

ShapeForce: Low-Cost Soft Robotic Wrist for Contact-Rich Manipulation

Jinxuan Zhu^{*1,2}, Zihao Yan^{*1,2}, Yangyu Xiao^{1,2},
 Jingxiang Guo¹, Chenrui Tie¹, Xinyi Cao¹, Yuhang Zheng¹ and Lin Shao^{†1,2}

Abstract—Contact feedback is essential for contact-rich robotic manipulation, as it allows the robot to detect subtle interaction changes and adjust its actions accordingly. Six-axis force-torque sensors are commonly used to obtain contact feedback, but their high cost and fragility have discouraged many researchers from adopting them in contact-rich tasks. To offer a more cost-efficient and easy-accessible source of contact feedback, we present ShapeForce, a low-cost, plug-and-play soft wrist that provides force-like signals for contact-rich robotic manipulation. Inspired by how humans rely on relative force changes in contact rather than precise force magnitudes, ShapeForce converts external force and torque into measurable deformations of its compliant core, which are then estimated via marker-based pose tracking and converted into force-like signals. Our design eliminates the need for calibration or specialized electronics to obtain exact values, and instead focuses on capturing force and torque changes sufficient for enabling contact-rich manipulation. Extensive experiments across diverse contact-rich tasks and manipulation policies demonstrate that ShapeForce delivers performance comparable to six-axis force-torque sensors at an extremely low cost. More details of this project can be found at our project page: <https://shapeforce.github.io/>.

I. INTRODUCTION

Contact-rich manipulation is crucial for moving robots beyond simple pick-and-place capabilities toward performing real-world, skillful tasks. Tasks such as peg insertion, assembly, and screw tightening often require detecting subtle interaction transitions, such as whether a peg has aligned with a hole or whether a screw has been fully tightened. In these scenarios, visual perception alone is insufficient for monitoring task progress, verifying success, or perceiving occluded events. In contrast, contact feedback provides a complementary sensing modality, allowing the robot to infer interaction states and decide when to adapt its action policies.

Traditionally, contact sensing has been achieved using dedicated six-axis force-torque (FT) sensors [1]–[3]. These devices rely on rigid mechanical structures with strain gauges or piezoelectric elements to deliver high-precision force and torque signals. However, they are typically priced in the thousands of U.S. dollars and are prone to damage under intensive contact forces, further increasing maintenance and replacement costs. These limitations highlight the need for a




Fig. 1: **ShapeForce**: a low-cost, plug-and-play sensing system that provides force-like signals, enabling contact-rich tasks with comparable performance of force–torque sensors.

cost-effective sensing solution that can provide reliable contact feedback and support contact-rich manipulation tasks.

Several prior works [4]–[12] have explored alternative approaches to low-cost contact-modality sensing. However, they remain difficult to widely reproduce and apply in contact-rich tasks for two main reasons: (1) they require specialized electrical components [4]–[8]; and (2) they rely on force–torque sensors for data collection and calibration [9]–[12]. In both cases, these requirements mainly serve to obtain precise ground-truth force and torque measurements, which creates a clear conflict between the need for accurate sensing and the goal of low-cost and widely reproducible fabrication.

Taking inspiration from human manipulation, we observe that when inserting a key or tightening a cap, humans rarely need exact force values. Instead, we rely on perceiving relative changes in force and torque to detect interaction events and time our actions. This suggests that precise ground-truth measurement is not strictly necessary for contact-rich manipulation; capturing reliable trends and relative changes suffices. By prioritizing relative changes over absolute values, we eliminate the dependence of expensive force-torque sensors for calibration and specialized electronics for measurement, enabling a low-cost, easily reproducible design that still delivers reliable contact feedback.

Guided by this idea, we present ShapeForce, a low-cost, plug-and-play soft wrist that converts external forces and torques into measurable deformations of its compliant core. These deformations are subsequently captured by a wrist-

* Equal contribution.

† Corresponding author: Lin Shao (linshao@nus.edu.sg).

¹ School of Computing, National University of Singapore, Singapore.

² RoboScience, Beijing, China.

This project is supported by the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2024)

mounted RGB camera and transformed into six-dimensional force-like signals via marker-based pose tracking. Our design removes the need for expensive force-torque sensors or specialized electronics, instead relying only on a 3D-printed material and a wrist-mounted RGB camera commonly available in fine-grained robotic manipulation setups. This makes ShapeForce extremely easy to fabricate, deploy, and integrate into existing systems, while still preserving the contact information necessary to detect contact events.

We evaluate ShapeForce on a diverse suite of contact-rich tasks, including peg insertion, USB insertion, bottle cap tightening, toy desk assembly, whiteboard wiping, and maze exploration. We integrate it with classical algorithms for contact-rich tasks, including search routines (e.g., hole-finding) and force-control strategies, collectively denoted as search-and-control policy, as well as learning-based policy. Across all tasks and policy settings, ShapeForce consistently improves task success compared to baselines without contact feedback and matches the performance of systems using force-torque sensors, while offering significantly lower fabrication cost and a simple, reproducible design.

Our key contributions are as follows:

- We present **ShapeForce**, a low-cost, plug-and-play soft robotic wrist that captures its proprioceptive deformation and converts it into force-like signals to support contact-rich manipulation tasks.
- Inspired by human perception, we resolve the accuracy-simplicity trade-off by providing a task-effective force representation based on relative changes, further reducing the reproduction and deployment burden.
- We integrate our sensing system with both classical search-and-control and learning-based policies, and show experimentally that it consistently improves task success across diverse contact-rich scenarios, demonstrating the effectiveness of our force-like signals.

