A SOLID Case for Active Bayesian Perception in Robot Touch

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

Sheffield Center for Robotics (SCentRo), University of Sheffield, UK
n.lepora@sheffield.ac.uk

Abstract. In a series of papers, we have formalized a *Bayesian perception* approach for robotics based on recent progress in understanding animal perception. The main principle is to accumulate evidence for multiple perceptual alternatives until reaching a preset belief threshold, formally related to sequential analysis methods for optimal decision making. Here, we extend this approach to active perception, by moving the sensor with a control strategy that depends on the posterior beliefs during decision making. This method can be used to solve problems involving Simultaneous Object Localization and IDentification (SOLID), or 'where and what'. Considering an example in robot touch, we find that active perception gives an efficient, accurate solution to the SOLID problem for uncertain object locations; in contrast, passive Bayesian perception, which lacked sensorimotor feedback, then performed poorly. Thus, active perception can enable robust sensing in unstructured environments.

Keywords: Active perception, tactile sensing, localization, robotics.

1 Introduction

Twenty five years after Bajcsy's landmark paper on active perception [1], it remains the case that most machine perception involves static analysis of passively sampled data. Certainly, there has been progress on passive approaches to pattern recognition in relation to machine learning and uncertainty, and there is a diverse body of work on active vision; nevertheless, a search through recent progress in robot vision, audition or touch reveals the majority of papers still rely on wholly forward perceptual processes without any sensorimotor feedback.

Why this slow uptake, when early arguments for active control of perception were compelling [1, 2] and, as Bajcsy said, it should be axiomatic that perception is active? One factor might be the required complexity of the robot hardware, which must involve actuated sensors and sensorimotor control loops. However, this should not be a barrier, because the technology is readily available and many standard robots have these capabilities, *e.g.* the iCub [3]. A more likely explanation is that researchers have focussed on sensing problems, such as identification, that can be solved adequately in many scenarios without introducing active methods for sensorimotor control. That being said, conventional robotics is reaching an impasse with present methods, such as poor performance in unstructured environments, which is preventing wider robot utilization beyond traditional factory settings [4].

N.F. Lepora et al. (Eds.): Living Machines 2013, LNAI 8064, pp. 154–166, 2013.

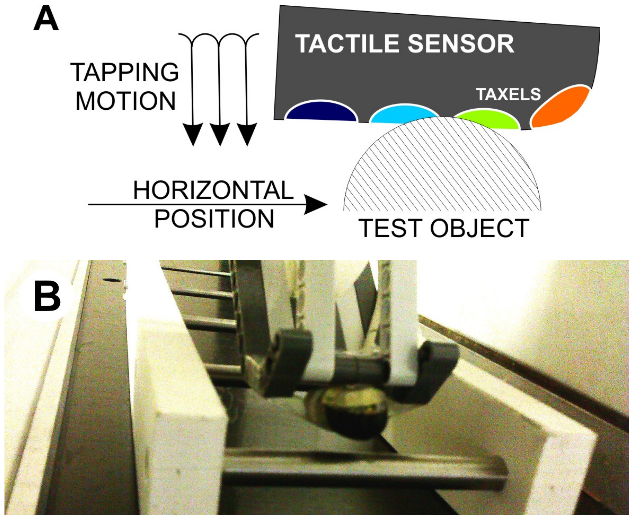

Fig. 1. Experimental setup. (A) Schematic of tactile sensor tapping against a cylindrical test object: the fingertip taps down and then back up again to press its pressure-sensitive taxels (colored) against the test object; each tap is then followed by a small horizontal move to sample object contacts over a range of positions. (B) Forward view of the experiment showing the fingertip mounted on the arm of the Cartesian robot. This experimental setup is ideal for systematic data collection to characterize the properties of the sensor interacting with its environment.

In a series of papers [5–10], we have formalized an approach for robot perception based on recent progress in understanding animal perception [11, 12]. The main principle is to accumulate evidence for multiple perceptual alternatives until reaching a preset belief threshold that triggers a decision, formally related to Bayesian sequential analysis methods for optimal decision making [13]. Here we describe how this perception approach extends naturally from passive to active perception and some implications of this theory of active Bayesian perception.

Our proposal for active Bayesian perception is tested with a simple but illustrative task of perceiving the location (horizontal position) and identity (diameter) of a test rod using tapping movements of a biomimetic fingertip at unknown contact location (Fig. 1; the colored regions are the pressure-sensitive taxels). We demonstrate first that passive perception can solve this task, but the perceptual acuity and reaction time depend strongly on the location of the fingertip relative to the rod. We then show that an active 'fixation point' control strategy can substantially improve the robustness, accuracy and speed of the perception, by moving the fingertip to locations with good perception independent of the starting position. Thus we demonstrate that active perception can enable appropriate perceptual decision making in an unstructured environment.

