Proceedings Paper:
Lepora, N.F., Sullivan, J.C., Mitchinson, B. et al. (3 more authors) (2012) Brain-inspired Bayesian perception for biomimetic robot touch. In: Robotics and Automation (ICRA), 2012 IEEE International Conference on. 2012 IEEE International Conference on Robotics and Automation, 14-18 May 2012, St Paul, MN, USA. IEEE , pp. 5111-5116. ISBN 978-1-4673-1403-9

https://doi.org/10.1109/ICRA.2012.6224815

Brain-inspired Bayesian perception for biomimetic robot touch

Nathan F. Lepora, J. Charlie Sullivan, Ben Mitchinson, Martin Pearson, Kevin Gurney, Tony J. Prescott

***Abstract*— Studies of decision making in animals suggest a neural mechanism of evidence accumulation for competing percepts according to Bayesian sequential analysis. This model of perception is embodied in a biomimetic tactile sensing robot based on the rodent whisker system. We implement simultaneous perception of object shape and location using two psychological test paradigms: first, a free-response paradigm in which the agent decides when to respond, implemented with Bayesian sequential analysis; and second an interrogative paradigm in which the agent responds after a fixed interval, implemented with maximum likelihood estimation. A benefit of free-response Bayesian perception is that it allows tuning of reaction speed against accuracy. In addition, we find that large gains in decision performance are achieved with unforced responses that allow null decisions on ambiguous data. Thus, free-response Bayesian perception offers benefits for artificial systems that make them more animal-like in behavior.**

I. INTRODUCTION

A view being advanced in psychology and neuroscience is that human and animal perception corresponds to statistically optimal inference from noisy and ambiguous sensory data [1], [2]. One notable series of experiments considers neuronal activity in parietal cortex as monkeys make perceptual judgements of the direction of motion for drifting random dots, and finds individual neurons that noisily ramp-up their firing rates until reaching a threshold when a decision is made [3], [4]. These processes appear well described by the statistical approach of sequential analysis [5], [6], which applies Bayes' rule to accumulate evidence for competing perceptual hypotheses over time series of sensory data until a preset threshold for the posterior probability is reached [7].

This study implements these principles of biological perception in a biomimetic robot based on the rodent whisker system [8]. There are two main motivations for this research. First, animals far surpass present-day robots in their perceptual abilities. Thus, by utilizing biological principles one hopes to develop better methods for robot perception. Second, uncovering the principles underlying biological perception is a key goal of the biosciences. Hence, the development of biomimetic robot experiments to test these biological theories could give new, complementary insights to those from psychological and neuroscientific investigations of humans and animals. For both of these reasons, we use a robot that closely mimics the biological system in its interaction with objects that it is sensing.

Our principal hypothesis is that this type of Bayesian perception leads to flexible, efficient and accurate sensory processing in robots and animals. In this study, we demonstrate that Bayesian sequential analysis gives an effective approach for simultaneous perception of object shape and location. These are important percepts for manipulating objects – for example, knowing object shape is key for grasping, while the object location is required to position a manipulator. The present study also builds on recent work in which we showed that similar probabilistic methods for classifying textural stimuli gave improved performance over non-probabilistic methods commonly used in robotics [9].

A main benefit of Bayesian perception is that it allows the decision making procedure to respond more fully to the sensory processing. In particular, the tradeoff between speed and accuracy is set directly to permit caution when penalties for mistakes are high and fast reactions when delays are critical. Then, once this tradeoff is set, reaction times are faster for unambiguous percepts and slower with ambiguity. Bayesian perception also allows the option of a null decision on ambiguous data after a deadline, which we find can greatly increase both reaction speed and accuracy. Given that these behaviors are characteristic of animals, we expect that these Bayesian methods can improve the sensing abilities of artificial systems in a way that mimics animal perception.

