



Lepora, N. F., Church, A., De Kerckhove, C., Hadsell, R., & Lloyd, J. (2019). 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*, 4(2), 2101-2107. [8641397]. <https://doi.org/10.1109/LRA.2019.2899192>

Peer reviewed version

License (if available):
Other

Link to published version (if available):
[10.1109/LRA.2019.2899192](https://doi.org/10.1109/LRA.2019.2899192)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). The final published version (version of record) is available online via IEEE at <https://doi.org/10.1109/LRA.2019.2899192>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available: <http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

From pixels to percepts: Highly robust edge perception and contour following using deep learning and an optical biomimetic tactile sensor

Nathan F. Lepora*, Alex Church, Conrad De Kerckhove, Raia Hadsell, John Lloyd*

Abstract—Deep learning has the potential to have the impact on robot touch that it has had on robot vision. Optical tactile sensors act as a bridge between the subjects by allowing techniques from vision to be applied to touch. In this paper, we apply deep learning to an optical biomimetic tactile sensor, the TacTip, which images an array of papillae (pins) inside its sensing surface analogous to structures within human skin. Our main result is that the application of a deep CNN can give reliable edge perception and thus a robust policy for planning contact points to move around object contours. Robustness is demonstrated over several irregular and compliant objects with both tapping and continuous sliding, using a model trained only by tapping onto a disk. These results relied on using techniques to encourage generalization to tasks beyond which the model was trained. We expect this is a generic problem in practical applications of tactile sensing that deep learning will solve.

Index Terms—Force and Tactile Sensing; Biomimetics; Deep Learning in Robotics and Automation

I. INTRODUCTION

ROBOT touch differs from robot vision: to touch, an agent must physically interact with its environment, which constrains the form and function of its tactile sensors (for example, to be robust, compliant and compact). Likewise, tactile perception differs from visual perception: to perceive touch, an agent interprets the deformation of its sensing surface, which depends on the shape and mechanics of the sensor, unlike vision where the eye does change what can potentially be seen [1].

Therefore, the application of deep learning to robot touch will be different from robot vision, just as robot vision poses different research questions than computer vision [2]. Thus far, there have been relatively few studies of deep learning for tactile perception compared with the explosion of work on robot vision. Those studies have mainly considered one particular device – the Gelsight [3], an optical tactile sensor that images the shading from 3 internal RGB-colored LEDs to transduce surface deformation (and complements this with

Manuscript received: September 10, 2018; Revised December 6, 2018; Accepted January 23, 2019.

This paper was recommended for publication by Editor Dan Popa upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by an award from the Leverhulme Trust on 'A biomimetic forebrain for robot touch' (RL-2016-39). * NL and JL contributed equally to this work.

¹NL, AC, CDK and JL are with the Department of Engineering Mathematics and Bristol Robotics Laboratory, University of Bristol, Bristol, U.K. {n.lepora, ac14293, cd14192, j115313}@bristol.ac.uk

²RH is with Google DeepMind. raia@google.com

Digital Object Identifier (DOI): see top of this page.




Figure 1: Tactile robotic system, showing: (a) the arm-mounted tactile sensor next to a few objects; (b) the (TacTip) optical tactile sensor; and (c) a labelled image from the tactile sensor.

painting markers to detect shear [4]). Use of an optical tactile sensor seems an appropriate starting place for applying deep learning to touch, given the rapid progress for vision.

In this paper, we apply deep learning to an optical biomimetic tactile sensor, the TacTip [5], [6], which images an array of 3D pins inside its sensing surface (Fig. 1c). These pins mimic the papillae that protrude across the dermal/epidermal layers of tactile skin, within which mechanoreceptors sense their displacement [7]. In past work, we used image-processing to track the pin positions [8], [9]. These were then passed to a distinct perception system, using statistical models [6]–[12], dimensionality reduction [13] or support vector machines [14]. Here we use a deep Convolutional Neural Network (CNN) for direct (end-to-end) perception from the tactile images.

Our main result is that application of a deep CNN can give robust edge perception and contour following in tasks beyond which the model was trained. We define a policy for planning contacts that moves the sensor while maintaining its pose relative to the edge. We test the robustness of this policy by whether the robot can successfully complete closed contours: first, with a tapping motion around a disk under variations in contact trajectory; second, around irregular objects differing in

local shape (Fig. 1); and third, for continual sliding motion. In all cases, the model was trained by tapping on a region of the disk. The policy met these tests, struggling only when sliding around sharp corners.

II. BACKGROUND AND RELATED WORK

Initial applications of deep learning to artificial tactile perception were with taxel-based sensors. The first used a four-digit robot hand covered with 241 distributed tactile skin sensors (with another 71 motor angles, currents and force/torque readings) to recognize 20 grasped objects, using a CNN and ~ 1000 training samples [15]. There have since been several other studies with taxel-based sensors [16]–[19].

More recently, the Gelsight optical tactile sensor has found a natural match with deep learning, beginning with applications of CNNs to shape-independent hardness perception [20] and grasp stability [21], then considering visuo-tactile tasks for surface texture perception [22] and slip detection [23]. The original Gelsight domed form-factor has been modified to a slimmer design for better integration with a two-digit gripper [24], enabling study of tactile grasp readjustment [25] and visuo-tactile grasping [26]. The majority of these studies use a CNN trained with several thousand examples, sometimes including an LSTM layer for sequential information.

