Proceedings Paper:

Martinez-Hernandez, U., Dodd, T.J. orcid.org/0000-0001-6820-4526, Natale, L. et al. (3 more authors) (2013) Active contour following to explore object shape with robot touch. In: World Haptics Conference (WHC), 2013. 2013 World Haptics Conference (WHC), April 14 - 18, 2013, Daejeon, Korea. IEEE , pp. 341-346. ISBN 978-1-4799-0087-9

https://doi.org/10.1109/WHC.2013.6548432

Active contour following to explore object shape with robot touch

Uriel Martinez-Hernandez*
University of Sheffield
Sheffield, UK

Tony J. Dodd†
University of Sheffield
Sheffield, UK

Lorenzo Natale‡
Italian Institute of Technology
Genoa, Italy

Giorgio Metta§
Italian Institute of Technology
Genoa, Italy

Tony J. Prescott¶
University of Sheffield
Sheffield, UK

Nathan F. Lepora‖
University of Sheffield
Sheffield, UK

Abstract
In this work, we present an active tactile perception approach for contour following based on a probabilistic framework. Tactile data were collected using a biomimetic fingertip sensor. We propose a control architecture that implements a perception-action cycle for the exploratory procedure, which allows the fingertip to react to tactile contact whilst regulating the applied contact force. In addition, the fingertip is actively repositioned to an optimal position to ensure accurate perception. The method is trained off-line and then the testing performed on-line based on contour following around several different test shapes. We then implement object recognition based on the extracted shapes. Our active approach is compared with a passive approach, demonstrating that active perception is necessary for successful contour following and hence shape recognition.

Keywords: perception-action, active tactile perception, tactile exploration, naive Bayes classifier, contour following, robotic finger.

1 Introduction

Despite rapid advances in robotics, most robots are still designed to work in controlled and structured environments with minimal human interaction. The challenge is to design robots that can perform appropriately in unstructured, human-centered environments alongside human colleagues. In order to achieve this aim, robots need to make best use of their senses, perceiving their environment whilst managing uncertainty similarly to how humans make perceptual decisions. The sense of touch plays an important role in our ability to safely explore and interact with the world. Humans rely on touch in situations where vision is partial or occluded (for instance, in darkness), through appropriate finger or hand movements such as palpation and sliding [1, 2, 3, 4]. Therefore, in this study, we are interested in developing methods for robots that make best use of their sense of touch, to enable them to interact tactually with their environment in a robust and accurate way.

In this work, we present an approach for active tactile perception using a biomimetic tactile sensor based on the human fingertip [5, 6]. Shape and position sensing properties of this fingertip have been characterised recently, and shown to naturally give perceptual hyperacuity [7]. Here we concentrate on a contour following task to give an example of a type of procedure that requires a perception-action cycle to complete. The proposed active control method has two aspects: (1) reacting to contact detection to ensure safe exploration and (2) repositioning the fingertip sensor in an optimal position to ensuring good perception. To perform this task successfully, the methods for perception will need to manage uncertainty appropriately, for which we use a probabilistic method for robot perception [7, 8] that relates to the neuroscience of human and animal decision making [9, 10]. The motivation for the current study is that a perception-action cycle based on this approach, which involves making perceptual decisions to execute appropriate actions, has not been examined or implemented in a robot.

*e-mail: uriel.martinez@sheffield.ac.uk
†e-mail: t.j.dodd@sheffield.ac.uk
‡e-mail: lorenzo.natale@iit.it
§e-mail: giorgio.metta@iit.it
¶e-mail: t.j.prescott@sheffield.ac.uk
‖e-mail: n.lepora@sheffield.ac.uk

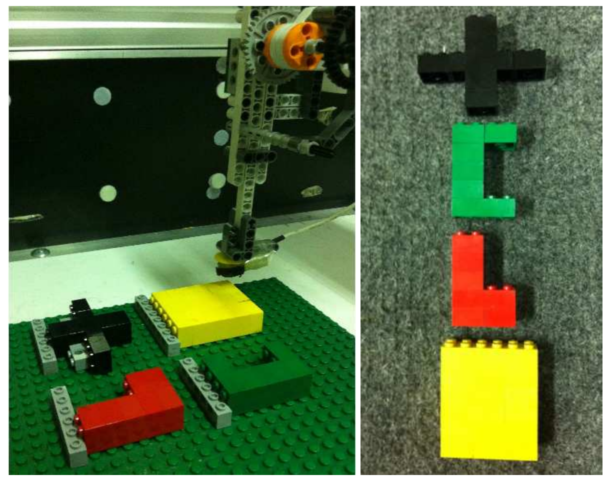

(a) iCub fingertip mounted on the exploratory robot platform

(b) Four test objects used for contour following

Figure 1: The contour following method is performed by (a) the iCub fingertip mounted on a exploratory platform composed of a Cartesian robot and Mindstorms NXT-robot. (b) The Four different objects used for the on-line testing (a cross-shape, a C-shape, an L-shape and a rectangle).

