Gas Source Declaration with a Mobile Robot

Achim Lilienthal, Holger Ulmer, Holger Fröhlich, Andreas Stützle, Felix Werner, Andreas Zell
University of Tübingen, WSI, Sand 1, D-72076 Tübingen, Germany

***Abstract*— As a sub-task of the general gas source localisation problem, gas source declaration is the process of determining the certainty that a source is in the immediate vicinity. Due to the turbulent character of gas transport in a natural indoor environment, it is not sufficient to search for instantaneous concentration maxima, in order to solve this task. Therefore, this paper introduces a method to classify whether an object is a gas source or not from a series of concentration measurements, recorded while the robot performs a rotation manoeuvre in front of a possible source. For three different gas source positions, a total of 288 declaration experiments were carried out at different robot-to-source distances. Based on these readings, two machine learning techniques (ANN, SVM) were evaluated in terms of their classification performance. With learning parameters that were optimised by grid search, a maximal hit rate of approximately 87.5% could be obtained using a support vector machine.**

I. Introduction

The ability to classify an object depending on whether it is a source of gas or not can be useful for mobile robots for several reasons. First, it is an essential part of gas source localisation – a task, which is important for applications such as automatic humanitarian demining, or surveillance tasks including the localisation of toxic gas leaks, leaking solvents or a fire at its initial stage ("electronic watchman"). Second, the classification capability of gas source declaration itself is of potential use for rescue and security missions even if the full gas source localisation problem cannot be accomplished using a sense of smell only (because of a too low concentration at locations distant from the source, for example). An object that is to be classified could be located using other sensor modalities, and attributed based on gas sensor measurements. For example, suspicious items could be identified as containing explosive materials or a rescue robot could determine whether a victim is alive by assessing whether that person is a source of carbon dioxide. Note that CO_2 emmision belongs also to the characteristics defined in the RoboCup Rescue scenario [1] by which the simulated victims display signs of life. While in rescue scenarios other sensor modalities will also be used to check for vital signs [2], a mobile robot that is equipped with gas sensors would be able to monitor the possibly contaminated air at an emergency site. Thus, the robot can prevent rescue workers from being harmed or killed due to explosions, asphyxiation or toxication [3]. Furthermore, such a rescue robot could assemble a map of the spatial gas distribution [4], providing an incident planning staff with information to support rational decision making.

A. Gas Source Declaration Using Gas Sensors Only

Throughout the animal kingdom, many examples can be found where olfactory information plays an important role for performing different localisation tasks. These tasks include finding a mate guided by pheromones, like the male silkworm moth *Bombyx mori* for instance [5], or locating a host as is painfully known from bloodsucking mosquitoes, which are attracted by a specific mixture of human scents, often including carbon dioxide [6] and lactic acid [7]. Host-oriented behaviour of animals however, often relies on a combination of cues determined from a variety of sensor modalities. Mosquitoes and other flying haematophagous insects are also attracted by particular colours [8], while the *Bombyx* males use the local wind direction as an approximation of the direction of a pheromone source [9]. The same is probably true for the mechanisms applied to identify a host.

This paper is concerned with the classification performance that can be achieved using gas sensors only. In contrast to previous works on gas source localisation ([10], [11], [12], [13]), the environment was not artificially ventilated for the experiments presented in this work to produce a strong unidirectional airflow (see Section II-C). Without a strong artificial airflow, the detection limits of the available wind measuring devices (anemometers) are not low enough to measure weak convective airflows. With state-of-the-art anemometers based on the cooling of a heated wire [14], the bending of an artificial whisker [15] or the influence on the speed of a small rotating paddle [12], reliable readings can be obtained only for wind speeds in the order of at least 10 cm/s.

The approach that is suggested here does not depend on sufficiently high wind speeds. It tries to classify the inspected object by recognising a pattern within a series of gas sensor readings that represent temporally as well as spatially sampled concentration data. Such a pattern is determined in this work by applying machine learning techniques to a set of experiments carried out in an uncontrolled indoor environment. To the authors' knowledge there is no physically justified model available yet to establish the required pattern in case of a natural environment by analytical means.

B. Gas Distribution in Natural Indoor Environments

Due to the low diffusion velocity of gases at room temperature [16], the dispersal of an analyte gas is dominated by turbulence and the prevailing air flow rather than diffusion [17] in an uncontrolled indoor environment. The gas distribution therefore reveals many discontinous patches of local eddies [18] and the absolute maximum of the instantaneous distribution is usually *not* located near the gas source if this source has been active for some time [19]. It is therefore not sufficient to search for maxima of the instantaneous concentration distribution in order to solve the gas source declaration task.