In this work, we consider tactile edge perception and contour following. In humans, edges and vertices are the most salient local features of 3D-shape [27], [28], and are thus a key sensation for artificial tactile systems such as robot hands or prosthetics [29]. Work in robotic tactile contour following dates back a quarter century [30]–[32], with a more recent approach adopting a control framework for tactile servoing [33]. However, these studies have relied on applying image processing techniques (*e.g.* image moments) to planar taxel arrays. For curved biomimetic sensors such as the iCub fingertip, another approach is to use a non-parametric probabilistic model of the taxel outputs [34]–[36].

The TacTip tactile sensor has been used for contour following, using a combination of servo control and a probabilistic model of the pin displacements [10]. After tuning the control policy, the robot could tap around shapes such as a circle, volute and spiral. However, the trajectories [10, Figs 7-10] were not robust to parameter changes and failed when applied to sliding rather than tapping motion. Controlled sliding using touch is a challenge because the sensing surface undergoes motion-dependent shear due to friction against the object surface [37]. Training data thus differs from the sensor output during the task, which will cause supervised learning methods to fail unless shear invariance is somehow applied.

III. METHODS

A. Robotic system: Tactile sensor mounted on a robot arm

1) *Tactile sensor*: We use an optical biomimetic tactile sensor developed in Bristol Robotics Laboratory: the TacTip [5], [6]. The version used here is 3D-printed with a 40 mm-diameter hemispherical sensing pad (Fig. 1b) and 127 tactile pins in a triangular hexagonal lattice (Fig. 1c). Deformation of the sensing pad is imaged with an internal camera (ELP




Figure 2: System and software infrastructure, from the TacTip optical tactile sensor to controlling the ABB robot arm.

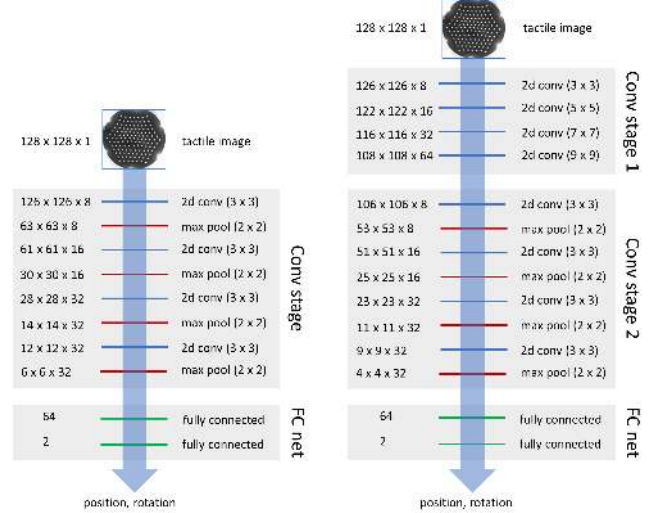


Figure 3: The two CNN architectures used in this work. The first (a) is based on a standard CNN pattern [43] and the second (b) has an additional input stage without max pooling.

USB 1080p module; used at 640×480 pixels and 30 fps). The pin deflections can accurately characterize contact location, depth, object curvature/sharpness, edge angle, shear and slip. For more details, we refer to recent studies with this tactile sensor [6]–[14], [38]–[42].

2) *Robot arm mounting*: The TacTip is mounted as an end-effector on a 6-DoF robot arm (IRB 120, ABB Robotics). The removable base of the TacTip is bolted onto a mounting plate attached to the rotating (wrist) section of the arm, then the other two modular components (central column and tip) are attached by bayonet fittings (Fig. 1a).

3) *Software infrastructure*: Our integrated sensorimotor framework has four components (Fig. 2): (1) Stacks of tactile images are collected in the OpenCV library, then preprocessed and either saved (for use in training) or used directly for prediction; (2) These images are first cropped and subsampled to (128×128) -pixel grey-scale images (Fig. 1c) and then passed to the Deep Learning system in Keras; (3) The resulting predictions are passed to control and visualization software in MATLAB; (4) The computed change in sensor pose is sent to a Python client that interfaces with a RAPID API for controlling the robot arm. Training is on a Titan Xp GPU hosted on a Windows 10 PC. The components run in real-time on the tasks (cycle time: ~ 1 ms for prediction and ~ 100 ms for control).

B. Deep learning system

Two types of CNN architecture are used here (Fig. 3): the first for tapping-based experiments; and the second, more complex architecture, to cope with the additional challenges associated with continuous-contact sliding motion.




Figure 4: Example tactile images from tapping on the disk. The angle is visible from the broader pin spacing above the edge. Further examples are given in the supplemental video.

The first architecture (Fig. 3a) passes a 128×128 grey-scale image through a sequence of convolutional and max-pooling layers to generate a high-level set of tactile features. These features are then passed through a fully-connected regression network to make predictions. This configuration is based on a simple convolutional network pattern [43], which was scaled and regularized by restricting the number of filters/ReLUs in each layer and using a dropout of 0.25 between the convolutional stage and fully-connected net.

The second architecture (Fig. 3b) was introduced to cope with the effects of non-linear shear on the sensor pin positions during continuous sliding motion. Under this type of motion, the pins tend to be displaced from the positions they occupy normally during tapping and thus their relative positions convey more useful information than their absolute positions.

