Ward-Cherrier, B. P. J., Cramphorn, L., & Lepora, N. F. (2017). Exploiting Sensor Symmetry for Generalized Tactile Perception in Biomimetic Touch. *IEEE Robotics and Automation Letters*, 2(2), 1218-1225. <https://doi.org/10.1109/LRA.2017.2665692>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1109/LRA.2017.2665692](https://doi.org/10.1109/LRA.2017.2665692)

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

This is the final published version of the article (version of record). It first appeared online via IEEE at <https://doi.org/10.1109/LRA.2017.2665692> . 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/>

Exploiting Sensor Symmetry for Generalized Tactile Perception in Biomimetic Touch

Benjamin Ward-Cherrier, Luke Cramphorn, and Nathan F. Lepora

Abstract—Tactile perception methods rely on creating mappings from tactile data to percepts. In many approaches to artificial tactile perception, this involves extensive sampling of the object during the training phase. We introduce here a method to instead generalize tactile features across different orientations. This method is applied to the TacTip v2, a three-dimensional printed optical tactile sensor with internal pins acting as taxels arranged with a 12-fold rotational symmetry. By rotating a small sample of tactile images, we are able to generalize tactile stimuli to new orientations. The method was validated across several tactile stimuli on an edge orientation classification task. Data were then generalized for a combination of orientations and locations of one of these stimuli, and this dataset was used as the basis for an exploratory control task: contour following around a circular disk. The generalization method leads to a strong reduction in the time needed to gather training data and only a moderate increase in classification error, and is particularly suited for multidimensional tactile data sampling in complex tasks such as tactile manipulation. We expect the method to generalize well to other tactile stimuli.

Index Terms—Biomimetics, force and tactile sensing.

I. INTRODUCTION

THE sense of touch is essential for humans to interact effectively with their environment, particularly in fine manipulation tasks [1]. The standard approach in tactile perception is to gather training data to obtain a model that maps data to percepts [2]–[5]. For example, in work on biomimetic active touch [6] this has relied on exhaustive sampling of stimuli over 4000 samples spanning 40 mm to demonstrate superresolution [7]. A problem is that for each new stimulus, training is a time consuming process (sometimes taking hours). Flexible manipulation and perception tasks based on tactile perception could thus strongly benefit from methods that reduce the initial quantity of training data to be gathered.

We introduce here a method for reducing training data gathering time by exploiting sensor symmetries to generate new

Fig. 1. The TacTip v2 sensor (left), and a novel tip with a 12-fold rotationally symmetric pin layout (right).

tactile data. As an example of this method, we demonstrate its performance on a 3d-printed optical tactile sensor: an improved version of the TacTip [8]. The TacTip v2 optically tracks pins on its deformable skin, which are here arranged with a 12-fold rotational symmetry (Fig. 1). This was chosen since 12 factorises nicely over 360° (one of the reasons clocks are divided into 12 hours). This symmetry in the pin layout is exploited by gathering only a small sample of training data and rotating that set of pin deflections 12-fold to obtain the various oriented tactile stimuli.

The method is first validated by performing edge orientation classification on a series of tactile stimuli. Data generalization only moderately increases orientation perception error (by approximately 1°), even though the data gathering time is significantly reduced to under a minute from almost ten minutes originally. Next, the method is applied to a set of orientations of a stimulus and locations across that stimulus, and the generalized data is used as the basis for an exploratory task. The task consists of contour following around a circular disk, and is demonstrated using generalized orientation and location tactile data in conjunction with an active perception control policy which displaces the sensor tangentially to the detected edge. The improvement in training data gathering time becomes more significant in this 2-dimensional dataset (from close to 20 minutes to approximately 1 minute and a half), and the moderate increase in classification error does not impede task completion.

