

GelSight Wedge: Measuring High-Resolution 3D Contact Geometry with a Compact Robot Finger

Shaoxiong Wang, Yu She, Branden Romero, Edward Adelson
Massachusetts Institute of Technology

<wang_sx, yushe, brromero>@mit.edu, adelson@csail.mit.edu

<http://gelsight.csail.mit.edu/wedge/>

Abstract—Vision-based tactile sensors have the potential to provide important contact geometry to localize the objective with visual occlusion. However, it is challenging to measure high-resolution 3D contact geometry for a compact robot finger, to simultaneously meet optical and mechanical constraints. In this work, we present the GelSight Wedge sensor, which is optimized to have a compact shape for robot fingers, while achieving high-resolution 3D reconstruction. We evaluate the 3D reconstruction under different lighting configurations, and extend the method from 3 lights to 1 or 2 lights. We demonstrate the flexibility of the design by shrinking the sensor to the size of a human finger for fine manipulation tasks. We also show the effectiveness and potential of the reconstructed 3D geometry for pose tracking in the 3D space.

I. INTRODUCTION

During manipulation, robots inevitably occlude the objective from vision, for example, when buttoning a shirt, or tying a shoelace. The occlusion makes it challenging for control policies with vision alone. In contrast, humans also use touch for daily manipulation tasks.

We aim to design a compact robot finger with tactile sensing to provide high-resolution 3D contact geometry, which can be used to localize [1] and estimate the pose of the in-held object with visual occlusion [2, 3]. We leverage GelSight [4] technology in this work for its high-resolution spatial signals and 3D reconstruction.

Previous GelSight sensors [5] provide high-resolution 3D contact geometry. However it is too bulky for robot fingers. GelSlim [6] optimizes the sensor shape, making it more compact for robotic tasks. But it is not optically designed for high-resolution 3D reconstruction.

The major challenge is to meet optical requirements for photometric stereo, with mechanical constraints for a compact sensor. Optically, it requires to estimate gradients accurately from color signals, in different directions. Mechanically, it requires a thin and slim shape like human fingers, and having the sense of the tip to pick up small objects on the table.

In this work, we present the GelSight Wedge sensor, shown in Fig. 1, which meets the mechanical requirements for a compact robot finger, while maintaining high-resolution 3D reconstruction. We introduce the related works on GelSight sensors and corresponding 3D reconstruction algorithms in Section II. We then describe the design details of the sensor in Section III. We evaluate 3D reconstruction with different light configurations for future design reference in Section IV.




Fig. 1. (a) A gripper equipped with two GelSight Wedge sensors making contact with a plate on the table; (b) A close-up view of the bottom sensor squeezing through the gap between the plate and the table; (c) The tactile imprint provides high-resolution contact geometry and forces.

Finally, we combine the reconstructed 3D point cloud with Iterative Closest Point (ICP) [7] to demonstrate its effectiveness and potential for pose tracking in the 3D space.

II. RELATED WORK

GelSight sensors are one type of vision-based tactile sensors [8, 9, 10, 11, 12, 13, 14]. Typical vision-based tactile sensors obtain high-resolution contact information by watching the deformation of soft gels with cameras when in touch. Contact information includes forces and geometry. Tracking the motion of visual features on the gel provides contact forces [4, 8, 11]. In addition, GelSight sensors use reflective silicone paints and directional lights to get high-resolution 3D contact geometry applying photometric stereo [4].

Robots can use contact geometry and forces for various robotic tasks. Contact geometry has been applied for pose estimation [1, 3], material perception [15], grasp adjustment [16, 17], contour following [18, 19], and ball manipulation [20, 12]. Researchers also combine contact geometry with forces for slip detection [21], hardness measurement [22], cube manipulation [23], and cable manipulation [24].

In this work, we follow the path of GelSlim [6], to make a compact robot finger, but keep measuring high-resolution 3D contact geometry, to provide more accurate information for robotic applications.

