Science Advances
AAAS

Supplementary Materials for

Muscle coactivation-inspired stiffness-driven control strategy for antagonistic soft robots

Jie Pan *et al.*

Corresponding author: Yifan Wang, yifan.wang@ntu.edu.sg

Sci. Adv. **12**, eaeg3346 (2026)
DOI: 10.1126/sciadv.aeg3346

The PDF file includes:

Supplementary Notes S1 to S15
Tables S1 to S8
Figs. S1 to S15
Legends for movies S1 to S4

Other Supplementary Material for this manuscript includes the following:

Movies S1 to S4

Note S1. Definition and computation of performance metrics

The performance metrics used to evaluate different control strategies are summarized in Table S1. All metrics are computed from experimentally measured physical quantities, including tracking error (°), angular velocity (°/s), and temporal response characteristics (s), under identical experimental conditions to ensure fair comparison. Tracking accuracy is quantified using the mean absolute error (MAE) between the desired trajectory and the measured trajectory. The reference trajectory is depicted in Fig. 2D. Disturbance rejection performance is evaluated by applying an additional 50 g payload at a fixed target angle of 20°. The experimental setup is shown in Fig. 2E. Vibration suppression is evaluated based on the mean angular velocity of the robot during small-deformation tracking tasks. Lower average angular velocity indicates reduced oscillations. Response speed is quantified as the average time required for the robot to first enter a tolerance band of ±1° around the target angle during large-deformation motions. Shorter response times indicate faster dynamic behavior. Learning speed (min) is quantified as the total offline time required before deployment, including experimental data collection and controller training when required. For methods without a learning stage, such as open-loop and PD control, this metric is marked as not applicable (NA). Computational cost (ms) is quantified as the average online computation time per control cycle under the same implementation conditions, where a lower value indicates a lighter real-time computational burden. Cross-configuration tracking accuracy (°) is quantified using the MAE between the desired and measured trajectories when the model identified from the original soft robot is directly applied to robots with different structural dimensions and task configurations, including the pump robot and the bridge robot. Cross-configuration vibration suppression (°/s) is quantified by the mean angular velocity in the corresponding cross-configuration tracking experiments under small-deformation conditions. The open-loop strategy uses predefined input signals without feedback control and serves as a baseline for comparison. For visualization purposes, all metrics are normalized with respect to the maximum value of each metric across the five strategies, enabling a relative comparison of their performance. The score is computed as follows:

$$\text{score} = 1 - \frac{q_i}{\max\left(q_i\right)} \tag{S1}$$

The corresponding raw metric values are reported in Table S1.

Note S2. Implementation details of the RL baseline

The RL baseline in this work is implemented by first identifying a Multi-Layer Perceptron (MLP) surrogate dynamics model from experimental data, and then training a Twin Delayed Deep Deterministic Policy Gradient (TD3) controller offline in the learned surrogate environment before deploying the policy to the Raspberry Pi for real-time control. The MLP uses three hidden layers with widths of 128, 128, and 64 and ReLU activation functions. Training is performed offline on a Windows 10 workstation equipped with an Intel CPU, 31.7 GB RAM, and an NVIDIA GeForce RTX 4060 Laptop GPU, using the Adam optimizer with a learning rate of 10^{-3}, weight decay of 10^{-6}, a batch size of 512, and 300 training epochs. Under this configuration, MLP training takes 30.66 s. To avoid data leakage, the dataset is split by trajectory, with 80% of the trajectories used for training and 20% for validation. Input and output variables are standardized based on the training set statistics before model fitting. The validation results of the MLP surrogate model are shown in Fig. S3. The model accurately reproduces the experimentally observed robot motion over the training and testing trajectories, indicating that the learned surrogate captures the dominant

input-output dynamics within the data distribution used for model identification. To ensure a comparable data budget between the RL baseline and the proposed controller, the MLP surrogate is trained using the same experimental dataset used for Koopman model identification.

The TD3 policy uses a two-layer MLP with hidden widths of 256 and 256, and is trained for 2×10^{6} timesteps with a learning rate of 5×10^{-4}, a replay buffer size of 2×10^{5}, a batch size of 256, a discount factor of 0.99, and a target smoothing coefficient of 5×10^{-3}. TD3 training takes 202.27 min. To improve training stability and obtain a more meaningful baseline, several practical settings are adopted, including a short observation history window, incremental action updates with a maximum physical step size of 0.01 per control step, and a simple curriculum schedule in which the target range is gradually expanded from ±15° to ±25°, and finally to ±40°. The incremental action constraint and reward function are designed to improve training stability and ensure hardware-compatible actuation during real-time deployment. The maximum step size of 0.01 per control cycle is the same as the input increment setting used in the stiffness-driven solver, and is therefore not expected to artificially limit the achievable response speed of the RL policy. In the training environment, this full maximum step size is allowed when the tracking error is larger than 8°, whereas smaller effective step size limits are used near the target, with the limit reduced to 0.006 for errors between 3° and 8° and to 0.003 for errors below 3°. The reward function used for the TD3 baseline balances global tracking accuracy, local tracking accuracy, and input smoothness by penalizing both normalized tracking errors and squared incremental input changes, which is consistent with the main performance considerations in our evaluation. The reported configuration represents the best-performing one among several tested reward formulations. This configuration reflects a trade-off between smooth convergence near the target and aggressive response during large-deformation tracking.

Despite the good predictive accuracy of the surrogate model, the deployed TD3 policy shows inferior closed-loop performance compared with the proposed method. This result should not be interpreted as a fundamental limitation of model-free reinforcement learning on antagonistic soft robots. Rather, it is mainly associated with the limited coverage of the offline dataset used to construct the surrogate environment and the above reward design trade-off. The TD3 policy often approaches large targets slowly or fails to enter the ±1° target band within the evaluation window, increasing the response time. The larger errors during large-deformation motions also increase the overall tracking error, which contributes to the poorer summarized performance of the RL baseline compared with open-loop control. These observations suggest that the reported TD3 baseline represents a conservative offline RL comparison under the same data budget and real-time deployment constraints, rather than an optimized upper bound of RL performance. Further improvements may be possible through broader exploration data, alternative reward shaping, stiffness-aware policy objectives, or different policy training strategies.

Note S3. Comparison with stiffness-aware antagonistic control methods

To clearly illustrate how the proposed method differs from representative existing control methods for antagonistic soft robots, Table S2 compares both implementation requirements and reported performance metrics. The implementation-related comparison includes analytical-model requirement, explicit stiffness-model requirement, explicit constraint handling, and deployment effort. To address performance-related aspects, we further include reported tracking accuracy, disturbance rejection, and response speed when available from the corresponding studies. Given

the differences in robotic platforms, tasks, and experimental conditions, the reported performance metrics are not intended as direct head-to-head comparisons. Instead, they are included to provide additional context for the trade-off between performance and deployment effort. Overall, the comparison shows that the proposed method eliminates the need for analytical modeling and explicit stiffness formulations, while retaining explicit constraint handling. It also provides favorable reported performance and supports rapid deployment on antagonistic soft robots with difficult-to-model dynamics.

Note S4. Ablation study

To clarify the respective roles of the stiffness constraint (SC) and the deep Koopman feedforward compensator (DK-FFC), we performed an ablation study in the path-tracking task, as summarized in Fig. S4A and Table S3. Two intermediate variants are evaluated based on the stiffness-free K-MPC baseline: K-MPC+SC, which includes only the stiffness constraint, and K-MPC+DK-FFC, which includes only the feedforward compensator. Compared with K-MPC, K-MPC+SC substantially improves tracking accuracy, disturbance rejection, and vibration suppression, although with a slower response. By contrast, adding DK-FFC improves tracking accuracy and response speed, but yields weaker disturbance rejection. Notably, the DK-FFC is trained on data collected from trajectories initialized with high antagonistic inputs, and part of its performance gain may therefore reflect stiffness-related characteristics embedded in the training data. This is consistent with the disturbance-rejection test, where at the 20° target, K-MPC+DK-FFC produces a representative input pair of (28, 59) kPa, compared with (33, 80) kPa for K-MPC+SC. The resulting lower input sum corresponds to a lower stiffness level under this operating condition, which helps explain its weaker disturbance rejection. When both components are combined, the full controller achieves the best overall performance across all evaluated metrics. This result supports that the two modules act in a complementary manner, leading to a more balanced control strategy in terms of accuracy, robustness, vibration suppression, and response speed.

Note S5. Robustness analysis under load and contact-induced perturbations

Since the learned model is identified from data collected under nominal conditions, external loading may cause the actual robot response to deviate from the nominal input-deformation relationship captured by the model. In the present framework, this effect is mitigated implicitly through closed-loop feedback and online re-optimization. In addition, stiffness regulation enhances resistance to payload-induced deformation and helps maintain motion stability under load. Similar robustness trends under loading are also observed in the tendon-driven antagonistic robot, where increasing the antagonistic tendon force improves motion accuracy under both tip-load and silicone-oil disturbance conditions. These results support the role of stiffness regulation in preserving motion performance when the actual response deviates from the nominal learned behavior.

