

RainbowSight: A Family of Generalizable, Curved, Camera-Based Tactile Sensors For Shape Reconstruction

Megha H. Tippur and Edward H. Adelson
Massachusetts Institute of Technology

mhtippur@csail.mit.edu, adelson@csail.mit.edu

Abstract—Camera-based tactile sensors can provide high resolution positional and local geometry information for robotic manipulation. Curved and rounded fingers are often advantageous, but it can be difficult to derive illumination systems that work well within curved geometries. To address this issue, we introduce RainbowSight, a family of curved, compact, camera-based tactile sensors which use addressable RGB LEDs illuminated in a novel rainbow spectrum pattern. In addition to being able to scale the illumination scheme to different sensor sizes and shapes to fit on a variety of end effector configurations, the sensors can be easily manufactured and require minimal optical tuning to obtain high resolution depth reconstructions of an object deforming the sensor’s soft elastomer surface. Additionally, we show the advantages of our new hardware design and improvements in calibration methods for accurate depth map generation when compared to alternative lighting methods commonly implemented in previous camera-based tactile sensors. With these advancements, we make the integration of tactile sensors more accessible to roboticists by allowing them the flexibility to easily customize, fabricate, and calibrate camera-based tactile sensors to best fit the needs of their robotic systems.

I. INTRODUCTION

As roboticists aim to accomplish a wider range of more complex dexterous manipulation tasks, tactile sensors have become a popular choice to improve a robotic end effector’s ability to intelligently interact with its environment. Much like humans, manipulators can greatly benefit from the positional, geometric, and textural information these types of sensors provide. In cluttered or highly dynamic scenes, occlusions to the external, global vision system can occur by other objects or even the manipulator itself. In these situations, it is important for the end effector to be able to obtain local sensory information to complete its task. In some cases, no visual feedback is even available, such as when rummaging in a bag or on the top shelf of a cabinet. Here, the robot must rely solely on its tactile feedback to succeed in its exploration. A variety of transduction mechanisms can be used to provide tactile feedback, including resistive [1]–[3], piezoelectric [4]–[6], capacitive [7]–[9], barometric [10], [11], and optical [12]–[15]. In addition, the diversity in size, material composition, and form factor of tactile sensors makes it possible to incorporate them onto different end effector configurations. In particular, camera-based tactile sensors, such as GelSight, have been widely adopted due to their ability to provide high resolution tactile information, such as contact position, object geometries, and applied forces. However, even though much work has been done to




Fig. 1. Two omnidirectional RainbowSight fingers mounted on a parallel-jaw gripper holding an M7 screw. (A) Rainbow illuminated addressable RGB LED ring mounted at the base of the sensors. To the right is an example of a scaled-down version of the sensor with a diameter similar to that of a dime (~ 20 mm). (B) Example difference image captured by camera showing the deformation of the elastomer coating. (C) Depth reconstruction of the sensor surface viewed from sensor base. (D) Point cloud of sensor surface viewed from camera base.

obtain accurate contact localization and depth reconstructions with GelSight and GelSight-based sensors, most employ a flat sensing surface with a bulky form factor to obtain these results [16]–[18]. Though a flat or curved half-sensing surface may be sufficient for parallel-jaw grippers, end effectors of varying sizes, finger configurations, and DOFs attempting to perform more complicated tasks may not get the most benefit out of camera-based tactile sensors with these shapes. In addition, previous attempts at building curved sensors based on GelSight principles required extensive hardware testing to ensure a proper illumination throughout the entire surface of the sensor, making it difficult to change the shape and size of the sensor. We build off our previous work in [19] to introduce RainbowSight, a family of curved, omnidirectional and half-configuration sensors which, to our knowledge, is the first of its kind to utilize a novel rainbow illumination strategy. This strategy allows the user to customize their tactile sensors to fit the size and shape requirements of their specific end effector. Specifically, we show the following contributions and improvements to our

previous curved, omnidirectional sensor design:

- A novel rainbow illumination scheme using a semi-specular coating to produce a blended, color gradient desired for photometric stereo techniques. Unlike in [19], there are no occlusions of the sensor’s surface.
- A simplified fabrication procedure utilizing the rainbow illumination scheme that allows customization of shape and size, while maintaining the quality of the final tactile image.
- Improved depth reconstructions of the contact deformations occurring on the sensor’s skin.