For 3D reconstruction algorithms, the previous GelSight sensor [5] uses the photometric stereo technique. It maps the color to surface gradients with a lookup table, and integrates



Fig. 2. The compact Wedge sensor (top) can squeeze into narrow spaces for robotic tasks, compared with the previous GelSight sensor [5] (bottom).

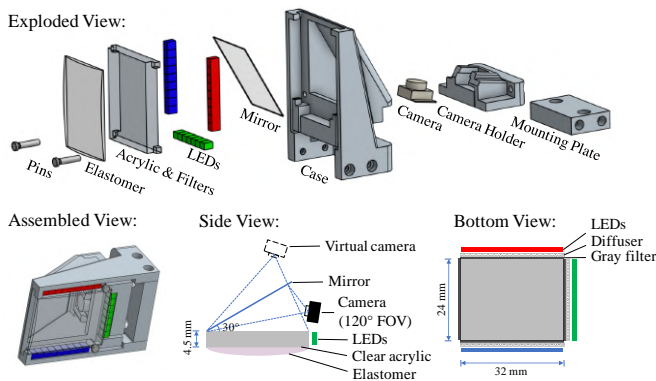


Fig. 3. Sensor design, including the exploded view showing inner components, the assembled CAD model, and the schematics.

the gradients with a fast Poisson solver [25] to get the depth. The mapping requires distinguishable color signals for different surface gradients, and the fast Poisson solver requires surface gradients in both horizontal and vertical directions. The current GelSlim design can not reach the two requirements, so the authors apply a neural network for 3D reconstruction [3]. Since it is a challenging task to reconstruct 3D from raw images, it requires thousands of calibration images to learn a model that can generalize to different shapes. In this work, we design the sensor with care to meet the two requirements for photometric stereo. We demonstrate that it achieves higher-fidelity 3D reconstruction with only 30 calibration images.

III. DESIGN AND FABRICATION

Our design goal is to have a compact shape for robotic tasks, as shown in Fig. 2, while maintaining high-resolution 3D reconstruction. To achieve this, we comprise some extent of optical properties for better shape. Fig. 3 shows components and schematics of the sensor design. Following, we describe the design principles and learned lessons for each sensor component.

Lighting The ideal lighting for GelSight 3D reconstruction would be 3 directional color lights evenly spread around 360 degrees, as [5, 26]. It can measure surface gradients from all directions.




Fig. 4. The effects of gray filters, diffusers, and curved gels. Gray filters increase the color contrast to better map to the surface gradients; Diffusers remove shadows and enable color gradients on the boundary (e.g. without diffusers, there is no blue lights for the bottom sphere); Curved gels produce more uniform global illumination.

We sacrifice the light at the tip to pick up small objects, which discard half of the horizontal gradients. To compensate that, we put lensless LED arrays directly on the 3 sides of the acrylic to generate directional lights while keep the sensor thin, and apply filters to estimate the gradients from other 3 directions as accurate as possible. We will demonstrate that a combination of these components can still reconstruct 3D with high-fidelity in Section IV. We use lensless LED with brand “Chanzone 3528” (3.5×2.8 mm) for our sensor.

Filters We find gray filters and diffusers can effectively increase the contrast and uniformity of tactile imprints. We stick the filters to the side of the acrylic block, as shown in Fig. 3, bottom view. The effects are shown in Fig. 4.

The gray filter needs to be optically coupled to the acrylic block so that there is no reflection from the air interface. One easy way is to use automobile window tint film, which comes with a layer of clear adhesive on one side. We attach the filter to the side of the acrylic block and press it carefully in place so there are no air bubbles. We use a brand called “VViViD Air-Tint Dark Air-Release Vinyl Wrap Film” for the sensor.

Gray filters can reduce internal reflections. When lights reflect inside the acrylic, it decreases the contrast of contact regions and the sensitivity for mapping from RGB color to the surface gradient. Gray filters can alleviate the problem. A gray filter absorbs light, which reduces the illumination level and also reduces internal reflections. The internal reflected light has to pass through the filter twice before returning, while the direct light from the LED’s has to pass through it only once. Thus if the gray filter passes 1/4 of the light, then the internal reflections will be reduced to 1/16.