Initially, we tried to encourage the original network architecture (Fig. 3a) to make use of relative pin positions by using data augmentation to introduce randomly shifted copies of the sensor images into the training data (shifting each image randomly by 0-2% in the horizontal and vertical directions on each presentation). However, this alone was not enough to achieve good performance for continuous-contact contour following around objects other than the disk.

We then extended the network by adding an initial convolution stage without any max-pooling layers before the original network (Fig. 3b), which did achieve better performance. Combined with the data augmentation, this allowed the network to learn broader features over larger groups of pixels, which allows the network to capture the spatial relationship between groups of adjacent pins. Once again, we used a dropout of 0.25 between the second convolutional stage and the fully-connected net to help over-regularize the model with respect to the validation and test data.

Both network architectures were trained using the Adam optimizer, with learning rate 10^{-4} and learning rate decay 10^{-6} . All networks were trained from scratch, using the default Keras weight initializers ('glorot_uniform'). Limiting the number of filters/ReLUs in each layer, using 0.25 dropout before the fully-connected net and early stopping (patience parameter 5) all helped prevent overfitting.

C. Task: Tactile servoing along a contour

Here we consider tasks in which a tactile sensor moves along a continuously-extended tactile feature such as the edge




Figure 5: Radial position and angular rotation prediction errors for the CNN in Fig. 3a, over 2000 random taps of 7 frames.

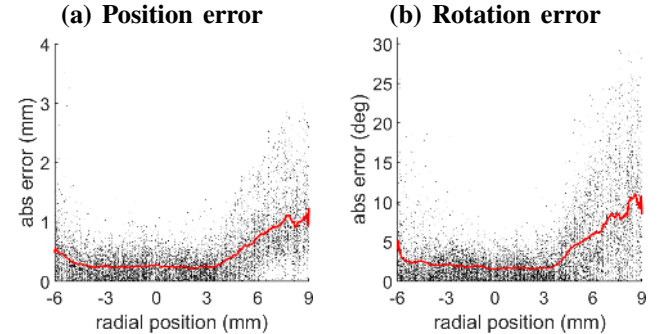


Figure 6: CNN prediction errors from Fig. 5 over radial position only. The red curve is a mean calculated from a 500-point (1/28 of the range) moving-average filter.

of an object, while rotating to maintain alignment with that feature. The control policy plans the next point to move along the contour around the object edge.

These tasks are performed on a range of objects chosen for a variety of shapes and material properties: a 105 mm-diameter circular disk, a tear-drop and a clover (all 3D-printed in ABS plastic); a lamina volute with radii of curvature 30, 40, 50, 60 mm in 90 deg segments and a 5 mm-wide ridge in a volute spiral with radii of curvature 20, 30, 40, 50, 60 mm in 180 deg segments (both laser-cut from acrylic sheet); and two objects from the YCB object set [44], one chosen to be compliant (a soft rubber brick) and the other irregular (a plastic banana). Only the circular disk is used to gather training data (at the 12 o'clock position).

We use a policy based on tactile servoing of the sensor pose to along the predicted edge orientation (r, θ) , with three components [10]: (i) a radial move Δr along the (predicted) normal to the edge; (ii) an axial rotation of the sensor $\Delta \theta$; and (iii) a translation Δe along the (predicted) tangent to the edge. This can be represented by a proportional controller

$$\Delta r = g_r (r_0 - r), \quad \Delta \theta = g_\theta (\theta_0 - \theta), \quad (1)$$

with unit gains $(g_r, g_\theta) = (1, 1)$, set-point $(r_0, \theta_0) = (0 \text{ mm}, 0 \text{ deg})$ (the edge pose in these coordinates) and we choose a default step $\Delta e = 3 \text{ mm}$ used before in ref. [10].

This control policy plans points for following an edge or contour. In the simple case considered here, the actions $(\Delta r, \Delta \theta) = (-r, -\theta)$ correspond to the training-data labels

and the tangential step Δe is assumed sufficiently small to not lose a curved edge. The policy may then be considered as end-to-end from the tactile images to controlled actions.

D. Data collection

For training, the tactile robotic system samples a local region of the edge of one object (105 mm disk at 12 o'clock position) over a range of radial positions r and axial (roll) rotation angles θ . Here we used 2000 uniform random tapping contacts sampled over ranges $(-6, +9)$ mm and $(-45, +45)$ deg, with each tap from ~ 1.5 mm above the object, down 5 mm, over ~ 0.7 sec (~ 20 frames). The origin (0 mm, 0 deg) has the sensor tip centred on the edge with camera and pin lattice aligned normal to the edge.

The training set was split into 1600 samples for learning the weights and 400 for hyperparameter optimization. Another dataset of 2000 contacts over new random positions and angles was used for evaluating perceptual performance. We used the 7 frames around the peak displaced frame (measured by the change in RMS pixel intensity from the initial frame) to capture data near the deepest contact part of the tap, including some variation over depth but excluding non-contact data. The images were then cropped and subsampled to a (128×128) -pixel region containing the pins (Fig. 4).

