



This is a repository copy of *Active Control for Object Perception and Exploration with a Robotic Hand*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/108451/>

Version: Accepted Version

Proceedings Paper:

Martinez-Hernandez, U., Lepora, N.F. and Prescott, T.J. (2015) Active Control for Object Perception and Exploration with a Robotic Hand. In: Biomimetic and Biohybrid Systems. 4th International Conference, Living Machines 2015, July 28 - 31, 2015, Barcelona, Spain. Lecture Notes in Computer Science, 9222 . Springer International Publishing , pp. 415-428. ISBN 978-3-319-22978-2

https://doi.org/10.1007/978-3-319-22979-9_42

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Active control for object perception and exploration with a robotic hand

Uriel Martinez-Hernandez¹, Nathan F. Lepora² and Tony J. Prescott¹

¹ Sheffield Robotics Laboratory and the Department of Psychology, University of Sheffield, Sheffield, U.K.

² Department of Engineering Mathematics, The University of Bristol, and the Bristol Robotics Laboratory (BRL), Bristol, U.K.

Abstract. We present an investigation on active control for intelligent object exploration using touch with a robotic hand. First, uncertainty from the exploration is reduced by a probabilistic method based on the accumulation of evidence through the interaction with an object of interest. Second, an intrinsic motivation approach allows the robot hand to perform intelligent active control of movements to explore interesting locations of the object. Passive and active perception and exploration were implemented in simulated and real environments to compare their benefits in accuracy and reaction time. The experiments for validation of the proposed method were performed with an object recognition task, using a robotic platform composed by a three-fingered robotic hand and a robot table. The results demonstrate that our method permits to the robotic hand to achieve high accuracy for object recognition with low impact on the reaction time required to perform the task. These benefits make the proposed method suitable for perception and exploration in autonomous robotics.

1 Introduction

The intelligent exploration of the environment performed by humans requires the use of exploratory procedures and intelligently controlled movements of their hands and fingers [1],[2]. The exploratory procedures are employed according with the information of interest from the environment, whilst active perception permits to decide where to move the hand and fingers to explore interesting locations and extract useful information [3],[4]. These are important features required for the development of intelligent robots capable to explore and interact with their environment in the presence of uncertainty.

In this work, we present a method for intelligent perception and exploration of objects using a three-fingered robotic hand. First, an exploratory procedure is developed to allow the robotic hand to explore various objects moving the hand and fingers around them, extracting both tactile and proprioceptive information. Second, reduction of uncertainty is implemented with a Bayesian approach, which has been employed for object shape extraction [5],[6] and simultaneous object localisation and identification with a biomimetic fingertip sensor [7],[8]. The reduction of uncertainty is achieved by the accumulation of evidence based on the continuous interaction of the robotic hand with the environment.

An active exploration behaviour, similar to the one employed by humans, is performed by an intrinsic motivation approach, which permits the robotic hand to move towards interesting locations to extract useful information. This approach studied by psychology and cognitive sciences [9],[10], have found that intrinsic motivation is essential for cognitive development in humans, and also required for robust exploration and manipulation in robotics [11],[12],[13].

We implemented our methods in a sensorimotor architecture for the intelligent control of the exploration movements with a robotic hand for an object exploration and recognition task. Object recognition has been studied using tactile feedback with a simulated robotic arm showing accurate results [14]. The use of proprioceptive information from a five-fingered robotic hand has allowed to develop an object recognition task [15]. A fixed number of exploratory movements with a Self-Organising Map (SOM) approach were proposed for object recognition with a three-fingered robotic hand [16]. A drawback from these methods is that they are based on a single contact and passive exploration modality, where the robotic hand is not permitted to moved to interesting locations to reduce uncertainty. This contrasts with our method for active object exploration, which allows the robotic hand to intelligently move and improve perception from the object by the continuous interaction with the environment.

