

Learning to Live Life on the Edge: Online Learning for Data-Efficient Tactile Contour Following

Elizabeth A. Stone¹, Nathan F. Lepora² and David A.W. Barton³

Abstract—Tactile sensing has been used for a variety of robotic exploration and manipulation tasks but a common constraint is a requirement for a large amount of training data. This paper addresses the issue of data-efficiency by proposing a novel method for online learning based on a Gaussian Process Latent Variable Model (GP-LVM), whereby the robot learns from tactile data whilst performing a contour following task thus enabling generalisation to a wide variety of tactile stimuli. The results show that contour following is successful with comparatively little data and is robust to novel stimuli. This work highlights that even with a simple learning architecture there are significant advantages to be gained in efficient and robust task performance by using latent variable models and online learning for tactile sensing tasks. This paves the way for a new generation of robust, fast, and data-efficient tactile systems.

I. INTRODUCTION

Often the success of tactile sensing methods is judged by the accuracy or the speed with which a task can be completed, with little or no focus on the data efficiency of the methods. The requirement to collect large training sets for every new task or new sensor means these methods can be impractical and time consuming. This paper focuses on a data-efficient method of tactile sensing based on a latent variable model.

Current methods for tactile sensing try to build an explicit sensor model in a separate phase preceding the testing of the model on a task, therefore needing to train across a large data space to be representative. If any configurations are missed, or a variable is simply not accounted for, then the robot is likely to fail at its task when unexpectedly encountering these novel states during testing (e.g. in [1] changes in height were not accounted for, leading to a failure in contour following in some cases). Further problems arise when slow data collection is combined with the well-known curse of dimensionality [2]: the more dimensions included in the models the less representative the data becomes of all states leading to the need for considerably more data to compensate and therefore a considerably increased time collecting data.

To overcome these problems of data efficiency this work proposes the use of a latent variable model, which needs only a single labelled data point for the initial model and

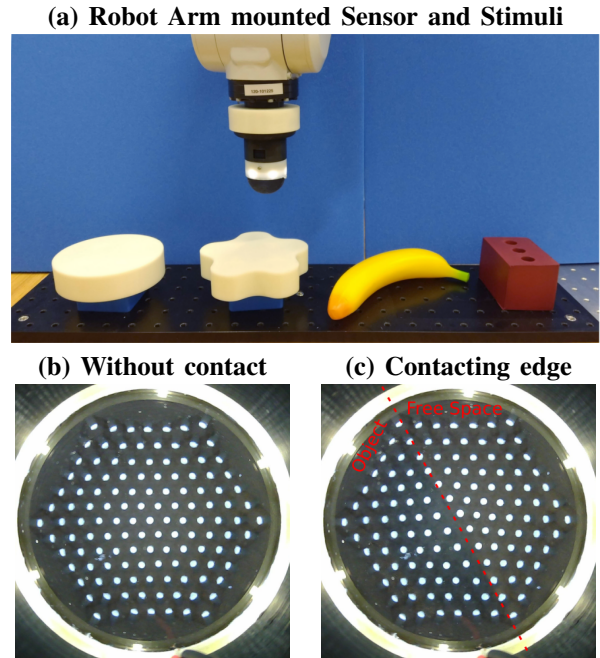


Fig. 1: (a) Experimental setup showing TacTip mounted on robotic arm with circle, flower, banana and brick shaped stimuli. (b) View from TacTip internal camera showing pins in neutral position (c) View from TacTip internal camera showing pins when contacting an object edge (location indicated by the dotted line).

uses an intelligent data collection policy to collect sufficient data throughout a task to autonomously build a sufficiently representative model. In this case the data collected will always be relevant to the task at hand, to some extent resolving the curse of dimensionality as it is no longer necessary to collect data over all possible dimensions of variation.

With online learning, each new reading is considered for addition to the model to improve it, unlike in offline methods where the model is fixed beforehand and new data is ignored, despite the new data being representative of, and most relevant to, the task at hand. As such, this work demonstrates: 1) the novel use of data-efficient methods for tactile sensing, 2) the use of online learning in the control of a tactile sensor, and 3) the use of latent variable models in dimensionality reduction of tactile data.

