

NLiPsCalib: An Efficient Calibration Framework for High-Fidelity 3D Reconstruction of Curved Visuotactile Sensors

Xuhao Qin*, Feiyu Zhao*, Yatao Leng, Runze Hu, and Chenxi Xiao[†]

Abstract—Recent advances in visuotactile sensors increasingly employ biomimetic curved surfaces to enhance sensorimotor capabilities. Although such curved visuotactile sensors enable more conformal object contact, their perceptual quality is often degraded by non-uniform illumination, which reduces reconstruction accuracy and typically necessitates calibration. Existing calibration methods commonly rely on customized indenters and specialized devices to collect large-scale photometric data, but these processes are expensive and labor-intensive. To overcome these calibration challenges, we present NLiPsCalib, a physics-consistent and efficient calibration framework for curved visuotactile sensors. NLiPsCalib integrates controllable near-field light sources and leverages Near-Light Photometric Stereo (NLiPs) to estimate contact geometry, simplifying calibration to just a few simple contacts with everyday objects. We further introduce NLiPsTac, a controllable-light-source tactile sensor developed to validate our framework. Experimental results demonstrate that our approach enables high-fidelity 3D reconstruction across diverse curved form factors with a simple calibration procedure. We emphasize that our approach lowers the barrier to developing customized visuotactile sensors of diverse geometries, thereby making visuotactile sensing more accessible to the broader community.

I. INTRODUCTION

Robots often adopt biomimetic or functional designs in both morphology and sensory systems [1]. In line with this trend, robotic tactile sensors have evolved beyond conventional flat architectures to encompass diverse form factors. For example, in humanoid robots, curved and rounded fingertip-inspired sensors are now widely used in dexterous robotic hands, as they mimic the geometry of fingertips and finger pads, enabling stable, conformal contact with objects and supporting omnidirectional perception [2]–[5]. Beyond humanoid robots, application-specific tactile sensors are also emerging, including cylindrical designs for industrial arms [6] and slim-profile sensors for minimally invasive robotic surgery [7], [8]. Thus, the capability to design and fabricate tactile sensors in customized shapes is essential for broadening the scope of tactile perception and manipulation.

* Equal contribution.

[†] Corresponding author.

This work was supported by the Natural Science Foundation of Shanghai (Grant No. 25ZR1402370), and partially by Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI), MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration (KLIP-HuMaCo). The experiments were supported by the Core Facility Platform of Computer Science and Communication, SIST, ShanghaiTech University.

Xuhao Qin, Feiyu Zhao, Yatao Leng, Runze Hu, and Chenxi Xiao are with the School of Information Science and Technology, ShanghaiTech University. {qinxh2024, zhaofy12024, lengyt2023, hurz2024, xiaochx}@shanghaitech.edu.cn

Project webpage: <https://nlipscalib.github.io/>




Fig. 1: NLiPsCalib, a new calibration pipeline for curved visuotactile sensors that requires no specialized devices. It leverages casual presses with everyday objects, followed by near-light photometric stereo to obtain accurate geometries for building the calibration dataset. This dataset further enables training a neural network for real-time normal inference.

Among these technologies, visuotactile sensors have gained prominence due to their high-resolution 3D shape sensing capabilities and inherently customizable form factors [9]. Over the past decade, numerous designs have been introduced and deployed on robotic manipulators. Representative examples range from flat sensors like GelSight [10], 9DTact [11], and DelTact [12] to curved, finger-shaped sensors such as the Soft Finger Sensor [13] and GelSight Svelte [14].

Despite this progress, reliably extracting high-fidelity depth information from a sensor’s output remains a persistent challenge. The standard pipeline for this task leverages photometric stereo: first, illumination patterns are mapped to surface normals; then, the normals are numerically integrated to reconstruct the 3D shape [10], [15]. A critical bottleneck of this approach lies in establishing an accurate mapping from illumination to normals. This mapping is complex because the internal illumination is inherently non-uniform: light intensity varies across the surface due to the elastomer’s curvature and near-field effects from embedded light sources, which cause distance-dependent attenuation.

To model these compound lighting effects, existing meth-

ods typically employ data-driven approaches that learn the mapping from collected intensity images to ground-truth surface normals. However, acquiring this ground-truth data is non-trivial. It requires either constructing a high-fidelity digital twin of the sensor for simulation [16]–[18] or using precisely calibrated hardware such as CNC-machined probes or professional indentation devices [19]. Both approaches are costly in time, effort, and resources, directly undermining the goal of rapid and accessible customization for visuotactile sensors.

In this paper, we introduce NLiPsCalib, a novel calibration technique for visuotactile sensors that eliminates the need for specialized equipment (Fig. 1). Our key insight is that the physics-based Near-Light Photometric Stereo (NLiPs) model [20] is highly consistent with the internal illumination conditions of curved visuotactile sensors. This enables high-quality shape calibration using only the sensor’s multiple internal light sources. The calibration process is simple: it requires only a few physical interactions, such as pressing the sensor against textured everyday objects. For each pressed indentation, NLiPsCalib performs a single run of Near-Light Photometric Stereo to capture the resulting image data, from which it directly estimates the indented geometry and surface normals without requiring the ground-truth shape of the indenting object, thereby significantly simplifying the creation of a calibration dataset.

We evaluate the NLiPsCalib framework on a physical sensor. To this end, we introduce NLiPsTac, a modular tactile sensor featuring individually controllable light sources and an adaptable structure for various elastomer form factors. Performing calibration on this sensor requires around 50 casual presses on everyday objects, and the resultant calibration quality is comparable to state-of-the-art device-based methods in terms of reconstruction error [6], [10], [21]. Furthermore, we validate our approach on elastomers of different shapes to demonstrate its adaptability. By significantly reducing the effort and expertise required for calibration, our work makes the design and deployment of custom visuotactile sensors more practical and accessible.