Thus we propose that tactile data can be reliably generated through our method, in effect giving a generalization from known encountered stimuli to novel stimuli from the symme-

Manuscript received September 10, 2016; accepted January 16, 2017. Date of publication February 8, 2017; date of current version March 1, 2017. This letter was recommended for publication by Associate Editor F. Lizaralde and Editor J. Wen upon evaluation of the reviewers' comments. The work of B. Ward-Cherrier was supported by an EPSRC DTP studentship. The work of L. Cramphorn was supported by a studentship in the EPSRC CDT in Future Autonomous and Robotic Systems (FARSCOPE). The work of N. Lepora was supported in part by a grant from the Engineering and Physical Sciences Research Council (EPSRC) on 'Tactile Superresolution Sensing' (EP/M02993X/1).

The authors are with the Department of Engineering Mathematics, University of Bristol and Bristol Robotics Laboratory, University of Bristol, Bristol BS16 1QY, U.K. (e-mail: bw14452@bristol.ac.uk; ll14468@bristol.ac.uk; n.lepora@bristol.ac.uk).

Color versions of one or more of the figures in this letter are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LRA.2017.2665692

try of the tactile sensor. This could be of particular benefit in complex tactile perception tasks involving multidimensional datasets, in which data gathering time is currently a bottleneck. The method should in principle apply to a range of tactile stimuli and sensors, and could be of particular value in generalizing over stimulus contact location for tactile skins.

II. BACKGROUND AND RELATED WORK

Feature extraction has been an important topic of research in vision for many years, with a host of algorithms designed to identify salient affine-invariant features [9], [10].

In tactile perception, computer vision techniques can be applied to optical tactile sensors such as the GelSight [11] or TacTip [12], but improvements in the resolution of tactile sensors have led to the production of tactile images by many types of sensors, encouraging a more general use of these algorithms in tactile perception. Use of techniques such as bag-of features [4], [13] have become common, as well as the application of computer vision techniques to time-series of tactile data [3]. As is the case for vision, these techniques tend to be applied to object recognition, and thus seek to identify features that are invariant to affine transformations of the tactile images. For instance, Luo *et al.* extracted tactile features that characterise local shape in a rotation and translation-invariant manner [14]. Here we take an alternative approach that seeks to generalize tactile features across different orientations, while still maintaining the ability to discriminate between them.

Our sensor is based on the TacTip [8], an optical tactile sensor which has undergone several iterations [7], [15], [16]. The TacTip tracks the deflection of pins on its inside surface, and this latest version (v2) is modular, allowing us to easily design and 3d-print a specific tip with a 12-fold rotationally symmetric pin layout. Tactile data gathered with this tip can then be generalized through rotations to obtain new orientations of tactile features.

The motivation for generalizing tactile data comes from our use of Bayesian approaches to perception [17]. These methods require training, and in some cases [18] the training data is sampled over many dimensions. Means of reducing training could thus become very useful, and in some cases essential as dimensions are scaled up.

III. METHODS

A. TacTip Design and Fabrication

The TacTip is a cheap, robust, 3d-printed optical tactile sensor developed at Bristol Robotics Laboratory [8]. In this newest iteration of the TacTip sensor (TacTip v2), we have introduced several improvements relative to the previous design [7], most notably 3d-printing the tip in a multi-material 3d-printer, introducing modularity to allow for easy tip replacement (Fig. 2) and creating a novel pin layout (Fig. 1) as described in more detail below.

The TacTip v2 sensor can be divided into two parts :

Tip: This is made up of a hard plastic casing and a black rubber hemisphere, 3d-printed as one piece using a multi-material printer (materials: Vero White and Tango Black+). The rubber

Fig. 2. Design of the sensor. The TacTip v2 has a rubber membrane filled with clear silicon gel, with a rotationally symmetric pattern of white-tipped pins on its inside surface. These are illuminated by the LED circuit and their displacements are tracked by a Microsoft Cinema HD webcam mounted on the TacTip v2's base.

Fig. 3. Edge stimuli used to validate our edge orientation generalization method.

skin has white-tipped pins on its inside surface (~ 1 mm dia.) which are displaced when an object is contacted. Pins are separated from each other by 3 mm on the sensor image and their displacements characterize and amplify skin deformations from object contact. The dome-shaped skin is filled with RTV27905 clear silicon gel, allowing the TacTip v2 to be compliant at its tip. A 3 mm thick circular acrylic lens keeps the gel separate from electronic components (LED circuit and webcam).