Although some techniques exist that allow robots to follow contours, none have used the probabilistic active perception approach developed here. For instance, in [11, 12, 13], tactile images were used for edge recognition, applying image processing techniques such as median filter, Hu transform and geometric moments. Meanwhile, another approach for contour following was based on contact force and orientation with respect to the surface [14, 15].

Here we propose a perception-action architecture for contour following based on the tactile feedback from a biomimetic fingertip. The tactile information, collected by tapping against the object, is used to perceive edge orientation using an active perception strategy that optimizes the position of the fingertip. To implement active tactile exploration, the control architecture moves the fingertip to a location that should improve the perception of the edge orientation. In addition, the tapping motion of the fingertip reacts actively to contact detection, to control the contact force and protect the fingertip. For these reasons, we collected tactile data with taps whose contact force is actively controlled. Other approaches have used types of contact, for example grasping to perceive object shape [16, 17].

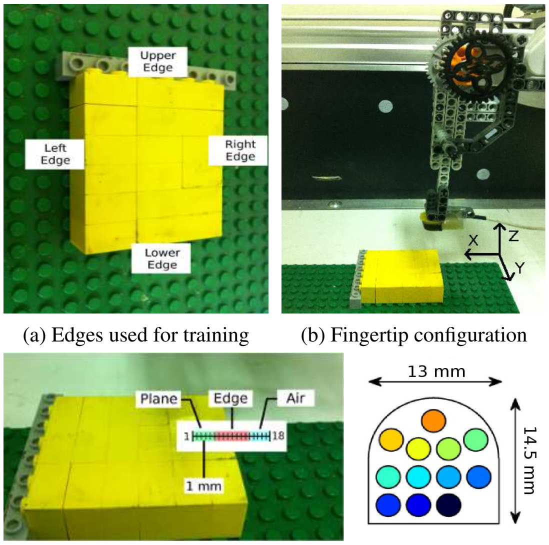

(a) Edges used for training (b) Fingertip configuration

(c) Position classes (d) Taxel layout

Figure 2: (a) The four edges of a rectangle were used for training the active perception approach. (b) Configuration of the iCub fingertip with its orientation fixed relative to the object. (c) Position classes used for classification. These positions are divided into three regions (plane, edge, air). (d) Fingertip shape and taxel layout.

Our approach for active tactile perception was first validated off-line and then tested on-line to track the edges of four different objects (Figure 1b). The extracted shapes of these objects could then be discriminated using a histogram method. Finally, the performance of the contour following with active and passive perception were compared. These results demonstrate that only by using active perception could a successful contour following behaviour and hence shape classification be achieved.

2 Methods

2.1 Tactile fingertip sensor

The tactile fingertip sensor used in this work was taken from the tactile sensory system of the iCub humanoid. This humanoid has a tactile sensory system composed of artificial skin in the palms, fingertips, forearms and torso [5, 6]. This sense of touch allows the iCub humanoid to perceive and interact with its environment.

The dimensions and rounded shape of the iCub fingertip are similar to the human fingertip (see Figure 2d). The fingertip uses a capacitive sensor technology, with contact pressure measured from the capacitance of 12 contact pads called taxels (tactile elements). These taxels perform akin to the mechanoreceptors in the human fingertips, which respond to mechanical stimuli such as pressure. In the iCub fingertip, the taxel separation is about 4 mm and the measurements are sampled at 50 Hz. The data from the taxels are digitised locally in the fingertip with a capacitive-to-digital converter. The output of the digitised data is in the range of 0 and 255, and is sent to a central computer to be processed. Pre-filtering operations such as sensory noise reduction and thermal drift compensation are managed with a supplied 'drift compensation' module [18].

In previous work, the perceptual abilities of the iCub fingertip sensor have been systematically explored and shown to be very good. In particular, position perception at sub-millimeter accuracy has been demonstrated, giving the first example of hyperacuity in robot touch [7].

2.2 Integrated robot platform

An integrated robot platform was constructed to enable shape exploration with the iCub fingertip. A Cartesian robot (Yamaha XY-X series) and a Mindstorms NXT Lego robot were integrated to achieve movements in the x-, y- and z-axes (Figure 2). The Cartesian robot allows precise positioning movements along the x- and y-axes to an accuracy of about 20 μm, and has been previously utilized for characterising the properties of tactile sensors, for example whiskers [19, 20]. Given the physical characteristics of the tactile sensor (in that it could be damaged by rubbing repeatedly along a surface), a tapping procedure was chosen for gathering sensory data. The taps were performed along the z-axis, and executed with a custom-built NXT Lego robot system (Figures 1 and 2). This NXT robot was responsible for controlling and executing periodic taps, with the iCub fingertip mounted at an appropriate angle to contact a horizontal surface.

