MoiréSkin: Ultra-Sensitive Sensor–Actuator Visuo-Tactile Skin using Moiré Patterns

Linzhen Zhu
University of Michigan
Ann Arbor, MI, USA
lzzhu@umich.edu

Liuxi Xing
University of Michigan
Ann Arbor, MI, USA
xingliux@umich.edu

Runqiu Wang
University of Michigan
Ann Arbor, MI, USA
runqiu@umich.edu

Ke Sun
University of Michigan
Ann Arbor, MI, USA
kesuniot@umich.edu

Abstract

Robots increasingly must perform dexterous object manipulation in real-world environments, yet existing VTS offer rich camera-based spatial readout while rigid optics limit force sensitivity, force range, and active compliance regulation. We present `MoiréSkin`, an inflatable visuo-tactile skin using flexible Moiré optics to unify ultra-sensitive tactile sensing and active compliance regulation. Dual-sided submillimeter gratings in a soft membrane amplify microscopic contact-induced strain into large, high-contrast Moiré fringe variations for tactile readout. To balance reliability and sensitivity, our tailored soft-material grating pipeline combines internal illumination with fluorescent doping to suppress ambient-light interference. Self-supervised calibration compensates for air-pressure-dependent sensing drift during pneumatic regulation, while a neural pipeline uses pre- and post-contact fringes to estimate contact position and force. Experiments show that `MoiréSkin` achieves 0.094 mm position and 0.0038 N contact force MAE, outperforming prior VTS by an order of magnitude. It also reaches 99.7% 10-shot classification accuracy over 7 ultra-soft materials spanning 8.73-21.72 kPa Young's modulus, and measures shear-force-related liquid mass variation with ~4 g precision, demonstrating a new path toward ultra-sensitive sensor-actuator visuo-tactile skins.

CCS Concepts

• **Hardware** → **Sensors and actuators**; • **Computer systems organization** → *Robotic components*; • **Computing methodologies** → *Vision for robotics.*

Keywords

Visuo-tactile sensor, Moiré patterns, Contact-rich manipulation, Sensor–actuator integration

MobiCom '26, Austin, TX, USA

ACM ISBN 979-8-4007-2505-0/2026/10
https://doi.org/10.1145/3795866.3844463

ACM Reference Format:
Linzhen Zhu, Liuxi Xing, Runqiu Wang, and Ke Sun. 2026. MoiréSkin: Ultra-Sensitive Sensor–Actuator Visuo-Tactile Skin using Moiré Patterns. In *The 32nd Annual International Conference on Mobile Computing and Networking (MobiCom '26), October 26–30, 2026, Austin, TX, USA.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3795866.3844463

1 Introduction

Robots are increasingly expected to perform everyday dexterous object manipulation in real-world environments, e.g., picking up fragile groceries, opening containers, inserting connectors, and handling deformable packaging. Yet these tasks are ultimately governed by *contact*: forces depend on geometry and frictional surface properties, contact patches evolve during motion, and brief transients can turn a seemingly correct grasp into slip or damage. Achieving reliable dexterity therefore requires closing the loop at the interface: robots must not only *sense* contact forces and slip risk, but also actively *shape* contact conditions in real time.

Human fingertips illustrate the level of precision and reactivity required. First, *dense tactile innervation* provides rich, localized feedback: a single fingertip is innervated by roughly 10^3 low-threshold tactile afferents, with mechanoreceptive unit densities reaching ~241 units/cm^2 [7, 19]. Second, *fine-grained force sensing* supports manipulation in the low-Newton regime: humans can scale and discriminate forces typical of manipulation (~0.1-5 N) [51], where changes on the order of *tens of mN* can be behaviorally meaningful for grasp stability. Third, *active compliance regulation* stabilizes grasps under friction uncertainty: during precision grip, humans maintain grip force only slightly above the slip threshold rather than relying on conservative over-squeezing [2, 50]. Together, these properties motivate a clear design target for robotic tactile skin: localized, high-sensitivity sensing paired with contact-aware actuation.

Traditional resistive/capacitive tactile sensor arrays are difficult to scale to dense, fine-grained tactile skins as higher density increases wiring/packaging complexity, reduces conformability, and exacerbates shear- and wear-induced drift [9]. Recent Visuo-Tactile Sensing (VTS) instead uses cameras to infer contact properties from surface deformation. However,

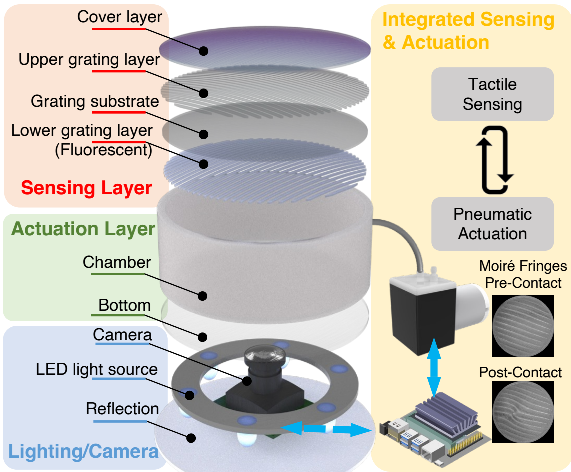

Figure 1: `MoiréSkin` architecture. It embeds Moiré patterns into a flexible film to enable ultra-sensitive integrated visuo-tactile sensing and actuation.

existing VTS designs suffer from a fundamental *sensitivity-compliance paradox*. Rigid VTS offers high spatial resolution but relies on stiff elastomers that saturate under large loads and lack the compliance needed for safe, irregular contact [42, 59]. Conversely, non-rigid VTS adopts compliant, air-filled interfaces but is bottlenecked by low spatial resolution, often remaining at sub-Newton precision due to the limits of sparse marker-tracking [23]. Furthermore, both classes decouple tactile sensing from active compliance regulation, limiting integrated sensor-actuator capabilities for dexterous manipulation. Tab. 1 summarizes their limitations.

To address these, we propose `MoiréSkin`, a novel non-rigid inflatable visuo-tactile skin based on Moiré-based sensing, designed for four unique capabilities: sub-gram-level sensitivity, air-inflated compliance, integrated sensing and actuation, and see-through visual access before contact (Tab. 1). This combination makes `MoiréSkin` particularly suitable for gram-scale loads (Fig. 14), fragile or deformable objects (Sec. 7.3.1), interactions with changing loads that require active pressure regulation (Sec. 7.3.2), and tasks that benefit from observing an object immediately before contact (Sec. 8).

Our key idea is to embed *dual-sided, sub-millimeter optical gratings* into a flexible film to generate *Moiré fringes*, low-frequency beat patterns arising from microscopic geometric mismatches. These fringes provide massive geometric amplification, magnifying tiny relative grating displacements into visible Moiré fringe distortions during external contact or internal pneumatic actuation. By decoding these distortions, `MoiréSkin` estimates both contact force and actuation state, enabling closed-loop control that can simultaneously *measure* and *shape* contact during manipulation.

`MoiréSkin` faces four technical challenges.

How to balance sensor sensitivity with fabrication reliability? `MoiréSkin` relies on sub-millimeter optical gratings whose pitch and contrast must remain uniform after transfer onto a compliant substrate. The two grating layers must be precisely aligned to generate stable Moiré fringes. We design a precise simulation tool to characterize the key fabrication parameters and guide the design space. Based on these insights, we establish a fabrication pipeline that preserves grating fidelity and alignment at scale while maintaining mechanical durability. We further tailor the soft material properties to achieve a favorable balance among compliance, deformability, stretchability, and optical transparency.

How to make `MoiréSkin` *robust under ambient lighting conditions without losing "see-through" functionality?* Existing structured-illumination-based VTS typically use an opaque cover layer to shield the sensing region from ambient light interference, which prevents the camera from observing the object before contact. In contrast, `MoiréSkin` aims to preserve "see-through" capability by using a transparent cover layer. However, this design introduces the challenge of ensuring that the observed Moiré fringes are not corrupted by the appearance of the contact object. We develop a fluorescence-assisted fabrication process in which one high-frequency grating excites fluorescent emission, thereby enhancing the contrast of the generated Moiré fringes under diverse ambient lighting conditions. In addition, we co-design the internal illumination to produce a uniform light field while minimizing glare and lens flare.

How to sense and calibrate the actuation state across different atmospheric pressure? `MoiréSkin` uses pneumatic actuation to enable closed-loop control. Without accurate real-time estimation of the actuation state, such a control loop cannot reliably achieve the desired deformation. However, accurate state estimation is challenging because both material properties and atmospheric pressure significantly affect the actuation response. To address it, we first perform a one-time material calibration by learning a Moiré-fringe-to-actuation-state conversion model. We then propose a self-supervised calibration algorithm that automatically identifies the zero-actuation state, i.e., when the internal chamber pressure equals the external atmospheric pressure.

How to achieve accurate contact position and force sensing across different actuation states? The Moiré fringe response depends not only on the contact itself, but also on the instantaneous actuation state of `MoiréSkin`. To address this challenge, we design a neural network framework that captures local Moiré fringe variations by taking the pre-contact and post-contact fringe patterns as input. A two-stage model is designed to refine tactile sensing results for higher spatial resolution and force estimation accuracy. Finally, we introduce a series of data augmentation strategies to improve model robustness across diverse sensing conditions.