Base: This 3d-printed section holds the ring-shaped circuit of 6 LEDs used to illuminate the pins. It also connects to the webcam mount, which houses a Microsoft Cinema HD webcam used to track pin displacements.

In the fabrication of the sensor, emphasis is placed on straight-forward manufacture and ease of assembly. The TacTip v2's base and tip are entirely 3D-printed, with the lasercut lens, gel, LED circuit and webcam all being straight-forward to assemble. Rapid prototyping keeps costs down, and allows us to easily test different versions for performance optimization. The tip is entirely 3d-printed and attached with a simple bayonet mount. This means that tips can be interchanged and tested with the same base, allowing for convenient testing and redesign of the sensor tip, as well as further cost reduction through the use of a single base. Through a trial and error process, we searched for a skin thickness that would balance skin robustness (avoiding cracks) while maximising pin displacement to obtain more

Fig. 4. Structure of the code for the experimental setup.

Fig. 5. TacTip v2 mounted on the ABB IR120 6-dof robot arm above the square edge stimulus.

informative tactile data. We choose a skin thickness of 1 mm and pin length of 2.3 mm in our final design.

The pin distribution is designed according to a rotationally symmetric pattern, with 49 pins arranged in 12 lines of 4 pins emanating from the centre of the sensor and one central pin. The pins form a 12-fold rotational symmetry (Fig. 1), which is exploited to generalize training data at different orientations. We chose a 12-fold symmetry in order to maximise time gains in the data gathering phase while allowing pins to be arranged in a simple, compact layout. 12 is also a good number for factorising over 360° (one of the reasons we use it to display hours on a clock). This particular layout was chosen to make evident the symmetry inherent in the sensor, however our method could be applied to any rotationally symmetric layout (series of circles, hexagonal, square, triangular) with any number of symmetries.

B. Experimental Setup

The TacTip v2 sensor is mounted vertically on the end of an ABB IRB120 6-DOF industrial robot arm (Fig. 5). The TacTip v2 is initially trained over a series of edge stimuli in the one-dimensional validation phase, which we describe as: sharp edge, square edge, large curve and small curve (Fig. 3). A more in-depth 2-dimensional validation (edge orientation and location across the edge) is then performed for the large curve object, followed by a practical demonstration of contour following on the large and small curve objects.

The ABB arm is controlled through an IronPython interface, and data is gathered from the TacTip v2 sensor with python OpenCV and analysed in MATLAB.

1) *Data Collection:* Data is collected from the TacTip v2 as a time-series of x - and y - pin coordinates during the vertical taps against each edge stimulus. Each tap lasts approximately 1.5 seconds, and gathers on average 30 frames.

a) *Training:* During training, the TacTip v2 performs vertical taps, which consist of a 5 mm downwards move followed by a 5 mm upward move. These taps contact the edge stimulus, and the sensor is rotated 1° after each tap to obtain 360 edge orientations over a range of 360° (Fig. 5).

For the contour following experiment, 72 angles are collected over a 360° range in 5° increments. For each angle, the sensor is moved across the edge over a 20 mm range, gathering data in 1 mm increments. We thus obtain a 2-dimensional dataset which includes 'where' (location across the edge) and 'what' (orientation) information.

b) *Testing:* For validation of our method, we use offline testing. A cross-validation method is used on 2 sets of collected data to assess the classification performance of edge orientation using the TacTip v2. After collecting 2 sets of data (training and test), samples from the test set are classified through a comparison with the training set, as described in Section III-C

2) *Data Pre-Processing:* In the TacTip v2 the webcam records images at approx. 20 fps. The x - and y -coordinates of pin centres (Fig. 1) are extracted using contour detection in opencv (<http://opencv.org/>). Pins are identified in each new image based on their proximity to a default set of positions. To deal with failures of the opencv contour detection in which pins can go undetected in a given frame, if no pin is found within a 2 mm radius of its default, the previous frame's pin position is repeated for the current image. This ensures that pins are assigned a position in their appropriate region of deflection even in frames in which they are not detected. The x - and y - coordinates of pin centres are treated as separate taxels and are sent to Matlab for analysis (Fig. 4). We thus have $49 \times 2 = 98$ taxels.