II. RELATED WORK

A. Mechanical Design of Robotic Wrists

The human wrist’s dexterity has inspired numerous robotic designs [13]–[17] exploring compact and variable-stiffness architectures to improve workspace and flexibility. However, actively actuated designs increase complexity, cost, and fragility under collision.

To address these drawbacks, passively compliant wrists [18]–[21] reduce cost while enhancing flexibility, though balancing compliance with precision remains challenging. Recent works like BiFlex [18] leverage controlled buckling [22] to achieve bimodal stiffness, inspiring our parameter tuning for optimal compliance-stability trade-offs.

Beyond these prior efforts, our approach uniquely exploits passive compliance as a sensing modality to estimate external wrenches, enabling contact-aware policies without dedicated force-torque sensors.

B. Vision-based Deformation Perception

Vision-based Deformation Perception (VBDeformP) is a perception paradigm that leverages visual sensing to esti-

mate the deformation of objects or structures, which are subsequently used for force sensing [9], [12], [23], [24], reconstruction [11], [25], and tactile perception [10], [26]. These deformable structures are usually made of thermo-plastic polyurethane (TPU) and can be manufactured by 3D printing, providing a low-cost perceptual approach.

A major challenge lies in representing the deformation of a soft body, which theoretically has infinite degrees of freedom. Prior works address this by using fiducial markers [27] or segmentation models such as SAM [24] to extract features, followed by regression models trained with paired force-torque data for deformation-to-force mapping. In contrast, we constrain deformation using a hard-soft-hard sandwich-like structure and encode it as the relative pose of two rigid ends, which is subsequently turned into single-marker pose tracking. The structural design enables a easier deformation estimation and benefits the deformation as representation.

C. Reproducible and Accessible Sensor Designs

Several prior works propose reproducible and accessible contact-sensing solutions through different mechanisms. eFlesh [6] combines parameterized cut-cell microstructures with off-the-shelf Hall effect sensors to build a highly customizable magnetic tactile sensor. Vision-based tactile sensors such as DIGIT [28] use a compact camera and a transparent elastomer to visually capture surface deformations. Building on this foundational idea, ReSkin [29], AnySkin [30], and the TacTip family [31] have further expanded the capabilities of vision-based tactile sensing.

However, most of these methods focus on recovering absolute force values or detailed tactile signals, which typically require calibration with ground-truth data or additional specialized electronics. In contrast, our design can be fabricated with a standard laboratory 3D printer with easy-accessible materials and uses only a wrist-mounted RGB camera to perceive, a setup that is already common in robotic manipulation. This simple and reproducible design lowers the fabrication and deployment barrier and increases its potential for adoption in both research and real-world applications.

III. METHODOLOGY

Our methodology consists of three components. We first present the mechanical design and fabrication, focusing on how the compliant core achieves desired deformation behavior (Sec. III-A). We then describe capturing deformation via marker-based pose tracking and converting it into force-like signals (Sec. III-B). Finally, we show how these signals are integrated into search-and-control policies and learning-based policies to enhance contact-rich tasks (Sec. III-C).

A. Mechanism Design and Fabrication

We followed two key principles for design and fabrication. First, the sensor is designed to provide a high-quality representation of external force and torque with sufficient sensitivity, while maintaining stability and precision under dynamic motion. Second, the system should be low-cost and simple to fabricate using easy-accessible materials, and also