There is modest scope for improving the results by fine-tuning these experiment parameters. However, trial and error (and experience with the experiment [10]) showed these to be reasonable and natural choices. The only non-obvious choice was to include a random shift between ± 1 deg of the sensor yaw/pitch in all data: this reduced specialisation in the trained network to small but noticeable non-normal sensor alignments that otherwise biased the angle predictions.

IV. RESULTS

A. End-to-end edge perception from tactile images

In a major departure from recent work with the TacTip optical biomimetic sensor, here we predict the percepts directly from tactile images with a deep CNN. Prior work with this sensor has used specialised preprocessing to detect then track the pin positions [6]–[14], [38]–[42]. Here, this preprocessing is subsumed into the trained neural network.

We report the performance of the first CNN architecture (Fig. 3a). During our preliminary investigations, we found that networks with more filters/ReLU's in each layer and less regularization achieved better performance on the validation and test data collected at the same point on the disk as where the training data was collected. However, they failed to generalize well to other regions of the disk or to other objects. Over-regularizing the network beyond the point required for good generalization on the test data helped solve this problem and produced models that perform well on a broader range of tapping-based contour following tasks.

The overall perceptual performance is then most accurate near the central positions for all rotations (Fig. 5). In this region (-3 to $+3$ mm), errors are generally less than 1 mm and 5 deg. Overall, the contacts are less informative further from the edge (9 mm into free space; -6 mm onto the disk),

consistent with the edge being no longer visible in those tactile images. Considering only the dependence on position (Fig. 6), the mean absolute errors are ~ 0.3 mm and ~ 2 deg in the central region where perception is most accurate (red curve), appropriate for the contour-following tasks below.

B. Robust contour following around a disk

The deep CNN model for predicting edge pose angle and radial position is now applied to contour following around the disk with tapping contacts. A servo control policy (Eq. 1) plans contact points to maintain the relative pose to the edge.

The completed trajectories are near-perfect circles around the disk under a range of conditions (Fig. 7; Table I). Trajectories are repeatable, as indicated by running the experiment from different starting positions relative to the edge (Fig. 7a), with a small offset in the radial displacement (0-1 mm) and sensor angle (6-9 deg) relative to the edge. Increasing the policy step size from 3 mm keeps the trajectory within 2 mm of the edge, but offsets the angle by another 4-11 deg consistent with the sensor moving from its predicted location (Fig. 7b; $\Delta e = 6, 9$ mm). Taking a set-point radius inside or outside the edge keeps a circular trajectory without contributing to the angle offset (Fig. 7c; $r_0 = -3, +6$ mm).

The circular trajectories are also robust to changing the tapping depth (Fig. 7d, depth change $\Delta = -1.5, +2.5$ mm), from shallow taps (2.5 mm above, down 5 mm) to deep taps (-1.5 mm above, down 5 mm). Shallow taps advanced the sensor angle by 8 deg and deep taps lag the angle by 20 deg. The additional is because the task data is at a different depth from that used to train the model; although the model is most accurate at the training depth, its performance has declined gracefully to still complete the contour following task.

These results are a major improvement over those obtained in a previous study on the same task with a probabilistic model of pin displacements [10] (comparison in Table I), which failed to complete the task in many circumstances.

experiment	parameter variation	probabilistic model [10]	deep CNN
tapping contact (Fig. 7a)	$r_{\text{init}} = -6$ mm	1 mm, 23 deg	0 mm, 6 deg
	$r_{\text{init}} = 0$ mm	2 mm, 25 deg	1 mm, 6 deg
	$r_{\text{init}} = 9$ mm	2 mm, 41 deg	0 mm, 9 deg
tapping contact (Fig. 7b)	$\Delta e = 6$ mm	2 mm, 25 deg	1 mm, 10 deg
	$\Delta e = 9$ mm	fail	1 mm, 17 deg
tapping contact (Fig. 7c)	$r_0 = -2$ mm	fail	1 mm, 6 deg
	$r_0 = +6$ mm	3 mm, 16 deg	3 mm, 7 deg
tapping contact (Fig. 7d)	$\Delta = -1.5$ mm	fail	0 mm, 8 deg
	$\Delta = +2.5$ mm	4 mm, 25 deg	3 mm, 20 deg
sliding contact (Fig. 9a)	$r_{\text{init}} = -6$ mm	fail	1 mm, 15 deg
	$r_{\text{init}} = 0$ mm	fail	1 mm, 12 deg
	$r_{\text{init}} = 9$ mm	fail	1 mm, 11 deg
sliding contact (Fig. 9b)	$\Delta e = 6$ mm	fail	1 mm, 15 deg
	$\Delta e = 9$ mm	fail	3 mm, 34 deg
sliding contact (Fig. 7c)	$r_0 = -3$ mm	fail	2 mm, 25 deg
	$r_0 = +2$ mm	fail	2 mm, 9 deg
sliding contact (Fig. 7d)	$\Delta = -1$ mm	fail	1 mm, 18 deg
	$\Delta = +3$ mm	fail	1 mm, 4 deg

Table I. Accuracy of exploration around the disk, showing the mean absolute errors of radial position and rotation angle for the trajectories in Figs 7,9. A comparison is shown with the original probabilistic model [10] for each experiment.