Our proposed methods were validated in simulation and real environments with an object exploration task. First, for the simulated environment we used the datasets collected from 6 test objects. Next, for the real

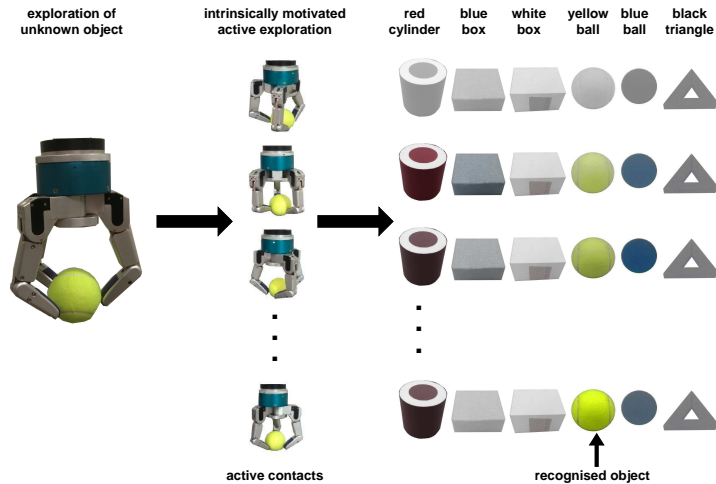


Fig. 1. Object recognition by the robotic hand using an intrinsically motivated active exploration approach. The robotic hand is actively moved towards interesting locations to improve perception. This process is repeated N times until a belief threshold is exceeded, permitting to make better decisions about the object being explored.

environment, we used a robotic platform composed by a robotic hand and a positioning robotic table for exploration of various test objects. For both environments, the exploration was performed in passive and active modalities to compare their performance. Results demonstrate that our approach for active control of object perception and exploration permits to achieve higher perception accuracy over passive exploration modality, which offers a method suitable for autonomous robotics.

2 METHODS

2.1 Robotic platform

This study employs a robotic platform composed by a three-fingered robotic hand mounted on a positioning robotic table shown in Figure 2.

The three-fingered robotic hand from Barrett Hand has 4-DoF, with 1-DoF in each finger for its opening and closing, and 1-DoF for spreading the fingers around the palm of the hand (see Figure 2a). The robotic hand also is integrated with tactile and force sensors. Each finger is composed by 22 taxels (tactile elements), whilst the palm has 24 taxels of 12 bit resolution. The strain sensors are located in each finger which permit to detect when a tactile contact has exceeded a force threshold. Also, it is possible to obtain proprioceptive information from the robotic hand in real-time. In this work, we focus our study on the use of proprioceptive information for the proposed methods described later in this section.

The positioning robotic table has 4-DoF that permit precise movements in x -, y - and z axes, and rotations on θ (see Figure 2b). The three-fingered robotic hand is mounted on the positioning robotic table to allow a larger set of exploration movements: 1) opening and closing of fingers; 2) spreading of fingers around the palm; 3) rotation of the wrist (θ); and 4) displacements of the robotic hand (x -, y - and z axes). Thus, this configuration permits the exploration of a large variety of objects by synchronising and controlling the fingers and table movements.

We developed a controller embedded in a microcontroller Arduino for the positioning robotic table. The data collection and exploratory movements performed by the robotic platform are controlled in real-time by tactile feedback and perceptions from the proposed methods. The synchronisation of the modules of software and hardware that compose the robotic platform is based on the YARP (Yet Another Robot Platform) middleware developed for robot control [17].

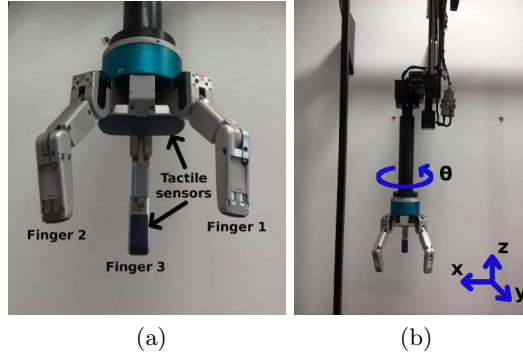


Fig. 2. Robotic platform used for data collection and validation of the proposed method. (a) Robotic hand with 4-DoF from Barrett Hand. (b) Positioning robotic table that provides mobility to the robotic hand.