Diffusers can lead to more uniform global illumination. Without diffusers, regions on the boundary are not illuminated, as shown in Fig. 4. Contact surfaces lose part of the gradient information inside these regions, causing errors in 3D reconstruction. Diffusers help to spread the light smoothly over the whole sensing area. We use “3M Diffuser 3635-70” for the sensor.

Camera We use a Raspberry Pi mini camera, with a high




Fig. 5. Comparison of different covering. The medical tape reserves more details, has lower friction forces, and medium-level of protection. It is more suitable for sliding motions for inspection and light-weight manipulation tasks. The fabric has higher friction forces, and is more rugged, but loses some extent of sensitivity. It is more suitable for heavy-weight grasping and manipulation tasks that do not require fine details. It would be ideal to switch “gloves” automatically for different tasks in the future.

FOV of 120 degrees, to make the sensor compact. The camera only costs \$16. We stream the video by `mjpg_streamer` on the Raspberry Pi, and parse the `mjpeg` video stream from python on the computer. It works at 60 FPS for 640×480 pixels, and 90 FPS for 320×240 pixels. The default `mjpg_streamer` setting can cause flickering images. To get stable and well-blend color images, it is important to turn off the AWB (auto white balance), and tune fixed gains for `awbgainR` (red channel), `awbgainB` (blue channel), and EV (exposure compensation for RGB channels).

Elastomer The clear silicone gel can be coated with gray (Lambertian) or silver (Semi-specular) paints. In this work, we focus on 3D reconstruction with Lambertian for it provides more accurate gradient information, as suggested in [5]. We use “Silicone Inc. XP-565” for clear silicone, and “Print-On Silicone Ink” for paints.

The shape of the silicone gel can be flat or curved. Optically, the curved gel provides more uniform illumination, as shown in Fig. 4. Because the curve make the lights shoot from tilted angles to more direct angles as distance grows. This compensates for the decreased LED lights over distance. Mechanically, flat gel can pick up smaller objects at the tip, while curved gel is more tolerant to flat and concave surfaces since it does not require the sensed surface to be perfectly aligned to the sensor surface. In our experiment, the shape of curved gel is a part of an ellipsoid ($a=b=60\text{mm}$, $c=100\text{mm}$).

Covering To improve the durability and reduce friction, the gel can be covered with “gloves”. Besides the fabric in GelSlim [6], we also find medical tapes helpful. Medical tapes are smooth and sensitive. Compared to fabric, the medical tape reserves much more contact details, as shown in Fig. 5. The medical tape has low friction coefficient, making it not suitable for heavy lifting application, but suitable for sliding motions for inspection or manipulation tasks. It provides medium-level protection compared to fabric. We use “3M Tagaderm” medical tapes for our experiments.

Assembly For preparation, we 3D printed the sensor case, camera holder, and mold (for curved gel) with PLA materials. We polish the mold by sand papers to be smooth. We then




Fig. 6. Workflow of 3D reconstruction. We first unwarped the raw imprint to rectangular shape; then subtract it with the blank background frame to get the difference image; finally calculate the 3D shape by mapping the difference image to surface gradients and applying fast Poisson solver for integration.

laser cut the clear acrylic block, mirror, filters, and diffusers.

The silicone gel is attached to the acrylic with Silpoxy or A-564 Medical Silicone Adhesive. Gray filters and diffusers are attached to the side of the acrylic. The acrylic can be either press-fit and fixed by pins or glued into the sensor case.

The camera is press-fit into the camera holder. The camera holder is attached to the sensor case with M2 screws. The LEDs are soldered in arrays and taped into the sides of the sensor case. The front-surface mirror is glued inside the case, which is tilted by 30 degrees.

To improve robustness, we follow GelSlim[6], and use a CSI-to-HDMI connector, which converts the fragile CSI ribbon cables to HDMI cables. The 3.3V from the connector can be used to supply the LEDs. We use an external PCB to connect resistors to the LEDs, to save space.