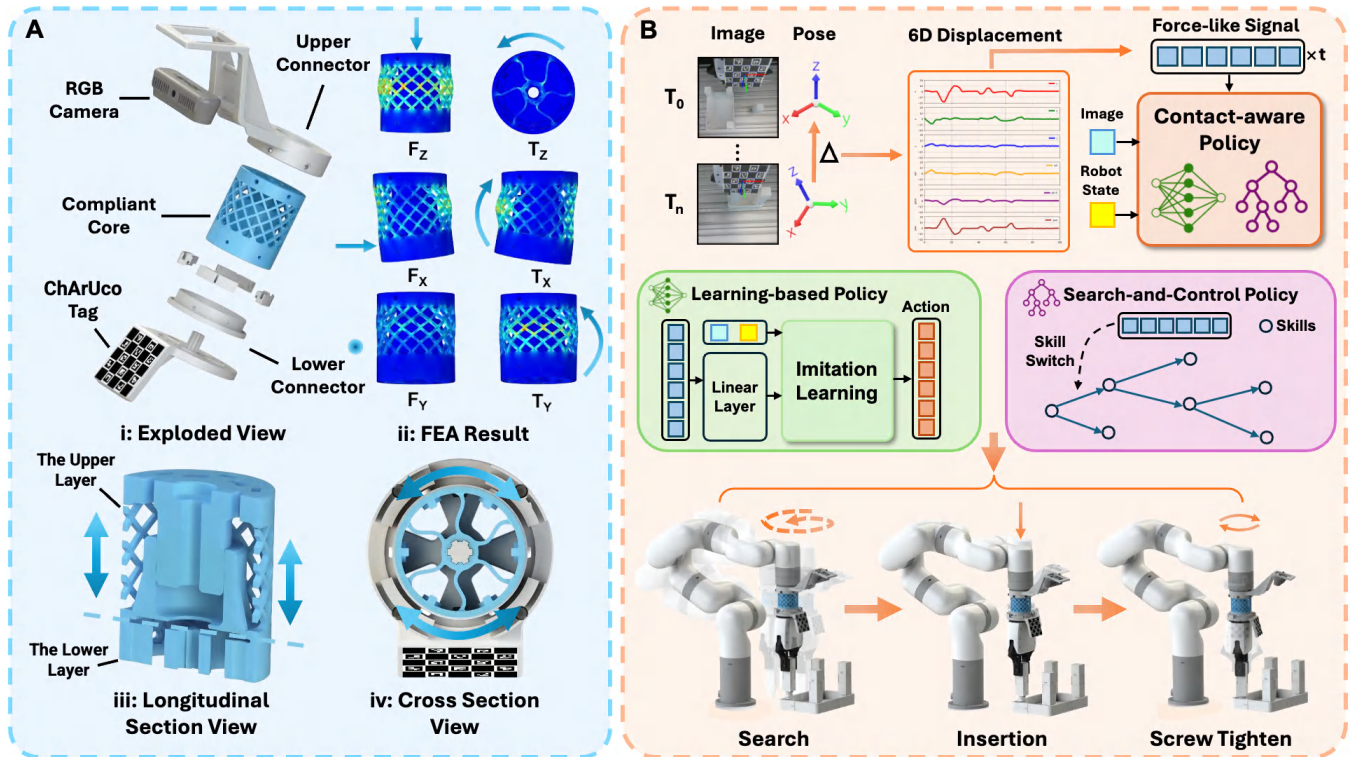


Fig. 2: **Overview of ShapeForce:** (A) Mechanical design of ShapeForce, including (i) exploded view, (ii) finite element analysis (FEA) results, (iii) longitudinal section, (iv) cross section; (B) Perception-to-control pipeline. Marker-based pose tracking converts deformation into 6D displacement, producing force-like signals that inform contact-aware policies. Both search-and-control policies and imitation-learning policies are supported, enabling execution of contact-rich tasks.

easy to assemble. Based on these two principles, the detailed mechanism design and fabrication process are as follows:

1) *Mechanism Design:* The mechanism design is illustrated in the exploded view in Fig. 2A(i). ShapeForce consists of two main components: (1) a central compliant core (blue) that converts external forces and torques into measurable deformations, and (2) rigid connectors (white) that both link the core to the robot and constrain its deformation.

The compliant core is the centerpiece of our system, designed to translate external wrenches into measurable proprioceptive deformation. In typical contact-rich manipulation with a wrist-mounted gripper, the dominant interaction wrench components are the axial force F_z (e.g., during insertion or pressing) and the torsional moment T_z (e.g., during screwing or twisting). External forces and torques along other directions can often be transformed through suitable robotic arm pose adjustments into equivalent effects along F_z or T_z of the compliant core. Our design therefore emphasizes these components, mechanically decoupling F_z and T_z as illustrated in Fig. 2A(iii)(iv), providing high-quality representation while allowing their sensitivities to be independently tuned to match task requirements. The remaining wrench components are tuned to maintain sufficient stiffness for stability and precise motion, while preserving passive compliance that improves robustness during contact. Building on this design principle, the compliant core is divided into two functional layers that decouple and emphasize

the axial force F_z and torsional moment T_z .

The upper layer adopts a hollow cylindrical shell with a lattice of square cutouts as the primary deformable structure [32], which provides reliable deformation while reducing weight and material usage. A centrally placed flexible sleeve (Fig. 2A, iii) increases stiffness against lateral forces F_x and F_y , while remaining compliant to F_z . This selective stiffening balances deflection, suppresses undesired lateral motion, and mitigates oscillations induced by gripper inertia, ultimately improving stability and positioning accuracy.

The lower layer consists of two concentric rings of different radii connected by wave-like support structures, as shown in Fig. 2A(iv). The inner ring is mechanically coupled to the rigid base via a spline connection, while the outer ring is attached to the rigid circle frame and constrained by a circle rail. When an external torsional moment T_z is applied, the inner and outer rings undergo relative rotation, producing an angular displacement that is quasi-linear to the applied torque. This configuration enables a relatively stable and measurable torsional deformation.

2) *Fabrication:* All components can be fabricated using a commercial-level 3D printer and assembled with only a few screws. The compliant core is printed with thermoplastic polyurethane (TPU), and the rigid connectors are made of polylactic acid (PLA), with a total material consumption of approximately six U.S. dollars (excluding the camera).

B. Deformation Sensing and Force Representation

In this section, we introduce how the deformation of the compliant core is captured and transformed into force-like signals. Theoretically, we can denote that the deformation of the compliant core as $\Delta T \in SE(3)$, which correlates with the external force F_{ext} and torque T_{ext} applied at the end-effector. There exists an approximately linear mapping:

$$\phi : \Delta T \leftrightarrow (F_{\text{ext}}, T_{\text{ext}}) \quad (1)$$

Benefiting from the hard-soft-hard sandwich structure, the state of the compliant core is fully characterized by the relative pose between its two rigidly connected ends. Because the robot flange and the fiducial tag are rigidly attached to the upper and lower connectors, respectively, their poses naturally serve as reference frames to compute this transform.

Let ${}^F T_{T,0} \in SE(3)$ denote the initial transform from the flange frame to the tag frame in the absence of external loads, and ${}^F T_{T,t} \in SE(3)$ the corresponding transform at time step t under external forces and torques. The resulting deformation transform is then computed as:

$$\Delta T_t = ({}^F T_{T,0})^{-1} {}^F T_{T,t}, \quad (2)$$

where $\Delta T_t \in SE(3)$ compactly represents the relative motion of the connectors induced by external loading. Estimating the external wrench thus reduces to estimating ${}^F T_{T,t}$.

