

Active Bayesian perception for angle and position discrimination with a biomimetic fingertip

Uriel Martinez-Hernandez, Tony Dodd, Tony J. Prescott and Nathan F. Lepora

Abstract—In this work, we apply active Bayesian perception to angle and position discrimination and extend the method to perform actions in a sensorimotor task using a biomimetic fingertip. The first part of this study tests active perception off-line with a large dataset of edge orientations and positions, using a Monte Carlo validation to ascertain the classification accuracy. We observe a significant improvement over passive methods that lack a sensorimotor loop for actively repositioning the sensor. The second part of this study then applies these findings about active perception to an example sensorimotor task in real-time. Using an appropriate online sensorimotor control architecture, the robot made decisions about *what to do next* and *where to move next*, which was applied to a contour-following task around several objects. The successful outcome of this simple but illustrative task demonstrates that active perception can be of practical benefit for tactile robotics.

I. INTRODUCTION

To operate autonomously, robots must make decisions and actions in the presence of sensory uncertainties. Inspiration can be taken from humans and animals, who adopt an active sensing strategy of directing their senses towards locations that reduce uncertainty about perceptual decisions [1]. For instance, humans actively move their hands and fingers to improve perception using sensorimotor feedback about the object in relation to the task being performed [2], [3]. Given such behaviour is widespread in biological systems, one may infer that robotics could also benefit from these principles.

In this study, we show how active Bayesian perception can enable a robot to perform a real-time sensorimotor task involving decisions about both *what to do next* and *where to move next*. These decisions are made by linking active perception to action in a sensorimotor loop that has two functions: first, to control the actions of the robot in achieving the task, and second to improve sensing accuracy through active perception. We apply these methods to a biomimetic fingertip perceiving both the angle and position of the edge of an object at different orientations and distances perpendicular to that edge. We also demonstrate how active perception can then enable a robot to achieve an example sensorimotor task, specifically to accurately trace the edge of various objects (Figure 1).

Our work is based on recent progress in applying active Bayesian perception to robot touch for texture [4] and shaped

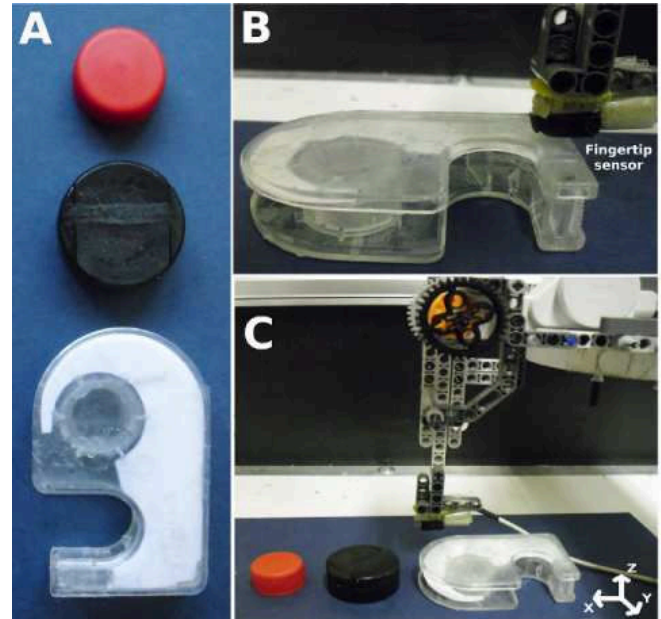


Fig. 1. (A) Different shaped and sized objects used for the sensorimotor task. (B) Tactile sensor in contact with the edge of an object at one angle and position. (C) Tactile sensor mounted on robotic platform to allow mobility in the x -, y - and z -axes.

stimuli [5], [6]. A series of papers [4], [5], [7], [8], [9], has formalised a *Bayesian perception* approach for robotics based on recent progress in understanding animal perception. That formalism extended naturally to active perception, by moving the sensor with a control strategy based on evidence received during decision making. Benefits of active Bayesian perception include: (i) robust perception in unstructured environments [4]; (ii) an order-of-magnitude improvement in acuity over passive methods [6] where movements for accumulating evidence to improve perception are not allowed; and (iii) a general framework for simultaneous object localisation and identification, or ‘where’ and ‘what’ [6]. However, the utility of active Bayesian perception has not yet been assessed on a practical sensorimotor task, which is the purpose of this study.

