

A Soft Wrist with Anisotropic and Selectable Stiffness for Robust Robot Learning in Contact-rich Manipulation

Steven Oh^{1*}, Tomoya Takahashi^{1*}, Cristian C. Beltran-Hernandez¹, Yuki Kuroda¹, and Masashi Hamaya¹

Abstract— Contact-rich manipulation tasks in unstructured environments pose significant robustness challenges for robot learning, where unexpected collisions can cause damage and hinder policy acquisition. Existing soft end-effectors face fundamental limitations: they either provide a limited deformation range, lack directional stiffness control, or require complex actuation systems that compromise practicality. This study introduces CLAW (Compliant Leaf-spring Anisotropic soft Wrist), a novel soft wrist mechanism that addresses these limitations through a simple yet effective design using two orthogonal leaf springs and rotary joints with a locking mechanism. CLAW provides large 6-degree-of-freedom deformation (40mm lateral, 20mm vertical), anisotropic stiffness that is tunable across three distinct modes, while maintaining lightweight construction (330g) at low cost ($\sim \$550$). Experimental evaluations using imitation learning demonstrate that CLAW achieves 76% success rate in benchmark peg-insertion tasks, outperforming both the Fin Ray gripper (43%) and rigid gripper alternatives (36%). CLAW successfully handles diverse contact-rich scenarios, including precision assembly with tight tolerances and delicate object manipulation, demonstrating its potential to enable robust robot learning in contact-rich domains. Project pages and CAD files will be released upon acceptance.

I. INTRODUCTION

Autonomous manipulation in real-world environments represents one of robotics' most significant challenges, particularly when robots must perform contact-rich tasks such as assembly, insertion, and surface interaction. The increasing deployment of robots beyond structured factory settings into unstructured environments, such as homes, laboratories, and collaborative workspaces, has intensified the need for manipulation systems that can robustly handle unexpected contacts due to environmental variations and uncertainties.

Recent learning approaches, including imitation learning [1], [2] and reinforcement learning [3], show considerable promise for enabling robots to acquire complex manipulation skills in environmental variations. However, these paradigms face challenges of data collection and policy execution in uncertain scenarios. During learning, traditional rigid robots, while precise, can pose substantial risks; for example, partially trained policies or human demonstrators may issue control commands that lead to hazardous collisions, potentially damaging both the robot and its surroundings.

Existing approaches to address this challenge fall into two main categories. Force-controlled robots [4] provide one solution by enabling compliant interaction, but require sophisticated control algorithms and expensive hardware that

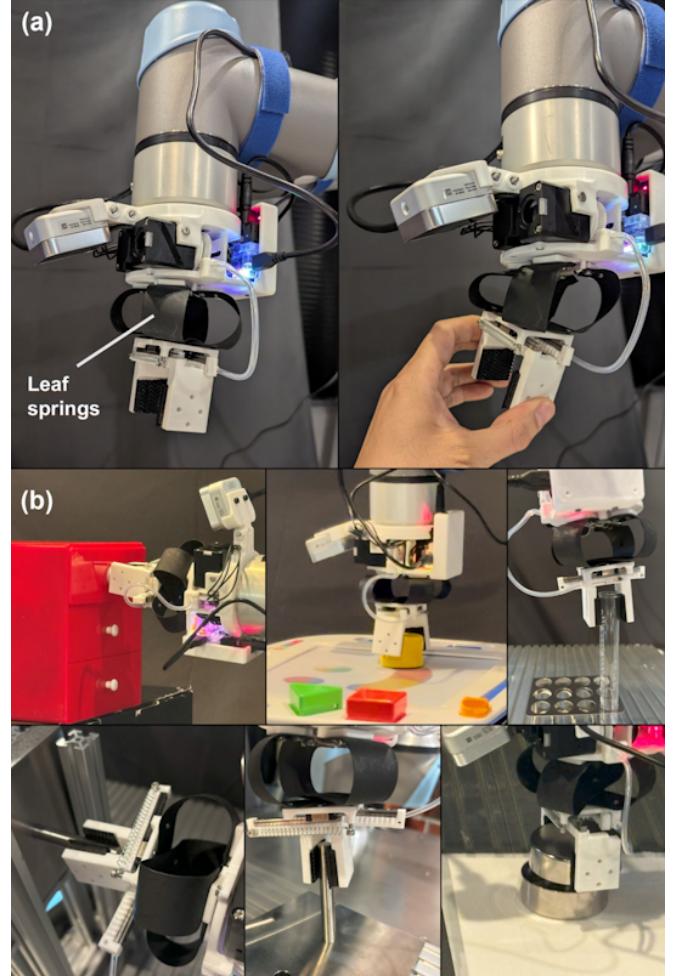


Fig. 1. The CLAW soft wrist: (a) A novel soft wrist featuring two orthogonal leaf springs to achieve anisotropic stiffness and enabling 6-DoF compliance. (b) CLAW's extensive evaluation across a range of challenging contact-rich manipulation tasks using imitation learning.

limit their accessibility. Alternatively, physically soft end-effectors can be integrated with standard position-controlled robots to provide passive safety through mechanical compliance [5], [6], [7]. Despite significant advances in soft robotics, existing soft end-effectors struggle with satisfying the following requirements:

a) Large and multi-directional deformation:

Contact-rich tasks inevitably involve unexpected collisions. To absorb such impacts, the end-effector must be able to deform in 6-degree-of-freedom (6-DoF), and a larger deformation range enables it to accommodate greater uncertainties.

*Equal contributions

¹ OMRON SINIC X Corp., Bunkyo-ku, Tokyo, 113-0033, Japan.
masashi.hamaya@sinicx.com

b) Anisotropic and adjustable stiffness: Direction-dependent stiffness is desirable. For instance, low vertical stiffness facilitates gentle surface contact, whereas moderate horizontal stiffness both absorbs lateral collisions and prevents the end-effector from sagging when oriented horizontally. Also, adjustable stiffness can broaden the range of tasks.

c) Lightweight and durable construction: Reducing mass not only increases the robot's payload capacity but also diminishes oscillations during abrupt accelerations and decelerations. Moreover, lowering the end-effector's weight reduces gravitational sag. Because training involves many trials, the end-effector should be robust and durable.

d) Low cost and simple implementation: To encourage widespread adoption, the design should minimize part count, rely on readily available components, and be inexpensive and straightforward to assemble.