In practice, the tag pose is first obtained in the camera frame, ${}^{\text{cam}} T_T$, using OpenCV's ChArUco marker detection. Because the camera is mounted on the robot flange, we use hand-eye calibration to obtain the extrinsic matrix K_E , which maps the camera frame to the flange frame:

$${}^F T_{\text{cam}} = K_E.$$

The tag pose in the flange frame can then be expressed as:

$${}^F T_{T,t} = {}^F T_{\text{cam}} {}^{\text{cam}} T_T. \quad (3)$$

To ensure that the signal is physically meaningful, we align the translation of ${}^F T_{\text{cam}}$ with the geometric center of the lower connector and its rotation with the end-effector pose using a single rigid-body transformation from the CAD model. The resulting transform $\Delta T_t \in SE(3)$ is then decomposed into translation and Euler angles $\Delta \mathbf{T}_t \in \mathbb{R}^6$ to obtain a 6D representation directly correlated with the force-torque components. This provides an intuitive and physically consistent description of the compliant core deformation state, which is subsequently used as a force-like signal for contact-aware control and policy learning.

C. Integrating Force-like Signals into Policies

Given the deformation signal $\Delta \mathbf{T}_t$, we aim to validate its utility across a wide range of policies for contact-rich manipulation. To this end, we consider two representative paradigms: classical **search-and-control policy**, which combines search strategies and control algorithms using $\Delta \mathbf{T}_t$ as both a state indicator and a reference signal. **learning-based policy**, which integrates $\Delta \mathbf{T}_t$ with other observations as an additional input to enhance policy training.

1) *Search-and-Control Policy*: Prior works on contact-rich manipulation [33], [34] employ a wide range of search and control strategies, including random and spiral search trajectories, tilt-based alignment methods, and hybrid position/force controllers. These approaches remain widely used in real-world assembly and have proven effective in solving contact-rich tasks. We adopt several of these classical methods as the skill library of our search-and-control policy, denoted by $\pi = \{\pi_1, \pi_2, \dots, \pi_n\}$, where each π_i represents a standard search or control skill. Our force-like signal $\Delta \mathbf{T}_t$ is then used either as a trigger for policy switching via simple thresholding, or as a reference for force control to regulate and adapt the robot's actions.

2) *Learning-Based Policy*: Recent advances [35]–[37] in learning-based methods, such as imitation learning, have shown strong capability for contact-rich tasks by learning from demonstrations, thereby reducing the need for manually tuning parameters and thresholds. We adopt popular imitation learning policies, such as the Action Chunking Transformer (ACT) [38] and the Diffusion Policy (DP) [39], as the backbone and train a contact-rich policy with our force-like signal $\Delta \mathbf{T}_t$ in conjunction with other observations. We then compare it against the use of a force–torque sensor or the absence of contact sensing, to further evaluate its performance in learning-based tasks.

During training the imitation policy, the policy parameters θ are optimized to minimize the discrepancy between the predicted actions \hat{a}_t and the expert demonstrations a_t :

$$\hat{a}_t = \pi_{\theta}(\Delta \mathbf{T}_t, I_t, \mathbf{o}_t^{\text{prop}}), \quad (4)$$

$$\theta^* = \arg \min_{\theta} \ell(\hat{a}_t, a_t), \quad (5)$$

where $\ell(\cdot)$ is a supervised imitation loss, I_t is the visual image of the wrist camera, $\mathbf{o}_t^{\text{prop}}$ is the proprioception of the robotic arm, such as joint angles, and $\Delta \mathbf{T}_t$ is our force-like signal embedded by a linear layer.

Once trained, the policy uses the optimized parameters θ^* during deployment to output actions directly:

$$a_t = \pi_{\theta^*}(\Delta \mathbf{T}_t, I_t, \mathbf{o}_t^{\text{prop}}). \quad (6)$$

Further details of the deployment and performance of the learning-based policy can be found at Section IV-B.

IV. EXPERIMENT

In this section, we perform a series of experiments aimed at answering the following questions.

- Q1: Can ShapeForce provide a high-quality force-torque representation, as well as sufficient sensitivity and stability for contact-rich manipulation? (Sec. IV-A)
- Q2: How does ShapeForce perform in real-world contact-rich tasks with different tasks and policies? (Sec. IV-B)
- Q3: Can ShapeForce maintain long-term durability, and ensure that pretrained policies perform consistently under extensive use? (Sec. IV-C)

A. Mechanical Properties and Performance

To validate the force-torque representation quality, sensing sensitivity as well as stability of the mechanical design, we conduct the following experiments:

1) *Correlation between Wrench and Deformation*: The relationship between the wrench (force and torque) \mathbf{F}_t and the deformation-based force-like signal $\Delta\mathbf{T}_t$ generated by ShapeForce can be approximated by a linear stiffness model:

$$\mathbf{F}_t \approx K \Delta\mathbf{T}_t, \quad (7)$$

where $K \in \mathbb{R}^{6 \times 6}$ is the equivalent stiffness matrix of the compliant core, and $\mathbf{F}_t = (F_x, F_y, F_z, T_x, T_y, T_z)^T \in \mathbb{R}^6$ is the six-dimensional ground-truth force and torque.

To validate this approximation, we mounted a commercial force and torque sensor in series with the ShapeForce sensor to collect synchronized paired data. By recording \mathbf{F}_t from a commercial force and torque sensor and force-like signal $\Delta\mathbf{T}_t$ from ShapeForce, we estimated the equivalent stiffness matrix K via linear regression. We report the coefficient of determination (R^2) for each wrench component ($F_x, F_y, F_z, T_x, T_y, T_z$) in Table I. The mean R^2 of 0.9577 demonstrates that the general relationship between $\Delta\mathbf{T}_t$ and the ground-truth wrench is highly linear. This near-linear correspondence suggests that our representation can serve as an effective proxy for ground-truth force-torque signals: in the context of search-and-control policies, it provides contact feedback that correspond proportionally to real interactions, enabling accurate event triggering and reference tracking; when applied within learning-based policies, it can be embedded through a simple linear layer to convey temporal contact information, making it a practical, highly cost-efficient alternative to physical force-torque sensors.