Here we develop and apply a sensorimotor architecture that controls actions by the outcome of decisions made using active Bayesian perception, extending a previous architecture to represent active perception [10]. For the first part of this study, we test angle and position discrimination using an off-line Monte Carlo validation of our active perception method. The results demonstrate a successful angular and position discrimination accuracy of 3.3 degs and 0.2 mm respectively,

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

UMH and TD are with the Department of Automatic Control and Systems Engineering, University of Sheffield, UK.

UMH, TD, TP and NL are with the Sheffield Centre for Robotics (SCentRo), University of Sheffield, UK.

Email: {uriel.martinez, t.j.dodd, t.j.prescott, n.lepora}@sheffield.ac.uk

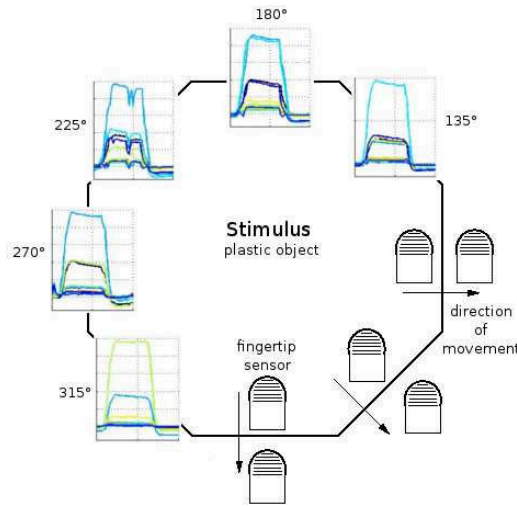


Fig. 2. Data collection for an edge stimulus over a range of orientations and positions. The object was turned at every 5 deg steps and the sensor was tapping along 18 mm at 0.2 mm steps for each angle. For visualisation, we show only five examples from a total of 72 angle and 18 position classes.

over 72 angular classes with resolution of 5 degs each and 18 position classes with resolution of 0.2 mm each. As part of the analysis, we evaluate the reaction times to make a decision with small classification errors. In the second part of this study, we apply this architecture to an online task in real-time, with a simple but illustrative problem of contour following the edges of various shapes. For each contact location, the fingertip sensor actively perceives the angle and position class by repositioning itself relative to the edge while integrating evidence up to decision threshold. After reaching the threshold, the robot then makes a movement parallel to the perceived edge direction to trace the outside edge of the object. For demonstration, we present traced contours from three differently shaped and sized objects, validating the method.

II. METHODS

A. Tactile data acquisition

This study considers tactile sensing using a biomimetic fingertip sensor designed for the iCub humanoid robot. The dimensions and shape of the fingertip sensor are similar to the human fingertip and are intended to allow the iCub humanoid to manipulate and interact with objects with its hands [11]. The sensor is built with a capacitive technology composed of 12 taxels (tactile elements) that react to contact pressure. The capacitive pressure values are digitised locally with a capacitive to digital converter (CDC) and sampled at 50 Hz. The conversion output is sent to a main computer through a CAN (Controller Area Network) bus.

We are interested in demonstrating angle and position perception with this fingertip. Therefore, we initially collected tactile data from the edge of a plastic square oriented at different angles and positions with respect to the fingertip. For data collection, we chose brief tapping motions, for two main reasons. First, to minimize damage to the sensor,

in contrast with sliding motions that would wear away the rubber coating. Second, to give a reliably repeatable motion for tactile perception, analogous to palpation and whisking performed by humans and rodents respectively.

For collecting angular classes, we turned manually the stimuli object with 5 deg steps over a total of 72 angular classes. Then, for each of these angular classes, the fingertip was moved perpendicularly to the edge along 18 mm at 0.2 mm steps, with a total of 90 taps. We built position classes in groups of 5 taps for each angle obtaining 18 classes of 1 mm each. In total, we thus obtained 1296 perceptual classes. The data collection was repeated two times, to give one set for testing another for training.