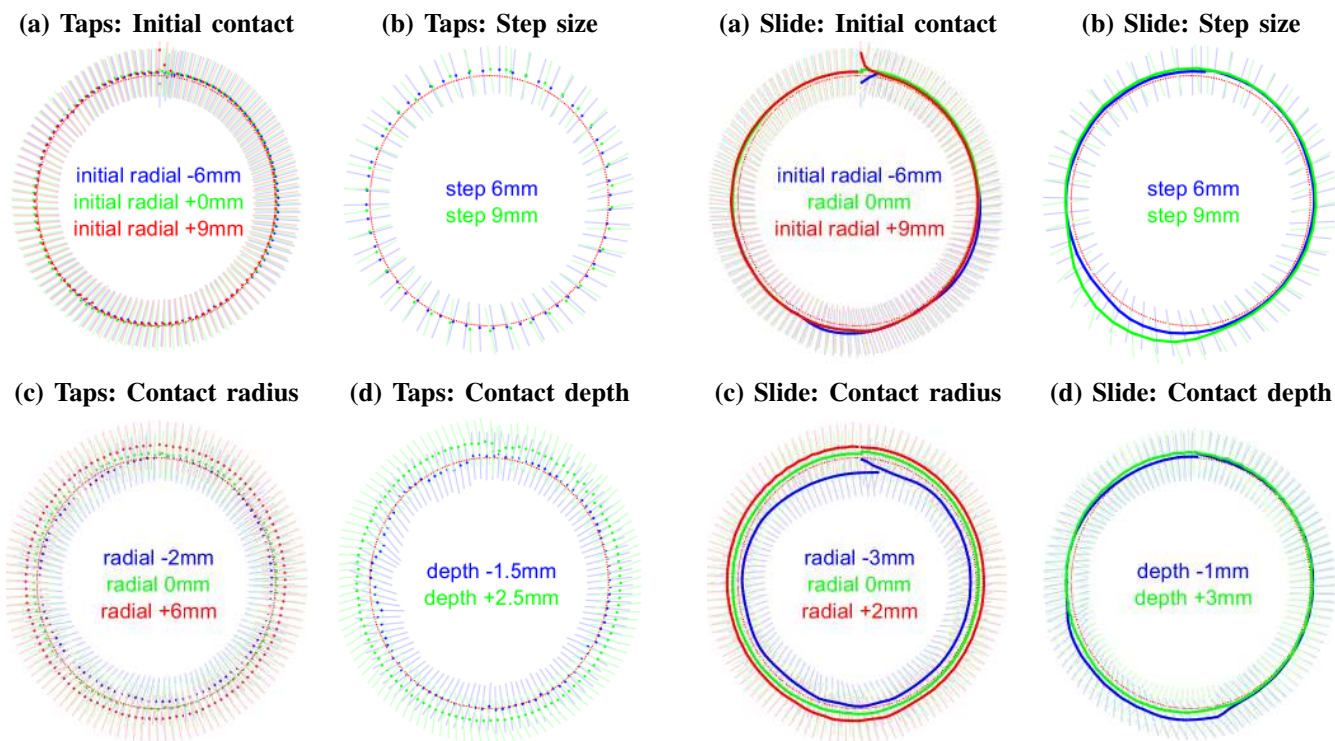


Figure 7: Robust tapping around a uniform disk to the contour changes in Table I (trained at the 12 o'clock position).

Figure 9: Robust sliding around a disk to the contour changes in Table I (trained with taps at the 12 o'clock position).