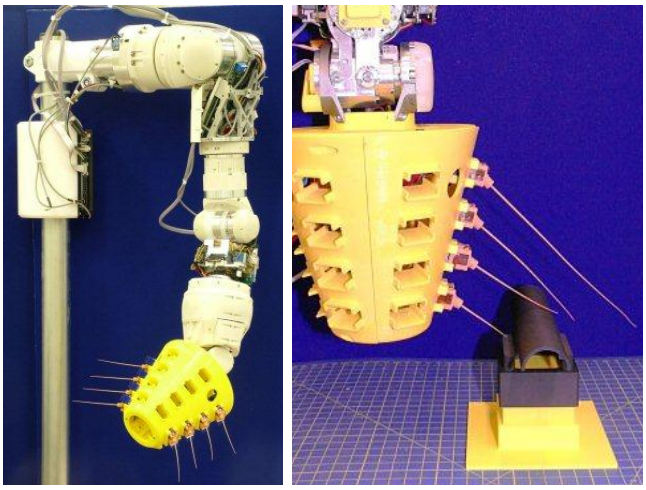

Fig. 1. (A) The BIOTACT sensor (yellow module) is mounted as an end-effector to a 7-dof robot arm. (B) Close-up of the sensor and whisker modules contacting a test object (30 mm diameter hemi-cylinder).

Principal contributions: NFL analyzed the data and wrote the paper. JCS and MP designed the whisker sensors which were constructed at BRL. BM wrote and implemented the control software. Everyone planned and designed the experiments which were performed by JCS.

This work was supported by EU Framework projects BIOTACT (ICT-215910) and EFAA (ICT-270490).

N. Lepora, B. Mitchinson, K. Gurney and T. Prescott are with the Adaptive Behavior Research Group, Department of Psychology, University of Sheffield, Western Bank, Sheffield S10 2TN, UK.
Email: {n.lepora, b.mitchinson, k.gurney, t.j.prescott}@sheffield.ac.uk.

J. Sullivan and M. Pearson are with the Bristol Robotics Laboratory, Du Pont Building, Bristol Business Park, Coldharbour Lane, Frenchay, Bristol BS16 1QD, UK. Email: {charlie.sullivan, martin.pearson}@brl.ac.uk

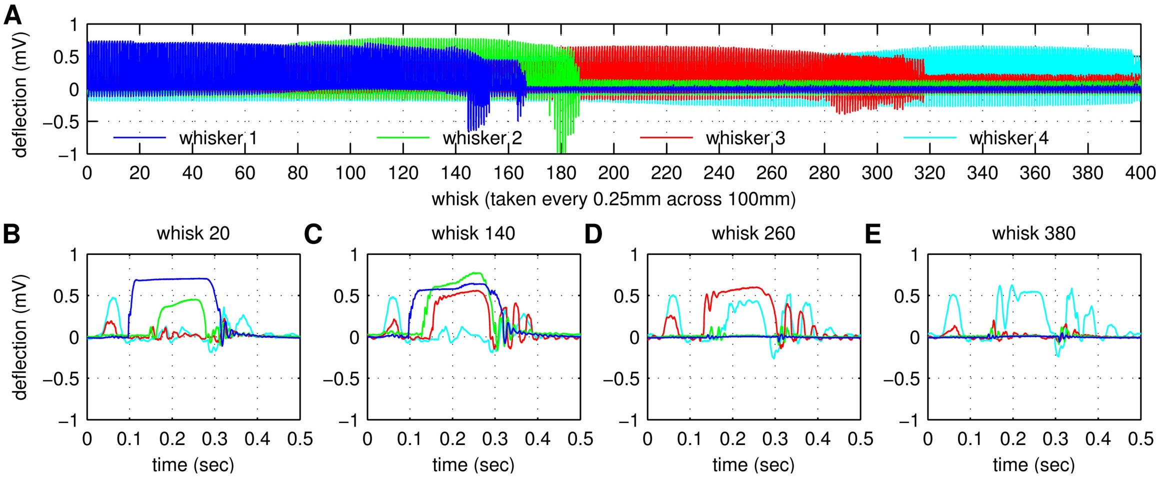

Fig. 2. (A) Whisker deflection data recorded as the robot head moves slowly over a test hemi-cylinder (diameter 50 mm) while actively whisking against the surface at a constant frequency of 2 whisks/s. The range of head positions spanned 100 mm over 202 s, giving 404 whisks spaced every 0.25 mm. Tickmarks are shown every 5 mm displacement, or 20 whisks. Data from the four whiskers are represented in distinct colors. (B-E) Examples of deflection data for individual whisks taken from panel (A) at 5 mm, 35 mm, 65 mm and 95 mm head displacement (whisk number 20, 140, 260 and 380).

