virtual, October, pp. 11995–11997.
- [20] Thuruthel, T. G., Falotico, E., Renda, F., and Laschi, C., 2018, "Model-Based Reinforcement Learning for Closed-Loop Dynamic Control of Soft Robotic Manipulators," *IEEE Trans. Rob.*, **35**(1), pp. 124–134.
- [21] Gan, Y., Li, P., Jiang, H., Wang, G., Jin, Y., Chen, X., and Ji, J., 2022, "A Reinforcement Learning Method for Motion Control With Constraints on an HPN Arm," *IEEE Rob. Autom. Lett.*, **7**(4), pp. 12006–12013.
- [22] Centurelli, A., Arleo, L., Rizzo, A., Tolu, S., Laschi, C., and Falotico, E., 2022, "Closed-Loop Dynamic Control of a Soft Manipulator Using Deep Reinforcement Learning," *IEEE Rob. Autom. Lett.*, **7**(2), pp. 4741–4748.
- [23] Escande, C., Chettibi, T., Merzouki, R., Coelen, V., and Pathak, P. M., 2014, "Kinematic Calibration of a Multisection Bionic Manipulator," *IEEE/ASME Trans. Mechatron.*, **20**(2), pp. 663–674.
- [24] Webster III, R. J., and Jones, B. A., 2010, "Design and Kinematic Modeling of Constant Curvature Continuum Robots: A Review," *Int. J. Rob. Res.*, **29**(13), pp. 1661–1683.
- [25] Shi, J., Abad Guaman, S., Dai, J., and Wurdemann, H., 2023, "Position and Orientation Control for Hyper-elastic Multi-segment Continuum Robots," *IEEE/ASME Trans. Mechatron.*, **29**(2), pp. 995–1006.
- [26] Satheeshbabu, S., Uppalapati, N. K., Fu, T., and Krishnan, G., 2020, "Continuous Control of a Soft Continuum Arm Using Deep Reinforcement Learning," International Conference on Soft Robotics, Virtual, May, pp. 497–503.
- [27] Wu, Q., Gu, Y., Li, Y., Zhang, B., Chepinskiy, S. A., Wang, J., Zhilenkov, A. A., Krasnov, A. Y., and Chernyi, S., 2020, "Position Control of Cable-Driven Robotic Soft Arm Based on Deep Reinforcement Learning," *Information*, **11**(6), p. 310.
- [28] Lin, D., Chen, W., He, K., Jiao, N., Wang, Z., and Liu, L., 2022, "Position and Orientation Control of Multisection Magnetic Soft Microcatheters," *IEEE/ASME Trans. Mechatron.*, **28**(2), pp. 907–918.
- [29] Haggerty, D. A., Banks, M. J., Kamenar, E., Cao, A. B., Curtis, P. C., Mezić, I., and Hawkes, E. W., 2023, "Control of Soft Robots With Inertial Dynamics," *Sci. Rob.*, **8**(81).
- [30] Bruder, D., Fu, X., Gillespie, R. B., Remy, C. D., and Vasudevan, R., 2020, "Data-Driven Control of Soft Robots Using Koopman Operator Theory," *IEEE Trans. Rob.*, **37**(3), pp. 948–961.
- [31] Bruder, D., Fu, X., Gillespie, R. B., Remy, C. D., and Vasudevan, R., 2021, "Koopman-Based Control of a Soft Continuum Manipulator Under Variable Loading Conditions," *IEEE Rob. Autom. Lett.*, **6**(4), pp. 6852–6859.
- [32] Katschmann, R. K., Thieffry, M., Goury, O., Kruszewski, A., Guerra, T.-M., Duriez, C., and Rus, D., 2019, "Dynamically Closed-Loop Control of Soft Robotic Arm Using a Reduced Order Finite Element Model With State Observer," International Conference on Soft Robotics, Seoul, South Korea, April, pp. 717–724.