The configuration of the edge stimulus during angle and position data collection is shown in Figure 2. To aid visualisation, we show the object rotated only every 45 degs and tactile data only when the fingertip is positioned on the edge.

B. Robot platform

We constructed a robot to provide the movement capability for the biomimetic fingertip. The overall robot is itself formed of two individual robots: a Cartesian robot (Yamaha XY-x series) and a Mindstorms NXT Lego robot. The Cartesian robot allows precise positioning movements in the x - and y -axes with an accuracy of $\sim 20\mu\text{m}$. The NXT Lego robot was built to allow controlled movements in the z -axis to tap against the object.

The fingertip sensor was mounted on the robot platform as shown in Figure 1. Movements are well-controlled in the x - and y -axes owing to the capabilities of the Cartesian robot. On the other hand, the NXT Lego robot is not able to perform very precise movements, but can nevertheless be controlled by the tactile feedback from the fingertip sensor. This highlights two important aspects of our work: (1) the system reacts in real-time to tactile detection and (2) the system works under uncertainty in the height of the tactile contact along the z axis. The degrees of freedom provided by the robot platform do not allow rotations of the fingertip sensor with respect to the object, and therefore the fingertip kept the same orientation during all experiments.

The procedure for data collection is as follows. First, the robot is moved to a start position. Next, a tap is performed by the NXT Lego robot, which stops and then quickly returns when the fingertip sensor detects a contact (with a fixed pressure threshold). Then, the tactile data is saved and the fingertip sensor is prepared for the next tap. We set a 5 secs delay between taps to ensure that any contact-induced transients decayed. Finally, the Cartesian robot moves to the next position on the perpendicular trajectory to the current angle of the edge. The complete system works in real-time, is encoded in C/C++ with the YARP library and is used for communication and control [12]. Additionally, we developed some subroutines using the NXTPP library to control the NXT Lego robot.

C. Active Bayesian perception

Our study is based on previous work on active Bayesian perception that was applied to robot touch perception of

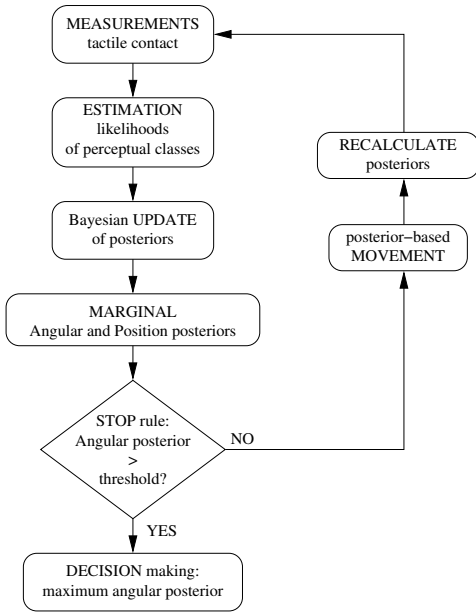


Fig. 3. Flowchart for active angle and position perception. In active Bayesian perception, the posteriors are updated by accumulating evidence represented as likelihoods of each tap while also using these posteriors to determine how to move the sensor. Once one of the angle classes crosses a preset threshold, a decision is made.

texture [4] and shape [7]. Here we explore tactile perception for a large number of angle and position classes and use active perception to improve the classification results. The method is based on the Bayes' rule via the method of sequential analysis, which recursively updates the posterior probabilities from the prior probabilities and the likelihoods found from a measurement model of the tactile data.

Priors: We assume uniformly distributed priors for all perceptual classes $P(c_n) = P(c_n|z_0) = 1/N$. These define the posteriors at time $t = 0$, and will be recursively updated with the likelihoods obtained from each tap of the fingertip.