Related arguments have been presented in two other papers: the active perception method has been applied to texture identification under unknown contact

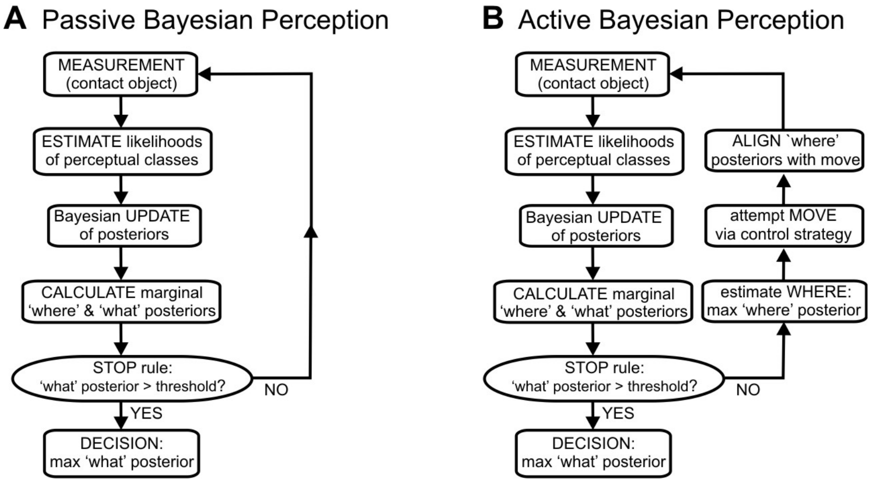

Fig. 2. Passive and active Bayesian perception applied to simultaneous object localization and identification. (A) Passive Bayesian perception has a recursive Bayesian update to give the marginal 'where' and 'what' posterior beliefs, with decision termination at sufficient 'what' belief. (B) Active Bayesian perception has the same recursive belief update, while also actively controlling the sensor location according to a strategy based on those beliefs; furthermore, when the sensor moves, it is necessary to re-align the 'where' component of the beliefs with the new sensor location. The two algorithms differ only in the sensorimotor control loop for active Bayesian perception.

depth [9], and a more detailed, systematic treatment of SOLID is given in [10] for a 2D (horizontal and vertical) 'where' and 'what' scenario.

2 Methods

The main goal of this work is to advance our understanding of the role of active perception for situated agents that seek to determine the 'where' and 'what' properties of objects. We refer to the computational task that must then be solved by Simultaneous Object Localization and IDentification (SOLID), to emphasize a similarity with SLAM of having two interdependent task aims, in that knowledge of location aids the computation of identity (mapping) and similarly that knowledge of object identity (mapping) aids localization.

Passive Bayesian perception accumulates belief for distinct 'where' and 'what' classes by making successive taps against a test object until at least one of the marginal 'what' posterior beliefs crosses a belief threshold, when a 'where' and 'what' decision is made. The passive nature of the perception means that the 'where' position class is constant over this process (Fig. 2A).

Active Bayesian perception also accumulates belief for the 'where' and 'what' perceptual classes by successively tapping until reaching a predefined 'what' belief threshold. In addition, it utilizes a sensorimotor loop to move the sensor according to the online marginal posterior beliefs during the perceptual process (Fig. 2B). For example, the sensor could be controlled with a 'fixation point' strategy, in which the marginal 'where' beliefs are used to infer a best estimate for current location and thus a relative move towards a preset 'good' target position on the object to improve the perceptual decision making.

2.1 Algorithms for Bayesian Perception

Our algorithm for active Bayesian perception is based on including a sensorimotor feedback loop in an existing method for passive Bayesian perception [5]. Both methods assume that the sensor makes a discrete contact measurement (here a tap) onto an object, from which the joint likelihoods of object location and identity are used to update the prior to posterior beliefs for those perceptual classes. In active perception, a control strategy repositions the sensor before each contact, taking input from the updated beliefs and outputting the sensor move.

Because these methods are applicable to any simultaneous object localization and identification task, this section is presented in a general 'where' and 'what' notation. A general SOLID task has $N_{\rm loc}$ distinct 'where' location classes x_l and $N_{\rm id}$ distinct 'what' identity classes w_i, totalling $N = N_{\rm loc}N_{\rm id}$ joint 'where-what' classes $c_n = (x_l, w_i)$. Each contact against a test object gives a multi-dimensional time series of sensor values $z = \{s_k(j) : 1 \le j \le N_{\rm samples}, 1 \le k \le N_{\rm channels}\}$, with indices j, k labeling the time samples and sensor channels. The tth contact in a sequence is denoted by z_t with $z_{1:t-1} = \{z_1, \cdots, z_{t-1}\}$ its contact history.