- [33] Fischer, O., Toshimitsu, Y., Kazempour, A., and Katschmann, R. K., 2023, "Dynamic Task Space Control Enables Soft Manipulators to Perform Real-World Tasks," *Adv. Intell. Syst.*, **5**(1), p. 2200024.
- [34] Gillespie, M. T., Best, C. M., Townsend, E. C., Wingate, D., and Killpack, M. D., 2018, "Learning Nonlinear Dynamic Models of Soft Robots for Model Predictive Control With Neural Networks," International Conference on Soft Robotics, Livorno, Italy, April, pp. 39–45.
- [35] Tariverdi, A., Venkiteswaran, V. K., Richter, M., Elle, O. J., Tørresen, J., Mathiassen, K., Misra, S., and Martinsen, Ø. G., 2021, "A Recurrent Neural-Network-Based Real-Time Dynamic Model for Soft Continuum Manipulators," *Front. Rob. AI*, **8**, p. 631303.
- [36] Li, M., Kang, R., Branson, D. T., and Dai, J. S., 2017, "Model-Free Control for Continuum Robots Based on an Adaptive Kalman Filter," *IEEE/ASME Trans. Mechatron.*, **23**(1), pp. 286–297.
- [37] Doroudchi, A., and Berman, S., 2021, "Configuration Tracking for Soft Continuum Robotic Arms Using Inverse Dynamic Control of a Cosserat Rod Model," International Conference on Soft Robotics, New Haven, CT, April, pp. 207–214.
- [38] Thuruthel, T. G., Falotico, E., Renda, F., and Laschi, C., 2017, "Learning Dynamic Models for Open Loop Predictive Control of Soft Robotic Manipulators," *Bioinspir. Biomim.*, **12**(6), p. 066003.
- [39] Alqumsan, A. A., Khoo, S., and Norton, M., 2019, "Robust Control of Continuum Robots Using Cosserat Rod Theory," *Mech. Mach. Theory*, **131**, pp. 48–61.
- [40] Jitsoho, R., Lum, T. G. W., Okamura, A., and Liu, K., 2023, "Reinforcement Learning Enables Real-Time Planning and Control of Agile Maneuvers for Soft Robot Arms," Conference on Robotic Learning, Atlanta, GA, November, pp. 1131–1153.
- [41] Nazeer, M. S., Bianchi, D., Campinoti, G., Laschi, C., and Falotico, E., 2023, "Policy Adaptation Using an Online Regressing Network in a Soft Robotic Arm," International Conference on Soft Robotics, Singapore, April, IEEE, pp. 1–7.

- [42] Nazeer, M. S., Laschi, C., and Falotico, E., 2024, "RI-Based Adaptive Controller for High Precision Reaching in a Soft Robot Arm," *IEEE Trans. Rob.*, **40**, pp. 2498–2512.
- [43] Piqué, F., Kalidindi, H. T., Fruzzetti, L., Laschi, C., Menciassi, A., and Falotico, E., 2022, "Controlling Soft Robotic Arms Using Continual Learning," *IEEE Rob. Autom. Lett.*, **7**(2), pp. 5469–5476.
- [44] Preiss, J. A., Millard, D., Yao, T., and Sukhatme, G. S., 2022, "Tracking Fast Trajectories With a Deformable Object Using a Learned Model," International Conference on Robotics and Automation (ICRA), Philadelphia, PA, May, IEEE, pp. 1351–1357.
- [45] Huang, Y., Hofer, M., and D'Andrea, R., 2021, "Offset-Free Model Predictive Control: A Ball Catching Application With a Spherical Soft Robotic Arm," International Conference on Intelligent Robotics and Systems, Prague, Czech Republic, September, IEEE, pp. 563–570.
- [46] Céspedes, A., Terreros, R., Morales, S., Huamaní, A., Fabián, J., and Canahuire, R., 2024, "Model Predictive Control of a Soft Laparoscope Using Neural Networks," International Conference on Mechatronics, Control and Robotics, Jeju, South Korea, February, IEEE, pp. 57–61.