II. RELATED WORKS

Camera-based tactile sensors have proven to be particularly beneficial in many types of manipulation tasks due to their ability to provide high-resolution tactile information by interpreting images of the sensor’s surface deformation when in contact with an object. Many of the current camera-based tactile sensors are based on principles introduced in GelSight, which uses a soft, gel elastomer covered in an opaque, reflective coating to convert geometric and pressure information about a contact region to image data [16], [20]. A camera housed behind the elastomer, along with directionally unique red, green, and blue illumination sources allows photometric stereo methods to be used to estimate the surface normals. These gradient estimates can ultimately be used with normal integration methods to produce 3D reconstructions of the sensor’s contact deformation [16], [17], [20]. Tactile information, such as depth maps or force sensing, have shown to be useful in many manipulation tasks requiring contact localization [14], [21], [22], slip detection [23], [24], and object geometry estimation [25], [26].

Much of the initial work on GelSight-based tactile sensors focused on using flat sensing surfaces to improve the spatial resolution, depth reconstruction quality, and force sensing for use in manipulation tasks [16]–[18]. More recently, many camera-based tactile sensors with curved, unconventional geometries have been developed to augment a variety of robotic gripper configurations [19], [27], [28]. Some curved omnidirectional sensors, such as [19] convert traditional GelSight processing methods to work on curved surfaces in order to produce depth maps of the deformed contact region. Other, half-configuration sensors, such as [29] use a hemispherical gel dome geometry and collect a variety of different probing shapes to directly translate an RGB difference image to a depth estimation. By adding a randomized pattern to the surface of the gel, the authors were also able to extract 6-axis F/T data from their sensor [29]. [30] uses a force sensitive spring system along with pad-printed markers on the gel’s surface to acquire multi-axis force data and contact geometry.

Though half sensor configurations provide sufficient sensing area for tasks using parallel-jaw grippers, more dexterous grippers may further benefit from all-around tactile sensors, similar to the curved sensing geometry of human fingertips [21], [31]–[33]. Due to their large, all-around sensing surface, omnidirectional sensors can sense multiple contact regions, speed up exploration time, prevent damaging

collisions, and perform complex in-hand manipulation tasks that would be difficult with flat or even one-sided curved sensors. Many camera-based, omnidirectional tactile sensors have been introduced in recent years. The OmniTact sensor introduced a multi-camera scheme in an elastomer-covered, finger-like shape that was able to estimate the contact angle of objects [34]. The GelTip sensor shared a similar shape to [34] and used structured red, green, and blue light with a camera mounted at the base of the internal plastic shell to observe deformations in the sensor skin. Insight, a soft, conically-shaped thumb size sensor, used information from internal texture and structured lighting to train a neural network to estimate contact localization and force estimates within sub-millimeter accuracy [35]. Later, a smaller, more fingertip-shaped version of the sensor, Minsight, was introduced and showed the sensor’s performance in a tactile servoing and lump classification task [32]. Similarly, the AllSight sensor is able to calculate extremely accurate position and F/T data using a red, blue, green lighting scheme and black markers painted on the elastomer’s surface [36]. Additionally, they demonstrate zero-shot transferability of their state estimation models to user-manufactured sensors [36]. Finally, GelSight360 used an LED crossing structure to recreate the unique directionality constraints of the colored lights in order to use photometric stereo techniques to provide high resolution depth reconstructions of the contact regions [19]. In addition, the generalizability of the cross LED illumination scheme was demonstrated by implementing it in a variety of omnidirectional finger shapes [19].

In this work, we build on many of the methods introduced in [19] with an alternate illumination scheme to propose a new family of curved tactile sensors of different sizes and shapes. In addition, we make improvements to the sensor hardware and image processing methods to produce more accurate depth reconstructions of deformed regions.

III. SENSOR DESIGN AND FABRICATION

A. Design Criteria

We focus on introducing a novel illumination system that can be generalized to a variety of curved sensor shapes, along with fitting the sensor and electronics into a compact form factor that can be modularly mounted to a variety of gripper configurations [19], [21]. We improve on GelSight360 by replacing the cross LED structure previously mounted in the body of the sensor with a single rainbow-illuminated addressable RGB LED ring housed at the bottom of the sensor. This prevents occlusions in the sensing surface without sacrificing the integrity of the tactile image. In turn, since the improved illumination scheme is heavily implemented in software rather than hardware, the process to fabricate the sensors is eased, and the sensor sizes and shapes can be scaled down and generalized to account for a variety of curved shapes (including both omnidirectional and half-sensor configurations). This section further describes the illumination system implemented and the fabrication methods used to build the sensors.




Fig. 2. Exploded view of sensor.

B. Sensor Illumination Strategy

Past flat GelSight and GelSlim sensors use red ($\sim 630\text{nm}$), green ($\sim 525\text{nm}$), and blue ($\sim 470\text{nm}$) light to illuminate a soft, clear gel elastomer coated with an opaque, matte, aluminum-based pigment to capture the gel’s deformation when in contact with an object [20]. Because the three colors illuminate the sensor from different directions, photometric stereo techniques can be used to linearly map the RGB intensities collected by a camera housed behind the gel to the surface gradients of the sensor using a look-up-table [16], [20]. These horizontal (G_x) and vertical (G_y) gradients can then be integrated using the Fast Poisson Solver to produce a 3D height map of the deformation of the gel [16]–[18].