We propose the novel soft end-effector CLAW (Compliant Leaf-spring Anisotropic soft Wrist) for robust robot learning, shown in Fig. 1 (a). CLAW features a two-finger gripper and a soft wrist composed of two orthogonal leaf springs and rotary joints with a locking mechanism. CLAW exhibits large 6-DoF deformations and anisotropic stiffness and can switch between three stiffness modes. We also propose using a convex tape for the leaf spring, resulting in a stable, compact wrist structure. The total construction cost is approximately \$550, and it only weighs 330g.

The main contributions of this study are the design of a novel soft wrist and its evaluation in learning contact-rich tasks. We compared CLAW with widely-used baselines, a TPU-printed soft Fin Ray gripper, and an off-the-shelf rigid parallel gripper. In the deformation evaluation, CLAW demonstrated higher compliance in a wider range of directions than the baselines. Then, we evaluated CLAW within an imitation learning framework, leveraging the sample-efficient method Action Chunking with Transformer (ACT) [1] to train manipulation policies on benchmark peg-insertion tasks. CLAW outperformed the baseline grippers. Beyond benchmarking, CLAW demonstrated robust performance across a variety of challenging contact-rich tasks, including delicate object manipulation and tight-tolerance insertions, showcasing its potential to enable robot learning in unstructured environments, shown in Fig. 1 (b).

II. RELATED WORK

A wide variety of mechanisms have been proposed to realize compliance in robotic manipulation. In this section, we review prior work on soft grippers and soft wrists for contact-rich manipulation.

Numerous soft grippers composed of flexible materials have been developed. One approach places springs along each principal axis to impart tunable, potentially anisotropic stiffness to the hand [8], although this inevitably increases part count and mechanical complexity. Tendon-driven, under-actuated rigid fingers [6], [9] and hands [10] have demonstrated robust and dexterous manipulation with relatively

TABLE I
MECHANICAL CHARACTERIZATION OF CLAW AND OTHER SOFT WRISTS

Specification	Value			Unit
	Ours 6-DoF	Wrist [18] 6-DoF	Wrist [19] 2-DoF	
Weight	330 (w/ gripper, wrist, camera)	Not reported	110 (w/o gripper)	g
Size	155 × 150 × 85 (W × H × L)	ϕ 75 × 60	ϕ 63 × 45	mm
Deformation	X/Y	40	6.35	mm
	Z	20 (comp.), 10 (ext.)	11	mm
	Roll/ Pitch	15	8	Deg
	Yaw	30 (Full/Half-lock), 45 (Free)	Not reported	N/A
	Maximum grip force	7.5	N/A	N

few actuators, but their deformation directions remain constrained. Fin Ray grippers, inspired by the deformation mechanics of fish fins [11], have been employed in insertion tasks [12], [13] and pick-and-place operations [14]. When coupled with a backdrivable motor, they enable adaptive grasping and other contact-rich actions [5]. Despite the advantages of rapid 3D printing and the ability to tune stiffness via geometry and material choice, Fin Ray grippers remain limited in both their range and directions of deformation. Pneumatic grippers and pneumatic tactile sensors [15], [16] allow large, multi-directional deformations, but their reliance on external pneumatic hardware poses a practical constraint. Similarly, tendon-driven silicone fingers can perform delicate, contact-rich tasks such as hair grooming [17], but they are less suitable for high-force interactions.

Given the limitations of soft grippers, researchers have also explored soft wrists as an alternative approach of achieving compliance, enabling large deformations while working in conjunction with a rigid gripper to deliver sufficient grasping force. As early as the 1980s, the Remote Center Compliance (RCC) device was introduced as a spring-based mechanism to passively compensate for positional errors in assembly [20]. However, RCC offers only a limited range of motion [18]. Soft wrists [18], [21] enable large deformations in 6-DoF, but they are prone to sagging under gravity when the manipulator is oriented horizontally. Counteracting this sag typically requires stiffer or additional springs, which reduces compliance or increases structural complexity. More recent designs seek variable stiffness. A sheet-type flexure combined with tendon actuation enables a wrist with adjustable stiffness and has demonstrated diverse contact-rich tasks [19]. Nevertheless, this design requires complex structures, including three servomotors and tendons to regulate stiffness. A 3D-printed wrist incorporating a honeycomb lattice has also been proposed to achieve bimodal stiffness responses [22]. Although straightforward to fabricate, its deformation range remains narrow. Tab. I introduces the mechanical characterization between CLAW and the other wrists [18], [19].

Overall, while they show significant potential, none fully resolve the challenges outlined in the introduction. Existing grippers and wrists require mechanically complex structures, restrict the range or direction of deformation, or struggle to maintain stable grasping in contact-rich tasks.

III. CLAW DESIGN CONCEPT

The objective of this study is to develop a mechanism for robust learning in contact-rich manipulation. To this end, we propose a novel soft wrist mechanism, CLAW, which integrates a two-finger parallel gripper. The mechanical design of CLAW is presented in Fig. 2. The overall system comprises three principal components: the finger unit, the compliant module, and the wrist base. The compliant module constitutes the core of the mechanism, for which we introduce a convex tape-based leaf-spring structure. This structure has the following features:

a) Large and multi-directional deformation: Leaf springs made of thin metallic tapes exhibit a wider deformation range and higher fatigue resistance compared to plastic flexure materials.

b) Anisotropic and adjustable stiffness: Thin metallic structures are highly compliant in bending and torsion, while remaining stiff in the shear direction. The stiffness is adjustable by a lock mechanism.

c) Lightweight construction: Metallic materials possess a high specific strength, and the thickness of the metallic tapes is less than 1 mm.

d) Low cost and simple implementation: By employing convex tapes as leaf springs, the compliant module inherently preserves a stable shape without the need for additional supporting components.