C. Related Work

To the authors' knowledge, there exist only a few suggestions to solve the gas source declaration problem up to now. Hayes et al. [20] propose an algorithm that tries to identify a source of gas by searching for a transition between high and low intensities in upwind direction. Their gas source localisation algorithm consists of an upwind movement (surge), which is performed for a set distance whenever an "odour packet" is encountered, followed by a spiral searching behaviour for other gas patches. At the head of a plume, the Spiral Surge Algorithm tends to surge into an area of low concentration, and then spiral back to the origin of the surge before receiving another "odour hit". The vicinity to a gas source is thus expected to appear as a series of small distances between the locations where the robot senses consecutive "odour hits". However, this might occur also at locations distant from the source. Furthermore, the approach relies on a sufficiently strong and constant airflow that enables to use an anemometer and that gives a low probability for patches of gas to occur on the upwind side of the source.

As a further solution to the problem of gas source declaration in an uncontrolled environment, Lilienthal and Duckett suggest an indirect localisation strategy based on exploration and concentration peak avoidance [19]. Here, a gas source was located by exploiting the fact that local concentration maxima occur more frequently near the gas source compared to distant regions. Finally, the concentration mapping technique introduced by Lilienthal and Duckett [4] provides another possibility for gas source declaration. By combining gas sensor readings with location estimates, the suggested algorithm is able to create a representation of the average relative concentration of a detected gas in a gridmap structure. As demonstrated in [4] the position of the maximum in the representation of the average relative concentration of a detected gas can often be used to estimate the approximate location of the source. However, the latter two approaches suffer from similar drawbacks. Aside from an increased time consumption (though this can be reduced by using multiple robots) it is not guaranteed that a good estimate of the source location can be obtained with these techniques and there is yet no method available to determine the certainty of this estimate.

This paper introduces a direct declaration strategy that tries to determine whether a gas source is located in the immediate vicinity of the robot from a series of concentration measurements, recorded while the robot performed a rotation manoeuvre in front of a possible gas source (see Section III for details). A similar approach was used by Duckett et al. [21] to learn the direction to a gas source from a series of sensor readings. The gas source tracing problem is not considered here. It is rather assumed that the source appears as an obstacle to the robot, which is to be analysed after it has been detected using other sensor modalities.

The rest of this paper is structured as follows: next, the experimental setup is described in Section II and the applied declaration strategy is introduced in Section III. Then, the preprocesssing of data is detailed (Section IV) and corresponding results are discussed (Section V), followed by conclusions and suggestions for future work (Section VI).

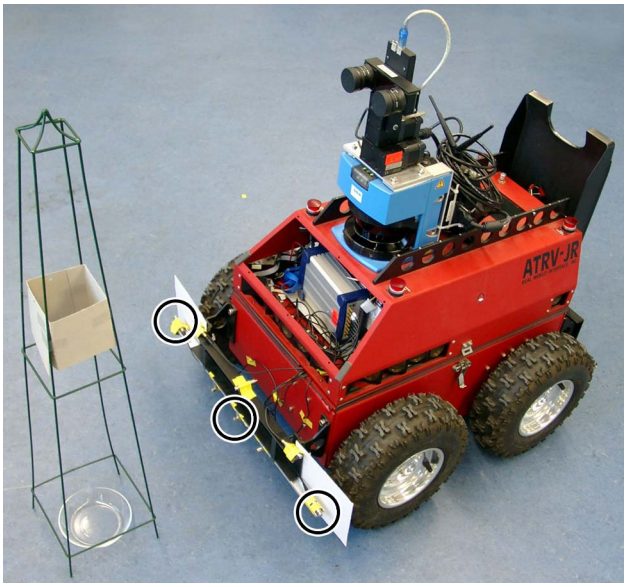

Fig. 1. The gas-sensitive mobile robot Arthur in front of the gas source. This distance was considered as being directly in front of the source. The three marked gas sensors were used for the declaration experiments presented in this work.

II. Experimental Setup

A. Robot

The gas source declaration strategy that is introduced in Section III was implemented on the gas-sensitive mobile robot "Arthur" (length = 80 cm, width = 65 cm, height without laser range scanner = 55 cm) that is based on the model ATRV-Jr. from iRobot (see Fig. 1). The robot is equipped with several external sensors. However, for the experiments presented in this work only odometry data were used in addition to the gas sensitive system. The data from the SICK laser range scanner were used to determine the position of the robot for evaluation purposes.