3) *Data Generalization:* Our method for data generalization considers a section of training data, applying rotations to tactile data based on their rotational symmetry to generalise across edge orientation. For the specific tip used here, the pins are laid out with 12-fold rotational symmetry (Fig. 1). Note however that the method could apply in principle to any sensor layout that has some level of rotational symmetry, to generalise features across all orientations.

We train the TacTip v2 to detect edge orientation by rotating over a 360° range in 1° increments before tapping down on a fixed edge stimulus. To generalize across orientation, we train only on the first 30 edge orientations (0–30°), and rotate each set of tactile data gathered 30° around the central pin to obtain the next 30 orientations (30–60°).

Pins from the rotated frame are reassigned to the identity of their nearest pin in the original non-rotated frame. Pin positions are recalculated by translating their x - and y -coordinates to a coordinate system centred on the TacTip v2's central pin, converting them to polar coordinates, incrementing the polar angle θ by 30° and re-converting to Cartesian coordinates. Thus we ensure that the deflections of pins from their default positions are rotated 30° and occur in the appropriate direction for the next 30 orientations.

By repeating this step 12 times, we are able to obtain all 360 estimated tactile images for the whole 360° range. Note however that the first 30 tactile images will remain as real, physically obtained data, since they are used to generate the rest of the range.

C. Passive Perception - Validation

Orientation detection on the TacTip v2 is achieved through a probabilistic method based on maximum likelihood. These methods have previously been applied to the TacTip [19]. We summarize the approach here, referring to previous work for a more detailed account [7], [19].

Training data is collected (or generalized) for each of the orientation values: $\theta_1 = 1, 2, 3 \dots 360$ as a multi-dimensional time series of taxel values z :

$$z = \{s_k(j) : 1 \leq j \leq N_{\text{samples}}, 1 \leq k \leq N_{\text{dims}}\}, \quad (1)$$

with j the time sample (frame number) and k the taxel number (x - and y - deflections of pins). Orientation likelihoods $P(z|\theta_1)$ are constructed using a log likelihood model of the training contact data

$$\log P(z|\theta_1) = \sum_{k=1}^{N_{\text{dims}}} \sum_{j=1}^{N_{\text{samples}}} \frac{\log P_k(s_k(j)|\theta_1)}{N_{\text{samples}} N_{\text{dims}}} \quad (2)$$

This model assumes statistical independence between data dimensions and time samples. Sensor values s_k are then binned into equal intervals I_b , $1 \leq b \leq N_{\text{bins}}$ using a histogram method

$$P_k(s_k(j)|\theta_1) = \frac{n_{kl}(b) + \epsilon}{\sum_{b'=1}^{N_{\text{bins}}} n_{kl}(b')}, \quad (3)$$

where $n_{kl}(b)$ is the number of samples in bin b for taxel k at orientation θ_1 . $\epsilon \ll 1$ is added to the equation to avoid an ill-defined likelihood with empty bins.

To classify a feature (orientation here) maximum likelihood is then applied, such that the classified orientation is

$$\theta_1 = \underset{1}{\operatorname{argmax}}(\log P(z|\theta_1)) \quad (4)$$

Validation of performance for both the real and generalized training data is carried out through an offline testing process. We train the sensor and save a training dataset, then retrain to obtain

a distinct test set. Contact data z is then sampled from the test set and classified for orientation as described above (1000 iterations per class). The mean absolute orientation classification error

$$e_\theta(\theta_1) = \langle |\theta_1(\text{classified}) - \theta_1(\text{actual})| \rangle \quad (5)$$

is reported for each class.