2.2 Data collection

Our work is focused on object recognition with robotic hands using proprioceptive information. For this purpose, we collected information from the position and orientation of the fingers and hand in real-time for each contact performed on the set of test objects.

Figure 3 shows the sequence of movements performed by the robotic hand around two test objects. First, each finger moves independently towards the unknown object. They stop as soon as a contact is detected by exceeding the tactile pressure and force thresholds. The fingers keep in contact with the object for 1 sec, collecting 50 samples of proprioceptive information from the hand per contact. Second, the fingers are opened to a predefined home position, and then the wrist is rotated to collect data from a new orientation of the hand. The wrist is rotated in 12 degrees steps covering 360 degrees to explore the complete object. This process was repeated 5 times per object to have one dataset for training and four datasets for testing.

The data collected is stored in a 50×5 matrix per contact. The first three columns contain the positions of contacts detected by each finger, the fourth column contains the value of the spread motor, and the fifth column contains the angle orientation of the hand for each contact detected.

Sequence of exploration movements for data collection




Fig. 3. Sequence of movements performed by the robotic hand around the test objects for data collection. For each contact, proprioceptive information from the robotic hand was recorded. A total of 30 contacts were performed for each object, which was repeated five times, thus having one dataset for training and four datasets for testing. For visualization, here we only show a sequence of four contacts.

2.3 Bayesian estimator

Robotics has made use of Bayesian methods to develop a variety of applications and estimate an state given the observations. Here, we use a Bayesian approach to estimate the most likely object been explored by using proprioceptive information from a robotic hand.

This probabilistic approach uses the Bayes' rule with a sequential analysis method, estimating the posterior probabilities recursively updated from the prior probabilities and likelihoods obtained from a measurement model. Then, the robotic hand makes a decision once the belief threshold about the object being explored is exceeded. This method has been tested for object shape extraction [5],[6] and simultaneous object localisation and identification [7],[8] using the fingertip sensors from the iCub humanoid robot [18].

Prior: an initial uniform prior probability is assumed for all the test objects to be explored. The initial prior probability for an object exploration process is define as follows:

$$P(c_n) = P(c_n|z_0) = \frac{1}{N} \quad (1)$$

where $c_n \in C$ is the perceptual class to be estimated, z_0 is the observation at time $t = 0$ and N is the number of objects used for exploration.

Measurement model and Likelihood estimation: each contact performed by the robotic hand during the object exploration task provides proprioceptive information from M motors: position and spread of the three fingers, and orientation of the hand. This information is used to construct the measurement model with a nonparametric estimation based on histograms. The histograms are used to evaluate a contact z_t performed by the robotic hand at time t , and estimating the likelihood of a perceptual class $c_n \in C$. The measurement model is obtained as follows:

$$P(s|c_n, m) = \frac{h(s, m)}{\sum_s h(s, m)} \quad (2)$$

where $h(s, m)$ is the number of observed values s in the histogram for motor m . The observed values are normalised by $\sum_s h(s, m)$ to have properly probabilities that sum to 1. Evaluating Equation (2) over all the motors, we obtained the likelihood of the contact z_t as follows:

$$\log P(z_t|c_n) = \sum_{m=1}^{M_{\text{motors}}} \sum_{s=1}^{S_{\text{samples}}} \frac{\log P(s|c_n, m)}{M_{\text{motors}} S_{\text{samples}}} \quad (3)$$

where $P(z_t|c_n)$ is the likelihood of a perceptual class c_n given the measurement z_t from M motors at time t .

Bayesian update: the posterior probabilities $P(c_n|z_t)$ are updated by the recursive implementation of the Bayes' rule over the N perceptual classes c_n . The likelihood $P(z_t|c_n)$ at time t and the prior $P(c_n|z_{t-1})$ obtained from the posterior at time $t - 1$ are combined as follows:

$$P(c_n|z_t) = \frac{P(z_t|c_n)P(c_n|z_{t-1})}{P(z_t|z_{t-1})} \quad (4)$$

Properly normalised values are obtained with the marginal probabilities conditioned from previous contact as follows:

