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

Naive Bayes novelty detection for a moving, whiskered robot

Nathan F. Lepora, Martin J. Pearson, Ben Mitchinson, Mat Evans, Charles Fox,
Tony Pipe, Kevin Gurney and Tony J. Prescott

Abstract— **Novelty detection would be a useful ability for any autonomous robot that seeks to categorize a new environment or notice unexpected changes in its present one. A biomimetic robot (SCRATCHbot) inspired by the rat whisker system was here used to examine the performance of a novelty detection algorithm based on a 'naive' implementation of Bayes rule. Naive Bayes algorithms are known to be both efficient and effective, and also have links with proposed neural mechanisms for decision making. To examine novelty detection, the robot first used its whiskers to sense an empty floor, after which it was tested with a textured strip placed in its path. Given only its experience of the familiar situation, the robot was able to distinguish the novel event and localize it in time. Performance increased with the number of whiskers, indicating benefits from integrating over multiple streams of information. Considering the generality of the algorithm, we suggest that such novelty detection could have widespread applicability as a trigger to react to important features in the robot's environment.**

I. Introduction

A major goal of robotics is to build autonomous devices that can explore and recognize their environment. Novelty detection could be a key component for achieving these tasks, by allowing the robot to notice environmental features that it has never experienced before; these features can then be prioritized for further exploration, after which they are recognizable and thus no longer novel. Another aspect of novelty detection is that, like their biological counterparts, successful biomimetic robots should be able to identify and react to a changing environment. Novelty detection offers a general way to notice such changes and trigger an appropriate change in behavior, such as avoidance or approach.

The present approach to novelty detection is based on a previous study of texture classification [1], which used data from a moving robot with biomimetic whiskers that sensed various floor textures [2]. Given training data from each class of texture, a test dataset was then classified into one of the previously experienced categories using a naive Bayes classifier. This classifier used Bayes rule to calculate the (posterior) probabilities of the test data being from each class of training data, 'naively' assuming statistical independence to simplify the computation. Such classifiers relate to several influential proposals for how humans and animals integrate evidence to make decisions [3]–[6] and are renowned for being effective despite their computational simplicity [7], [8]. This effectiveness was confirmed in the above study, with the classifier achieving success rates close to 90% and many of the errors actually due to data artifacts (such as collisions with objects) rather than misclassifications.

Here we adapted the naive Bayes classifier to novelty detection and applied it to data from a biomimetic robot (SCRATCHbot) inspired by the rat whisker system [9]–[11]. The key idea is that whereas classification recognizes which previous event a new experience is most like, novelty detection finds when a new experience is *not* like a previous one. Novelty, or outlier, detection has many practical applications, such as discovering financial fraud or computer hacking, using a wide variety of statistical, neural network and machine learning methods [12]–[15]. The naive Bayes method is appropriate for making efficient decisions over multiple channels of time series data, motivating its present use in a mobile robot with multiple whisker sensors.

To test the applicability of naive Bayes novelty detection for robot task, we used an experimental approach in which a robot was first trained in a familiar environment and then tested in a novel situation (Fig. 1). In training, the

N. Lepora, B. Mitchinson, C. Fox, M. Evans, K. Gurney and T. Prescott are with the Adaptive Behavior Research Group, Department of Psychology, The University of Sheffield, Western Bank, Sheffield S10 2TN, UK. (email: n.lepora@shef.ac.uk, b.mitchinson@shef.ac.uk, c.fox@shef.ac.uk, mat.evans@shef.ac.uk, k.gurney@shef.ac.uk, t.j.prescott@shef.ac.uk).

M. Pearson and T. Pipe are with the Bristol Robotics Laboratory, Du Pont Building, Bristol Business Park, Coldharbour Lane, Frenchay, Bristol BS16 1QD, UK (email: martin.pearson@brl.ac.uk, tony.pipe@brl.ac.uk).

