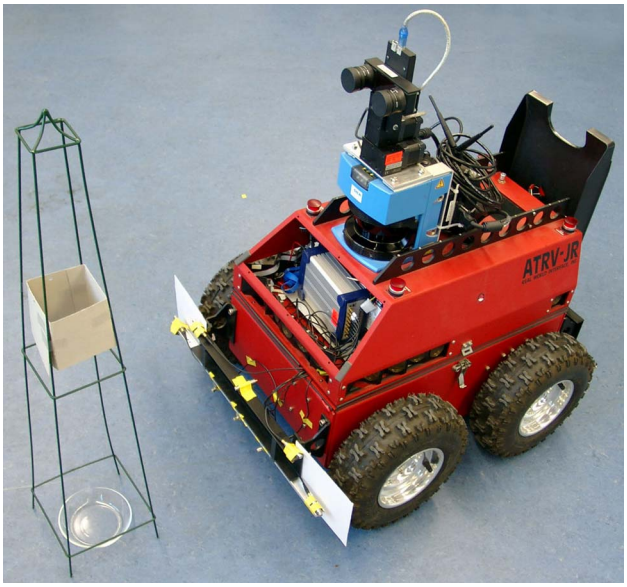

Fig. 1. The gas-sensitive mobile robot Arthur in front of the gas source. This distance was considered as being directly in front of the source.

edge about the features of a gas source that appear to other sensor modalities might also assist the declaration step. However, this information will often not be available in a typical application scenario. It is thus desired to utilise more general characteristics to determine proximity to a gas source. In an environment with a sufficiently strong unidirectional airflow, a gas source could be identified by determining a concentration drop on the upwind side of the source [6], [18], [19] or by detecting a reducing plume width while approaching the source [17]. However, because detailed experimental evaluations are not available, it remains an open question as to what amount of temporal averaging is necessary to extract these characteristics from the turbulent concentration field. An interesting alternative would be to utilise positional clues in the fine structure of a turbulent distribution [20], [17]. However, to the best of the authors' knowledge, a detailed analysis of the gas source declaration performance based on experimental tests is not yet available.

While a few experiments have been published where a gas source localisation strategy based on an analytical model was applied that includes the declaration step [21], [22], no evaluation of the corresponding declaration performance by statistical means is available either. A further solution to the problem of gas source declaration, which falls in the last defined category, is provided by the reactive localisation strategy based on exploration and concentration peak avoidance that was suggested by Lilienthal and Duckett in [16]. 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. Another possibility for gas source declaration is provided by the concentration mapping technique introduced by Lilienthal and Duckett 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 [23]. 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.

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). For the experiments presented in this work only odometry data were utilised in addition to the concentration measurements. 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 [24]. For the gas source declaration experiments seven metal oxide sensors were utilised which were placed on the robot as shown in Fig. 2. Metal oxide gas sensors are discussed in [25] especially referring to their use on a mobile robot. Five TGS 2620 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, $\pm$16 cm, and $\pm$40 cm. Two additional sensors of type TGS 2600 were mounted at a height of 16 cm with a distance of $\pm$32 cm to the centre. 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

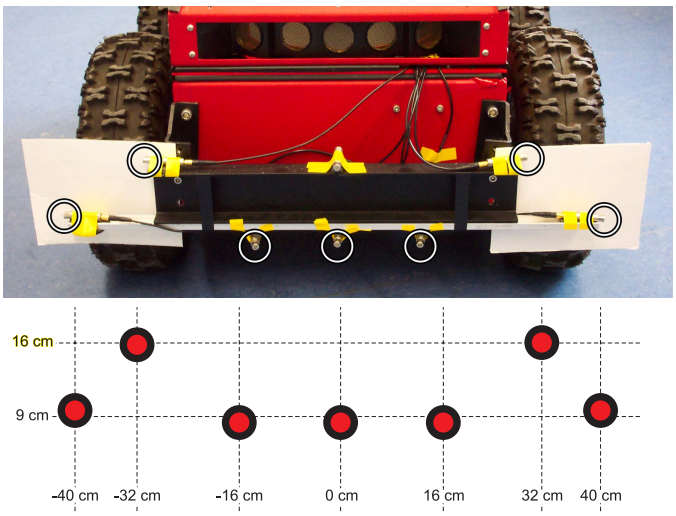

Fig. 2. Setup of the gas sensor array.