$$P(z_t|z_{t-1}) = \sum_{n=1}^N P(z_t|c_n)P(c_n|z_{t-1}) \quad (5)$$

Stop decision for object recognition: the accumulation of evidence with the Bayesian update process stops once a belief threshold is exceeded, making a decision about the object being explored. The object perceptual class is obtained using the *maximum a posteriori* (MAP) estimate as follows:

$$\begin{aligned} &\text{if any } P(c_n|z_t) > \theta_{\text{threshold}} \text{ then} \\ &c_{\text{decision}} = \arg \max_{c_n} P(c_n|z_t) \end{aligned} \quad (6)$$

where the object estimated at time t is represented by c_{decision} . The belief threshold θ_{decision} permits to adjust the confidence level for the decision making process. Here, we have defined the belief threshold to the set of values $\{0.0, 0.05, \dots, 0.999\}$ to observe their effects on the accuracy of the object recognition process.

2.4 Intrinsic motivation for active exploration

Intelligent control of movements by an active exploration behaviour are achieved by the development of a computational method based on intrinsic motivation. It has been demonstrated by studies on cognitive development [9],[10] that intrinsic motivation is primordial to humans for engaging them to explore and manipulate their environment.

In this work, we use a predictive novelty motivation model, where interesting locations for exploration are those for which prediction errors are higher [12]. This is defined as follows:

$$I(\text{SM}(t)) = E_I(t - 1) \cdot E_I(t) \quad (7)$$

where the interesting location I for the sensorimotor state SM is obtained by the prediction error $E_I(t)$ at time t multiplied by the prediction error $E_I(t - 1)$ at time $t - 1$.

We define the prediction error $E_I(t)$ as the distance between the MAP from the Bayesian approach and the belief threshold value for making a decision:

$$E_I(t) = \arg \max_{c_n} P(c_n | z_t) - \theta_{\text{threshold}} \quad (8)$$

The active exploration performed by the robotic hand then is intelligently controlled by Equation 7, selecting the action for the highest SM state:

$$a = \arg \max_{\text{SM}} I(\text{SM}(t)) \quad (9)$$

where a is the action selected by the robotic hand. Figure 4 shows the process described to perform object exploration. This process composed by the Bayesian method and intrinsic motivation is repeated until the belief threshold is exceeded to make a decision about the object being explored.

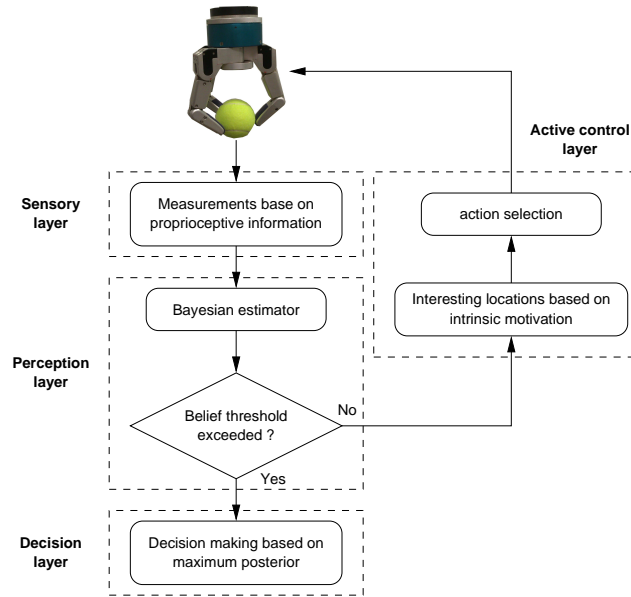


Fig. 4. Flow diagram with the steps required for the proposed intrinsically motivated active object exploration method. The robotic hand collects proprioceptive information from each contact performed. The robot is actively moved to interesting locations to improve perception based on an intrinsic motivations approach. Finally, a decision about the object being explored is made once the belief threshold is exceeded.




Fig. 5. Test objects used for the experiments in both simulated and real environments. The validation in simulated environment was performed using real data collected from these objects. For the validation in the real environment, the objects were placed and explored one at a time on a table.

3 RESULTS

In this section we present the results from the object exploration and recognition with passive and active modalities in simulated and real environments. Figure 5 shows the following objects used for validation: black triangle, red cylinder, blue ball, yellow ball, blue box and white box.