In addition to payload-induced deviation, the gripper grasping experiments may also involve contact-induced interactions among fingers that are absent in free-space operation. In the present framework, without additional contact or force sensing, these effects are not modeled in the controller as explicit inter-finger coupling states. Instead, they appear implicitly as local disturbance variations experienced by each finger. The controller remains effective because closed-loop feedback can correct moderate model mismatch online, while higher antagonistic

coactivation increases effective finger stiffness and improves robustness against contact-induced loads.

Note S6. Heuristic stiffness enhancement experiments

To further examine the role of stiffness regulation in PD control, we considered two heuristic stiffening strategies. In the first strategy, one actuator is fixed at a constant input u_c (1%, 2%, and 5% PWM duty cycle), while the opposing actuator is controlled by PD. In the second strategy, both actuators are given the same antagonistic preload u_s (10%, 20%, 50% and 100% PWM duty cycle), while the PD command is applied differentially, with one actuator receiving an additive term and the other a subtractive term to generate motion. The second strategy is introduced to better evaluate the effect of stiffness regulation under a more balanced antagonistic actuation scheme, where both actuators contribute to motion control.

The representative results are shown in Fig. S4, and the quantitative comparison is summarized in Table S4. For the first strategy, introducing a small fixed antagonistic input already improves tracking performance. Compared with the stiffness-free PD controller, u_c = 1% reduces the tracking error from 2.37° to 1.96°, while also lowering the vibration index and improving disturbance rejection, although the response becomes slower. As u_c increases further, vibration and disturbance sensitivity are further reduced, but the tracking accuracy deteriorates and the response becomes progressively slower. For the second strategy, a moderate symmetric antagonistic preload provides a better trade-off. In particular, u_s = 10% improves the tracking accuracy to 1.77°, while also reducing the vibration and improving disturbance rejection compared with the stiffness-free PD controller. As u_s increases further, vibration and disturbance sensitivity are further reduced, but the response becomes slower and the tracking accuracy deteriorates.

These results show that increasing antagonistic input can enhance effective stiffness and improve motion stability under a simple PD framework. However, although heuristic stiffening can suppress vibration and improve disturbance tolerance, it also introduces clear trade-offs in tracking accuracy and response speed, especially when the antagonistic preload becomes excessively large.

Note S7. Model performance under different excitation data coverage conditions

To further examine model applicability beyond the nominal identification setting, we tested Koopman models identified from datasets with different excitation coverage. Relative to the full excitation set, models identified from reduced-amplitude excitation sets (20% and 50%) produce similar closed-loop tracking errors (3.48° and 3.49°, respectively, versus 3.21° for the full set), indicating that moderate amplitude reduction alone does not substantially degrade control performance when the main excitation structure is preserved. By contrast, models identified from a single excitation trajectory generally lead to much larger tracking errors (11.23°, 13.77°, 15.04°, and 14.01°), except for one case (5.00°). This comparatively better case corresponds to a trajectory spanning a broader antagonistic coactivation range, with $u_1 + u_2$ covering approximately 20%-180%, suggesting that the practical boundary of the training distribution is governed more by the richness of the excitation data and the coverage of the relevant state-input region than by amplitude alone. Therefore, re-identification is required when the robot geometry, installation condition, sensor mounting, loading condition, or operating regime introduces response patterns that are insufficiently represented in the original identification dataset.

Note S8. Cross-configuration transfer performance

Preliminary cross-configuration transfer experiments are performed by applying the model identified from a single soft robotic arm to soft robotic systems with different geometric configurations, including the pump and the bridge configurations. As summarized in Table S1, the proposed framework remains effective under cross-configuration transfer and outperforms the other control methods. These results suggest that the learned representation is able to capture shared actuation-deformation features across related embodiments. In practice, the dynamics can still be affected by fabrication variability, differences in initial shape, IMU mounting orientation, robot installation configuration, and gravity-induced loading, which may introduce robot-specific offsets and dynamic variations. Therefore, fast robot-specific identification remains a more reliable and practical choice in the current framework. This requirement does not impose a heavy burden, because model construction is fast, and for multi-robot systems the data acquisition stage can be performed simultaneously, while the offline identification time scales approximately linearly with the number of robots. Future improvements in cross-configuration generalization may be achieved through morphology or parameter-conditioned model structures, or through parameterized Koopman formulations that explicitly account for varying system parameters.

Note S9. Model prediction performance under different training data quantities

The effect of training data quantity on the prediction performance of the identified Koopman model is analyzed using the five excitation trials described in the Methods as the candidate training set. Koopman models are identified using different combinations of these five trials, corresponding to different training data quantities. For each data quantity, all possible combinations of the corresponding number of training trajectories are considered. The identified models are then evaluated on five additional unseen datasets that are not used for model identification. For each model, the 1-step, 5-step, and 10-step prediction errors are computed on these unseen test datasets. The corresponding results are summarized in Fig. S5. As the amount of training data increases, the prediction error generally decreases, with a more pronounced effect for longer prediction horizons.

Note S10. History-dependent behavior in the reconfigurable bridge experiment

The incomplete recovery observed after unloading in the reconfigurable bridge experiment suggests that the soft antagonistic structure exhibits history-dependent behavior, likely arising from the internal friction between the scales and the base layer. Such behavior is common in soft robotic systems, where deformation may depend not only on the current actuation input but also on the recent loading history. This further highlights the importance of dynamic closed-loop control, rather than relying solely on an instantaneous input-deformation relationship. The proposed controller does not assume a unique static mapping between actuator input and bending angle. Instead, it is formulated based on a discrete-time dynamic model and updated in closed loop using real-time state feedback, so that history-dependent effects can be compensated online during path tracking.

Note S11. Extension to different stiffness levels

Although stiffness cannot be directly measured in real time, we can qualitatively modulate the desired stiffness level by adjusting the sum of antagonistic actuation inputs. To illustrate this capability, we conducted additional experiments using the same pneumatic antagonistic robot. In one set of experiments, we commanded the robot to maintain a constant bending angle of 20° and applied external perturbations under different desired stiffness levels (50%, 100%, and 150%), as

shown in Fig. S6A. Under a 20 g external load, the maximum tracking error decreases from -4.11° at 50% antagonistic input to -1.59° at 100% and further to -1.03° at 150%. Similarly, under a 50 g load, the maximum tracking error decreases from -2.10° at 50% to -1.29° at 100% and -0.44° at 150%. These results indicate that increasing the desired stiffness level consistently reduces the maximum tracking error under external disturbances. In addition, we performed sinusoidal tracking experiments at fixed stiffness levels. As shown in Fig. S6B, the robot stably tracks sinusoidal reference trajectories while maintaining stiffness levels corresponding to antagonistic input sums of 50% and 100%, further demonstrating the capability of the proposed framework to support stable motion control across different stiffness settings.

Note S12. Extension to tendon-driven antagonistic robot

The mechanical design of the tendon-driven antagonistic robotic platform is illustrated in Fig. S12. The robot consists of four tendons arranged in an antagonistic configuration, a ring-shaped particle-based structure, and an onboard inertial measurement unit (IMU). Each tendon is actuated independently by a motor, and motor rotations directly control the robot's motion. Unlike antagonistic pneumatic robots, where the actuation inputs are inherently nonnegative, the inputs of tendon-driven antagonistic robots can take both positive and negative values. In such systems, even small motor commands may result in a rapid increase in tendon tension, which can potentially cause mechanical damage to the robotic platform. To ensure safe operation and reliable stiffness modulation, tension sensors are integrated into each tendon. Based on real-time tendon tension measurements, the stiffness-related constraints in the control framework are reformulated as

$$J = \sum_{i=1}^{N_p} \left(\left\| \boldsymbol{s}^{e}_{k+i|k} \right\|^2_{\boldsymbol{Q}} + \left\| \Delta \boldsymbol{u}_{k+i-1|k} \right\|^2_{\boldsymbol{P}} - \left\| \boldsymbol{\kappa} \boldsymbol{F}_{k+i-1|k} - \boldsymbol{F}_d \right\|^2_{\boldsymbol{R}} \right) \tag{S2}$$

where $\boldsymbol{F}$ represents the tension force and $\boldsymbol{F}_d$ denotes the desired tension force.

On this platform, we employed real-time trajectory tracking as a benchmark task to assess the robustness and adaptability of the proposed control strategy across different environmental conditions and stiffness configurations. The reference trajectory was defined as a five-point star.

$$\text{Task: } \begin{cases} \phi = 6\cos\left(\dfrac{\pi}{5}t\right) + 15\cos\left(\dfrac{2\pi}{15}t\right) \\ \theta = 6\cos\left(\dfrac{\pi}{5}t\right) - 15\cos\left(\dfrac{2\pi}{15}t\right) \end{cases} \tag{S3}$$

where ϕ denotes yaw angle and θ represents pitch angle. All experiments were repeated three times, and the control interval was set to 30 ms using the Raspberry Pi 4B platform. The environmental conditions included a disturbance-free condition, a 100-g load, and silicone oil with a viscosity of 100000 cSt. The tracking error for different scenarios under various antagonistic forces is depicted in Table S8.