II. METHODS

A. Robot experiments

In rats, the long facial whiskers (macrovibrissae) form a grid on each side of the snout, with each whisker mounted in a specialized hair follicle that sits within an area of dense muscular tissue. From an engineering perspective, the facial whiskers resemble tapered elastic beams that deform easily upon contact with objects. Mechanoreceptors at the base of the whisker shaft respond when it is deflected and are transduced into patterns of neural firing that may encode information about the nature of the contact. When the rat is exploring its environment, the long facial whiskers are swept back and forth many times per second in a motion that is actively modulated by contacts with impeding objects.

The G1 (Generation 1) BIOTACT sensor was designed to mimic these aspects of the rat whisker system [8]. The robot consisted of a truncated conical head holding up to 24 whisker modules arranged in 6 radially symmetric rows of 4, oriented normally to the cone surface. The head is mounted as the end-effector of a 7-degree-of-freedom (dof) robot arm (Fig. 1A) that allow it to be moved to positions of interest. The whiskers were swept back-and-forth to repeatedly contact surfaces akin to animal whisking behavior [10]. Whiskers towards the front of the head were shorter than at the back (lengths 50 mm, 80 mm, 115 mm and 160 mm), and were designed with a taper from a 1.5 mm base to a 0.25 mm tip. The deflections upon contacting objects were measured with a hall effect sensor at the base of the whisker [11], with processing/controlling software executed under the BRAHMS Modular Execution Framework [12].

For the present experiments, the head was fitted with a total of 4 whiskers in one row appropriate for sensing axially symmetric shapes such as cylinders aligned perpendicular to the whisker row (Fig. 1B). Six rigid plastic hemi-cylinders with diameters of curvature 30 mm, 40 mm, 50 mm, 60 mm, 70 mm and 80 mm were used as test objects. They were mounted with their curved surfaces facing upwards on bases of appropriate height to ensure equal vertical dimensions. The orientation of the BIOTACT sensor head was such that the four whiskers contacted the cylinders along the horizontal and perpendicular to the cylinder axis (Fig. 1B). When contacting a horizontal surface, the whisker tips were approximately 30 mm apart, spanning a distance from 100 mm to 190 mm perpendicular to the head axis. The depth of the contacts was arranged to be equal for all trials and only contacting the curved sections of the test objects.

The data sets were collected while moving the robot head horizontally at constant speed of 0.5 mm/sec over each test object. In each case a range of 100 mm was considered, with constant whisk frequency of 2 whisks per second. This gave 404 whisks per object at increments of 0.25 mm. Example whisker deflection data for the four whiskers contacting the cylinder of 50 mm diameter is shown in Fig. 2. One training and one test set was collected for each of the six cylinders.

B. Analysis

Bayes' rule states that the posterior probability $P(\mathrm{C}_l|x_n)$ for a new measurement x_n being from one of the N classes C_l is equal to likelihood $P(x_n|\mathrm{C}_l)$ times the prior normalized by the marginal probability. In sequential analysis, the prior is the posterior probability for the preceding measurement, giving a sequential update rule

$$P(\mathrm{C}_l|x_n) = \frac{P(x_n|\mathrm{C}_l)P(\mathrm{C}_l|x_{n-1})}{P(x_n|x_{n-1})}. \tag{1}$$

The likelihoods $P(x|\mathrm{C}_l)$ are assumed identically distributed and independent of when the measurements are taken. The marginal probabilities are explicitly denoted as conditioned on the preceding measurement and are found from summing over all N classes

$$P(x_n|x_{n-1}) = \sum_{l=1}^{N} P(x_n|\mathrm{C}_l)P(\mathrm{C}_l|x_{n-1}) \tag{2}$$

to satisfy the constraint $\sum_{l=1}^{N} P(\mathrm{C}_l|x_n) = 1$. Taking a sequence of measurements $x_1, \cdots, x_n$ thus results in sequence of posteriors $P(\mathrm{C}_l|x_1), \cdots, P(\mathrm{C}_l|x_n)$ for each class, which are calculated by iterating over the relations (1,2) starting from uniform priors $P(C_l) = P(\mathrm{C}_l|x_0) = 1/N$.

Here we compare and contrast decision making based on two paradigms typical of psychological experiments on decision making [6]. The first involves fixed duration stimuli after which participants are expected to answer, thus constraining their reaction times. This case is referred to as an interrogation paradigm. The second paradigm is termed free-response, with participants allowed to respond in their own time. In the latter, both error rates and reaction times vary as participants implicitly choose a speed-accuracy tradeoff.