II. BACKGROUND AND RELATED WORK

The TacTip [3], as used in this paper, is a biomimetic tactile sensor, consisting of a soft, deformable rubber-like gel-filled skin with 127 raised pins on the internal surface, the tips of which are tracked in 2D using an internal camera. Since its development in 2009 the TacTip has been applied

¹ES is a PhD student at the EPSRC Centre for Doctoral Training in Future Autonomous and Robotic Systems (FARSCOPE) at the Bristol Robotics Laboratory. lizzie.stone@brl.ac.uk

²NL is with the Dept. of Engineering Maths and Bristol Robotics Laboratory, University of Bristol, UK, and is supported by the Leverhulme Trust (RL-2016-39). n.lepora@bristol.ac.uk

³DB is with the Dept. of Engineering Maths, University of Bristol, UK. david.barton@bristol.ac.uk

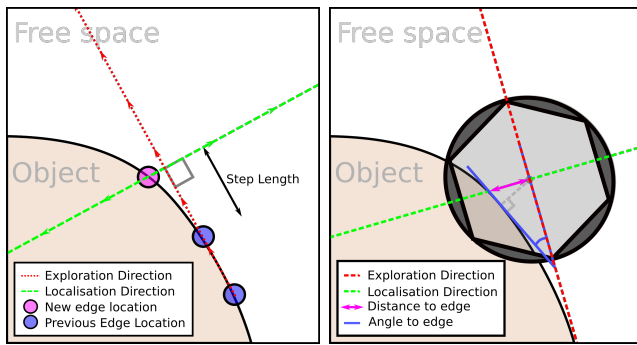


Fig. 2: Left: How the direction of Exploration is found based on previously found edge locations. The aim of Localisation is to find the new edge location. Right: Close up of TacTip location at end of Exploration phase, where “distance to edge” (pink) must be found during Localisation, irrespective of the orientation of the sensor (blue).

to many different tasks. One of the biggest challenges is decoding the meaning of the rich sensor data, which has 254 dimensions (2D displacement of each pin), into a simple enough form to be usable in control tasks. So far all of the methods applied to this challenge are offline methods: histogram methods [1], Principle Component Analysis (PCA) [4] and deep CNNs [5]. These methods are successful to varying degrees, with deep CNNs greatly outperforming histogram methods, but all rely on the collection of large data sets for training; specifically a set of at least 180 taps for histograms and 2000 contacts for deep CNNs.

These methods have then been used to perform a variety of control tasks, including object manipulation [6], [7], slip detection [8], [9], [10], surface following [11], [12] and contour following [1], [5], [13], [14]. Contour following is the tracking of a prominent feature, such as a well-defined object edge or ridge, often in a complete loop to discern the shape of an unknown object. This has been considered as an example task many times with sensors such as the TacTip [1], [5], [14] and iCub fingertip [13]. The TacTip is particularly well suited to contour following due to its ability to deform around edges. Some of the biggest challenges in contour following include the objects not being uniform in shape or height, objects being made of varying materials that deform in different ways when touched, and every day objects being highly irregular. Contour following is often done with 2D contours as this is the simplest case, with displacement in the third dimension kept constant.

The literature concerning the combination of online learning with tactile sensing is sparse, which highlights the opportunities in this area. Previous work has focused on the use of tactile hands in object classification, using Gaussian Processes (GPs) [15], [16] and Passive Aggressive online learning [17], and object localisation, using particle filters [18].

III. METHODS

A. Experimental Setup

1) *Hardware*: A TacTip is attached to an IRB120 ABB 6-DOF robotic arm, with stimuli rigidly attached to a plate

in front of the robot. This setup replicates that in [14].

2) *Software*: The main algorithms are implemented in MATLAB, which is run on Windows. The robotic arm is controlled through a separate code base which also handles the detection of pin positions using computer vision. Individual pin locations are detected using OpenCV as described in [19], and the average maximum pin displacement during a tap is used as the taxel data.

B. Algorithms

1) *Task Setup*: The task of contour following when restricted to 2D displacements, as in the case of following the edge of a planar object, is broken down into two phases. The first is Exploration which gives a prediction of the edge having moved further around the object. The second is Localisation which, having moved according to the prediction in Exploration, corrects the error in this movement (i.e. estimates the displacement from the edge perpendicular to the direction moved during Exploration). The contour following is achieved by repeating these two phases as shown in Algorithm 1.

Algorithm 1: High-level Contour Following

1. Bootstrap;
 2. **while** *distance to start location* > *1 step-length* **do**
 3. Exploration;
 4. Localisation;
 5. **end**
-

This paper focuses on optimisation of the Localisation phase and therefore only a naive Exploration phase is implemented: linear extrapolation of the two previously found edge locations, as shown in Figure 2 and Algorithm 2.