The control and synchronisation of this robot was based on using tactile feedback, with the movements in the x- and y-axes and the taps along the z-axis coordinated by the signal from the iCub fingertip. Modules for communication and control for both robots (Cartesian robots and NXT Lego robot) and also for the iCub fingertip were developed in the C/C++ language and the standard library for the iCub platform, YARP [21].

2.3 Active positioning for contour following

In this work, contour following is achieved using a perception-action cycle in which a tap of the iCub fingertip is processed and classified to decide the next movement. This procedure is repeated so that the fingertip follows the edge of an object. These movements of the exploratory platform are active in two ways.

First, the fingertip responds actively to tactile contacts according to contact detection that causes the NXT Lego robot to stop a tap whenever the contact pressure crosses a predefined threshold. The 20 ms latency for this task is sufficient to detect an object and avoid collisions that could potentially damage the fingertip. After the fingertip has been stopped, it returns to a set height. This active contact strategy enables objects with different heights to be perceived equivalently without explicitly controlling the height parameter for tapping. It also avoids issues with unstructured environments in which robots do not have prior information about object position.

Our second use of active touch optimizes the edge perception by analysing if the fingertip is placed in a valid position range for perceiving the current edge orientation class (Figure 3). The valid ranges for positioning the fingertip relative to the edge are from the 8 mm to 11 mm position classes (with these ranges obtained from the off-line validation). When the fingertip position is in this valid range, it continues following the contour along the direction defined by the perceived edge orientation. If the position is outside this range, then an active repositioning step is executed to move the fingertip towards the desired range for optimal perception. The direction of the movements for both following the edge and repositioning relative to the edge are defined according to the current perceived edge orientation.

2.4 Control architecture for contour following

The control architecture for contour following with active tactile perception (Figure 4) is based on how humans perceive and act with their sense of touch. When humans feel physical pain through the sense of touch, they react with an avoidance movement of the part contacted. This is done first by a direct reflex arc to the muscles, while concurrently, the signal is sent to the brain to then make more appropriate avoidance and other actions concerning the sensation. Similarly, the architecture proposed for our robot reacts when a tactile contact is detected with the fingertip, regulating the contact force and moving it to a predefined position. Concurrently, the sensory data from the tactile contact is sent to other modules that decide the appropriate next movement for the contour following task. The modules in this architecture are described below.

Firstly, the *tactile sensor* module receives the information about the environment, here by performing taps against an object surface. The contact pressure is compared with a predefined threshold value,

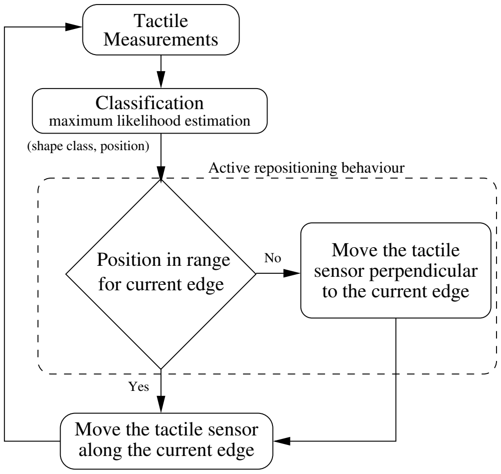

Figure 3: Active repositioning module. The repositioning step analyses whether the fingertip position is the optimal range according to the current perceived edge. If the fingertip is out of the optimal range, then it is moved until the optimal range is reached. Concurrently, the fingertip continues following along the perceived direction of the current edge.

and if the pressure is higher than the threshold the *contact detection* module sends a signal to stop the tapping movement. At the same time, the *classification* module perceives the current edge orientation and position class of the contact with respect to the tactile sensor. The *decision making* module decides the next movement for the fingertip using the outputs from the *classification* and *short-term memory* modules. In this work, the *decision making* module is based on a set of *if-then* rules, with possible output directions of movement either left, right, forward and backward. The *repositioning* module then places the fingertip in the optimal position for perception with respect to the current edge of the object (see previous section on active contour following), using outputs from the *decision making* and *short-term memory* modules. The information about the last state is stored in the *short-term memory* module and both updated and used by the *decision making* and *repositioning* modules. Finally, the *motor command* module sends the corresponding commands for the robot to perform the next movement.

2.5 Classification method

Data collected across four edges of a rectangular object were used to train the classifier (Figure 2a). A group of eighteen taps was obtained across each edge, from which we grouped and separated the data into regions (plane, edge orientation and air) per edge as shown in Figure 2c. Taking advantage of this separation of regions, it was possible to obtain from the classifier one of the following shape classes: lower edge, upper edge, right edge, left edge, plane and air. In addition, the contact position of the edge relative to the iCub fingertip was also returned from the classifier to give a pair (shape class, position) for each tap. This information was used by the control architecture to give the contour following behaviour (see previous section).

