The 2009 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 11-15, 2009 St. Louis, USA

Online Classification of Gases for Environmental Exploration

Marco Trincavelli, Silvia Coradeschi and Amy Loutfi

Abstract— **In this paper we investigate how a mobile robot equipped with tin dioxide gas sensors and an anemometer can use an online classification algorithm in order to improve the exploration strategy. The purpose of the platform is to establish the character of a gas source with accuracy while minimizing the time required for exploration. For this to be possible, the output of the classification algorithm is probabilistic, feeding in a sequence of posterior probabilities to a path planner. To further assist path planning, a 3d-ultrasonic anemometer is available which give indication on the average wind speed and direction. In addition to evaluating different olfaction driven path planning strategies, experimental validations also evaluate the classification algorithms and its application to different environments with varying characteristics.**

I. INTRODUCTION

Environmental monitoring and inspection of dangerous areas are important applications for mobile robots. Equipped with the appropriate sensor technology, mobile robots can be dispatched to evaluate and inspect air and soil quality. Particularly, in cases where hazardous contaminants are involved mobile robots can play an important role in assessing the presence of dangerous substances, identifying their character and if possible quantifying their concentration. Within the field of mobile robot olfaction, gas sensor technology (often tin dioxide based) has shown potential to be used for environmental inspection or exploration of areas where toxic contaminants might be present [1]. An important aspect to be considered for such a platform is the ability to detect specific odours when more than one substance might be present. The classification should also cope with the properties of the area of inspection which may include a turbulent airflow and a patch-like distribution of the gas. Moreover the system is most likely deployed in a different location from the one it has been trained and therefore it needs to be robust to variations in the environmental conditions.

To date, many algorithms for gas classification using metal oxide sensors have been developed under the assumption of a controlled sampling process and with offline data processing [2], [3]. In these cases, a standard three phase sampling process is used where first a baseline phase measures the sensors' reaction to a reference gas, a sampling phase brings in an unknown gas and exposes the sensors until a steady state is achieved, and a recovery phase flushes the unknown gas from the array. Although the three phase technique is adequate for a static sensory array, the applicability is limited for cases where the sensor is mounted on a mobile platform and the surrounding air flow is turbulent. Recent work has demonstrated the feasibility of classifying offline different substances using the data collected from gas sensors as the robot is continuously moving in a predefined path [4]. A necessary step to achieve a platform for realistic environmental inspection, that includes both exploration and classification, is to consider an online classification algorithm and to couple the output of such an algorithm to the movement of the robot in order to optimize the exploration strategy. Moreover it is also necessary to evaluate the robustness to the system with respect to variations of the environment in which the system is deployed.

AASS Mobile Robotics Lab Örebro University, Örebro, Sweden
firstname.lastname@aass.oru.se

A. Related Work

Related work in mobile robot olfaction has primarily considered only the amplitude of the gas sensor response in order to control the robot movement, often for odor source localization [5]. In the context of classification however, the sensor array contains gas sensors of partially overlapping selectivities. Consequently, it is the pattern of response which characterizes the gas source and not necessarily the amplitude of each individual sensor. Also, the difficulty for an odour detection robot to navigate and classify substances is enhanced by a number of challenges with respect to the sensing technology and the configuration of the environment. Firstly, both the gas sensor and the anemometer share the common feature of taking point measurements. Therefore, unlike range sensors, only interaction at the surface of the sensor can be measured and in the case of tin dioxide sensors, surface area is approximately 1 cm^2 [6]. Thus the mobile robot serves an important role to bring the sensor system to the region or area of interest in order to obtain measurement. Secondly, there is a delay in the response of the gas sensor and consequently the measurement value does not necessarily correspond to the concentration occurring at a physical location. Thirdly, in non-controlled environments, gas diffusion is generally dominated by turbulence. Already, in indoor environments, the resulting distribution of gas is chaotic and characterized with intermittent patches [7]. This distribution is further complicated in outdoor environments exposed to wind gusts and deflections from nearby buildings.