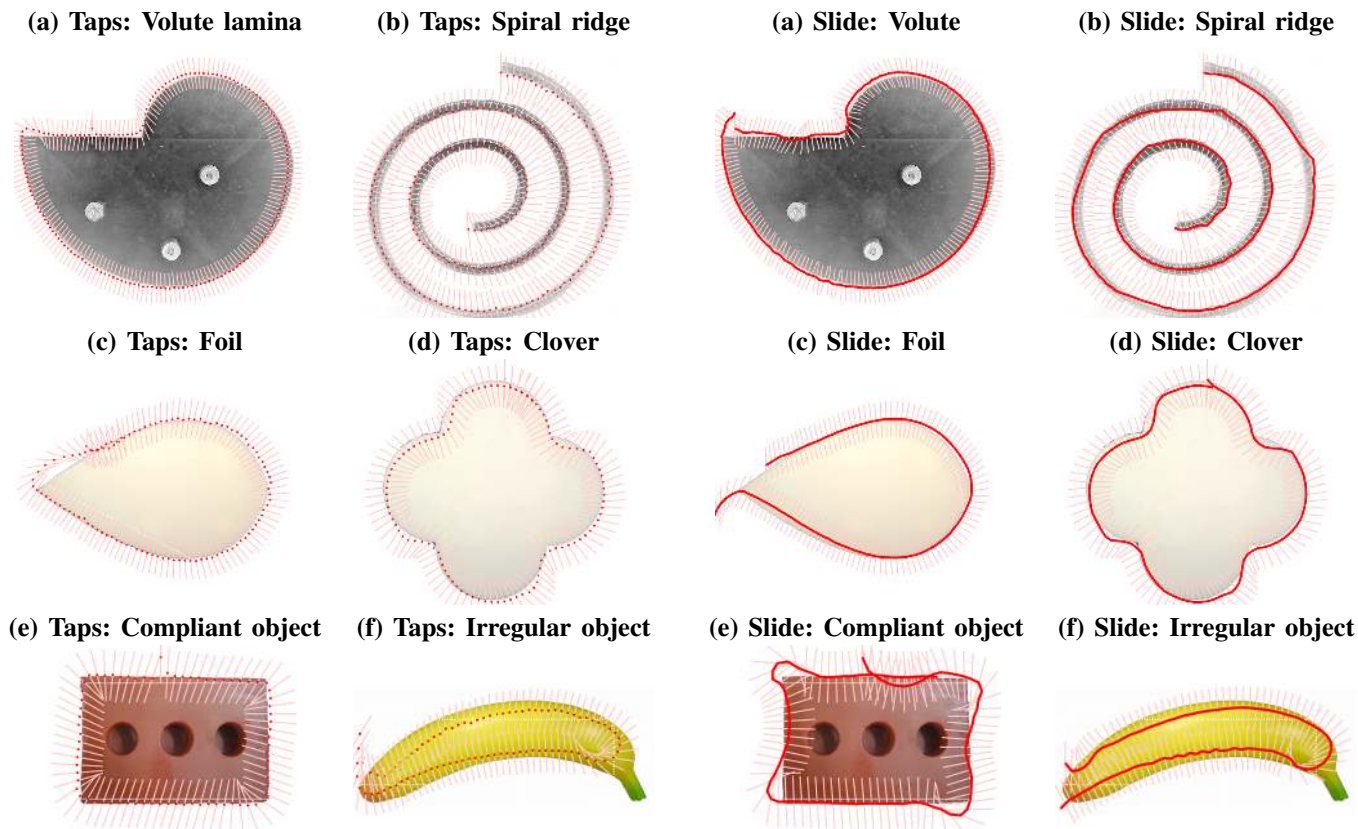


Figure 8: Robust tapping around non-uniform objects (trained with taps at the 12 o'clock position on the disk).

Figure 10: Robust sliding around non-uniform objects (trained with taps at the 12 o'clock position on the disk).

C. Robustness to non-uniform object shapes

To demonstrate further robustness, the task is extended to four fabricated planar shapes chosen to have non-uniform curvature (a volute lamina, spiral ridge, tear drop and clover) and two household objects from the YCB object set [44]: one compliant (a soft rubber brick) and one irregular (a plastic banana). Task completion shows generalization to novel contours differing from the disk edge used in training.

For the four fabricated shapes, the completed trajectories match the object shapes (Figs 8a-d). When the radius of curvature is close to that of the disk, the sensor angle aligns to the edge normal. For smaller or larger curvatures, there is offset in the sensor angle (advancing for smaller and lagging for higher). The large changes in orientation at corners on the volute and foil cause overshoots, but the task still completes.

The task also completes on the compliant object (Fig. 8e, rubber brick). Again, the policy advances the edge angle on the straight sections. There is less overshoot at the compliant corners than those of the rigid objects; however, the angle offset on the straight edges would make turning easier.

The task does not quite complete on the irregular object (Fig. 8f, banana). The task was challenging because there are only slight ridges and the height varies by a few millimetres. However, the policy managed to traverse most of the object, failing only at the tip where there is both a sharp change in orientation and no well-defined edge to follow.

D. Robustness to sliding contact

A far more demanding test of robustness is to use a continuous sliding motion with the same training data from tapping the disk. We repeat the above experiments, dropping the sensor 3 mm (1.5 mm into the object) and collecting data between the exploratory movements (~ 0.15 sec duration; ~ 5 frames). The second CNN architecture (Fig. 3b) was used, which has an initial convolution stage without any max-pooling layers to help generalize over tactile features.

Under a sliding motion, the task completion was robust to changes in starting point (Fig. 9a), step size (Fig. 9b; $\Delta e = 6, 9$ mm), set-point radius (Fig. 9c; $r_{\text{fix}} = -3, +2$ mm) and contact depth (Fig. 9d; -1 mm shallower, $+3$ mm deeper). The range of depths where the task completes is likely greater, but concerns about damaging the sensor limited further testing. The angle offset improved to 4 deg with the deepest contact, which is the best overall trajectory.

Three objects (the spiral, clover and compliant brick) were successfully traversed with sliding (Figs 10b,d,e). The two rigid objects have similar sliding and tapping trajectories (*c.f.* Figs 8b,d and 10b,d). For the rubber brick, the angle offset on the straight edges appears less than for tapping motion, but there is a larger overshoot around the corners (*c.f.* Figs 8e and 10e). We interpret this as a consequence of the sliding motion inducing a shear of the sensing surface that makes corners more challenging.

The policy failed on the other three objects (the volute, tear drop and banana) at the sharp corners (Figs 10a,c,f), but was successful otherwise. The first two were objects with large overshoots at the corners when tapping (Figs 8a,c), which

appear to have caused the failure under the more demanding condition of sliding. The policy also failed at the tip of banana for both tapping (Fig. 8f) and sliding motion (Fig. 10f) where there is no well-defined edge to follow.

V. DISCUSSION

This work is the first application of deep learning to an optical biomimetic tactile sensor and the first to edge perception and contour following. We found robust generalization of the contour-following policy to tasks beyond which the model was trained, such as continuous sliding around compliant or irregular objects after training with taps on part of a disk.

