Classification of Odours for Mobile Robots Using an Ensemble of Linear Classifiers

Marco Trincavelli, Silvia Coradeschi and Amy Loutfi

AASS Research Center, Örebro University
Fakultetsgatan 1, 70182 Örebro
E mail: name.surname@aass.oru.se

Abstract. This paper investigates the classification of odours using an electronic nose mounted on a mobile robot. The samples are collected as the robot explores the environment. Under such conditions, the sensor response differs from typical three phase sampling processes. In this paper, we focus particularly on the classification problem and how it is influenced by the movement of the robot. To cope with these influences, an algorithm consisting of an ensemble of classifiers is presented. Experimental results show that this algorithm increases classification performance compared to other traditional classification methods.

Keywords: Odour Classification, Mobile Robotics
PACS: 01.30.Cc

INTRODUCTION

Mobile robots equipped with an array of gas sensors can be a valuable tool in scenarios like inspection of hazardous areas or environmental monitoring [1], particularly in cases where toxic contaminants are involved. Mobile robots can play an important role in assessing the presence of dangerous substances, identifying their character, providing a map of their distribution and if possible quantifying their concentration [2]. One essential feature of such a robot is to be able to discriminate substances in a fast and reliable way. Up to date the classification of odours with an array of gas sensors has been done in systems that perform a controlled sampling process (three-phase sampling) and then process the data offline [3]. For mobile robotics application the bulky and expensive equipment needed for controlling the sampling process is often unsuitable. Therefore it is more convenient for the array of sensors to be exposed directly to the external environment and sample it continuously. Under these conditions, however, the signal from the array does not have the three characteristic phases of the controlled sampling process but shows a series of sensor responses which vary depending on the interaction between the nose and the odour plume. An example signal collected during an experiment is shown in Figure 1. Only few works have dealt with odour classification problem in such conditions [4, 5].

Since the gas diffusion in uncontrolled environments is dominated by turbulence [6], the odour plume has a very articulated shape with many patches and meanders. Therefore the interaction of the nose with the odour plume is complicated and hard to model. Despite the difficulties in providing an exact model of this interaction we can investigate if some controllable factors, like the movement of the robot or the location of the experiment, introduce some regularity in the data set. The main contribution of this work is to discover such regularities and exploit them to enhance the classification performance. The approach proposed is the creation of an ensemble of specialized classifiers. The performance of the proposed method is then compared to two standard classification algorithm, namely the Support Vector Machine (SVM) and the Radial Basis Function (RBF) [7].

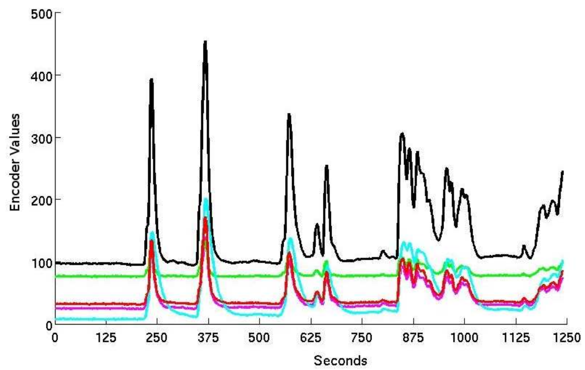

Figure 1. Signal collected during an experimental run.

EXPERIMENTAL AND METHODS

The robot used in the experiments is an ATRV-JR all terrain robot equipped with the Player Robot Device Interface [8]. Player provides both the interface to the sensors and the actuators, and high level algorithms to address robotic tasks such as localization (*amcl* driver) and navigation (*vfh* and *wavefront* drivers). The robot is equipped with an electronic nose, an actively ventilated aluminum tube containing an array of five metal oxide gas sensors, mounted in front of the robot at a height of 0.1 m on the ground. The sensors present in the array are listed in Table 1 together with their target gases.

Table 1. Gas sensors used in the electronic nose.