Table 1: Benchmarking MoiréSkin against SOTA VTS (F.: Normal Force; P.: Position; Sense &Act: Integrated Sensing & Actuation Ability)

Sensor	F. MAE (N)	P. MAE (mm)	Compliance	Sense &Act
[55]	–	–	Elastic skin on rigid frame	×
[54]	–	–	Active mobile surface	✓
[36]	–	$0.4_{<0.5}$	Rigid core with cast gel	×
[59]	$1.9_{\geq 0.1}$	$0.03_{<0.1}$	Rigid base w/ soft elastomer	×
[23]	$0.8_{\geq 0.1}$	–	**Non-rigid (Air-inflated)**	×
[10]	$0.41_{\geq 0.1}$	$0.3633_{<0.5}$	Elastomer on rigid base	×
[28, 29]	$0.307_{\geq 0.1}$	$0.0462_{<0.1}$	Soft silicone on rigid base	×
[43]	$0.25_{\geq 0.1}$	–	Gratings on rigid substrate	×
[42]	$0.04_{<0.1}$	$0.96_{\geq 0.5}$	Rigid finger base	×
[44]	$0.03_{<0.1}$	$0.4_{<0.5}$	Single elastomer on cone	×
[13, 49]	$0.01_{<0.1}$	$0.2_{<0.5}$	**Non-rigid (3D dome)**	×
ours	$0.0038_{<0.01}$	$0.094_{<0.1}$	**Non-rigid (Air-inflated)**	✓

We implement an end-to-end hardware/ software prototype of MoiréSkin, as shown in Fig. 1. The hardware architecture consists of three components: (1) a sensing layer that generates Moiré fringes, (2) an actuation layer formed by a sealed chamber connected to a pneumatic actuator, and (3) a lighting and imaging layer comprising an RGB camera and LED illumination. Both the pneumatic actuator and the camera are interfaced with a shared control and computation unit, which captures the Moiré fringes and enables integrated sensing and actuation. To evaluate the system, we collect a dataset of more than 20,000 samples spanning diverse contact forces and positions. We further conduct two case studies: (1) ultra-soft material stiffness classification and (2) mass estimation via shear force, to demonstrate the benefits of integrated sensing and actuation. Extensive experiments show that MoiréSkin achieves 0.094 mm position and 0.0038 N contact force MAE. It also achieves 99.7% 10-shot classification accuracy across seven ultra-soft materials with Young's moduli ranging from 8.73 to 21.72 kPa, and measures shear-force-induced liquid mass variation with 4 g precision. Project webpage: https://ambient-intelligence-lab-umich-eecs.github.io/MoireSkin-Page/. Contributions:

- We introduce MoiréSkin, the first inflated VTS system to leverage stretchable dual-sided high-frequency gratings for generating ultra-sensitive Moiré fringes.
- We develop a specialized soft-material fabrication pipeline to ensure high resolution and ambient-light suppression.
- We propose a physics-informed framework that calibrates material properties and compensates for atmospheric-pressure variation, along with a two-stage neural inference pipeline for simultaneous contact position and force estimation.
- Extensive experiments show that MoiréSkin achieves SoTA tactile sensing performance, outperforming prior visuo-tactile sensing systems by an order of magnitude, while enabling integrated sensing and actuation.

2 RELATED WORKS

Robotic Tactile Sensing. Early robotic tactile systems mainly relied on discrete point-contact arrays using piezoresistive or capacitive transduction [8, 57]. Although effective for basic contact detection, these architectures suffered from complex wiring and limited ability to capture continuous and accurate deformation or shear-force distributions [46]. Vision-based tactile sensors (VTS) addressed these limitations by placing cameras behind deformable sensing layers and recasting tactile measurement as an image-based inference problem [3, 20, 21, 31, 40]. GelSight, for example, demonstrated high-resolution tactile sensing based on 3D surface reconstruction via photometric stereo [59], while later systems extended the design toward dense force sensing, slip detection, and compact embodiments [11, 24, 27, 42, 45, 58]. MoiréSkin preserves the benefits of high-dimensional VTS but departs from their reliance on complex structured illumination design. Instead, it uses a thin inflatable membrane with Moiré fringes to jointly enable ultra-sensitive optical tactile sensing and active compliance regulation.

Sensitivity–Compliance Trade-off in VTS. Translating contact-induced deformation into visual cues introduces a fundamental trade-off between sensitivity and compliance [6]. Optical intensity-mapping mechanisms that achieve the highest spatial resolution rely on stiff substrates to avoid out-of-focus visual cues as well as unstable or chaotic optical distortions [10, 14, 16, 29, 30, 41, 47, 59, 60]. However, such stiffness reduces physical compliance and limits the sensor's ability to safely conform to irregular objects. Alternatively, tracking the kinematic movement of discrete fiducial markers allows systems to utilize softer substrates, but their spatial resolution is strictly bottlenecked by marker density [13, 15, 26, 49, 55, 63]. MoiréSkin fundamentally addresses this sensitivity-compliance trade-off by decoupling the Moiré-based optical amplification mechanism from the mechanical stiffness of the contact medium. As a result, MoiréSkin achieves both extreme volumetric compliance and ultra-high tactile sensitivity while enabling integrated sensing and actuation, as shown in Tab. 1.

Moiré-Based Sensing. Moiré patterns have been widely used in visual sensing because small physical displacements can produce large and measurable changes in fringe patterns [4, 33–35, 38, 53, 61, 64]. These work falls into two mechanism classes. Stacked rigid gratings support tracking and tangible interaction through markers and widgets [4, 38, 53]. Camera-screen or color-filter-array interference instead encodes pose, motion, and vibration [33–35]. Moiré patterns have also supported other sensing applications: mID uses screen–camera interference for photo forensics [5], MoiréEar captures sound-induced motion [61], and MoiréLens visualizes refractive disturbances caused by airflows, heat plumes,

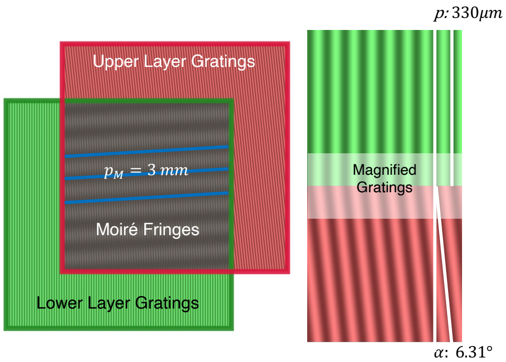

Figure 2: Moiré Fringes Generation Model.

and gas leaks [64]. Concurrent work, MoiréTac, further extends this principle to tactile sensing [43]. However, its rigid, planar structure is not well suited to capturing non-rigid surface deformation, limiting both compliance and the range of supported physical interactions. In contrast, `MoiréSkin` embeds opposing gratings in a highly compliant inflatable film, converting contact-induced deformation into tactile signals. The resulting fringe rotation also provides feedback for pneumatic state control, allowing Moiré patterns to serve as both the tactile observation channel and the feedback signal for active compliance regulation.

3 PRIMER AND PRELIMINARY STUDY

First, we introduce the basic primer underlying `MoiréSkin`, which uses distortions in Moiré fringes on a soft elastomeric film to transduce contact-induced deformation for tactile sensing. Second, we present simulation-based preliminary studies to characterize the trade-off among spatial sensitivity, deformation sensitivity, and fabrication reliability.

3.1 Moiré Fringes for Tactile Sensing

Moiré Generation. Moiré fringes are low-frequency patterns formed by superimposing two similar high-frequency gratings with a small geometric mismatch [1]. Fig. 2 illustrates the Moiré fringe generation model used in `MoiréSkin`. Specifically, `MoiréSkin` patterns identical periodic gratings on the upper and lower layers of a thin, compliant elastomeric film. The grating pitch p denotes distance between adjacent grating lines. When the two gratings are superimposed with a relative angular mismatch α, they produce low-frequency Moiré fringes with pitch

$$p_M = \frac{p}{2\sin\left(\frac{\alpha}{2}\right)}. \tag{1}$$

When α is small, the Moiré fringes p_M become much larger than the underlying grating pitch p, visually amplifying microscopic geometric mismatch into prominent large-scale fringes. Fig. 2 shows a zoomed-in example: when $p = 330\ \mu$m and $\alpha = 6.31°$, the Moiré fringe pitch is $p_M \approx 3.0$ mm.

Deformation-Induced Moiré Fringes Change. This geometric amplification makes Moiré fringes well suited for tactile sensing. When contact is applied to the compliant film, the membrane deforms locally. Because the gratings are patterned on opposite sides of the film and separated by its thickness, they do not deform identically. Their local strain, relative registration, and effective pitch can therefore diverge, perturbing the local Moiré fringe generation condition. For example, Fig. 3 shows the simulated Moiré fringe rotation under different levels of global stretching. `MoiréSkin` leverages these amplified Moiré fringes to infer tactile quantities such as contact location and contact normal/shear force.

3.2 Sensitivity Preliminary Study

Next, we conduct simulation-based studies on key design parameters, i.e., grating pitch p and grating angular mismatch α, to understand the trade-off among spatial sensitivity, deformation sensitivity, and fabrication robustness.

Spatial Sensitivity. *Smaller grating pitch p and larger angular mismatch α improve the spatial sensitivity by producing denser Moiré fringes within the same sensing area.*

Spatial sensitivity characterizes the ability of `MoiréSkin` to localize contact-induced deformation. It is mainly determined by the number of Moiré fringes captured within the camera region of interest (RoI), since denser fringes provide a finer spatial representation of local deformation. For consistent analysis, we define the default RoI as a circular area with diameter of 30 mm (~707 mm^2). Although our analysis is based on this default scale, the same principles generalize to other form factors. To reduce boundary artifacts, the compliant elastomeric film is designed to be slightly larger than the RoI, with a diameter of 35 mm. This margin mitigates edge effects and minimizes interference from the surrounding rigid frame, which could introduce nonlinear distortion.

Using Eq. 1, we analyze how p and α jointly determine the Moiré fringe pitch p_M. We sweep both parameters and convert the resulting p_M into the number of Moiré fringes within the RoI to visualize spatial sensitivity (Fig. 4 (b)). The results show that smaller p and larger α lead to higher spatial sensitivity. In practice, however, the minimum achievable grating pitch is constrained by our current soft-lithography fabrication pipeline to 330 μm (discussed in Sec. 4). We therefore fix $p = 330\ \mu$m and further visualize the generated Moiré fringes under different α in Fig. 4 (c). *With a fixed p, a higher α gives* `MoiréSkin` *better spatial sensitivity by producing denser Moiré fringes.*

Deformation Sensitivity. *Deformation sensitivity is primarily governed by the angular mismatch α. Smaller α improves deformation sensitivity by producing larger Moiré fringe distortions under the same mechanical deformation.*

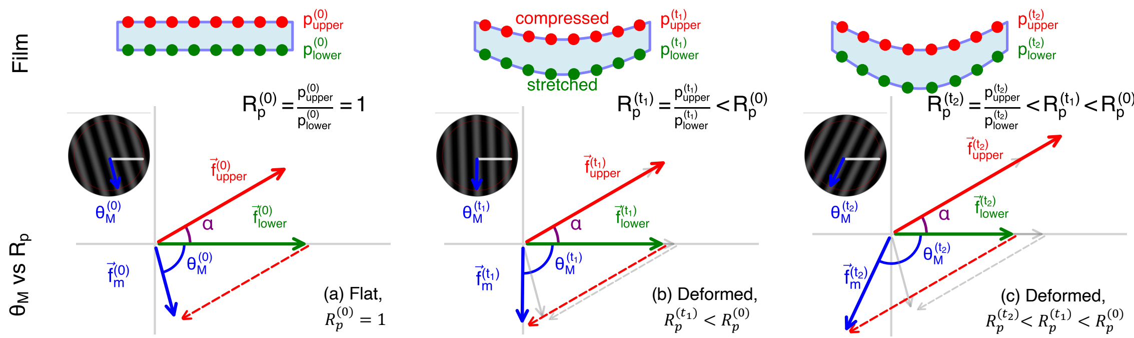

Figure 3: Relationship between Moiré fringes θ_M orientation angle and film deformation.