B. Approach

The key to successful exploration given the above challenges is to be able to balance efficient movement of the robot and the collection of significant data. This is a multi-step process requiring methods for 1) online detection of the presence of a gas, which may indicate the presence of a plume 2) identification of gas character which includes an indication of the quality of the classification 3) moving of

the robot to the next best location for sampling 4) stopping the exploration once classification is satisfactory.

In this paper we investigate the ingredients required for each of these steps. The main contribution is an extension of the classification algorithm to work in an online setting. This allows the path planning algorithm to take decision based on the output of the classification module in order to try to increase the classification performance and to shorten the exploration time in case a reliable classification has been obtained in an early phase of the exploration. Another contribution is a preliminary evaluation of the robustness of the system to varying environmental conditions. This evaluation has been performed by collecting training samples indoor and then deploying the robot outdoor.

We begin this paper with a description of the robot and sensor system in Section II. In Section II-A the classification algorithm is described. In Section II-B, the path planning strategies are explained. Finally, in Section III and IV the experimental results are presented and the system is validated in both outdoor and indoor environments.

Fig. 1. The mobile robot used to build the exploration system.

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

TABLE I

GAS SENSORS USED IN THE ELECTRONIC NOSE.

II. THE EXPLORATION SYSTEM

The robot that is used in the experiments is an iRobot ATRV-JR that is particularly suited for rough terrains, see Figure 1. The software used to interface with the sensors and the actuators is the Player Robot Device Interface [8]. Player provides an easy way to communicate with the underlaying hardware, and provides high level algorithms to address robotic tasks such as localization and navigation. In particular, in our experiments the *amcl* driver is used for localization and the *wavefront* and *vfh* drivers are used respectively for path planning and obstacle avoidance. The robot is equipped with a SICK LMS200 laser range scanner that is used by the *amcl* and *vfh* drivers.

The software architecture we use in our exploration system is shown in Figure 2. In addition to the laser scanner, two other sensor modalities are included: an electronic nose and an anemometer. The electronic nose is an array of five gas sensors. Table I shows the models of the sensors and their target gases. The e-nose is mounted on the robot at a height of 0.01 m. It is close to the ground because the analytes we consider, namely ethanol, acetone and isopropyl, are heavier than air and therefore propagate at ground level. The anemometer is a Young 81000 Ultrasonic Anemometer that gives as output a 3D vector that expresses direction and intensity of the airflow. This anemometer has a range from 0.02 m/s up to 40 m/s with a resolution of 0.01 m/s. The classification and path planning algorithms that are at the core of this work are described in the next two subsections.

A. Classification Algorithm

The classification algorithm is articulated into five phases, namely baseline subtraction, signal segmentation, feature extraction, data normalization 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 short term drift. The baseline is subtracted from the output value of the sensors in order to limit the effect of these factors. After performing this first transformation the signal is smoothed using an average filter 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 3. 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 [9]. A response to a patch is considered to be a rise in the sensors' signal followed by a decay. An example of the sensors' output collected during a part of an experiment is shown in Figure 4 and illustrates the different phases of the signal response.

The isolated response is then passed to the feature extraction module that calculates the Discrete Wavelet Transform (DWT) of the signal. The DWT is a multilevel decomposition technique that gives a description of the signal in the time-scale domain and is suited for the analysis of highly dynamic signals since it is able to capture abrupt changes in the signal. Then the feature vector is normalized in order to reduce the sample to sample variation using the Dimension Auto-Scaling (DAS) approach. This normalization method transforms every feature in order to give it zero mean and unitary standard deviation over the all training set. Finally, the normalized sample is classified using a Relevance Vector Machine (RVM) [10]. The advantage of using the RVM rather than the more popular Support Vector Machine (SVM) is the possibility to obtain an estimation of the posterior probability of a sample belonging to a certain class. This feature gives us the possibility of introducing a rejection class

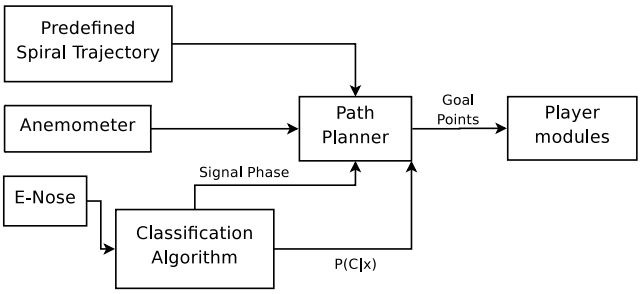

Fig. 2. Software architecture of the exploration system.