Likelihood estimation: From each tap performed by the tactile sensor we obtain a time series of digitised pressure values from the K taxels (12 taxels). The measurement model is constructed off-line using a nonparametric estimation based on the histograms of the sensor values from training data sets. The likelihood of a perceptual class $c_n \in C$ of a test tap z_t is then evaluated with these histograms. The measurement model is obtained from the probabilities

$$P_k(s|c_n) = \frac{h_k(s)}{\sum_s h_k(s)}, \quad (1)$$

where $h_k(s)$ is the number of observed values s in the histogram for taxel k . The histograms were uniformly constructed over 100 bins. The value from equation (1) then gives the likelihood of a tap evaluated over all samples s_j :

$$P(z_t|c_n) = \sqrt[J]{\prod_{j=1}^J \prod_{k=1}^K P_k(s_j|c_n)}, \quad (2)$$

where $J = 100$ and $K = 12$ are the number of data and taxels sampled for each tap respectively. The model treats all

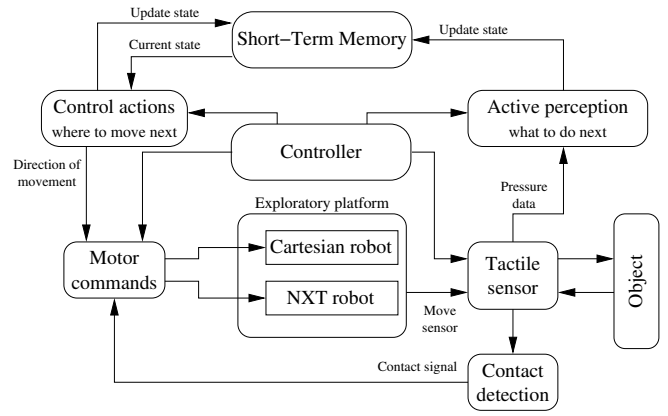


Fig. 4. Sensorimotor architecture developed for applying active Bayesian perception to an online task. The tactile sensor provides the feedback from the world, whilst the mobility in x -, y - and z axes is provided by the robotic platform. The *what to do next* and *where to move next* decisions are implemented by the 'active' perception and 'control actions' modules, respectively.

samples as independent and identically distributed for each tap and taxel.

Bayesian update: For updating the posterior probabilities $P(c_n|z_{1:t})$, we use a recursive implementation of Bayes' rule over all N perceptual classes c_n with the likelihoods $P(z_t|c_n)$ of the measured tap z_t . The prior takes the value of the posterior at time step $t - 1$ resulting in an update,

$$P(c_n|z_{1:t}) = \frac{P(z_t|c_n)P(c_n|z_{1:t-1})}{\sum_{n=1}^N P(z_t|c_n)P(c_n|z_{1:t-1})}. \quad (3)$$

Marginal angle and position posteriors: Each class c_n corresponds to a (w_i, x_l) pair where w_i and x_l are the angle and position for each perceptual class respectively. The posteriors are the joint distributions over these joint classes, then the beliefs over individual angular and position perceptual classes are given by the marginal posteriors

$$P(w_i|z_{1:t}) = \sum_{l=1}^L P(x_l, w_i|z_{1:t}), \quad (4)$$

$$P(x_l|z_{1:t}) = \sum_{i=1}^I P(x_l, w_i|z_{1:t}), \quad (5)$$

with the angle beliefs summed over all position classes and the position beliefs summed over all angle classes.

Stop decision for angle posteriors: A threshold crossing rule is used to stop accumulating evidence and make a final decision about the angle and position classes. The *maximum a posteriori* (MAP) estimate is used to decide the angle and position perceptual class when the angle belief passes a threshold

$$\text{if any } P(w_i|z_{1:t}) > \theta_w \text{ then } w_{\text{MAP}} = \arg \max_{w_i} P(w_i|z_{1:t}) \quad (6)$$

In previous work with this fingertip sensor, we have found that the best perception is obtained for contacts at its centre [5], [10]. This assumption is supported with results here in section III-B. Thus, for active perception we assume there