Deformation sensitivity characterizes the ability to resolve mechanical deformation of the elastomeric film. In our system, this sensitivity is primarily reflected by the orientation angle of the Moiré fringes, denoted by θ_M. For simplicity, we first consider global deformation, which causes a global change in θ_M across the RoI. Under localized contact, the same principle applies locally and produces spatially varying θ_M. As shown in Fig. 3, the Moiré fringe orientation θ_M can be modeled as

$$\theta_M = \arctan\left(\frac{p_{\text{upper}}\sin(\alpha)}{p_{\text{upper}}\cos(\alpha) - p_{\text{lower}}}\right), \quad (2)$$

where p_{upper} and p_{lower} denote the grating pitches of the upper and lower layers, respectively. As shown in Fig. 3 (a), Before deformation, the two pitches are identical, i.e., $p^{(0)}_{\text{upper}} = p^{(0)}_{\text{lower}}$. After deformation, however, depth separation causes the two layers to experience different strain, resulting in $p^{(t)}_{\text{upper}} \neq p^{(t)}_{\text{lower}}$ at time t (Fig. 3 (b), (c)). This pitch mismatch perturbs the Moiré fringe orientation and converts subtle deformation into amplified Moiré fringes variations.

To quantify this effect, we define the deformation sensitivity as $S_d = \frac{d\,\Delta\theta_M}{dR_p}$, where $R_p = p_{\text{upper}}/p_{\text{lower}}$ is the deformation-induced pitch ratio and $\Delta\theta_M$ is the corresponding fringe orientation change. For a given change in pitch ratio R_p, a larger variation in $\Delta\theta_M$ indicates higher deformation sensitivity. Fig. 4 (a) presents the simulated deformation sensitivity S_d. The results reveal two key insights: *(1) the geometric amplification gain increases as the angular mismatch α approaches zero; and (2) the deformation sensitivity is primarily determined by α and is decoupled from the grating pitch p.*

Sensitivity Trade-off and Design Choice. Our preliminary studies reveal a fundamental design trade-off in MoiréSkin. With the grating pitch $p = 330\ \mu$m fixed by fabrication constraints, the angular mismatch α becomes the main parameter for balancing spatial and deformation sensitivity. A larger α improves spatial sensitivity but reduces deformation sensitivity. MoiréSkin *can be made task-adaptive by selecting α during fabrication according to the desired sensing objective.* In our prototype, the commercial mold and transfer fidelity fix $p = 330\ \mu$m. We therefore choose $\alpha = 6.31°$, yielding $p_M = 3$ mm (ten fringes across the 30 mm RoI), as a balanced design point that jointly supports spatial localization and tiny-force deformation amplification (Fig. 4(c)).

4 Hardware Design

We next describe the hardware construction of MoiréSkin. As shown in Fig. 1, the sensing layer consists of multiple flexible film layers that generate high-contrast Moiré fringes, which amplifies contact-induced deformation for tactile sensing. The actuation layer contains a sealed pneumatic chamber for actuation and active compliance. The lighting/ camera layer integrates a commercial RGB camera with dedicated illumination for stable visual readout. The camera and pneumatic actuator are connected to a shared controller and computation unit, enabling integrated sensing and actuation. This section presents three key hardware components.

4.1 Materials and Fabrication Design

We first design the material stack and fabrication process to form high-fidelity Moiré fringes on a soft substrate. Our key contribution is a parametric design and controllable fabrication framework that registers high-frequency gratings on opposite sides of a single compliant film while preserving the optical contrast required for Moiré amplification.

4.1.1 Mold and Material Preparation. **Preparation of casting process.** We prepare a custom optical grating mold that transfers the grating geometry to the soft layers, with a uniform pitch of $p = 330\ \mu$m. Because the fidelity of this transfer directly affects the quality of the resulting Moiré fringes, the mold surface must be clean and uniform to support reliable sub-millimeter-scale fabrication. We thus sequentially rinse the mold with acetone, isopropyl alcohol, and deionized water, followed by nitrogen drying, then treated with plasma for approximately 2 min to improve pattern transfer reliability. **Preparation of liquid silicone rubber (LSR).** We next prepare the soft LSR used throughout MoiréSkin. The elastomer forms the mechanical interface, grating carrier, and optical path. We select it for compliance and deformability under tiny contact forces, stretchability under pneumatic strain, optical transparency, grating-transfer fidelity, and

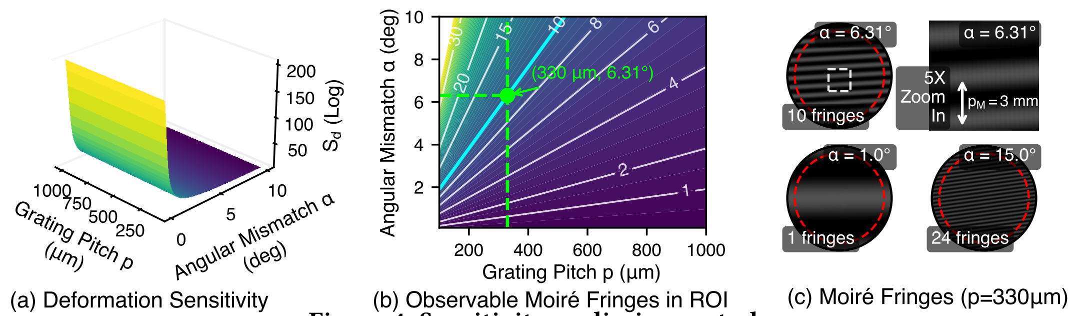

Figure 4: Sensitivity preliminary study.

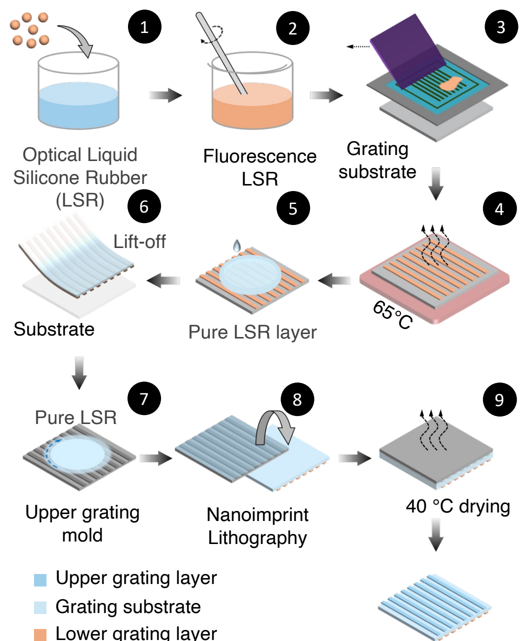

Figure 5: Dual-side grating fabrication pipeline.

soft-lithography compatibility. Our formulation mixes Silicones XP-565 base, Sylgard 527 softening agent, and curing agent at a 9:1:1 mass ratio [25, 37, 55, 62].

4.1.2 Dual-side Grating Fabrication Pipeline. To produce Moiré fringes that amplify contact-induced deformation, `MoiréSkin` patterns high-frequency gratings on dual-sided layers. Fig. 5 illustrates the detailed fabrication pipeline of the dual-sided grating structure in Fig. 1, including the grating substrate and the upper and lower grating layers.

Grating Substrate Fabrication. We first fabricate the grating substrate. The LSR is manually stirred for 5 min until homogeneous, followed by 2 min of ultrasonication. To guarantee the fabrication quality, the mixture is then placed in a vacuum desiccator, and degassed until visible bubbles are removed. Then, the degassed LSR was cast onto a flat mold to form a uniform layer and cured in room temperature to obtain the grating substrate. This casting process enables the fabrication of grating substrates with thicknesses exceeding 100μm, allowing substrate thickness control by adjusting the casting volume.

Single-Layer Grating Casting. Next, we fabricate a single-layer grating film by casting the degassed grating substrate onto the patterned mold, following a spin coating-imprinting-curing strategy processes (Fig. 5 Step 7-9). At Step 7, a volume of 2–3 mL LSR is dispensed onto cover the grating patterned mold, and spin-coated to form a homogeneous thin film. The coated layer is then subjected to a preliminary low-temperature curing step to bring the material into a semi-cured state, to stabilize the geometry of the replicated grating. At Step 8, the mold with the thin film is flipped so that the patterned surface faces downward and is brought into conformal contact with the previously prepared structure. A microscale imprinting process is then applied by exerting gentle pressure, allowing the semi-cured LSR film to conformally fill the microcavities of the mold and replicate the patterned features. A vacuum treatment is applied to remove trapped bubbles and improve both pattern fidelity and optical quality, then thermally cured in a vacuum oven At Step 9, The film is then soft-baked at 65 °C under vacuum for 10–15 min to stabilize the surface relief, after which the patterned elastomer film is gently peeled off. The resulting film carries a faithful negative replica of the grating pattern while maintaining the flexibility and optical transparency required for Moiré-based sensing.

Double-Layer Grating Casting. A straightforward way to form the second grating would be to repeat the same process on the opposite side with a small angular offset α. Instead, we adopt a modified fabrication design for the bottom grating layer so that it can additionally provide fluorescent reflection, which reduces interference from ambient-light interference. We describe this bottom grating layer fabrication strategy (Step 2-6) in the next section.

Assembly. Finally, the dual-grating films are inspected under an optical microscope to verify structural integrity. Samples with delamination, incomplete pattern transfer, or residual bubbles are discarded. Completed films are stored in a dust-free Petri dish at room temperature and remain stable for at least three months without noticeable degradation. Across more than 20,000 tactile sensing and actuation samples collected using the same `MoiréSkin` cell throughout our experiments (Sec. 7), the membrane and gratings remained visually intact, providing preliminary evidence of durability under the tested conditions. We leave a comprehensive evaluation of long-term mechanical durability for future work. To complete the MoiréSkin cell, we integrate the dual-sided grating structure with a protective top cover and a sealed pneumatic chamber underneath, as shown in Fig. 1. The top cover is formed using a commercial spray-on TPU elastomer, which provides a thin protective coating while preserving the optical readout and mechanical compliance of the sensor.