In the following subsections, we describe the deformation principle, provide theoretical analyses, and detail the remaining components of the system.

A. Deformation Principle of Compliant Module

The compliant module contains two orthogonally arranged, looped leaf springs coupled to rotary joints pre-loaded by two coil springs, shown in Fig. 3 (a). With only four compliant components, it provides 6-DoF and anisotropic compliance. The middle of the leaf springs is fixed to the finger unit, whereas both ends are attached to the rotary joints on the wrist base. Anisotropic stiffness arises from the combined bending and twisting of the leaf springs. Furthermore, variable stiffness in X, Y, and yaw directions is obtained by selectively locking the rotary joints as shown in Fig. 2 (b). CLAW has three lock modes: 1) Free mode releases all rotational joints, 2) Half-lock mode locks two joints in the X direction, and 3) Full-lock mode locks all four joints. The following explains the deformation mechanism for each degree of freedom.

X and Y deformation When an external force is applied along the X- or Y-axis, the springs deform as illustrated in Fig. 3 (b): the two bending (semi-circular) sections translate, shifting the finger-unit. When forces are applied in the shear direction to the leaf spring, displacement occurs

by the rotary joints and the springs' bending and torsion, as shown in Fig. 3 (a). If the rotary joints are locked while a force is applied (e.g., along the X-axis), the spring oriented along Y resists shear deformation because of its stiffness, thereby limiting displacement in the X direction.

Z, roll, and pitch deformation As shown in Fig. 3 (c) and (d), displacement along the Z-axis and rotation in roll or pitch occur due to changes in the radius of curvature of the bending parts (semi-arc segments) of the leaf springs. Also, in roll and pitch deformation, the leaf springs accommodate these movements through their twisting (center of Fig. 3 (d)).

Yaw deformation As shown in Fig. 3 (e), yaw deformation occurs when the rotary joints rotate in the same direction. This causes the leaf springs to undergo combined bending and torsional deformation, enabling the finger unit to rotate. When the rotary joints are locked, the shear stiffness of the leaf springs restricts yaw rotation.

B. Deformation analysis of Compliant Module

Expanding the compliant module's range of motion requires a corresponding increase in the looped leaf spring's initial lateral width. Below, we discuss the relationship between the size of CLAW's loop structure and its range of motion in the XY plane.

Definition of leaf-spring lengths and loop width The deformations occur in the semi-arc segment of the leaf springs. Their lengths are defined as shown in Fig. 4. Points *A* and *A'* correspond to the axes of the rotational joints, and *C* and *C'* denote the clamped sections to the Finger unit. $L_{\alpha-\beta}$ represents a distance measured along the spring between two points α and β , and thus $L_{a-a'}$ is the total length of a leaf spring. Given that the loop shape is approximated as a combination of semicircular arcs and straight segments, the lateral width of the loop D is given by the following expression:

$$D = \frac{L_{a-a'} + d}{2} - \pi R + 2R, \quad (1)$$

where d is the inter-axial distance between the a and a' rotational joints, determined by clearances with other parts.

Range of motion in XY plane The maximum deformation in the XY plane is determined by the length of the leaf spring. When the gripper is displaced in the XY direction, the boundary between the semicircular segment and the straight segment first coincides with point *B*. As the displacement increases further, the radius of curvature of the semicircular arc gradually increases while deforming in the *X* direction. For simplicity, it is assumed that when the gripper reaches its maximum displacement, the leaf spring becomes completely straight. Under this assumption, the maximum displacement of the gripper X_{max} can be expressed as:

$$X_{max} = \sqrt{L_{b-c}^2 - 4R^2} - X_0. \quad (2)$$

where X_0 is the initial position of *C*. Additionally, L_{b-c} and $L_{a-a'}$ exhibit a linear relationship:

$$L_{b-c} = \frac{L_{a-a'} - L_{c-c'}}{2} - L_{a-b}. \quad (3)$$

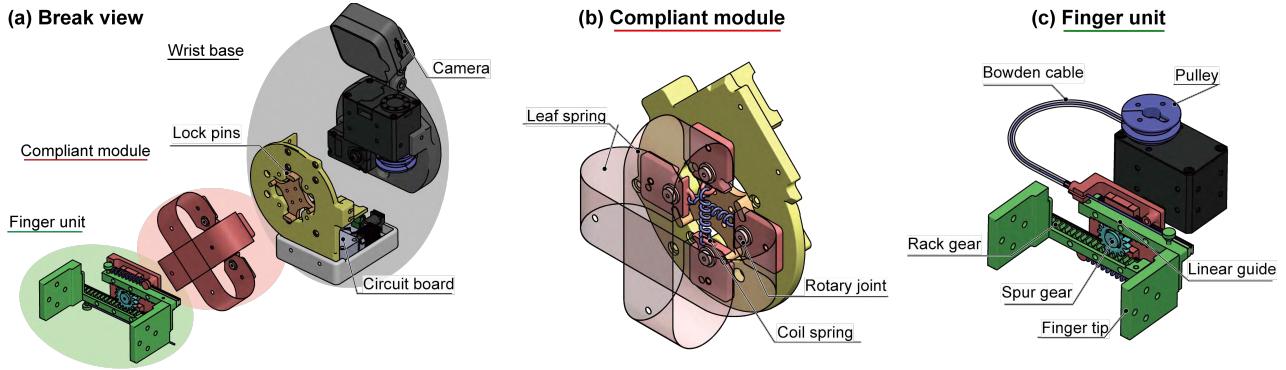


Fig. 2. **Mechanical design of the CLAW:** Orthogonal leaf springs and rotary joints provide 6-DoF compliance with anisotropic stiffness. A simple locking mechanism enables mode-switchable stiffness while preserving low weight and structural simplicity.