IV. MEASURING 3D GEOMETRY

In this section, we discuss the performance of 3D reconstruction with GelSight Wedge sensor. We first describe the approach for photometric stereo, and demonstrate the results with 3 lights in Section IV-A. We then extend the method to designs with 1 or 2 lights, and compare the results with different light configurations in Section IV-B. We finally show the effect of disturbance in Section IV-C, and potential capabilities in Section IV-D.

A. 3D reconstruction with 3 lights

As described in [4], GelSights 3D reconstruction consists of two steps: 1) calculating the mapping from color (RGB) to horizontal and vertical surface gradients (G_x , G_y); 2) applying fast Poisson solver [25] to integrate gradients and get the depth.

Color-to-gradient mapping To get the data for color-to-gradient mapping, we 1) press a calibration ball, with a known radius (2.4 mm in our case), on the sensor; 2) calculate the center and the radius of the contact circle in pixels, with Hough Circle Transform and human adjustment; 3) press a caliper on the sensor to calculate the number of pixels per millimeter. With these parameters, we can calculate the gradient G_x , G_y for each pixel in contact, and gather the pixel’s color and position: R , G , B , X , Y . Because the LED diminishes when travelling, we add the pixel location, X and Y , to compensate the light attenuation across the sensor.

To learn the mapping from color (RGBXY) to gradient (G_x , G_y), we simplify the lookup table [5] to a small neural




Fig. 7. From top row to bottom: visual images, GelSight imprints, and inferred depth of a ball array, a Lego block, a rubber tyre, a screw cap, a M2 screw.

network, a multi-layer perceptron (MLP) with 3 hidden layers (5-32-32-32-2), with tanh activation. Previous lookup table requires discretization to optimize efficiency and it is difficult to incorporate position. In contrast, a small neural network produces continuous results and can consider color changes across the sensor.

We collect 32 images for training, and 8 images for testing. Besides the pixels in contact, we add 5% of pixels not in contact (whose gradients should be zeros) to balance the data.

3D reconstruction with fast Poisson solver After getting the gradients G_x , G_y , we use a 2D fast Poisson solver to calculate the depth. We use a python version[25]. It takes in G_x , G_y , and boundary conditions and returns the depth. We set the boundaries to zeros for best approximation.

Fig. 7 shows some qualitative results of reconstructed 3D of different objects: a ball array, a Lego block, a rubber tyre, a screw cap, and a M2 screw. Compared to previous work on GelSlim [3], we can reconstruct the surface with more fidelity. Following, we analyze and compare the 3D reconstruction results with different light configurations.

B. 3D reconstruction with 1 or 2 lights

It can be difficult to place all 3 lights for some devices, e.g. GelFlex [27], where one axis of the sensor is too long and deformable for directional lights to travel. We aim to still estimate 3D geometries with the limited lights.

Perpendicular lights The perpendicular lights, such as RG, provide gradients for both axes, allowing us to keep using fast Poisson solver. Although it misses half of gradients for each of horizontal and vertical axes, the fast Poisson solver can still integrate the gradients reasonably. It generates comparable results to RGB lights qualitatively as shown in Fig. 9, and quantitatively as shown in Table I.

However, if we miss the light completely from one axis, it becomes more challenging to 3D reconstruct. Following, we will describe the challenges and possible solutions.

Challenge The fast Poisson solver requires the gradient of both directions, G_x , and G_y . When missing lights in one axis, we lose the corresponding gradient information. Without losing generality, assume we lose the green lights for our sensor, which corresponds to horizontal gradients G_x .

A naive approximation would be assuming the gradient G_x is all zeros, and sending it to fast Poisson solver. As shown in Fig. 8, this causes the reconstructed 3D to be “flattened” in the x-axis, and generates negative depth to meet the constraints.

Alternatively, we can leverage neural networks (NN). One way would be directly mapping from raw images to depth, as described in [3]. However, 3D reconstruction from raw images is complex, which involves both learning the color-to-gradient mapping, and solving partial derivative equations. As a result, it requires many tactile imprints pressing on various shapes to generalize to different shapes and sensors.

We hope to find a more data-efficient method that uses only a handful of calibration data. We approach this by combining NN with fast Poisson solver and Sim2real techniques.