This algorithm is also applied to a 2-dimensional tactile dataset to identify orientation and location. In this case, location classification is performed as described in (1) to (4) above, with orientation θ_1 being replaced by location x_1 . The errors in location and orientation are then similarly found with a sampling procedure as

$$e_\theta(x_1, \theta_1) = \langle |\theta_1(\text{classified}) - \theta_1(\text{actual})| \rangle \quad (6)$$

$$e_x(x_1, \theta_1) = \langle |x_1(\text{classified}) - x_1(\text{actual})| \rangle \quad (7)$$

for each location and orientation class. These errors can then be reduced to a single dimension (location or orientation) by simply averaging the errors over the other dimension.

Orientation errors are cyclical, in that the error between a 359° and 0° angle needs to be reported as 1° and not 359°. To implement this, we take the inverse cosine of the cosine of the errors, thus always obtaining an orientation error in the correct range of 0–180°

D. Active Perception - Contour Following

As a further test of the generalization method, we consider a contour following task, in which the sensor moves along the edge of a circular disk. The sensor regularly taps vertically downwards (5 mm) on the edge and then returns to its original height. Contour following is not a focus of this letter, but a very useful application of the generalization method in a practical scenario. For this reason and for space considerations, we refer to the original reference for details of the algorithm implementation [20] and describe its conceptual basis here.

In the contour following experiment, 2-dimensional datasets over both edge orientations (360° range) and location of the sensor across the edge (20 mm range) are used. Along with the generalization method, an active perception policy is introduced to move the sensor around the outer edge of the circular stimulus. After each tap, both sensor location and edge orientation are classified, and based on this estimation, the sensor is moved tangentially 1 mm along the (estimated) edge. With each move, the sensor also adjusts its location along the 20 mm range of x locations to realign the edge to its centre. The contour following task is demonstrated on both the large curve and small curve objects described in the one-dimensional validation. It requires both location and angle perception to succeed, and the generalization method allows for a strong reduction in data gathering time with this 2-dimensional dataset.

IV. RESULTS

A. Inspection of Data

Data is collected from the TacTip v2 sensor as time-series of repeated vertical taps over an edge stimulus (Fig. 5), each tap consisting of approximately 30 frames over 98 taxels

Fig. 6. Real and generalized training data for the TacTip v2: Left panels show real training data, right panels correspond to generalized data. In each panel, the top graph represents pin deflections in the x -direction and the bottom graph in the y -direction (treated as separate taxels). The middle panel identifies pins according to their position on the sensor skin.

(x - and y - deflections for each pin). Between each tap, the TacTip v2 is rotated 1° over the full 360° range thus training over 360 equally spaced edge orientations. In the case of generalized training data, only the first 30 angle orientations are physically gathered, the rest of the range being obtained through 30° rotations of the tactile images. It is not immediately apparent that the real training data gathered (Fig. 6, left panel) represents a series of different edge orientations. We can however observe a sinusoidal dependence of deflection magnitude on angle, relating to the rotation between x - and y - coordinates of the pins. The generalized data (Fig. 6, right panel) is very similar to the real training data to the naked eye, with minor differences detectable at certain angles (for instance positive y -deflections around the 240° mark, Fig. 6) corresponding to minor inaccuracies from the data rotation.

A root mean square (RMS) difference comparison of the generalized and real datasets gives 0.28 px. This value represents the RMS difference in pin-deflections averaged over all frames, taxels and orientations, and is in the same range as when comparing 2 sets of real data: 0.19 px, signifying strong similarity between generalized and real data, and leading us to expect low errors in the validation of our method.

B. Validation of Tactile Generalization Method

Validation of the method is performed by running offline testing using a pre-collected test data set as described in Section III-B1b and comparing the classification accuracy of real training data with that of generalized training data. This is repeated for 4 different kinds of edges: sharp edge, square edge, large curve and small curve (Fig. 3) to reinforce the validation. Orientation classification errors for each stimulus are described in Table I below, for both the original and the generalized data.