Note S13. Implementation details of the model predictive control

Here, we show the detailed MPC. The robot state space can be represented by

$$\begin{aligned} \boldsymbol{\Psi}_{k+1} &= \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{0} & \boldsymbol{I} \end{bmatrix} \boldsymbol{\Psi}_k + \begin{bmatrix} \boldsymbol{B} \\ \boldsymbol{I} \end{bmatrix} \Delta \boldsymbol{u}_k \\ \boldsymbol{s}_{k+1} &= \boldsymbol{D} \boldsymbol{\Psi}_{k+1} \end{aligned} \tag{S4}$$

where $\boldsymbol{D} = \mathrm{diag}(\boldsymbol{D}_1, \ldots, \boldsymbol{D}_n)$, and $\boldsymbol{D}_i$ is a row unit vector. After N_p times iteration, the prediction state can be calculated by

$$\boldsymbol{\Xi}_k = \boldsymbol{H}\boldsymbol{\Psi}_k + \boldsymbol{L}\Delta\boldsymbol{U}_k \tag{S5}$$

where $\boldsymbol{\Xi}_k = [\boldsymbol{s}_{k+1|k}^T, \boldsymbol{s}_{k+2|k}^T, \ldots, \boldsymbol{s}_{k+N_p|k}^T]^T$, $\Delta\boldsymbol{U}_k = [\Delta\boldsymbol{u}_{k|k}^T, \Delta\boldsymbol{u}_{k+1|k}^T, \ldots, \Delta\boldsymbol{u}_{k+N_p-1|k}^T]^T$,

$$\begin{aligned}
\boldsymbol{H} &= \begin{bmatrix} \boldsymbol{D}\tilde{\boldsymbol{A}} \\ \cdots \\ \boldsymbol{D}\tilde{\boldsymbol{A}}^{N_p} \end{bmatrix} \\
\boldsymbol{L} &= \begin{bmatrix} \boldsymbol{D}\tilde{\boldsymbol{B}} & \boldsymbol{0} & \cdots & \boldsymbol{0} \\ \cdots & \cdots & \ddots & \cdots \\ \boldsymbol{D}\tilde{\boldsymbol{A}}^{N_p-1}\tilde{\boldsymbol{B}} & \boldsymbol{D}\tilde{\boldsymbol{A}}^{N_p-2}\tilde{\boldsymbol{B}} & \cdots & \boldsymbol{D}\tilde{\boldsymbol{B}} \end{bmatrix} \\
\tilde{\boldsymbol{A}} &= \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{0} & \boldsymbol{I} \end{bmatrix} \\
\tilde{\boldsymbol{B}} &= \begin{bmatrix} \boldsymbol{B} \\ \boldsymbol{I} \end{bmatrix}
\end{aligned} \tag{S6}$$

Then, the cost function can be simplified as a quadratic programming

$$\begin{aligned}
\Delta\boldsymbol{U}_k &= \arg\min_{\Delta\boldsymbol{U}_k} J \\
&= \arg\min_{\Delta\boldsymbol{U}_k} \left(\frac{1}{2}\Delta\boldsymbol{U}_k^T \tilde{\boldsymbol{L}}\Delta\boldsymbol{U}_k + \tilde{\boldsymbol{H}}_k\Delta\boldsymbol{U}_k \right)
\end{aligned} \tag{S7}$$

where $\tilde{\boldsymbol{L}} = 2(\boldsymbol{L}^T\tilde{\boldsymbol{Q}}\boldsymbol{L} + \tilde{\boldsymbol{P}} - \boldsymbol{\Theta}^T\tilde{\boldsymbol{R}}\boldsymbol{\Theta})$, $\boldsymbol{H}_k = 2(\boldsymbol{H}\boldsymbol{\Psi}_k - \boldsymbol{\Xi}_r)^T\tilde{\boldsymbol{Q}}\boldsymbol{L} - 2(\boldsymbol{u}_{k-1}^T\boldsymbol{\kappa}^T)\otimes\boldsymbol{1}_{1\times Np}\tilde{\boldsymbol{R}}\boldsymbol{\Theta}$, $\tilde{\boldsymbol{P}} = \mathrm{diag}(\boldsymbol{P}, \ldots, \boldsymbol{P})$, $\tilde{\boldsymbol{Q}} = \mathrm{diag}(\boldsymbol{Q}, \ldots, \boldsymbol{Q})$, $\tilde{\boldsymbol{R}} = \mathrm{diag}(\boldsymbol{R}, \ldots, \boldsymbol{R})$, $\boldsymbol{\Theta} = \boldsymbol{E}\otimes\boldsymbol{\kappa}$, $\boldsymbol{E}$ is a unit lower triangular matrix. Thus, we can get optimal robot input after solving the quadratic programming function

$$\begin{aligned}
\boldsymbol{u}_k &= \boldsymbol{u}_{k-1} + \Delta\boldsymbol{u}_{k|k} \\
\boldsymbol{u}_b &= \boldsymbol{u}_k
\end{aligned} \tag{S8}$$

Note S14. Mechanism underlying the stiffness variation

To analyze the stiffness variation observed in Fig. 1B, we examined the FE simulation results in detail, focusing on the contact forces experienced by the base layer in the antagonistic robot. Specifically, we compared two distinct actuation scenarios with similar bending angles but drastically different negative pressure levels. The color gradient in Fig. S15 represents the magnitude of friction force, with light orange indicating higher friction values. The results reveal significant differences in contact mechanics between these two actuation modes. When comparing unilateral actuation to bilateral actuation, the latter generates much greater friction forces at the contact interfaces. This is particularly evident at three key locations labeled as points 1, 2, and 3 along the structure, where bilateral actuation produces approximately four times higher friction forces compared to unilateral actuation, as shown in the inset plots in Fig. S15. This enhanced friction under bilateral actuation engage the scales more firmly with the base layer, which directly contributes to the increased bending stiffness observed in the robot. Moreover, under bilateral actuation, the density of the contact points between the scales and the base layer are higher, which also contributes to the enhanced stiffness, similar to the stiffening mechanism observed in the variable stiffness chainmail structure.

Note S15. Design rationale for the antagonistic robot

This robot design is selected due to its dual capability of achieving controlled deformation and stiffness variation using a single pneumatic input for each actuator chamber. When negative pressure is applied to either chamber, the programmed surface patterns on the scales engage with those on the base layer, causing the robot to morph to a predefined shape while concurrently increasing stiffness. This dual functionality is achieved through the same actuation mechanism, significantly simplifying the control architecture and eliminating the need for separate actuators or complex mechanical systems. In contrast, other variable stiffness mechanisms, such as geometric locking mechanisms or jamming mechanisms, while effective for stiffness modulation, typically provide only stiffness variation without inherent shape-changing capabilities. These approaches would require integration with additional actuators capable of active shape control to achieve the multifunctionality demonstrated in our work. Another advantage of our scale design is its ability to achieve continuous shape change and continuous stiffness variation, since the working mechanism relies not solely on geometric coupling but also on friction forces. In contrast, other geometric locking mechanisms can only lock into a limited number of discrete positions, with stiffness essentially switching between high and low states without intermediate levels.

Table S1. Comparison of multiple performance metrics for different control methods.

Metrics	Open-loop	PD	K-MPC	RL	**Our work**
Tracking accuracy (°)	3.09	2.37	3.21	5.25	**1.58**
Disturbance rejection (°)	6.95	6.46	2.59	6.56	**1.11**
Vibration suppression (°/s)	0.25	2.09	3.45	1.44	**0.43**
Response speed (s)	5.48	3.67	3.25	12.40	**2.09**
Learning speed (min)	NA	NA	4.83	202.78	**7.67**
Computational cost (ms)	0.24	0.51	11.34	12.24	**11.61**
Cross-configuration tracking accuracy (°)	6.44	2.61	3.34	6.26	**2.00**
Cross-configuration vibration suppression (°/s)	0.33	3.03	2.59	0.94	**0.46**

Table S2. Comparison of implementation requirements and reported performance metrics for representative stiffness-aware antagonistic control methods.

Metrics	Bruder (*17*)	Gillespie (*25*)	Yi (*27*)	Best (*45*)	**Our work**
Analytical model required	Yes	Yes	Yes	Yes	**No**
Explicit stiffness model required	Yes	Yes	Yes	Yes	**No**
Explicit constraint handling	No	Yes	No	Yes	**Yes**
Deployment effort	Re-model	Re-model	Re-model	Re-model	**Re-sample**
Tracking accuracy (°)*	†	‡	0.04-0.43§	4.14	**1.58**
Disturbance rejection (°)*	†	†	‡	†	**1.11**
Response speed (s)*	†	0.56	‡	‡	**2.09**

* Reported performance metrics are taken from the corresponding studies and are not intended as direct head-to-head comparisons because the methods were evaluated on different platforms and tasks.

† Not reported.

‡ Qualitative performance was demonstrated graphically, but no numerical metric was reported.

§ Steady-state tracking errors reported over the interval [4, 8] s under different control parameters.

Table S3. Comparison of multiple performance metrics for different variants of the proposed controller.

Methods	Tracking accuracy	Disturbance rejection	Vibration suppression	Response speed
K-MPC	3.21°	2.59°	3.45°/s	3.25 s
K-MPC+SC	1.94°	1.13°	0.43°/s	4.98 s
K-MPC+DK-FF	1.68°	2.12°	0.43°/s	2.80 s
Our work	**1.58°**	**1.11°**	**0.43°/s**	**2.09 s**

Table S4. Comparison of multiple performance metrics for different heuristic PD-based controller.