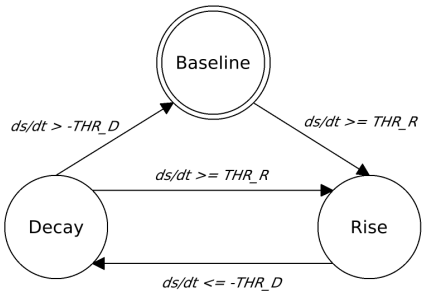

Fig. 3. Finite State Machine that illustrates the segmentation algorithm.

for all the samples for which a definitive membership can not be determined. The output label from the classifier can be expressed as:

$$L_x = \begin{cases} \underset{k}{argmax} P(C_k|x) & \text{if} \quad P(C_k|x) \geq \Gamma \\ rejected & \text{if} \quad P(C_k|x) < \Gamma \end{cases} \quad (1)$$

where L_x is the output label of the classifier for input vector x and $P(C_k|x)$ is the posterior probability of class k for input vector x. Γ is the rejection threshold. A more detailed description of the classification algorithm can be found in [4].

To use the classification algorithm online we output both the signal phase obtained from the segmentation and the estimation of the posterior probabilities. It is important to notice that the only part of the algorithm that executes for every sample is the segmentation that is computationally inexpensive. All the other parts are executed only when a complete response to a patch is detected and therefore the execution time of the complete classification algorithm is not a crucial issue.

B. Path Planning Algorithm

Three different path planning strategies have been considered in this work. The simplest one consists in following a predefined spiral trajectory and going back to the initial position when the whole trajectory has been performed. In this strategy all the goal points are generated at the beginning of the experiment and the path is therefore independent from the sensor readings collected during the exploration. The spiral has been chosen as the basic movement since it can be seen as a systematic strategy that starts by exploring the boundaries of the area of interest and continues inwards, gradually covering all the region to be inspected. The second strategy, that is displayed in the block diagram in Figure 5,

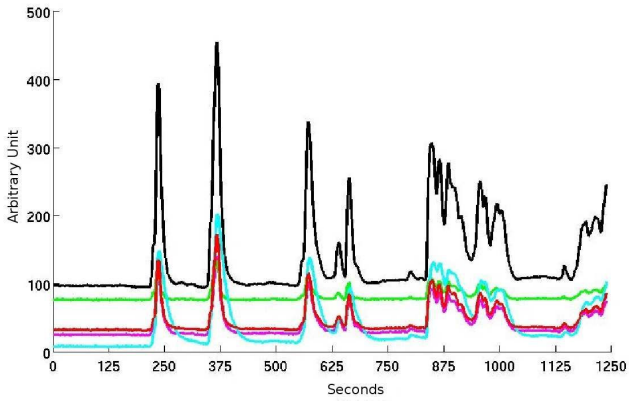

Fig. 4. Example signal collected during a part of an experimental run.

takes into account the output of the classification algorithm in order to determine the next action. In particular the robot is able to stop when a rising signal is detected, wait for a decay in the sensor response and perform classification. Once the classification is done, if the posterior probability of the selected class is greater than a preselected threshold, that in our case is set to $THR = 0.95$, the robot goes back to the initial position. If the threshold is not surpassed the robot continues to follow the spiral trajectory.