Measurement model and likelihood estimation: The likelihoods of all perceptual classes are found using a measurement model of the contact data, which we find by applying a histogram method to training examples for each perceptual class [5, 6]. First, the sensor values s for channel k are binned into $N_{\rm bins} = 100$ intervals, with sampling distribution for each perceptual class c_n given by the normalized histogram over all training data in that class:

$$P(b|c_n, k) = \frac{h(b,k)}{\sum_{b=1}^{N_{\rm bins}} h(b,k)}, \qquad 1 \le k \le N_{\rm channels}, \tag{1}$$

where $h(b, k)$ is the histogram count for bin b ($1 \le b \le N_{\rm bins}$) in sensor channel k. Then, given a test tap z, we construct a measurement model from the mean log likelihood over all samples in that tap

$$\log P(z|c_n) = \frac{1}{N_{\rm samples}N_{\rm channels}} \sum_{k=1}^{N_{\rm channels}} \sum_{j=1}^{N_{\rm samples}} \log P(b_k(j)|c_n, k), \tag{2}$$

where $b_k(j)$ is the bin occupied by sample $s_k(j)$. Technically, this measurement model becomes ill-defined if any histogram bin is empty, which is easily fixed by regularizing the bin counts with a small constant ($\epsilon \ll 1$), giving $h(b, k) + \epsilon$.

Bayesian update: Bayes' rule is used after each successive test contact z_t to recursively update the posterior beliefs $P(c_n|z_{1:t})$ for the perceptual classes with the estimated likelihoods $P(z_t|c_n)$ of that contact data

$$P(c_n|z_{1:t}) = \frac{P(z_t|c_n)P(c_n|z_{1:t-1})}{P(z_t|z_{1:t-1})}, \tag{3}$$

from background information given by the prior beliefs $P(c_n|z_{1:t-1})$. The marginal probabilities are also conditioned on the preceding contacts $z_{1:t-1}$ and given by

$$P(z_t|z_{1:t-1}) = \sum_{n=1}^{N} P(z_t|c_n)P(c_n|z_{1:t-1}). \tag{4}$$

Iterating (3,4), a sequence of contacts $z_1, \cdots, z_t$ results in a sequence of posteriors $P(c_n|z_1), \cdots, P(c_n|z_{1:t})$ initialized from uniform priors $P(c_n|z_0) := P(c_n) = \frac{1}{N}$.

Marginal 'where' and 'what' posteriors: For the following methods, we will need the posterior beliefs for just location or identity, rather than the joint beliefs considered so far. Because each class $c_n = (x_l, w_i)$ has a 'where' location x_l and 'what' identity w_i component, these beliefs can be found by marginalizing

$$P(x_l|z_{1:t}) = \sum_{i=1}^{N_{\rm id}} P(x_l, w_i|z_{1:t}), \tag{5}$$

$$P(w_i|z_{1:t}) = \sum_{l=1}^{N_{\rm loc}} P(x_l, w_i|z_{1:t}), \tag{6}$$

with the 'where' location beliefs given from summing over all 'what' identity classes w_i and the 'what' identity beliefs over all 'where' location classes x_l.

Final decision on the 'what' posteriors: Here we follow sequential analysis methods for optimal decision making that recursively update beliefs up to a threshold that triggers the final decision [13], as used in passive Bayesian perception [5]. The update stops when the marginal 'what' identity belief passes a threshold, giving a final decision from the maximal *a posteriori* (MAP) estimate

$$\text{if any } P(w_i|z_{1:t}) > \theta_{\rm id} \text{ then } w_{\rm id} = \arg\max_{w_i} P(w_i|z_{1:t}). \tag{7}$$

This belief threshold $\theta_{\rm id}$ is a free parameter that adjusts the balance between decision speed and accuracy. For $N = 2$, this speed-accuracy balance can be proved optimal [13]; optimality is not known for the many perceptual choices considered here, and so we make a reasonable assumption of near optimality [5].

Move decision on the 'where' posteriors: Analogously to the stop decision, a sensor move requires a marginal 'where' location belief to cross its own decision threshold, with the MAP estimate giving the 'where' location decision

$$\text{if any } P(x_l|z_{1:t}) > \theta_{\rm loc} \text{ then } x_{\rm loc} = \arg\max_{x_l} P(x_l|z_{1:t}). \tag{8}$$

Here we consider two particular cases (Figs 2A,B), termed:

(A) passive perception: $\theta_{\rm loc} = 1$ (never moves)

(B) active perception: $\theta_{\rm loc} = 0$ (always tries to move).