TABLE I: R^2 between $\Delta\mathbf{T}_t$ and Ground-truth Wrench.

	F_x	F_y	F_z	T_x	T_y	T_z
R^2	0.962	0.9681	0.9587	0.9739	0.9675	0.9159

Moreover, if ground-truth wrench values are required, our system can reconstruct them by applying a straightforward linear transformation based on the estimated stiffness matrix K . A visualization example is presented in Figure 3, demonstrating that our method can also reconstruct ground-truth force values via a simple linear transformation with high accuracy, achieving high fidelity in practice.

2) *Sensing Sensitivity and Stability*: Sensitivity is a key metric for evaluating force-torque sensor performance. Since ShapeForce relies on an RGB camera for perception, its sensitivity is fundamentally limited by the visual resolution. To evaluate this for ShapeForce, we refer to the fiducial-based sensitivity analysis in prior work [40] to calculate our theoretical sensitivity. First, the pixel-level resolution d_R is set to 1/4 based on prior work. Using the physical tag width w_{tag} (mm) and its image width w_{img} (px), the translational sensitivities are computed as:

$$s_z = s_x = \frac{w_{\text{tag}}}{w_{\text{img}}} d_R \quad (\text{mm}).$$

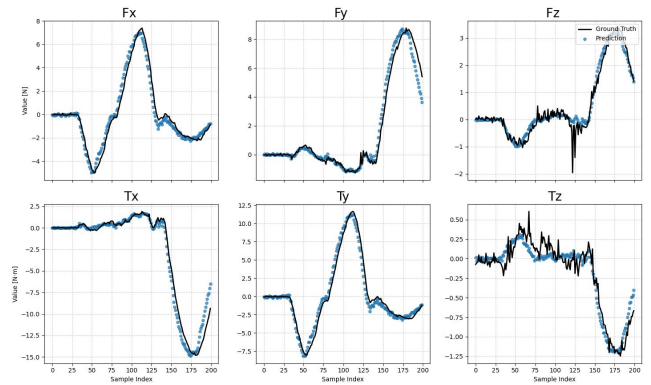


Fig. 3: **Visualization of predicted vs. ground-truth signals**: the blue dots represent ShapeForce’s predictions, and the black line represents ground-truth data.

And rotational sensitivity s_{θ_y} is obtained from the chord-length relation:

$$l_{\text{chord}} = 2r \sin\left(\frac{\theta}{2}\right),$$

where r is the half-diagonal of the image patch, and set θ to be 45 degrees, and the minimum detectable angle s_{θ} satisfies:

$$s_{\theta} = \frac{\theta}{l_{\text{chord}}} d_R$$

The axial sensitivity s_{θ_z} , s_{θ_x} , and s_y are estimated via similar-triangle geometry. Stacking the translational and rotational sensitivities into $\mathbf{s} = [s_x, s_y, s_z, s_{\theta_x}, s_{\theta_y}, s_{\theta_z}]^T$, the minimum detectable wrench is estimated by:

$$\mathbf{F}_{\min} \approx K \mathbf{s},$$

where $K \in \mathbb{R}^{6 \times 6}$ is the calibrated stiffness matrix mentioned before. Our result shows that ShapeForce achieves a minimum detectable wrench of:

$$\mathbf{F}_{\min} = [0.41, 0.45, 0.87, 0.13, 0.12, 0.03]^T \text{ N or N} \cdot \text{m},$$

The \mathbf{F}_{\min} results showcase the sensitivity of our system to subtle forces and torques, and this result is align closely with real-world measurements, which proves that our design is capable of detecting fine-grained contact information to support a wide range of delicate contact-rich tasks.

Also, we evaluated the mechanical stability by measuring the lateral deflection of the end-effector’s tip relative to the central axis under a static payload. We mounted a payload of 0.8kg and the lever arm of a force is 6.8 cm to simulate the gravity effect of grasping an object and the end-effector. The result is shown in Table II. The maximum lateral deflection at the tip under a 0.8kg load is 3.312 mm. The stability results show that the soft wrist provides sufficient stiffness to support the payload and maintain accurate motion.

Overall, these results indicate that ShapeForce provides both high sensitivity, allowing the detection of subtle forces and torques, and sufficient mechanical stability to maintain precise motion under realistic payloads. Together, these properties ensure that the system delivers reliable and physically meaningful feedback for a variety of contact-rich tasks.

TABLE II: Displacement under Different Angles

Orientation Angle (°)	30	45	60	90
Value(mm)	1.497	2.418	2.857	3.312




Fig. 4: **Robot Experiment Setup:** We use Ufactory xArm7 robotic arm for manipulation and an Intel RealSense D435 for visual perception. ShapeForce is mounted on the robotic wrist to obtain the contact information.

B. Effectiveness in Real-World Manipulation Tasks

To directly validate ShapeForce’s task performance, we evaluate performance across six representative contact-rich manipulation tasks with the aforementioned two types of policies: search-and-control policy and learning-based policy. The details of the implementation and results are as follows:

1) *Search-and-Control Policy:* In this setting, we use ShapeForce as a six-axis force–torque sensor to provide contact feedback for search-and-control policies. For insertion tasks, we employ search-based algorithms to locate the target, while classical controllers such as PID are used to track desired forces during tasks like whiteboard wiping. Other state transitions are triggered by simple threshold-based events.

For each task, we evaluate the performance of ShapeForce and a 6-axis force-torque sensor from UFactory, which provides ground-truth measurements. Performance is measured by success rate, with outcomes assessed by a human expert due to the explicit nature of the tasks. We do not compare with a setting without any contact feedback, as search-and-control policies cannot be designed without it, and the success rate ought to be zero. We evaluate each both setup in task for 20 times to ensure statistical reliability.