Model	Gases Detected	Quantity
Figaro TGS 2600	Hydrogen, Carbon Monoxide	2
Figaro TGS 2602	Ammonia, VOC, Hydrogen Sulfide	1
Figaro TGS 2611	Methane	1
Figaro TGS 2620	Organic Solvents	1

The experiments have been performed in three different locations using four different moving strategies. In all the experiments the robot was moving with a speed of 0.05 m/s. The odour source was a cup full of the analyte placed on the ground. The first location that has been considered is a large close room in which the robot followed a sweeping trajectory with two orthogonal orientations that we name N-S and E-W. The second set of experiments has been carried out in a small classroom whose door has been left open. In this environment the robot performed two different kinds of spiral path: a spiral without any stops from the beginning to the end of the experiment and a spiral with stops when an odour is detected, at which point the robots stands static until enough information is obtained to perform a classification. The last experimental location was a courtyard with an uneven grass surface. In this case the robot performed a spiral movement stopping when a gas is detected similar to the one performed in the classroom. Table 2 summarizes the five different experimental configurations. The experiments have been repeated multiple times with three different substances, namely ethanol, acetone and isopropyl, that are the target substances for our classification problem. During one experimental run multiple responses were collected, for a total of 592 responses evenly distributed among the three classes.

The classfication algorithm is articulated into five phases, namely baseline subtraction, signal segmentation, feature extraction, dimensionality reduction and classification. The baseline is the value that a gas sensor gives as output when it is exposed to clean air. This value depends on temperature, humidity and long/short term drift [9]. The baseline is subtracted from the output value of the sensors in order to limit the effect of these factors using differential baseline subtraction. The baseline value is measured for 60 s at the beginning of every experiment, when the odour source is closed and the robot is standing still. After performing this first transformation the signal is smoothed using an average filter of dimension 5 in order to suppress the noise due to sampling and quantization. The smoothed signal is then segmented into three different phases, namely baseline, rise and decay according to the value of the first derivative. The segmentation procedure can be easily explained using a finite state machine as shown in Figure 2.

Table 2. Summary of the experimental conditions in which the data have been collected.

Data Set	Location	Moving Strategy	Number of Runs
1	Large Room	Sweep N-S	15
2	Large Room	Sweep E-W	15
3	Classroom	Spiral	18
4	Classroom	Spiral with Stops	72
5	Garden	Spiral with Stops	16

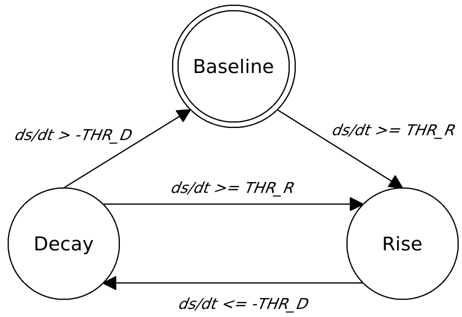

Figure 2. Finite State Machine that illustrates the segmentation algorithm.

In this figure the first derivative is denoted as ds/dt and the threshold for the rise and decay are THR_R and THR_D respectively. Two different thresholds are needed since the rise and decay phase are best described using a first-order model and the time constant for the rising phase is smaller [10]. A complete response to an odour patch is considered to be the ensemble of a consecutive rise and decay phase. The isolated response is then passed to the feature extraction module that fits a 2^{nd} degree polynomial to the points in the response. The choice of the polynomial degree that is more suitable to the sensor response that consists of a consecutive rise and decay. Moreover a parabula provides a sufficient degree of fitting without overfitting the signal. The feature vector is built by concatenating the 3 coefficients obtained by fitting each of the five sensors, obtaining a 15 dimensions vector. The dimensionality of the feature vector is then reduced applying the Linear Discriminant Analysis

(LDA) technique, that projects the vectors into a lower dimensional space in which the distance between samples of different classes is maximized and the distance between samples of the same class is minimized [7]. In contrast to the Principal Component Analysis (PCA) which searches for the directions in which the variance of the data is maximized neglecting the fact that samples belong to different classes, the LDA searches for the directions on which to project the data taking into consideration the membership of the samples to the different classes. The first two LDA Components are kept and passed to the subsequent classification algorithm. The classification algorithm we propose is an ensemble of Linear Probabilistic Discriminative Models [7]. This model provides as output the posterior probabilities that the sample belongs to each of the classes. In particular, we want to investigate the possibility of breaking down the odor classification problem into a set of simpler classification tasks/subproblems. These tasks can be solved efficiently by a set of linear classifiers. The global solution is then obtained recombining the solutions of the simpler tasks by marginalizing with respect to the subproblem according to the following equation:

$$P(C_k|\mathbf{x}) = \sum_{m=1}^{M} P(C_k|\mathbf{x},m)P(m|\mathbf{x}) \qquad (1)$$

Where $P(C_k|\mathbf{x})$ is the probability that sample $\mathbf{x}$ belongs to class C_k, $P(C_k|\mathbf{x},m)$ is the probability that sample $\mathbf{x}$ belongs to class C_k given that sample $\mathbf{x}$ belongs to subproblem m and $P(m|\mathbf{x})$ is the probability that sample $\mathbf{x}$ belong to subproblem m. This method requires the training of $M+1$ classifiers: one for each of the subproblems and one for assigning the membership of the sample to the subproblem.