B. Gas Sensors

The gas sensing system is based on the commercially available device VOCmeter-Vario (AppliedSensor), which is described in detail in [22]. For the experiments presented in this paper three metal oxide sensors (Figaro, TGS 2620) were utilised. These sensors were symmetrically mounted at a height of 9 cm above the floor on the front bumper of the robot. The distance of these sensors to the middle of the bumper was 0 cm and ± 32 cm. The distance between the outer sensors and the front wheels is very small. In order to avoid a corruption of the results due to an additional airflow created by the wheels, a shield made of cardboard was placed inbetween the wheels and the sensors (see Fig. 1).

Metal oxide sensors comprise a heating element coated with a sintered semiconducting material. The measured quantity is the resistance R_S of the surface layer at an operating

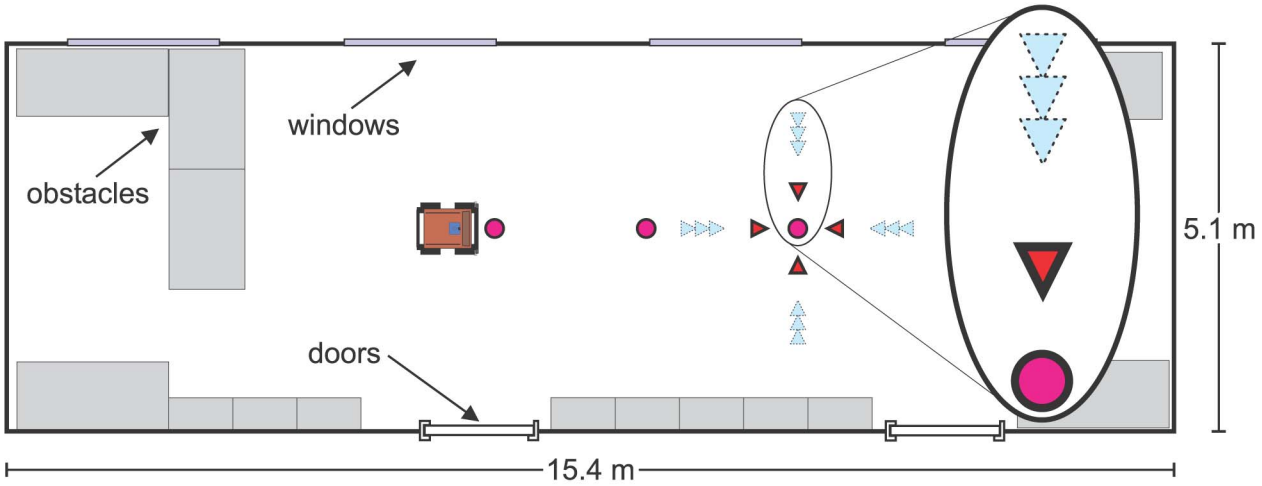

Fig. 2. Floor plan of the laboratory room in which the experiments were performed. Also indicated are the windows at the upper and the doors at the lower side as well as the obstacles in the room (cupboards and desks). In addition, the tested locations of the gas source are indicated by circles. Beneath the source on the left side, the robot is sketched in a position that was considered as being directly in front of the source. Further on, all the tested robot positions are shown for the rightmost source location using triangles that indicate the centre of the robot and its initial heading. Light triangles with a dotted border indicate positions that were considered as being not in the immediate vicinity of the source.

temperature of between 300°C and 500°C [23]. Exposed to a reducing gas, the potential barrier at the grain boundary is lowered, and thus the resistance of the surface layer decreases. In consequence of the measurement principle, metal oxide sensors exhibit some drawbacks. Namely the low selectivity, the comparatively high power consumption (caused by the heating device) and a weak durability. Furthermore, metal oxide sensors are subject to a long response time and an even longer decay time [24]. However, this type of gas sensor is most often used for mobile noses because it is inexpensive, highly sensitive and relatively unaffected by changing environmental conditions like room temperature or humidity.