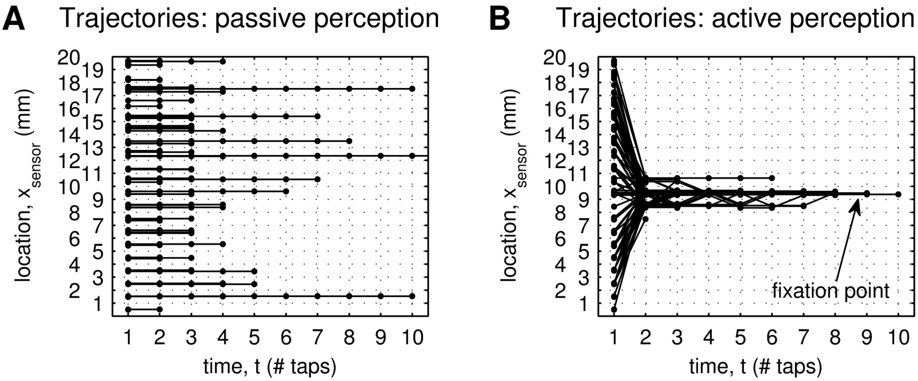

Fig. 3. Example trajectories for passive and active perception. 100 trajectories were selected randomly for each case. (A) Passive perception, with location (x-position) constant over time. (B) Active perception, with trajectories converging rapidly on the central fixation point (10 mm location class) independent of starting position.

For simplicity, we consider a basic movement strategy in which the sensor move Δ depends only on estimated location $x_{\rm loc}$, although more complex strategies are encompassed by the formalism. Whatever the strategy, the marginal 'where' location belief should be kept aligned with the sensor by shifting the joint 'where-what' posterior beliefs upon each move

$$P(x_l, w_i|z_{1:t}) \leftarrow \begin{cases} P(x_l - \Delta(x_{\rm loc}), w_i|z_{1:t}), & 1 \leq x_l - \Delta(x_{\rm loc}) \leq N_{\rm loc}, \\ p_0, & \text{otherwise}, \end{cases} \tag{9}$$

where we recalculate the beliefs p_0 lying outside the original range by assuming they are uniformly distributed and the shifted beliefs sum to unity. The left arrow denotes that the quantity on the left is replaced with that on the right.

Active control strategy: The final component of the active perception algorithm is to define the control strategy for moving the sensor based on the posterior beliefs. For simplicity, here we consider a 'fixation point' strategy motivated by orienting movements in animals: the sensor attempts to move to a predefined fixation point $x_{\rm fix}$ relative to the object assuming it is at the estimated location $x_{\rm loc}$ on the object, with each move resulting in

$$x_{\rm sensor} \leftarrow x_{\rm sensor} + \Delta\left(x_{\rm loc}\right), \quad \Delta(x_{\rm loc}) = x_{\rm fix} - x_{\rm loc}, \tag{10}$$

where $x_{\rm sensor}$ is the actual (unknown) location of the sensor. In practise, only the move Δ need be found, to instruct the sensor how to change location. Example trajectories resulting from this active control strategy are shown in Fig. 3B.

2.2 Data Collection and Analysis

The tactile sensors used in this study have a rounded shape that resembles a human fingertip [14], of dimensions 14.5 mm long by 13 mm wide. They consist

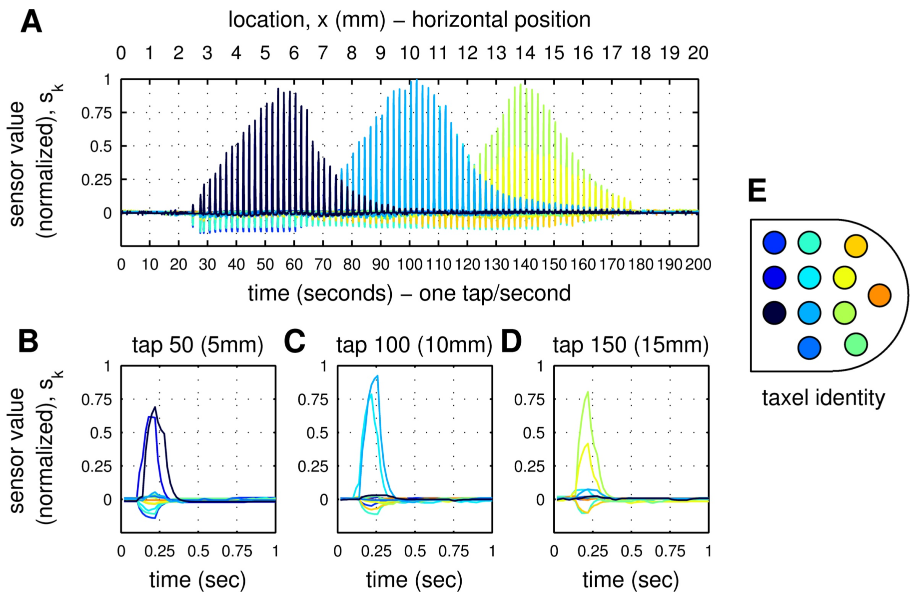

Fig. 4. Tactile dataset (for test rod of diameter 4 mm). (A) Entire dataset, with 200 taps over positions spanning 20 mm. Taps are every 0.1 mm displacement. (B-D) Individual tap data taken from panel A. (E) Taxel layout with color-code for plots A-D.