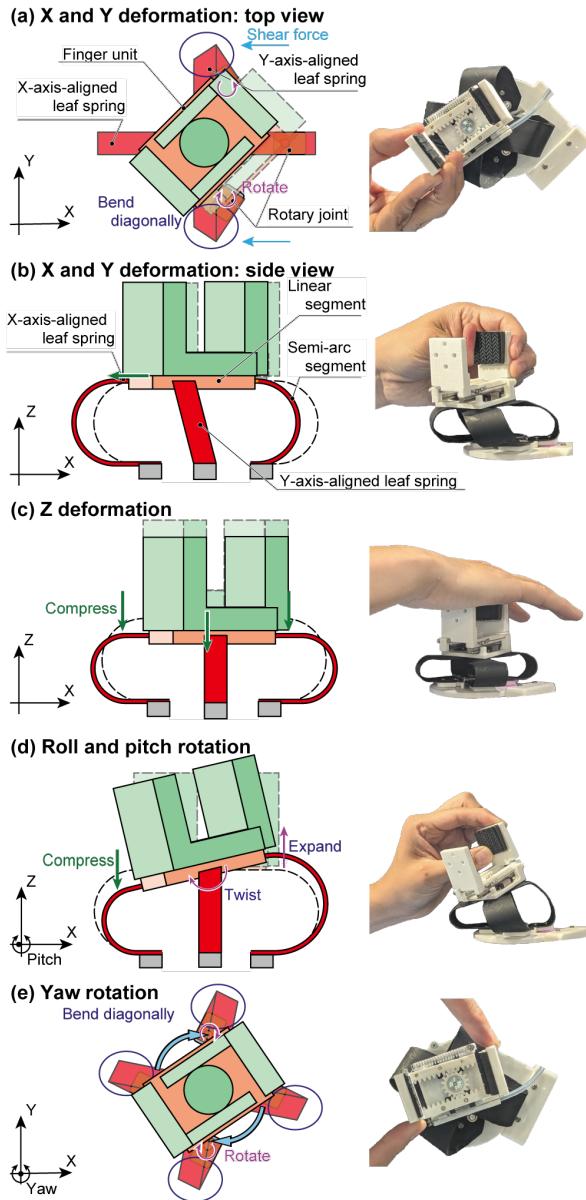


Fig. 3. **Deformation characteristics of the CLAW compliant module:** Two orthogonal looped leaf springs with rotary joints and coil spring preload enable 6-DoF anisotropic compliance.

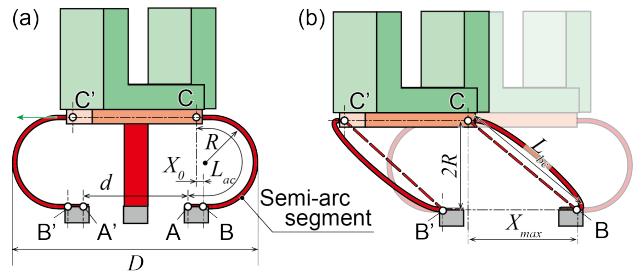


Fig. 4. **Length and range of motion of the leaf spring:** (a) Fixed section of the leaf spring and dimensions of each part, (b) Maximum deformation state in the XY direction.

In practical design, D is first determined by size constraints, which in turn define $L_{a-a'}$. We adopt a leaf spring with $R = 15$ mm and $L_{a-a'} = 180$ mm, yielding $D = 90$ mm, which corresponds to 1.2 times the UR5e's flange diameter. Additionally, the system is operated such that the fingertip position error remains within X_{max} . Deformation to this maximum value may lead to mechanical failure at the joints or other parts of the leaf spring. Therefore, in this study, the maximum allowable displacement is defined as 80% of X_{max} , resulting in an allowable displacement of 38 mm.

C. Convex-tape

While a recent study proposes convex tapes for a gripper [23], we use them for the wrist compliant module. The tapes' elastic restoring forces stabilize the finger unit's pose while providing compliance in multiple directions. A convex tape is a metal strip with a uniform, arched cross-section that is inherently stable in a straight configuration. When a convex tape is bent, the cross-section at the bending location flattens locally, allowing deformation only at that specific region. Consequently, a loop formed from this material is not circular; instead, it has an oval-slot shape composed of two semicircular arcs connected by straight segments. This geometry shortens the structure in the Z' -direction and lowers the moment transmitted to the springs by the finger unit's weight when the end-effector is oriented horizontally.

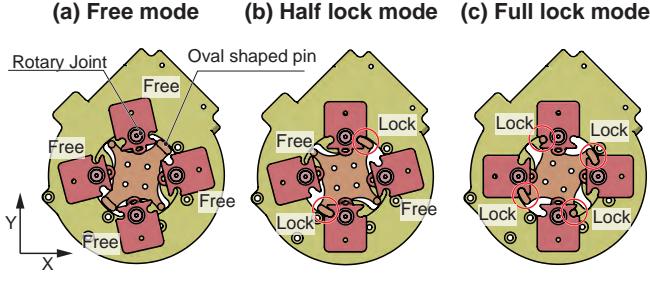


Fig. 5. **Locking mechanism of the CLAW:** A single actuator moves four pins with different profiles on a carrier to engage grooves in joint links. This system enables three stiffness modes: (a) Free mode with all joints rotatable, (b) Half-lock mode restricting motion along the X-axis by locking two joints, and (c) Full-lock mode locking all joints to increase stiffness in X, Y, and yaw while retaining compliance in other directions.

D. Lock Mechanism

Fig. 5 illustrates the proposed lock mechanism. Each joint link contains a groove, and locking is achieved by inserting pins into these grooves. All four pins are mounted on a single carrier and therefore move synchronously under a single actuator. Because the pins have different profiles, the joints that lock depend on the actuator's direction of rotation, enabling three distinct stiffness modes.

Free mode (Fig. 5(a)) None of the pins engage their grooves, so all four joints remain fully rotatable.

Half-lock mode (Fig. 5(b)) Only the slot-shaped pins engage, locking two opposing joints; motion along the X-axis is restrained, while compliance in the Y-axis is preserved. This mode is advantageous when gravity acts along X, as it increases stiffness solely in the gravity direction.