When a semi-specular aluminum coating is used to paint the elastomer (as opposed to the matte, gray Lambertian paint used in [16]–[18]), a one-to-one mapping between the RGB intensities and surface reflectance function no longer exists, making it more difficult to estimate the surface gradients. Additionally, due to the reflective properties of the material, the transitions between red, green, and blue are extremely harsh, and do not produce the diffused, blended color gradient observed with Lambertian coated surfaces. However, because of the semi-specular material’s sensitivity to smaller changes in the normals, it can actually provide higher resolution signals than the Lambertian case [16].

We attempt to leverage the high-resolution characteristics of the semi-specular aluminum coating for our curved sensors without sacrificing the desired, smooth, rainbow color gradient observed in sensors using a gray Lambertian coating by utilizing addressable RGB LEDs in our sensors. By weighting the intensity of red, green, or blue light in each LED, the hue around the perimeter of the board appears as a rainbow spectrum. When an object is pressed into the surface of a curved sensor, this blended rainbow gradient, as in Figure 4, can be observed when comparing the contact and reference image (later referred to as the difference image). However, for omnidirectional sensors, only the top section of the sensor body is able to achieve this gradient due to the nature of the sensor geometry. Because the light source sits at the base, on the sides of the sensor, the different hues are unable to blend, and the rainbow LEDs act more

as colored point sources with the semi-specular coating. Recently, Wang et al. showed that it was possible to leverage deep learning methods to predict relatively accurate gradient estimates and depth reconstructions even when only two colored lights were shined parallel to each other; we employ similar methods to resolve our gradient estimation issues [17].

C. Sensor Fabrication

By using only a single rainbow LED ring to satisfy the photometric stereo illumination requirements for the sensor, we simplify its fabrication process. Many of the materials and methods used are similar to those in [19] but have been altered to account for improvements in sensor illumination, deformation sensitivity, and overall sensor rigidity. The sensor body consists of three components: 1) a thin, internal resin shell, 2) a transparent, gel elastomer, and 3) an opaque semi-specular coating on the elastomer’s surface.

Transparent Resin Shell: Because of the length and vertically-oriented nature of curved, omnidirectional sensors, an internal, rigid skeleton is often chosen to provide a rigid mounting base and to prevent the gel from delaminating or tearing off [32], [34], [36], [37]. For these reasons, we use a 1mm-thick epoxy resin shell as the internal support for the sensors. Smooth-On’s EpoxAcast690 resin is chosen since it produces a transparent, bubble-free shell that is resistant to yellowing caused by UV light over time (unlike other clear SLA resins). In addition, to further the mechanical resiliency of the sensor hardware, we cast the epoxy shell directly into the sensor mounting piece (printed on FormLabs Tough 2000 Resin), rather than gluing it into the base as in [19]. A two-part, gravity molding structure is used to create the epoxy shell. The mold negatives are printed with a FormLabs Form2 SLA printer. After polishing the negatives, Smooth-On MoldStar 20T is used to cast the silicone molds.

Painted Gel Elastomer: Once the epoxy is cured, the shell is removed from the mold using the top covering piece; this piece is re-used when casting the soft gel to ensure the resin shell and elastomer remain aligned on all axes. Unlike in [19], we chose to couple the base LED ring with the silicone gel rather than the epoxy to minimize ray refractions and reflections that might be caused by polishing imperfections or refractive index mismatches between the resin and clear elastomer. To ensure there is no air-interface between the LEDs and the elastomer, we press-fit the soldered LED ring into the PCB holder in the sensor mount to allow the silicone to cure around the LEDs. The resin shell is primed to promote adhesion to the elastomer. To create the semi-specular coating, a mixture (by weight) of 1 part Print-On Clear Silicone Ink Base (Raw Materials): 0.1 part Print-On Silicone Ink Catalyst: 0.2 part aluminum flake powder: 3 parts of NOVOCSS Gloss (Smooth-On) is prepared and sprayed into the mold. A similar procedure described in [19] is used to prepare the soft, clear elastomer.

Sensor Electronics: We design custom RGB LED PCBs in different ring sizes and shapes to fit the desired footprints of both omnidirectional and half-sensor shapes. Even though

RGB LED strips and rings are available off-the-shelf, they come in limited shapes and sizes and are too bulky for our applications. Additionally, we need as many LEDs as spatially possible on the ring in order to successfully mimic a continuous color gradient with the discretized LEDs. Depending on the size of the board being designed, we use either the NeoPixel Addressable 2020 (for larger sensors) or 1515 (for smaller sensors) RGB LEDs with an integrated driver chip (WS2812B or SK6805 respectively). A $0.1\mu\text{F}$ decoupling capacitor is used for every two RGB LEDs to prevent possible current fluctuations in the IC chip if there are voltage drops or surges. Anywhere between 21 – 28 LEDs are able to be fit on a single side of the various shaped boards.