The last path planning strategy considers also the airflow measured by the anemometer in addition to the output of the classification module. As illustrated in Figure 6, when the first classification has been performed and the posterior probability threshold is not met a casting behaviour similar to the one described in [11] is performed. Casting is a biologically inspired behaviour which consists of either a zig-zag or spiralling motion each time increasing the explored area. To perform casting, the robot needs to stop in order to collect wind measurements so that the movement of the robot does not interfere with the anemometer readings. The robot stops for 10 s to measure the wind direction. An estimation of the variance of the direction of the airflow is made, and in cases of high wind variance the spiral trajectory is resumed since the assumption is made that a directional plume is not present. If, on the other hand, the variance of the wind direction is low, the robot plans a path perpendicular to the average wind direction in an attempt to reenter the plume. To estimate the variance of the wind direction we use a measure where the sum of the modula of each wind vector is compared to the modulus of the resulting sum vector in the following:

$$\omega = \frac{\|V_1 + \ldots + V_N\|}{\|V_1\| + \ldots + \|V_N\|} \quad (2)$$

where $V_1 \ldots V_N$ are the wind readings from the anemometer. ω varies from 0 (when the wind vectors cancel each other) to 1 (when the wind vectors are all parallel). The threshold on ω for deciding whether to perform casting, is at *THR_W = 0.5*.

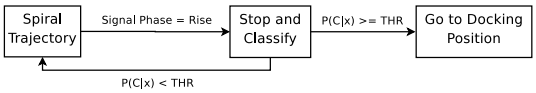

Fig. 5. Path planning strategy in which the robot stops when the classification algorithm detects a rising signal and navigates to a docking station when a classification with a high confidence is obtained. Initially the robot follows a spiral path.

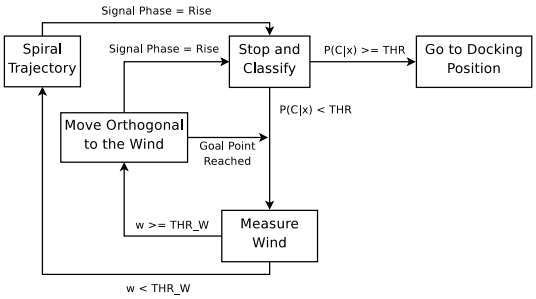

Fig. 6. Path planning strategy in which the robot performs casting moves in order to try to reenter a plume and perform another classification attempt. Initially the robot follows a spiral path.

III. EXPERIMENTAL ANALYSIS

The experiments presented in this work evaluate the benefits that the exploration system obtains from an online classification algorithm. First, we assess if stopping when a signal is detected brings some benefit to the classification performance. Then we evaluate the robustness of the system with respect to environmental condition in which the training samples are collected and the system is deployed. Lastly, we analyze if the casting behaviour, a technique often used for odour based navigation, can be used to optimize the exploration time when classification is the primary task.

In order to address these issues two different experiment locations have been set up, one indoor and one outdoor. In both the locations the path of the robot was covering an area of roughly 16 m^2. The indoor location was a room with a single opening, the door, and a ventilation system. The ventilation was regulated by a central system, activated at unpredictable intervals, and could not be changed. The outdoor environment was a courtyard between two buildings with uneven grass surface, see Figure 7. There has been no effort in trying to control environmental variables such as airflow or temperature in either of the two environments. The odour source is a cup that has been placed in the middle of the experimental area and contains one of three substances, ethanol, acetone or isopropyl.

IV. RESULTS

Table II summarizes all the exploration trials indicating both the location and the adopted path planning strategy. A lower number of trials have been carried out in the outdoor environment due to weather conditions. To collect the training data for the classification algorithm the spiral path planning strategy was used. In this strategy a number of transients were collected without stopping the robot when a rise phase was detected. In the 18 runs approximately 160 transients were collected. The remaining path planning

Fig. 7. Picture taken during an outdoor experiment.