Full-lock mode (Fig. 5(c)) All pins engage, locking every joint. The compliant unit thus becomes stiff in X, Y, and yaw, whereas stiffness about Z, roll, and pitch remains unchanged. This configuration suppresses vibration of the finger unit during free-space motions while still permitting absorption of impact forces along Z during contact.

E. Finger Unit and Wrist Base

To prevent gravity-induced sagging of the springs, the finger unit is actuated via a remote Bowden-cable transmission. The wrist base houses two Dynamixel motors, including an XM430 for finger actuation and an XL330 for stiffness switching, together with the control electronics and an Intel RealSense D405 depth camera.

IV. EXPERIMENTAL EVALUATION

To validate our system, we performed real robot experiments. Our goal is to answer the following questions: (1) Is CLAW more compliant in multiple directions (in Sec. IV-B), (2) Does the adjustable stiffness benefit in contact-rich manipulation (in Sec. IV-C), (3) Does CLAW complete contact-rich tasks more successfully than the other end-effectors (in Sec. IV-D), and (4) Is our system capable of executing real-world contact-rich manipulation tasks (in Sec. IV-E).

To address these questions, first, we conducted deformation tests to characterize CLAW's compliance by comparing

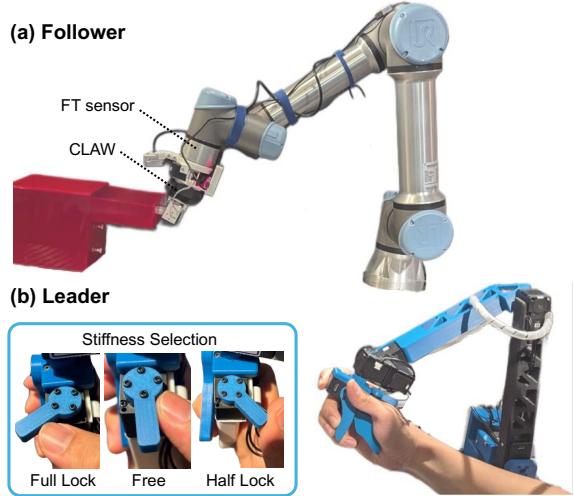


Fig. 6. **Teleoperation setup:** A haptic teleoperation leader device is employed to control a UR5e robot arm. On the leader device, user can select stiffness mode by switching a lever. FT sensor on the follower arm relay external force information to the leader, which is then mapped to joint torques to create haptic feedback.

different end-effectors. Second, we tested the adjustable stiffness effects in a door-opening task. Third, we evaluated its learning performance in a peg-in-hole benchmarking task against the end-effectors. Finally, we deployed CLAW on a variety of real-world manipulation tasks.

A. Experimental Setting

Hardware We used a collaborative robot arm (UR5e, Universal Robots A/S, Denmark) equipped with two cameras (RealSense D405, Intel Corporation, USA) and a desktop PC powered by a CPU AMD Ryzen Threadripper 7980X and a GPU NVIDIA RTX 6000 Ada. In the deformation and benchmark evaluations, we compared CLAW with two baselines: a rigid gripper (2F-85, Robotiq, Canada) and a TPU-printed Fin Ray gripper, a common soft end-effector that is easy to fabricate. The policy was executed at 50 Hz.

Teleoperation system We developed a haptic feedback teleoperation system inspired by GELLO [24] (leader robot) and its extension [25], [26], which has an equivalent kinematic structure to the follower robot (UR5e), see Fig. 6. The frame can be 3D-printed, and the servo motors are low-cost and can operate with current control. While the prior study utilized the leader robot for a torque control-based robot [25], [26], our system was designed for a position-control robot by the two-channel teleoperation scheme from [27]: the leader streams joint positions to the follower, while the follower returns force feedback to the leader. We also added a lever switch for the user to easily change the stiffness of CLAW.

Imitation learning We employed ACT [1], a widely used imitation learning method known for its sample efficiency. We collected demonstrations using the teleoperation system. ACT's inputs were images from wrist-mounted and fixed-view cameras, force-torque (FT) sensor readings, joint positions, and stiffness modes, and the outputs were action chunks: sequences of joint positions and stiffness modes.

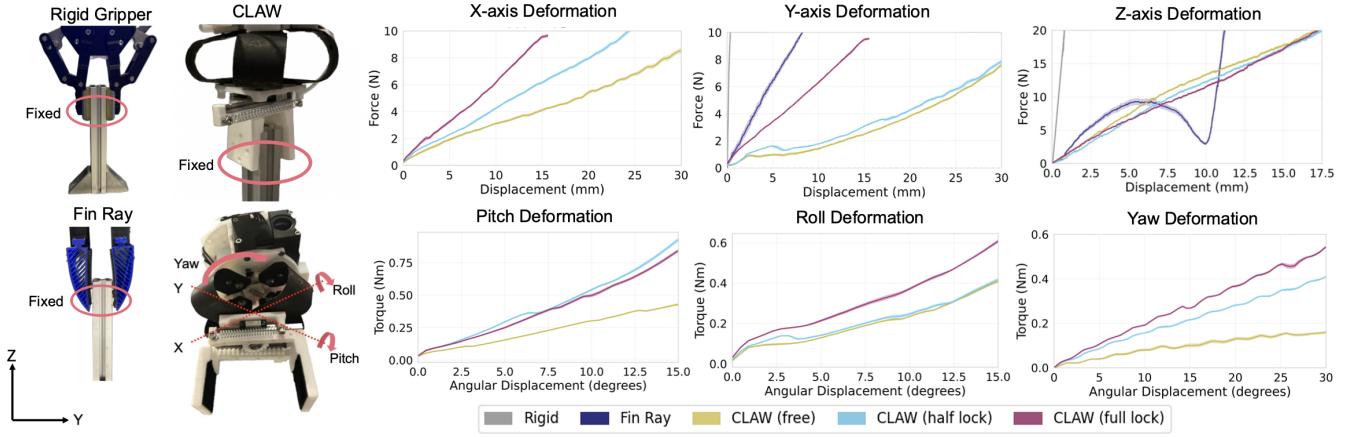


Fig. 7. **Deformation comparison of CLAW, Fin Ray, and a rigid gripper:** Each end-effector's tool center point was fixed to a vertical pole, as shown on the left side. For the Z-axis deformation on the Fin Ray, we placed the gripper on top of the vertical pole. Force readings were measured via the UR5e's onboard FT sensor.