To control the colors of the addressable LEDs, the ring is connected to an Adafruit™ Pro Trinket Microcontroller (5V Logic). The microcontroller is chosen due to its small, compact form factor, allowing it to fit in a variety of housing configurations, rather than being housed off the robot entirely since it is best practice to keep the length of the data wire below 1m to minimize possible signal reflections. In addition, when powered with only a microUSB cable, the Pro Trinket has enough onboard memory and current output capability to power up to 5 LED rings at the settings desired for RainbowSight sensors. A 300 ohm resistor is placed on the digital signal line to prevent voltage spikes which could cause communication issues. The FastLED library is chosen to tune the ratios of the red, green, and blue diodes in each of the LEDs since it provides full rainbow HSV color support and brightness controls on the Arduino platform.

Camera: Depending on the sensor size, we chose between two wide-angle cameras to view the entire sensing surface. For larger sensors, we use the 175° FOV 8MP Sony IMX219 camera, since the inner diameter of the resin shell is large enough to fit its lens. Smaller sensors utilize the 160° FOV 5MP OV5647 camera since it has a smaller form factor. Both cameras are run off a Raspberry Pi using a CSI connector, and the 640×480 images are streamed at 30 fps to a local desktop using mjpg-streamer.

IV. ALTERNATE SENSOR SHAPES

In addition to fabricating the sensor shape shown in Figure 1, we demonstrate the novel illumination scheme’s generalizability to different sensor geometries (both half-sided and omnidirectional) with the same manufacturing technique described in Section III. Figure 3 shows the sizes of the sensor, the board shape used, and the resulting images produced when a 4mm sphere and M7 screw (respectively) are pushed into each of the sensors. The shapes and approximate sizes of the half-sensor configurations are inspired by the differing geometries of the commercially available SynTouch BioTac sensors. In these shapes, even though the sensor is not radially symmetric with respect to the vertical axis, the desired blended, rainbow effect is still achieved even though a semi-specular coating is used. With the removal of the cross LED structure used in [19], fabricating half-sensor configurations and truly fingertip-sized sensors becomes possible. One

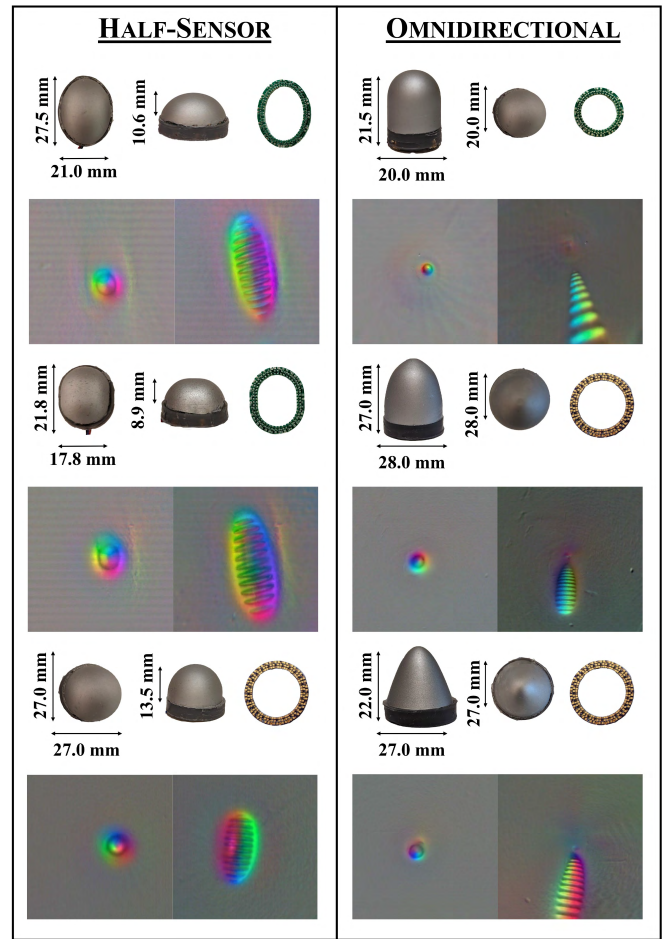


Fig. 3. Example difference images collected when a 4mm ball and M7 screw are pressed into the various sensor shapes and sizes. In the half-sensor configurations and tops of omnidirectional sensors, we are able to achieve the desired rainbow gradient illumination pattern over a large area of the main sensing surface. For omnidirectional sensors, the sensor sides provide adjacent two-color parallel light.