In summary, our main contributions are:

- 1) **NLiPsCalib**: A physics-based and data-efficient calibration framework for curved visuotactile sensors that requires no specialized external hardware.
- 2) The adaptation and validation of the Near-Light Photometric Stereo (NLiPs) model for 3D shape reconstruction in tactile sensing.
- 3) **NLiPsTac**: A hardware platform designed for developing and testing near-light visuotactile sensors.
- 4) A comprehensive evaluation of the proposed technique, along with discussions and insights for promoting broader usage.

II. RELATED WORKS

A. Visuotactile Sensors

Tactile sensing is essential for robotic manipulation, as it provides contact information that vision alone cannot

TABLE I: Methods and devices for calibrating curved visuotactile sensors often require external machines or specialized parts, whereas our sensor achieves calibration without them.

Sensor	Calibration Method or Device
GelSight360 [19]	CNC machine indentation
RainbowSight [21]	CNC machine indentation
DenseTact 2.0 [26]	CNC machine indentation
GelSplitter3D [17]	Robotic arm indentation
HumanFT [27]	3D-printed calibration rig
R-Tac [25]	3D-printed calibration rig
NLiPsTac (Ours)	NLiPs + Everyday objects

capture. Among various designs, visuotactile sensors have gained particular attention for their ability to deliver high-resolution contact geometry. Early work mainly focused on planar sensors, such as GelSight [10], 9DTact [11], and DelTact [12], which demonstrated strong performance in dexterous manipulation tasks. However, planar structures are most suitable for parallel grippers and are less compatible with biomimetic mechanisms such as humanoid fingertips [22], arms [23], or other curved morphologies.

More recently, curved elastomer geometries have been increasingly adopted to improve contact coverage and stability. Examples include the Soft Finger Sensor [13], OmniTact [5], GelSight Svelte [14], GelStereo Tip [24] and R-Tac [25], which feature fingertip-inspired designs that enable stable, multi-directional contact. Building on this trend, we introduce **NLiPsTac**, a curved visuotactile sensor inspired by the human fingertip, designed to provide omnidirectional contact perception.

B. Reconstruction and Calibration for Visuotactile Sensing

Accurate 3D reconstruction of contact geometry is fundamental to a wide range of tactile perception tasks, empowering the estimation of object’s shape [28], [29], pose [30], [31], and texture [32]. Traditionally, reconstruction has been performed using different principles. For example, GelSight employs photometric stereo, while DTact relies on light intensity under single-light illumination [33]. However, these methods depend on mappings between light intensity and either surface shape or normals, and the quality of such mappings directly affects reconstruction accuracy. Moreover, while mappings can be reliably established on flat surfaces, curved-surface sensors inherently suffer from uneven illumination and near-field effects [34], which invalidate parallel-light assumptions and compromise reconstruction stability.

To improve mapping fidelity, most existing approaches perform sensor calibration prior to reconstruction. For example, mechanical indentation is a widely used method [10], [17], [19], [21], typically involving a standard ball probe driven by a CNC machine or a robotic platform. Some studies have also employed customized 3D-printed parts for controlled indentation [25]–[27]. Across these works, the key requirement is knowledge of the indenter’s pose and geometry, which makes the ground-truth surface normals

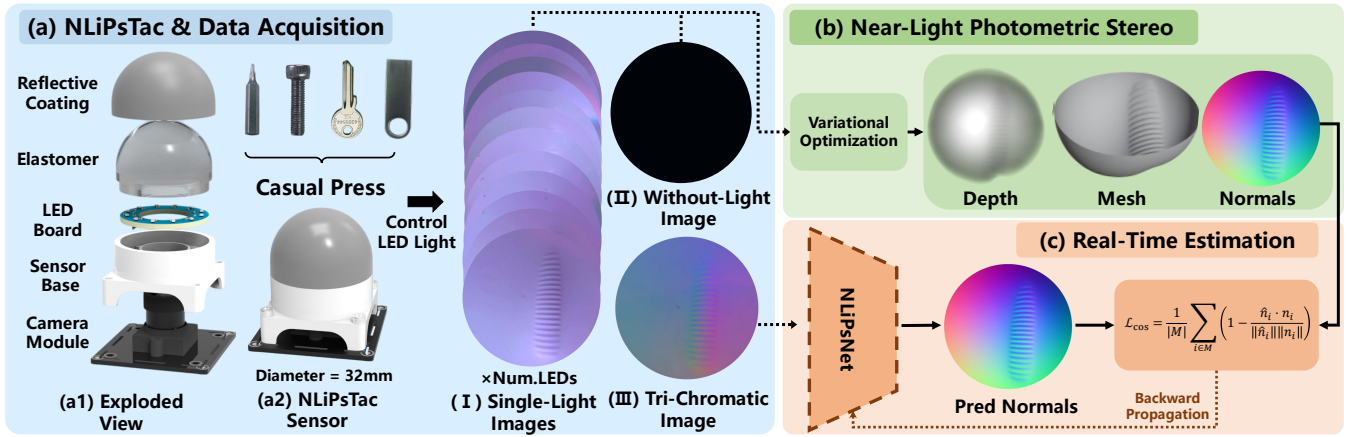


Fig. 2: System pipeline. Using the proposed NLiPsTac tactile sensor, the framework collects a calibration dataset with NLiPs, enabling the training of NLiPsNet, a network designed for real-time 3D shape inference under trichromatic illumination.

accessible and thus enables pairing with RGB observations. Although effective, these calibration procedures are often complex, requiring specialized hardware (e.g., CNC machines or spherical probes) or extensive datasets, which increases cost and limits scalability.