Methods	Tracking accuracy	Disturbance rejection	Vibration suppression	Response speed
PD	2.37°	6.46°	2.09°/s	3.67 s
PD (u_c = 1%)	1.96°	5.31°	1.66°/s	5.14 s
PD (u_c = 2%)	3.11°	3.86°	0.67°/s	7.62 s
PD (u_c = 5%)	5.25°	3.02°	0.34°/s	10.74 s
PD (u_s = 10%)	1.77°	3.23°	0.47°/s	5.28 s
PD (u_s = 20%)	1.97°	2.78°	0.33°/s	5.82 s
PD (u_s = 50%)	3.09°	2.03°	0.32°/s	6.13 s
PD (u_s = 100%)	6.04°	1.15°	0.30°/s	7.05 s
Our work	**1.58°**	**1.11°**	**0.43°/s**	**2.09 s**

Table S5. Shape and weight parameters of the objects.

Object	Weight (g)	Characteristic size (mm)
Sponge	1	49 × 71
Ping pong ball	3	40
Transparent box	10	103
Pill box	12	35.5
Polylactic acid ring	13	75
Masking tape	25	85
Toy bear	27	55.5
Silicone ring	28	70
Strawberry	32	40.5
Cake	34	50 × 67
Mangosteen	47	47.5
Tennis ball	58	65
Toolbox	76	60 × 110
Polylactic acid sculpture	101	147
Polyimide tape	107	87
Apple	205	79
Glass cup	126	81
Beaker (small)	47	69
Beaker (large)	69	65

Table S6. Parameters of the antagonistic robot used in different applications.

Application	Target shape	Target radius (mm)	Length (mm)
Gripper / stiffness sensor	Arc	31.83	66.67
Liquid pump	Arc	63.66	133.33
Reconfigurable bridge	Arc	63.66	133.33

Table S7. Parameters of the FE model.

Parameter	Value
Elastic modulus of PLA	2681 MPa
Poisson's ratio of PLA	0.48
Density of PLA	1.2 g/cm^3
Elastic modulus of latex	1.1 MPa
Poisson's ratio of latex	0.48
Density of latex	1 g/cm^3
Coefficient of friction between PTFE	0.09
Coefficient of friction between PLA and latex	0.87
Coefficient of friction between latex	1

Table S8. The tracking error for different scenarios under various antagonistic forces.

Scenarios	Desired antagonistic force (N)	Position (°)	Yaw velocity (°/s)	Pitch velocity (°/s)
Disturbance free	0	0.260±0.014	2.601±0.131	1.823±0.071
	10	0.107±0.002	0.934±0.052	0.895±0.069
	40	**0.100±0.002**	**0.824±0.044**	**0.758±0.026**
Tip load	0	1.799±0.036	23.634±0.686	9.709±0.966
	10	0.176±0.027	1.203±0.123	1.272±0.125
	40	**0.101±0.002**	**0.908±0.050**	**0.886±0.079**
Silicone oil	0	0.610±0.015	1.616±0.042	2.081±0.103
	10	0.147±0.029	0.769±0.066	0.868±0.063
	40	**0.106±0.003**	**0.757±0.050**	**0.677±0.048**

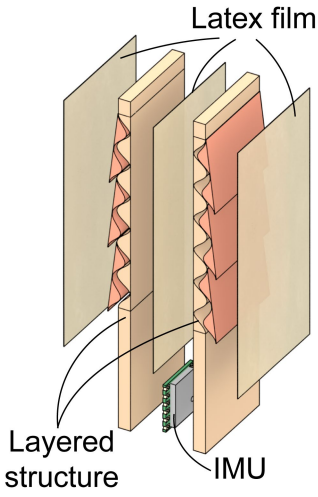


Fig. S1. Exploded view of the antagonistic soft robot, showing two units of the layered structure separated by latex films. An IMU is placed at the end of the robot to record its bending angle.

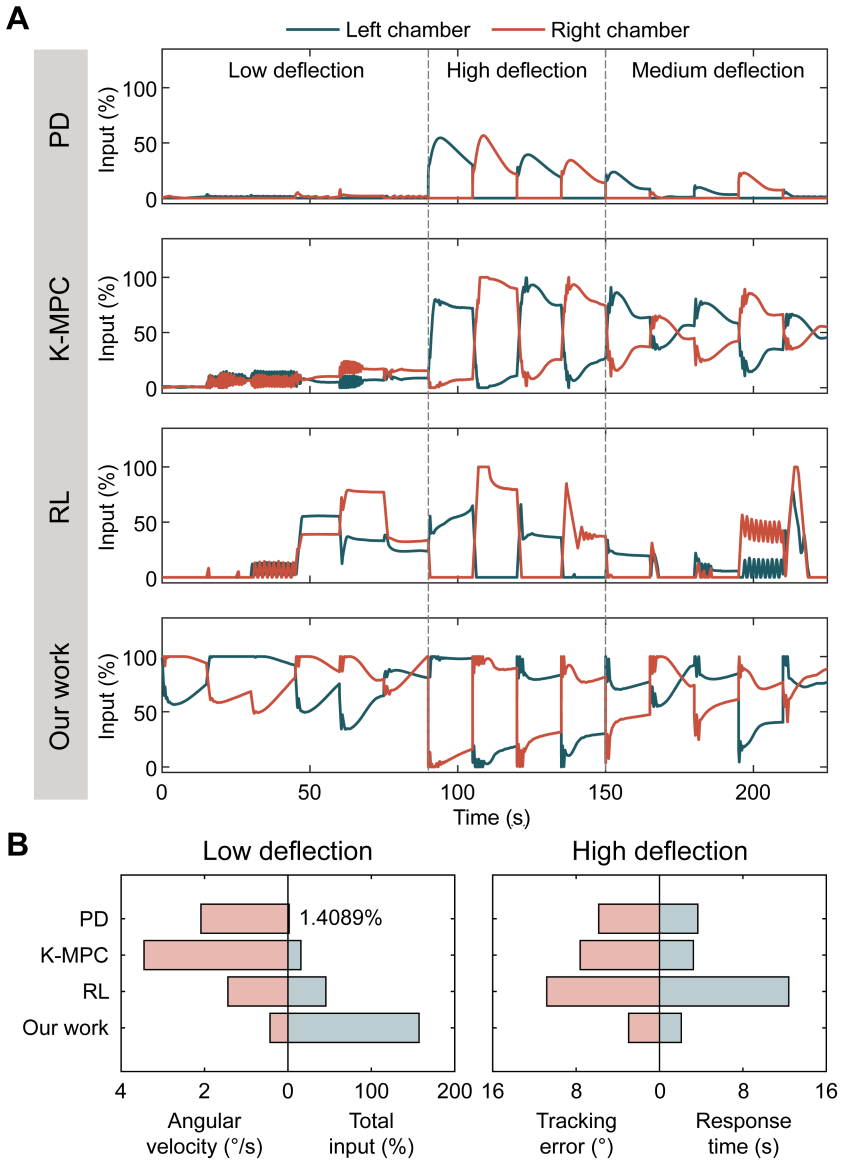


Fig. S2. Real-time path tracking experiments of the antagonistic soft robot. (A) Input comparison. The input comparison of the real-time path tracking experiments. **(B)** Performance comparison under different scenarios.

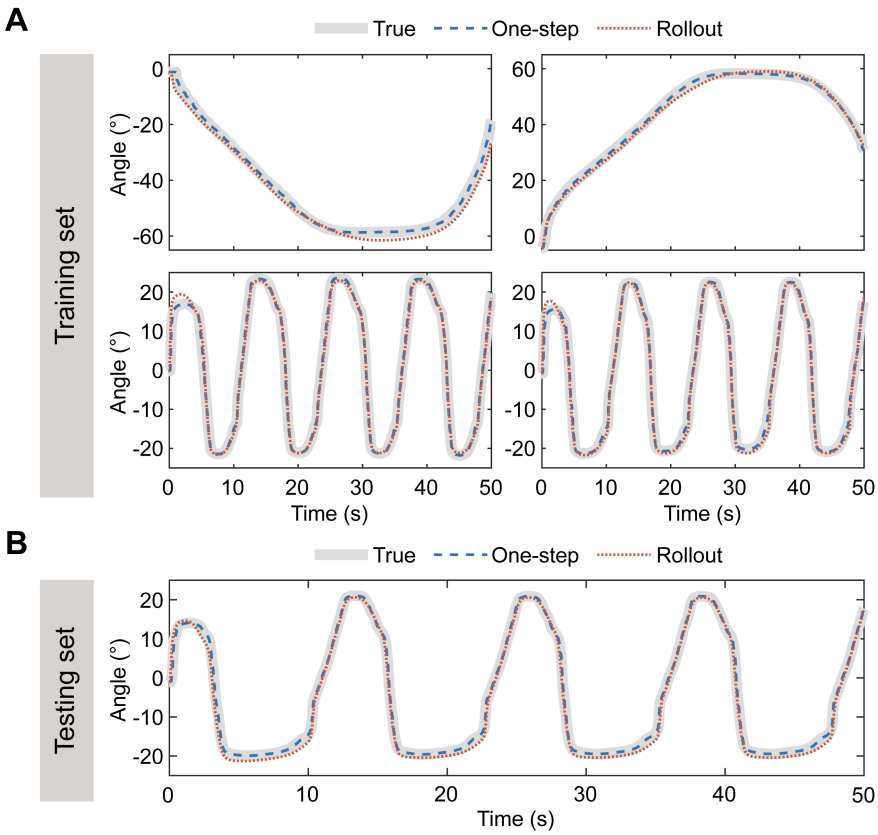


Fig. S3. Prediction performance of the identified MLP dynamics model on training and testing trajectories. (A) One-step and rollout prediction results of the MLP dynamics model on training trajectories. (B) Prediction performance on the testing trajectory.