These paradigms for decision making are implemented here by supplementing the posterior update rule (1,2) with a stopping condition for when to report the decision:
(I) *Interrogation paradigm*: A fixed number of measurements n are specified in advance (here the number of whisks).
(II) *Free-response paradigm*: The number of measurements n are determined in reaction to the posterior values. Here we use the condition that at least one posterior crosses a preset threshold $P(\mathrm{C}_l|x_n) > p$ for any class C_l.
Note that a decision can be aborted when using a free-response paradigm if, for example, a preset 'boredom' deadline is reached before crossing the probability threshold. Following the psychological terminology, we refer to situations where decisions must always be made as *forced* choices, and those where null decisions are allowed as *unforced* choices.

It is also necessary to specify an appropriate decision rule for which choice to make. For simplicity, we use the maximal *a posteriori* (MAP) estimate common in Bayesian analysis

$$\mathrm{C} = \arg\max_{\mathrm{C}_l} P(\mathrm{C}_l|x_n), \tag{3}$$

where $\arg\max$ refers to the argument (the class) that maximizes the posterior.

In this study, the above analysis is applied over multiple whisker contacts of a test object. Therefore we consider the likelihoods $P(x_n|\mathrm{C}_l)$ as single probability values derived over the nth whisk of the four whiskers. There are many ways of estimating this probability, depending upon how the contact information is compared with training data. Here we use a method based on a probabilistic classifier demonstrated previously to accurately discriminate whisker contacts under various circumstances [8], [9], [13], [14].

The whisker deflections have a time series $y(1), \cdots, y(T)$ over a whisk, where T is the number of samples per whisk. The histogram of deflection values then defines a probability distribution for each whisker k

$$P_k(y) = \frac{n_y}{\sum_y n_y}, \tag{4}$$

where n_y is the total number of times that the value y occurs for whisker k. (Implicit in this description is a binning of the measurements y, here taken at 1 mV intervals.) This transformation from time series to probability distribution is used to estimate the likelihoods in the Bayesian update

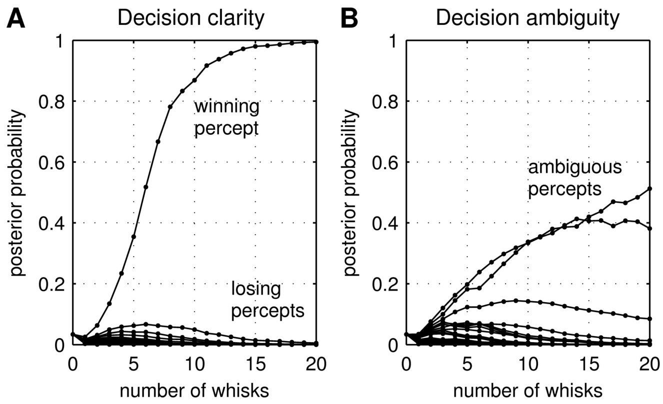

Fig. 3. Examples of the accumulating/depreciating posterior probabilities for the (120) distinct percept classes as whisker deflection data from more whisks is included. The (winning) percept after 20 whisks is shown in green. Panel (A) shows an example with a clear winner and panel (B) of ambiguity.

rule (1), over two phases:
(I) *Training phase*: The conditional probability distributions $P_k(y|C_l)$ for class C_l are estimated by taking the mean over the probability distributions (4) from the whisks in a training data set for that class. This training is repeated for all classes.
(II) *Testing phase*: Each new test whisk (denoted x_n in the Bayesian update rule) is used to estimate the log likelihood as a mean over the log probability distributions

$$\log P(x_n|C_l) = \frac{1}{4T} \sum_{k=1}^{4} \sum_{i=1}^{T} \log P_k(y(i)|C_l). \tag{5}$$

Taking the sum over logs is equivalent to multiplying the probabilities, by treating the samples $y(i)$ as independently and identically distributed. The $1/4T$ factor is from using the average log likelihood as an estimator and assures that the multiplication over sample probabilities $P_k(y(i)|C_l)$ does not produce a vanishingly small likelihood [9].

Note that with this likelihood definition (5), the Bayesian update rule (1) with interrogation paradigm is formally equivalent to the naive Bayes classification considered previously for robot whiskers [8], [13], [14]. The marginal term in Bayes rule is then no longer relevant for decision making over a fixed number of whisks because it is a common normalization for all posteriors, and for flat priors the method becomes equivalent to maximum likelihood estimation. This contrasts with the free-response paradigm, where the division over the marginal term is crucial for comparing the probability threshold with (appropriately normalized) posteriors calculated from Bayes' rule.