drawback of the cross-LED was that because the thickness of the PCBs and LEDs remained constant, as the diameter of the sensor body was reduced, a greater percentage of the sensing area would become severely occluded, limiting the advantages of an omnidirectional sensor. The camera’s perspective model also prevented the cross-LED structure to be properly implemented in half-sensor configurations, since it heavily occluded the main sensing surface due to how close it was to the camera lens. However, since the illumination is now housed at the base of the sensor, a wider range of shapes can be designed without sacrificing either sensing area visibility or the colored illumination’s directionality and blending.

V. MEASURING 3D GEOMETRY

Similar to calibration methods used for previous GelSight sensors in [16], [17], [19], a lookup table mapping the image’s RGB pixel intensities to the surface gradients of the sensor surface must be constructed. These gradients can then be integrated to reconstruct the deformations in the sensor’s surface. Unlike in [29], which uses an array of different shapes that must be probed into the sensor’s surface, our calibration method requires only a spherical probe, making

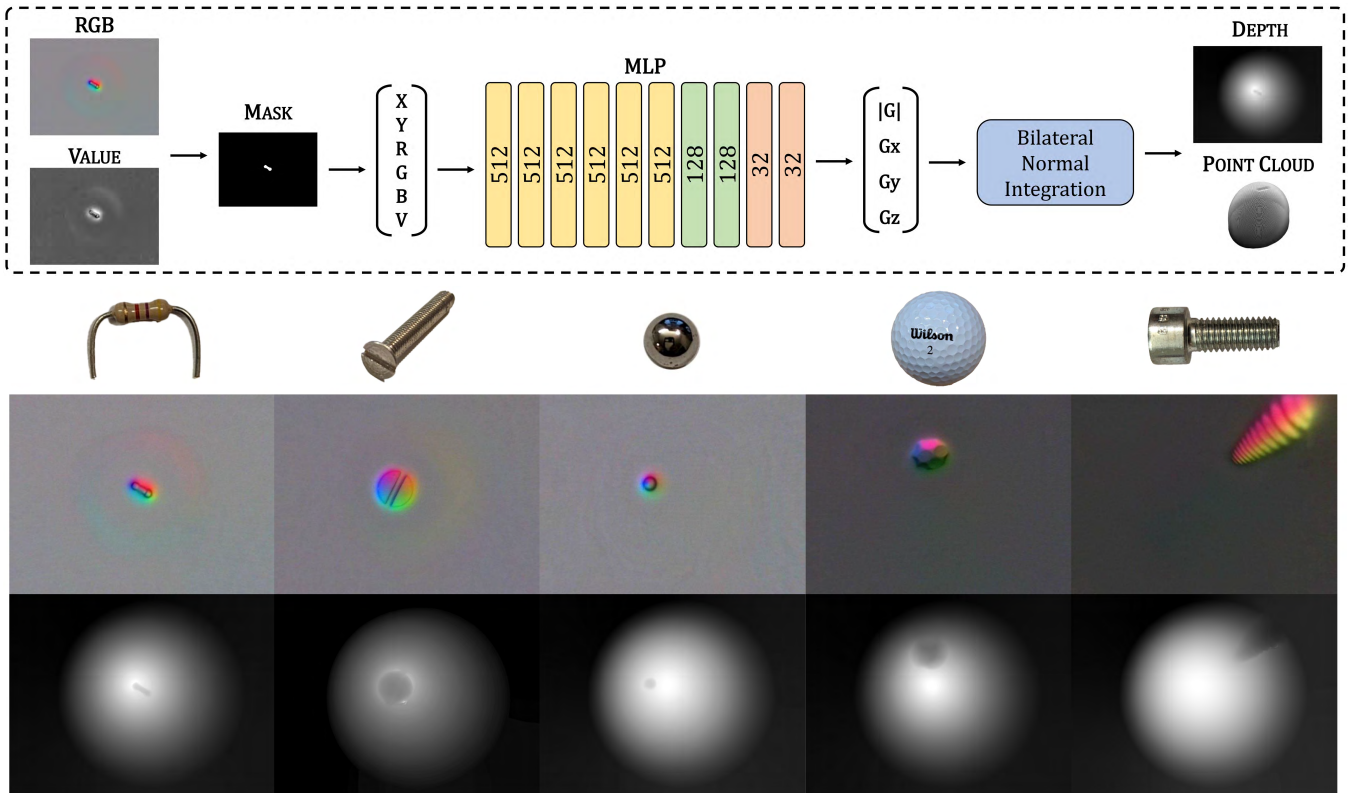


Fig. 4. **Top Dashed Box:** Depth reconstruction pipeline. **Bottom:** Tactile signals collected when different objects are pressed at various locations on the omnidirectional cylinder with hemisphere top sensor shape. **Top Row:** Objects pressed into the sensor surface. **Middle Row:** Tactile difference images of the contact regions. **Bottom Row:** Estimated depth map of the imprinted object in the sensor skin.

it easier for our calibration method to be generalized to the different shapes and sizes of the sensors tested in this work.