of an inner support wrapped with a flexible printed circuit board containing $N_{\text{channels}} = 12$ conductive patches for the touch sensor 'taxels'. These are coated with non-conductive foam and conductive silicone layers that together comprise a capacitive touch sensor that detects pressure by compression. Data were collected at 8 bit resolution and 50 cycles/sec then normalized and high-pass filtered [14].

The present experiments test the capabilities of the tactile fingertip mounted on a Cartesian robot. This robot moves the sensor in a horizontal/vertical plane in a precise and controlled way onto various test stimuli (~20 μm accuracy), and has been used for testing various tactile sensors [15]. The fingertip was mounted at an orientation appropriate for contacting axially symmetric shapes such as cylinders aligned along an axis perpendicular to the plane of movement (Fig. 1). $N_{\text{id}} = 5$ smooth steel rods with diameters 4,6,8,10,12 mm were used as test objects, mounted with their centers offset vertically (by 4,3,2,1,0 mm) to align their closest point of contact with the fingertip in the direction of tapping.

Touch data were collected while the fingertip tapped vertically onto and off each test object, followed by a horizontal move $\Delta x = 0.1$ mm across the closest face of the object (Fig. 1A). The fingertip was oriented so that it initially contacted the rod at its base and finally at its tip. A horizontal x-range of 20 mm was used, giving 200 taps for each of the $N_{\text{id}} = 5$ objects, or 1000 taps in total. From each tap of the fingertip against the object, a 1 sec time series of pressure readings ($N_{\text{samples}} = 50$) was extracted for all $N_{\text{channels}} = 12$ taxels (Fig. 4). All data were collected twice to give distinct training and test sets.

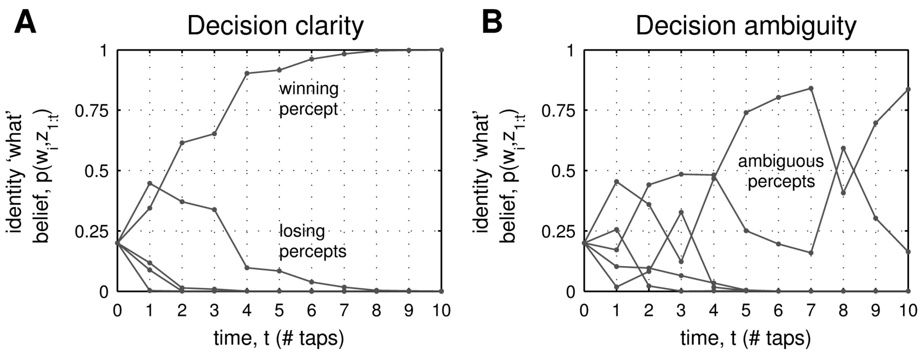

Fig. 5. Example 'what' belief update for perceptual decision making. Evidence from successive taps is integrated to result in accumulating/depreciating marginal beliefs for the $N_{\rm id} = 5$ distinct identity 'what' percepts. The examples show: (A) one clear winning percept and (B) two ambiguous percepts. Using a belief threshold to trigger the decision results in the appropriate number of taps to have a clear winner.

For analysis, the data were separated into $N_{\rm loc} = 20$ distinct location classes, by collecting groups of 10 taps each spanning 1 mm of the 20 mm x-range (tick-marks on Fig. 4A). In total, there were thus $N = N_{\rm loc}N_{id} = 100$ distinct 'where-what' perceptual classes. These were used to set up a 'virtual environment' in which methods for perception could be compared off-line on identical data. A Monte Carlo validation ensured good statistics, by averaging perceptual acuities over many test runs with taps drawn randomly from the perceptual classes (typically 20000 runs per data point in results). Perceptual acuities $e_{\rm loc}$, $e_{\rm id}$ were quantified using the mean absolute error (MAE) between the actual $x_{\rm test}$, $w_{\rm test}$ and classified values $x_{\rm loc}$, $w_{\rm id}$ of object location and identity over the test runs.