Note that the sharp edge has a 180° symmetry, therefore to make the edge orientations for this stimulus distinguishable over a full 360° range, we shift the sensor to be 1 mm off-centre. Note also the very low errors in the first 30° for generalised data (Fig. 7) correspond to real training data from which data has

TABLE I
ORIENTATION IDENTIFICATION ERRORS FOR REAL AND GENERALISED DATA FOR ALL STIMULI

Stimulus	Real data error	Generalised data error
Sharp edge	$\bar{e}_\theta = 1.5^\circ$	$\bar{e}_\theta = 2.6^\circ$
Square edge	$\bar{e}_\theta = 1.6^\circ$	$\bar{e}_\theta = 2.8^\circ$
Large curve	$\bar{e}_\theta = 1.5^\circ$	$\bar{e}_\theta = 3.1^\circ$
Small curve	$\bar{e}_\theta = 1.6^\circ$	$\bar{e}_\theta = 3.2^\circ$

been generalized, therefore mean errors have been calculated ignoring this initial range.

Errors in orientation are thus moderately increased for the generalized data ($e_\theta = 2.8^\circ$ average over all stimuli cf $e_\theta = 1.55^\circ$ for real data), but remain within a reasonable range based on the orientation being sampled every 1° . We argue that this notable increase in errors of edge orientation classification is countered by the strong 12-fold reduction in training data gathering times (from 10 minutes 24 seconds to 52 seconds) from this tactile feature generalization method, making it attractive for tactile perception tasks, particularly in cases in which multi-dimensional training is required. In the next section we extend this validation for the large curve stimulus with a 2-dimensional dataset comprising both edge orientation and location across the edge.

C. Two-Dimensional Validation of Tactile Generalization Method

Here we analyse data representing orientations and locations across the large curve stimulus in order to further test the generalization method. Data is gathered for 72 angles over the 360° range in 5° increments, and a separate dataset is generalized as described in Section III-B3 from the first 6 angles. At each angle, 20 locations x_1 are considered across the edge, in 1 mm intervals, such that the full range of locations is centred on the edge. This 2-dimensional dataset is analysed (Fig. 8, top panels)

Fig. 7. Orientation classification accuracy with angle for the different edge stimuli. Average errors over all orientations are specified in the legend of each graph, for both real and generalized data.

Fig. 8. Orientation identification accuracy with location and orientation (top panels). Orientation error averages over all orientations are shown for each location in the bottom two panels. Panels A and B display results for real data and panels C and D for generalized data.

Fig. 9. Contour following for the large curve (panels A, B) and small curve (panels C, D) objects. For each object, results for the real dataset are displayed on the left panel (A, C) and generalized data on the right panel (B, D).

and the errors are averaged for each location over all orientations (Fig. 8, bottom panels).

We find only a slight increase in orientation errors ($e_\theta(x) \approx 4^\circ$ cf $e_\theta(x) \approx 2^\circ$) for the generalized set over the middle of the location range (6-13 mm). At the lower and higher end of the range (1-5 mm, 13-19 mm), performance is notably worse in the generalized case, however these areas corresponds to the sensor not being in contact with the edge or contacting the flat section of the stimulus and thus are ambiguous by nature. For this task, training data gathering time falls to 1 minute 37 seconds from 19 minutes 25 seconds for the generalized data.

D. Application: Contour Following

After demonstrating our generalization method using offline data, we seek to validate its performance on a real task. We apply the method to contour following around 2 circular disk (the large curve and small curve stimuli).

The 2-dimensional large curve dataset analysed previously is used in conjunction with an active perception control policy described in Section III-D to perform contour following on the large curve stimulus. Both real data (Fig. 9, left) and generalized data (Fig. 9, right) successfully follow the shape contour. This procedure is repeated for the small curve stimulus, and average error in orientation classification is used as a performance metric for the task (see Table II).

We note that although perception remains accurate enough to complete the task in all four cases, validating the use of our generalization method, it is less accurate in orientation identification around the 9 o'clock mark on the large disk (Fig. 9), leading to a larger error for the generalized data/large curve case. This decrease in performance corresponds to the 60–120°

TABLE II
AVERAGE ORIENTATION IDENTIFICATION ERRORS OVER THE CONTOUR FOLLOWING TASK FOR BOTH DISK SIZES

Object	Real data error	Generalised data error
Large curve	$\bar{e}_\theta = 2.1^\circ$	$\bar{e}_\theta = 6.4^\circ$
Small curve	$\bar{e}_\theta = 1^\circ$	$\bar{e}_\theta = 0.8^\circ$

range of orientation, which can be seen to have larger errors in the validation phase (Fig. 8, top right). This could be the result of an artefact of the stimulus or experimental setup.