Control We employed a forward dynamics compliance controller [28] to reduce oscillations and enable more stable teleoperation and policy execution during contact. The controller parameters were manually fine-tuned to minimize oscillations while preserving responsiveness for each gripper. The control loop ran at 500 Hz.

B. Deformation Evaluation

We imposed incremental deformations by applying small translational (X, Y, Z) and angular (roll, pitch, yaw) motions and recorded the reaction wrench using the UR5e six-axis FT sensor. Fig. 7 shows the deformation comparison. Only the deformations of interest (Y- and Z-axis translations) were evaluated for the Fin Ray and the rigid gripper; therefore, the X-axis and rotational plots in Fig. 7 show CLAW only. The Fin Ray is not compliant along the other axes, and the rigid gripper is stiff along all axes.

CLAW exhibits anisotropic stiffness. The Free mode is the softest, and the Full-lock mode is the stiffest. The Half-lock mode behaves similarly to Free along the Y-axis but is noticeably stiffer along the X-axis. For rotations, Half-lock is close to Full-lock in pitch but remains close to Free in roll. Along the Z-axis, the three CLAW modes have similar compliance. By contrast, the rigid gripper is highly stiff; Fin Ray's response is non-monotonic and unpredictable, stiffening initially, softening near 10 mm, and then stiffening again. In the Y-axis, the Fin Ray reaches about 10 N at 8 mm displacement, and the rigid gripper shows almost no compliance. Within CLAW, Full-lock requires roughly two times the force of Free at 15 mm in the Y-axis; in yaw, Full-lock delivers about three times the torque of Free at 30°, whereas Half-lock delivers about two times the torque.

Overall, CLAW enables compliant interaction in more directions than the baseline grippers, which are compliant only in a limited set of directions. With these stiffness modes being selectable on demand, CLAW can be advantageous for robust and reliable contact-rich manipulation.

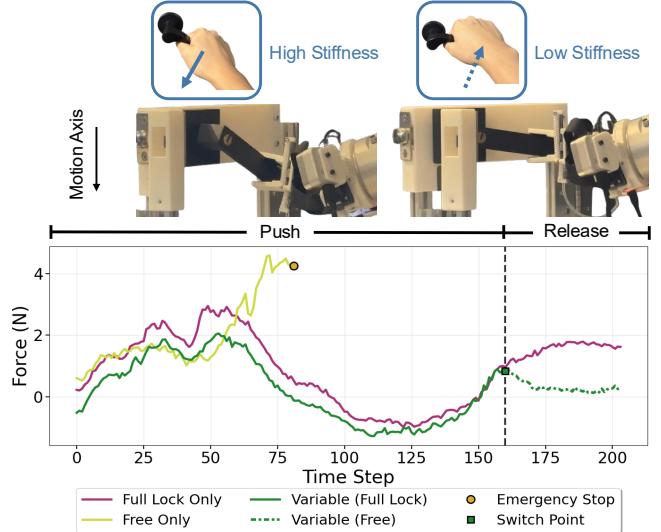


Fig. 8. **Door opening test:** Top figures show the motivation for variable stiffness in door opening as humans turn doorknobs with high wrist stiffness and release with low wrist stiffness. The bottom figure shows a comparison of the force trajectories of variable stiffness, free mode only, and full-lock only. Free mode triggered the emergency stop as the force threshold was exceeded in a different axis.

C. Benefit of Variable Stiffness

To validate the benefit of the variable stiffness mechanism, we examine whether switching stiffness mode mid-task provides a benefit to performance in a door-opening task.

In this experiment, we teleoperate CLAW using the system reported in Sec. IV-A to open a spring-loaded door handle. The door latch releases only after the handle is rotated by approximately 45°. As illustrated in Fig. 8 (top), we follow a human-like strategy: a high-stiffness (Full-lock) mode during the push/rotation phase to transmit torque efficiently, then an immediate switch to low-stiffness (Free) mode once the latch releases. In Free mode, the wrist back-drives, so the door handle's internal return spring sets the final pose without the robot resisting it. Then, we replayed the recorded episode

in three stiffness conditions: (1) the recorded condition in which CLAW switches from Full-lock mode to Free mode, (2) only using Full-lock mode, and (3) only using Free mode.

Fig. 8 (bottom) shows the force profiles during the motion replays with the three conditions. CLAW, with only Free mode, could not transmit enough force through the motion and triggered the emergency stop. Both Full-lock and variable stiffness modes were able to complete the task. By looking at the force trajectory along the primary motion axis, it can be observed that in Full-lock, elevated force persists even after the handle has rotated, indicating the wrist keeps resisting the handle's return spring during the release phase. In contrast, variable stiffness switches to Free immediately at release, producing a sharp drop of force as it lets the handle's internal spring set the terminal pose.

This stiffness switch is advantageous for the door opening. By handing control back to the environment after opening, the policy no longer needs to plan or time a precise “rotate-back” motion; Free mode's low stiffness passively accommodates the external spring. This reduces the need to actively exert forces (fewer protective stops).

D. Learning Performance

To test whether CLAW has superior learning performance than off-the-shelf end-effectors, we conducted a benchmarking experiment using the assembly board from a Functional Manipulation Benchmark for Generalizable Robotic Learning (FMB) [29]. We evaluated the three end-effectors on three peg-in-hole geometries representing the best-, median-, and worst-performing cases in the benchmark [29]. Since in different insertion paths the end effector may need to be compliant in different directions, we collected 25 human demonstrations from paths both parallel and perpendicular to the gripper palm (in Fig. 9(a)), yielding 50 demonstrations per task. We set a timeout of 60 seconds for inference.