The probabilistic classifier was based on a Bayesian formalism developed for tactile perception [7, 9]. A simplification of this formalism produced a maximum likelihood estimation appropriate for this work, making decisions over a single tap of data [8]. The decision of perceived class C_l (shape class, position) was then found from the maximum likelihood estimate over the $T = 50$ samples in a single test tap and the $K = 12$ taxels across the fingertip surface

$$C = \arg\max_{C_l} \left[\prod_{k=1}^{K} \prod_{t=1}^{T} P(s_{t,k}|C_l) \right], \quad (1)$$

assuming the samples $s_{t,k}$ are independent and identically distributed. The likelihoods $\log P(x_i|C_l)$ were estimated from the training data using a histogram method. The histogram of pressure readings s of the training data for percept class c_n defined a sample distribution

$$P_k(s|c_n) = \frac{h_k(s)}{\sum_s h_k(s)}, \quad (2)$$

where $h_k(s)$ was the number of measurement values s in the histogram occurring for taxel k. We extracted 12 histograms (one for each taxel) for each of the training classes. Pressure values were binned to construct the histogram, uniformly partitioning each pressure range over 100 bins.

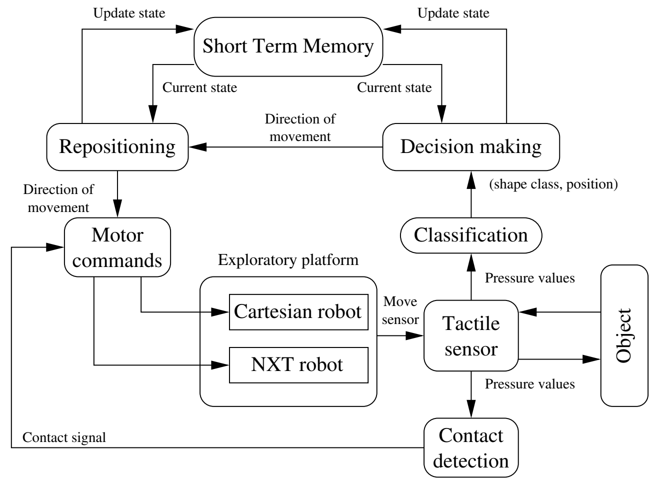

Figure 4: Perception-action architecture for the active contour following task. The next fingertip movement is obtained by the decision-making, repositioning and short-term memory modules. The fingertip also reacts to the tactile contact to directly regulate the contact force to avoid damage.

3 Results

3.1 Off-line validation of object shape and position

A plastic rectangle (dimensions 48 mm $\times$ 58 mm) was used to gather distinct sets of training and testing data, by using taps to sample an 18 mm region across each of its four edges (lower, upper, left, right) (Figure 2c). The 18 mm region was sampled with taps at 1 mm steps, obtaining 18 taps across each edge. The contact of the fingertip with the current object lasted $\sim$500 ms with a 2 s delay between taps to avoid possible transients such as from visco-elastic deformation of the sensor. Five datasets were collected for each edge, using one dataset for training and the rest for testing. The orientation of the iCub fingertip was fixed, in that it could not rotate relative to the edge orientation. Thus, the four edges represented four different orientations relative to the fingertip. These orientations are 0, 90, 180 and 270 degrees, as shown in Figure 2b.

Figure 5 shows the data collected from taps over the four different edges. The fingertip started the taps from the plane surface, then passed to the edge and finished in the air region. These movements passing across the different regions are indicated by change in contact pressure across the taxel positions according to the current edge and region. This information inputs the classification module returning the pair (shape class, position) which indicates the current region (lower edge, upper edge, left edge, right edge, plane or air) and position with respect to the fingertip.

Figure 6 shows the off-line classification results averaged over the four test datasets. The classification is averaged over the four collected datasets that were used for testing. We observe that the smallest classification errors for edge orientation and position are between the 8th and 11th position classes. Thus the region from