Object exploration in simulated environment: We developed an object exploration and recognition task using the data collected from the test objects (see Section 2.2) in a simulated environment. One dataset was used for training and four datasets for testing. The objects were randomly drawn from the testing datasets with 10,000 iterations for each belief threshold in the set of values $\{0.0, 0.05, \dots, 0.999\}$.

Passive object exploration: First, the simulated robot moved the hand and fingers around the object to obtain an initial belief of the object being explored. Next, the hand and fingers were randomly moved, accumulating evidence from each contact and making a decision once the current belief threshold was exceeded. The perception accuracy and reaction time were evaluated for each belief threshold.

Figure 6a shows the results in perception accuracy for the object exploration process with passive perception (red curve). It is observed that the robotic hand achieved the minimum perception error of 60% for a belief threshold of 0.75. Similarly, the reaction time which refers to the number of contacts required for making a decision with passive perception (red curve) is shown in Figure 6b. The number of contacts increased for large belief thresholds, where a maximum of ~ 2 contacts were required to make a decision for a belief threshold of ~ 0.999 . The results for perception accuracy and reaction time shown in Figure 6a and Figure 6b were obtained by averaging all perceptual classes over all trials for each belief threshold.

The confusion matrices (top panels) shown in Figure 7 permit to observe the performance of the classification accuracy with passive perception for each object and for different belief thresholds. These results show a slightly improvement of the classification accuracy with 68.28%, 71.77% and 76.18% for the belief thresholds of 0.0,

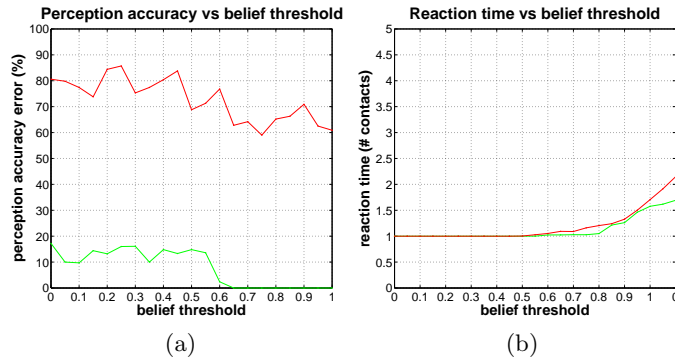
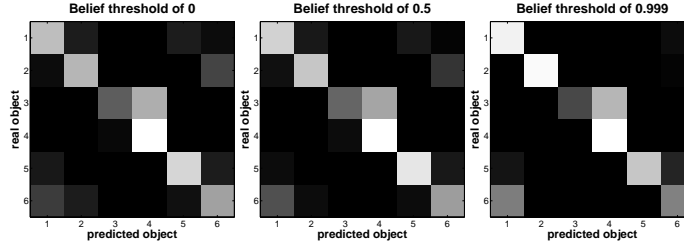


Fig. 6. Results from the passive and active object recognition in a simulated environment. Passive object recognition is presented by the red dotted-line. Active object recognition is presented by the green dotted-line. The experiment was performed for the set of belief threshold of $\{0.0, 0.05, \dots, 0.999\}$ with 10,000 iterations each. Results show the superiority of active over passive perception for object recognition with the robotic hand.

Passive object exploration in simulated environment



Active object exploration in simulated environment

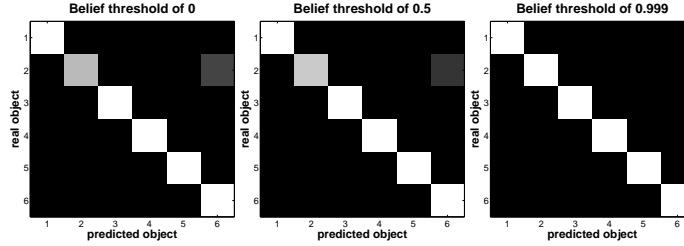


Fig. 7. Confusion matrices from the object recognition process with passive (top panels) and active (bottom panels) exploration modalities. The test objects used for the experiment are: 1) black triangle, 2) red cylinder, 3) blue ball, 4) yellow ball, 5) blue box and 6) white box. Validation results with passive perception shows a small improvement in the object recognition for large belief thresholds. Validation results with active perception show higher perception accuracy over passive perception.