C. Environment and Gas Source

All experiments were carried out in a 15.4 m × 5.1 m room at the University of Tübingen. A floor plan is shown in Fig. 2, including doors, windows, cupboards and desks. In addition, the tested gas source positions are indicated by circles. A total of 288 declaration trials were performed using three different source locations and four different orientations with respect to the source as indicated in Fig. 2. For each source position, 48 experiments were carried out directly in front of the gas source ($d = d_0$) alternating with 48 trials at a randomly chosen larger distance of $d = d_0 + \Delta d$ with Δd = 60 cm, 80 cm and 100 cm, respectively. After each trial, the robot was stopped for 60 s in order to avoid disturbance from the preceding measurements due to the long decay time of the sensors. All the robot positions tested are shown for the right source position, using triangles that indicate the centre of the robot and its initial heading. Light triangles with a dotted border indicate positions that were considered as being not in the immediate vicinity of the source.

With regard to real world applications, the environment was not modified for this investigation. The unventilated room was also used as an office during the experiments, with up to two persons working, moving and sometimes leaving or entering the room. Although the windows were kept closed and the persons were told to be careful, this indoor environment can be considered uncontrolled to some extent.

The gas source was chosen to be a bowl with a diameter of 140 mm and a height of 20 mm filled with Single Malt Whiskey (40% alcohol), which was used because it is non-toxic, less volatile than pure ethanol and easily detectable by metal oxide sensors. In order to be recognisable by the laser range scanner, a frame made of wire with a cardboard marking mounted on top was placed above the container (see Fig. 1).

III. Gas Source Declaration Strategy

Due to the properties of gas distribution in real world environments discussed in Section I-B, single concentration measurements do not contain enough information to allow determination of the proximity to a gas source. It was instead considered most promising to apply a strategy that provides temporally as well as spatially sampled concentration data.

Therefore, the gas sensor readings were acquired while the robot performs a rotation manoeuvre containing three successive rotations: 90° to the left, then 180° to the right (without stopping, in order to minimise self-induced disturbance of the gas distribution) and finally 90° to the left again (see Figure 3). Initially, the robot was oriented towards the suspected object as indicated in Fig. 3. This manoeuvre is easy to implement, requires little space and does not involve periods of backward motion where the ATRV-Jr robot offers only a limited obstacle avoidance capability. The rotation was performed with an angular speed of approximately 4°/s corresponding to a total time of approximately 90 s to complete the manoeuvre. Simultaneously, sensor readings were acquired at the maximum rate of almost 4 Hz, resulting in a total of Q readings per experiment with $Q \in$ [349,362].

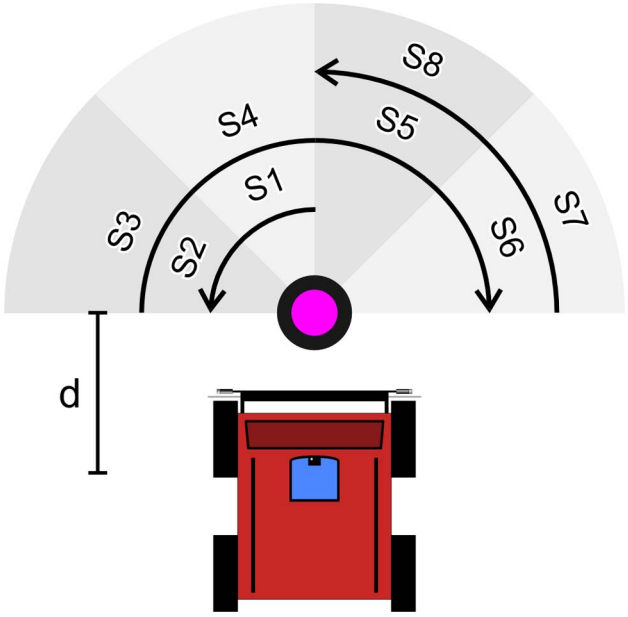

Fig. 3. Rotation manoeuvre performed to collect sensor data for gas source declaration. Indicated are the initial robot position, the gas source, the three successive rotations (given by arrows starting with the innermost one) and the sectors for which the mean and standard deviation is calculated as a feature.

IV. Data Pre-Processing

To evaluate the performance of the two machine learning methods tested (artificial neural network and support vector machine), the recorded data were first pre-processed by means of feature extraction (Section IV-A) and normalisation (Section IV-B). Next, an output value was added to each data set, indicating whether the corresponding experiment was performed directly in front of a gas source (+1) or not (–1). The robot was considered as being in the "immediate vicinity of a source" only in the case of minimal distance between the robot and the gas source, corresponding to a laser scanner reading of $d = d_0 = 50$ cm (see Fig. 3). Here, the trajectory of the sensors just avoids hitting the object under inspection at the point of closest approximation. By contrast, all the positions with a larger distance $d \geq d_0 + \Delta d^{ns}_{min}$ were considered as being "not in the immediate vicinity of a source". The minimal distance Δd^{ns}_{min} for the negative examples was 60 cm and the average distance $\overline{\Delta d^{ns}}$ was 80 cm.