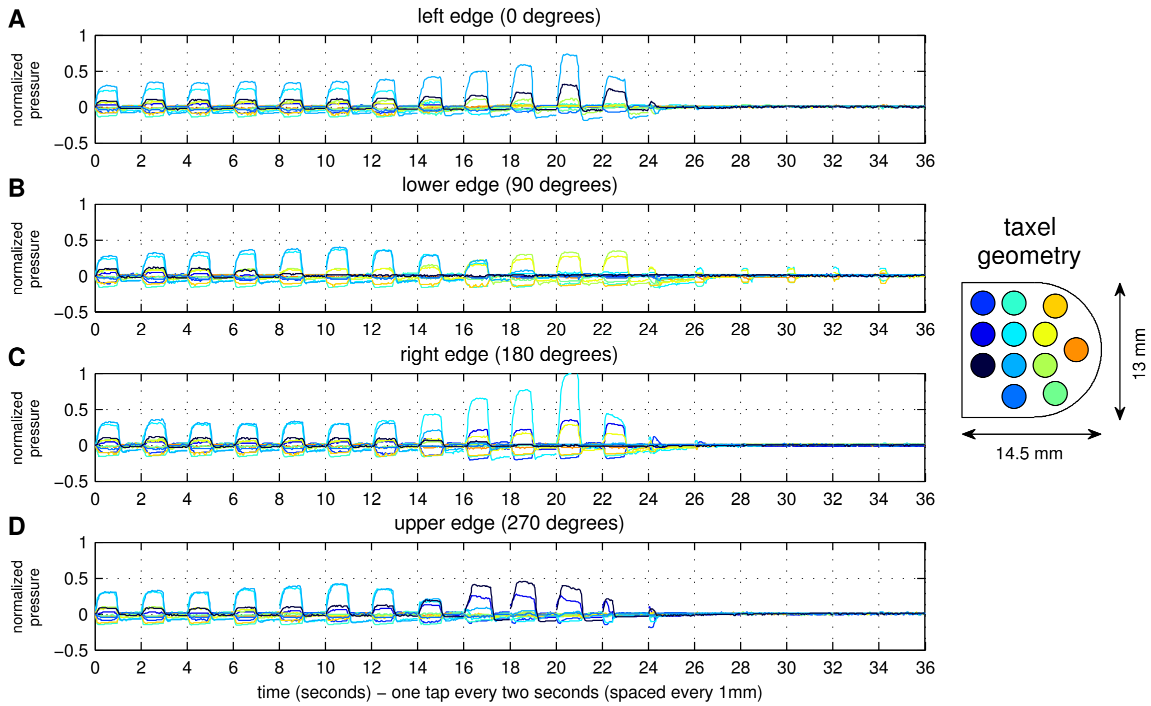

Figure 5: Taps performed by the iCub fingertip along an 18 mm displacement perpendicular to each edge in 1 mm steps. The colours denote the taxel in contact with the object (layout shown on fingertip diagram). The taps start on the plane region, then move onto the edge, and finally finish on the air region. The pressure values from taps are normalised to have maximum value one. These taps are used as input for training the perception method.

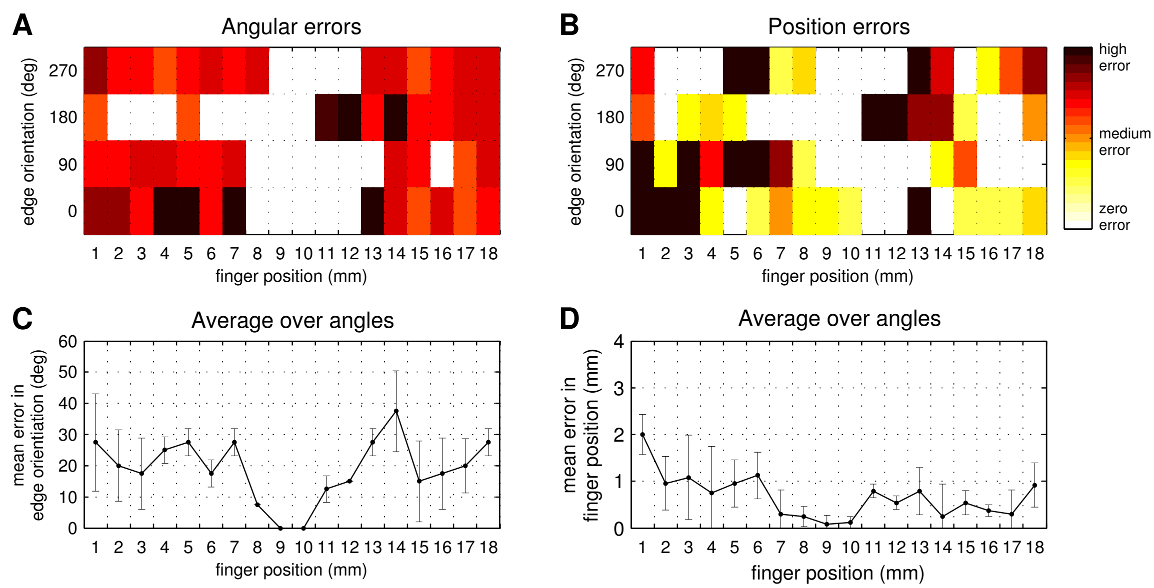

Figure 6: Angle and position errors averaged over four test datasets. The mean angle and position errors are shown for the four different edge orientations and eighteen positions (panels A, B). The overall edge orientation errors and fingertip position errors averaged over the edge orientations are shown, with bars at one standard deviation (panels C, D). Note that the smallest errors are between the 8-11 mm positions, which are used for the optimal locations for active contour following.