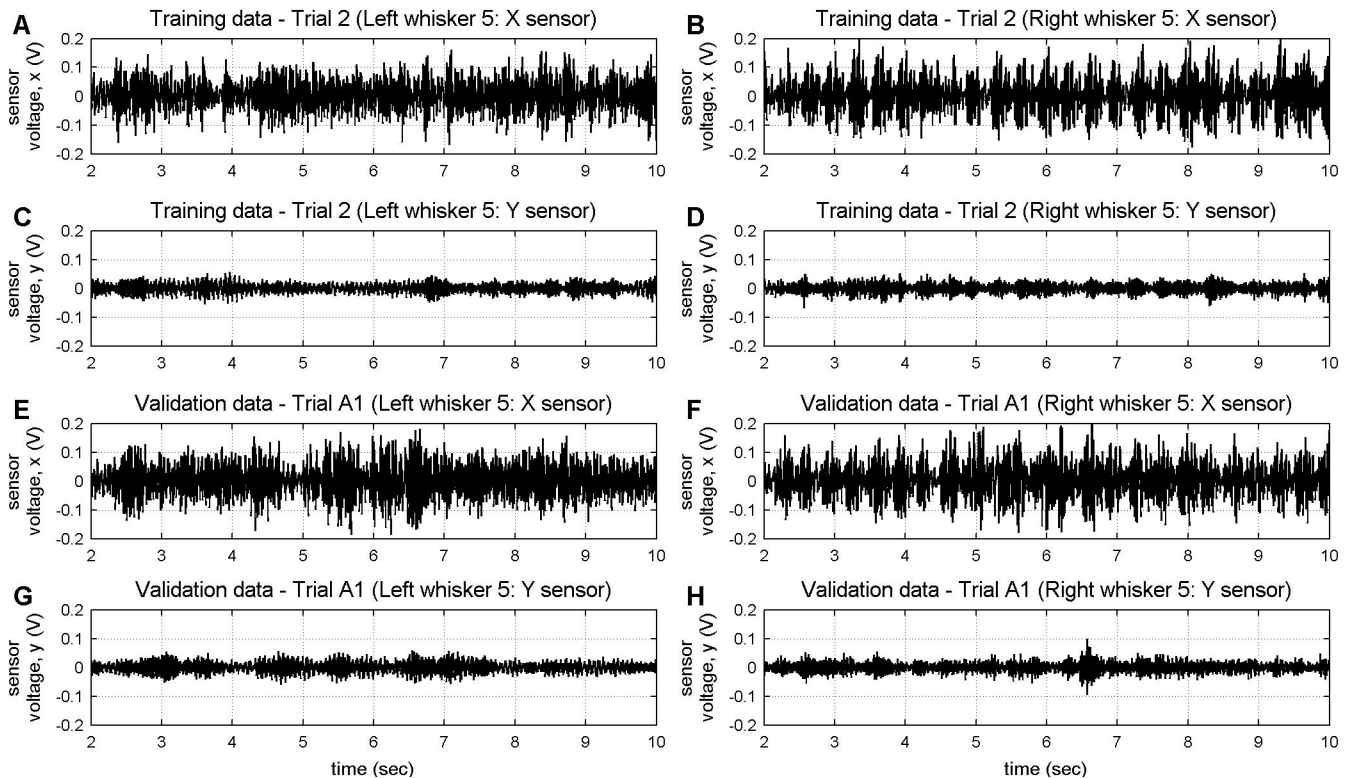Fig. 1. Novelty detection with SCRATCHbot.
Panel A shows a control trial, where the robot traversed a plain vinyl floor while sweeping its whiskers across the surface. Panel B shows a test trial, where a textured strip was placed in the robot's path.

Fig. 2. Example training and test data.
Panels A-D show the X- and Y-sensor measurements for one left and one right whisker in trial 2 of the training data. Panels E-H show the corresponding measurements for the same whisker in trial A1 of the test data. The data was preprocessed to remove the low-frequency artifact due to whisker self-motion.

moving robot performed some trials in which it swept its whiskers across an empty floor. For testing, the robot then had to detect a textured strip placed in its path. In all cases, the detection algorithm distinguished the novel event based only on previous experience of the familiar situation, and accurately localized this event in time.