Algorithm 2: Exploration

1. Extrapolate in a straight line between the locations of the previous two least dissimilar points (call this the Exploration line);
 2. Rotate the sensor to maintain alignment with Exploration line;
 3. Move along Exploration line by step length;
-

2) *GP-LVM*: To efficiently achieve Localisation, accurate estimation of the displacement of the sensor, r , from the edge along any axis, irrespective of any other variables, is required given only the taxel pattern, y , from a tap. Hence a Gaussian Process Latent Variable Model (GP-LVM) [20] is used to predict the displacement of the edge from the sensor irrespective of any other variables. The orientation of the sensor is not estimated directly, but a proxy for angle and all other variations (e.g. edge height, compliance or sharpness), termed ϕ is inferred as part of the method; this variable is the latent variable or unlabelled variable in the problem. The input to the GP-LVM is thus (r, ϕ) and the output is y .

Algorithm 3: Localisation

1. Tap and estimate the distance to the edge, along the Localisation line (perpendicular to the Exploration line), using GP-LVM;
 2. Move the estimated distance along the Localisation line;
 3. Tap again and estimate the new distance to the edge;
 4. **if** distance to edge > tolerance (i.e. 2mm) **then**
 5. Collect a series of evenly spaced taps along the Localisation line;
 6. Find displacement of taps from the edge using the dissimilarity measure;
 7. Optimise ϕ of taps;
 8. Add this labelled data to the model;
 9. **end**
-

A GP-LVM is similar to a Gaussian Process (GP) [21] in that it is built on the assumption that outputs (observations) follow a multi-variate Gaussian distribution. It differs in that a GP is used to calculate a novel output, y (i.e. taxel pattern), given a novel input, $x := (r, \phi)$ (i.e. displacement and orientation), whereas a GP-LVM does the opposite, inferring the input (the latent variables) given a novel output. This is achieved via optimisation of the $D \times n$ input using a set of known input-output pairs (the finding of such pairs is described later). The objective function to optimise is given by

$$L = -\frac{D}{2} \text{tr} \left(K^{-1} \frac{1}{D} yy^T \right) - \frac{D}{2} \log |K| - \frac{Dn}{2} \log 2\pi \quad (1)$$

which is the log marginal likelihood [20] for GP-LVMs. Here K is the covariance matrix where the covariance function, k , is the squared exponential [21]

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left[- \left(\frac{(r - r')^2}{2l_r^2} + \frac{(\phi - \phi')^2}{2l_\phi^2} \right) \right]. \quad (2)$$

This depends on the hyperparameters $\theta := (\sigma_f, l_r, l_\phi)$ — the variance of the (noise free) signal, characteristic length-scale of displacement and characteristic length-scale of ϕ respectively — which must be optimised for this problem. The signal noise is represented by another parameter, σ_n , in K and must also be tuned to the problem, but this value is assumed problem independent for a given setup as it represents the system noise (e.g. sensor noise and robot positioning noise). It was calculated from representative data beforehand to be $\sigma_n = 1.14$.

Note that the value of (1) is dependent on both the inputs, x , and the outputs, y . The GP-LVM can then therefore be used to infer the novel input x given a novel output y .

The model consists of both optimised hyperparameters θ and known input-output pairs (x, y) representative of the data; when data is collected during testing, the set of input-output pairs is augmented. Hence, the accuracy of the model relies on the accuracy of the automatic labelling of new data. In the case of r , the true (relative) displacements between

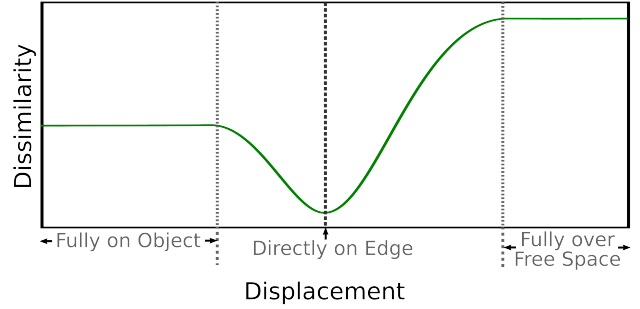


Fig. 3: Diagram of how dissimilarity should change with displacement from an edge, given that the reference tap was taken directly on the edge and assuming the object is flat topped with uniform height.

taps on the same Localisation line are known, so the only inaccuracy in r occurs in identifying the offset from the edge. The true value of ϕ is not known (being the representation of all variation not caused by r), however it is assumed that ϕ is constant along any given Localisation line and all variation in the output along that line is driven by r .

As the optimisation of the log marginal likelihood requires calculating the inverse of K , the complexity is $\mathcal{O}(n^3)$. Therefore the quantity of data in the model is kept to a minimum by not adding every tap taken during an experiment. Instead, only taps collected according to the following data collection policies are added.