Results are summarized in Table III. Our system achieves a success rate comparable to that of a 6-axis force-torque sensor, demonstrating that the proposed force-like signal can serve as an effective substitute for ground-truth measurements. Unlike commercial sensors, however, our approach requires no expensive hardware and is seamlessly integrated into the robot, offering a low-cost yet highly reliable alternative for search-and-control manipulation policies.

The only noticeable drop occurs in the USB insertion,

TABLE III: Task Success Rate of Search-and-Control Policy

Task	With F/T Sensor	ShapeForce
Peg Insertion	100%	100%
USB Insertion	100%	90%
Toy Desk Assembly	65%	80%
Bottom Cap Screw Tightening	100%	100%
Whiteboard Wiping	100%	100%
Maze Exploration	100%	100%

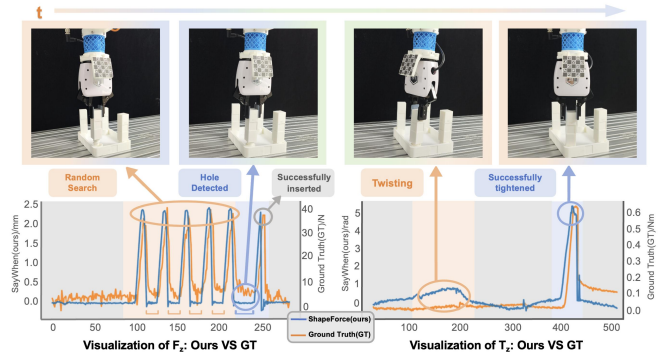


Fig. 5: **Visualization of Toy Desk Assembly:** Comparison of the F_z and T_z components of our force-like signal with ground-truth measurements. After proportional scaling, the signals exhibit close alignment, confirming that our representation preserves the essential contact information.

where the inherent compliance of our design introduces slight uncertainty. Despite our efforts to mitigate this effect, tasks that demand extreme precision in contact detection, such as USB insertion, remain more challenging. In particular, distinguishing between the correct and incorrect orientations during insertion requires detecting differences within a tolerance of about 1 mm, which can be easily falsely triggered by noise or compliance-related uncertainty.

However, our system outperforms the 6-axis force-torque sensor in the toy desk assembly task, which requires highly precise alignment of the insertion direction to properly engage the threads. Here, the compliant nature of our design becomes advantageous: even with slight orientation errors, the passive deformation of the wrist compensates for misalignment, allowing smooth insertion and tightening. In contrast, a rigid wrist would cause large contact forces under the same condition, often leading to slippage between the leg and the gripper and ultimately task failure. This highlights the inherent robustness provided by passive compliance. Overall, ShapeForce performs a comparable performance with force-torque sensor in classical search-and-control policy.

2) *Learning-based Policy:* To validate that our force-like signal ΔT_t can be applied in learning-based policy therefore further prove that it can be an effective alternative of ground-truth force-torque signals, we adopt two imitation learning policies Action Chunking Transformer (ACT) [38] and the Diffusion Policy (DP) [39] as our backbones to evaluate. Importantly, the same wrist-mounted RGB camera serves dual purposes: it captures ChArUco images for deformation sensing (ΔT_t) and provides raw workspace images as visual

TABLE IV: Task Success Rate of Learning-based Policy

Task	Without Sensor	With F/T Sensor	ShapeForce
Peg Insertion (ACT)	60%	85%	90%
Peg Insertion (DP)	65%	75%	85%
USB Insertion (ACT)	40%	85%	75%
USB Insertion (DP)	35%	80%	75%
Whiteboard Wiping (ACT)	25%	95%	95%
Whiteboard Wiping (DP)	30%	90%	90%
Toy Desk Assembly (ACT)	45%	65%	80%
Toy Desk Assembly (DP)	40%	50%	70%

input to ACT and DP policies (Fig. 4).

We construct a 15-step history of $\Delta\mathbf{T}_t$, encode it via a linear layer, and concatenate it with other observations to train the policies, enhancing their contact-awareness. Expert demonstrations (50 trajectories per task) are collected via Gello teleoperation [41]. During data collection and evaluation, we add small random variations to object poses to improve robustness. We benchmark on four tasks: peg insertion, USB insertion, whiteboard wiping, and toy desk assembly, comparing against two baselines: (i) policies without force input, and (ii) policies with ground-truth force-torque signals.

The results of our experiments are summarized in Table IV. Across all tasks, policies trained without any force or deformation signal perform significantly worse. These policies can only replicate demonstrated motion patterns without perceiving contact states, which causes them to learn an “average” behavior that fails under dynamic variations. This limitation is particularly evident in tasks requiring sustained contact, such as whiteboard wiping. In this task, visual occlusion from the eraser often prevents the policy from accurately estimating the surface distance by image. Meanwhile, the board’s orientation and height distance vary across trials. Successfully wiping requires consistent contact pressure. Without force-related signals, success rates for both ACT and DP policies drop below 30%. In contrast, when incorporating either ground-truth force or our deformation signal $\Delta\mathbf{T}_t$, both policies exceed 90% success rate. This huge leverage demonstrate the critical role of force-related observations for contact-rich imitation learning.

Comparing ShapeForce with the ground-truth sensor, our deformation signal $\Delta\mathbf{T}_t$ achieves success rates comparable to those of the force sensor, further demonstrating the linear relationship between $\Delta\mathbf{T}_t$ and the wrench, which is approximately equivalent for the learning-based method. And in some cases, like the toy desk assembly task, ShapeForce enables the ACT policy to reach 80% success, compared to 65% with the force sensor, highlighting the advantages of passive compliance. Overall, these results show that ShapeForce generalizes well in imitation learning and serves as an effective, low-cost alternative to force–torque sensors.