II. MATERIALS AND METHODS

A. Data collection

The SCRATCHbot platform [9]–[11] consists of three main components: a head, onto which the whisker arrays are mounted; a neck, that allows the head to move independently from the body; and a body that contains the computing resources, locomotion systems and power supply. The processing architecture of the robot is based on the neural pathways identified in the rat whisker system. Neural structures such as the trigeminal sensory complex, superior colliculus and basal ganglia are modeled and developed in software and integrated into a unified system for testing using the BRain And Head Modelling System (BRAHMS) execution framework [16].

The head was designed to carry six independent columns of three whiskers, with each column driven in the anterior-posterior axis by a small motor. These columns are arranged in three-by-three arrays, projecting from the left and right sides of the head. To measure deflections of the whisker shaft caused by environmental contact, a small magnet is bonded

to the base of each whisker, so that a Hall effect sensor can sample the whisker displacements in two directions (denoted here by X and Y). The sensory information from each array is passed to the computing resources on the platform body via serial buses. For the present experiment, the data was sampled and recorded at 2kHz.

Each of the bilateral arrays of whiskers has an associated microcontroller to sample all nine whiskers and to control the rotation of the three columns and tilt angle of the array. This rotation was controlled with the Whisking Pattern Generator, based on a model of the activity underlying the rhythmic whisker motion observed in behaving rats [11], [17]. This self-motion introduces a low-frequency artifact in the whisker signal that will make object detection more difficult. Here we preprocessed the whisker data by simply subtracting the low-pass filtered signal off line (using a gaussian smoother of width 5ms). Note, though, that more general and effective adaptive noise cancelation methods are becoming available for this platform that can achieve the subtraction online [18].

The experimental setup for investigating novelty detection consisted of a length of vinyl flooring over which SCRATCHbot traveled whilst sweeping all eighteen whiskers across this surface. As the robot moved forwards, its head also angled from side-to-side to more fully sample the environment in front of it. In addition to this periodic left-right motion, the whiskers also whisked back and forth according to their pattern generator. This generated a total of

thirty six data streams (18 whiskers each with two directions) about the floor surface over which the robot was traveling.

Three experimental situations were considered, using eight trials of 10 seconds. The first two trials were the control situation, and consisted of just normal vinyl flooring (Fig. 1A). The next three trials (A1-A3) had the robot travel over a thin strip of rough textured surface placed over the vinyl floor, occurring between about four and seven seconds after the beginning of the trial (Fig. 1B). The final three trials (B1-B3) were similar to the previous three, but instead used a smoother strip that was harder to distinguish from the floor.

The goal of the novelty detection task is for the robot to notice the textured strips in trials A1-A3 or B1-B3 based only on its previous experience of the two control trials. In this sense, the two control trials are considered as *training* data on which to tune the novelty detection algorithm, and the other six trials are *test* data for validating its performance.

B. Analysis methods

1) Probability distributions: The methods described here rely on using the probability distributions of the measured time series values, calculated from the empirical frequency with which each value occurs in training data of the control (familiar) situation. These sensor measurements are sampled in time at frequency f_s and their range is binned into N equal-width intervals.

Denoting the training data of the familiar situation by F, the conditional probability of a quantized measurement q being from it is

$$P(x|\mathrm{F}) \equiv P(q(x)|\mathrm{F}) = \frac{n_q}{\sum_{q=1}^{N} n_q}, \qquad (1)$$

where n_q is the total number of times that the value q occurs in the quantized time series. The conditional probability $P(x|\mathrm{F})$ is commonly referred to as a likelihood of the sensor measurement x occurring.

Note that to reduce sampling bias from having too few samples in the distribution tails, the normalized frequencies n_q were convolved with a Gaussian smoother of width 10 quantized intervals before applying Eq. 1, as described in [1].