We have included a supplementary MPEG video clip (9.03 MB in size) which shows this task being performed.

V. DISCUSSION

This study applied a method of training reduction for tactile perception to an optical tactile sensor with a symmetrical morphology, and compared the performance of the sensor with real and generalized training data. The generalization entails a rotation of a subset of tactile images to generate the full dataset of all edge orientations. Results indicate that although classification errors are moderately increased using our method, the 12-fold gain in time to gather training data makes the method attractive for use in tactile perception, particularly in the case of multi-dimensional perceptual data. This was demonstrated with a contour following task of two circular disks of different diameters, which involved both orientation and location identification.

The method proposed here could be applied to other more complex tasks, such as slip detection, since it enables training over a small range of slip orientations with similar savings in data gathering time as here. Other applications could include generalizing over the location that a stimulus contacts the tactile sensor, of particular use for flat tactile sensors or tactile skins such as roboskin [21] with its hexagonal symmetry. We expect our method to present a trade-off between gains in time to gather training data and precision in a given task. This trade-off could be explored to obtain a more accurate understanding of the optimal parameter values (number of symmetries, number of taxels..) to use for any given application of the method.

The generalization method proposed is contingent on sensors being designed in a symmetrical way, and considers the shape of the sensor to implement a form of low-level perceptual learning through feature generalization. Applications of our method to robotic tactile skins could represent a straightforward way to generalize both across tactile stimulus orientation and location of the contact on the robot (using translational symmetry). This could have implications in fields such as human-robot interaction and service robotics, improving autonomy by allowing robots to quickly and flexibly adapt to their environment.

Applying the generalization method to a tactile skin could also be a first step towards mimicking tactile stimuli generalization across the skin's surface. This has been shown to be performed in humans, with an interval discrimination task

generalizing across the hand's surface [22] and pressure and roughness discrimination being learned in adjacent and homologous fingers [23], suggesting that low-level sensory generalization could be an important component of perceptual learning.

The method proposed here for tactile stimuli generalization strongly reduces training data gathering time, and can effectively sample multi-dimensional tactile data. This method could be applied to tasks such as slip detection or surface classification, as well as more complex tasks including manipulation in higher dimensions. The method's data gathering time savings become more pronounced for these more complex tasks, and could make certain tasks tractable where long training phases would otherwise prove intractable.

VI. CONCLUSION

This study demonstrates a new method to generalize tactile data by exploiting sensor symmetries. Orientation classification performance was only moderately affected by the use of generalized rather than physically gathered training data, and contour following was successfully demonstrated with generalized data. This tactile feature generalization method can thus be applied to significantly reduce training data gathering time in tactile manipulation tasks.

ACKNOWLEDGMENT

The authors would like to thank N. Pestell, M. E. Giannaccini, and K. Aquilina for their help. The data used in this letter is available at <http://doi.org/bzpr>.