To address these challenges, we propose **NLiPsCalib**, a calibration framework based on Near-Light Photometric Stereo (NLiPs) [20]. By explicitly modeling attenuation and the spatial distribution of near-field illumination, NLiPsCalib enables efficient and high-quality reconstruction while greatly simplifying calibration, thereby reducing reliance on specialized hardware and large datasets. This approach facilitates fast, low-cost, and high-fidelity tactile perception in curved-surface visuotactile sensors. Comparison between NLiPsCalib and previous works is shown in Tab. I.

III. OVERVIEW

The goal of calibration is to establish the mapping between light intensities and surface geometry, serving as a prerequisite for high-fidelity 3D reconstruction. Specifically, the calibration process aims to obtain the mapping function f between pixel intensities and surface normals, expressed as:

$$\mathbf{n} = f(u, v, r, g, b), \quad (1)$$

where (u, v) denote pixel coordinates and (r, g, b) represent the corresponding light intensities in R, G, B color channels, respectively.

A key challenge in calibration lies in obtaining accurate geometry \mathbf{n} at the indented location, enabling the derivation of function f for normal inference during test time.

To address this challenge, we introduce **NLiPsCalib**, a unified sensor calibration framework that derives ground-truth geometries directly from multi-source photometric cues. The design goal of **NLiPsCalib** is to infer surface normals \mathbf{n} directly using near-light photometric stereo (Sec. IV-A). This approach offers two key advantages. First, it enables direct estimation without requiring costly equipment such as CNC-machined probes or robotic arms. Second, it avoids spatial misalignment between ground-truth normals \mathbf{n} and intensity measurements (u, v, r, g, b) , thereby enabling the creation

of high-fidelity paired calibration datasets for learning f . Based on the dataset, we further present **NLiPsNet**, a neural network for real-time normal prediction during deployment (Sec. IV-B).

Finally, we describe the design of a novel tactile sensor that serves as a testbed for validating the proposed framework (Sec. V). The overall system pipeline is illustrated in Fig. 2.

IV. SENSOR CALIBRATION FOR NORMAL ESTIMATION

A. Geometry Calibration with Near-light Photometric Stereo

Our calibration pipeline generates high-fidelity surface geometry directly from sensor images, eliminating the need for external measurement hardware. To achieve this, we employ NLiPs framework, which is originally developed for high fidelity 3D reconstruction tasks such as 3D scanning for surface micro-textures [20]. In this work, we adapt NLiPs to the optical environment of visuotactile sensors. The core contribution lies in its novel use for generating high-fidelity ground-truth data without external devices. In contrast to parallel illumination commonly used in previous works [10], NLiPs is more suitable for non-uniform illumination, thus naturally fits curved surfaces.

The goal of NLiPs is to compute a dense, pixel-wise depth map, $z(\mathbf{p})$, of the deformed surface. From this depth map, we can then derive the surface normals, estimate the surface albedo, and directly reconstruct the 3D surface points as a dense point cloud. The process involves two key components: a physical model of light intensity and an optimization procedure to estimate the depth map.

a) The Physical Model of Light Intensity: The NLiPs model is built on a physical description of how light from a known point source reflects off an unknown surface and is captured by the camera. We define the sensor surface in the camera's coordinate system. For each pixel $\mathbf{p} = (u_p, v_p)$, the corresponding 3D point on the surface is $\mathbf{x}(\mathbf{p}) = (u, v, z(\mathbf{p}))$. Here, (u, v) are the pixel coordinates converted to metric units using the camera's intrinsic parameters, and $z(\mathbf{p})$ is the unknown depth value we aim to recover.

For the i -th LED, located at a pre-calibrated 3D position \mathbf{x}_s^i , the intensity $I_{i,c}(\mathbf{p})$ observed at a pixel \mathbf{p} in color

channel $c \in \{R, G, B\}$ is given by:

$$I_{i,c}(\mathbf{p}) = \Psi_{i,c} \rho_c(\mathbf{p}) \left[\frac{\mathbf{n}_s^i \cdot (\mathbf{x}(\mathbf{p}) - \mathbf{x}_s^i)}{\|\mathbf{x}(\mathbf{p}) - \mathbf{x}_s^i\|} \right]^{\mu^i} \times \frac{\{(\mathbf{x}_s^i - \mathbf{x}(\mathbf{p})) \cdot \mathbf{n}(\mathbf{p})\}_+}{\|\mathbf{x}_s^i - \mathbf{x}(\mathbf{p})\|^3} \quad (2)$$

In this model, the unknowns are the surface geometry and its reflectance. The geometry is described by the 3D positions $\mathbf{x}(\mathbf{p})$ and the unit surface normals $\mathbf{n}(\mathbf{p})$, while reflectance is described by the albedo $\rho_c(\mathbf{p})$. Note that both $\mathbf{x}(\mathbf{p})$ and $\mathbf{n}(\mathbf{p})$ are fundamentally determined by the depth map $z(\mathbf{p})$. The remaining terms are known parameters of the light sources: $\Psi_{i,c}$ is the calibrated intensity, \mathbf{n}_s^i is the principal direction (Fig. 3), and μ^i is the anisotropy parameter. The operator $\{\cdot\}_+ = \max(\cdot, 0)$ accounts for self-shadowing.

b) Calculating the Depth Map: With the intensity model in place, the task is to recover the surface geometry (represented by the depth map $z(\mathbf{p})$) from multiple intensity images captured under different lighting conditions. In our case, these conditions are generated by sequentially activating the LEDs inside the sensor (refer to Sec. V). Each lighting condition provides data for solving Eq. (2).