8 mm to 11 mm with respect to the fingertip are the optimal positions to detect the edge and contact position.

This position range was used by the repositioning module to place the fingertip at an optimal location for edge perception (Figure 3). Thus, we can attempt to move the fingertip to achieve better edge perception in order to improve the contour following. A perception-action architecture was developed to implement this contour following task (Figure 4). Note that the repositioning module of this control architecture can be enabled/disabled to enact active/passive tactile perception. Both methods, the passive and the active, are described below

3.2 Contour following: passive tactile perception

The passive approach for contour following was tested with a plastic rectangle object over two trials. The approach is passive in the

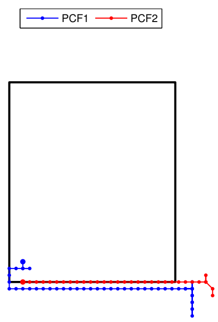

Figure 7: The two trials of the passive contour following (PCF) behaviour were not able to successfully finish the task.

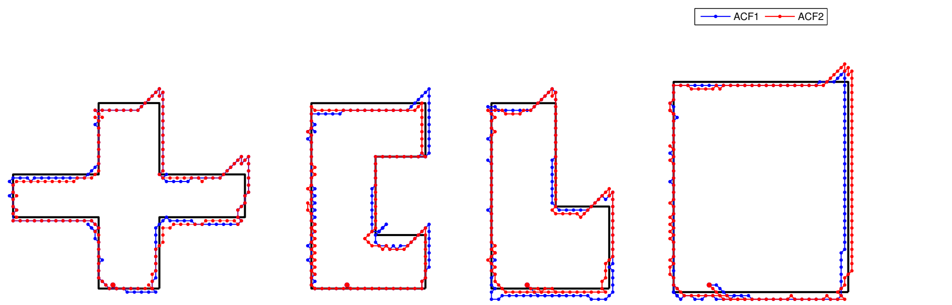

(a) Cross-shape (b) C-shape (c) L-shape (d) Rectangle

Figure 8: Results for contour following using active tactile perception. Two trials of the active contour following (ACF) were performed for each object. The active approach enabled the robot to successfully complete the task.

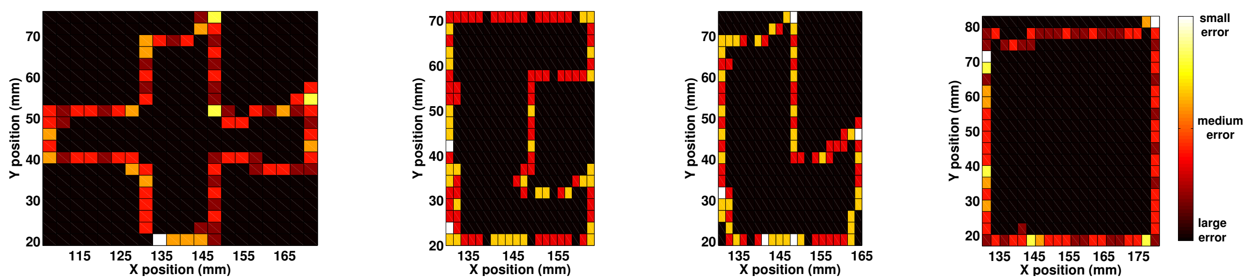

(a) Cross-shape (b) C-shape (c) L-shape (d) Rectangle

Figure 9: Object classification using the 2D-histogram of each object.

sense that the iCub fingertip does not determine the contact position of the fingertip relative to the edge, and thus makes movements only along the perceived orientation of the edge. This contrasts with the approach for active perception below, where corrective moves to reposition the centre of the fingertip onto the edge are made.

Figure 7 shows the Passive Contour Following (PCF) executed by the iCub fingertip. The dots represent the taps performed by the tactile sensor whilst the lines correspond to the displacements in x- and y-directions. The object is represented by the black lines. We observe that the fingertip can follow the initial edge at which it is placed, but it becomes lost at the corners in both test trials PCF1 (blue) and PCF2 (red). In consequence, the robot failed to to complete the contour following task.

3.3 Contour following: active tactile perception

The active approach for contour following was first tested with the plastic rectangular object used for passive perception. The approach is active in the sense that the iCub fingertip recognises the position of the contact, in order to reposition the iCub fingertip so that the edge lies in its centre where the perception was found to be best (using the off-line results from section 3.1). These corrective movements to reposition the centre of the fingertip onto the edge are made concurrently with following the orientation of the edge. The length of displacement between taps is 2 mm and 1 mm for repositioning.

Using active perception, the fingertip could now successfully trace around the entire outside edge of the rectangle. In particular, it had no problems at the corners, as shown by the two test trials ACF1 and ACF2 in Figure 8d.