Gradient estimation for the missing axis The fast Poisson solver generates inaccurate results when receiving horizontal gradients G_x as zeros, as shown in Fig. 8. But we have the vertical gradients G_y . We can improve G_x estimation with the information of G_y .

Inferring G_x from G_y is a generic problem independent of real or synthesized data. Therefore, we can use synthesized data to reduce the required real data. We synthesized 10k random depth images, and calculated gradients (G_y as input, G_x as output) for training a neural network model. We then send the inferred G_x along with G_y to Fast Poisson Solver. As shown in Fig. 8, it significantly improves the 3D reconstruction with the inferred G_x , without any additional real data.

We use U-Net [28] as our neural network model, which has been widely used for semantic segmentation. U-Net follows the standard encoder-decoder framework, and adds multi-scale skip passes to provide more details.

For the synthesis process, we generate random surfaces to

TABLE I
GRADIENT ESTIMATION ERRORS WITH DIFFERENT LIGHT CONFIGURATIONS

	R (w/o NN)	R (w/ NN)	RB (w/o NN)	RB (w/ NN)	RG (w/o NN)	RGB (w/o NN)
G_x error	0.104	0.055	0.106	0.051	0.045	0.043
G_y error	0.055	0.055	0.045	0.045	0.044	0.042
θ_x error (degs)	5.598	2.899	5.706	2.656	2.331	2.248
θ_y error (degs)	2.889	2.897	2.349	2.351	2.293	2.187




Fig. 8. 3D reconstruction with two lights from opposite directions (RB). Without the horizontal gradients, fast Poisson solver reconstruct depth with horizontal blur and negative value. With NN-estimated horizontal gradients, it significantly improve the 3D reconstruction results, which are close to the results of RGB lights, without extra real data. Depth images are added with a constant value in this figure to show the negative depth for RB (Poisson).

learn a generic model. We flatten the boundaries and remove negative depth to match the real data. For details, we generate random arrays ranging from -1 to 1 and cube the values for magnification; set boundaries to zeros; apply Gaussian filters with random kernel sizes; clip the negative depth; and scale to random maximal height within the range of real sensors. We find it important to set boundaries to zeros to match the real sensors and provide constraints. Otherwise, it can generate artifacts like small bumps between contact regions to the boundary.

Results We compare the 3D reconstruction results with different methods and different light configurations.

Fig. 8 shows the effectiveness of applying neural networks for estimating the gradients G_x . Without neural networks, the fast Poisson solver can only take G_x as zeros, and generates depth with horizontal blur and negative values. With neural networks to estimate missing gradients, it significantly improves the 3D reconstruction, without any extra real data.

Fig. 9 shows the qualitative comparison of 3D reconstruction with different light configurations. We find that the RG can perform surprisingly well with the fast Poisson solver alone. The half gradients of G_x and G_y already provide many constraints for the fast Poisson solver. In contrast, the RB requires neural networks to estimate the G_x from G_y to achieve comparable results. It is more challenging to 3D reconstruct




Fig. 9. 3D reconstruction with different light configurations. With the perpendicular lights (e.g., RGB and RG), Poisson alone can 3D reconstruct with high fidelity. Without the perpendicular lights (e.g., RB and R), neural networks can be applied to estimate the horizontal gradients and improve the reconstruction results.

with only 1 light, it tends to generate artifacts like small bumps between the contact region to the boundary due to too many missing gradients. We hope to refine the results for 1 light with a small amount of real data in the future.

Table I shows the quantitative results for gradient estimation with different light configurations. We use the collected 40 real images with a calibration ball, and split 32 for training, 8 for testing. The UNet uses 10k synthesized data with the described process. The results agree with the qualitative evaluation. In summary, the UNet can significantly improve the gradient estimation for the missing axis without extra real data. The preference order for light configurations is: RGB (3 lights) > RG (2 perpendicular lights) > RB (2 opposing lights) > R (1 light).