2) Bayes novelty detection (single measurement): Given a new set of test data, Bayes Theorem states that the (posterior) probability $P(\mathrm{F}|x)$ for a measurement x being drawn from the training data is proportional to the likelihood of that measurement $P(x|\mathrm{F})$ estimated from the training data

$$P(\mathrm{F}|x) = \frac{P(x|\mathrm{F})P(\mathrm{F})}{P(x)}, \qquad (2)$$

where $P(\mathrm{F})$ is the (prior) probability of the data being familiar and $P(x)$ is the (marginal) probability of measuring x given no other information.

Here we consider a novelty detector that finds when the posterior probability for a measurement x being familiar passes below a fixed threshold θ weighted by the marginal probability. By Bayes theorem, a measurement x is novel if

$$P(\mathrm{F}|x) = \frac{P(x|\mathrm{F})P(\mathrm{F})}{P(x)} < \frac{\theta}{P(x)} \quad \Rightarrow \quad \text{Novel.} \qquad (3)$$

For the following arguments, it is convenient to use the logarithm of the posterior probability, which since $\log(x)$ (monotonically) increases with x does not affect the inequality in Eq. 3. Then with a little algebra, the above novelty condition can be rewritten in terms of the log-likelihood

$$\log P(x|\mathrm{F}) < \eta \quad \Rightarrow \quad \text{Novel}, \qquad (4)$$

where the novelty threshold $\eta = \log[\theta/P(\mathrm{F})]$ implicitly contains also the contribution from the prior probability.

How is the novelty threshold determined? A simple criterion followed here is to suppose that it is the minimum value of η for which the novelty detector in Eq. 4 would always infer familiarity when tested on the training data. This threshold lies just above the log-likelihood values for all single measurements x,

$$\eta = \sup_{x \in \mathrm{F}} \log P(x|\mathrm{F}). \qquad (5)$$

where sup denotes the supremum (least upper bound) over the familiar training set F.

3) Naive Bayes novelty detection (many measurements): A potentially more powerful method for detecting novelty is to make the decision over many measurement values, either in time and/or in parallel across different whisker sensors. The above arguments using Bayes theorem are unaffected if the single measurement x is replaced by a time series $\boldsymbol{x}_1, ..., \boldsymbol{x}_n$, of vectors $\boldsymbol{x}_i = (x_i^{(1)}, ..., x_{(i)}^K)$. Then the novelty detection criterion from Eq. 4 becomes

$$\log P(\boldsymbol{x}_{n_s}, \cdots, \boldsymbol{x}_{n_f}|\mathrm{F}) < \eta, \qquad (6)$$

and the novelty threshold from Eq. 5 is now

$$\eta = \sup_{\boldsymbol{x}_i \in \mathrm{F}} \log P(\boldsymbol{x}_{n_s}, \cdots, \boldsymbol{x}_{n_f}|\mathrm{F}), \qquad (7)$$

with $n_s \leq i \leq n_f$ and n_s, n_f the window start and finish.

An important simplification occurs if all these measurements are assumed independent, because the overall conditional probability factorizes into a product of individual conditional probabilities,

$$P(\boldsymbol{x}_{n_s}, \cdots, \boldsymbol{x}_{n_f}|\mathrm{F}) = \prod_{k=1}^{K} \prod_{i=n_s}^{n_f} P(x_i^{(k)}|\mathrm{F}). \qquad (8)$$

Consequently, the classification in Eq. 6 can be rewritten as

$$\sum_{k=1}^{K} \sum_{i=n_s}^{n_f} \log P(x_i^{(k)}|\mathrm{F}) < \eta, \qquad (9)$$

and the novelty threshold is now

$$\eta = N \sum_{k=1}^{K} \sup_{x_i^{(k)} \in \mathrm{F}} \log P(x_i^{(k)}|\mathrm{F}), \qquad (10)$$