3) *Dissimilarity*: In the beginning only a single labelled data-point is given, defining the desired state of the sensor during contact with an edge (named the “reference tap”). A way of labelling new data collected at unknown displacements and orientations is therefore needed to add useful information to the model. To accomplish labelling of displacement we make use of a dissimilarity measure, namely the Euclidean distance

$$d_{Euc}(\mathbf{a}, \mathbf{b}) = \sqrt{\sum_{i=1}^n (a_i - b_i)^2} \quad (3)$$

to quantify the dissimilarity between unlabelled data and the reference tap.

Using this measure the edge of the object is identified as the point with the smallest dissimilarity with the reference tap. The Euclidean distance is scalar so the direction of the edge from a single tap cannot be found, but with multiple taps in a straight line the dissimilarity profile will have a clear minimum where the edge is encountered (Figure 3).

4) *Online Policies*: To build the GP-LVM model online a data collection policy is defined, capable of collecting relevant data autonomously. As the aim is to be efficient in terms of number of taps, new data is only collected and added to the model when the model is inaccurate. This inaccuracy was assessed by doing two taps; one at the intersection of the Exploration line and the Localisation line, and one after moving the predicted displacement. This second tap should be directly on the edge if the model is accurate, giving a new prediction of 0 mm to the edge (within a tolerance – here 2 mm was used), otherwise the model is giving contradictory

predictions and is therefore inaccurate.

Having found the model is inaccurate, new data needs to be collected and added to the model to improve the predictions, as it is assumed the current data in this region is anomalous or non-existent. The new data consists of an evenly spaced line of taps along the Localisation line between ± 10 mm from the intersect with the Exploration line.

By finding the dissimilarity for these taps, which were taken at known intervals of displacement, the dissimilarity profile is found which will have a distinct minimum at the displacement where the edge is (i.e. in the location most similar to the edge), as shown in Figure 3. The location of the minima can be estimated with greater resolution than the displacement interval through use of a separate GP trained only on the new data. Knowing the relative location of this minima enables the real displacements from the edge to be calculated for this line of taps.

The value of the latent variable ϕ for this particular Localisation line is inferred using the GP-LVM. The data was collected with the sensor at the same orientation to the edge, assuming all other variables are constant. The value of ϕ is therefore the same for the entire line of taps, meaning only a single value needs inferring. This approach shares similarities with the alignment learning method proposed in [22], however that method has been applied only in an abstract setting and does not exploit all the structure of this problem.

Obviously this procedure cannot start on its own as two previous edge points must have been collected for the initial Exploration step. An initialisation phase, termed Bootstrap, is used to collect the first two edge locations. This involves the steps outlined in Algorithm 4.

This leads to the full control policy described by Algorithms 1, 2, 3, and 4.

Algorithm 4: Bootstrap

1. Collect 3 taps in a right-angled triangle;
 2. Calculate the dissimilarity to the reference tap for all 3 taps;
 3. Calculate the direction of greatest decrease in dissimilarity;
 4. Tap along this direction collecting evenly spaced taps and label (in the same fashion as in Algorithm 3);
 5. Identify and move to edge location;
 6. Move perpendicular to direction of greatest decrease in dissimilarity by one step length;
 7. Localisation (Algorithm 3);
-

IV. RESULTS

A. Offline Testing

Offline testing was done to evaluate the methods on a representative dataset. Having been given only 5 lines of taps to build the model, simulating data collection at 5 different angles, GP-LVM was able to predict the displacement and other variations, ϕ , of a dataset of taps taken over a range of

TABLE I: Table of offline results. An anomaly is defined as a r prediction with magnitude greater than 12 mm (as the maximum should be 10 mm), and for ϕ predictions of magnitude greater than 3 (maximum should be 2).

| No. of training lines | Error in r (mm) | | Error in ϕ | | No. of Excluded Anomalies | |
|-----------------------|-------------------|------|-----------------|------|---------------------------|--------|
| | Mean | SD | Mean | SD | r | ϕ |
| 1 | 1.02 | 1.56 | 1.05 | 1.29 | 7 | 3 |
| 3 | 0.48 | 0.70 | 0.17 | 0.32 | 1 | 6 |
| 5 | 0.48 | 0.62 | 0.13 | 0.21 | 2 | 2 |

displacements (-10 to +10 mm) and angles (-45° to 45°) from an edge. In this case ϕ is taken as a direct proxy of angle as no other variables change. Here ϕ is ranged between -2 and 2, corresponding to -45° and 45° respectively. The predictions are good, except at the extremes of displacement where the ϕ predictions converge towards 0; this is expected because here the sensor is fully on a flat surface or fully over free space, both of which are invariant to rotation.