C. Durability and Sensing Consistency

As a low-cost sensing system made from 3D-printed compliant materials, durability is crucial for practical use, so we performed an aging test to evaluate the durability of the sensor. We fix the lower connector to a fixture and move the robotic arm to apply force in various directions to

TABLE V: Task Success Rate after Aging Test

Task	1k	5k	10k	20k
Whiteboard Wiping (S&C)	100%	100%	100%	95%
Toy Desk Assembly (S&C)	80%	80%	80%	80%
Peg Insertion (ACT)	85%	85%	85%	85%
Toy Desk Assembly (ACT)	80%	80%	60%	60%
Toy Desk Assembly (ACT-finetune)	85%	80%	80%	80%

the compliant core. Generally, aging of the elastic material can cause some plastic deformation and degradation of its linear response characteristics. To comprehensively evaluate the influence, we test parameter-tuned search-and-control policies and a pretrained ACT model after 1k, 5k, 10k, and 20k aging cycles. Each cycle pushes the sensor against a rigid fixture, generating forces up to 50N and visibly deforming the elastic structure at 1 Hz.

Results are shown in Table V. For search-and-control policies, performance was evaluated on whiteboard wiping and toy desk assembly. Stiffness degradation lowers the force threshold, reducing the applied force under the same conditions. This causes a minor drop in whiteboard wiping performance after 20k repetitions, but the magnitude remains small, demonstrating the robustness and consistency of our system under search-and-control policy. For learning-based policy, we tested the pretrained ACT model directly on peg insertion and toy desk assembly. The result demonstrates that ACT can perform consistently in relatively simple tasks like peg insertion, but in complicated tasks like toy desk assembly, our pretrained policy fails due to the degradation of linearity and elastic property. However, fine-tuning with just five demonstrations quickly restores performance to near-new levels. In a nutshell, the aging test shows ShapeForce maintains high durability and policy consistency under extensive use, with rapid recovery through minimal fine-tuning.

V. CONCLUSION AND DISCUSSION

In this paper, we introduce ShapeForce, a low-cost, plug-and-play soft robotic wrist that provides force-like signals enabling contact-rich tasks. We proved the efficiency of this force-like signal through extensive experiments on various contact-rich tasks with different policies, where ShapeForce significantly improves success rates and is comparable to commercial force-torque sensors. This work not only offers a cost-efficient, easily reproducible alternative to traditional force-torque sensors but also provides a fundamental insight: the trade-off between accuracy and simplicity in sensing systems can be effectively optimized by leveraging task-relevant features to construct an equivalent representation.

However, ShapeForce has some limitations. Its passive compliance makes precise motion difficult under high payloads, and its material properties are not fully optimized. These limitations stem from its original design goals of being low-cost, accessible, and reproducible. Despite these limitations, we hope this work will inspire more researchers to study further and explore the contact-rich tasks.