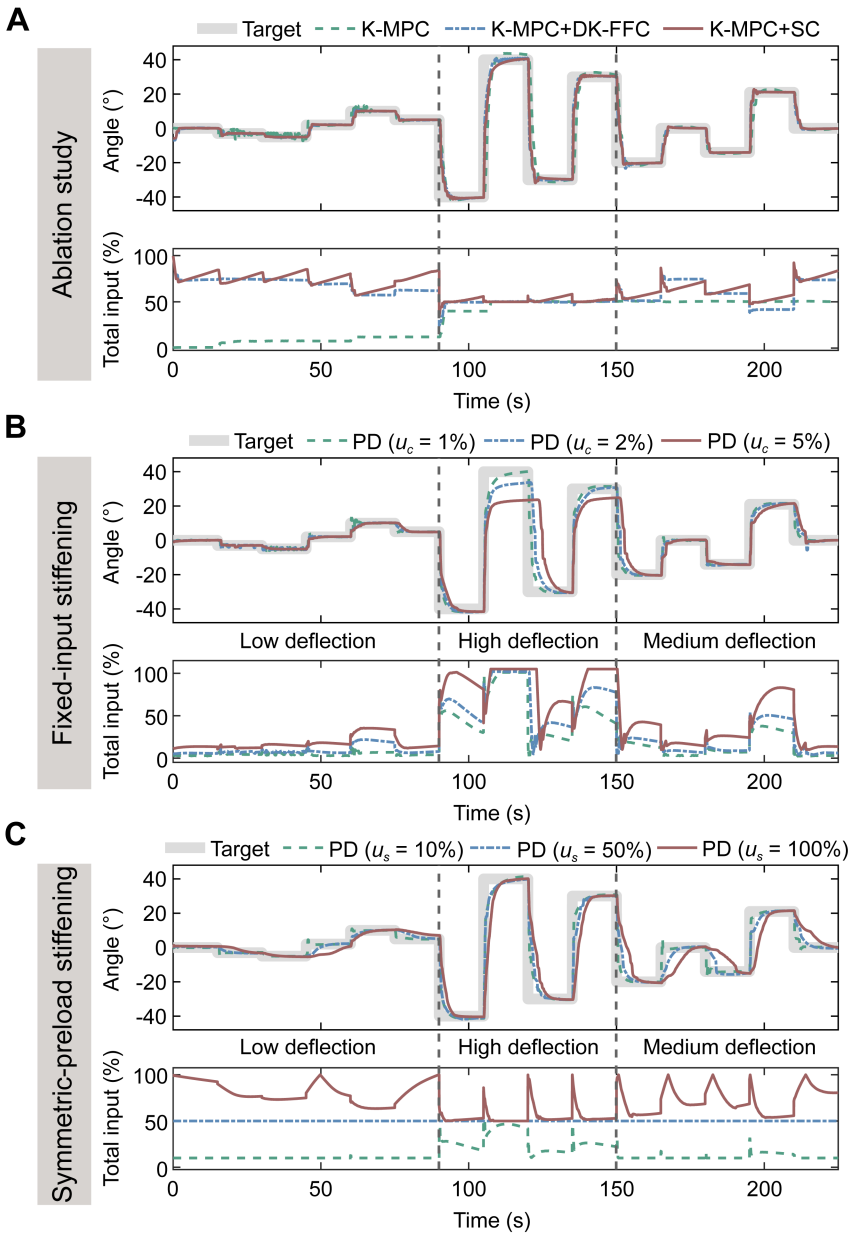


Fig. S4. Real-time path-tracking experiments of the antagonistic soft robot under ablation variants and different stiffening strategies. (A) Ablation study. **(B)** Fixed-input stiffening. **(C)** Symmetric-preload stiffening.

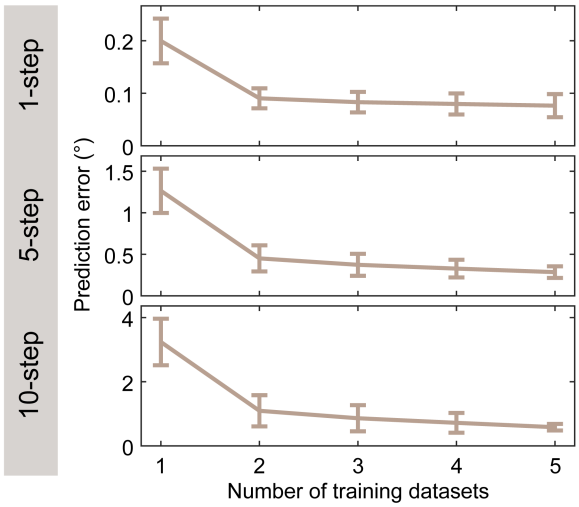


Fig. S5. Effect of training data quantity on Koopman model prediction performance.

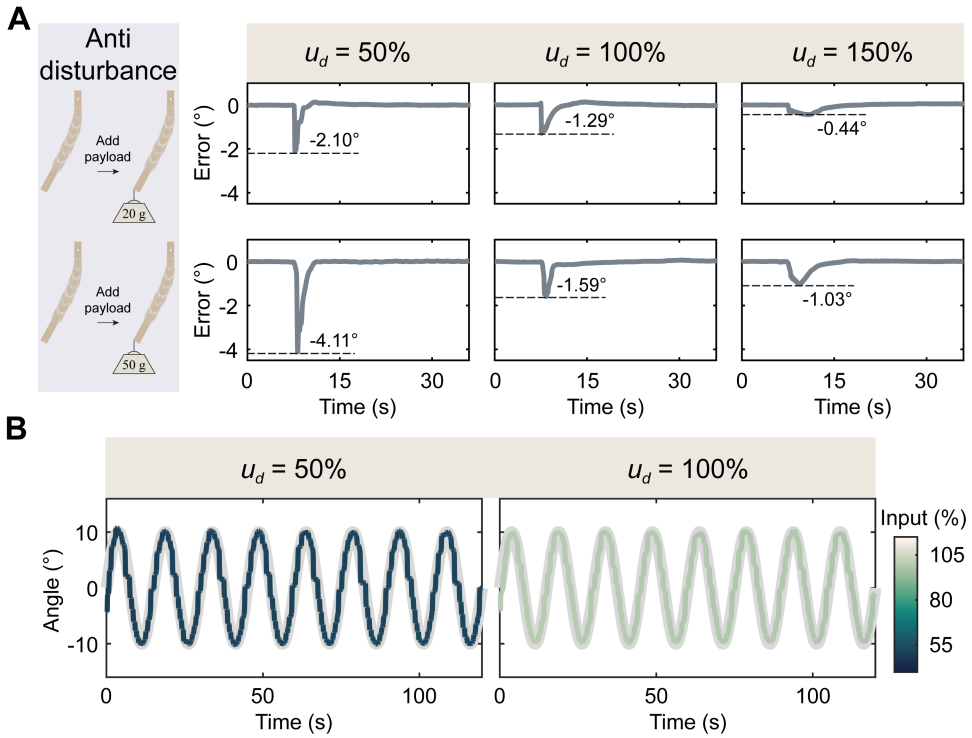


Fig. S6. Tracking performance under different stiffness levels. (A) Anti-disturbance test. Tracking performance under external loads of 20 g and 50 g while maintaining the bending angle at 20°. **(B)** Sinusoidal angle tracking under different stiffness levels. The color represents the sum of antagonistic inputs.