Even with fewer training lines the error in displacement predictions are small, with the highest mean error being 1.02 mm when only one training line (and therefore only one orientation) was given. Even this precision is enough to complete contour following, although the standard deviation is relatively large meaning this accuracy may not always be maintained.

These results indicate that the use of a single training line should be sufficient for task completion, therefore the online methods were designed to collect only a single line during start up. When this is no longer sufficient more data is collected automatically, using the algorithm as described.

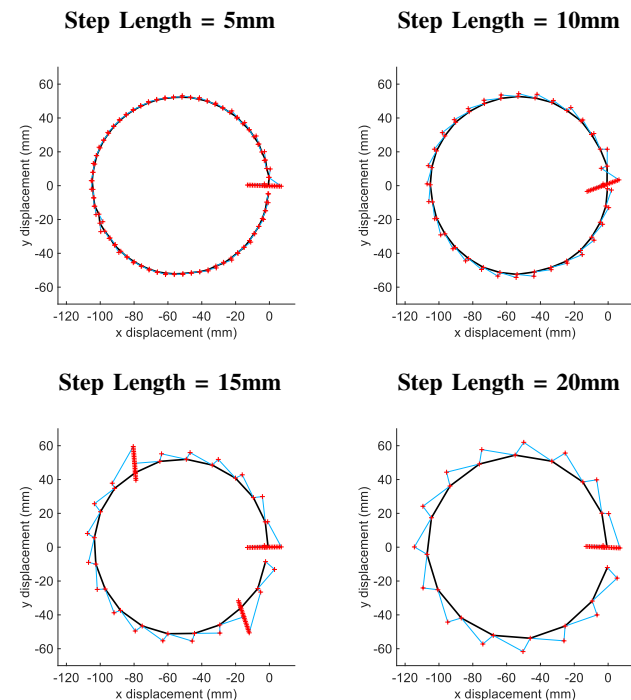


Fig. 4: Varying step length with the circle stimulus. Red crosses are locations where the robot tapped, blue lines show the movement between taps, and the black line joins the locations identified as the object edge. Start location was (0,0) and movement is anticlockwise.

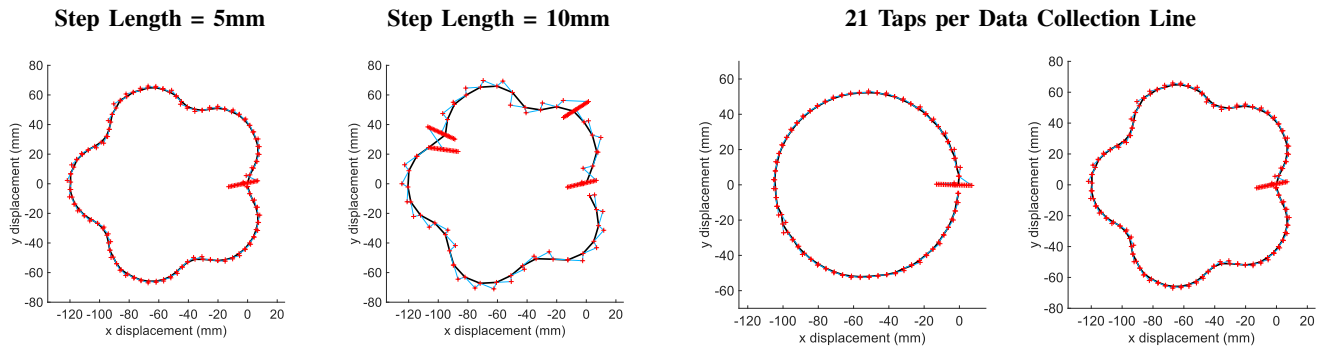


Fig. 5: Varying step length with flower stimulus. Red crosses are locations where the robot tapped, blue lines show the movement between taps, and the black line joins the locations identified as the object edge. Start location was (0,0) and movement is anticlockwise.

TABLE II: Mean and Standard Deviation (SD) of the shortest distance between actual circle and predicted circle location with varying Exploration step length.

| Step Length (mm) | Mean distance (mm) | SD (mm) |
|------------------|--------------------|---------|
| 5 | 0.72 | 0.39 |
| 10 | 0.47 | 0.42 |
| 15 | 1.33 | 0.63 |
| 20 | 1.07 | 0.53 |

B. Online Testing

The online system was tested on a task of contour following around a circle and a flower-like stimuli. The circle is the simplest object as the edge is uniform in curvature. The flower is designed to test the system on varying edge curvatures, with both positive and negative curvatures.

Given only one line of 21 taps, at 1 mm intervals, the robot was able to complete the task on both stimuli, tracking the object edge until the start location was successfully reached again. Following this success the system was pushed to its limits with the following experiments.