with $N = n_f - n_s + 1$ the length of the time-series. The novelty detector finds when the summed log-likelihoods passes a novelty threshold, given by N times the sum of the individual thresholds (Eq. 5) over all dataset dimensions.
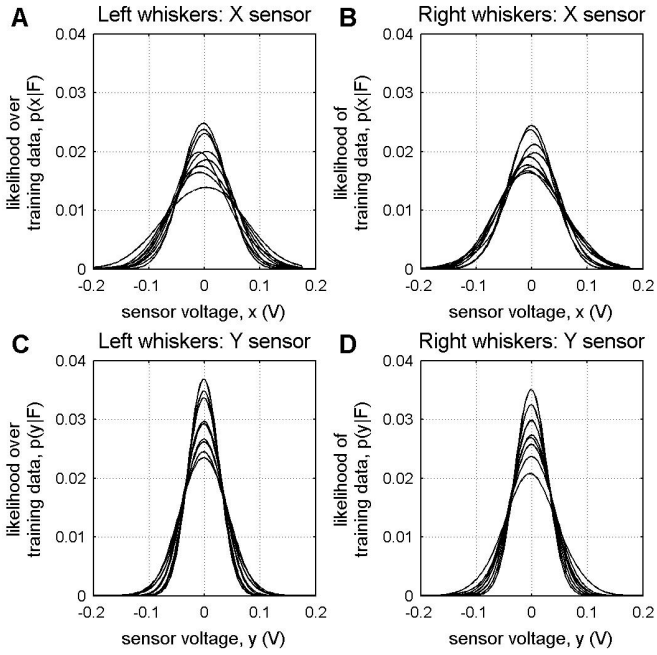
Fig. 3. Training data probability distributions.
These probabilities were calculated from the empirical frequencies with which the measurement values occurred in the training data. In total, there were 36 probability distributions, here grouped by sensor (X or Y) and whisker position (left or right).

III. RESULTS

A. Properties of the (familiar) training data

Two control trials for training were considered as examples of the familiar situation, in which the moving robot swept its whiskers over a plain vinyl floor (top panel of Fig. 1). This data was preprocessed to reduce the low frequency artifact due to the whiskers' periodic motion relative to the floor. Example traces are shown in Figs 2A-D for the X- and Y-sensor measurements of the whisker deflection, showing one left and one right whisker. The peak reading for the X-sensor was about 0.2V and that for the Y-sensor about 0.05V, consistent with the greater displacement of the whiskers in the horizontal plane during whisking. Visual inspection of recordings from the other whiskers and others trial revealed no obvious qualitative differences from those in Fig. 2.

Data from the two training trials were pooled to determine the probability distributions of the measured X- and Y-sensor voltages for each of the 18 whiskers. The range of sensor voltages were binned into widths of 2mV and the number of values in each bin totaled. These totals are the empirical frequencies of the measurements, from which the probability distributions of the sensor values are estimated (Methods, Eq. 1). All 36 probability distributions resembled Gaussian profiles centered on zero deflection (Fig. 3). These probabilities are interpreted as the likelihoods of the X- and Y-sensor measurement for each whisker.

B. Novelty detection over (non-familiar) test data

Given test data that may include a new event, the conditional probability distributions over the familiar training data (Fig. 3) can be used to estimate the occurrence of a novel event. This estimation uses a novelty score calculated with the naive Bayes' rule (Methods, Eq. 2), which represents the log-probability that the test data was drawn from the same distribution as the training data. (More precisely, this score is the log-posterior probability estimated from the log-likelihood for each test data measurement.) The utility of the naive Bayes assumption is that allows the evidence for novelty to be summed over both the time-window and different whisker sensors (Methods, Eq. 9). In the following, the temporal window was fixed at 0.5 seconds to be consistent with related work on texture classification [1].