4.2 Optical Design

A common challenge for VTS is that ambient light outside the cover layer can corrupt the captured tactile sensing image. To mitigate this challenge, many structured-light-based VTS designs use opaque cover layers [59]. However, doing so blocks the camera's ability to see through the sensing surface, preventing pre-contact perception and limiting shape observation during the transition from non-contact to contact [18]. MoiréSkin addresses ambient-light interference while preserving "see-through" capability through a high-contrast *fluorescent* Moiré fringes generation and optical co-design.

4.2.1 Fluorescent Fabrication and Casting. As shown in Fig. 5 Step 2, long-afterglow luminescent materials are incorporated into LSR to form the lower grating layer. The luminescent powder is firstly dried at 80 °C for 12 h to remove adsorbed moisture. After cooling to room temperature in a desiccator, the powder was mixed with LSR at a mass ratio of 1:20. To ensure uniform dispersion and stable material properties, the powder is gradually added under continuous stirring at low temperature. Step 3-6 involves patterned printing, thermal curing, LSR casting, and transfer printing. First, the fluorescent LSR is transferred onto the substrate by patterned printing, forming the desired grating pattern. The printed structures are then thermally cured at 65 ∘C to stabilize the printed features and improve their mechanical integrity. Then LSR is cast onto the patterned surface, allowing the LSR to infiltrate and conformally cover the micro-structured features. After curing, the fluorescent patterned structures will be transferred onto the elastomeric film, then the entire layer is peeled off from the substrate for the following process. This process enables the bottom grating layer to emit fluorescent signals that enhance Moiré visibility while reducing sensitivity to external illumination.

4.2.2 Cover Layer Design. The cover layer is co-designed to further enhance the contrast of the fluorescent Moiré fringes. Specifically, it is engineered with spectral filtering properties that attenuate long-wavelength red light (approximately 620–750 nm) from the environment while allowing the desired fluorescent emission to pass. By suppressing unwanted background components, this design improves the visibility of Moiré fringes without sacrificing transparency.

4.2.3 Illumination Co-Design. The illumination system in MoiréSkin serves two purposes: (1) exciting the embedded fluorescent gratings, and (2) providing a stable optical baseline for camera exposure. To this end, MoiréSkin uses a programmable WS2812 LED ring tuned to a blue excitation band (455–485 nm), which matches the excitation spectrum of the fluorescent material embedded in the lower grating layer. As shown in Fig. 1, the LED ring is oriented downward toward a diffuse reflection plate at the base of the module. This configuration converts discrete LED sources into a more uniform internal light field, reducing local hotspots on reflective elastomer surfaces. The LED ring is also mounted below the camera optics so that the surrounding structure acts as a baffle, blocking direct light paths into the lens and suppressing glare and lens flare. Together, these optical components enable MoiréSkin to retain "see-through" functionality while remaining robust to ambient-light interference.

4.3 Actuation Hardware Design

To tightly couple tactile sensing with active compliance regulation, MoiréSkin integrates a sealed pneumatic chamber beneath the dual-layer grating structure, as shown in Fig. 1. The chamber is connected to an air pump, and both the pump and the camera interface with the same backend computing unit, enabling closed-loop sensing and actuation. Through this shared backend, MoiréSkin can regulate contact conditions based on real-time tactile feedback. For example, when the sensing module detects incipient slip from shear-induced Moiré changes, the controller can increase chamber pressure to stiffen the contact interface and raise the available friction force. Conversely, the pressure can be reduced to soften the interface for gentler contact, such as when interacting with fragile objects.

Such sensing-actuation integration is difficult to achieve with conventional structured-light-based rigid VTS [59]. First, rigid VTS typically relies on a fixed camera-surface geometry and stable illumination path to preserve photometric consistency for tactile reconstruction. Introducing pneumatic deformation would perturb this geometry and violate the optical conditions required for accurate structured-light sensing. Second, actuator-induced surface displacement can move the sensing layer away from the nominal camera focal plane, leading to blur and degraded fine-grained texture or color cues. In contrast, MoiréSkin encodes deformation into low-frequency, large-scale Moiré fringes that remain visually salient across different actuation states, including under moderate defocus, as shown in Fig. 7 (a), (b). This property makes MoiréSkin naturally compatible with integrated pneumatic actuation, enabling a sensor-actuator VTS that supports both tactile perception and active compliance control.

5 Software Design

Next, we introduce our software design, which is on a control and computation unit connecting to both the camera and actuator. We first introduce the actuation and sensing module separately, and then we introduce the overall process.

5.1 Sensing Layer State Estimation

A key requirement for closed-loop sensing and actuation is *accurate and real-time estimation of the sensing layer actuation state, i.e., whether the layer is flat, inflated, or deflated.* This state estimation serves as the feedback for the controller to regulate air pumping and drive the sensing layer toward a desired actuation configuration. Physically, the actuation state is determined by the pressure difference $\Delta B = B_{\text{in}} - B_{\text{air}}$, where B_{in} is the chamber pressure and B_{air} is the ambient atmospheric pressure. We propose a calibration method that infers the pneumatic state ΔB from the Moiré fringe geometry θ_M without requiring additional pressure sensors.

5.1.1 Fringe-to-Pressure Modeling. Motivated by our preliminary study in Sec. 3.2, we use Moiré fringe rotation, θ_M, to infer ΔB directly. Assuming no external contact, we define the actuation "zero-state" when $\Delta B = 0$, i.e., $B_{\text{in}} = B_{\text{air}}$. At this state, the upper and lower grating layers remain flat and have identical pitch$p^{(0)}_{\text{upper}} = p^{(0)}_{\text{lower}}$, with corresponding fringe orientation $\theta^{(0)}_M$ (Fig. 3 (a)). Our goal is to estimate $\widehat{\Delta B}$ from the observed orientation change $\Delta\theta_M = \theta_M - \theta^{(0)}_M$.

However, the mapping from $\Delta\theta_M$ to ΔB is monotonic but inherently nonlinear based on our preliminary studies. This nonlinearity arises from two factors. First, as shown in Fig. 3 (b) and (c), the upper and lower gratings are separated by a finite thickness, so inflation or deflation deforms them differently, yielding $p^{(t)}_{\text{upper}} \neq p^{(t)}_{\text{lower}}$, which leads to a time-varying pitch ratio $R_p = p_{\text{upper}}/p_{\text{lower}} \neq 1$. Since the Moiré fringe orientation is determined by the vector difference between the two grating frequencies, Eq. 2 implies a nonlinear optical response as the geometry changes. Second, the elastomeric membrane exhibits strain hardening, so the marginal deformation induced by additional pressure decreases as ΔB increases [12, 22]. As a result, the fringe-to-pressure relationship depends on the material properties of the sensing layer, including thickness, compliance, and stretchability, and must be calibrated empirically.

Material Calibration. To calibrate this mapping, we temporarily instrument MoiréSkin with two commercial MEMS barometers (Adafruit MPRLS Ported Pressure Sensor): one connected to the sealed chamber and the other exposed to the ambient air. These barometers are used only during calibration and are removed during normal operation. We randomly sample target pressure differences ΔB over the chamber's operating range (−11 to +12 hPa), and regulate B_{in} using a 5 V DC mini air pump together with vacuum control until

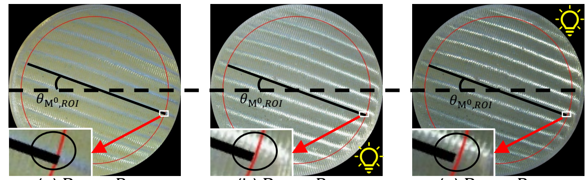

(a) $B_{in} = B_{air}$ Normal lighting θ_M Stable (b) $B_{in} = B_{air}$ Lower-right lighting θ_M Changed (c) $B_{in} = B_{air}$ Upper-right lighting θ_M Changed

Figure 6: Fringe stability under different lighting

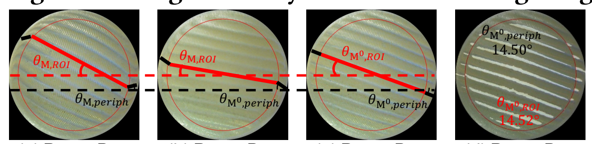

(a) $B_{in} < B_{air}$ Intrafocal (b) $B_{in} > B_{air}$ Extrafocal (c) $B_{in} = B_{air}$ In-focus (d) $B_{in} = B_{air}$ Processed

Figure 7: "Zero-State" detection

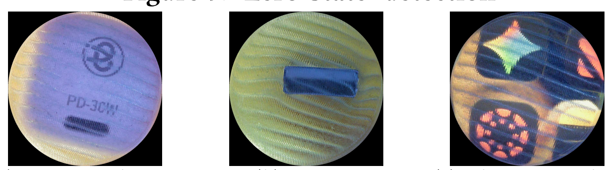

(a) Power Adapter **(b) USB-A** **(c) Phone Display**

Figure 8: "See-through" functionality.

the desired ΔB is reached. Once the pressure stabilizes, the camera captures an image to record the corresponding Moiré fringe orientation θ_M. We repeat this process 100 times to collect samples across the full actuation range. We then identify the reference orientation $\theta^{(0)}_M$ at the zero-state ($\Delta B = 0$), compute $\Delta\theta_M = \theta_M - \theta^{(0)}_M$, and fit the mapping from $\Delta\theta_M$ to ΔB. For our current material configuration, this relationship is well approximated by a third-order polynomial:

$$\widehat{\Delta B(\Delta\theta)} = -0.0108\Delta\theta^3 + -0.1748\Delta\theta \tag{3}$$

With this calibrated regression function, MoiréSkin can directly infer $\widehat{\Delta B}$ from the observed Moiré fringe rotation, eliminating the need for embedded barometers during normal operation while still achieving precise actuation control.