1) *Robustness to Larger Edge Displacements*: By increasing the step length during exploration, the greater distance the edge is likely to have deviated from the Exploration line, testing the capabilities of the Localisation (which corrects these perturbations) over a larger range. With the circle stimulus the step length was increased from 5 mm in steps of 5 mm until 20 mm. All of these trials were successful using only a single line of data except for 15 mm which needed 2 extra lines of data (automatically collected by the algorithm), as shown in Figure 4. The shape of the circle is clearly visible in all cases, despite the 20 mm case needing very few Exploration steps to complete the task. As can be seen in Table II the error in estimation of the location of the circle edge is very small, with the smallest mean distance being 0.47 mm and the greatest being 1.33 mm.

The same was attempted with the flower stimulus but owing to the tighter curves of the “petals”, the step length could not be increased as much. The results can be seen in Figure 5, where it is obvious the large changes in edge direction compared with the step length caused the system to collect more data as necessary; 4 lines of data were needed to build a sufficient model to complete the task and the flower

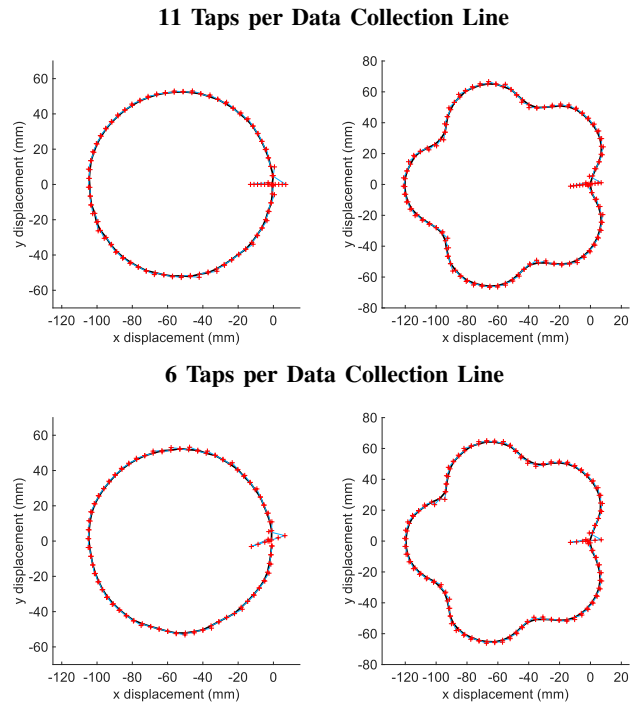


Fig. 6: Using only 6 taps per data collection line on the flower and circle stimuli. Red crosses are locations where the robot tapped, blue lines show the movement between taps, and the black line joins the locations identified as the object edge. Start location was (0,0) and movement is anticlockwise.

TABLE III: Mean and Standard Deviation of the shortest distance between actual circle and predicted circle location with varying number of taps per data collection line.

| Number of taps | Mean distance (mm) | SD (mm) |
|----------------|--------------------|---------|
| 21 | 0.72 | 0.39 |
| 11 | 0.74 | 0.42 |
| 6 | 1.10 | 0.47 |

shape is clearly visible.

2) *Robustness to Sparse Data Collection*: The main aim of this system is to reduce the quantity of data needed to complete the task. By reducing the number of samples along the localisation line when collecting tap data to add to the model, the number of taps can be significantly reduced. To test this the distance between taps was doubled from 1 mm to 2 mm to 4 mm, giving a total of 21, 11 and 6 taps per line respectively. This should be easily handled due to the use of GP regression for estimating the minimum dissimilarity, and the use of GP-LVM for estimating displacement and ϕ . In all three cases the task was successfully completed. Even with

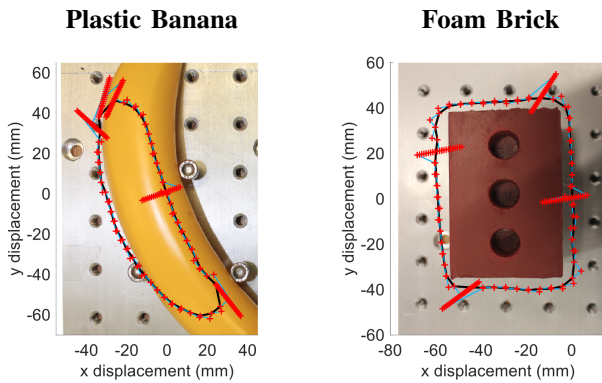


Fig. 7: Contour following on a banana and foam brick. Red crosses are locations where the robot tapped, blue lines show the movement between taps, and the black line joins the locations identified as the object edge. Start location was (0,0) and movement is anticlockwise. (Note, location of background image is approximate.)

the use of only 6 taps per line, the edge was closely tracked with only 1 line of data on both the circle and the flower, as shown in Figure 6. For the circle, it can be seen in Table III that even when the number of taps per data collection line is reduced to 6 the location error is only 1.10 mm, only 0.38 mm greater than with 21 taps.