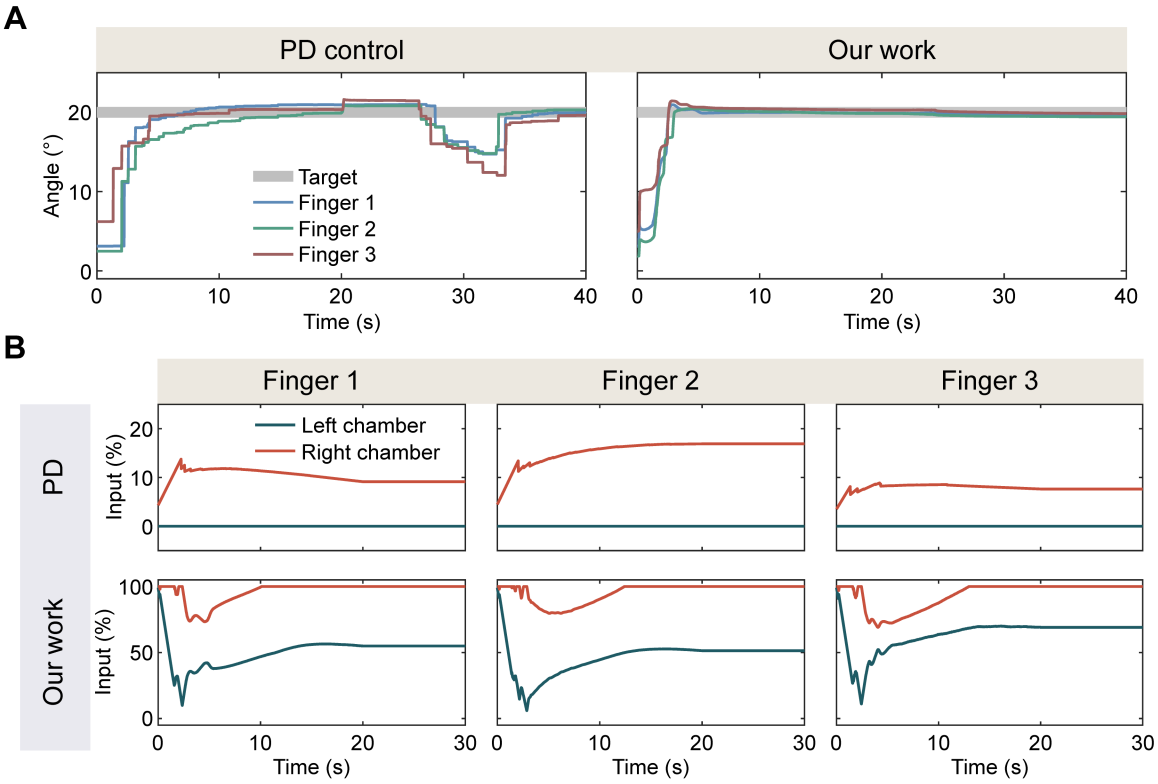


Fig. S7. Comparison of powering rigid object between PD control with our work. (A) The angle variation of each finger. **(B)** The input of each finger.

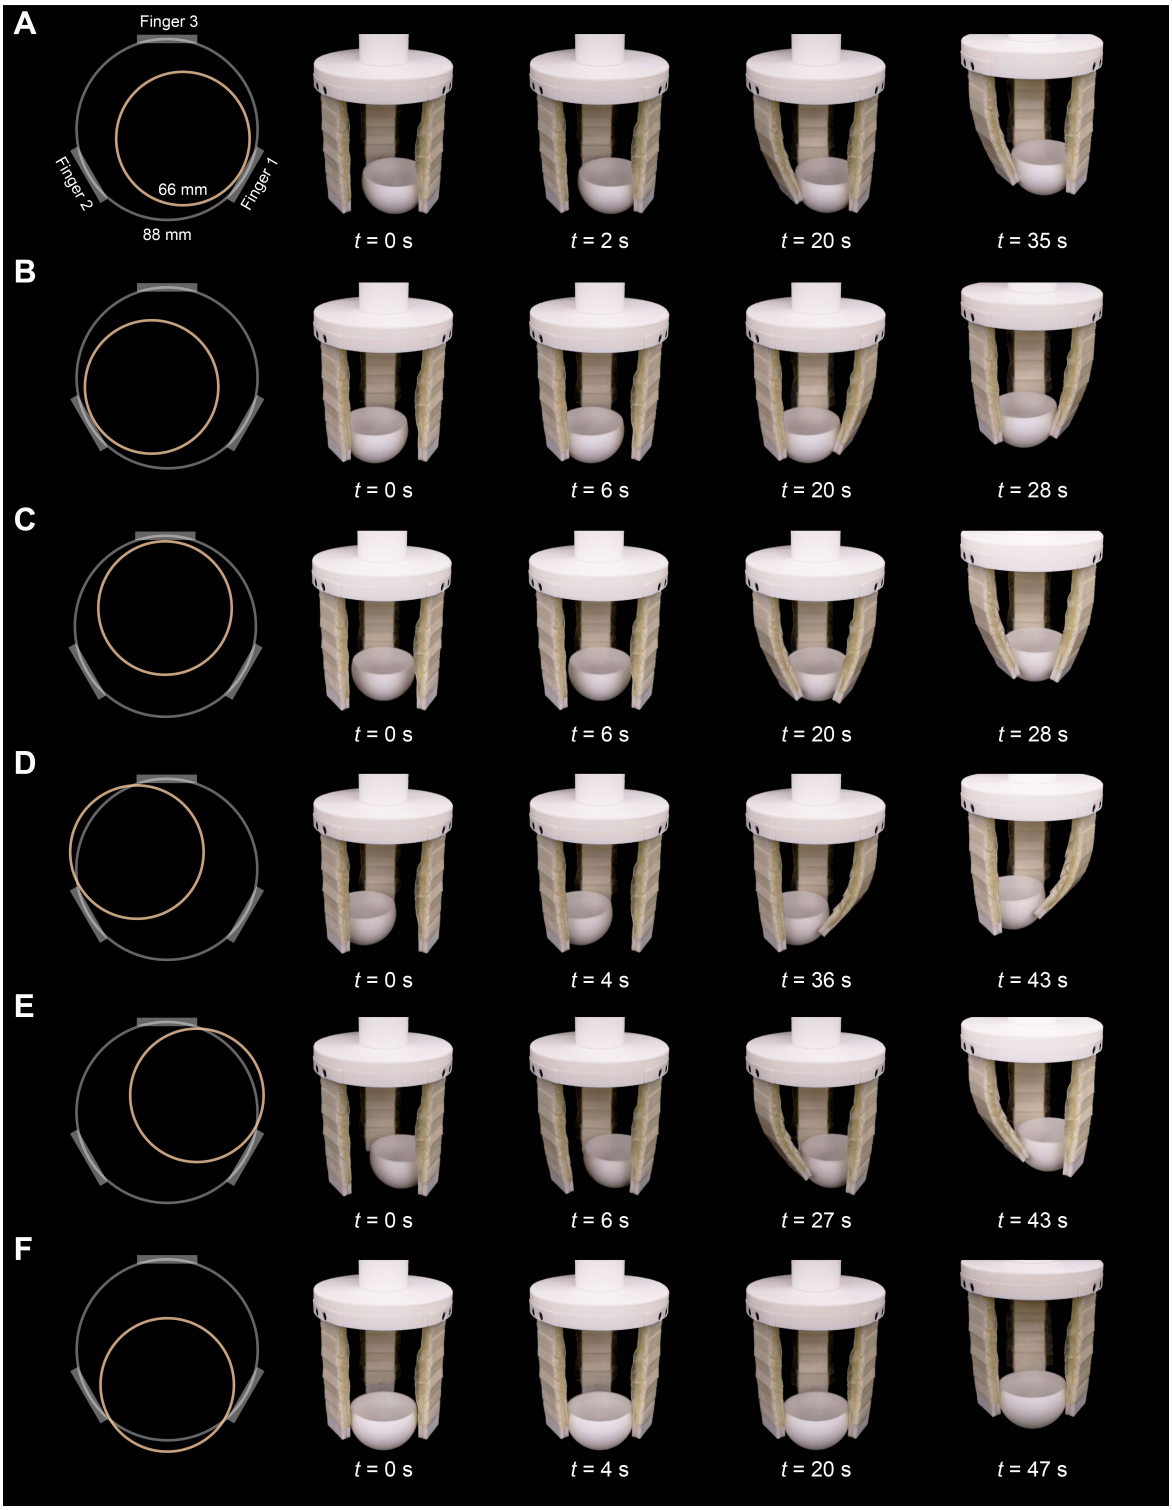


Fig. S8. Positioning and grasping details of the cup-shaped object during pinching. (A) The cup near finger 1. **(B)** The cup near finger 2. **(C)** The cup near finger 3. **(D)** The cup between fingers 2 and 3. **(E)** The cup between fingers 1 and 3. **(F)** The cup between fingers 1 and 2.

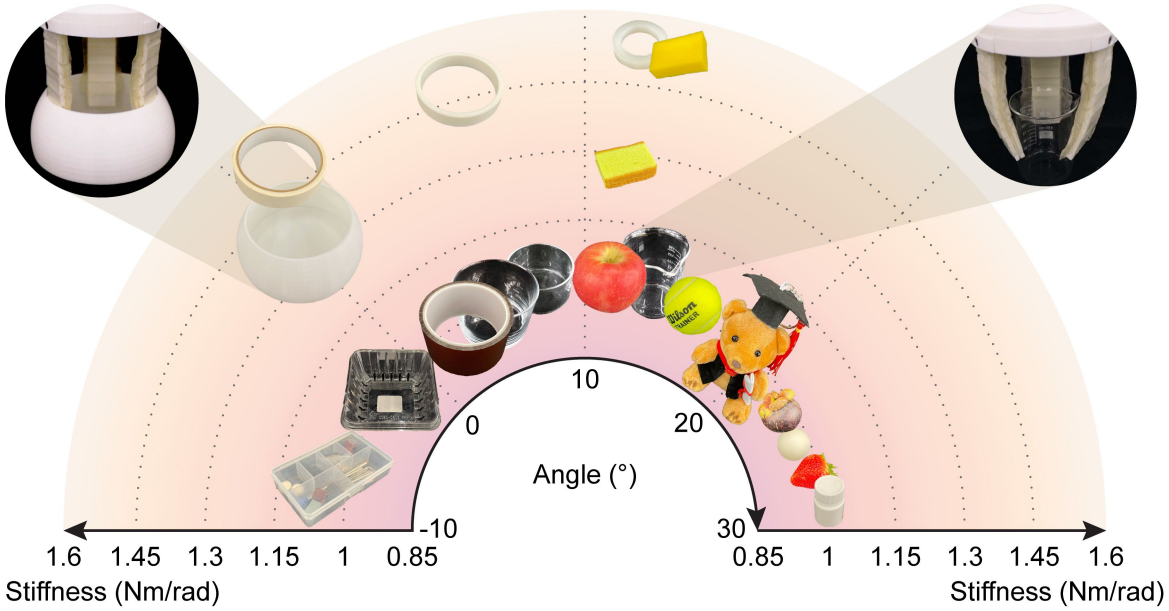


Fig. S9. Universality of the gripper. The gripper can handle various objects with different shapes, stiffness, sizes, and weights under the stiffness-based control strategy. The angle and stiffness represent the average angle and stiffness of the three fingers when grasping the object. All objects and gripper images shown in this figure are original photographs taken by the authors.