3 Results

3.1 Evidence Accumulation for Robot Perception

The 'where' and 'what' perceptual task is to find the location (x-position) of a rod and its identity (diameter) using tactile fingertip data over a sequence of test taps. Example 'what' perceptual beliefs $P(w_i|z_{1:t})$ for tap sequences $z_{1:t}$ of clear and ambiguous data are shown above (Fig. 5). These beliefs begin at equality corresponding to uniform prior beliefs and then evolve smoothly with some rising gradually towards unity and others falling towards zero. In the first example (Fig. 5A), the decision given by the largest perceptual belief remains the same after applying 2 taps or more, while the second example (Fig. 5B) flips between the two leading choices.

There are two common methods for making decisions from sequential data of this type: (i) set in advance the number of taps that will be used, or (ii) set in advance a belief threshold θ that will trigger the decision, so that the decision (reaction) time is a dynamic quantity that depends on the data received.

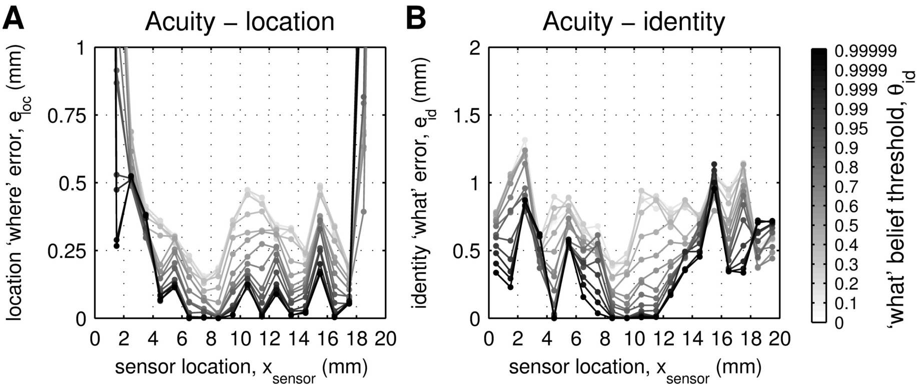

Fig. 6. Dependence of passive perception on sensor location and 'what' belief threshold. The 'where' location errors are shown in (A) and the 'what' identity errors in (B), plotted against sensor location x_{sensor}. The gray-scale denotes the 'what' belief threshold. Each data point corresponds to 1000 decision trials. Perceptual performance improves in the center of the sensor location range and at greater belief thresholds.

Recent progress in perceptual neuroscience strongly supports that animals use a belief threshold to make decisions [11], consistent with the brain implementing sequential analysis for optimal decision making [13]. In accordance, a comparison of these two methods on tactile robot data found that the belief threshold method gave superior performance in perceptual acuity [5]. This can be seen intuitively from Fig. 5: if, for example, a deadline of 10 taps was set in advance, then the decision is unnecessarily slow in situations of clarity (Fig. 5A) and too quick in situations of ambiguity (Fig. 5B). Instead, setting a belief threshold allows the decision time to adjust dynamically to the uncertainty of the evidence.

Both the passive and active methods for perception considered here update beliefs from successive test taps to threshold θ_{id}. They differ, though, in how the sensor responds during the decision process: for passive perception its location is fixed, whereas for active perception it can control changes in location.

3.2 Passive Perception of Location and Identity

This section considers the application of passive Bayesian perception to the 'where' and 'what' perceptual task of identifying rod location (x-position) and identity (diameter). Results are generated with a Monte Carlo procedure using test data as a virtual environment (Sec. 2.2), such that each contact tap passively remains within its initial location class on the object (examples in Fig. 3A).

The 'where' and 'what' decisions for perceiving object location and identity were evaluated over identity belief thresholds θ_{id} from 0.1 to 0.99999 (Fig. 6, colorbar). Initial locations spanned all horizontal position classes x_l. Mean perceptual acuities over all objects improved with identity threshold and towards the center (10 mm) of the location range (Fig. 6), giving minimal errors near zero for identity e_{id} and location e_{loc}. The number of taps to reach a decision also