0.5 and 0.999. These errors still can be improved by intelligent movements to interesting locations to reduce uncertainty.

Active object exploration: For the object recognition process with active perception, the robotic hand performed an exploration around the object to have an initial belief of the object being explored, similar to passive perception. Next, the robotic hand was actively moved, based on the proposed intrinsic motivation approach, towards interesting places around the object to improve perception. The active exploration process was repeated until the belief threshold was exceeded to make a decision. Similar to passive perception, the objects to be recognised were randomly drawn from the testing datasets with 10,000 iterations for each belief threshold in the set of values $\{0.0, 0.05, \dots, 0.999\}$.

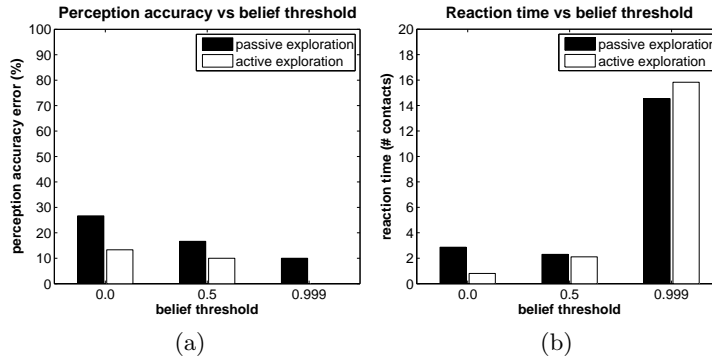


Fig. 8. (a) Perception accuracy and (b) reaction time results from the passive and active object recognition in a real environment. The experiment was performed with the belief thresholds of 0.0, 0.5 0.999. Passive perception was able to achieve the smallest error of 10% with 15 contacts for the belief threshold of 0.999. In contrast, active perception was able to achieve an error of 0% with 16 contacts for the belief threshold of 0.999. These results validate the performance of our proposed method for exploration.




Fig. 9. Confusion matrices from the object recognition process with passive (top panels) and active (bottom panels) exploration modalities in real environment. The test objects used for the experiment are: 1) black triangle, 2) red cylinder, 3) blue ball, 4) yellow ball, 5) blue box and 6) white box. Validation results with passive perception shows a small improvement in the object recognition for large belief thresholds, achieving an accuracy of 90% for the belief threshold of 0.999. Active perception shows higher perception accuracy of 100% for the belief threshold of 0.999.

The perception accuracy results from active exploration are represented by the green curve in Figure 6a. It is clearly observed the improvement in accuracy by actively moving the robotic hand towards interesting locations for exploration, achieving an error of 0% for the belief thresholds of 0.65 to 0.999. This result validates our proposed method for active exploration, and also shows its superiority over passive perception. The reaction time required for making a decision with active perception also is represented by the green curve in Figure 6b. We observe that the reaction time increases for large belief thresholds, where ~ 2 contacts are required for making a decision with a belief threshold of ~ 0.999 . These results were obtained by averaging all perceptual classes over all trials for each belief threshold.

The classification accuracy for each object is presented by the confusion matrices (bottom panels) in Figure 7 for different belief thresholds. It is observed how the accuracy is gradually improved, achieving a 95.49%, 96.41% and 100.0% for the belief thresholds of 0.0, 0.5 and 0.999 respectively. The accuracy obtained by actively exploring an object is clearly superior to the passive exploration process.

Object exploration in real environment: To validate our methods in a real environment, we implemented the object exploration and recognition task with the robotic platform described in Section 2.1. For this experiment, we used the objects from the validation in simulated environment.

Passive object exploration: For the passive object exploration and recognition, the test objects were placed on a table one at a time. The robotic hand performed an exploration around the object through a fixed set of movements, building an initial belief of the object being explored. Next, the robotic hand started the random action selection of exploration movements, accumulating evidence to reduce uncertainty from the object being explored. The exploration process was repeated until the belief threshold was exceeded, making a decision about the current object.