- [47] Jensen, S., Salmon, J. L., and Killpack, M. D., 2024, "Model Evolutionary Gain-Based Predictive Control (MEGA-PC) for Soft Robotics," International Conference on Soft Robotics, San Diego, USA, April, IEEE, pp. 816–823.
- [48] Tang, Z., Xin, W., Wang, P., and Laschi, C., 2024, "Learning-Based Control for Soft Robot–Environment Interaction With Force/Position Tracking Capability," *Soft Rob.*, **11**(5), pp. 767–778.
- [49] Hyatt, P., Wingate, D., and Killpack, M. D., 2019, "Model-Based Control of Soft Actuators Using Learned Non-linear Discrete-Time Models," *Front. Rob. AI*, **6**, p. 22.
- [50] Singh, G., and Krishnan, G., 2017, "A Constrained Maximization Formulation to Analyze Deformation of Fiber Reinforced Elastomeric Actuators," *Smart Mater. Struct.*, **26**(6), p. 065024.
- [51] Rumelhart, D. E., Hinton, G. E., and Williams, R. J., 1985, "Learning Internal Representations by Error Propagation," Tech. Rep., California Univ San Diego La Jolla Inst for Cognitive Science.
- [52] Haykin, S., 1998, *Neural Networks: A Comprehensive Foundation*, Prentice Hall PTR, Englewood Cliffs, NJ.
- [53] Hochreiter, S., and Schmidhuber, J., 1997, "Long Short-Term Memory," *Neural Comput.*, **9**(8), pp. 1735–1780.
- [54] Voelker, A., Kajić, I., and Eliasmith, C., 2019, "Legendre Memory Units: Continuous-Time Representation in Recurrent Neural Networks," Advances in Neural Information and Processing Systems, Vancouver, Canada, December, pp. 15570–15579.
- [55] Kingma, D. P., and Ba, J., 2015, "Adam: A Method for Stochastic Optimization," International Conference for Learning Representations, San Diego, CA, May.
- [56] Sambharya, R., Hall, G., Amos, B., and Stellato, B., 2024, "Learning to Warm-Start Fixed-Point Optimization Algorithms," *J. Mach. Learn. Res.*, **25**(166), pp. 1–46.
- [57] Kamtikar, S., Marri, S., Walt, B., Uppalapati, N. K., Krishnan, G., and Chowdhary, G., 2022, "Visual Servoing for Pose Control of Soft Continuum Arm in a Structured Environment," *IEEE Rob. Autom. Lett.*, **7**(2), pp. 5504–5511.
- [58] Garcia, C. E., Prett, D. M., and Morari, M., 1989, "Model Predictive Control: Theory and Practice—A Survey," *Automatica*, **25**(3), pp. 335–348.
- [59] Williams, G., Aldrich, A., and Theodorou, E. A., 2017, "Model Predictive Path Integral Control: From Theory to Parallel Computation," *J. Guid. Control Dyn.*, **40**(2), pp. 344–357.
- [60] Alonso, C. A., and Tseng, S.-H., 2022, "Effective GPU Parallelization of Distributed and Localized Model Predictive Control," International Conference on Control and Automation, Naples, Italy, June, IEEE, pp. 199–206.
- [61] Adabag, E., Atal, M., Gerard, W., and Plancher, B., 2024, "Mpcgpu: Real-Time Nonlinear Model Predictive Control Through Preconditioned Conjugate Gradient on the GPU," International Conference on Robotics and Automation (ICRA), Yokohama, Japan, May, IEEE, pp. 9787–9794.
- [62] Lee, Y., Choi, K. H., and Kim, K.-S., 2024, "GPU-Enabled Parallel Trajectory Optimization Framework for Safe Motion Planning of Autonomous Vehicles," *IEEE Rob. Autom. Lett.*, **9**(11), pp. 10407–10414.