Based on captured data, directly solving Eq. (2) for $z(\mathbf{p})$ is difficult because the problem is highly non-linear, mainly due to the depth-dependent distance terms in the denominator. In addition, it is not possible to solve the surface normal $\mathbf{n}(\mathbf{p})$ independent of the depth $z(\mathbf{p})$, since this breaks the geometric constraint that the normal is the gradient of the surface, which may result in inconsistent results.

To address these challenges, we adopt the variational optimization framework from [20]. This framework resolves this by framing the optimization solely in terms of the depth map $z(\mathbf{p})$. In other words, the surface normal $\mathbf{n}(\mathbf{p})$ is not an independent unknown. Instead, it is explicitly computed as a function of the depth map's spatial derivatives at each iteration. For a surface defined by $\mathbf{x}(\mathbf{p}) = (u, v, z(\mathbf{p}))$, the normal is derived from its partial derivatives:

$$\mathbf{n}(\mathbf{p}) \propto \left(-\frac{\partial z}{\partial u}, -\frac{\partial z}{\partial v}, 1 \right). \quad (3)$$

By optimizing over $z(\mathbf{p})$ alone, this approach ensures that the recovered depth and normals are always geometrically consistent, producing a valid, integrable surface.

For stable and efficient optimization, we also follow [20] and [35] to introduce a log-depth parameterization. This is achieved by setting our optimization variable to be $\tilde{z}(\mathbf{p}) = \log z(\mathbf{p})$. This transformation offers two key advantages: it removes the physical constraint of positive depth ($z > 0$) by allowing the unconstrained variable $\tilde{z}(\mathbf{p})$ to span all real numbers, and it helps improve numerical stability by linearizing the image formation model.

With this re-parameterization, the calibration problem can be reformulated as minimizing a global energy functional $E(\tilde{z}, \tilde{\rho})$, which jointly estimates the log-depth map \tilde{z} and the ‘‘effective albedo’’ $\tilde{\rho}$ from the observed images:




Fig. 3: Optical path of near-light photometric stereo calibration, showing illumination from embedded point light sources to the sensor surface, which is then observed by the camera. Controllable LED point light sources are arranged in a ring ($S_0 - S_k$) with known coordinates, and illumination from a selected source (e.g., S_1) defines the shaded region.

$$\min_{\tilde{z}, \tilde{\rho}} E(\tilde{z}, \tilde{\rho}) = \sum_{\mathbf{p}, i, c} \left(I_{i,c}^{\text{obs}}(\mathbf{p}) - I_{i,c}^{\text{pre}}(\mathbf{p}; \tilde{z}, \tilde{\rho}) \right)^2 + \zeta \sum_{\mathbf{p}} (\tilde{z}(\mathbf{p}) - \tilde{z}_0(\mathbf{p}))^2, \quad (4)$$

where the first term penalizes the squared difference between the observed intensities (I^{obs}) and those predicted by our reformulated model (I^{pre}), and the second term is a regularization term with respect to an initial prior \tilde{z}_0 , whose strength is controlled by the weight ζ . Following [20], the prior \tilde{z}_0 is initialized using the nominal camera-surface distance. Although this initialization influences the optimization process, we found the reconstructed local shape is robust to variations of \tilde{z}_0 .

To solve this optimization problem, we follow [20] and use an Alternating Reweighted Least Squares (ARLS) algorithm. This method iteratively refines the solution by alternating between two main steps:

- 1) **Albedo Update:** With the current log-depth map \tilde{z} held fixed, the problem of finding the optimal effective albedo $\tilde{\rho}_c$ simplifies to a pixel-wise linear least-squares problem, which has a closed-form solution.
- 2) **Depth Update:** With the albedo $\tilde{\rho}_c$ fixed, the energy functional is linearized with respect to the log-depth \tilde{z} . This results in a large but sparse system of linear equations, which is solved efficiently for an updated \tilde{z} using a Gauss-Newton step with a preconditioned conjugate gradient (PCG) solver.

By iteratively applying these two steps, the energy functional decreases monotonically until convergence. The final output is a high-fidelity logarithmic depth map \tilde{z} , from which we recover the physical 3D surface $z = \exp(\tilde{z})$ and its corresponding normal field, serving as the ground-truth geometry for our calibration dataset.

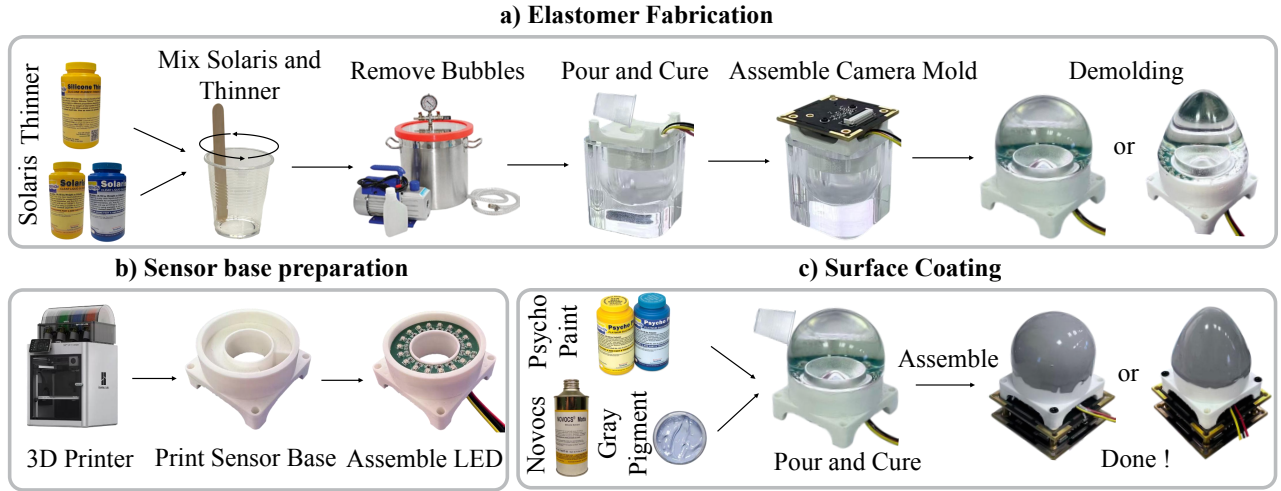


Fig. 4: Fabrication of the NLiPsTac sensor, including (a) casting a clear elastomer for light transmission and contact support, (b) assembling the 3D-printed base with the LED board, and (c) casting a coating layer to form the reflective surface.