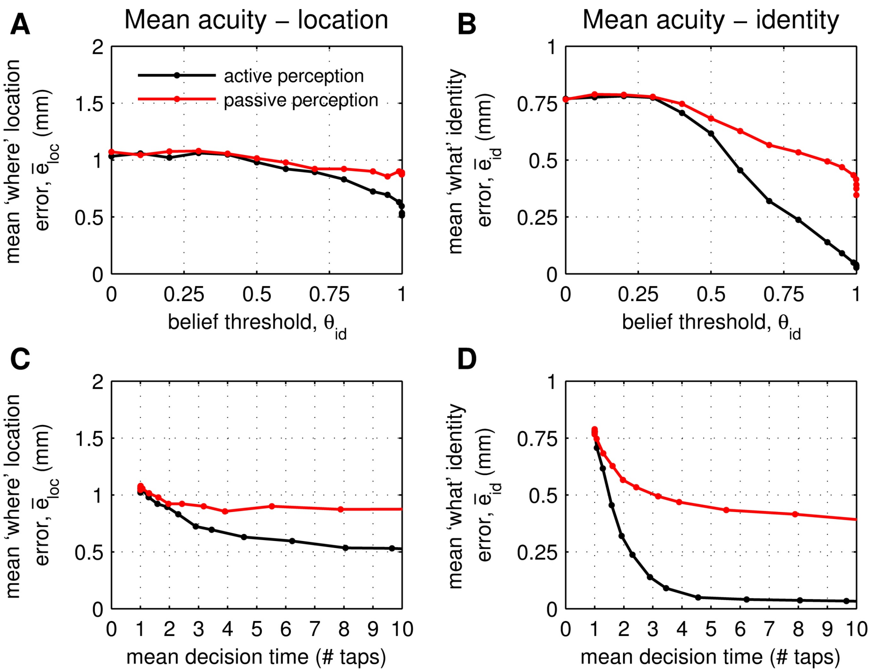

Fig. 7. Comparative perceptual acuity for active and passive Bayesian perception. (A,B) Dependence of the mean absolute errors of location $\bar{e}_{\rm loc}$ (x-position) and identity $\bar{e}_{\rm id}$ (rod diameter) upon the identity belief threshold θ_{id}. Passive perception is shown in red and active perception in black. (C,D) Dependence of these perceptual errors upon mean decision time (with threshold an implicit parameter). Active performs better than passive perception, and both improve with increasing 'what' belief threshold.

varied, such that the mean tap number (reaction time) increased with identity threshold and decreased towards the central location range (Figure not shown).

Here we consider situations in which the initial contact location is unknown, with performance measure the mean perceptual errors over all such locations, consistent with location being selected randomly. These mean errors reached their best values $\bar{e}_{\rm loc} \sim 0.9\,{\rm mm}$ and $\bar{e}_{\rm id} \sim 0.4\,{\rm mm}$ at the largest belief thresholds and reaction times (Fig. 7, red plots). Not unexpectedly, the perception is poorer than when choosing just the central contact location (on Fig. 6 at $x_l = 10\,{\rm mm}$), emphasizing that passive perception performs poorly because it cannot control contact location.

3.3 Active Perception of Location and Identity

This section considers active Bayesian perception in the same scenario as for passive perception, with a 'where' and 'what' perceptual task of identifying rod location and identity. Results are again generated using a Monte Carlo procedure over a virtual environment (Sec. 2.2), now with an active control strategy that tries to re-locate the sensor to a fixation point in the center of the location range (10 mm) where the passive perception was best (example trajectories in Fig. 3B).

The 'where' and 'what' decisions for perceiving object location and identity were evaluated over the same range of belief thresholds ($\theta_{\rm id}$ from 0.1 to 0.99999) and locations (x_l from 0 to 20 mm) as for passive perception, to permit direct comparison of the active and passive approaches. We again considered the location as uncertain, and thus measured performance by the mean perceptual errors $\bar{e}_{\rm loc}$, $\bar{e}_{\rm id}$ over all initial contact locations. These location-averaged errors reached their best acuities for object location of $\bar{e}_{\rm loc} \sim 0.5$ mm and identity $\bar{e}_{\rm id} \sim 0.1$ mm at the largest belief thresholds and reaction times (Fig. 7, black plots).

Comparing active perception with passive perception, the best mean perceptual errors improve from $\bar{e}_{\rm loc} \sim 0.9$ mm to 0.5 mm for object location (over a 20 mm x-position range) and $\bar{e}_{\rm id} \sim 0.4$ mm to 0.1 mm for object identity (over a 4-12 mm diameter range). Thus, active perception gives far finer perceptual acuity than passive perception when compared under similar conditions of uncertain object location and identity.

4 Discussion

In this paper, we compared active and passive Bayesian perception methods for Simultaneous Object Localization and IDentification (SOLID), or perceiving 'where' and 'what'. We considered a task in which a biomimetic fingertip taps against a smooth steel rod to simultaneously perceive its location (horizontal position) and identity (rod diameter). Active perception can control changes in location of the sensor during the decision making process, whereas for passive perception the location is fixed at where the sensor initially contacted the object. We found that active perception gives far more accurate perception in situations of uncertain object location and identity than passive perception. Thus, active perception is appropriate for sensing in unstructured environments where location is uncertain, and improves performance by compensating the uncertainty in initial sensor placement to, in effect, structure an unstructured environment.