A. Feature Extraction

The features used for classification were derived by calculating the first two statistical moments (mean and standard deviation of the sensor measurements) for each of the 8 consecutive 45° sectors covered by the rotation manoeuvre. These sectors are denominated by S1 – S8 in Figures 3 and 4. Depending on the number M of gas sensors utilised, a maximum of $M\times16$ features was extracted. Either the full $M\times16$-dimensional input vector was utilised for training and testing, or only the $M\times8$ mean or standard deviation values. Examples of feature vectors obtained in the experiments are depicted in Fig. 4. Here, the vertically normalised mean values of the leftmost, middle and rightmost sensor are plotted in

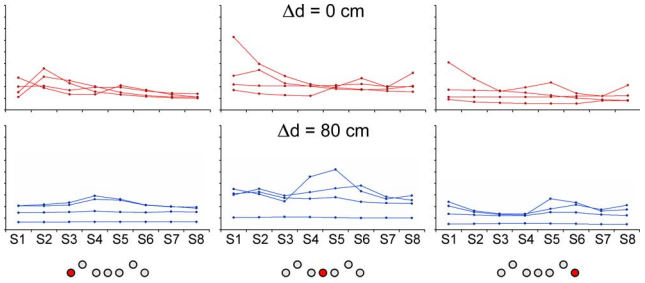

Fig. 4. Examples of mean values obtained at the indicated distance from the gas source. For each distance, four vertically normalised feature vectors are shown for the leftmost, middle and rightmost sensor (indicated by the iconic front view below each column).

order to indicate the relative strength of the sensor responses. Four examples are depicted recorded at a distance of d_0 (in the immediate vicinity of the gas source) and at a distance of d_0 + 80 cm that was considered as being not in the immediate vicinity of the gas source.

B. Normalisation

The set of feature vectors $\vec{F}_i$ (corresponding to the desired classification t_i of the i-th experiment) creates a matrix F_{ij} ($j \in [1, M\times8]$ or $j \in [1, M\times16]$ and $i \in [1,N]$ with the number of experiments N and the number of sensors M). Before training and testing, this matrix is normalised *vertically*, meaning that each column is mapped linearly to the range of [0,1] as

$$f^v_{ij} = \frac{F_{ij} - min_i\{F_{ij}\}}{max_i\{F_{ij}\} - min_i\{F_{ij}\}}. \quad (1)$$

Note that this kind of normalisation cannot be applied in the same way for classification of a single trial because it is necessary to know all N experiments in order to establish the normalisation range. It might be also problematic to apply the normalisation factors obtained from the training data in a test experiment in the case of varying environmental conditions that cause a shift of the sensor values, such as a different temperature or humidity. Finally, the vertical normalisation factors contain knowledge about the intensity of the gas source used in the training phase, and could thus be misleading in the case of a different source. For online evaluation of a single experiment rather *horizontal* normalisation could be used:

$$f^h_{ij} = \frac{F_{ij} - min_j\{F_{ij}\}}{max_j\{F_{ij}\} - min_j\{F_{ij}\}}. \quad (2)$$

While in the case of vertical normalisation, the available information about the strength of the sensor response (relative to the range experienced in all the experiments) is included in the feature vector, a horizontally normalised feature vector represents the relative intensity of the sensor response with respect to the values that occur during the rotation manoeuvre. Therefore, examples have to be classified in the latter case based on the relative course of the concentration measurements only. For real world applications, however, the concentration measurements collected before the rotation manoeuvre started can also be used to acquire an approximation of the range that is used for vertical normalisation. A similar classification

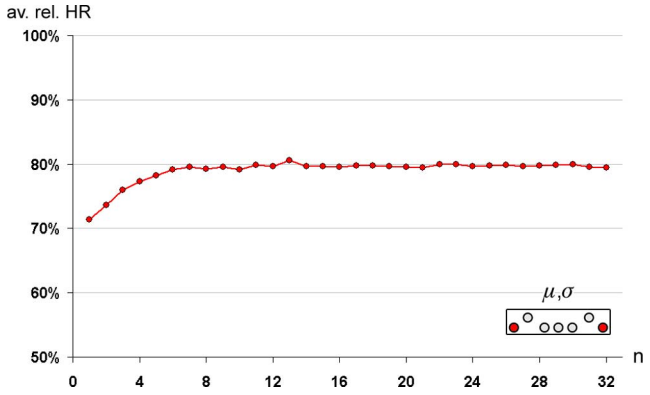

Fig. 5. Total hit rate of a multilayer feedforward network depending on the number n of hidden neurons. The input layer contains 32 neurons corresponding to the mean and standard deviation of the readings of the left- and rightmost sensor.

performance as obtained after vertical normalisation can thus be expected for real world applications where the robot collects gas sensor readings on its way to inspected objects.