REFERENCES

- [1] R. Johansson and R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks," *Nature Rev. Neuroscience*, vol. 10, no. 5, pp. 345–359, 2009.
- [2] A. Russell, "Object recognition by a smart tactile sensor," in *Proc. Australian Conf. Robot. Autom.*, 2000, pp. 93–98.
- [3] G. Heidemann and M. Schopfer, "Dynamic tactile sensing for object identification," in *Proc. Int. Conf. Robot. Autom.*, IEEE, 2004, vol. 1, pp. 813–818.
- [4] Z. Pezzementi, E. Plaku, C. Reyda, and G. Hager, "Tactile-object recognition from appearance information," *IEEE Trans. Robot.*, vol. 27, no. 3, pp. 473–487, Jun. 2011.
- [5] A. Khasnobish *et al.*, "Object-shape recognition from tactile images using a feed-forward neural network," in *Proc. Int. Joint Conf. Neural Netw.*, IEEE, 2012, pp. 1–8.
- [6] N. Lepora, U. Martinez-Hernandez, and T. Prescott, "Active touch for robust perception under position uncertainty," in *Proc. IEEE Int. Conf. Robot. Autom.*, IEEE, 2013, pp. 3020–3025.
- [7] N. Lepora and B. Ward-Cherrier, "Superresolution with an optical tactile sensor," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2015.
- [8] C. Chorley, C. Melhuish, T. Pipe, and J. Rossiter, "Development of a tactile sensor based on biologically inspired edge encoding," in *Proc. Int. Conf. Advanced Robot.*, IEEE, 2009, pp. 1–6.
- [9] D. Lowe, "Object recognition from local scale-invariant features," in *Proc. Int. Conf. Comput. Vision*, IEEE, 1999, vol. 2, pp. 1150–1157.
- [10] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (SURF)," *Comput. Vision Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008.
- [11] R. Li and E. Adelson, "Sensing and recognizing surface textures using a GelSight sensor," in *Proc. IEEE Conf. Comput. Vision Pattern Recog.*, 2013, pp. 1241–1247.
- [12] T. Assaf, C. Roke, J. Rossiter, T. Pipe, and C. Melhuish, "Seeing by touch: Evaluation of a soft biologically-inspired artificial fingertip in real-time active touch," *Sensors*, vol. 14, no. 2, pp. 2561–2577, 2014.
- [13] A. Schneider, J. Sturm, C. Stachniss, M. Reiser, H. Burkhardt, and W. Burgard, "Object identification with tactile sensors using bag-of-features," in *Proc. Int. Conf. Intell. Robots Syst.*, 2009, pp. 243–248.
- [14] S. Luo, W. Mou, M. Li, K. Althoefer, and H. Liu, "Rotation and translation invariant object recognition with a tactile sensor," in *Proc. SENSORS*, IEEE, 2014, pp. 1030–1033.
- [15] B. Winstone, G. Griffiths, C. Melhuish, T. Pipe, and J. Rossiter, "TacTip-tactile fingertip device, challenges in reduction of size to ready for robot hand integration," in *Proc. Int. Conf. Robot. Biomimetics*, IEEE, 2012, pp. 160–166.
- [16] B. Ward-Cherrier, L. Cramphorn, and N. Lepora, "Tactile manipulation with a TacThumb integrated on the open-hand M2 gripper," *Robot. Autom. Lett.*, vol. 1, no. 1, pp. 169–175, 2016.
- [17] N. Lepora, C. Sullivan, B. Mitchinson, M. Pearson, K. Gurney, and T. Prescott, "Brain-inspired Bayesian perception for biomimetic robot touch," in *Proc. Int. Conf. Robot. Autom.*, 2012, pp. 5111–5116.
- [18] N. Lepora, U. Martinez-Hernandez, and T. Prescott, "Active Bayesian perception for simultaneous object localization and identification," *Proc. Robot.: Sci. Syst.*, 2013.
- [19] N. Lepora, U. Martinez-Hernandez, M. Evans, L. Natale, G. Metta, and T. Prescott, "Tactile superresolution and biomimetic hyperacuity," *IEEE Trans. Robot.*, vol. 31, no. 3, pp. 605–618, 2015.
- [20] 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 *Proc. World Haptics Conf.*, 2013, pp. 341–346.
- [21] A. Billard *et al.*, "The ROBOSKIN project: Challenges and results," in *Romansy 19—Robot Design, Dynamics and Control*, New York, NY, USA: Springer, 2013, pp. 351–358.
- [22] S. Nagarajan, D. Blake, B. Wright, N. Byl, and M. Merzenich, "Practice-related improvements in somatosensory interval discrimination are temporally specific but generalize across skin location, hemisphere, and modality," *J. Neuroscience*, vol. 18, no. 4, pp. 1559–1570, 1998.
- [23] J. Harris, I. Harris, and M. Diamond, "The topography of tactile learning in humans," *J. Neuroscience*, vol. 21, no. 3, pp. 1056–1061, 2001.