Location	Path Planning Strategy	Number of Explorations	Number of Successes
Classroom	Spiral	18	training
Classroom	Docking on success	36	34
Classroom	Casting	36	34
Courtyard	Docking on success	9	8
Courtyard	Casting	7	6

TABLE II
NUMBER OF EXPLORATIONS PERFORMED WITH THE DIFFERENT PATH PLANNING STRATEGIES.

strategies stopped the robot when a rising phase was detected. In these 88 (36+36+9+7) runs also approximately 160 transients were collected giving equally sized data sets when stopping the robot on detection of gas vs. non stopping. Notice that performing a predefined trajectory without docking on success facilitates the collection of more transients. Thus fewer experimental runs are needed for collecting the same amount of responses.

To examine if in fact stopping the robot when a signal is detected improves the classification performance we perform a leave one out cross validation. The results of the classification for a range of different rejection threshold are shown in Figure 8. We can observe how the transients collected by stopping the robot obtain a slightly better classification rate and, more evidently, a lower rejection rate. This means that in general it will be possible to classify more transients with a high confidence degree if the robot is stopped while the response is being collected. This is a desirable property since it means that, in average, a lower number of classification attempts would be needed in order to achieve a result with a high degree of confidence. It is most likely that the higher classification performance when stopping the robot is achieved since the robot is remaining in the plume for a longer period of time. This is however difficult to confirm since there is no ground truth about the dynamics of the plume.

The overall performance of the classification algorithm is summarized in the fourth column of Table II. In all these cases the rejection threshold has been set to 0.95. When the robot performs a classification with a posterior probability greater or equal than the threshold, the exploration

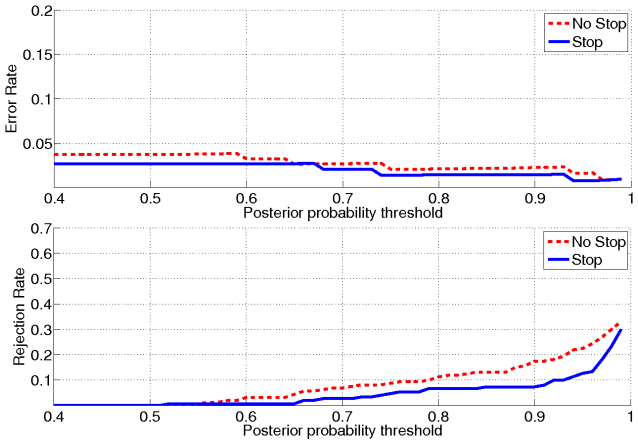

Fig. 8. Performance of the classification algorithm with a varying rejection threshold. The dashed line represents data collected without stopping the robot, while the solid line represents data collected stopping the robot every time a sensor response was detected.

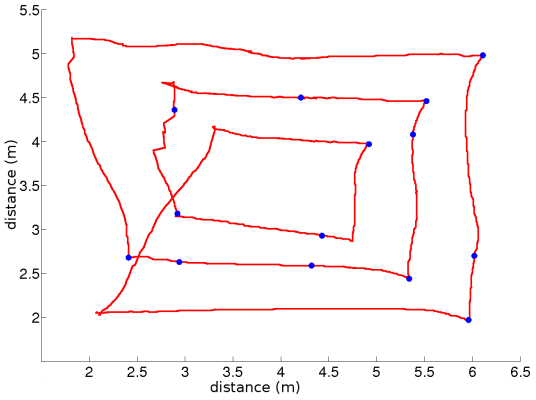

Fig. 9. Example of experiment in which the path planning strategy was a predefined spiral trajectory. The line is the trajectory followed by the robot, the dots are the locations in which responses have been detected.

is concluded and the robot goes back to docking position. A success is considered when the final classification output corresponds to the true substance placed in the environment. Although the training set has been collected without stopping the robot it is still applicable when performing experiments stopping the robot when rise phase is detected. Furthermore the training data has been collected in an indoor environment and used also in the outdoor experiment. The number of explorations in which the gas was correctly classified was high and only in few cases the robot misclassified between the acetone, ethanol and isopropyl. This is an important feature since in a real application the environmental conditions of the location where the system will be deployed will be most probably different from the ones where the system has been trained.