V. Results

At each of the gas source positions indicated in Fig. 2, four experiments were carried out at four different directions (north, east, south, west) and three different distances Δd, alternating with four experiments in the direct vicinity of the gas source (see Fig. 2). Thus, a total of $N = 3 \times 4 \times 4 \times 3 \times 2 = 288$ declaration trials were performed including 144 trials (50%) in the immediate vicinity of the source ($d = d_0$) and 144 experiments (50%) at a larger distance of $d = d_0 + \Delta d$ with $\Delta d \geq 60$ cm. Using the obtained data set, the two pattern recognition algorithms were evaluated by means of 5-fold cross-validation. In order to increase the accuracy of the evaluation, the hit rate (the percentage of correctly classified examples) was calculated by averaging over fifteen 5-fold cross-validation runs.

A. Artificial Neural Network

One way to solve non-linear classification problems is provided by artificial neural networks. Here, a multilayer feedforward (MLFF) network with a sigmoidal activation function was used, containing an input layer with $N_{in} = M \times 8$ or $M \times 16$ neurons, an output layer of one unit and a hidden layer with a variable number of $n \in [1, N_{in}]$ neurons. The achieved classification rate, however, was found to reach an approximately stable level below $n \approx 8$ hidden neurons. This can be seen in Fig. 5, which shows the hit rate depending on the number of hidden neurons using the mean and standard deviation of the two outermost sensors as a 32-dimensional feature vector. Training was performed using conjugate gradient descent [25] with 100 training cycles and a weighting of $\alpha = 0.1$ for the momentum term, which was found to yield good results in initial tests. The weights were initialised with a randomly chosen set of values at the beginning of each run.

B. Support Vector Machine

Support Vector Machines (SVM) [26] have become established during the recent years as a major state-of-the-art classification method. The main idea is to construct a so-called optimal separating hyperplane between two classes A and B, which maximises the margin between the convex hulls of A and B. It is understood that a larger margin leads to a better generalization performance of the SVM [27]. Learning thus corresponds to a dual optimisation problem, for which a unique and global optimal solution can be obtained by quadratic programming. If A and B are not linearly separable, the so called *kernel trick* is applied. All input patterns are mapped via a nonlinear function Φ into a high dimensional *feature space* where the problem is linearly separable again. It is, however, usually not necessary to perform this transformation directly. Instead, the *kernel function* k can be used, which can be thought of a similarity measure between two patterns x and y, representing a dot product in feature space. As a popular kernel function the radial basis function is used here as

$$k_\gamma(x,y) = \exp\left(-\frac{\|x-y\|^2}{\gamma^2}\right). \quad (3)$$

In order to find suitable learning parameters, a grid search was carried out in the two-dimensional search space spanned by the kernel parameter γ and the parameter C that determines the extent to which outliers are penalised. At this, 2205 points were sampled for each set of feature vectors at $\gamma = 2^{-3}, 2^{-2.75}, \ldots, 2^{8}$ and $C = 2^{-6}, 2^{-5.75}, \ldots, 2^{6}$, which is the parameter range where all the optimal combinations were found in initial tests.

Corresponding results are itemised in Table I, which shows a comparison of the classification performance achieved with the artificial neural network and the support vector machine. The first two columns specify the used feature vector, including the considered sensors and statistical moments (mean "μ" and/or standard deviation "σ"). Then, the best result in terms of maximum hit rate (HR) obtained with a MLFF neural network (third and fourth column) and the support vector machine (fifth and sixth column) is given. In addition, the corresponding average cross-validation rate of false positives (FP) and false negatives (FN) is also itemised. Finally, the corresponding