The contour following test was then repeated on three other shapes: (a) a cross-shape; (b) a C-shape; and (c) and L-shape. The paths resulting from the active contour following are presented in Figure 8(a-d), with two Active Contour Following (ACF1 (blue) and ACF2 (red)) tests repeated for each object. Once again, the coloured dots correspond to the movements of the fingertip given by the perception-action control architecture (methods) and the solid black lines are the ground truth positions of the edge of the test shape. We observe that the active contour following method based on active tactile perception is able to consistently complete a circuit around the edge of all four test shapes.

Interestingly, the active contour following approach can change direction appropriately at corners without an explicit model derived from corner data. This is an emergent solution for corner classification that is a consequence of using an active perception strategy for edge detection, where the tactile sensor is repositioned in an optimal location for perceiving the edge while also moving along the edge.

3.4 Active contour following for shape recognition

Finally, the objects can be classified based on the extracted shapes from remembering the history of fingertip positions as it followed the edge of the objects. In this work, we used the Euclidean distance measured between a 2D histogram of where the fingertip tapped against the object and the 2D histogram of where a perfect edge following of each shape would have been, as shown in Figure 9 for the paths. For these histograms we used 20 bins over both the x- and y-axes. The values are normalised from 0 to 1. The minimum Euclidean distance between the histograms gives the classified object. As may expected from Figures 8 and 9, using this classification we obtained 100% classification accuracy for the shapes from the paths for active contour following. This result validates that active contour following can be used for shape recognition.

4 Conclusion

In this work, we presented an approach for active tactile perception, which we applied to contour following around the edges of several distinct test shapes. The method is appropriate for situations in which visual information is not available, so that the procedure needs to be purely tactile. Our approach was tested by whether it could successfully perform the contour following task. Data were collected by positioning the sensor in the *x*- and *y*-plane whilst tapping along the *z*-axis against an horizontal object. A tapping procedure was chosen to simplify the perception methodology and to avoid damaging the fingertip. To allow precise movements, a platform consisting of two robots (Yamaha XY series Cartesian robot and a Mindstorms NXT Lego robot) was built. The classification of tactile data was based on a probabilistic approach for robot perception [8, 7, 9] that has been applied previously to various forms of tactile perception [22, 23, 24].

Our approach was active in two aspects: (1) the fingertip reacts to each tactile contact by stopping the current tap and moving the fingertip back up to a predefined height; and (2) a repositioning module places the fingertip in an optimal position to perceive the orientation of the edge. Reacting to the contact detection also allowed the fingertip to protect itself against high pressures that could damage it. Moreover, this type of active tactile perception is useful in unstructured environments where robots do not have prior information about the position of objects.

Active perception was demonstrated using the contour following behaviour. This strategy allowed the robot to perform robust and successful contour following by moving the sensor actively to the optimal region to ensure it does not lose the edge of the object. The approach was applied to a range of different shapes over multiple test trials (Figure 8). Conversely, a passive contour following strategy, where the fingertip follows the edge without repositioning perpendicularly to it, failed at the corners of the objects (Figure 7). In consequence, we claim that active perception is necessary to successfully contour follow around the edge of objects with the perception algorithms and biomimetic fingertip hardware used here.

Object classification was performed using the path traced out by the contour following, analogously to how humans judge shape by following the contour of an object [1, 2]. The final shape classification resulted in 100% accuracy over two trials of four distinct test objects (a cross-shape, a C-shape, an L-shape and a rectangle).

Interestingly, we obtained an emergent solution from the active tactile repositioning to successfully navigate corners. Even though we did not define an explicit model for corners in the classification, the active approach was able to change the direction of movement when the fingertip reached a corner. This is another feature that highlights the utility of the active tactile perception approach.

In this initial study, only objects composed of right angles were used to show that active perception is necessary to complete the task. We expect, however, that our proposed approach can be extended straightforwardly to finer angular resolution, allowing the contour following method to apply to more complicated shapes. The current approach for edge and position perception is another aspect that could be improved; for instance, by using more than one tap for classification, as fits naturally within Bayesian perception framework [7, 9]. In general, we expect the proposed method can be improved to apply to completely general object shapes.

Acknowledgements

This work was supported by EU Framework project EFAA (ICT-270490) and the Mexican National Council of Science and Technology (CONACyT).

References

[1] S. Lederman and R. Klatzky. Haptic perception: A tutorial. *Attention, Perception, and Psychophysics*, 71:1439–1459, 2009. 10.3758/APP.71.7.1439.

[2] S. J. Lederman and Klatzky RL. Hand movements: A window into haptic object recognition. *Cognitive Psychology*, 19(19):342–368, 1987.

[3] R. Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, aug 1988.

[4] T. J. Prescott, M. E. Diamond, and A. M. Wing. Active touch sensing. *Phil. Trans. R. Soc. B*, 366:2989–2995, 2011.