Perception accuracy results are shown in Figure 8a for different belief thresholds. We observe that the error achieved for the object recognition process is improved with 26.66%, 16.66% and 10.0% for the belief threshold of 0.0, 0.5 and 0.999 respectively. The reaction time results required for making a decision are presented in Figure 8b. This result shows that for achieving the smallest error of 10% with passive perception, it was required ~ 15 contacts, whilst for the largest error of 26.66% it was required ~ 3 contacts by the robotic hand. These results still can be improved by the use of our proposed method for exploration.

The classification accuracy for each object based on passive perception is presented by the confusion matrices (top panels) in Figure 9. The exploration task achieved the perception accuracies of 73.33%, 83.33% and 90.0% for the belief threshold of 0.0, 0.5 and 0.999 respectively.

Active object exploration: For the validation of the active exploration in a real environment, the test objects were placed on a table and explored by the robotic hand through a fixed set of movements. This step permitted to construct an initial belief of the object being explored. On the contrary to passive perception, here the robotic hand was able to selected the next action movement towards an interesting location around the object to improve perception. A decision about the object being explored was made once the evidence accumulated exceeded the belief threshold.

Figure 8a shows the perception accuracy results for the active exploration. We observe that the errors achieved for the object recognition process is improved with 13.33%, 10.0% and 0.0% for the belief thresholds of 0.0, 0.5 and 0.999 respectively. The reaction times required for making a decision are presented in Figure 8b. It is clearly observed that to achieve the best error of 0.0% it was required 16 contacts, whilst the error of 13.33% was obtained with 1 contact.

The classification accuracy for each object based on active perception is presented by the confusion matrices (bottom panels) in Figure 9. The exploration task achieved the perception accuracies of 86.66%, 90.0% and 100.0% for the belief thresholds of 0.0, 0.5 and 0.999 respectively. These results are improved over the accuracies obtained by passive perception. On the one hand, these results in simulated and real environments demonstrate the benefits of active over passive perception. On the other hand, they also validate the accuracy of our proposed method for tactile perception and exploration in autonomous robotic systems.

4 CONCLUSIONS

In this work we presented a method for object recognition using active exploration with a robotic hand under the presence of uncertainty. Our active exploration method, composed by a probabilistic method and an intrinsic motivation approach, was able to achieve accurate results.

We used a set of test objects for training and testing our methods in simulated and real environments. Tactile sensing was used for contact detection, whilst proprioceptive information composed by the position of the fingers and orientation of a robotic hand was used for object recognition. The robotic hand performed 30 contacts around each test object, which was repeated five times, to have one training dataset and four testing datasets.

A Bayesian method for uncertainty reduction through the interaction with an object was presented. This approach, together with a sequential analysis method, permitted the robotic hand to autonomously decide, based on exceeding a belief threshold, when to finish the exploration and make a decision about the object being explored.

Active exploration behaviour was obtained with an intrinsic motivation approach by moving the robotic hand towards the more interesting locations for exploration. Interesting locations were represented as the locations with large variances, obtained from the distance between the posterior probability from the Bayesian approach and the belief threshold. The combination of Bayesian and intrinsic motivation approaches allowed to develop an active exploration behaviour, accumulating evidence and reducing uncertainty by exploring the most interesting locations of the object.

Our method was validated in simulated and real environments using passive and active exploration modalities. In simulated environment and active exploration the robotic hand achieved a perception error of 0% for belief thresholds from 0.65 to 0.99. These results contrast with the error of 60% for the belief threshold of 0.75 with passive exploration (Figure 6a). We did not observed large differences for the reaction time with both exploration modalities, where ~ 2 contacts were required to make a decision for the smallest perception errors (Figure 6b).

The validation in a real environment also shows the high accuracy achieved by the robotic hand using our proposed method. For active perception, the smallest error of 0% was achieved by the robotic hand with a belief threshold of 0.999 (Figure 8a). For passive perception, the smallest error of 10% was achieved for the belief threshold of 0.999. Similar to the validation in the simulated environment, the reaction time required to make a decision for the best accuracies did not present large differences, with 15 and 16 contacts for passive and active perception respectively. The validations from simulated and real environments show the benefits of our proposed method for object exploration.