5.1.2 "Zero-State" Calibration. Although the fringe-to-pressure model is defined relative to $\theta^{(0)}_M$, this reference cannot be fixed once and reused indefinitely, because the observed Moiré pattern drifts across environments. Two external factors are particularly important. First, the ambient atmospheric pressure B_{air} can vary substantially with weather, temperature, humidity, and altitude. Even at a fixed location, B_{air} may fluctuate by approximately 10–20 hPa over a day [32], which is comparable to the full actuation range of MoiréSkin (−11 to +12 hPa). As a result, the sensing layer may appear inflated or deflated even without active pumping. Second, ambient illumination also affects the measured fringe orientation. Unlike conventional rigid VTS designs that use opaque structures to isolate the sensing region from ambient light, MoiréSkin adopts a transparent cover layer to

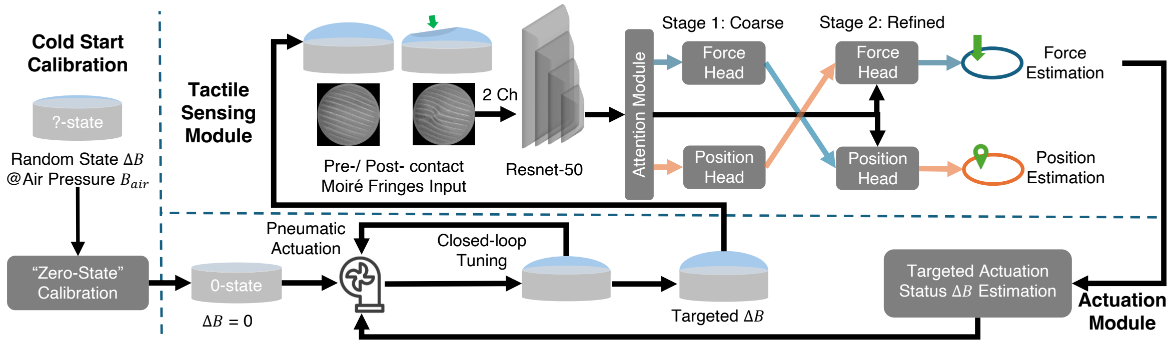

Figure 9: `MoiréSkin` software pipeline.

support the "see-through" functionality described in Sec. 4.2. This transparency makes the optical signal sensitive to external disturbances such as moving shadows and changes in light-source position or intensity. Even under constant pressure, slight illumination changes can shift the apparent fringe orientation by several degrees, as shown in Fig. 6(b) and (c). Therefore, `MoiréSkin` cannot rely on a fixed preset value of $\theta_M^{(0)}$ to represent the "zero-state".

To address this issue, *we design a dedicated calibration module to identify the "zero-state" ($\Delta B = 0$) by exploiting a unique physical property: at the "zero-state", the sensing layer is perfectly flat.* As shown in Fig. 7, the observed Moiré fringes can be divided into two regions: a central ROI and a peripheral region near the clamped boundary. Because the outer annulus of the membrane is mechanically constrained by the rigid frame [56], the peripheral region experiences a different strain profile from the central ROI during inflation or deflation. Consequently, the fringe orientations in these two regions evolve differently. As shown in Fig. 7 (c), only when the chamber film is perfectly flat, the fringes in the periphery region (black line) $\theta_{M,\text{Periph}}$ and central regions (red line) $\theta_{M,\text{ROI}}$ align into a single, continuous straight-line distribution ($\theta_{M,\text{Periph}} = \theta_{M,\text{ROI}}$). We use this observation to detect the "zero-state".

To estimate $\theta_{M,\text{Periph}}$ and $\theta_{M,\text{ROI}}$, we first convert the captured image to LAB color space and extract the L channel. We then apply Contrast Limited Adaptive Histogram Equalization (CLAHE) [39] to the L-channel to enhance local contrast, followed by Gaussian blurring to suppress high-frequency texture while preserving the low-frequency Moiré fringes. Next, an adaptive Gaussian-weighted threshold is used to isolate locally dark fringe regions. A contour-based principal component analysis is used to estimate the dominant fringe orientation in each region to determine $\theta_{M,\text{ROI}}$ and $\theta_{M,\text{Periph}}$. Because the peripheral region is more susceptible to edge effects and non-uniform illumination near the rigid frame, as shown in Fig. 7(d), we further apply area-weighted averaging to obtain a robust estimation of the dominant peripheral orientation $\theta_{M,\text{Periph}}$.

During calibration, `MoiréSkin` continuously monitors the orientation difference $\Delta\theta_{M,\text{ROI-Periph}} = \theta_{M,\text{ROI}} - \theta_{M,\text{Periph}}$. When the sensing layer is concave, $\theta_{M,\text{ROI}} < \theta_{M,\text{Periph}}$, indicating that $B_{\text{in}} < B_{\text{air}}$ and the chamber should be inflated, as shown in Fig. 7 (a). Conversely, when the sensing layer is convex, $\theta_{M,\text{ROI}} > \theta_{M,\text{Periph}}$, indicating $B_{\text{in}} > B_{\text{air}}$ and that the chamber should be deflated. Based on this principle, we implement a PID controller that uses $\Delta\theta_{\text{M, ROI-Periph}} = \theta_{M,\text{ROI}} - \theta_{M,\text{Periph}}$ as the control error and modulates the air pump duty cycle accordingly. This feedback loop autonomously drives the sensing layer to the "zero-state" within 2 s.

After reaching the zero-state, the system records its fringe orientation as $\theta_M^{(0)}$ and estimates subsequent $\widehat{\Delta B}$ using Eq. 3. We perform this cold-start procedure to compensate for environmental drift caused by atmospheric-pressure and ambient-light changes.

5.2 Tactile Sensing Module Design

The tactile sensing module uses contact-induced Moiré fringe distortions to estimate contact position and normal force in real time. It is challenging because the optical response depends on the unknown contact geometry and force as well as the sensing layer's actuation state. Unlike calibration (Sec. 5.1), real-world contacts produce complex local distortions that vary across the nonplanar surface. Pneumatic actuation further changes the surface geometry, causing identical contacts to produce different, shifted, or defocused patterns across locations and actuation states (Fig. 7). We therefore develop a neural network for simultaneous contact position and normal force estimation.

5.2.1 Model Input and Output. **Pre-/Post- Contact Inputs.** The model takes pre- and post-contact images as input. Each image is applied by a circular mask to isolate the sensing ROI and converted from BGR to the CIELAB L-channel, separating Moiré fringe distortions from color variations. The mean ROI brightness is then normalized to 128, reducing sensitivity to LED and ambient illumination changes (Fig. 12).
Model Output. Following prior work [48], the model predicts two quantities: (1) contact position, defined as the center of the contact region; and (2) normal contact force.

5.2.2 Backbone Neural Network. MoiréSkin adopts a modified ResNet-50 architecture [17] with a custom weight initialization strategy to adapt pre-trained features to single-channel geometric fringe distortions. To bridge the domain gap between natural images and tactile patterns, we initialize the first convolutional layer by summing pre-trained RGB kernels and replicating them across all input channels, preserving edge and orientation detectors while removing color-specific bias. A Convolutional Block Attention Module (CBAM) [52] is integrated before global pooling to emphasize localized fringe distortions and suppress background interference from the transparent sensing layer.

5.2.3 Two-Stage Tactile Sensing Refinement. As discussed above, the inflatable geometry and varying actuation states of the sensing layer lead to position-dependent force responses and can also bias contact position. These nonlinear effects are difficult to capture with a single-stage predictor. We therefore design a two-stage, physics-informed refinement framework, as shown in Fig. 9

Stage 1: Coarse-Grained Estimation. The latent features extracted by the ResNet-50 backbone are first processed by two regression heads: a coarse position head and a coarse force head. These heads generate initial estimation of the contact position and normal force magnitude.

Stage 2: Fine-Grained Refinement. We then use the Stage 1 outputs as physics-informed conditioning variables for refinement. For force refinement, the coarse position estimate is concatenated with the backbone latent features and fed into a refined force head. This allows the model to adapt force estimation to the location-dependent stiffness of the membrane. For position refinement, the coarse force estimate is concatenated with the same latent features and passed to a refined position head. This stage compensates for position shifts caused by actuation-induced geometric deformation. Together, these two refinement paths improve both force and position accuracy by modeling their mutual dependence.

5.2.4 Training Data Augmentation. During model training, we apply a physics-informed data augmentation pipeline to improve the model robustness.

Actuation-State Augmentation. To simulate the geometric variability introduced by pneumatic actuation, we apply random perspective transformation with distortion scale 0.15 and probability 0.5. This augmentation approximates the stretching and fringe-orientation changes that arise across different actuation states, helping the network learn features that are less sensitive to global shape variation.

Ambient-Light Augmentation. To improve robustness to changing optical conditions, we apply color jitter with brightness 0.4 and contrast 0.4 at probability 0.8. This models illumination and contrast variation caused by ambient light changes, shifting light paths, and partial light occlusion.

Table 2: MoiréSkin per-cell material cost.

Part	Cost (USD)	Description
Part 1: Sensing Layer		
Cover layer	0.05	Commercial TPU
Upper grating layer	0.25	Pure LSR
Grating substrate	0.1	Pure LSR
Lower grating layer	0.65	Fluoresent LSR
Part 2: Actuation Layer		
Chamber	0.4	Commercial elastomer
bottom	0.1	Transparent rigid material
PETG housing	0.2	PETG HF, Bambu Lab

Low-Light Robustness: We further inject additive Gaussian noise with standard deviation 0.01 at probability 0.5 to simulate sensor noise under low-light conditions. This improves robustness when fringe observations are degraded by insufficient illumination or camera noise.

6 SETUP AND IMPLEMENTATION

6.1 Hardware Setup

Hardware Setup. Fig. 1 shows the hardware setup; the sensing and actuation layers follow Sec. 4.1. A 5 V pump with vacuum control is driven by an Arduino-controlled L298N motor driver for inflation/ deflation switching, and PWM-based speed regulation, with an external supply providing its high current. An Arducam IMX291 RGB camera (2 MP, 1920×1080) with a 160° lens monitors the sensing layer. A WS2812B 8-bit RGB ring (B: 455–485 nm) excites the lower-layer fluorescent dopants, while a frosted base film diffuses illumination and reduces hotspots. A PETG housing rigidly aligns the chamber, camera, and illumination. Arduino microcontroller and the camera are connected to a MacBook Pro for integrated sensing and actuation.

Fabrication Cost and Replicability. Tab. 2 summarizes the per-cell material costs of the sensing and actuation layers. The total material cost of each cell is less than US$2. This estimate excludes reusable fabrication equipment, including the patterned molds, plasma cleaner, ultrasonicator, spin coater, vacuum oven, and microscope. The fabrication process does not require a specialized cleanroom. Including the RGB camera, LED ring, pump, and driver, the total system cost is approximately US$50. Repeatability is supported by fixed process parameters and reusable molds. Each dual-grating film is microscopically inspected before integration and accepted only if both patterns are fully transferred without delamination or residual bubbles.