REFERENCES

- [1] L.-P. Chao and K.-T. Chen, "Shape optimal design and force sensitivity evaluation of six-axis force sensors," *Sensors and Actuators A: Physical*, vol. 63, no. 2, pp. 105–112, 1997.
- [2] M. Y. Cao, S. Laws, and F. R. y Baena, "Six-axis force/torque sensors for robotics applications: A review," *IEEE Sensors Journal*, vol. 21, no. 24, pp. 27238–27251, 2021.
- [3] Z. Lu, P. C. Chen, and W. Lin, "Force sensing and control in micro-manipulation," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 36, no. 6, pp. 713–724, 2006.
- [4] S. Wang, Y. She, B. Romero, and E. Adelson, "Gelsight wedge: Measuring high-resolution 3d contact geometry with a compact robot finger," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6468–6475.
- [5] S. Wang, M. Lambeta, P.-W. Chou, and R. Calandra, "Tacto: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3930–3937, 2022.
- [6] V. Pattabiraman, Z. Huang, D. Panozzo, D. Zorin, L. Pinto, and R. Bhirangi, "eflesh: Highly customizable magnetic touch sensing using cut-cell microstructures," *arXiv preprint arXiv:2506.09994*, 2025.
- [7] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, "Cable manipulation with a tactile-reactive gripper," *The International Journal of Robotics Research*, 2021.
- [8] M. Neto, P. Ribeiro, R. Nunes, L. Jamone, A. Bernardino, and S. Cardoso, "A soft tactile sensor based on magnetics and hybrid flexible-rigid electronics," *Sensors*, vol. 21, no. 15, p. 5098, 2021.
- [9] C. Dong, X. Dai, Y. Pan, W. Qiu, S. Li, T. Wu, Y. Jin, H. Wang, C. Song, and F. Wan, "Teaching oral care via vision-based deformable perception," *Soft Science*, vol. 5, no. 3, p. 36, 2025.
- [10] N. Guo, X. Han, S. Zhong, Z. Zhou, J. Lin, J. Dai, F. Wan, and C. Song, "Proprioceptive state estimation for amphibious tactile sensing," *IEEE Transactions on Robotics*, vol. 40, pp. 4684–4698, 2024.
- [11] X. Liu, X. Han, W. Hong, F. Wan, and C. Song, "Proprioceptive learning with soft polyhedral networks," *The International Journal of Robotics Research*, vol. 43, no. 12, pp. 1916–1935, 2024.
- [12] S. Li, C. Dong, C. Song, and F. Wan, "Activespn: Active soft polyhedral networks with pose estimation for in-finger object manipulation," *IEEE Robotics and Automation Letters*, vol. 0, no. 0, p. 0, 2025, accepted.
- [13] M. Peticco, G. Ulloa, J. Marangola, and P. Agrawal, "Dexwrist: A robotic wrist for constrained and dynamic manipulation," *arXiv preprint arXiv:2507.01008*, 2025.
- [14] F. Negrello, S. Mghames, G. Grioli, M. Garabini, and M. G. Catalano, "A compact soft articulated parallel wrist for grasping in narrow spaces," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3161–3168, 2019.
- [15] H. Sun, S. Park, and D. Hwang, "Compact modular robotic wrist with variable stiffness capability," *IEEE Transactions on Robotics*, 2024.
- [16] S. Lemerle, M. G. Catalano, A. Bicchi, and G. Grioli, "A configurable architecture for two degree-of-freedom variable stiffness actuators to match the compliant behavior of human joints," *Frontiers in Robotics and AI*, vol. 8, p. 614145, 2021.
- [17] S. Taheri, A. Yousefi-Koma, and F. Yazdani, "Design, analysis, and construction of a 2rss-1u compact, singularity free humanoid robot wrist mechanism," in *2023 11th RSI International Conference on Robotics and Mechatronics (ICRoM)*. IEEE, 2023, pp. 897–903.
- [18] 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*, 2025.
- [19] S. Kurumaya, B. T. Phillips, K. P. Becker, M. H. Rosen, D. F. Gruber, K. C. Galloway, K. Suzumori, and R. J. Wood, "A modular soft robotic wrist for underwater manipulation," *Soft robotics*, vol. 5, no. 4, pp. 399–409, 2018.
- [20] 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 *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 8752–8757.
- [21] Z. Zhang, G. Chen, W. Fan, W. Yan, L. Kong, and H. Wang, "A stiffness variable passive compliance device with reconfigurable elastic inner skeleton and origami shell," *Chinese Journal of Mechanical Engineering*, vol. 33, no. 1, p. 75, 2020.
- [22] B. Budiansky, "Theory of buckling and post-buckling behavior of elastic structures," *Advances in applied mechanics*, 1974.
- [23] F. Wan, X. Liu, N. Guo, X. Han, F. Tian, and C. Song, "Visual learning towards soft robot force control using a 3d metamaterial with differential stiffness," in *Conference on Robot Learning*. PMLR, 2022, pp. 1269–1278.
- [24] Y. Zhu, M. Hao, X. Zhu, Q. Bateux, A. Wong, and A. M. Dollar, "Forces for free: Vision-based contact force estimation with a compliant hand," *Science Robotics*, vol. 10, no. 103, p. eadq5046, 2025.
- [25] N. Guo, X. Han, S. Zhong, Z. Zhou, J. Lin, F. Wan, and C. Song, "Reconstructing soft robotic touch via in-finger vision," *Advanced Intelligent Systems*, vol. 6, no. 10, p. 2400022, 2024.
- [26] W. Xu, G. Zhou, Y. Zhou, Z. Zou, J. Wang, W. Wu, and X. Li, "A vision-based tactile sensing system for multimodal contact information perception via neural network," *IEEE Transactions on Instrumentation and Measurement*, 2024.
- [27] T. Wu, Y. Dong, X. Liu, X. Han, Y. Xiao, J. Wei, F. Wan, and C. Song, "Vision-based tactile intelligence with soft robotic metamaterial," *Materials & Design*, vol. 238, p. 112629, 2024.
- [28] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer et al., "Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation," *IEEE Robotics and Automation Letters*, 2020.
- [29] R. Bhirangi, T. Hellebrekers, C. Majidi, and A. Gupta, "Reskin: versatile, replaceable, lasting tactile skins," *arXiv preprint arXiv:2111.00071*, 2021.
- [30] R. Bhirangi, V. Pattabiraman, E. Erciyes, Y. Cao, T. Hellebrekers, and L. Pinto, "Anyskin: Plug-and-play skin sensing for robotic touch 2024," URL <https://arxiv.org/abs/2409.08276>.
- [31] B. Ward-Cherrier, N. Pestell, L. Cramphorn, B. Winstone, M. E. Giannaccini, J. Rossiter, and N. F. Lepora, "The tactip family: Soft optical tactile sensors with 3d-printed biomimetic morphologies," *Soft robotics*, vol. 5, no. 2, pp. 216–227, 2018.
- [32] A. Seharing, A. H. Azman, and S. Abdullah, "A review on integration of lightweight gradient lattice structures in additive manufacturing parts," *Advances in Mechanical Engineering*, vol. 12, no. 6, p. 1687814020916951, 2020.
- [33] S. R. Chhatpar and M. S. Branicky, "Search strategies for peg-in-hole assemblies with position uncertainty," in *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)*, vol. 3. IEEE, 2001, pp. 1465–1470.
- [34] D. Austin and B. McCarragher, "Force control command synthesis for assembly using a discrete event framework," in *Proceedings of International Conference on Robotics and Automation*, vol. 2. IEEE, 1997, pp. 933–938.
- [35] W. Liu, J. Wang, Y. Wang, W. Wang, and C. Lu, "Forcemimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 1105–1112.
- [36] H. Xue, J. Ren, W. Chen, G. Zhang, Y. Fang, G. Gu, H. Xu, and C. Lu, "Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation," *arXiv preprint arXiv:2503.02881*, 2025.
- [37] Y. Hou, Z. Liu, C. Chi, E. Cousineau, N. Kuppusswamy, S. Feng, B. Burchfiel, and S. Song, "Adaptive compliance policy: Learning approximate compliance for diffusion guided control," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 4829–4836.
- [38] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *arXiv preprint arXiv:2304.13705*, 2023.
- [39] 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.
- [40] R. Ouyang and R. Howe, "Low-cost fiducial-based 6-axis force-torque sensor," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1653–1659.
- [41] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, "Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024.