III. RESULTS

A. Initial observations

As described in the methods, whisker deflection data was collected for six distinct hemi-cylinders (diameters 30 mm, 40 mm, 50 mm, 60 mm, 70 mm and 80 mm) by passing the robot head slowly over each object whisking at a constant rate (configuration shown in Fig. 1B). Over a 100 mm head displacement, this resulted in 404 distinct whisker deflection profiles. For the example data shown in Fig. 2A (50 mm radius hemi-cylinder), at the start of the movement whiskers 1 and 2 are in contact (Fig. 2B), then whisker 3 contacts and (Fig. 2C), then whisker 4 contacts and whiskers 1 then 2 detach (Fig. 2D), and finally whisker 3 detaches (Fig. 2E). Notice that the pattern of whisker deflections depends on both the curvature of the surface being contacted and the position of the head over the object, permitting the simultaneous classification of object shape and location.

B. Training and test sets

The training and test data were each separated into 120 distinct percept classes, composed of the 6 hemi-cylinders by 20 groups of similarly-positioned whisks (here over 5 mm intervals, with 20 whisks each). We experimented with other groupings, but found this choice a good tradeoff for giving a large numbers of percepts with plenty of whisks in each class. The tick-marks in Fig. 2A show the class delineations.

From the training data, the probability distributions over the sensor values $P_k(y|C_l)$ for each percept class C_l and whisker k were estimated by averaging over all 20 whisks in each class. This resulted in a total of 480 (4 whiskers and 120 classes) distinct probability distributions.

The test data was then used to construct sequences of whisks to validate the Bayesian classifiers. A Monte Carlo procedure was employed, in which percepts with known shape and position were selected randomly from the test data with a a data sequence drawn randomly from the 20 whisks in that test percept class. Two methods of random sampling were used: with replacement and without replacement, differing by whether whisks could be repeated in the test sequence. In total, 2000 randomly generated whisk sequences were constructed for each data point of the following results and about 80000 sequences for each plot.

C. Interrogative perception

Interrogative perception is when a decision is forced after a preset number of whisks (see Methods: Analysis). The perceptual task is to identify which hemi-cylinder is being sensed at what position using whisker deflection data over a sequence of test whisks. The dependence on whisk number was probed up to 20 whisks, corresponding to the number of whisks in each percept class of the test data.

Examples of probabilities derived from whisk sequences from two different test classes are shown in Fig. 3, which plots the posteriors for each putative percept against the whisk number. The probabilities begin at equality corresponding to uniform priors and then evolve smoothly with some rising gradually and others falling. In the first example (Fig. 3A), the perceptual decision given by the largest posterior remains the same after interrogating with 2 whisks or more, while the second example (Fig. 3B) flips between the two leading choices.

Decision errors after interrogating with 1 to 20 whisks are shown in Fig. 4, giving accuracies for both shape (hemi-cylinder diameter) and position. These errors decrease strongly as more whisks are used, as expected by the greater

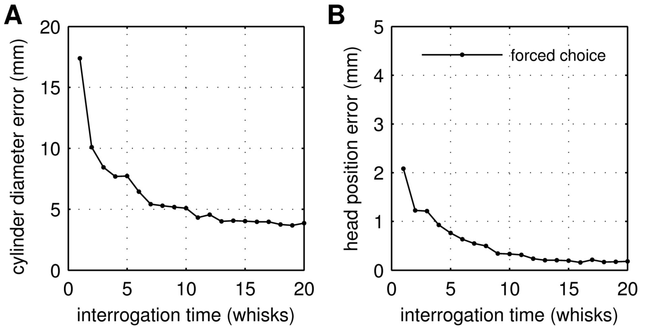

Fig. 4. Decision errors for interrogative perception. Mean classification errors over all percept test classes are plotted against interrogation time counted in whisks. Errors for shape classification are shown in panel (A) and for head position in panel (B).