B. Real-Time Normal Estimation

Near-light photometric stereo excels in reconstruction quality but is not well suited for real-time applications. This limitation arises primarily because it requires capturing multiple images, which increases acquisition time, and because the optimization process is computationally expensive. To address this issue, we first construct a training dataset using the approach introduced in Sec. IV-A, and then develop a neural network to accelerate the inference process.

a) Dataset Acquisition: The calibration dataset is collected through indentation experiments. Unlike previous works, our approach (Sec. IV-A) allows direct indentation of the sensor’s elastomer using textured everyday objects.

For each indentation, we record a set of calibration images. Using a sensor with 12 LEDs as an example, we collect: (i) 12 single-light RGB images, each with one LED activated under *white* illumination; (ii) one *dark* image with all LEDs off; and (iii) one tri-chromatic RGB image with all 12 LEDs activated simultaneously in three color groups (1–4 red, 5–8 green, 9–12 blue). The NLiPs model uses the single-light and dark images to reconstruct surface normals of the deformed region, which serve as ground-truth supervision. This enables the network to learn a mapping from tri-chromatic RGB images to surface normal ground truths.

b) Network Architecture: We then propose *NLiPsNet*, which infers the normal vector \mathbf{n} from a 5-dimensional input (u, v, r, g, b) (Eq. (1)). The network takes as input a single tactile image captured with all LEDs activated, following the color configuration in [36].

NLiPsNet is implemented as a lightweight multilayer perceptron with three hidden layers (256–256–128) and an output layer predicting a 3-dimensional normal vector. Batch normalization, ReLU activation, and dropout ($p = 0.2$) are applied to improve training stability and reduce overfitting. During both training and inference, the predicted normals are normalized to unit length.

c) Training Objective: We use a cosine similarity loss to train the network, enforcing angular alignment between

the predicted normal $\hat{\mathbf{n}}_i$ and the ground-truth normal \mathbf{n}_i for each valid pixel $i \in M$:

$$\mathcal{L}_{\cos} = \frac{1}{|M|} \sum_{i \in M} \left(1 - \frac{\hat{\mathbf{n}}_i \cdot \mathbf{n}_i}{\|\hat{\mathbf{n}}_i\| \|\mathbf{n}_i\|} \right). \quad (5)$$

V. SENSOR DESIGN AND FABRICATION

A. Design Goals

We present NLiPsTac, a visuotactile sensor designed to validate our NLiPsCalib framework. NLiPsTac includes three key features: (i) integration of multiple individually controllable LEDs, providing directional illumination of the coating; (ii) does not include a light diffuser. Unlike sensors that approximate parallel lighting, each LED in our design acts as a point light source; and (iii) housing both the camera and LEDs within a single optical medium, minimizing refraction effects. These design choices align the sensor’s optical behavior with the physical assumptions of the NLiPs model, reducing model mismatch and simplifying algorithmic adaptation. The fabrication process of NLiPsTac is shown in Fig. 4.

B. Details of Sensor Design and Fabrication

Camera. As the core imaging unit of the visuotactile sensor, we use an IMX274 camera module (30 FPS, Sony Inc.) equipped with an M14 lens. The 8 MP CMOS sensor supports a streaming resolution of 4K (3840×2160 pixels). M14 optical lens provides a 100° field of view for comprehensive visualization of the dome-shaped elastomer surface.

Sensor Base. The base is a 3D-printed PLA structure that supports the LED circuit board, elastomer, and the camera module.

LED Circuit Board. The LED circuit board provides illumination for the sensor. It is a printed circuit board (FR4 substrate) populated with WS2812 addressable LEDs. Since near-field photometric stereo requires illumination from multiple directions, each LED can be individually controlled to emit trichromatic light at a user-specified intensity.




Fig. 5: Normal estimation results for sphere and cube indenters. Columns show **Indenter**, **Pressed**, **Normals**, **GT (zoomed)**, **Normals (zoomed)**, and **Error Map**. The shared colorbar maps blue \rightarrow red to errors from 0° to 25° .

Elastomer. The elastomer is designed to transmit light to the surface, support the reflective coating, and withstand mechanical contact during deformation. Following [37], we used Solaris (Smooth-On Inc.) as the base gel material and incorporated Thinner (Smooth-On Inc.) to enhance deformability by reducing hardness. The gel casting process employed a custom acrylic mold to avoid surface mold lines. The A and B components of Solaris were mixed with Thinner in a 1:1:0.8 ratio, and the sensor base was secured in the mold before pouring. Following [27], a camera mold was carefully inserted into the reserved hole to ensure a seamless fit between the camera and elastomer.

Reflective Coating. The reflective coating is designed to approximate Lambertian reflectance. For fabrication, we used Psycho Paint and NOVOCS solvent (Smooth-On Inc.) with gray pigment. Following [25], the A and B components of Psycho Paint, NOVOCS solvent, and gray pigment were mixed in a 1:1:1:0.3 ratio, and the resulting solution was uniformly sprayed onto the demolded elastomer surface.