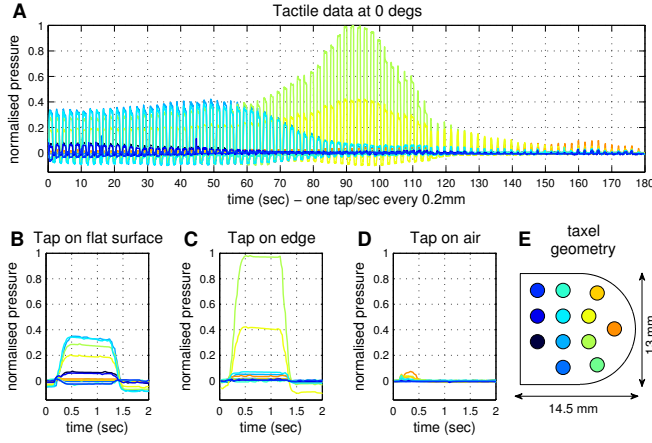


Fig. 5. (A) Tactile data collection at 0 degs along 18 mm perpendicular to the edge of the stimulus. (B) Example tap on flat surface, here the pressure is spread to most of the taxels. (C) Tap on edge where the pressure is concentrated in a small number of taxels. (D) Tap on air, the sensor is not in contact with the stimulus. (E) Colour code that shows the location of the taxels on the fingertip.

is a preset target fixation position x_{fix} at the centre of the sensor (position class at 9 mm). The movement represented by Δ is then determined from the position decision x_{loc} of the current sensor location,

$$x \rightarrow x + \Delta(x_{\text{loc}}), \quad (7)$$

$$\Delta(x_{\text{loc}}) = x_{\text{fix}} - x_{\text{loc}}. \quad (8)$$

The algorithm tries to move the sensor to a good location (its centre) relative to the object, thereby improving the tactile perception. The sequence of the different steps of the resulting algorithm are visualised as a flowchart in Figure 3.

Online active control: Online active control was achieved with a sensorimotor control architecture that applied the decisions from active perception to move the robot according to the task (Figure 4). Briefly, the tactile sensor module provides the feedback from contacting an object. To protect the sensor against dangerous pressures, the contact detection module sends a signal to stop the movements whenever the contact pressure crosses a preset threshold. At the same time, the tactile data are prepared and sent to be processed by the active perception module for classification about the current angle and position perceptions. Here the system decides what to do next: either accumulating more evidence or do a movement. Once a decision is made, the classification result is updated in the short-term memory module. The control actions module defines where to move next, corresponding to the current active perception outcome and prior position information stored in the short-term memory. These movements are to trace the contour of the object. All movements are managed by the motor commands module which is directly connected to the robotic platform. The controller is responsible for synchronizing and controlling all the processes involved in the experimental task.

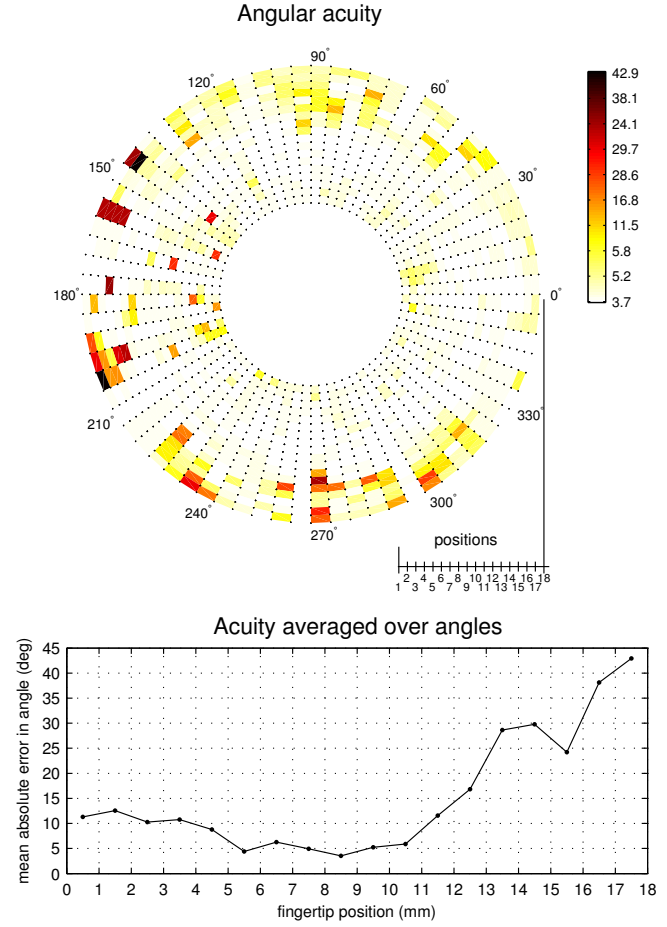


Fig. 6. Angle acuity using passive Bayesian perception. The top ‘ring’ plot shows the perceptual accuracy over all 72 angular and 18 position classes, with large errors shown in red and small errors in white. The bottom plot shows these results averaged over the angles, so that only the positional dependence is visible. Results were evaluated with a Monte Carlo method over 10000 iterations.