By plotting multiple rollout trajectories (in Fig. 9(b)) from the benchmarking experiment, it can be observed that CLAW more effectively absorbs and suppresses post-contact oscillation. As a result, insertion paths are smooth and consistent across trials, demonstrating CLAW's ability to maintain more precise, reliable motion in contact-rich scenarios.

The success rates (in Fig. 9(c)), were computed from policies trained on three random seeds, with each seed evaluated over 20 rollouts per approach direction (in total 120 trials). For the simpler circle peg, all grippers achieved relatively high success rates. For the square+circle peg, the Fin Ray gripper showed strongly divergent performance depending on approach direction. It achieved about four times higher success in the parallel direction, as its compliant axis helped with alignment and insertion along that path. The rigid gripper struggled in both directions, while CLAW maintained strong performance across both. For the three-prong peg, success rates dropped further due to the triangular layout (one peg at the top and two at the bottom). This configuration made the perpendicular approach particularly prone to collisions, resulting in a large discrepancy in approach directions. Despite this inherent geometric difficulty,

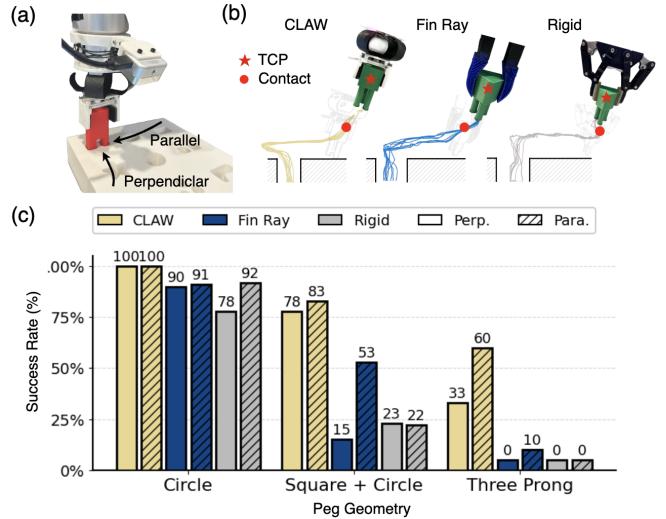


Fig. 9. **Learning performance benchmarking:** (a) Experimental setup with the FMB board and insertion peg. (b) Representative trajectory from the rollout insertion path of the three end-effectors. (c) Success rate comparison between the learning performance of CLAW, Fin Ray, and rigid gripper for different peg geometries and approach directions.

CLAW still outperformed the baselines with an overall 76% success rate across geometries, compared to 43% and 36% for the Fin Ray and rigid gripper, demonstrating robustness to approach directions and geometries.

CLAW's main failure modes were slight misalignment or exceeding the time limit for each trial. Fin Ray often failed due to slippage caused by contact deformation. The rigid gripper frequently exerted excessive lateral (X-Y) forces that led to emergency stops.

E. Task Versatility

To demonstrate the versatility of our system, we evaluated CLAW on a range of challenging, delicate, and contact-rich tasks. For operations on vertical surfaces, such as horizontal peg-in-hole and drawer pulling, we engaged Full-lock mode to stiffen the end-effector when oriented horizontally, preventing sagging. Half-lock mode was chosen for dragging metal pegs, as its anisotropic stiffness allowed the gripper to transmit force along the motion axis while remaining laterally compliant to prevent jamming. For more complex tasks, including door opening and key insertion, stiffness modes were varied during the task. Finally, Free mode was used for the remaining tasks. We collected 30 demonstrations for single-stiffness tasks and 50 for variable-stiffness tasks. Success rates are reported from 20 rollout episodes.

Results in Tab. II show the versatility of CLAW across a variety of real-world tasks by leveraging different stiffness modes. In peg-in-hole tasks with tight tolerances (approximately $50 \mu\text{m}$), CLAW enables precise and compliant insertion without applying excessive force. Likewise, in glass-tube insertion, it ensures stable contact and minimizes the risk of damaging delicate objects. By actively varying its stiffness, CLAW adapts seamlessly to tasks such as key insertion and door opening. For key insertion, it first operates in Free

TABLE II
TASK SUCCESS RATES FOR VARIOUS REAL-LIFE TASKS.

Task	Success Rate	Stiffness Mode
Vertical peg-in-hole	0.90	Free
Glass-tube insertion	0.85	Free
Dish-cap placement	0.68	Free
Drawing smiley face	0.47	Free
Dragging metal peg	1.00	Half-lock
Horizontal peg-in-hole	0.45	Full-lock
Drawer pulling	0.75	Full-lock
Key insertion	1.00	Free \rightarrow Full-lock
Door opening	1.00	Full-lock \rightarrow Free

mode to allow easier alignment, then switches to stiff mode to transmit sufficient torque for key rotation. For the door opening, as described in Sec. IV-C, stiff mode is engaged initially to rotate the handle, after which CLAW transitions to Free mode to robustly release it. The supplementary material includes a video of these task executions.

V. CONCLUSION

This study introduced a novel soft wrist, CLAW, enabling robust learning of contact-rich manipulation tasks. Experiments showed that CLAW achieved higher success rates on benchmark tasks than widely used baseline grippers and successfully executed more demanding, contact-rich tasks.

Limitations and Future Work Despite its promise, CLAW has several limitations. First, it occasionally failed because the grasped object slipped. This issue could be mitigated by increasing fingertip friction or by using a higher-torque motor. Also, it may sag under the weight of excessively heavy objects. Future work will apply CLAW to the other learning algorithms, such as reinforcement learning.