3) *Robustness to Every-day Objects*: The final test was using every-day objects which better reflect the irregularity of objects outside a laboratory setting. The two objects chosen were the foam brick and plastic banana from the YCB dataset (items ID60 and ID11 [23] respectively). These objects test robustness to previously unencountered compliances and edge sharpness in addition to novel edge curvature.

The right-angles of the brick theoretically pose a challenge for this system due to the Exploration and Localisation lines being perpendicular; if the edge is being tracked perfectly the Exploration step will reach the corner and overshoot, tapping over free-space and attempting to find the minimum dissimilarity of the taps will result in a random location being chosen as the edge, leading the robot to add bad data to the model and to move in a random direction. The reason this is not reflected in the results (see Figure 7) is because the TacTip sensor is large and compliant, meaning that with this step size (10mm) the corners are effectively smoothed off (rather than moving starkly from on edge to off edge).

The brick also adds the novelty of being made of very compliant foam which compresses when touched. This poses a challenge as the sensor readings vary with the compliance of the object. This system is not specifically designed to handle differences in compliance, but owing to the use of ϕ as a representation of all variation not caused by displacement, this system is shown to not be affected by the compliance of the brick.

The plastic banana poses a different challenge, testing the resilience to changing edge sharpness. The sensor is started on a prominent edge of the banana, but this only runs along one side of the banana and the height of it changes as this side of the banana curves downwards at both ends. This means when following a 2D contour, maintaining the same height above the object as in this case, the edge will change from

a sharp edge (along the length of the banana) into a very smooth edge (crossing the width of the banana) which is less an edge and more of a slope. This changes the sensor readings significantly.

As can be seen in Figure 7, both the brick and banana were successfully traced. As expected, more data was needed than in the experiments with the circle and flower, but still only a maximum of 106 taps (5 lines) were used for the model.

V. DISCUSSION & CONCLUSION

This work has shown that online learning with a latent variable model and intelligent data collection policy is a data efficient and robust approach to tactile exploration with a TacTip tactile sensor.

This method shows significant improvements in either accuracy or data efficiency over the previous methods with the TacTip. The histogram methods in [1] appear to use 180 taps for training, but are not very accurate in the tracking of the edge of objects. In our methods the highest number of taps in the model was 106, which is 40% less data but with more accurate tracking of the object edges. The deep CNN in [5] shows similar accuracy at tracking the edge (within 1 mm of the edge) but used 2000 taps for training, which is 285 times more data than our most efficient case, and 19 times our least data efficient case. In addition to better data efficiency, the online method is able to cope with variations in the contour: the histogram methods do not generalise well, and while the deep CNN generalises well from its training data (a right angled edge) to everyday objects such as a banana with rounded sides, it was still not able to cope when there is no edge to follow at all (e.g. over the flat top of the banana), whereas the online method presented here is able to relearn the model to fully complete the loop.

This work has shown that even with a simple architecture there are significant advantages in terms of data efficiency and robustness by using latent variable models and online learning methods in preference to offline learning methods when applied to tactile sensing tasks. Areas for future improvement on the contour-following task considered here include the use of arc localisation curves instead of straight lines to improve robustness to tighter curvatures, the use of more sophisticated exploration algorithms to improve the initial predictions at each step, and the use of raw camera pixels instead of taxel locations to retain more information about each tap and reduce the reliance on taxel detection algorithms.

There are other opportunities in extending these methods to 3D contour or surface exploration where the collection of training data for offline learning methods is more challenging but the tasks potentially much closer to human capabilities to interact with the environment via touch. Similarly, the application to sliding motion instead of tapping will open up many possibilities, but is also a challenge to gather training data and construct a robust model, and it would be interesting to compare the performance with deep CNNs that are able to successfully generalise over sliding motion. In principle, there is no reason why the methods proposed here should not apply also to multi-fingered tactile robotic hands, opening up

the possibility of online learning for the manipulation and exploration of complex objects.

In conclusion, this work shows clear benefits to online learning with tactile sensors and lays the foundations for further research into its more widespread use.

ACKNOWLEDGMENT

We thank Kirsty Aquilina and Noor Alakhawand for all their help.