III. RESULTS

A. Initial inspection of data

Figure 5 shows an example of the normalised tactile data for the fingertip sensor tapping at 0 degs along an 18 mm trajectory in 0.2 mm steps. Tap movements follow a perpendicular trajectory with respect to the edge of the test object. Example taps at the flat surface, edge and air are shown. For the flat surface, the fingertip is in contact with most of the taxels and then obtains an even pressure contact. The pressure contact gradually starts to concentrate in a small number of taxels whilst the fingertip approaches the edge of the object, thus having high pressure. Finally, when the fingertip leaves the object, there is no contact and just a small signal due to background noise is received by the sensor, such as due to the motion of the fingertip. The layout colour shows the taxel activations and pressure levels at each tap during the data collection. The pattern of these pressures and activations enables perception of both the angle and position of the fingertip relative to the edge.

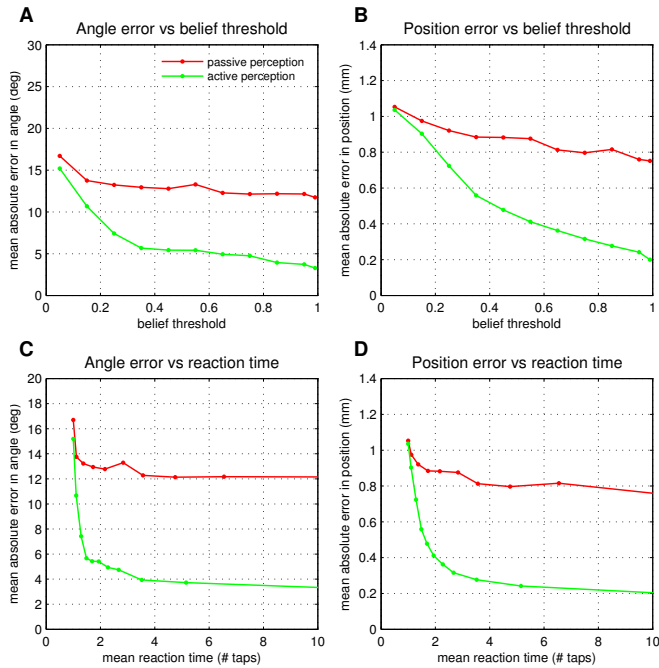


Fig. 7. Angle and position acuity for active and passive Bayesian perception. Results for passive perception are shown in red, and those for active perception in green. Angle errors are plotted against belief threshold (A) and reaction time (C), while position errors are also plotted against belief threshold (B) and reaction time (D).

B. Passive angle and position perception

For our first robot experiment, we investigated off-line how successive taps lead to successful perception via evidence accumulation. The tapping process was repeated until the belief for one of the angle classes crossed a threshold, which then triggered a decision about the angle and position where the fingertip is located. This approach is known as passive Bayesian perception, meaning that the fingertip sensor position is fixed until a threshold is crossed. In other words, the fingertip is not allowed to move to another location to improve perception. The validation was based on a Monte Carlo method of drawing random angle and position data from the test dataset. We generally used 10000 iterations for each data point in the following figures. To implement passive perception, an algorithm based on the flowchart in Figure 3 was used but with the movement rule disabled.

Figure 6 shows the results for passive angle perception over 72 angle classes with the smallest classification error of 3.7 degs. This result was obtained for a belief threshold of 0.45, corresponding to a mean reaction time of about 5 taps. We observe that for the position class at 9 mm (its centre) the fingertip has the best perception.

C. Active angle and position perception

For our next experiment, we examined active perception with a sensorimotor control loop that moves the fingertip to attain improved perception based on the tactile feedback (Figure 3). For this validation study, we use the training and testing datasets collected above (see section II-A).