REFERENCES

- [1] T. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” *Robotics: Science and Systems*, 2023.
- [2] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [3] H. Nguyen and H. La, “Review of deep reinforcement learning for robot manipulation,” in *IEEE International Conference on Robotic Computing*, 2019, pp. 590–595.
- [4] S. Haddadin, S. Parusel, L. Johannsmeier, S. Golz, S. Gabl, F. Walch, M. Sabaghian, C. Jähne, L. Hausperger, and S. Haddadin, “The franka emika robot: A reference platform for robotics research and education,” *IEEE Robotics & Automation Magazine*, vol. 29, no. 2, pp. 46–64, 2022.
- [5] G.-C. Jeong, A. Bahety, G. Pedraza, A. D. Deshpande, and R. Martín-Martín, “Bariflex: A robotic gripper with versatility and collision robustness for robot learning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024, pp. 4106–4113.
- [6] O. Azulay, D. M. Ramesh, N. Curtis, and A. Sintov, “Visuotactile-based learning for insertion with compliant hands,” *IEEE Robotics and Automation Letters*, vol. 10, no. 4, pp. 4053–4060, 2025.
- [7] C. Zhao, C. Jiang, L. Luo, S. Yuan, Q. Chen, and H. Yu, “Learning thin deformable object manipulation with a multi-sensory integrated soft hand,” *IEEE Transactions on Robotics*, vol. 41, pp. 4606–4621, 2025.
- [8] N. Fukaya, K. Yamane, S. Masuda, A. Ummadisingu, S.-i. Maeda, and K. Takahashi, “Four-axis adaptive fingers hand for object insertion: Faaf hand,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024, pp. 13 996–14 003.
- [9] A. S. Morgan, B. Wen, J. Liang, A. Boularias, A. M. Dollar, and K. Bekris, “Vision-driven compliant manipulation for reliable, high-precision assembly tasks,” in *Robotics: Science and Systems*, 2021.
- [10] T. Ko, “A tendon-driven robot gripper with passively switchable underactuated surface and its physics simulation based parameter optimization,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5002–5009, 2020.
- [11] W. Crooks, G. Vukasin, M. O’Sullivan, W. Messner, and C. Rogers, “Fin ray® effect inspired soft robotic gripper: From the robosoft grand challenge toward optimization,” *Frontiers in Robotics and AI*, vol. 3, p. 70, 2016.
- [12] K. Nie, W. Wan, and K. Harada, “An adaptive robotic gripper with l-shape fingers for peg-in-hole tasks,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 4022–4028.
- [13] R. M. Hartisch and K. Haninger, “High-speed electrical connector assembly by structured compliance in a finray-effect gripper,” *IEEE/ASME Transactions on Mechatronics*, vol. 29, no. 2, pp. 810–819, 2023.
- [14] X. Wang, B. Wang, J. Pinski, Y. Xie, J. Brett, R. Scalzo, and D. Howard, “Fin-bayes: A multi-objective bayesian optimization framework for soft robotic fingers,” *Soft Robotics*, vol. 11, no. 5, pp. 791–801, 2024.
- [15] S. Brahmabhatt, A. Deka, A. Spielberg, and M. Müller, “Zero-shot transfer of haptics-based object insertion policies,” in *IEEE International Conference on Robotics and Automation*, 2023, pp. 3940–3947.
- [16] M. Oller, M. P. i Lisbona, D. Berenson, and N. Fazeli, “Manipulation via membranes: High-resolution and highly deformable tactile sensing and control,” in *Conference on Robot Learning*, 2023, pp. 1850–1859.
- [17] U. Yoo, N. Dennler, E. Xing, M. Matarić, S. Nikolaidis, J. Ichnowski, and J. Oh, “Soft and compliant contact-rich hair manipulation and care,” in *ACM/IEEE International Conference on Human-Robot Interaction*, 2025, pp. 610–619.
- [18] F. von Drigalski, K. Tanaka, M. Hamaya, R. Lee, C. Nakashima, Y. Shibata, and Y. Ijiri, “A compact, cable-driven, activatable soft wrist with six degrees of freedom for assembly tasks,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 8752–8757.
- [19] H. Sun, S. Park, and D. Hwang, “Compact modular robotic wrist with variable stiffness capability,” *IEEE Transactions on Robotics*, vol. 41, pp. 141–158, 2025.
- [20] D. E. Whitney *et al.*, “Quasi-static assembly of compliantly supported rigid parts,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 104, no. 1, pp. 65–77, 1982.
- [21] Q. Zhang, Z. Hu, W. Wan, and K. Harada, “Compliant peg-in-hole assembly using a very soft wrist,” *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 17–24, 2023.
- [22] G.-C. Jeong, S. Dalla Gasperina, A. D. Deshpande, L. Chin, and R. Martín-Martín, “Biflex: A passive bimodal stiffness flexible wrist for manipulation in unstructured environments,” *IEEE Robotics and Automation Letters*, vol. 10, no. 10, pp. 10 783–10 790, 2025.
- [23] G. He, C. Sparks, and N. Gravish, “Grasping and rolling in-plane manipulation using deployable tape spring appendages,” *Science Advances*, vol. 11, no. 15, p. ead7505, 2025.
- [24] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, “Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024, pp. 12 156–12 163.
- [25] S. Sujit, L. Nunziante, D. O. Lillrank, R. F. J. Dossa, and K. Arulkumaran, “Improving low-cost teleoperation: Augmenting gello with force,” in *IEEE/SICE International Symposium on System Integration*, 2025, pp. 747–752.
- [26] J. J. Liu, Y. Li, K. Shaw, T. Tao, R. Salakhutdinov, and D. Pathak, “Factr: Force-attending curriculum training for contact-rich policy learning,” *arXiv preprint arXiv:2502.17432*, 2025.
- [27] B. Hannaford, “A design framework for teleoperators with kinesthetic feedback,” *IEEE Transactions on Robotics and Automation*, vol. 5, no. 4, pp. 426–434, 1989.
- [28] S. Scherzinger, A. Roennau, and R. Dillmann, “Forward dynamics compliance control (FDCC): A new approach to cartesian compliance for robotic manipulators,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 4568–4575.

[29] J. Luo, C. Xu, F. Liu, L. Tan, Z. Lin, J. Wu, P. Abbeel, and S. Levine, “FMB: A functional manipulation benchmark for generalizable robotic learning,” *The International Journal of Robotics Research*, vol. 44, no. 4, pp. 592–606, 2025.