The novelty score was first calculated over the familiar training data, using the likelihood determined from that same training data. For both training trials, this resulted in a constant log-posterior value overlaid with random fluctuations (Fig. 4A). A novelty threshold η was then determined as the minimum log-posterior value over both training trials, corresponding to the most unlikely window of training data. Note that the novelty scores plotted in Fig. 4 show the log-posterior normalized by this threshold, and thus the novelty threshold
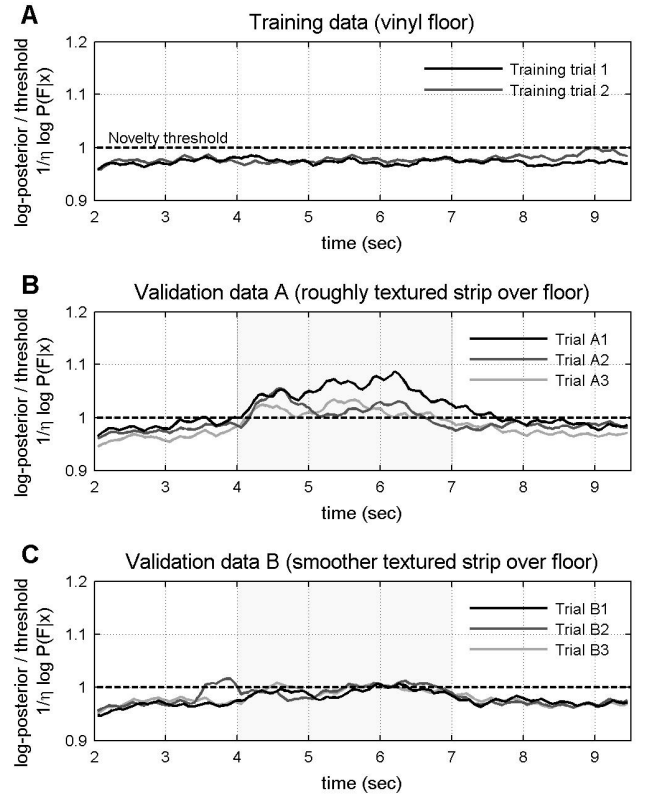

Fig. 4. Novelty score over training and test data.
Panel A shows the novelty score for the training data, while panels B and C are for the two types of test data. This novelty score was the log-posterior probability for a window of data being from the conditional probability distributions in Fig. 3, normalized by the threshold for novelty detection. The approximate times of encountering the novel texture are shown in gray.

Fig. 5.   Times of novelty detection.
Panel A shows the times when the novelty score in Fig. 4 passes the novelty threshold for the six test trials using all 18 whiskers. Panels B and C shows the corresponding times for only 9 whiskers grouped to the left or right. Novelty detections are depicted by the bold lines.

is depicted at unity (dashed line); in addition, because the novelty threshold is negative, the log-posterior/threshold ratio is less than unity over the training data.

This normalized novelty score was then calculated over the six trials of test data, again using the likelihood determined from the training data. As described in methods, this test data consisted of the vinyl floor (like the training data), but with the robot moving over a novel strip in the central portion of the trial (robot pictured in Fig. 1B). Two types of strip were considered: type A, roughly textured, and type B, a smoother texture more similar to the vinyl floor itself. There were three test trials, A1-3 and B1-3, for each strip.

In trials A1-A3, for the roughly textured strip, the normalized novelty score rose above the unit threshold in the central portion of each trial and was below it otherwise (Fig. 4B). Thus, by only having knowledge of the familiar situation, the robot is able to tell that a novel event has occurred and accurately localize it in time. A graphical depiction of the associated novelty detection is shown in Fig. 5A. Novelty detection is clearly a success in this situation.

In trials B1-B3, for the smoother strip, the normalized novelty score again rose above unit threshold in the center of the trial, but was less pronounced than for the textured strip (Fig. 4C). The characterization of novelty now seems more difficult, presumably because the strip is similar to the familiar situation. Examining the graphical depiction of when novelty detection occurs (Fig. 5A), the principal detection times are at the initial and final contacts with the strip. Because the strip is placed over the floor, it seems likely that the novelty score is detecting the sharp contact with its edges. In contrast, the rest of the strip is similar to the familiar situation and is thus not detected as novel.

C. Reliability of novelty detection with whisker number

Thus far the novelty detection has been both trained on and tested against data using all 18 whiskers. Given that each whisker sensor measures both the X and Y deflection, this detection uses a total of 36 distinct information streams.

To examine whether all this information is necessary, we checked whether novelty could be detected reliably using only single whiskers. For the present platform and experiments, this was found to not give a reliable way of detecting novelty, with many false negatives and false positives resulting. This outcome was expected because the novel whisker signals could not be obviously distinguished from the familiar signals by eye (e.g. Fig. 2).