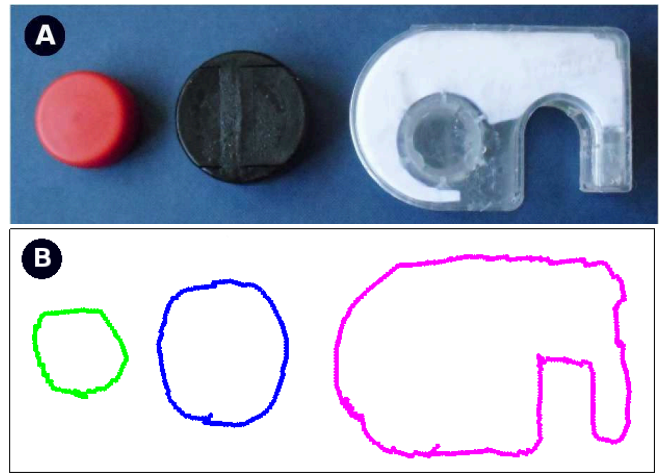


Fig. 8. Traced contours from applying active perception to an online contour following task. (A) Circles with 2 cm, 4 cm diameter and an asymmetric object (sellotape holder) used for real-time contour following. (B) Contours traced obtained as result of active perception applied to sensorimotor control.

Our active perception method requires a target fixation position to orient to. By the results in Figure 7, we selected a target fixation position class 9 mm, since this location has the smallest classification error and is thus optimal for perception. For the validation we used the range of belief thresholds [0.05, 0.10, ..., 0.999]. Mean angle and position errors for passive perception (red curve) and active perception (green curve) are shown in Figures 7A,B. We observe the smallest mean error of 12.2 deg and 0.8 mm for passive perception. On the other hand, for active perception, the best angle and position errors are 3.3 degs and 0.2 mm respectively, showing considerable improvement over passive perception. Finally, plots 7C,D show the mean absolute angular error and mean absolute position error against reaction times. From these results, we also observe that decision threshold leading to decisions at around 4-6 taps give the minimum classification error for both angle and position perception.

D. Application of active perception to sensorimotor task

For our final demonstration, we apply our sensorimotor architecture (Figure 4) based on active perception to contour following of various shapes in real time. We chose this task to illustrate how our method solves a problem where the robot has to decide both *what to do next* and *where to move next*. In particular, for each contact location, the fingertip actively perceives the angle and position class by repositioning itself relative to the edge. After reaching the decision threshold, the robot then makes a movement parallel to the perceived edge direction to trace the outside edge of the object. Thus, by implementing a real-time interaction between the control of the fingertip sensor and objects in its environment, we can demonstrate how active perception enables the robot to successfully complete the sensorimotor task.

In the online test, we used two circles of diameter 2 cm and 4 cm, and an asymmetric object (Figure 1A - a sellotape holder). These objects cover a wide range of angles and radii of curvature for discrimination, which provides a good test

scenario for our methods. For this test, we used a belief threshold of 0.8 for each decision. The fingertip was able to move 2 mm along the angle perceived at each location. Also, a mean of 6 taps per location were needed to cross the belief threshold and make a decision.

The first traced contour (Figure 8B, green line) is from tracking the edge of the 2 cm circular object. Evidently, the robot successfully traced the contour without losing contact with the edge. Some parts of the contour seem less accurate (*i.e.* right and top sides), which are apparently at angles where the base of the fingertip was in contact with the edge. However, the final contour is good enough to describe a circle. For the second traced contour (Figure 8B, blue line), which is around the 4 cm circle, the tracking was more accurate. In this case, the active perception was able to ensure that the centre of the tactile sensor was almost always in contact with the edge of object. Finally, for the traced contour around the asymmetric object (Figure 8B, purple line), we observe the trace corresponds well to the original object, even though there were some regions (bottom-left and right) where it was more difficult for the robot to perceive the angles due to the very small radii of curvature.