VI. EXPERIMENTS

In this section, we conduct experiments to validate the proposed calibration method, aiming to answer the following questions: (Q1) Can the NLiPsCalib approach obtain high-fidelity calibration results (Sec.VI-A)? (Q2) Can NLiPsNet achieve real-time, high-fidelity normal estimation (Sec.VI-B)? (Q3) Can the calibration pipeline be applied to elastomers of different form factors (Sec.VI-C)? In addition, we perform extensive experiments to investigate factors that influence calibration fidelity (Sec.VI-D).

A. Validation of NLiPs-Based Calibration

We first evaluate whether NLiPsCalib can produce high-fidelity calibration results (Q1). To this end, we consider two representative cases: a small sphere (curved surface) and a cube (planar surface), both with analytically known surface normals. Using NLiPsCalib, the reconstructed normals were compared against the analytical ground truth. As shown in Fig. 5, the calibrations closely match the analytical normals across the sensor.

For quantitative evaluation, we report accuracy using the Average Angular Error (AAE, ranging from 0° to 180°), which measures the angular deviation between reconstructed and ground-truth normals, and the Mean Absolute Error of normal components (MabsE), which mea-

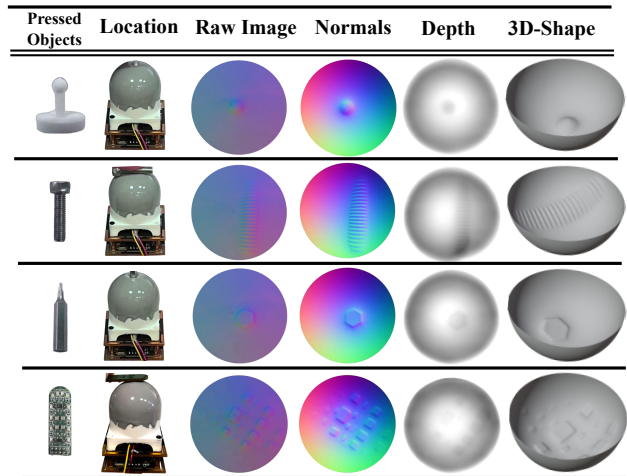


Fig. 6: Samples from the NLiPsNet training dataset. Each row shows one pressing instance: the input *probe* image, *indentation location*, NLiPs-recovered *normals*, and *3D shape*.

sures the absolute difference across the x , y , and z components: $\frac{1}{3|M|} \sum_{i \in M} \|\hat{\mathbf{n}}_i - \mathbf{n}_i\|_1$.

Our method achieves an AAE of 7.0415° and an MabsE of 0.0588. These results confirm that NLiPsCalib produces high-fidelity calibration results, validating its effectiveness for sensor calibration (Q1).

B. Validation of Real-Time Normal Inference

Dataset Preparation. A training dataset is required to train the normal inference network NLiPsNet. To prepare this dataset, we press three distinct objects at multiple sensor locations, resulting in **50** pressing conditions following the approach described in Sec. IV-B (a). Some samples from the collected dataset are shown in Fig. 6.

Normal Inference. We next evaluate the real-time normal inference capability of NLiPsNet. To assess both accuracy and generalization, we test the network using (1) indenters included in the training dataset but applied at unseen locations, and (2) entirely novel indenters. For all test cases, we also perform NLiPs-based reconstruction to obtain ground-truth normals, enabling both qualitative and quantitative comparisons. Example indentation observations, predicted normals, and NLiPs-derived ground truth are shown in Fig. 7. The predicted normal maps closely match the ground truth across all cases, demonstrating that NLiPsNet generalizes well to novel contact geometries.

Quantitative results (Fig.7) show that the network achieves an average angular error of 3.332° on training-set indenters and 3.113° on unseen indenters. These results highlight the strong generalization capability of NLiPsNet. The error levels are comparable to or better than many recent visuotactile sensors, which often report higher errors (e.g., GelRoller [6] reports 16.17°). This validates that NLiPsCalib can effectively generate training datasets using only a few everyday objects and a few dozen indentations (Q2).

C. Generalization Across Curved Elastomer Surfaces

We further demonstrate that the proposed techniques can be applied to visuotactile sensors with diverse curved geome-
















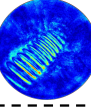
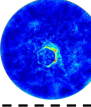
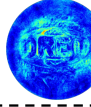
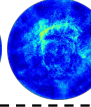
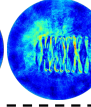
	Seen objects		Unseen objects		
	Screw	Driver	OREO	Foot	Indenter
Objects					
GT					
Normals					
Err. Map					
AAE	3.21°	2.97°	2.592°	3.257°	2.855°
MabsE	0.0217	0.0171	0.0252	0.0227	0.0285

Fig. 7: Evaluation results of NLiPsNet. Columns 1–2 correspond to objects seen during training, and Columns 3–5 to unseen objects. Rows show the pressed **Objects**, **GT** from NLiPsCalib, predicted **Normals**, and the corresponding **Err. Map** with ground truth from NLiPsCalib. The colorbar (blue to red) indicates errors from 0.00 to 0.05 in MabsE.

tries. For this purpose, we designed and fabricated elastomer domes with three distinct surface shapes. For each shape, we used NLiPsCalib to reconstruct the geometry and compared it against the corresponding ground-truth model, as shown in Fig. 8. Following the evaluation protocol in Sec. VI-A, we quantitatively measured calibration accuracy for each case. The results show that our approach consistently achieves less than 10° AAE calibration error while preserving fine-scale details across varying elastomer geometries. These findings confirm that the proposed calibration method maintains high accuracy and robustness regardless of the underlying elastomer curvature (Q3).