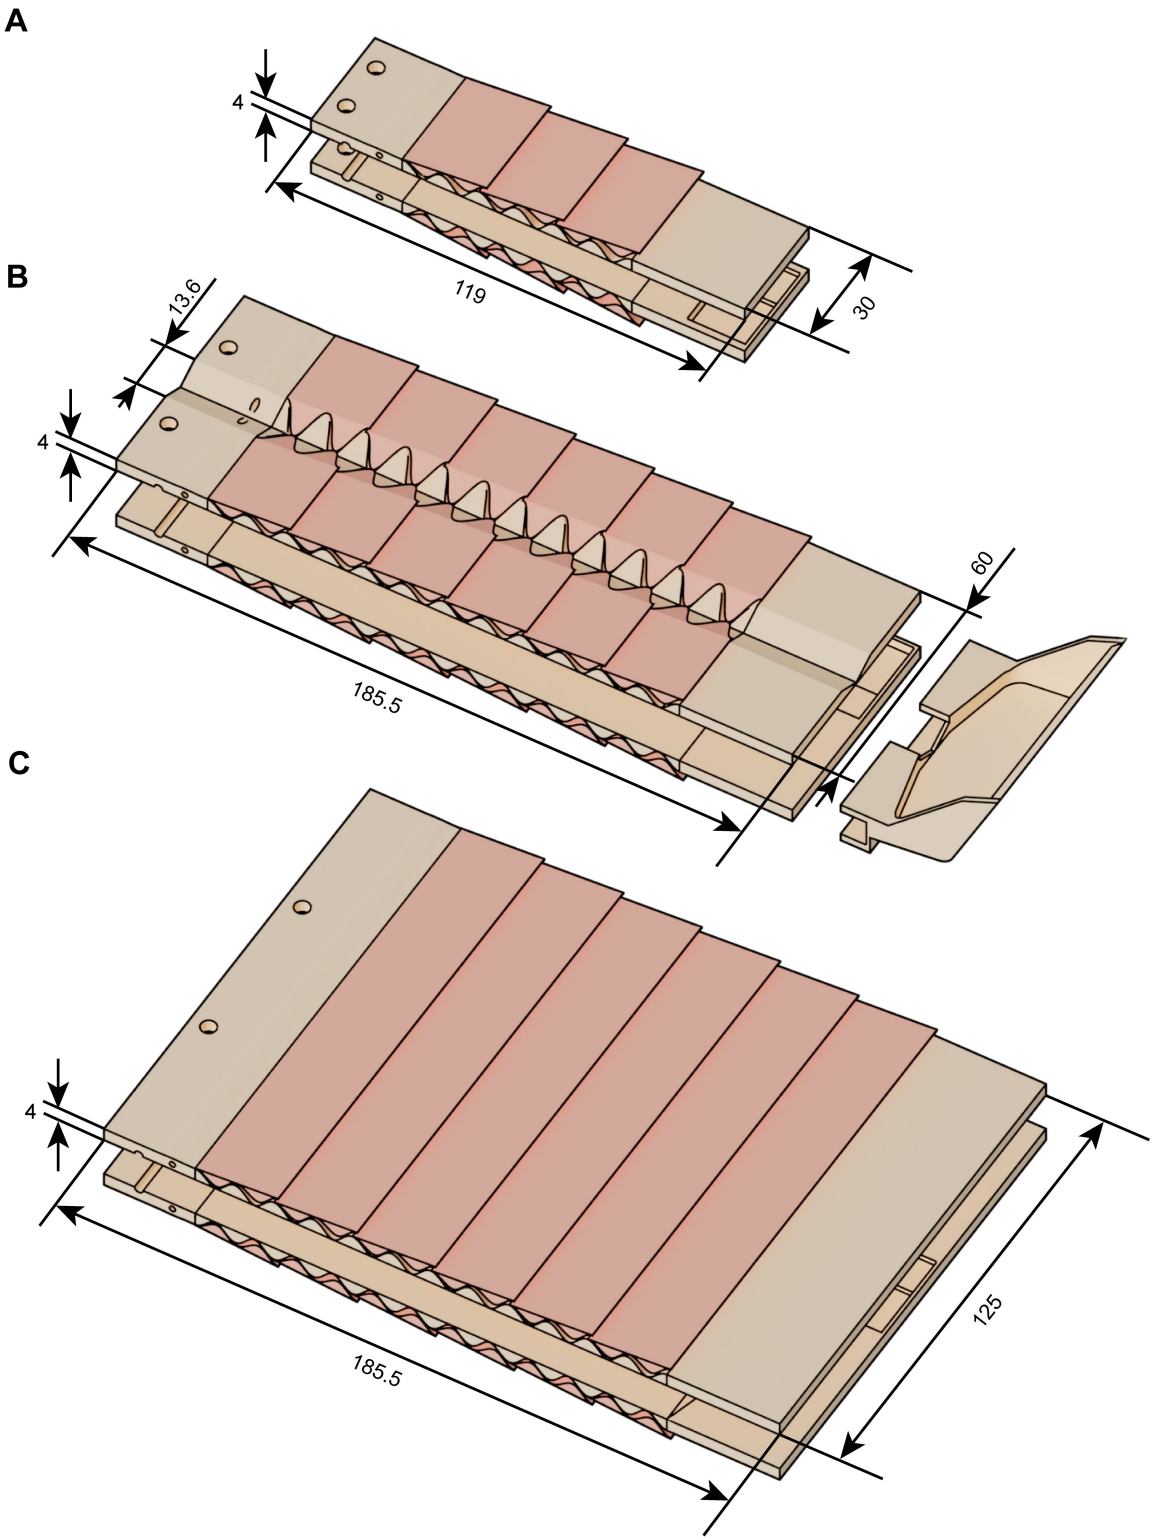


Fig. S10. Design parameters of antagonistic actuators for different applications. (A) Gripper and stiffness sensor. **(B)** Liquid pump. **(C)** Reconfigurable bridge. The base layer and the scales of the robot in each application are colored with yellow and red, respectively. Unit: mm.