In addition to the evaluation of the classification performace another issue to consider is the time needed to complete an exploration depending on the adopted path planning. Table III shows average exploration time and standard deviation for the experiments performed in the classroom and outdoor. We can see that concluding the experiment when a high degree of confidence is achieved shortens the exploration time. This result is expected as the use of online classification allows the possibility to interrupt the complete spiral. A generalized quantification of the decrease in the exploration time is difficult to obtain since a rigorous experimental validation is needed where the position of the source, the starting point of the robot, the spatial properties of the inspected area and many other parameters are varied systematically. Nonetheless a reasonable estimate can be seen in the experimental validation presented here where the exploration time is shortened by half.

Figure 9 - 11 show three example runs executed using respectively the spiral, the docking on success and the casting path planning strategies. Locations where classifications are made are indicated with dots, and in Figure 10 and 11 the posterior probabilities are also indicated. Note that when the posterior probability exceeds the threshold of 0.95 the robot returns to the docking station. In Figure 9 no posterior probabilities are indicated since the data collected with these runs have been used to train the system. With respect of docking on success vs. casting strategies, we can notice that casting gives us slightly better results.

In order to gain deeper insight into the relationsip between the wind measurements and the resulting casting behaviour, we performed an experiment in which the robot was moving with a sweeping trajectory and stopping every 0.8 m for 10 s to measure the airflow in the indoor environment. An airflow map is created which overlays the spatial information from the range sensors onto a gridmap representing the interpolated measurements from the anemometer. Figure 12 shows the resulting airflow map. As we can notice the wind does not have a well defined direction, an assumption under which the casting strategy has been developed, and therefore this strategy of moving the robot is limited. We can also observe how the two casting steps performed in the experiment displayed in Figure 11 are almost perpendicular. The values from the wind vectors measured are given in Table IV. This is again a problem with the point measurement of the anemometer and shows that in turbulent environments there can be large differences in the direction and intensity of the wind. Although the casting gave a marginal improvement in these experiments, the examination of the windflow suggests that a random movement could achieve the same performance.

Exploration Methodology	Number of Explorations	Average Exploration Time	Standard Deviation
Indoor			
Spiral	18	1041 s	23 s
Docking on success	36	578 s	168 s
Casting	36	471 s	179 s
Outdoor			
Docking on success	9	431 s	319 s
Casting	7	426 s	280 s

TABLE III

AVERAGE AND STANDARD DEVIATION OF THE EXPLORATION TIME DEPENDING ON THE ADOPTED PATH PLANNING STRATEGY.

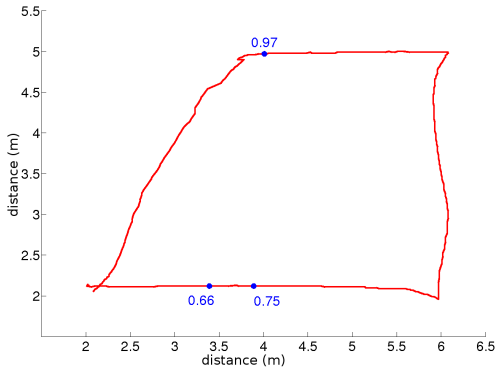

Fig. 10. Example of experiment in which the path planning strategy was stop and dock. The line is the trajectory followed by the robot, the dots are the locations in which classification attempts have taken place. Next to the dots the posterior probability of the selected class have been displayed.

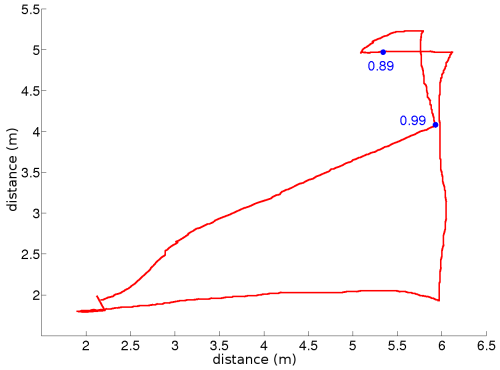

Fig. 11. Example of experiment in which the path planning strategy was casting. The line is the trajectory followed by the robot, the dots are the locations in which classification attempts have taken place. Next to the dots the posterior probability of the selected class have been displayed.