D. Influence of LED Channels

Since our approach leverages highly redundant LED illumination, we conducted an ablation study to analyze the impact of the number of LED channels on NLiPsCalib performance. We varied the number of active LEDs during calibration, testing configurations from 3 to 24 LEDs, and evaluated calibration accuracy by comparing the recovered geometry of a spherical indenter against the ground-truth shape. The results are summarized in Tab. II.

Results indicated that calibration accuracy depends on the number of LEDs. While acceptable results can be obtained with as few as 3 LEDs, increasing to 12 LEDs significantly improves both AAE and MabsE, indicating that denser illumination provides richer cues for surface normal estimation. Beyond 12 LEDs, however, improvements plateau: using 12 LEDs already yields a relatively low AAE, and 18 or 24 LEDs only slightly reduce MabsE. These results suggest that additional LEDs offer marginal benefits. Considering that more LEDs linearly increase calibration time with limited



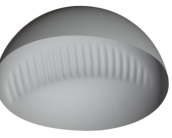


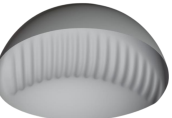

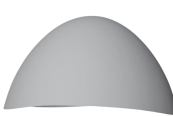

Surface	Rec. Surfs	Rec. Pressed Screw Surfs	Accuracy AAE/MabsE
			7.48°/0.0689
			9.58°/0.0817
			9.96°/0.0869

Fig. 8: The proposed approach is applicable to diverse curved elastomer surfaces. Each row corresponds to a different elastomer dome geometry. Columns show the fabricated **Surface**, reconstructed **Rec. Surfs**, reconstructed **Pressed Screws**, and quantitative **Accuracy** (AAE, MabsE).

TABLE II: Performance of NLiPsCalib with 3–24 active LEDs used for calibration and reconstruction.

#LEDs	3	6	9	12	15	18	21	24
AAE ↓	9.86	8.48	8.17	7.28	7.84	7.17	7.81	7.46
MabsE ↓	0.0858	0.0754	0.072	0.065	0.069	0.065	0.069	0.063

accuracy gains, we adopt 12 LEDs in our main experiments as a practical trade-off.

VII. CONCLUSION

We presented NLiPsCalib, an easy-to-use framework for calibrating curved visuotactile sensors. Our approach addresses a key bottleneck in the development of custom tactile sensors: the dependence on costly devices, and labor-intensive data acquisition procedures. Leveraging a NLiPs model, NLiPsCalib generates high-fidelity ground-truth geometry directly from the sensor’s internal illumination, requiring only a few simple contacts with everyday objects and eliminating the need for specialized hardware. Experiments on the NLiPsTac sensor validate the effectiveness of this approach, showing that NLiPsCalib achieves calibration efficiency and accuracy comparable to state-of-the-art methods while significantly reducing setup time and complexity. The resulting calibration data enables the lightweight neural network NLiPsNet to perform accurate, real-time 3D reconstruction from a single image, demonstrating strong generalization to unseen objects and contact geometries.

A current limitation is the computational cost of NLiPs optimization during offline calibration (3-4 minutes per reconstruction, and 3 hours for overall calibration without human intervention), as our implementation is CPU-based and not parallelized. Future work will focus on improving computational efficiency of NLiPs, potentially through GPU acceleration. Moreover, the selection of suitable objects for

calibration is also still empirical. Although the quantitative impact is difficult to measure, we provide the following guidance: objects should fully excite the sensor's normal patterns. To capture diverse normal conditions, we recommend using a few different objects that produce salient deformations and pressing them across locations covering the entire elastomer. In addition, when LEDs illuminate the surface in sequence, the object should remain stationary.