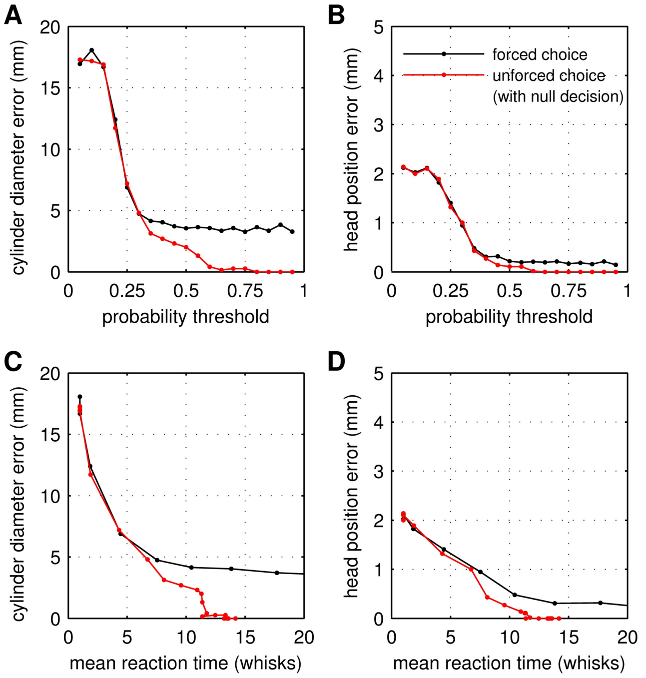

Fig. 5. Decision errors for free-response perception. Mean classification errors over all percept test classes are plotted against probability threshold to make a decision in panels (A,B) and mean reaction time in panels (C,D). A distinction is made between forced choices (black lines) and unforced choices (red lines) depending upon whether a null decision is permitted.

information received, reaching a steady minimum after about 15 whisks. The minimum error of 4 mm for shape (Fig. 4A) is about half the width of a cylinder class and considerably less than the 18 mm error expected at chance. The minimum error of 0.2 mm for position (Fig. 4B), compared with a 5 mm class width and 33 mm error at chance, indicates that position can be determined with high accuracy irrespective of whether errors occur for shape.

D. Free-response perception with forced choices

Free-response perception is when the decision is reported after a number of whisks that depend on the sensory data. Here we used a responsive condition that the decision was made when at least one of the posterior probabilities reached a preset threshold, corresponding to Bayesian sequential analysis [7]. A forced-choice situation was considered first, in which the decision maker kept sampling sensory data until a decision was made. Accordingly, a Monte Carlo sampling method with replacement was used to generate test whisk sequences of arbitrary length (Results: Training and test sets).

Decision errors for perceived shape and position were examined over a range of probability thresholds from 0.05 to 0.95 (Figs 5A,B; black curves). These errors decreased with higher probability thresholds, reaching a steady minimum of 4 mm for shape and 2 mm for position above threshold of about 0.5. The values values of these minimum errors were similar to those for interrogative perception (*cf.* Fig. 4).

In accordance with the perception being freely responsive, the number of whisks to reach a decision had a reaction time distribution (Fig. 6A). Increasing the probability threshold increased the mean reaction time (Fig. 6B; black curve) and decreased the decision errors. Treating the probability threshold as an implicit parameter gave a direct plot of decision error against mean reaction time (Figs 5C,D; black curves). The performance against mean whisk number was similar to interrogative perception (Fig. 4). However, free-response perception with forced choices has out-performed interrogative perception [9] for other data without replacement sampling, so the present results might improve if more whisks were available to sample. Moreover, free-response perception has other benefits over interrogative perception in allowing the speed-accuracy tradeoff to be set directly.

E. Free-response perception with unforced choices

Another type of free-response perception is when the decisions are unforced, in that a decision is not always made. We modeled this situation by including a null decision if the probability threshold was not reached before a deadline. In accordance to each percept class of the test data containing twenty whisks, we used a Monte Carlo sampling method without replacement to give this number of whisks as a natural deadline (Results: Training and test sets).

Unforced free-response perception behraved qualitatively like forced perception, with the mean reaction time increasing (Fig. 6B; red curve) and the shape and position error decreasing (Figs 5A,B; red curve) as the probability threshold was raised. This speed-accuracy tradeoff was determined directly by treating the probability threshold as an implicit parameter in a plot of decision error against mean reaction time (Figs 5C,D; red curves). In addition, unforced perception also has a null decision rate, which increases with the probability threshold because of the requirement of larger posterior probabilities (Fig. 7).

Evidently, the principal improvement of unforced over forced free-response perception is that the decisions both become faster (Fig. 6B) and more accurate (Fig. 5). The improvement in reaction speed is due to an effective deadline of 20 whisks being imposed on the tasks, while the improvement in accuracy is from permitting null decisions if a decision cannot be made by the deadline. It seems that null decisions occur on data that is hard to discriminate, so error rates improve when the probability threshold can be reached.