We used two techniques to encourage generalization. The first was to anticipate nuisance variables to marginalise out and then either modify the data collection (e.g. training over shifts in yaw/pitch) or augment the data set with artificially generated data (e.g. randomly shifting frames). The second technique was to over-regularize the architecture to avoid specialization to the training task; this may be because it encourages the development of simpler features throughout the network. In both cases, we introduce inductive bias into the model to improve performance in situations different from its training. This is necessary because generalizing beyond the task a model is trained on cannot be reliably achieved by trying to validate on data from the original task.

We emphasise the generalization from discrete tapping contacts to continuous sliding motion is a challenging test for the policy. Sliding causes a friction-dependent shear of the sensing surface that depends on the motion direction and recent history of the interaction [37]. Hence, the tactile data during a task can differ in complex ways from those during training. Although this caused our previous statistical model to fail [10], the deep learning model performed robustly.

The principal failure mode of the deep learning model was on sharp corners under sliding motion. This is unsurprising, as the model was only trained on edge data from the disk, so corners are both outside its experience and give a singularity in the prediction. The model degraded gracefully, with corners successfully followed with a tapping motion, albeit with some overshoot, and also for sliding around reflex angles and the compliant object. In principle, this limitation could be solved by crafting a more complete policy that can predict points around corners, *e.g.* by training on corners of various angles. A complementary method would be to adapt the step size of the policy based on the predicted curvature of the object.

In our view, the greatest benefit of using a deep CNN to learn a tactile control policy is its capability to generalize beyond the training data. Previous studies with the same biomimetic tactile sensor found good performance when the task and training were similar [8]–[10], [41]; however, we were aware that these results were fragile to small changes in the task (*e.g.* sensor orientation). Since practical applications of tactile sensing require robust performance in situations beyond those previously experienced, we expect this is a generic problem in robot touch that deep learning will solve.

Acknowledgements: We thank NVIDIA Corporation for the donation of the Titan Xp GPU used for this research