Δd^{ns}_{min} = 60 cm		MLFF neural network		SVM	
Sensors	Features	n^*	HR (FP, FN)[%]	C^*, γ^*	HR (FP, FN)[%]
∘°∘•∘°∘	μ (8)	6	75.7 (22.5, 26.1)	$2^{1.25}, 2^{3.25}$	77.1 (24.5, 21.3)
	σ (8)	8	75.5 (20.3, 28.7)	$2^{-0.25}, 2^{4.75}$	77.1 (25.4, 20.3)
	μ, σ (16)	9	75.6 (22.5, 26.3)	$2^{8.25}, 2^{4.5}$	78.1 (13.6, 30.2)
•°∘∘∘°•	μ (16)	16	76.8 (19.9, 26.4)	$2^{4.25}, 2^{4}$	80.0 (15.1, 24.9)
	σ (16)	13	76.3 (20.4, 27.0)	$2^{3.25}, 2^{1}$	78.3 (26.4, 17.1)
	μ, σ (32)	13	80.6 (19.5, 19.4)	$2^{3.75}, 2^{-0.25}$	83.0 (19.0, 15.0)
•°∘•∘°•	μ (24)	16	82.7 (16.3, 18.4)	$2^{5.5}, 2^{2.5}$	87.5 (10.3, 14.8)
	σ (24)	23	80.8 (15.6, 22.7)	$2^{5}, 2^{0.5}$	83.4 (18.7, 14.5)
	μ, σ (48)	24	84.1 (14.9, 16.9)	$2^{6.75}, 2^{0.5}$	86.4 (10.1, 17.2)

TABLE I: Comparison of the classification performance obtained with a MLFF network and the SVM for an average distance of 80 cm of negative examples from the source.

parameters for which the best classification performance was achieved (the number of hidden neurons n^* in the case of the MLFF network and the penalty and kernel parameter (C^*, γ^*) in the case of support vector machines) are also given.

With optimised learning parameters the support vector machine always yielded slightly better results compared to the MLFF neural net that was used in this investigation. The highest total hit rate of 87.5% was achieved with the SVM using the mean values of the three considered gas sensors as a feature vector.

VI. Conclusions

This paper is concerned with the task of gas source declaration. It introduces a classification method based on gas sensor readings only. In order to decide whether a gas source is in the direct vicinity, the robot collects gas sensor readings while it performs a rotation manoeuvre in front of a suspected object.

For the chosen average distance of 80 cm between positive and negative examples, a maximum hit rate of 87.5% could be achieved using a support vector machine, whereas the cross-validation rate of false positives was 10.3% and the rate of false negatives was 14.8%. This performance could be improved by combining single predictions obtained from data recorded at different positions. One possibility would be to repeat the rotation maneouvre three times at a different orientation with respect to the source and to use the majority vote as the new prediction. Thus, a majority vote is guaranteed while single estimates can be assumed to be independent due to the different direction of the convective airflow at different orientations of the robot in relation to the gas source. Under the assumption of independent estimates, such a combined classifier would yield a total hit rate of 95.6% (FP = 3.0%, FN = 5.9%). Moreover, it is expected, that the classification performance could be improved by using multiple series of measurements recorded at different positions (as one feature vector), because in this way the difference between the sensor response at the downstream and upstream side of the gas source could also be used for classification. Our ongoing work is concerned with this issue. Future work will also address the performance of online classification methods (using horizontal normalisation) and the dependance of the classification rate on the chosen distance from the gas source between positive and negative examples.

References

[1] H. Kitano, S. Tadokor, H. Noda, I. Matsubara, T. Takhasi, A. Shinjou, and S. Shimada, "RoboCup Rescue: Search and Rescue for Large Scale Disasters as a Domain for Multi-Agent Research," in *Proceedings of the IEEE Conference on Systems, Men, and Cybernetics*, 1999.

[2] J. Casper and R. Murphy, "Workflow Study on Human-Robot Interaction in USAR," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2002)*, 2002, pp. 1997 – 2003.

[3] R. Murphy, J. Casper, J. Hyams, M. Micire, and B. Minten, "Mobility and Sensing Demands in USAR," in *Proceedings of IECON 2000 (invited)*, 2000.

[4] A. Lilienthal and T. Duckett, "Creating Gas Concentration Gridmaps with a Mobile Robot," in *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, Las Vegas, USA, 2003, pp. 118 – 123.

[5] R. Kanzaki, "Behavioral and Neural Basis of Instinctive Behavior in Insects: Odor-Source Searching Strategies without Memory and Learning," *Robotics and Autonomous Systems*, vol. 18, pp. 33–43, 1996.