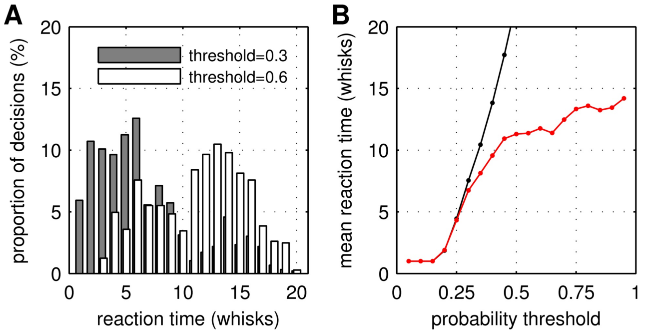

Fig. 6. Reaction times for free-response perception. Panel (A) shows example reaction time histograms with probability threshold $p = 0.3$ (grey), $p = 0.6$ (white) for unforced choices. Panel (B) plots the mean reaction time against probability threshold for forced (black curve) and unforced choices (red curve), corresponding to the histogram means from panel (A).

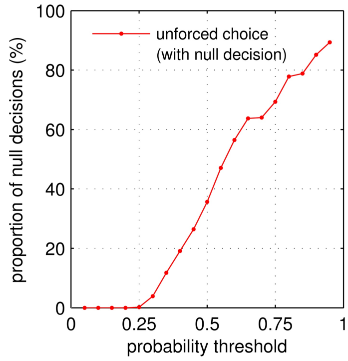

Fig. 7. Null decision rates against probability threshold. Results are shown for unforced free-response perception (red plot).

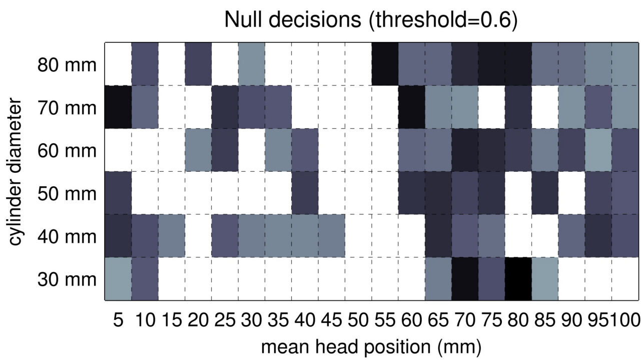

Fig. 8. Test percepts resulting in null decisions. Black regions show percepts with high null decision rates and white regions show low rates. Results are plotted for a probability threshold of 0.6.

These improvements impart near-perfect decision making for probability thresholds greater than 0.6 (corresponding to mean reaction times of 12 whisks or greater and null decision rates greater than 60%).

Examining where the robot head was placed for the test percepts leading to null decisions (Fig. 8) shows that these occur mainly at the extreme head placements. Most of these positions correspond to when only one or two outside whiskers were contacting the test object (Fig. 2). Such placements would be expected to be relatively information poor compared to the central positions where many whiskers are in contact and most test percepts result in clear decisions.

IV. CONCLUSIONS

We have implemented a biologically-inspired scheme of Bayesian perception in a biomimetic robot based on the rodent whisker system. The physical design of the robot is based on the biology of how rats actively palpate (whisk) their long facial whiskers against objects of interest [8]. The sensory processing system in the robot is grounded in proposals for perception from neuroscience using Bayesian decision theory and sequential analysis [1], [2], [5], [6]. Using this system we compare and contrast two approaches for perception from experimental psychology termed the interrogation and free-response paradigms [6]. In free-response perception, the agent reacts when it has received enough sensory information to make a decision; meanwhile, in interrogative perception, the decision is instead reported after a preset number of whisks.

The free-response perception implemented in this robot study used Bayes' rule to calculate the posteriors after each new whisk from priors given by the posteriors from the previous whisk, and relied on normalizing these posteriors with the marginals so they may be properly compared with preset probability thresholds [9]. Meanwhile, interrogative perception with uniform priors can be considered maximum likelihood estimation because it is based on comparing only the likelihoods across the entire data, and can be framed without using Bayes rule. (Note that this interrogative method is equivalent to the Naive Bayes method considered previously for robot whisker systems [8], [13], [14].) We also considered two types of free-response perception, depending upon whether the choices are *forced* or *unforced*. Forced choices necessitate that a decision be made, whereas unforced choices permit a null decision, which was modeled here by a deadline after which a null decision is reported.