We use a similar calibration and data-processing technique as in [19], as most of the improved depth reconstruction results come from the illumination improvements, network architecture, and integration methods employed. The calibration data for the sensor is collected using a tabletop 3020 CNC since it can precisely probe the sensors at the desired positions. The surface mesh of the sensor is sampled at $N = 5000$ probing points. The CNC autonomously pushes the sphere 4mm into the soft, sensing surface and records images of the deformation.

Because of the variety of sensors possible with this new illumination technique, we aim to implement a universal calibration procedure to collect ground truth data when training the learning-based lookup table, rather than using specified ray-casting methods or coordinate conversions as in [29], [37], respectively. Before the sensor is fully assembled, the fisheye camera is focused to the correct focal length, and a checkerboard calibration is done using OpenCV’s Fisheye Calibration procedure to obtain the camera intrinsic and distortion parameters. With the world coordinates obtained from the CNC probing and the probed sensor images, PnP RANSAC is used to find the camera’s location. Using the camera’s intrinsic and extrinsic parameters, along with the model of the sensor, the gradients of the sensor’s surface and spherical contact from the CNC probing points can be rendered using graphics software.

The data is prepared for training by first undistorting

the images. A difference image is found by subtracting the reference image from the contact image. Thresholding and blob detection are used to mask the probed area, and the pixel coordinates (x, y) along with RGB intensities are collected. The difference image is also converted into the Hue Saturation Value (HSV) color space, and the masked Value layer is collected to create an $M \times 6$ data input for training, where M represents all of the masked pixels from the probed CNC points. The same mask is applied to the rendered normal map to collect the ground truth training data. The gradients are broken into their magnitudes and normalized X, Y, Z components, which is used as the $M \times 4$ ground truth data for training. A multilayer perceptron network (MLP) is trained using ReLU activation with early stopping on $M = 1,500,00$ points.

The 3D surfaces are reconstructed using normal integration techniques. In previous, flat GelSight sensors, the Fast Poisson Solver was used to integrate the gradients to produce the depth maps [16], [17]. However, in [19], we found that due to the curved nature of our sensor, the Fast Poisson Solver caused distortion in our depth maps due to the lack of discontinuity preservation in the perspective projection case. To address this issue, we utilize Cao et al.’s Bilateral Normal Integration method to generate the depth maps and resulting point clouds [38]. When integrating a small portion of the gradients, the method can run at about 14Hz.

Results: The depth maps generated when objects of different sizes and textures are pressed into the elastomer skin of the calibrated sensor at different locations are shown in

TABLE 1:
COMPARISON OF DIFFERENT ILLUMINATION STRATEGIES













		RGB IMAGE	DIFFERENCE IMAGE	RMSE LOSS	EPOCHS
RainbowSight				0.0529	120
Red-Green-Blue Lighting				0.0575	200
Green Lighting				0.0959	150
GelSight360				0.0872	85

Figure 4. The general shape and orientation of the object can be recognized with significantly less warping than the results shown in [19]; this can mostly be attributed to 1) a more accurate gradient estimation network and 2) the use of a different normal integration technique. Although the finer details, such as the threads on the screw or the individual textures of the golfball, are lost due to the inherent smoothing applied by this integration technique, the pose and contact location are preserved.

To further evaluate the advantage of using the rainbow LED illumination, we run our calibration procedure on other lighting strategies previously introduced and commonly used in other curved tactile sensors. We compare the tactile images and resulting RMSE testing losses for the normal estimation network in Table 1. This includes the common red-green-blue illumination pattern introduced in [20] and adopted in [29], [36], [37], a single-colored light source [39], and the cross illumination scheme of GelSight360 [19]. We probe the sensor bodies using the lighting schemes described, train the gradient estimation network (with early stopping) for each, and compare their RMSE testing losses. The sensor shape with the cylindrical base and hemisphere top is chosen for testing since the sensor geometry allows for the 2-directional parallel lighting discussed in III-B on the sides of the sensor, and the traditional color gradient effect on the hemispherical portion of the sensor. The hardware and illumination methods of RainbowSight help it outperform GelSight360, which is an improvement considering there are no longer any occluded areas in the sensor. Additionally, the rainbow illumination scheme offers slightly better performance over the traditional red-green-blue pattern in estimating the gradients and uses fewer epochs to train. However, even though there is only a slight improvement in the accuracy of the gradient prediction for this shape, the rainbow illumination system can still be more advantageous in alternate curved sensor geometries, especially those that are not radially symmetric around the vertical axis. For example, in Figure 5, when the red-green-blue illumination pattern is used in the elliptical sensor shape, the harsh color directionality characteristic of the semi-