C. Influence of disturbances

Markers We can add markers on the gel to observe force distributions. However, the marker will influence the 3D reconstruction as shown in Fig. 10. Without extra processing, the markers result in small bumps in 3D. We can alleviate it by filling zeros gradients to the markers. It removes bumps and runs fast, but generates flat patches. We can improve the result by interpolating the gradients for the markers. In




Fig. 10. Comparison of interpolation methods for 3D reconstruction with markers.




Fig. 11. The gel deformation and shadows can influence the 3D reconstruction. We press a screw against the sensor with different forces. We can observe that the reconstructed shape does not exactly follow the object shape. It requires further processing from the perception side in the future.

our experiments, we use *griddata* in the *scipy* package [29] for interpolation. To run in real time, we interpolate with the ‘nearest’ method, and only sample the pixels around the markers. It generates reasonable interpolation results and only take 10 ms (for resolution of 200x150), allowing real-time feedback control. In contrast, ‘linear’/‘cubic’ generate better results but take 60/70 ms.

Gel deformation and shadows Fig. 11 shows the gel deformation and shadows can influence the 3D reconstruction, especially for sharp surfaces. It bring challenges to extract the surface of the object from reconstructed 3D shape, but also make the signals force-dependent.

D. Potential capabilities

Flexibility We demonstrate the design’s flexibility by shrinking it to the size of a human finger, while retaining high-resolution 3D reconstruction, as shown in Fig. 12. It is more suitable for fine manipulation tasks, and multi-finger hands in the future.

Pose tracking The reconstructed 3D point cloud makes it convenient to apply ICP for pose tracking [2, 9]. We demonstrate its effectiveness and potential by an example of tracking the pose of a cube by its corner. Fig. 13 shows that it accurately tracks the poses with various orientations. We assume the initial pose is given. We use the ICP algorithm from Open3D [30]. It runs at about 10 Hz. Detailed analysis of pose error and complex geometry is out of the scope of this paper.

V. CONCLUSION

In this work, we present the GelSight Wedge sensor, which can measure high-resolution 3D contact geometry in a compact




Fig. 12. Shrinking the sensor to the size of human fingers. Left shows the comparison with human fingers. Right shows the visual image, the GelSight difference image, and reconstructed 3D of touching a screw. The finger-size sensor can be applied for fine manipulation tasks, and multi-finger hands.




Fig. 13. Tracking the pose of a cube by corners with ICP and reconstructed 3D point clouds from touch alone.

shape. We illustrate the effect of the components in our design, and evaluate the 3D reconstruction with different light configurations. For applications, we successfully shrink the sensor to the size of a human finger to work for narrow environments and fine manipulation tasks in the future. We also demonstrate the effectiveness and potential of the reconstructed 3D point cloud for pose tracking.

In the future, we hope to further reduce the influence of markers on 3D reconstruction. It can be achieved by reducing the size of markers, or changing black dots to checkerboard of lighter and darker gray. The acrylic currently is 4.5mm thick, which is decided by the LEDs size. Replacing LED arrays with smaller sizes can further reduce the thickness of the finger tip. We hope to combine the sensor with robotic grippers and multi-finger hands towards more reactive and robust control for various grasping and manipulation tasks in the future.

ACKNOWLEDGEMENTS

Toyota Research Institute (TRI), and the Office of Naval Research (ONR) [N00014-18-1-2815] provided funds to support this work. The authors would also like to thank Achu Wilson, Alex Alspach for the help during prototyping, and Alberto Rodriguez, Siyuan Dong, Neha Sunil for insightful discussions.