We found that Bayesian perception provided a simple and effective method of simultaneously classifying both object shape and position. Results for both interrogative perception and forced free-response perception achieved mean accuracies near 0.2 mm for object localization over a 100 mm range and 4 mm for cylinder diameter over a 30-80 mm range, for sufficiently large decision thresholds or whisk counts. In previous work on texture perception with whiskers, we have found forced free-response perception to achieve better accuracies than interrogative perception [9]. However, the present study found little difference in maximal accuracy between these perceptual methods. The validation methods used in the previous study are distinct from the present ones though (see Sec. IIID), so another robot experiment with, say, more whisk data may reveal differences. It should also be appreciated that free-response perception has other benefits over interrogative perception, such as allowing the speed-accuracy tradeoff to be set directly and having the decision speed depend on the ease of discriminating sensory data.

Another main result was that unforced free-response perception resulted in considerable gains in both reaction speed and accuracy compared with forced decisions, at the expense of a null decision being reported on ambiguous data. Near-perfect decision making was obtained for sufficiently high decision thresholds and null decision rates greater than one-half. In free-response perception, ambiguous sensory data results in slower decisions, so avoiding decisions when there is ambiguity leaves the more easily discriminable data, resulting in lower error rates. We comment that humans and animals commonly avoid making decisions when sensory data is hard to discriminate, often because a more urgent situation arises that commands their attention.

Just as for biological organisms, a robot could find some situations urgent and need to make quick but possibly inaccurate perceptual decisions, other situations less time-critical but requiring greater accuracy, and some situations too ambiguous to make a decision in reasonable time. The potential to tune the desired tradeoff between reaction times, error rates and null decisions is an advantage of Bayesian perception that could be as crucial for robots as the resulting behaviors are for humans and animals interacting with their complex and ever-changing environments.

REFERENCES

[1] D.C Knill and A. Pouget. The bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12):712719, 2004.

[2] D. Kersten, P. Mamassian, and A. Yuille. Object perception as Bayesian inference. *Annu. Rev. Psychol.*, 55:271–304, 2004.

[3] A.C. Huk and M.N. Shadlen. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience*, 25(45):10420, 2005.

[4] M.L. Platt and P.W. Glimcher. Neural correlates of decision variables in parietal cortex. *Nature*, 400(6741):233–238, 1999.

[5] J.I. Gold and M.N. Shadlen. The neural basis of decision making. *Annu. Rev. Neurosci.*, 30:535–574, 2007.

[6] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J.D. Cohen. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700, 2006.

[7] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19(3):326–339, 1948.

[8] J.C. Sullivan, B. Mitchinson, M.J. Pearson, M. Evans, N.F. Lepora, C.W. Fox, C. Melhuish, and T.J. Prescott. Tactile Discrimination using Active Whisker Sensors. *IEEE Sensors*, 99:1, 2011.

[9] N.F. Lepora, C.W. Fox, M. Evans, M. Diamond, K. Gurney, and T.J. Prescott. Optimal decision-making in mammals: insights from a robot study of rodent texture discrimination. *J. R. Soc. Interface*, 2012.

[10] B. Mitchinson, M. Pearson, C. Melhuish, and T. Prescott. A model of sensorimotor coordination in the rat whisker system. *From Animals to Animats SAB 2006*.

[11] M. Pearson, B. Mitchinson, J. Welsby, T. Pipe, and T. Prescott. Scratchbot: Active tactile sensing in a whiskered mobile robot. *From Animals to Animats 11*, pages 93–103, 2010.

[12] B. Mitchinson, T.S. Chan, J. Chambers, M. Pearson, M. Humphries, C. Fox, K. Gurney, and T.J. Prescott. BRAHMS: Novel middleware for integrated systems computation. *Advanced Engineering Informatics*, 24(1):49–61, 2010.

[13] 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. *Proc. IEEE World Congress on Comp. Int. WCCI 2010*.

[14] N.F. Lepora, C.W. Fox, M. Evans, B. Mitchinson, A. Motiwala, J.C. Sullivan, M. Pearson, J. Welsby, T. Pipe, K. Gurney, and T.J. Prescott. A general classifier of whisker data using stationary naive Bayes: Application to BIOTACT robot. *Lecture notes in artificial intelligence, Proc. TAROS 2011*.