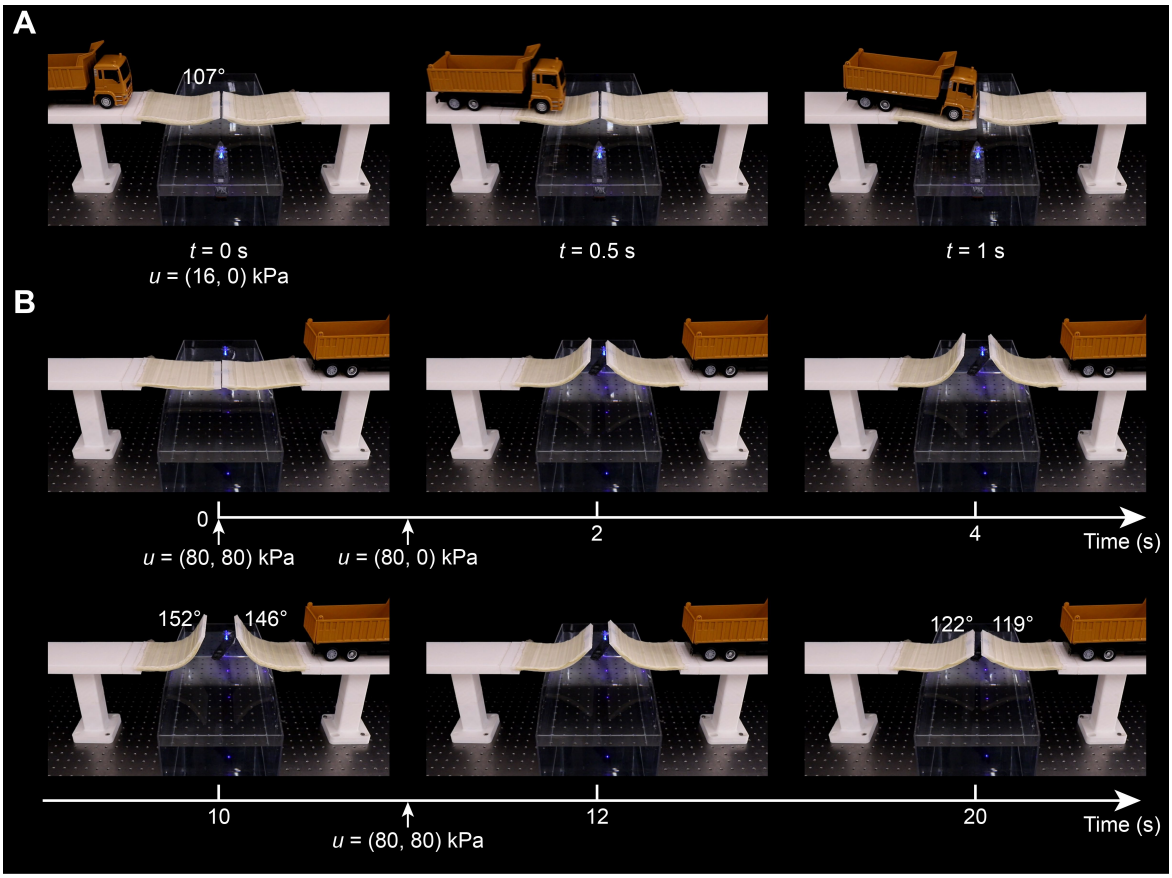


Fig. S11. Challenges of the reconfigurable bridge without control. (A) The bridge will sink into the water when the car passes without control. The input is (16, 0) kPa and the angle of the bridge tip is 107°. **(B)** The bridge initially holds horizontally under the input (80, 80) kPa. Then, the input turns to (80, 0) kPa to let the ship pass. After the ship passes, the bridge cannot recover to the horizontal position by the initial input (80, 80) kPa.

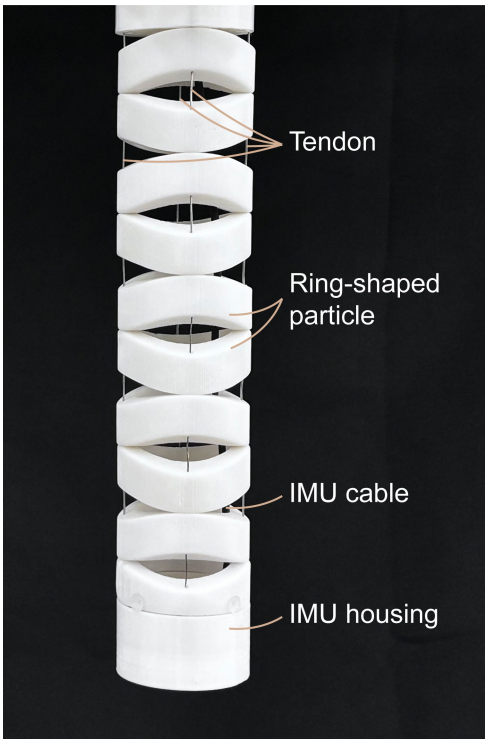


Fig. S12. Mechanical design of the tendon-driven antagonistic soft robot.

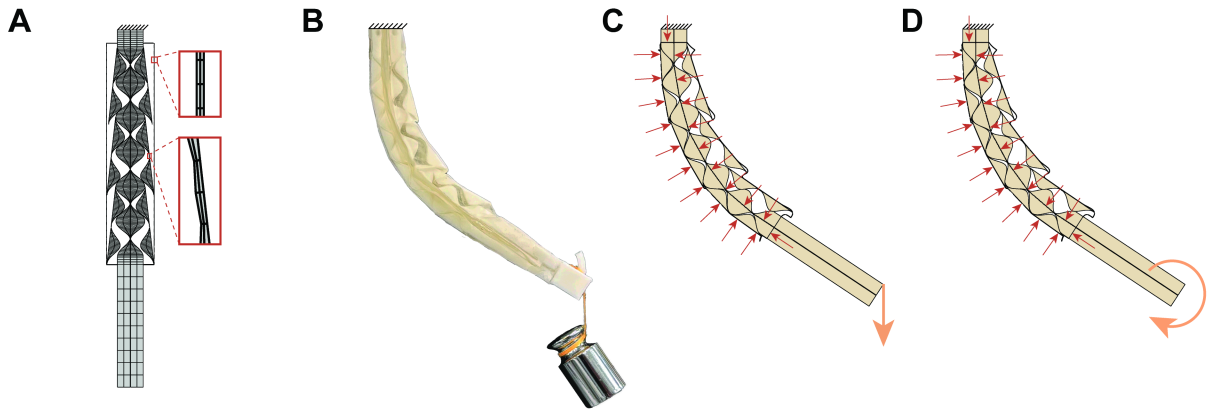


Fig. S13. Experimental setup and the corresponding finite element model. (A) Mesh configuration of the antagonistic soft robot showing the detailed finite element discretization with quadrilateral elements. **(B)** Experimental setup for finite element model validation under an added 30 g payload condition. **(C)** Model validation setup where negative pressure (red arrows) is applied while a 30 g vertical load (orange arrow) is applied at the robot tip to validate the FE model against experimental results. **(D)** Bending stiffness measurement setup where negative pressure is applied (red arrows) and a moment (orange circular arrow) is applied at the tip to determine the bending stiffness of the actuator.

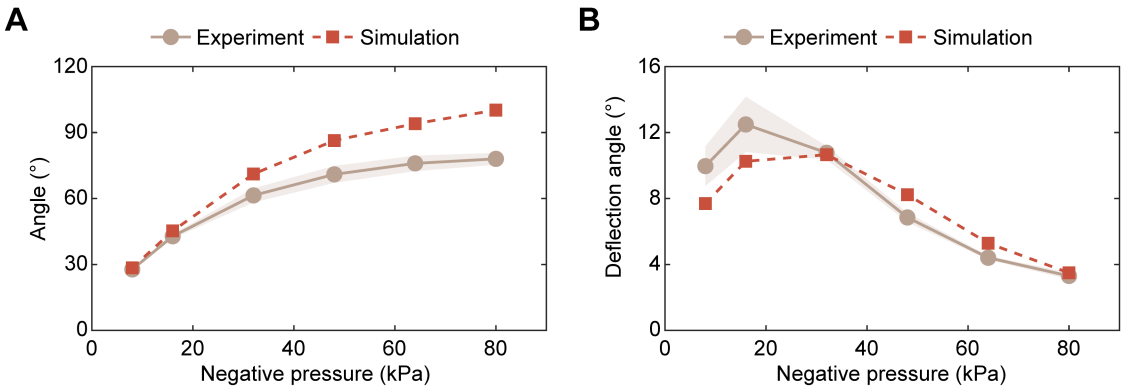


Fig. S14. Comparison of experimental and simulation results for the antagonistic actuator. **(A)** Bending angle θ of the actuator as a function of applied negative pressure in a single chamber. **(B)** Deflection angle $\Delta\theta_F$ caused by a 30 g-load at the end of the actuator, plotted against the applied negative pressure. The shaded areas represent standard deviations.

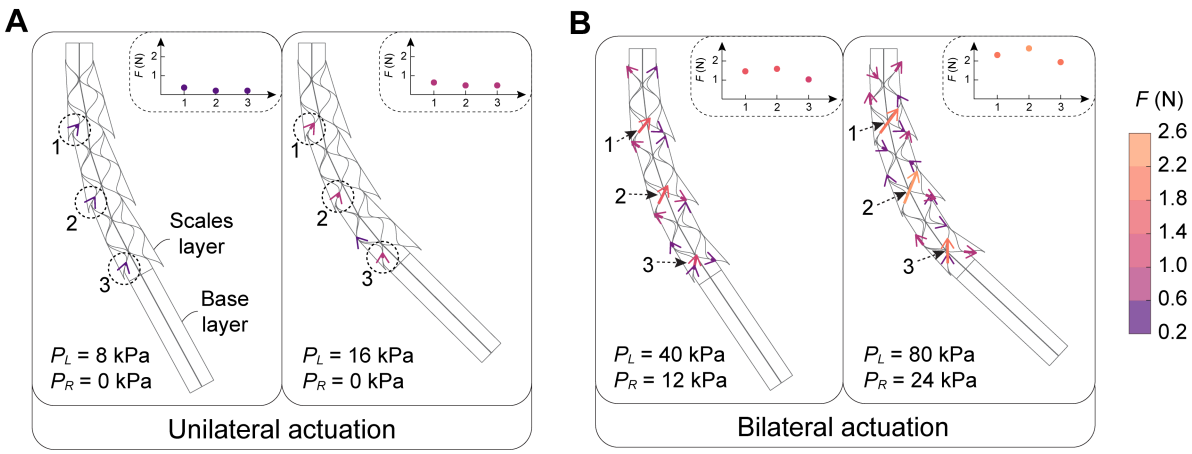


Fig. S15. The friction force distribution on the base layer. (A) Unilateral actuation. **(B)** Bilateral actuation.

Legends for Supplementary Movies S1 to S4

Movie S1. Path Tracking Experiment for Antagonistic Soft Robots

Movie S2. Versatile Grasping with an Antagonistic Soft Gripper

Movie S3. Versatile Applications Enabled by Stiffness-Driven Control

Movie S4. Generality Validation on a Tendon-Driven Robot