RESULTS

Figure 3 shows a plot of the two LDA components when the LDA is calculated with respect to the analyte. We can notice how the acetone samples are well clustered, while the ethanol and isopropyl samples are partially overlapping. This makes the classification problem non-trivial.

If instead of considering the LDA projection with respect to the analyte, we consider the LDA projection with respect to the experimental conditions in which the responses has been collected. This result displayed in Figure 4. Here we can observe clearly three data clusters that correspond to the different moving strategy of the robot, namely sweeping, spiraling and spiraling with stops. In Figure 5 the two LDA components with respect to the substance to classify for each single subgroup are displayed.

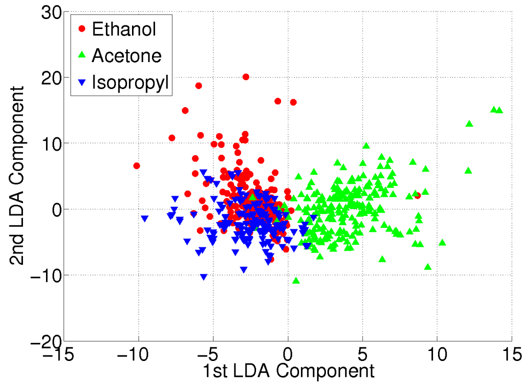

Figure 3. LDA Plot of the full classification problem.

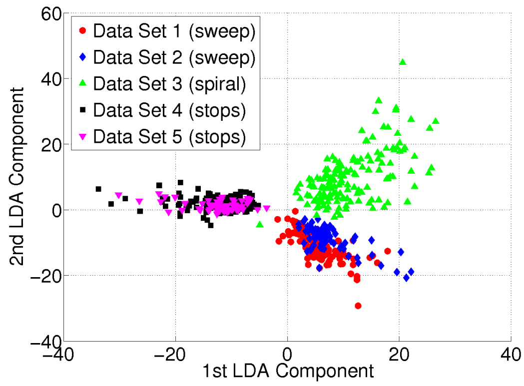

Figure 4. LDA plot of the data set classification problem.

The classification problem is much simpler in each of the groups than in the whole training set. Therefore an ensemble of *four* Linear Probabilistic Discriminative Models is used: one that calculates the probability $P(m|\mathbf{x})$ that a sample has been collected using moving strategy m, where m can take three values that correspond to sweeping, spiraling or spiraling with stops movement. The other three classifiers calculate the probability $P(C_k|m,\mathbf{x})$ that sample $\mathbf{x}$ belongs to class k, with k that can be ethanol, acetone or isopropyl, given that $\mathbf{x}$ has been collected performing movement m. The four decisions are then ensembled according to Equation (1) in order to obtain $P(C_k|\mathbf{x})$. The sample is then assigned to the class with the highest $P(C_k|\mathbf{x})$.

The obtained classifier has been tested with a 20 fold cross validation. As a term of comparison a SVM and a RBF have been trained and evaluated performing also a 20 fold crossvalidation on the whole data set. The performance of the classifiers together with the average training time is reported in Table 3. Table 4 reports the classification rates of the classifiers that form the ensemble. The proposed ensemble of classifiers clearly outperforms both the SVM and the RBF that operate on