REFERENCES

- [1] R. Li, R. Platt, W. Yuan, A. ten Pas, N. Roscup, M. A. Srinivasan, and E. Adelson, "Localization and manipulation of small parts using gelsight tactile sensing," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 3988–3993. 1
- [2] G. Izatt, G. Mirano, E. Adelson, and R. Tedrake, "Tracking objects with point clouds from vision and touch," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4000–4007. 1, 6
- [3] M. Bauza, O. Canal, and A. Rodriguez, "Tactile mapping and localization from high-resolution tactile imprints," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3811–3817. 1, 2, 4
- [4] 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. 1, 3
- [5] S. Dong, W. Yuan, and E. H. Adelson, "Improved gelsight tactile sensor for measuring geometry and slip," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 137–144. 1, 2, 3
- [6] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez, "Gelslim: A high-resolution, compact, robust, and calibrated tactile-sensing finger," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1927–1934. 1, 3
- [7] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611. International Society for Optics and Photonics, 1992, pp. 586–606. 1
- [8] 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. 1
- [9] A. Alspach, K. Hashimoto, N. Kuppuswamy, 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. 1, 6
- [10] C. Sferrazza, T. Bi, and R. D'Andrea, "Learning the sense of touch in simulation: a sim-to-real strategy for vision-based tactile sensing," *arXiv preprint arXiv:2003.02640*, 2020. 1
- [11] A. Yamaguchi and C. G. Atkeson, "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, 2016, pp. 1045–1051. 1
- [12] 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. 1
- [13] A. Padmanabha, F. Ebert, S. Tian, R. Calandra, C. Finn, and S. Levine, "Omniact: A multi-directional high resolution touch sensor," *arXiv preprint arXiv:2003.06965*, 2020. 1
- [14] B. Romero, F. Veiga, and E. Adelson, "Soft, round, high resolution tactile fingertip sensors for dexterous robotic manipulation," *arXiv preprint arXiv:2005.09068*, 2020. 1
- [15] W. Yuan, Y. Mo, S. Wang, and E. H. Adelson, "Active clothing material perception using tactile sensing and deep learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–8. 1
- [16] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine, "More than a feeling: Learning to grasp and regrasp using vision and touch," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3300–3307, 2018. 1
- [17] F. R. Hogan, M. Bauza, O. Canal, E. Donlon, and A. Rodriguez, "Tactile regrasp: Grasp adjustments via simulated tactile transformations," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 2963–2970. 1
- [18] N. F. Lepora, A. Church, C. De Kerckhove, R. Hadsell, and J. Lloyd, "From pixels to percepts: Highly robust edge perception and contour following using deep learning and an optical biomimetic tactile sensor," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2101–2107, 2019. 1
- [19] R. B. Hellman, C. Tekin, M. van der Schaar, and V. J. Santos, "Functional contour-following via haptic perception and reinforcement learning," *IEEE transactions on haptics*, vol. 11, no. 1, pp. 61–72, 2017. 1
- [20] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine, "Manipulation by feel: Touch-based control with deep predictive models," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 818–824. 1
- [21] S. Dong, D. Ma, E. Donlon, and A. Rodriguez, "Maintaining grasps within slipping bounds by monitoring incipient slip," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3818–3824. 1
- [22] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson, "Shape-independent hardness estimation using deep learning and a gelsight tactile sensor," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 951–958. 1
- [23] F. R. Hogan, J. Ballester, S. Dong, and A. Rodriguez, "Tactile dexterity: Manipulation primitives with tactile feedback," *arXiv preprint arXiv:2002.03236*, 2020. 1
- [24] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, "Cable manipulation with a tactile-reactive gripper," *arXiv preprint arXiv:1910.02860*, 2019. 1
- [25] J. Doerner, "Fast poisson reconstruction in python," <https://gist.github.com/jackdoerner/b9b5e62a4c3893c76e4c>,

2014. [2](#), [3](#), [4](#)

- [26] M. K. Johnson, F. Cole, A. Raj, and E. H. Adelson, “Microgeometry capture using an elastomeric sensor,” *ACM Transactions on Graphics (TOG)*, vol. 30, no. 4, pp. 1–8, 2011. [2](#)
- [27] Y. She, S. Q. Liu, P. Yu, and E. Adelson, “Exoskeleton-covered soft finger with vision-based proprioception and exteroception,” *arXiv preprint arXiv:1910.01287*, 2019. [4](#)
- [28] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241. [4](#)
- [29] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020. [6](#)
- [30] Q.-Y. Zhou, J. Park, and V. Koltun, “Open3D: A modern library for 3D data processing,” *arXiv:1801.09847*, 2018. [6](#)