As in other work on active perception [1, 2], the inspiration for our approach was from animal perception. The particular active perception strategy considered here was to fixate the sensor onto an object at a good contact location for perception (here centering the fingertip over the middle of the object); this strategy is analogous to orienting movements found in many perceptual modalities, such as saccadic foveation in vision and head turning for audition. In addition, we used an evidence accumulation method for Bayesian perception [5] that has close relation to leading models of perceptual decision making in neuroscience [11] and also relates to proposals for cortico-basal ganglia function [12]. Bayesian perception, where evidence is integrated to a belief threshold, leads to decision time being a dynamic quantity that depends on the quality of data. Although rare in contemporary approaches to machine learning, having this type of variation in reaction times is ubiquitous in the natural world.

Related methods to those presented here have enabled the first demonstration of hyperacuity in robot touch [8], giving localization acuity finer than the sensor resolution, as is common in animal perception including human touch [16]. Although our previous study also found that active perception helped attain hyperacuity, those methods now seem somewhat *ad hoc* in light of the present study, by not making best use of the 'where' and 'what' aspects of the problem. In the present work, we have developed a principled approach to active Bayesian perception of applicability to simultaneous object localization and identification. Given the taxel spacing is 4 mm, our results verify that passive Bayesian perception is capable of hyperacuity (mean localization error ~ 1 mm) and active Bayesian perception of stronger hyperacuity (mean error ~ 0.5 mm).

In seminal work on active perception, Bajcsy said it is axiomatic that perception (in animals) is active [1]. Robotics is currently in a state of transition from rigidly controlled tasks in predictable structured environments like factory assembly lines, to applications in unpredictable unstructured environments like our homes, hospitals and workplaces. In our opinion, robots will need active perception to accomplish these tasks in unstructured environments, and thus it may also become axiomatic that future robot perception will be active too.

Acknowledgments. We thank Ben Mitchinson and Mat Evans for comments on earlier drafts of this paper. This work was funded by the European Commission under the FP7 project grant EFAA (ICT-270490); UMH was supported by the Mexican National Council of Science and Technology (CONACyT).

References

1. Bajcsy, R.: Active perception. Proceedings of the IEEE 76, 966–1005 (1988)
2. Ballard, D.: Animate vision. Artificial Intelligence 48(1), 57–86 (1991)
3. Metta, G., Sandini, G., Vernon, D., Natale, L., Nori, F.: The icub humanoid robot: an open platform for research in embodied cognition. In: Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems, pp. 50–56 (2008)
4. Kemp, C.C., Edsinger, A., Torres-Jara, E.: Challenges for robot manipulation in human environments (grand challenges of robotics). IEEE Robotics & Automation Magazine 14(1), 20–29 (2007)
5. Lepora, N.F., Fox, C.W., Evans, M.H., Diamond, M.E., Gurney, K., Prescott, T.J.: 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)
6. Lepora, N.F., Evans, M., Fox, C.W., Diamond, M.E., Gurney, K., Prescott, T.J.: Naive bayes texture classification applied to whisker data from a moving robot. In: The 2010 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2010)
7. Lepora, N.F., Sullivan, J.C., Mitchinson, B., Pearson, M., Gurney, K., Prescott, T.J.: Brain-inspired bayesian perception for biomimetic robot touch. In: 2012 IEEE International Conference on Robotics and Automation (ICRA), pp. 5111–5116 (2012)
8. Lepora, N.F., Martinez-Hernandez, U., Barron-Gonzalez, H., Evans, M., Metta, G., Prescott, T.J.: Embodied hyperacuity from bayesian perception: Shape and position discrimination with an icub fingertip sensor. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 4638–4643 (2012)

9. Lepora, N.F., Martinez-Hernandez, U., Prescott, T.J.: Active touch for robust perception under position uncertainty. In: 2013 IEEE International Conference on Robotics and Automation (ICRA), pp. 3005–3010 (2013)
10. Lepora, N.F., Martinez-Hernandez, U., Prescott, T.J.: Active bayesian perception for simultaneous object localization and identification. In: Robotics: Science and Systems (2013)
11. Gold, J.I., Shadlen, M.N.: The neural basis of decision making. Annual Reviews Neuroscience 30, 535–574 (2007)
12. Lepora, N.F., Gurney, K.: The basal ganglia optimize decision making over general perceptual hypotheses. Neural Computation 24(11), 2924–2945 (2012)
13. Wald, A.: Sequential analysis. John Wiley and Sons, NY (1947)
14. Schmitz, A., Maiolino, P., Maggiali, M., Natale, L., Cannata, G., Metta, G.: Methods and technologies for the implementation of large-scale robot tactile sensors. IEEE Transactions on Robotics 27(3), 389–400 (2011)
15. Evans, M., Fox, C., Lepora, N., Pearson, M., Sullivan, J., Prescott, T.: The effect of whisker movement on radial distance estimation: a case study in comparative robotics. Frontiers in Neurorobotics 6 (2013)
16. Loomis, J.M.: An investigation of tactile hyperacuity. Sensory Processes 3, 289–302 (1979)