[5] A. Schmitz, M. Maggiali, L. Natale, B. Bonino, and G. Metta. A tactile sensor for the fingertips of the humanoid robot icub. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 2212–2217, Oct. 2010.

[6] A. Schmitz, P. Maiolino, M. Maggiali, L. Natale, G. Cannata, and G. Metta. Methods and technologies for the implementation of large-scale robot tactile sensors. *Robotics, IEEE Transactions on*, 27(3):389 –400, june 2011.

[7] N.F. Lepora, U. Martinez-Hernandez, H. Barron-Gonzalez, M. Evans, G. Metta, and T.J. Prescott. Embodied hyperacuity from bayesian perception: Shape and position discrimination with an icub fingertip sensor. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 4638–4643, 2012.

[8] N.F. Lepora, M. Evans, C.W. Fox, M.E. Diamond, K. Gurney, and T.J. Prescott. Naive bayes texture classification applied to whisker data from a moving robot. In *Neural Networks (IJCNN), The 2010 International Joint Conference on*, pages 1–8, July 2010.

[9] N.F. Lepora, Charles W. Fox, Mathew H. Evans, Diamond Mathew E., Kevin Gurney, and Tony J. Prescott. Optimal decision-making in mammals: Insights from a robot study of rodent texture discrimination. 9(72):1517–1528, 2012.

[10] N.F. Lepora and K.N. Gurney. The basal ganglia optimize decision making over general perceptual hypotheses. *Neural Computation*, pages 1–22, 2012.

[11] A.D. Berger and P.K. Khoslar. Edge dectection for tactile sensing. In *In Proc. NASA conf. on space telerobotics*, pages 163–172, 1989.

[12] Ning Chen, Ray Rink, and H Zhang. Efficient edge detection from tactile data. *Proceedings of the 1995 IEEERSJ*, pages 386–391, 1995.

[13] Tri Phung, Yong Ihn, Ja Koo, and Hyouk Choi. An enhanced edge tracking method using a low resolution tactile sensor. *International Journal of Control, Automation and Systems*, 8:462–467, 2010.

[14] Ning N Chen, Hong Zhang, and Ray E Rink. *Edge tracking using tactile servo*, volume 2, pages 84–89. IEEE Comput. Soc. Press, 1995.

[15] K. Suwanratchatamanee, R. Saegusa, M. Matsumoto, and S. Hashimoto. A simple tactile sensor system for robot manipulator and object edge shape recognition. In *Industrial Electronics Society, 2007. IECON 2007. 33rd Annual Conference of the IEEE*, pages 245–250, nov. 2007.

[16] L. Natale and Eduardo Torres-Jara. A sensitive approach to grasping. *Sixth international Workshop on Epigenetic Robotics*, setp 2006.

[17] A. Schneider, J. Sturm, C. Stachniss, M. Reisert, H. Burkhardt, and W. Burgard. Object identification with tactile sensors using bag-of-features. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 243 –248, oct. 2009.

[18] RobotCub. icub software repository documentation. Website, 2010. http://eris.liralab.it/iCub/main/dox/html/index.html.

[19] M. Evans, C.W. Fox, M.J. Pearson, N.F. Lepora, and T.J. Prescott. Whisker-object contact speed affects radial distance estimation. In *Robotics and Biomimetics (ROBIO), 2010 IEEE International Conference on*, pages 720–725. IEEE, 2010.

[20] M.H. Evans, C.W. Fox, N.F. Lepora, M.J. Pearson, J.C. Sullivan, and T.J. Prescott. The effect of whisker movement on radial distance estimation: a case study in comparative robotics. *Frontiers in Neurorobotics*, 6, 2012.

[21] Paul Fitzpatrick, Giorgio Metta, and Lorenzo Natale. Towards long-lived robot genes. *Robot. Auton. Syst.*, 56(1):29–45, jan 2008.

[22] J.C. Sullivan, B. Mitchinson, M.J. Pearson, M. Evans, N.F. Lepora, C.W. Fox, C. Melhuish, and T.J. Prescott. Tactile discrimination using active whisker sensors. *Sensors Journal, IEEE*, 12(2):350–362, Feb. 2012.

[23] N.F. Lepora, J.C. Sullivan, B. Mitchinson, M. Pearson, K. Gurney, and T.J. Prescott. Brain-inspired bayesian perception for biomimetic robot touch. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 5111–5116, may 2012.

[24] N.F. Lepora, M.J. Pearson, B. Mitchinson, M. Evans, C. Fox, A. Pipe, K. Gurney, and T.J. Prescott. Naive bayes novelty detection for a moving robot with whiskers. In *Robotics and Biomimetics (ROBIO), 2010 IEEE International Conference on*, pages 131–136, Dec. 2010.