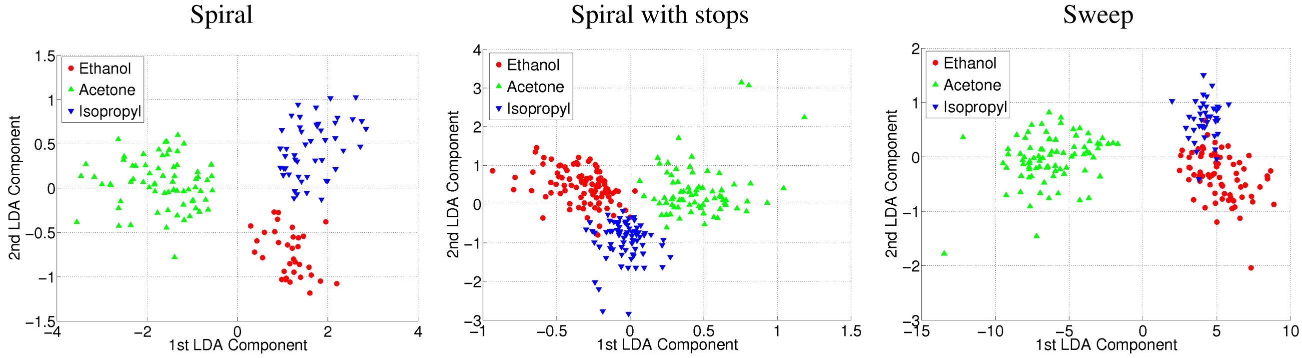

Figure 5. LDA plot of the classification problem when a single subgroup is considered.

the whole data set. Moreover the ensemble is also much quicker to train than the SVM.

Table 3. Average and confidence interval (95% confidence level) of the classification rate for the ensemble of classifiers and for the reference classifiers (SVM and RBF). In the third column the average training time for the three classifiers is reported.

Classifier	Accuracy	Avg. Train Time
Ensemble	89.36 ± 2.58	0.02 *s*
SVM	75.41 ± 4.59	14.52 *s*
RBF	68.64 ± 3.55	0.08 *s*

Table 4. Average and confidence interval (95% confidence level) of the classification rate for the classifiers that are composed into the ensemble.

Classifier	Accuracy
Movements	98.67 ± 0.93
Sweep group	88.47 ± 3.94
Spiral group	93.46 ± 4.93
Spiral with stops group	93.15 ± 2.91

CONCLUSIONS

Odour classification is a fundamental ability for a robot that has to monitor the pollution or explore an area and discover hazardous gases. In order to achieve reliable classification it is important to understand the factors that influence the shape and the properties of the sensors response. In this paper we have shown that the movement strategy of the robot clearly affects the properties of the signal and by taking into account this movement in the classification performance can be enhanced. A next step is to analyze exactly which properties of the signal response are affected by the different movement strategies. Also future work will continue the analysis of the mobile robotic platform and the environment in which the data is collected in order to find other patterns that can further improve the classification performance.

REFERENCES

1. M. Trincavelli, M. Reggente, S. Coradeschi, A. Loutfi, H. Ishida, and A. J. Lilienthal, "Towards Environmental Monitoring with Mobile Robots," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS 2008)*, 2008, pp. 2210–2215.
2. A. J. Lilienthal, A. Loutfi, and T. Duckett, *Sensors* **6**, 1616–1678 (2006).
3. R. Gutierrez-Osuna, A. Gutierrez-Galvez, and N. U. Powar, *Sensors and Actuators B: Chemical* **93**, 57–66 (2003).
4. M. Trincavelli, S. Coradeschi, and A. Loutfi, "Classification of Odours with Mobile Robots Based on Transient Response," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS 2008)*, 2008, pp. 4110–4115.
5. D. Martinez, O. Rochel, and E. Hughes, *Autonomous Robots* **20**, 185–195(11) (June 2006).
6. P. J. W. Roberts, and D. R. Webster, "Turbulent Diffusion," in *Environmental Fluid Mechanics - Theories and Application*, edited by H. Shen, A. Cheng, K.-H. Wang, M. Teng, and C. Liu, ASCE Press, Reston, Virginia, 2002.
7. C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, Springer, 2006, ISBN 0387310738.
8. B. Gerkey, R. T. Vaughan, and A. Howard, "The Player/Stage Project: Tools for Multi-Robot and Distributed Sensor Systems," in *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR)*, 2003, pp. 317–323.
9. Figaro Gas Sensor: Technical Reference, 1992. http://www.figaro.com.
10. A. J. Lilienthal, and T. Duckett, "A Stereo Electronic Nose for a Mobile Inspection Robot," in *Proceedings of the IEEE International Workshop on Robotic Sensing (ROSE)*, Örebro, Sweden, 2003.