REFERENCES

- [1] V Hayward. Is there a 'plenhaptic' function? ***Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1581):3115–3122, 2011.
- [2] N Sündlerhauf, O Brock, W Scheirer, R Hadsell, D Fox, J Leitner, B Upcroft, P Abbeel, W Burgard, M Milford, and P Corke. The limits and potentials of deep learning for robotics. *The International Journal of Robotics Research*, 37(4-5):405–420, 2018.
- [3] W Yuan, S Dong, and E Adelson. GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force. *Sensors*, 17(12):2762, 2017.
- [4] W Yuan, R Li, M Srinivasan, and E Adelson. Measurement of shear and slip with a GelSight tactile sensor. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 304–311, 2015.
- [5] C Chorley, C Melhuish, T Pipe, and J Rossiter. Development of a Tactile Sensor Based on Biologically Inspired Edge Encoding. In *International Conference on Advanced Robotics*, pages 1–6, 2009.
- [6] B Ward-Cherrier, N Pestell, L Cramphorn, B Winstone, M Giannaccini, J Rossiter, and N Lepora. The TacTip Family: Soft Optical Tactile Sensors with 3D-Printed Biomimetic Morphologies. *Soft Robotics*, 5(2):216–227, 2018.
- [7] L Cramphorn, B Ward-Cherrier, and N Lepora. Addition of a Biomimetic Fingerprint on an Artificial Fingertip Enhances Tactile Spatial Acuity. *IEEE Robotics and Automation Letters*, 2(3):1336–1343, 2017.
- [8] N Lepora and B Ward-Cherrier. Superresolution with an optical tactile sensor. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2686–2691, 2015.
- [9] N Lepora. Biomimetic Active Touch with Fingertips and Whiskers. *IEEE Transactions on Haptics*, 9(2):170–183, 2016.
- [10] N Lepora, K Aquilina, and L Cramphorn. Exploratory Tactile Servoing with Active Touch. *IEEE Robotics and Automation Letters*, 2(2):1156–1163, 2017.
- [11] B Ward-Cherrier, L Cramphorn, and N Lepora. Exploiting Sensor Symmetry for Generalized Tactile Perception in Biomimetic Touch. *IEEE Robotics and Automation Letters*, 2(2):1218–1225, 2017.
- [12] L Cramphorn, B Ward-Cherrier, and N Lepora. Tactile manipulation with biomimetic active touch. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 123–129, 2016.
- [13] K Aquilina, D Barton, and N Lepora. Principal Components of Touch. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4071–4078, 2018.
- [14] J James, N Pestell, and N Lepora. Slip Detection with an Optical Tactile Sensor. *IEEE International Conference on Robotics and Automation (ICRA)*, 3(4):3340–3346, 2018.
- [15] A Schmitz, Y Bansho, K Noda, H Iwata, T Ogata, and S Sugano. Tactile object recognition using deep learning and dropout. *IEEE-RAS International Conference on Humanoid Robots*, pages 1044–1050, 2014.
- [16] L Cao, R Kotagiri, F Sun, H Li, W Huang, Z Maung, and Maung Aye. Efficient Spatio-Temporal Tactile Object Recognition with Randomized Tiling Convolutional Networks in a Hierarchical Fusion Strategy. *AAAI Conference on Artificial Intelligence*, 2015.
- [17] M Meier, F Patzelt, R Haschke, and H Ritter. Tactile Convolutional Networks for Online Slip and Rotation Detection. *International Conference on Artificial Neural Networks (ICANN)*, 2016.
- [18] S Baishya and B Bäuml. Robust material classification with a tactile skin using deep learning. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8–15, 2016.
- [19] J Kwiatkowski, D Cockburn, and V Duchaine. Grasp stability assessment through the fusion of proprioception and tactile signals using convolutional neural networks. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 286–292, 2017.
- [20] W Yuan, C Zhu, A Owens, M Srinivasan, and E Adelson. Shape-independent hardness estimation using deep learning and a GelSight tactile sensor. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 951–958, 2017.
- [21] R Calandra, A Owens, M Upadhyaya, W Yuan, J Lin, E Adelson, and Sergey Levine. The Feeling of Success: Does Touch Sensing Help Predict Grasp Outcomes? *Conference on Robot Learning (CoRL)*, pages 1–10, 2017.
- [22] W Yuan, S Wang, S Dong, and E Adelson. Connecting Look and Feel: Associating the visual and tactile properties of physical materials. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4494–4502, 2017.
- [23] J Li, S Dong, and E Adelson. Slip Detection with Combined Tactile and Visual Information. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7772–7777, 2018.
- [24] 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)*, pages 1927–1934, 2018.
- [25] F Hogan, M Bauza, O Canal, E Donlon, and A Rodriguez. Tactile Regrasp: Grasp Adjustments via Simulated Tactile Transformations. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2963–2970, 2018.
- [26] R Calandra, A Owens, D Jayaraman, J Lin, W Yuan, J Malik, E Adelson, and S Levine. More Than a Feeling: Learning to Grasp and Regrasp Using Vision and Touch. *IEEE Robotics and Automation Letters*, 3(4):3300–3307, 2018.
- [27] M Plaisier, W Tiest, and A Kappers. Salient features in 3-D haptic shape perception. *Attention, perception & psychophysics*, 71(2):421–430, 2009.
- [28] S Lederman and R Klatzky. Relative Availability of Surface and Object Properties during Early Haptic Processing. *Journal of Experimental Psychology: Human Perception and Performance*, 23(6):1680–1707, 1997.
- [29] R Ponce Wong, R Hellman, and V Santos. Spatial asymmetry in tactile sensor skin deformation aids perception of edge orientation during haptic exploration. *IEEE Transactions on Haptics*, 7(2):191–202, 2013.
- [30] A Berger. On Using a Tactile Sensor for Real-Time Feature Extraction. *Master's thesis, Carnegie-Mellon University.*, 1988.
- [31] A Berger and P Khosla. Using Tactile Data for Real-Time Feedback. *The International Journal of Robotics Research*, 10(2):88–102, 1991.
- [32] N Chen, H Zhang, and R Rink. Edge tracking using tactile servo. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2:84–89, 1995.
- [33] Q Li, C Schurmann, R Haschke, and H Ritter. A control framework for tactile servoing. *Robotics: Science and Systems (RSS)*, 2013.
- [34] U Martinez-Hernandez, G Metta, T Dodd, T Prescott, L Natale, and N Lepora. Active contour following to explore object shape with robot touch. In *World Haptics Conference (WHC)*, 2013.
- [35] U Martinez-Hernandez, T Dodd, T Prescott, and N Lepora. Active Bayesian perception for angle and position discrimination with a biomimetic fingertip. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5968–5973, 2013.
- [36] U Martinez-Hernandez, T Dodd, M Evans, T Prescott, and N Lepora. Active sensorimotor control for tactile exploration. *Robotics and Autonomous Systems*, 87:15–27, 2017.
- [37] W Chen, H Khamis, I Birznieks, N Lepora, and S Redmond. Tactile Sensors for Friction Estimation and Incipient Slip Detection – Towards Dexterous Robotic Manipulation: A Review. *IEEE Sensors Journal*, 18(22):9049–9064, 2018.
- [38] B Ward-Cherrier, N Rojas, and N Lepora. Model-Free Precise in-Hand Manipulation with a 3D-Printed Tactile Gripper. *IEEE Robotics and Automation Letters*, 2(4):2056–2063, 2017.
- [39] B Ward-Cherrier, L Cramphorn, and N Lepora. Tactile Manipulation with a TacThumb Integrated on the Open-Hand M2 Gripper. *IEEE Robotics and Automation Letters*, 1(1):169–175, 2016.
- [40] L Cramphorn, J Lloyd, and N Lepora. Voronoi Features for Tactile Sensing: Direct Inference of Pressure, Shear and Contact Locations. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2752–2757, 2018.
- [41] N Lepora and B Ward-Cherrier. Tactile Quality Control with Biomimetic Active Touch. *IEEE Robotics and Automation Letters*, 1(2):646–652, 2016.
- [42] N Pestell, J Lloyd, J Rossiter, and N Lepora. Dual-Modal Tactile Perception and Exploration. *IEEE Robotics and Automation Letters*, 3(2):1033–1040, 2018.
- [43] I Goodfellow, Y Bengio, and A Courville. *Deep Learning*. MIT Press, 2016.
- [44] B Calli, A Wallsman, A Singh, and S Srinivasa. Benchmarking in Manipulation Research. *IEEE Robotics and Automation Magazine*, pages 36–52, 2015.