REFERENCES

- [1] N. F. Lepora, K. Aquilina, and L. Cramphorn, "Exploratory Tactile Servoing With Active Touch," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1156–1163, 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7837664/>
- [2] R. Bellman and R. Corporation, *Adaptive Control Processes: a Guided Tour: R-350*. Rand Corporation, 1961. [Online]. Available: <https://books.google.co.uk/books?id=AcJPAQAAMAAJ>
- [3] C. Chorley, C. Melhuish, T. Pipe, and J. Rossiter, "Development of a tactile sensor based on biologically inspired edge encoding," in *2009 International Conference on Advanced Robotics*. IEEE, 2009, pp. 1–6.
- [4] K. Aquilina, D. A. W. Barton, and N. F. Lepora, "Principal components of touch," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–8.
- [5] 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.
- [6] L. Cramphorn, B. Ward-Cherrier, and N. F. Lepora, "Tactile manipulation with biomimetic active touch," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 123–129.
- [7] 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," *arXiv preprint arXiv:1903.04128*, 2019.
- [8] Z. Su, K. Hausman, Y. Chebotar, A. Molchanov, G. E. Loeb, G. S. Sukhatme, and S. Schaal, "Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor," *IEEE-RAS International Conference on Humanoid Robots*, vol. 2015-Decem, pp. 297–303, 2015.
- [9] J. W. James, N. Pestell, and N. F. Lepora, "Slip detection with a biomimetic tactile sensor," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3340–3346, 2018.
- [10] F. Veiga, H. Van Hoof, J. Peters, and T. Hermans, "Stabilizing novel objects by learning to predict tactile slip," *IEEE International Conference on Intelligent Robots and Systems*, vol. 2015-Decem, pp. 5065–5072, 2015.
- [11] Z. Yi, R. Calandra, F. Veiga, H. Van Hoof, T. Hermans, Y. Zhang, and J. Peters, "Active tactile object exploration with Gaussian processes," *IEEE International Conference on Intelligent Robots and Systems*, vol. 2016-Novem, pp. 4925–4930, 2016.
- [12] G. Sutanto, N. Ratliff, B. Sundaralingam, Y. Chebotar, Z. Su, A. Handa, and D. Fox, "Learning latent space dynamics for tactile servoing," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3622–3628. [Online]. Available: <http://arxiv.org/abs/1811.03704>
- [13] U. Martinez-Hernandez, G. Metta, T. J. Dodd, T. J. Prescott, L. Natale, and N. F. Lepora, "Active contour following to explore object shape with robot touch," in *2013 World Haptics Conference (WHC)*. IEEE, 2013, pp. 341–346.
- [14] K. Aquilina, D. A. W. Barton, and N. F. Lepora, "Shear-invariant Sliding Contact Perception with a Soft Tactile Sensor," in *to be published in ICRA 2019*, 2019.
- [15] H. Soh and Y. Demiris, "Incrementally learning objects by touch: Online discriminative and generative models for tactile-based recognition," *IEEE transactions on haptics*, vol. 7, no. 4, pp. 512–525, 2014.
- [16] H. Soh, Y. Su, and Y. Demiris, "Online spatio-temporal Gaussian process experts with application to tactile classification," *IEEE International Conference on Intelligent Robots and Systems*, pp. 4489–4496, 2012.
- [17] M. Kaboli, A. T. De La Rosa, R. Walker, and G. Cheng, "In-hand object recognition via texture properties with robotic hands, artificial skin, and novel tactile descriptors," in *IEEE-RAS International Conference on Humanoid Robots*, vol. 2015-Decem. IEEE, 2015, pp. 1155–1160.
- [18] M. Chalon, J. Reinecke, and M. Pfanne, "Online in-hand object localization," *IEEE International Conference on Intelligent Robots and Systems*, pp. 2977–2984, 2013.
- [19] 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, p. soro.2017.0052, 2018. [Online]. Available: <http://online.liebertpub.com/doi/10.1089/soro.2017.0052>
- [20] N. D. Lawrence, "Gaussian process latent variable models for visualisation of high dimensional data," in *Advances in neural information processing systems*, 2004, pp. 329–336.
- [21] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*. MIT Press, 2006. [Online]. Available: <http://www.gaussianprocess.org/gpml/chapters/RW.pdf>
- [22] I. Kazlauskaitė, C. H. Ek, and N. Campbell, "Gaussian Process Latent Variable Alignment Learning," in *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019, pp. 748–757. [Online]. Available: <http://arxiv.org/abs/1803.02603>
- [23] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The YCB object and Model set: Towards common benchmarks for manipulation research," *Proceedings of the 17th International Conference on Advanced Robotics, ICAR 2015*, pp. 510–517, 2015.