6.2 Software and Inference Pipeline

The software stack is implemented in PyTorch and OpenCV, and all training and inference are performed on an NVIDIA RTX 6000 GPU. We train two models for evaluation: (1) core tactile sensing and (2) downstream shear-force application.

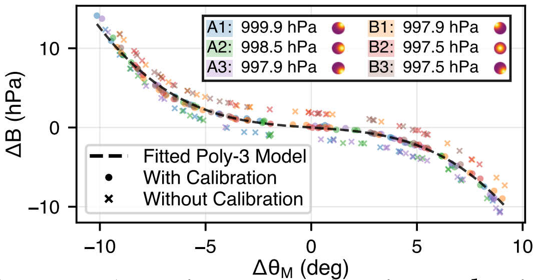

Figure 10: Actuation state estimation evaluation.

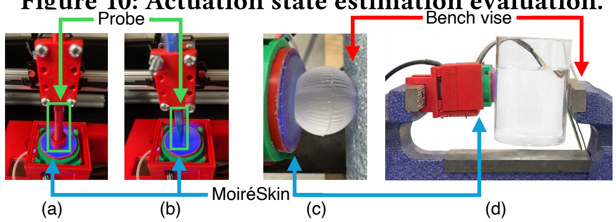

Figure 11: Experimental setups used in § 7.2 & § 7.3

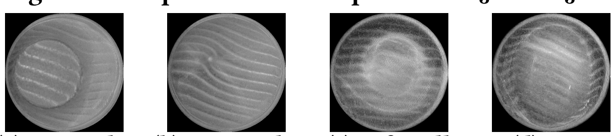

(a) Big Probe **(b) Tiny Probe** **(c) Soft Ball** **(d) Cup**

Figure 12: Processed Internal Sensor Views.

Tactile Sensing Model. For contact position and normal force estimation, we optimize the model using AdamW with a weight decay of 1×10^{-4} and an initial learning rate of 1×10^{-4}. We use a `ReduceLROnPlateau` scheduler (factor 0.5, patience 7), a batch size of 32, and early stopping with a patience of 50 over a maximum of 500 epochs. To improve robustness in the highly non-linear sensing space, we use Smooth L1 loss, automatic mixed precision (AMP), and gradient clipping with a maximum norm of 1.0. Its evaluation results are reported in Sec. 7.2.2. We also directly reuse this model, without additional fine-tuning, for ultra-soft material stiffness classification in Sec. 7.3.1.

Shear-Force Fine-Tuning. The shear-force refinement model (Sec. 7.3.2) is initialized from the pre-trained tactile sensing model and fine-tuned using AdamW with an initial learning rate of 5×10^{-5}. We use a batch size of 64, gradient clipping with a maximum norm of 5.0, and early stopping with a patience of 20 over a maximum of 200 epochs.

7 Evaluation

We evaluate `MoiréSkin` on (1) actuation-state estimation (Sec. 5.1), (2) contact position and normal force (Sec. 5.2), (3) integrated sensing–actuation through stiffness classification and shear-force-based mass estimation (Sec. 7.3), and (4) system latency and model size.

7.1 Actuation State Estimation Accuracy

We evaluate material and "zero-state" calibration accuracy (Sec. 5.1) in estimating the actuation state ΔB.

7.1.1 Evaluation Setup & Data Collection. We instrument `MoiréSkin` with two MEMS barometers (Sec. 5.1.1) to obtain ground-truth ΔB. A recursive EMA filter ($\alpha = 0.1$) suppresses electronic jitter and high-frequency noise in the raw pressure stream; a PID-controlled air pump regulates target states. We log each point after the rolling standard deviation remains below 0.2 hPa, then wait 2 s for pneumatic stabilization. We average 10 consecutive readings to reduce barometer uncertainty. We establish the baseline pressure-fringe relationship between ΔB and $\Delta\theta$ (Eq. 3) from 100 repetitions, then collect 120 test points in six sessions spanning two days at 2-hour intervals. Joint variations in ambient atmospheric pressure B_{air} and illumination evaluate long-term temporal and environmental stability.

7.1.2 Experimental Results. Fig. 10 shows that without "zero-state" calibration ("x" markers), R^2 falls to 0.772 under varying atmospheric conditions. The calibration restores R^2 above 0.994 and reduces RMSE by up to 85%. It also reduces ΔB MAE from 1.771 hPa to 0.151 hPa in high-drift Sessions B. Thus, *"zero-state" calibration guarantees stable actuation-state sensing without external pressure sensors, enabling closed-loop sensing and actuation.*

7.2 Tactile Sensing Accuracy

We evaluate contact-position and normal-force sensing at two scales, followed by an ablation of the design in Sec. 5.2.

7.2.1 Evaluation Setup & Data Collection. A `MoiréSkin` cell on a digital lab scale (1,000 g capacity, 0.01 g resolution) and an iDraw 2.0 XY plotter provide contact-position and normal-force ground truth for repeatable indentations. Pre- and post-contact image pairs provide model inputs for two configurations. *General-force* uses a 15 mm flat-ended probe (Figs. 11(a) and 12(a)) under randomized actuation states for normal forces up to 4 N ($\Delta B \in [-11, 12]$ hPa); *tiny-force* uses a 0.72 mm pointed-tip probe (Figs. 11(b) and 12(b)) at the actuation "zero-state" for maximum sensitivity below 400 mN. Each setup has 4,000 training/validation samples at 256 and 400 Poisson-disk positions, respectively, plus a disjoint 1,000-sample test set for unseen-position generalization. Separate models use 5-fold cross-validation on the training/validation data; we report normal-force and contact-position mean absolute error (MAE), root mean square error (RMSE), and coefficient of determination (R^2) on the test sets.

7.2.2 Experimental Results. Fig. 13 shows highly linear force response across 4 N: 94.9% of force errors are below 0.1 N despite optical shifts across randomized actuation states. Position MAE rises slightly above 3 N because of focal shifts, but overall accuracy remains high. In the tiny-force setting (Fig. 15), errors remain stable across loads, with force-error bars typically below 15 mN even near 400 mN. Position error remains submillimeter across the sensing area, rising slightly

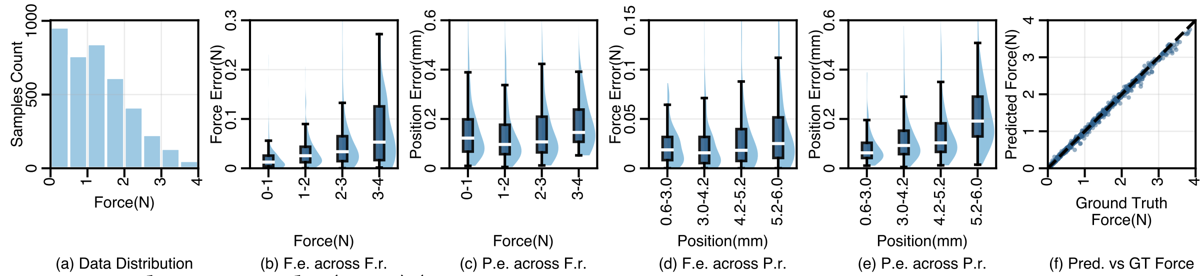

Figure 13: General-Force Results (0-4 N) (F.e.: Force error; F.r.: Force range; P.e.: Position error; P.r.: Position range. Position range uses an equal-area binning strategy, where each bin represents an identical spatial extent.)

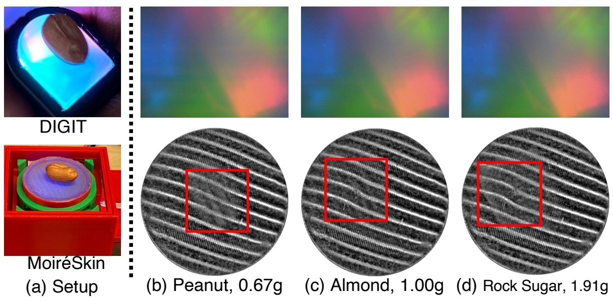

Figure 14: Qualitative Comparison: (a) Setup. (b–d) Sensor imaging with DIGIT and MoiréSkin.

near the frame (10.4–12.0 mm radius). MoiréSkin *reaches 0.094 mm position and 3.78 mN normal-force MAE, supporting gram-scale perception.*

7.2.3 Ablation Study. To quantify the impact of our key design choices, we compare the proposed model against two baselines: (1) *"1-Ch, W/o Aug."*, which uses only a single-channel post-contact image without data augmentation, and (2) *"1-Ch, W/ Aug."*, which adds data augmentation to the single-channel input. As shown in Tab. 3, our full design consistently improves sensing accuracy and reduces measurement variance in both the general-force and tiny-force configurations. In the general-force configuration, it reduces position and force MAE by 70.0% and 28.7%, respectively, compared to *"1-Ch, W/o Aug."*. In the tiny-force configuration, it reduces position and force MAE by 19.5% and 44.7%, respectively, relative to *"1-Ch, W/o Aug."*. The gain over both one-channel baselines suggests that the pre-contact image supplies a physical reference that complements augmentation by separating contact deformation from actuation-dependent background shifts. *Overall, these results show that our model design improves generalization across varying actuation states and lighting conditions.*

7.2.4 Comparison with DIGIT [24]. We qualitatively compare MoiréSkin and GelSight DIGIT [24] using gram-scale objects in Fig. 14. DIGIT shows no distinguishable change, whereas MoiréSkin produces clear localized Moiré fringe distortions. Thus, measured tiny-force sensitivity translates into visible gram-scale contact cues.

Table 3: Ablation Study: Quantifying the performance gap in contact position and force estimation by 5-fold cross validation. Red arrows (↓) indicate the relative performance degradation compared to our framework.