Next, we considered an intermediate test using half of the whiskers, grouped to either the left or right sides of the robot's head. The detection performance for the three whisker groups (all, left and right) is depicted graphically in Fig. 5. The top panel (Fig. 5A) shows the control group using all whiskers. The plotted bold lines show when the normalized novelty score rose above unit threshold for the six test trials from Figs 4B,C. As discussed above, the novel events in trials A1-A3 were easily distinguishable, whereas only the initial and final contacts were detectable in trials B1-B3.

For the left whiskers, the novelty detection (Fig. 5B) was slightly over-sensitive compared with the control group of all whiskers. In consequence, there were false-positives in trials A2 and A3, although these inaccuracies were somewhat compensated by more reliable detections in trials B1-3. For the right whiskers, the novelty detection (Fig. 5C) was far poorer that the control group. In particular, trial A1 was detected as novel across most of its duration and trials B1 or B3 had no novel events. Therefore the left whisker group was slightly over-sensitive to novelty, while the right group was a poor detector of novelty.

This unreliability with fewer whiskers indicates that the naive Bayes novelty detector performs best if it is given more information, as one might expect intuitively. Given there were no discernable difference between the left and right whiskers from the training data, there was also no obvious way to choose the better (left) group before testing. Instead, the best decisions were made by pooling over all data. Then if a misleading signal occurred in one data stream, it could be outweighed by reliable information from the other channels.

IV. Discussion

The performance of a naive Bayes novelty detection algorithm was examined in several experiments using a biomimetic robot (SCRATCHbot) based on the rat whisker system. These experiments consisted of two control trials,

in which the robot swept its whiskers over an empty floor to characterize a familiar situation. These were followed by six test trials in which a strip was placed across the robot's path, with the first three using a roughly textured strip easily distinguishable from the floor and the last three using a smoother strip more similar to the floor. Using information from all 18 whiskers, the detection algorithm was able to identify the novel event in all test trials. In general, the textured strip had a strong novelty signature that was easy to identify, whereas the smoother strip was mainly evident from the initial and final contacts with its edge.

A crucial aspect of the detection algorithm was that it could integrate information over multiple streams of whisker data. Initial results were thus obtained with all 36 channels (18 whiskers with two directional displacements each). To check the reliability of the detection if less data were used, novelty was reassessed using only 9 whiskers on either the left or right sides of the robot's head. Although the novel event was still noticeable in many test trials, the results were degraded with many false-positive and false-negative detections resulting. Meanwhile, testing with single whiskers gave no useful information about novelty. In general, robust decisions resulted from pooling as much data as possible, which allowed a misleading signal in one channel to be outweighed by the other evidence.

A. A common system for novelty detection and classification

The present approach to novelty detection is based on a classification method that was previously used to recognize textures [1]. Because the same mathematical structure underlies these two types of algorithm, both tasks can in principle be achieved with a common computational architecture. Thus a robot could test for recognition (classification) and non-recognition (novelty detection) of its environment using a single, hybrid classification/novelty system.

For example, suppose a robot navigates using whiskers to sense texture. It has previously been trained on some representative surfaces, say smooth vinyl and rough carpet, and the memory of each texture is stored as the log-likelihoods of the sensor measurements. On each set of training data, the largest novelty threshold that characterizes this data as familiar has also been calculated and stored. Now, as the robot moves around its environment, it continually streams data to two parallel modules, one for vinyl and one for carpet, that each calculates a log-posterior from their stored log-likelihoods and thresholds this value to test for novelty. The robot can then make various inferences from these results: to recognize vinyl or carpet, it identifies the largest log-posterior; or to determine novelty, it examine whether both modules conclude novelty. These decisions can trigger different behaviors, such as continued navigation in a familiar environment, or exploration/avoidance when there is novelty.