REFERENCES

- [1] J. Wang, W. Chen, X. Xiao, Y. Xu, C. Li, X. Jia, and M. Q.-H. Meng, "A survey of the development of biomimetic intelligence and robotics," *Biomimetic Intelligence and Robotics*, vol. 1, p. 100001, 2021.
- [2] 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.
- [3] H. Sun, K. J. Kuchenbecker, and G. Martius, "A soft thumb-sized vision-based sensor with accurate all-round force perception," *Nature Machine Intelligence*, vol. 4, no. 2, pp. 135–145, 2022.
- [4] P. Piacenza, K. Behrman, B. Schifferer, I. Kymissis, and M. Ciocarlie, "A sensorized multicurved robot finger with data-driven touch sensing via overlapping light signals," *IEEE/ASME Transactions on Mechatronics*, vol. 25, no. 5, pp. 2416–2427, 2020.
- [5] A. Padmanabha, F. Ebert, S. Tian, R. Calandra, C. Finn, and S. Levine, "Ommitact: A multi-directional high-resolution touch sensor," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 618–624.
- [6] Z. Zhang, H. Ma, Y. Zhou, J. Ji, and H. Yang, "Gelroller: A rolling vision-based tactile sensor for large surface reconstruction using self-supervised photometric stereo method," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 7961–7967.
- [7] Q. Yu, M. Moghani, K. Dharmarajan, V. Schorp, W. C.-H. Panitch, J. Liu, K. Hari, H. Huang, M. Mittal, K. Goldberg *et al.*, "Orbitsurgical: An open-simulation framework for learning surgical augmented dexterity," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 15 509–15 516.
- [8] W. Li, Z. Zhao, L. Cui, W. Zhang, H. Liu, L.-A. Li, and Y. Zhu, "Minitac: An ultra-compact 8 mm vision-based tactile sensor for enhanced palpation in robot-assisted minimally invasive surgery," *IEEE Robotics and Automation Letters*, vol. 9, no. 12, pp. 11 170–11 177, 2024.
- [9] A. C. Abad and A. Ranasinghe, "Visuotactile sensors with emphasis on gelsight sensor: A review," *IEEE Sensors Journal*, vol. 20, no. 14, pp. 7628–7638, 2020.
- [10] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.
- [11] C. Lin, H. Zhang, J. Xu, L. Wu, and H. Xu, "9dtact: A compact vision-based tactile sensor for accurate 3d shape reconstruction and generalizable 6d force estimation," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 923–930, 2023.
- [12] G. Zhang, Y. Du, H. Yu, and M. Y. Wang, "Deltact: A vision-based tactile sensor using a dense color pattern," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 778–10 785, 2022.
- [13] B. Romero, F. Veiga, and E. Adelson, "Soft, round, high resolution tactile fingertip sensors for dexterous robotic manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4796–4802.
- [14] J. Zhao and E. H. Adelson, "Gelsight svelte: A human finger-shaped single-camera tactile robot finger with large sensing coverage and proprioceptive sensing," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 8979–8984.
- [15] M. K. Johnson and E. H. Adelson, "Retrographic sensing for the measurement of surface texture and shape," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2009, pp. 1070–1077.
- [16] D. F. Gomes, P. Paoletti, and S. Luo, "Beyond flat gelsight sensors: Simulation of optical tactile sensors of complex morphologies for sim2real learning," *arXiv preprint arXiv:2305.12605*, 2023.
- [17] Y. Lin, X. Lu, J. Chen, and H. Yang, "3d vision-tactile reconstruction from infrared and visible images for robotic fine-grained tactile perception," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 16 023–16 029.
- [18] A. Agarwal, A. Wilson, T. Man, E. Adelson, I. Gkioulekas, and W. Yuan, "Vision-based tactile sensor design using physically based rendering," *Communications Engineering*, vol. 4, no. 1, p. 21, 2025.
- [19] M. H. Tippur and E. H. Adelson, "Gelsight360: An omnidirectional camera-based tactile sensor for dexterous robotic manipulation," in *2023 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2023, pp. 1–8.
- [20] Y. Quéau, B. Durix, T. Wu, D. Cremers, F. Lauze, and J.-D. Durou, "Led-based photometric stereo: Modeling, calibration and numerical solution," *Journal of Mathematical Imaging and Vision*, vol. 60, no. 3, pp. 313–340, 2018.
- [21] M. H. Tippur and E. H. Adelson, "Rainbowsight: A family of generalizable, curved, camera-based tactile sensors for shape reconstruction," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 1114–1120.
- [22] Y.-J. Kim, Y. Lee, J. Kim, J.-W. Lee, K.-M. Park, K.-S. Roh, and J.-Y. Choi, "Roboray hand: A highly backdrivable robotic hand with sensorless contact force measurements," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6712–6718.
- [23] Y.-J. Kim, "Anthropomorphic low-inertia high-stiffness manipulator for high-speed safe interaction," *IEEE Transactions on robotics*, vol. 33, no. 6, pp. 1358–1374, 2017.
- [24] B. Zhang, S. Cui, C. Zhang, J. Hu, R. Wang, and S. Wang, "Gelstereo tip: A spherical fingertip visuotactile sensor for multi-finger screwing manipulation," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 12 012–12 023, 2025.
- [25] P. Lin, Y. Huang, W. Li, J. Ma, C. Xiao, and Z. Jiao, "Pp-tac: Paper picking using tactile feedback in dexterous robotic hands," *arXiv preprint arXiv:2504.16649*, 2025.
- [26] W. K. Do, B. Jurewicz, and M. Kennedy, "Densetact 2.0: Optical tactile sensor for shape and force reconstruction," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 12 549–12 555.
- [27] Y. Wu, Y. Chen, Z. Zhu, X. Qin, and C. Xiao, "Humanft: A human-like fingertip multimodal visuo-tactile sensor," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 244–250.
- [28] J. Lu, Z. Wan, and Y. Zhang, "Tac2structure: Object surface reconstruction only through multi times touch," *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1391–1398, 2023.
- [29] H.-J. Huang, M. A. Mirzaee, M. Kaess, and W. Yuan, "Gelslam: A real-time, high-fidelity, and robust 3d tactile slam system," *arXiv preprint arXiv:2508.15990*, 2025.
- [30] M. Bauza, A. Bronars, and A. Rodriguez, "Tac2pose: Tactile object pose estimation from the first touch," *The International Journal of Robotics Research*, vol. 42, no. 13, pp. 1185–1209, 2023.
- [31] S. Suresh, H. Qi, T. Wu, T. Fan, L. Pineda, M. Lambeta, J. Malik, M. Kalakrishnan, R. Calandra, M. Kaess, J. Ortiz, and M. Mukadam, "Neural feels with neural fields: Visuo-tactile perception for in-hand manipulation," *Science Robotics*, p. adl0628, 2024.
- [32] R. Li and E. H. Adelson, "Sensing and recognizing surface textures using a gelsight sensor," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 1241–1247.
- [33] C. Lin, Z. Lin, S. Wang, and H. Xu, "Dtact: A vision-based tactile sensor that measures high-resolution 3d geometry directly from darkness," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 10 359–10 366.
- [34] J. Ji, Y. Liu, and H. Ma, "Model-based 3d contact geometry perception for visual tactile sensor," *Sensors*, vol. 22, no. 17, p. 6470, 2022.
- [35] R. Prados and O. Faugeras, "Perspective Shape from Shading" and Viscosity Solutions," in *Proceedings Ninth IEEE International Conference on Computer Vision*, 2003, pp. 826–831 vol.2.
- [36] D. F. Gomes, Z. Lin, and S. Luo, "Geltip: A finger-shaped optical tactile sensor for robotic manipulation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9903–9909.
- [37] G. Vitrani, B. Pasquale, and M. Wiertelwski, "Shadowtact: Dense measurement of shear and normal deformation of a tactile membrane from colored shadows," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 5004–5010.