[6] M. C. Pinto, D. H. Campbell-Lendrum, A. L. Lozovei, U. Teodoro, and C. R. Davies, "Phlebotomine Sandfly Responses to Carbon Dioxide and Human Odour in the Field," *Medical and Veterinary Entomology*, vol. 15, pp. 132–139, 2001.

[7] T. Dekker, B. Steib, R. T. Cardé, and M. Geier, "L-Lactic Acid: A Human-Sgnifying Host Cue for the Anthropophilic Mosquito *Anopheles Gambiae*," *Medical and Veterinary Entomology*, vol. 16, pp. 91–98, 2002.

[8] G. Gibson and S. J. Torr, "Visual and Olfactory Responses of Haematophagous Diptera to Host Stimuli," *Medical and Veterinary Entomology*, vol. 13, pp. 2–23, 1999.

[9] R. Kanzaki, "Coordination of Wing Motion and Walking Suggests Common Control of Zigzag Motor Program in a Male Silkworm Moth," *J Comp Physiol A*, vol. 182, pp. 267–276, 1998.

[10] H. Ishida, Y. Kagawa, T. Nakamoto, and T. Moriizumi, "Odour-Source Localization in the Clean Room by an Autonomous Mobile Sensing System," *Sensors and Actuators B*, vol. 33, pp. 115–121, 1996.

[11] R. A. Russell, D. Thiel, R. Deveza, and A. Mackay-Sim, "A Robotic System to Locate Hazardous Chemical Leaks," in *IEEE Int Conf. Robotics and Automation (ICRA 1995)*, 1995, pp. 556–561.

[12] R. A. Russell, L. Kleeman, and S. Kennedy, "Using Volatile Chemicals to Help Locate Targets in Complex Environments," in *Proceedings of the Australian Conference on Robotics and Automation*, Melbourne, Aug 30- Sept 1 2000, pp. 87–91.

[13] A. Hayes, A. Martinoli, and R.M.Goodman, "Swarm Robotic Odor Localization," in *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-01)*, vol. 2, Maui, Hawaii, USA, October 2001, pp. 1073–1078.

[14] 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.

[15] R. A. Russell and A. H. Purnamadjaja, "Odour and Airflow: Complementary Senses for a Humanoid Robot," in *IEEE Int Conf. Robotics and Automation (ICRA 2002)*, 2002, pp. 1842–1847.

[16] T. Nakamoto, H. Ishida, and T. Moriizumi, "A Sensing System for Odor Plumes," *Analytical Chem. News & Features*, vol. 1, pp. 531–537, August 1999.

[17] M. R. Wandel, A. Lilienthal, T. Duckett, U. Weimar, and A. Zell, "Gas Distribution in Unventilated Indoor Environments Inspected by a Mobile Robot," in *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, 2003, pp. 507–512.

[18] R. A. Russell, *Odour Sensing for Mobile Robots*. World Scientific, 1999.

[19] A. Lilienthal and T. Duckett, "Experimental Analysis of Smelling Braitenberg Vehicles," in *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, 2003, pp. 375–380.

[20] A. Hayes, A. Martinoli, and R. Goodman, "Distributed Odor Source Localization," *IEEE Sensors Journal, Special Issue on Electronic Nose Technologies*, vol. 2, no. 3, pp. 260–273, 2002, june.

[21] T. Duckett, M. Axelsson, and A. Saffiotti, "Learning to Locate an Odour Source with a Mobile Robot," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2001)*, Seoul, South Korea, May, 21–26 2001.

[22] A. Lilienthal, A. Zell, M. R. Wandel, and U. Weimar, "Sensing Odour Sources in Indoor Environments Without a Constant Airflow by a Mobile Robot," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2001)*, 2001, pp. 4005–4010.

[23] J. W. Gardner and P. N. Bartlett, *Electronic Noses - Principles and Applications*. Oxford: Oxford Science Publications, 1999.

[24] A. Lilienthal and T. Duckett, "A Stereo Electronic Nose for a Mobile Inspection Robot," in *Proceedings of the IEEE International Workshop on Robotic Sensing (ROSE 2003)*, Örebro, Sweden, 2003.

[25] D. E. Rumelhart and J. L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*. MIT Press, 1986.

[26] C. Cortes and V. Vapnik, "Support Vector Networks," *Machine Learning*, vol. 20, pp. 273 – 297, 1995.

[27] V. Vapnik, *The Nature of Statistical Learning Theory*. New York, NY: Springer, 1995.