B. Novelty-triggered learning

How might an autonomous robot use novelty detection to characterize a new environment? One method would be to use novelty detection to trigger an exploring behavior, where novel aspects of the environment are examined repeatedly to provide training data for future classification. This exploration ceases when the environment no longer appears novel, as judged from applying the novelty detector trained by the exploration data to the environment being explored. The robot can then return to its normal behavior until novelty is again encountered. Such a behavioral strategy would, in effect, intrinsically motivate the robot to reduce the perceived novelty of its environment.

References

[1] N. Lepora, M. Evans, C. Fox, M. Diamond, K. Gurney, and T. Prescott, "Naive Bayes texture classification applied to whisker data from a moving robot," *IEEE World Congress on Computational Intelligence* [In Press].

[2] M. Evans, C. Fox, M. Pearson, and T. Prescott, "Spectral Template Based Classification of Robotic Whisker Sensor Signals in a Floor Texture Discrimination Task," *Proceedings TAROS*, pp. 19–24, 2009.

[3] J. Gold and M. Shadlen, "Neural computations that underlie decisions about sensory stimuli," *Trends in Cognitive Sciences*, vol. 5, no. 1, pp. 10–16, 2001.

[4] J. Tenenbaum and T. Griffiths, "Generalization, similarity, and Bayesian inference," *Behavioral and Brain Sciences*, vol. 24, no. 04, pp. 629–640, 2002.

[5] K. Kording, "Decision theory: what 'should' the nervous system do?" *Science*, vol. 318, no. 5850, p. 606, 2007.

[6] R. Bogacz and K. Gurney, "The basal ganglia and cortex implement optimal decision making between alternative actions," *Neural computation*, vol. 19, no. 2, pp. 442–477, 2007.

[7] D. Hand and K. Yu, "Idiot's BayesNot So Stupid After All?" *International Statistical Review*, vol. 69, no. 3, pp. 385–398, 2001.

[8] I. Rish, "An empirical study of the naive Bayes classifier," in *IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence*, 2001, pp. 41–46.

[9] M. Pearson, A. Pipe, C. Melhuish, B. Mitchinson, and T. Prescott, "Whiskerbot: a robotic active touch system modeled on the rat whisker sensory system," *Adaptive Behavior*, vol. 15, no. 3, p. 223, 2007.

[10] T. Prescott, M. Pearson, B. Mitchinson, J. Sullivan, and A. Pipe, "Whisking with robots from rat vibrissae to biomimetic technology for active touch," *IEEE Robotics and Automation Magazine*, vol. 16, no. 3, pp. 42–50, 2009.

[11] M. J. Pearson, B. Mitchinson, J. Welsby, A. G. Pipe, and T. J. Prescott, "Scratchbot: Active tactile sensing in a whiskered mobile robot," *From Animals to Animats SAB2010* [In Press].

[12] M. Markou and S. Singh, "Novelty detection: a review–part 1: statistical approaches," *Signal Processing*, vol. 83, no. 12, pp. 2481–2497, 2003.

[13] ——, "Novelty detection: a review–part 2: neural network based approaches," *Signal Processing*, vol. 83, no. 12, pp. 2499–2521, 2003.

[14] V. Hodge and J. Austin, "A survey of outlier detection methodologies," *Artificial Intelligence Review*, vol. 22, no. 2, pp. 85–126, 2004.

[15] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys (CSUR)*, vol. 41, no. 3, p. 15, 2009.

[16] B. Mitchinson, T. Chan, J. Chambers, M. Pearson, M. Humphries, C. Fox, K. Gurney, and T. Prescott, "BRAHMS: Novel middleware for integrated systems computation," *Advanced Engineering Informatics*, vol. 24, no. 1, pp. 49–61, 2010.

[17] B. Mitchinson, M. Pearson, C. Melhuish, and T. Prescott, "A model of sensorimotor coordination in the rat whisker system," *From Animals to Animats 9*, pp. 77–88, 2006.

[18] S. R. Anderson, M. J. Pearson, A. G. Pipe, T. J. Prescott, P. Dean, and J. Porrill, "Enhanced detection of robot whisker contacts by adaptive cancelation of self-generated signals," *IEEE Transactions on Robotics*. [In Press].