Fig. 5. Comparison of the color gradients seen when illuminating a curved sensor shape that is not radially symmetric around the vertical axis. The lack of blending colors can be seen on the left when only red, green, and blue light illuminates the sensor. A more gradual gradation of colors can be seen in the rainbow illuminated sensor.

specular coating is observed. In contrast, when illuminated with the rainbow illumination system, a more gradual color gradient is observed, causing more variation in RGB image intensities across the surface of the sensor. Features such as this may be useful in improving training time and accuracy for unexplored shapes.

VI. CONCLUSION

In this work, we introduce a novel rainbow illumination method that can be implemented in a variety of shapes and sizes for curved, camera-based tactile sensors. Rather than illuminating the inside of the sensor with only red, green, and blue wavelengths, we employ an illuminated color gradient around the perimeter of the sensor body and photometric stereo techniques to produce 3D depth reconstructions of the deformation in the sensor surface when in contact with an object. This rainbow gradient is achieved by digitally varying the intensities of the red, green, and blue diodes in a ring of addressable RGB LEDs and is shown to produce satisfactory illumination for use with photometric stereo techniques with minimal hardware iterations or tuning. When compared to the previous cross-LED illumination strategy presented in [19], we show improvements in ease of manufacturing, depth reconstruction results, and overall design generalizability to a wider range of sensor configurations, now including both half-sensor and omnidirectional sensor shapes. Additionally, we are able to eliminate the occlusions to the sensing area (previously caused by the LED cross in the sensor body) without sacrificing the quality of the depth reconstructions by housing the rainbow illumination source only at the sensor base. By designing this family of curved, rainbow-illuminated sensors that can be easily manufactured in a variety of shapes, provide good tactile data with minimal optical hardware tuning, and be fitted on a variety of robotic end effectors, we aim to make tactile sensing more accessible to roboticists as they pursue more complex manipulation tasks.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship, Toyota Research Institute (TRI), and the Amazon Science Hub. The authors would also like to thank Sandra Q. Liu and Arpit Agarwal for helpful discussions on sensor improvements and depth reconstruction methods.