Overall, the classification accuracy was good enough to perceive a wide range of angles across a range of radii of curvature, even though for training we only used angles of 5 degs resolution. We note that the perceptual process is able to interpolate between the discrete angles from the training to give accurate control for tasks involving a continuum of angles. Thus, we observe with a real task how active perception in robotics enables good perceptual decisions that then lead to successful task performance.

IV. DISCUSSION

In this work we applied active Bayesian perception to angle and position discrimination with a biomimetic fingertip sensor, with application to autonomous robots making decision and actions in the presence of uncertainty. We showed that active perception allows improved performance over passive methods, to achieve mean angle and positions errors of just 3.3 degs and 0.2 mm respectively (Figure 7). We then developed a sensorimotor control architecture that applied active perception to robot control in a task requiring decisions about actions based on the perception. We chose a simple but illustrative sensorimotor task of contour following around three differently shaped and sized objects, which were chosen to cover a wide range of angles and radii of curvature (Figure 8A). The robot managed to successfully trace around all of the objects (Figure 8B).

We can observe how active perception linked to actions allows the robot to successfully behave and complete the task. The traced contours were achieved under the presence of uncertainty from different sources: (1) tactile measurements and (2) height of contact (z -axis). Thus, the robot accumulates evidence reducing these uncertainties to make good decisions. These results can be compared with previous work about passive perception making one tap per decision and where then, in the absence of active perception, the robot

failed to follow the contour of the object [10].

It is interesting to compare these results with previous studies of active perception for angle classification concerning humans using their index finger [13], [14]. For these experiments, people were asked to touch the edge of an object with their index finger to perceive the angle. During exploration, the subjects moved their hand and fingers in order to improve the perception in a manner analogous to the behaviour of the robot in the present experiments.

Overall, we observed that active perception can help develop autonomous robots able to make accurate decisions and consequently good actions under the presence of sensory uncertainty. We believe that the methods presented here can be extended and applied to tactile perception and sensorimotor control with more than one fingertip. For future work, we therefore expect to explore the methods developed here with a robotic hand to have a complete, tactile robot able to actively perceive, explore and manipulate its environment.

REFERENCES

- [1] R. Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, 1988.
- [2] S. J. Lederman and R. L. Klatzky. Hand movements: A window into haptic object recognition. *Cognitive Psychology*, 19(19):342–368, 1987.
- [3] T. J. Prescott, M. E. Diamond, and A. M. Wing. Active touch sensing. *Phil. Trans. R. Soc. B*, 366:2989–2995, 2011.
- [4] N.F. Lepora, U. Martinez-Hernandez, and T.J. Prescott. Active touch for robust perception under position uncertainty. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 3005–3010, 2013.
- [5] 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.
- [6] N.F. Lepora, U. Martinez-Hernandez, and T.J. Prescott. Active bayesian perception for simultaneous object localization and identification. In *Robotics: Science and Systems*, 2013.
- [7] 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, 2012.
- [8] N.F. Lepora, C.W. Fox, M.H. Evans, M.E. Diamond, K. Gurney, and T.J. Prescott. Optimal decision-making in mammals: Insights from a robot study of rodent texture discrimination. *Journal of The Royal Society Interface*, 9(72):1517–1528, 2012.
- [9] 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. *Neural Networks (IJCNN), The 2010 International Joint Conference on*, pages 1–8, 2010.
- [10] U. Martinez-Hernandez, T.J. Dodd, L. Natale, G. Metta, T.J. Prescott, and N.F. Lepora. Active contour following to explore object shape with robot touch. In *World Haptics Conference (WHC), 2013*, pages 341–346, 2013.
- [11] 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, 2011.
- [12] G. Metta, P. Fitzpatrick, and L. Natale. Yarp: Yet another robot platform. *International Journal on Advanced Robotics Systems*, 3(1):43–48, 2006.
- [13] Y. Jiajia and W. Jinglong. Human characteristics on tactile angle discrimination by object movement condition. In *Complex Medical Engineering (CME), 2011 IEEE/ICME International Conference on*, pages 514–519, 2011.
- [14] J. Voisin, G. Benoit, and C.E. Chapman. Haptic discrimination of object shape in humans: two-dimensional angle discrimination. *Experimental Brain Research*, 145:239–250, 2002.