Configuration	Pos. MAE [mm]	Degradation	Force MAE [N]	Degradation
General-force Configuration ($0-4$ *N*)				
1-Ch, W/o Aug.	0.1749 ± 0.0236	↓ 70.0%	0.0318 ± 0.0018	↓ 28.7%
1-Ch, W/ Aug.	0.2149 ± 0.0204	↓ 108.8%	0.0319 ± 0.0023	↓ 29.1%
2-Ch, W/ Aug.	**0.1029** ± **0.0032**	*Proposed*	**0.0247** ± **0.0005**	*Proposed*
Tiny-force Configuration ($\Delta B = 0$, $0-0.4$ *N*)				
1-Ch, W/o Aug.	0.1188 ± 0.0102	↓ 19.5%	0.0068 ± 0.0013	↓ 44.7%
1-Ch, W/ Aug.	0.1148 ± 0.0056	↓ 15.5%	0.0052 ± 0.0004	↓ 10.6%
2-Ch, W/ Aug.	**0.0994** ± **0.0075**	*Proposed*	**0.0047** ± **0.0004**	*Proposed*

7.3 Integrated Sensing & Actuation Studies

Next, we conduct two case studies using a fixed bench vise to emulate consistent robotic grasp control and demonstrate integrated sensing–actuation: (1) material-stiffness classification and (2) dynamic mass estimation from shear force. For both cases, MoiréSkin is mounted on one jaw of a bench vise, with the object grasped at a fixed vise position to isolate tactile changes induced by actuation and other object mechanics, as shown in Fig. 11 (c) and (d), with the internal sensor views shown in Fig. 12 (c) and (d) respectively. We use a higher actuation state, $\Delta B = 41.25$ hPa, to create the stiffer contact interface needed for stable grasping. This setting extends beyond the standard training range (-11 to $+12$ hPa); fine-tuning adapts the pretrained tactile model to this operating point.

7.3.1 Case Study 1: Soft Material Stiffness Classification. This study evaluates whether MoiréSkin can discriminate among seven ultra-soft materials that differ only in stiffness, with Young's moduli ranging from 8.73 kPa to 21.72 kPa, roughly covering tofu-like to skin-like softness, as shown in (Fig. 16). To mimic humans' active probing, each ball is grasped once at the fixed vise setting and imaged first at $\Delta B = 0$, then again after inflating the chamber to $\Delta B = +41.25$ hPa. This produces a two-state tactile image pair for each trial, and we collect 60 sessions in total.

For classification, we directly use the 2048-dimensional latent features extracted by our pre-trained general-force tactile model, and applied Linear Discriminant Analysis (LDA) to the features. Fig. 17 shows that MoiréSkin achieves 92.7%

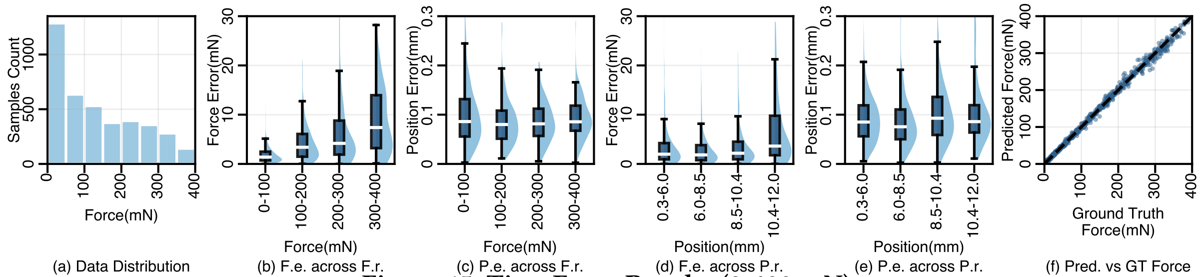

Figure 15: Tiny-Force Results (0-400 mN)

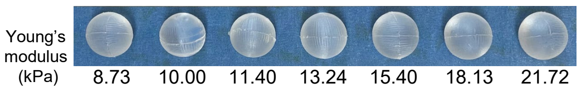

Figure 16: Seven Balls Soft Material Test Set

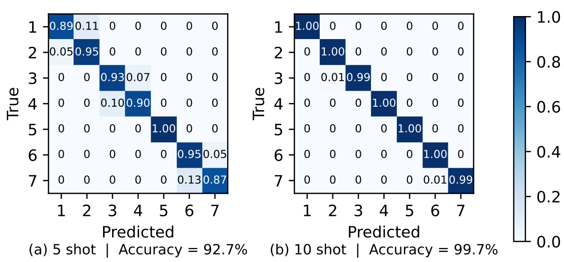

Figure 17: Stiffness Classification Confusion Matrices

and 99.7% accuracy across seven stiffness classes in the 5-shot and 10-shot settings, respectively, demonstrating that the pre-trained backbone can be directly applied to this task.

7.3.2 *Case Study 2: Dynamic Mass Estimation via Shear Force.*

This case study evaluates MoiréSkin's ability to sense shear force during manipulation by tracking the changing mass of water in a grasped cup, where added liquid produces a corresponding increase in gravity-driven shear load. The same cup is grasped at a fixed vise position and constant actuation state ΔB while water is gradually added from 0 to approximately 300 mL, and Moiré images are recorded at 10 Hz (Fig. 11(d)). We collect six sessions, using one unseen session for testing and the remaining five for cross-validation.

To adapt the model, we fine-tune the pre-trained ResNet-50 backbone from the general-force setting to capture subtle fringe bending caused by the liquid's downward pull, while training a new shear-force regression head from scratch. Fig. 18 shows stable mass-estimation errors across the full range, with an MAE of 4.1462 ± 0.5994 g and an RMSE of 4.7517 ± 0.5353 g. The error does not increase substantially even at larger masses (250–350 g), where shear stress is higher. These results show that Moiré fringes encode sufficient mechanical information for shear force estimation, allowing the cup to serve as a dynamic scale during grasping. In closed-loop operation, MoiréSkin can detect mass changes in real time and prompt the pneumatic controller to increase chamber pressure, stiffening the interface and increasing friction to prevent incipient slip.

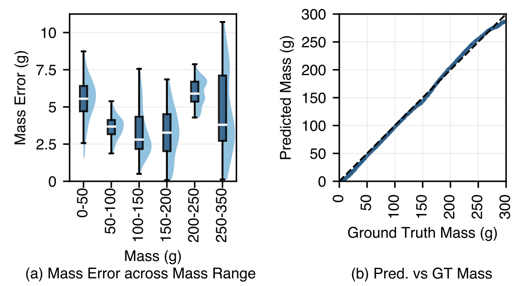

Figure 18: Dynamic Mass Estimation

7.4 System Evaluation

Actuation and Calibration Latency. Both pneumatic actuation and zero-state calibration converge within 2 s across all tested initial pressure states. The PID control loop introduces negligible computational overhead, with latency dominated by the pneumatic response.

Sensing Latency. The tactile sensing model achieves per-sample inference latencies of 0.756 ± 0.085 ms on an NVIDIA RTX 6000 GPU (≈ 1.3 kHz) and 22.06 ± 4.08 ms on an Apple M1 GPU (≈ 42 Hz), respectively.

Compute Footprint. The model contains 28.23M parameters and requires 107.69 MB of memory at FP32 precision.

Although the current model prioritizes accuracy over efficiency, these results support real-time tactile sensing. Future work will reduce end-to-end latency and computational cost for deployment on a robotic arm.

8 CONCLUSION

We presented MoiréSkin, an inflatable visuo-tactile skin that embeds dual-sided, sub-millimeter gratings in a soft membrane to combine ultra-sensitive tactile sensing, active compliance regulation, and "see-through" perception. It offers a new direction for developing sensitive sensor–actuator systems for fine-grained robotic manipulation and dynamic tactile interaction. Our current evaluation uses a fixed platform to characterize sensing and actuation accuracy. Future work will integrate MoiréSkin into dexterous robotic fingers with onboard computation and explore real-world applications, such as manipulating fragile and deformable objects with diverse geometries.

References

[1] Isaac Amidror. 2009. *The Theory of the Moiré Phenomenon: Volume I: Periodic Layers*. Vol. 38. Springer Science & Business Media, London, UK. doi:10.1007/978-1-84882-181-1

[2] A.-S. Augurelle, A. M. Smith, T. Lejeune, and J.-L. Thonnard. 2003. Importance of cutaneous feedback in maintaining a secure grip during manipulation of hand-held objects. *Journal of Neurophysiology* (2003). doi:10.1152/jn.00249.2002

[3] S. Begej. 1988. Planar and finger-shaped optical tactile sensors for robotic applications. *IEEE Journal on Robotics and Automation* (1988).

[4] D. Campos Zamora et al. 2024. MoiréWidgets: High-Precision, Passive Tangible Interfaces via Moiré Effect. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.

[5] Yushi Cheng, Xiaoyu Ji, Lixu Wang, Qi Pang, Yi-Chao Chen, and Wenyuan Xu. 2021. mID: Tracing Screen Photos via Moiré Patterns. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 2969–2986. https://www.usenix.org/conference/usenixsecurity21/presentation/cheng-yushi

[6] Qingzheng Cong, Steven Oh, Wen Fan, Shan Luo, Kaspar Althoefer, and Dandan Zhang. 2025. TacEva: A Performance Evaluation Framework For Vision-Based Tactile Sensors. *arXiv preprint arXiv:2509.19037* (2025). doi:10.48550/arXiv.2509.19037

[7] G. Corniani and H. P. Saal. 2020. Tactile innervation densities across the whole body. *Journal of Neurophysiology* 124, 4 (2020), 1229–1240. doi:10.1152/jn.00313.2020

[8] Ravinder S Dahiya, Giorgio Metta, Maurizio Vallejo, and Sandro Adami. 2010. Tactile Sensing—From Humans to Humanoids. *IEEE Transactions on Robotics* 26, 1 (2010), 1–20.

[9] Ravinder S Dahiya, Philipp Mittendorfer, Maurizio Valle, Gordon Cheng, and Vladimir J Lumelsky. 2013. Directions toward effective utilization of tactile skin: A review. *IEEE Sensors Journal* 13, 11 (2013), 4121–4138.

[10] Won Kyung Do, Bianca Jurewicz, and Monroe Kennedy III. 2023. DenseTact 2.0: Optical Tactile Sensor for Shape and Force Reconstruction. arXiv:2209.10122 [cs.RO] https://arxiv.org/abs/2209.10122

[11] S. Dong et al. 2017. Improved GelSight Tactile Sensor for Measuring Geometry and Slip. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

[12] James M Duncan and Chin-Yung Chang. 1970. Nonlinear analysis of stress and strain in soils. *Journal of the soil mechanics and foundations division* 96, 5 (1970), 1629–1653.

[13] W. Fan et al. 2024. ViTacTip: Design and Verification of a Novel Biomimetic Physical Vision-Tactile Fusion Sensor. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1056–1062.