REFERENCES

- [1] S. Sundaram, P. Kellnhofer, Y. Li, J.-Y. Zhu, A. Torralba, and W. Matusik, "Learning the signatures of the human grasp using a scalable tactile glove," *Nature*, vol. 569, no. 7758, pp. 698–702, 2019.
- [2] J. A. Fishel and G. E. Loeb, "Sensing tactile microvibrations with the biotac—comparison with human sensitivity," in *2012 4th IEEE RAS & EMBS international conference on biomedical robotics and biomechatronics (BioRob)*. IEEE, 2012, pp. 1122–1127.
- [3] M. Ntagios, H. Nassar, A. Pullanchiyodan, W. T. Navaraj, and R. Dahiya, "Robotic hands with intrinsic tactile sensing via 3d printed soft pressure sensors," *Advanced Intelligent Systems*, vol. 2, no. 6, p. 1900080, 2020.
- [4] J. Zhang, H. Yao, J. Mo, S. Chen, Y. Xie, S. Ma, R. Chen, T. Luo, W. Ling, L. Qin, *et al.*, "Finger-inspired rigid-soft hybrid tactile sensor with superior sensitivity at high frequency," *Nature communications*, vol. 13, no. 1, pp. 1–9, 2022.
- [5] S. Yong, J. Chapman, and K. Aw, "Soft and flexible large-strain piezoresistive sensors: On implementing proprioception, object classification and curvature estimation systems in adaptive, human-like robot hands," *Sensors and Actuators A: Physical*, vol. 341, p. 113609, 2022.
- [6] Z. Tang, Z. Wang, J. Lu, and G. Ma, "Design of robot finger based on flexible tactile sensor," *International Journal of Advanced Robotic Systems*, vol. 16, no. 5, p. 1729881419879853, 2019.
- [7] N. Zainuddin, N. F. Anuar, A. L. Mansur, N. I. M. Fauzi, W. F. Hanim, and S. H. Herman, "Resistive-based sensor system for prosthetic fingers application," *Procedia Computer Science*, vol. 76, pp. 323–329, 2015.
- [8] B. Heyneman and M. R. Cutkosky, "Biologically inspired tactile classification of object-hand and object-world interactions," in *2012 IEEE International Conference on Robotics and Biomimetics (RO-BIO)*. IEEE, 2012, pp. 167–173.
- [9] S. Mühlbacher-Karrer, A. Gaschler, and H. Zangl, "Responsive fingers—capacitive sensing during object manipulation," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 4394–4401.
- [10] A. SaLoutos, E. Stanger-Jones, M. Guo, H. Kim, and S. Kim, "Design of a multimodal fingertip sensor for dynamic manipulation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 8017–8024.
- [11] L. Epstein, A. SaLoutos, D. Kim, and S. Kim, "Bi-modal hemispherical sensors for dynamic locomotion and manipulation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 7381–7381.
- [12] A. Alspach, K. Hashimoto, N. Kuppawamy, and R. Tedrake, "Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation," in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2019, pp. 597–604.
- [13] 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.
- [14] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, *et al.*, "Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [15] K. Shimonomura, "Tactile image sensors employing camera: A review," *Sensors*, vol. 19, no. 18, p. 3933, 2019.
- [16] 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.
- [17] S. Wang, Y. She, B. Romero, and E. Adelson, "Gelsight wedge: Measuring high-resolution 3d contact geometry with a compact robot finger," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6468–6475.
- [18] I. H. Taylor, S. Dong, and A. Rodriguez, "Gelslim 3.0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 10781–10787.
- [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] 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.
- [21] S. Pai, T. Chen, M. Tippur, E. Adelson, A. Gupta, and P. Agrawal, "Tactofind: A tactile only system for object retrieval," *arXiv preprint arXiv:2303.13482*, 2023.
- [22] A. Yamaguchi and C. G. Atkeson, "Recent progress in tactile sensing and sensors for robotic manipulation: can we turn tactile sensing into vision?" *Advanced Robotics*, vol. 33, no. 14, pp. 661–673, 2019.
- [23] J. Li, S. Dong, and E. Adelson, "Slip detection with combined tactile and visual information," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7772–7777.
- [24] W. Yuan, R. Li, M. A. Srinivasan, and E. H. Adelson, "Measurement of shear and slip with a gelsight tactile sensor," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 304–311.
- [25] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, "Cable manipulation with a tactile-reactive gripper," *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1385–1401, 2021.
- [26] N. Sunil, S. Wang, Y. She, E. Adelson, and A. R. Garcia, "Visuotactile affordances for cloth manipulation with local control," in *Conference on Robot Learning*. PMLR, 2023, pp. 1596–1606.
- [27] S. Q. Liu, L. Z. Yañez, and E. H. Adelson, "Gelsight endoflex: A soft endoskeleton hand with continuous high-resolution tactile sensing," in *2023 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2023, pp. 1–6.
- [28] R. Patel, R. Ouyang, B. Romero, and E. Adelson, "Digger finger: Gelsight tactile sensor for object identification inside granular media," in *Experimental Robotics: The 17th International Symposium*. Springer, 2021, pp. 105–115.
- [29] 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. 12549–12555.
- [30] K. Althoefer, Y. Ling, W. Li, X. Qian, W. W. Lee, and P. Qi, "A miniaturised camera-based multi-modal tactile sensor," *arXiv preprint arXiv:2303.03093*, 2023.
- [31] M. A. Lin, R. Thomasson, G. Uribe, H. Choi, and M. R. Cutkosky, "Exploratory hand: Leveraging safe contact to facilitate manipulation in cluttered spaces," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5159–5166, 2021.
- [32] I. Andrussov, H. Sun, K. J. Kuchenbecker, and G. Martius, "Minsight: A fingertip-sized vision-based tactile sensor for robotic manipulation," *Advanced Intelligent Systems*, p. 2300042, 2023.
- [33] G. Khandate, S. Shang, E. T. Chang, T. L. Saidi, J. Adams, and M. Ciocarlie, "Sampling-based exploration for reinforcement learning of dexterous manipulation," *arXiv preprint arXiv:2303.03486*, 2023.
- [34] A. Padmanabha, F. Ebert, S. Tian, R. Calandra, C. Finn, and S. Levine, "Omniact: A multi-directional high-resolution touch sensor," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 618–624.
- [35] 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.
- [36] O. Azulay, N. Curtis, R. Sokolovsky, G. Levitski, D. Slomovik, G. Lilling, and A. Sintov, "Allsight: A low-cost and high-resolution round tactile sensor with zero-shot learning capability," *arXiv preprint arXiv:2307.02928*, 2023.
- [37] 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.
- [38] X. Cao, H. Santo, B. Shi, F. Okura, and Y. Matsushita, "Bilateral normal integration," in *European Conference on Computer Vision*. Springer, 2022, pp. 552–567.
- [39] 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)*. IEEE, 2023, pp. 10359–10366.