Coordinate	Average	Standard Deviation
Measurement Point 1		
X	-3.73 cm/s	2.35 cm/s
Y	0.02 cm/s	3.44 cm/s
Measurement Point 2		
X	0.47 cm/s	0.03 cm/s
Y	-2.96 cm/s	0.00 cm/s

TABLE IV

AVERAGE AND STANDARD DEVIATION OF THE WIND VECTOR MEASURED DURING THE TWO CASTING STEPS IN FIGURE 11.

V. CONCLUSION

Odour classification is a crucial capability for a mobile robot that has to explore an area and discover hazardous gases. If the classification is performed online the benefit is two-fold. First, it gives the possibility of stopping the robot when a signal is detected improving the classification performance. Second, by interrupting the exploration when a sufficient degree of confidence in the classification is achieved, the exploration time can be significantly reduced.Another important contribution is the analysis of the system behaviour when it is deployed in an area that has different properties from the one in which the system has been trained. In particular the system has been trained indoor

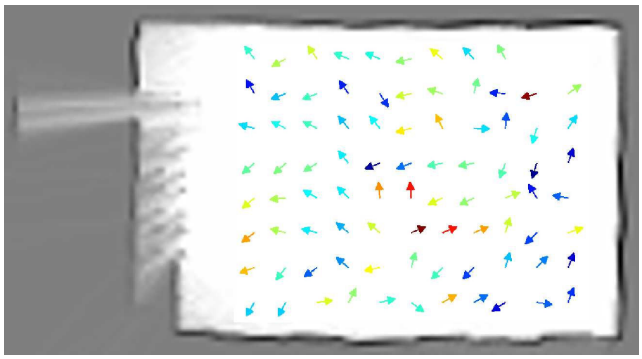

Fig. 12. Laser scan map of the classroom with overlayed arrows indicating the average wind directions. The arrows are coloured according to their relative strength ranging from blue (weak) to red (strong).

and deployed outdoor with satisfactory results. Moreover a preliminary analysis has been carried out on the possibility of exploiting the airflow information in an environment where no artificial airflow has been induced. As future work we intend to consider further experimentations. Furthermore, the system will be extended to a scenario in which more than one substance is present at the same time. Also, more complex strategy for exploiting the wind information will be considered.

VI. ACKNOWLEDGMENTS

This work has been supported by the Swedish Research Council - Vetenskapsrådet.

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] J. W. Gardner and P. N. Bartlett, *Electronic Noses: Principles and Applications*, O. U. Press, Ed. Oxford Science Publications, 1991.

[3] R. Gutierrez-Osuna, A. Gutierrez-Galvez, and N. U. Powar, "Transient response analysis for temperature modulated chemoresistors," *Sensors and Actuators B: Chemical*, vol. 93, no. 1-3, pp. 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] H. Ishida, K. Suetsugu, T. Nakamoto, and T. Moriizumi, "Study of autonomous mobile sensing system for localization of odor source using gas sensors and anemometric sensors," *Sensors and Actuators A*, vol. 45, pp. 153–157, 1994.

[6] Figaro Gas Sensor: Technical Reference, 1992. http://www.figaro.com.

[7] P. J. W. Roberts and D. R. Webster, "Turbulent Diffusion," in *Environmental Fluid Mechanics - Theories and Application*, H. Shen, A. Cheng, K.-H. Wang, M. Teng, and C. Liu, Eds. ASCE Press, Reston, Virginia, 2002.

[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] 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.

[10] M. E. Tipping, "Sparse bayesian learning and the relevance vector machine," *Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.

[11] H. Ishida, "Robotic systems for gas/odor source localization: Gap between experiments and real-life situations," in *Proceedings of the IEEE International Conference on Robotics and Automation, 2007. (ICRA 2007)*, 2007, pp. 3–8.