Overall, we have observed how active movements performed by the robotic hand to explore interesting locations, improve the perception accuracy and decision making for an autonomous exploration task. For future

work, we plan to extend our methods combining them with vision and implementing them with more complex robots to autonomously perceive and explore their environment.

Acknowledgements. This work was supported by EU Framework project WYSIWYD (FP7-ICT-2013-10). We thank to the technical staff and the equipment provided by the Department of Psychology and the Sheffield Robotics Laboratory at the University of Sheffield.

References

- [1] S. J. Lederman and R. L. Klatzky. Hand movements: A window into haptic object recognition. *Cognitive Psychology*, vol. 19, no. 19, pp. 342-368, 1987.
- [2] R. Bajcsy, S. Lederman and R. Klatzky. Object exploration in one and two fingered robots. *Computer Society Press*, 1987, pp. 1806-1810.
- [3] R. Bajcsy. Active perception. *Proceedings of the IEEE*, vol. 76, no. 8, pp. 966-1005, 1988.
- [4] T. Prescott, M. Pearson, B. Mitchinson, J. Sullivan and A. Pipe. Whisking with robots. *Robotics Automation Magazine, IEEE*, vol. 16, no. 3, pp. 42-50, 2009.
- [5] U. Martinez-Hernandez, T. Dodd, T. Prescott and N. Lepora. Active Bayesian perception for angle and position discrimination with a biomimetic fingertip. In *(IROS), 2013 IEEE International Conference on*, 2013, pp. 5968-5973.
- [6] U. Martinez-Hernandez, T. Dodd, L. Natale, G. Metta, T. Prescott and N. Lepora. Active contour following to explore object shape with robot touch. In *World Haptics Conference (WHC)*, 2013, pp. 341-346.
- [7] N. F. Lepora, U. Martinez-Hernandez and T. J. Prescott. A SOLID case for active Bayesian perception in robot touch. In *Biomimetic and Biohybrid Systems*. Springer, 2013, pp. 154-166.
- [8] N. Lepora, U. Martinez-Hernandez and T. Prescott. Active Bayesian perception for simultaneous object localization and identification. In *Robotics: Science and Systems*, 2013.
- [9] R. M. Ryan and E. L. Deci. Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology*, vol. 25, no. 1, pp. 54-67, 2000.
- [10] A. G. Barto, S. Singh and N. Chentanez. Intrinsically motivated learning of hierarchical collections of skills. In *Proc. 3rd Int. Conf. Development Learn*, 2004, pp. 112-119.
- [11] A. Baran'es and P.-Y. Oudeyer. R-iac: Robust intrinsically motivated exploration and active learning. *Autonomous Mental Development, IEEE Transactions on*, vol. 1, no. 3, pp. 155-169, 2009.
- [12] P.-Y. Oudeyer and F. Kaplan. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, vol. 1, 2007.
- [13] A. Stout, G. D. Konidaris and A. G. Barto. Intrinsically motivated reinforcement learning: A promising framework for developmental robot learning. *DTIC Document, Tech. Rep.*, 2005.
- [14] Z. Pezzementi, E. Plaku, C. Reyda and G. D. Hager. Tactile-object recognition from appearance information. *Robotics, IEEE Transactions on*, vol. 27, no. 3, pp. 473-487, 2011.
- [15] S. Ratnasingam and T. McGinnity. Object recognition based on tactile form perception. In *Robotic Intelligence In Informationally Structured Space (RiiSS)*, 2011 IEEE Workshop on, April 2011, pp. 26-31.
- [16] M. Johnsson and C. Balkenius. Neural network models of haptic shape perception. *Robotics and Autonomous Systems*, vol. 55, no. 9, pp. 720-727, 2007.
- [17] P. Fitzpatrick, G. Metta and L. Natale. Yet another robot platform. Website, <http://eris.liralab.it/yarpdoc/index.html>.
- [18] 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*, vol. 27, no. 3, pp. 389-400, 2011.