[14] W. Fan, H. Li, and D. Zhang. 2024. MagicTac: A Novel High-Resolution 3D Multi-layer Grid-Based Tactile Sensor. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 388–394.

[15] A. Fernandez et al. 2021. Visiflex: A Low-Cost Compliant Tactile Fingertip for Force, Torque, and Contact Sensing. *IEEE Robotics and Automation Letters* (2021).

[16] D. Gomes et al. 2020. GelTip: A Finger-shaped Optical Tactile Sensor for Robotic Manipulation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Sun Jian. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778.

[18] Francois Robert Hogan, Jean-Francois Tremblay, Bobak H Baghi, and Gregory Dudek. 2021. Seeing Through Your Skin: Recognizing Objects With a Novel Visuotactile Sensor. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 232–241.

[19] R. S. Johansson and Å. B. Vallbo. 1979. Tactile sensibility in the human hand: Relative and absolute densities of four types of mechanoreceptive units in glabrous skin. *The Journal of Physiology* 286, 1 (1979), 283–300. doi:10.1113/jphysiol.1979.sp012619

[20] Micah K Johnson and Edward H Adelson. 2009. Retrographic sensing for the measurement of surface texture and shape. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1070–1077.

[21] Kazuito Kamiyama, K. Vlack, T. Mizota, H. Kajimoto, N. Kawakami, and S. Tachi. 2005. Vision-based sensor for real-time measuring of surface traction fields. *IEEE Computer Graphics and Applications* 25, 1 (2005), 68–75.

[22] Robert L Kondner. 1963. Hyperbolic stress-strain response: cohesive soils. *Journal of the Soil Mechanics and Foundations Division* 89, 1 (1963), 115–143.

[23] N. Kuppuswamy et al. 2020. Soft-Bubble grippers for robust and perceptive manipulation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 9917–9924.

[24] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Brandon Benjamin, Dave Maurice, Roberto Pister, Roberto Calandra, Mayank Mudigonda, Ilge Akkaya, et al. 2020. DIGIT: A Novel Design for a Low-Cost Compact High-Resolution Tactile Sensor. *IEEE Robotics and Automation Letters* 5, 3 (2020), 3838–3845.

[25] Seoyeon Lee, Mohammad Amin Mirzaee, and Wenzhen Yuan. 2026. GelSphere: An Omnidirectional Rolling Vision-Based Tactile Sensor for Online 3D Reconstruction and Normal Force Estimation. *arXiv preprint arXiv:2603.14104* (2026).

[26] N. Lepora et al. 2022. DigiTac: A DIGIT-TacTip Hybrid Tactile Sensor for Comparing Low-Cost High-Resolution Robot Touch. *IEEE Robotics and Automation Letters* (2022).

[27] R. Li et al. 2014. Localization and Manipulation of Small Parts Using GelSight Tactile Sensing. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

[28] Changyi Lin, Ziqi Lin, Shaoxiong Wang, and Huazhe Xu. 2022. DTact: A Vision-Based Tactile Sensor that Measures High-Resolution 3D Geometry Directly from Darkness. arXiv:2209.13916 [cs.RO] https://arxiv.org/abs/2209.13916

[29] Changyi Lin, Han Zhang, Jikai Xu, Lei Wu, and Huazhe Xu. 2023. 9DTact: A Compact Vision-Based Tactile Sensor for Accurate 3D Shape Reconstruction and Generalizable 6D Force Estimation. *IEEE Robotics and Automation Letters* (2023).

[30] H. Lin et al. 2019. Sensing the Frictional State of a Robotic Skin via Subtractive Color Mixing. *IEEE Robotics and Automation Letters* (2019).

[31] H. Maekawa et al. 1993. A finger-shaped tactile sensor using an optical waveguide. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC)*.

[32] National Oceanic and Atmospheric Administration (NOAA). 2024. Air Pressure: JetStream - An Online School for Weather. https://www.noaa.gov/jetstream/atmosphere/air-pressure Accessed: 2026-03-09.

[33] J. Ning et al. 2022. MoiréPose: Ultra High Precision Camera-to-Screen Pose Estimation based on Moiré Pattern. In *Proceedings of the 28th Annual International Conference on Mobile Computing and Networking (MobiCom)*. 103–119.

[34] J. Ning et al. 2024. MoiréVib: Micron-level Vibration Detection based on Moiré Pattern. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking (MobiCom)*. 1393–1407.

[35] J. Ning et al. 2024. MoiréVision: A Generalized Moiré-based Mechanism for 6-DoF Motion Sensing. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking (MobiCom)*. 467–481.

[36] A. Padmanabha et al. 2020. OmniTact: A Multi-Directional High-Resolution Touch Sensor. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 618–624.

[37] Sung-Jin Park, Mattia Gazzola, Kyung Soo Park, Shirley Park, Valentina Di Santo, Erin L Blevins, Johan U Lind, Patrick H Campbell, Stephanie Dauth, Andrew K Capulli, et al. 2016. Phototactic guidance of a tissue-engineered soft-robotic ray. *Science* 353, 6295 (2016), 158–162.

[38] S. Qiu, H. Amata, and W. Heidrich. 2023. MoiréTag: Angular Measurement and Tracking with a Passive Marker. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–10.

[39] Ali M Reza. 2004. Realization of the contrast limited adaptive histogram equalization (CLAHE) for real-time image enhancement. *Journal of VLSI signal processing systems for signal, image and video technology* 38, 1 (2004), 35–44.

[40] K. Sato et al. 2008. Measurement of force vector field of robotic finger using vision-based haptic sensor. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

[41] C. Sferrazza et al. 2019. Design, Motivation and Evaluation of a Full-Resolution Optical Tactile Sensor. *Sensors* (2019).

[42] Andrea Sipos, William van den Bogert, and Nima Fazeli. 2024. GelSlim 4.0: Focusing on Touch and Reproducibility. arXiv:2409.19770 [cs.RO] doi:10.48550/arXiv.2409.19770

[43] K.-W. Sou et al. 2025. MoiréTac: A Dual-Mode Visuotactile Sensor for Multidimensional Perception Using Moiré Pattern Amplification. *arXiv preprint arXiv:2509.12714* (2025).

[44] Huanbo Sun, Katherine J. Kuchenbecker, and Georg Martius. 2022. A soft thumb-sized vision-based sensor with accurate all-round force perception. *Nature Machine Intelligence* 4, 2 (2022), 135–145. doi:10.1038/s42256-021-00439-3

[45] I. Taylor et al. 2022. GelSlim 3.0: High-Resolution Measurement of Shape, Force and Slip in a Compact Tactile-Sensing Finger. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[46] Mohsin I Tiwana, Stephen J Redmond, and Nigel H Lovell. 2012. A review of tactile sensing technologies with applications in biomedical engineering. *Sensors and Actuators A: Physical* 179 (2012), 17–31.

[47] S. Wang et al. 2021. GelSight Wedge: Measuring High-Resolution 3D Contact Geometry with a Compact Robot Finger. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[48] Yanzhe Wang, Haotian Guo, Hao Wu, and Huixu Dong. 2025. Flexible robotic hand harnesses large deformations for full-coverage human-like multimodal haptic perception. *Nature Communications* (2025).

[49] B. Ward-Cherrier et al. 2018. The TacTip Family: Soft Optical Tactile Sensors with 3D-Printed Biomimetic Morphologies. *Soft Robotics* 5, 2 (2018), 216–227.

[50] G. Westling and R. S. Johansson. 1984. Factors influencing the force control during precision grip. *Experimental Brain Research* (1984).

[51] H. E. Wheat, L. M. Salo, and A. W. Goodwin. 2004. Human ability to scale and discriminate forces typical of those occurring during grasp and manipulation. *Journal of Neuroscience* 24, 13 (2004), 3394–3401. doi:10.1523/JNEUROSCI.4822-03.2004

[52] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. 2018. CBAM: Convolutional Block Attention Module. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 3–19.

[53] C. Xiao and C. Zheng. 2021. MoiréBoard: A Stable, Accurate and Low-cost Camera Tracking Method. In *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology (UIST)*. 881–889.

[54] Jikai Xu, Lei Wu, Changyi Lin, Ding Zhao, and Huazhe Xu. 2024. DTactive: A Vision-Based Tactile Sensor with Active Surface. arXiv:2410.08337 [cs.RO] https://arxiv.org/abs/2410.08337

[55] Akihiko Yamaguchi and Christopher G. Atkeson. 2016. Combining Finger Vision and Optical Tactile Sensing: Reducing and Handling Errors While Cutting Vegetables. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, Cancun, Mexico, 1045–1051. doi:10.1109/HUMANOIDS.2016.7803400

[56] O. H. Yeoh. 1993. Some forms of the strain energy function for rubber. *Rubber Chemistry and Technology* 66, 5 (1993), 754–771.

[57] Hassan Yousef, Mehdi Boukallel, and Kaspar Althoefer. 2011. Tactile sensing for dexterous in-hand manipulation in robotics—A review. *Sensors and Actuators A: Physical* 167, 2 (2011), 171–187.

[58] W. Yuan et al. 2015. Measurement of Shear and Slip with a GelSight Tactile Sensor. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.

[59] Wenzhen Yuan, Siyuan Dong, and Edward H Adelson. 2017. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors* 17, 12 (2017), 2762.

[60] M. Zaid et al. 2022. Elastomer-Based Visuotactile Sensor for Normality of Robotic Manufacturing Systems. *Polymers* (2022).

[61] Hongqiang Zhang, Lupeng Zhang, Chengcheng Zhao, Yuanchao Shu, Peng Cheng, Jiming Chen, and Jie Xiong. 2026. MoiréEar: Moiré Can See What You Cannot Hear. In *Proceedings of the 2026 ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys ’26)*. ACM, 947–960. doi:10.1145/3774906.3802760

[62] Corey Zheng and Shu Jia. 2025. Bioinspired photoresponsive soft robotic lens. *Science Robotics* 10, 107 (2025), eadw8905.

[63] Y. Zheng et al. 2023. DotView: A Low-Cost Compact Tactile Sensor for Pressure, Shear, and Torsion Estimation. *IEEE Robotics and Automation Letters* (2023).

[64] Linzhen Zhu, Runqiu Wang, Yi Rong, and Ke Sun. 2026. MoiréLens: Bringing Schlieren Imaging into Real-World Environments Using Moiré Patterns. In *Proceedings of the 2026 ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys ’26)*. ACM, 1274–1287. doi:10.1145/3774906.3802765