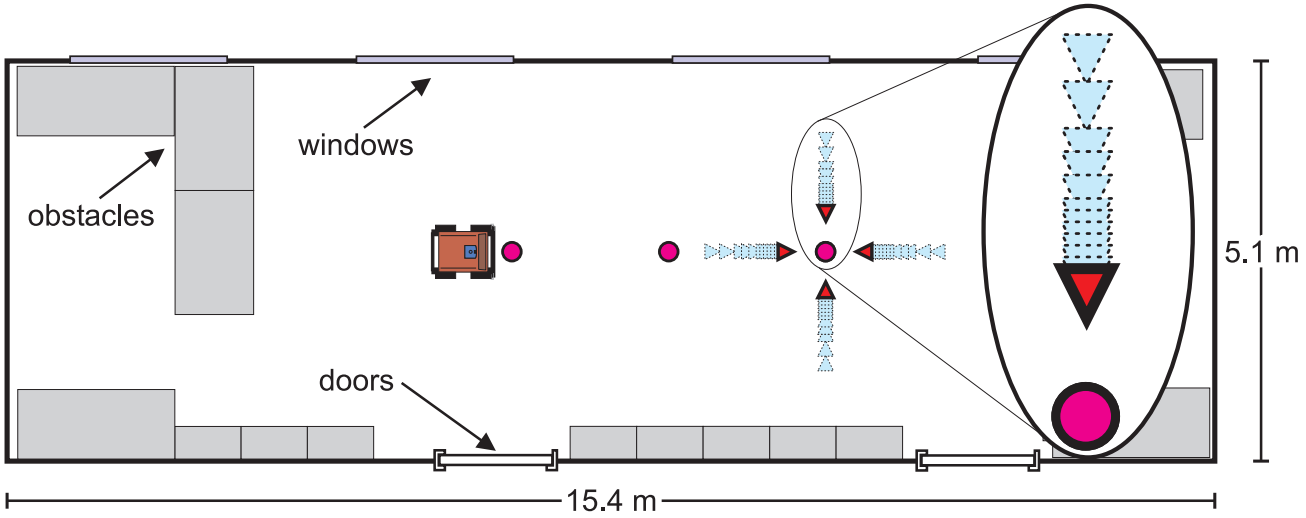

Fig. 3. Floor plan of the laboratory room in which the experiments were performed. Also indicated are the windows at the upper and the doors at the lower side as well as the obstacles in the room (cupboards and desks). In addition, the tested locations of the gas source are indicated by circles. Beneath 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.

was placed inbetween the wheels and the sensors (see Figs. 1 and 2).

C. Environment and Gas Source

All experiments were carried out in a 15.4 m $\times$ 5.1 m room at the University of Tübingen. A floor plan is shown in Fig. 3, including doors, windows, cupboards and desks. In addition, the tested gas source positions are indicated by circles. A total of $N = 1056$ declaration trials were performed using three different source locations and four different orientations with respect to the source as indicated in Fig. 3. For each source position, 176 experiments were carried out at a distance d directly in front of the gas source ($d = d_0$) alternating with 176 trials at a randomly chosen larger distance of $d = d_0 + \Delta d$ with Δd = 5 cm, 10 cm, 15 cm, 20 cm, 25 cm, 30 cm, 40 cm, 50 cm, 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-A, single concentration measurements do not contain enough information to allow determination of 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 4). Initially, the robot was oriented towards the suspected object as indicated in Fig. 4. 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 approximately 4 Hz, resulting in a total of Q readings per experiment with $Q \in [349,362]$.

IV. Data Pre-Processing

To evaluate the performance of the 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

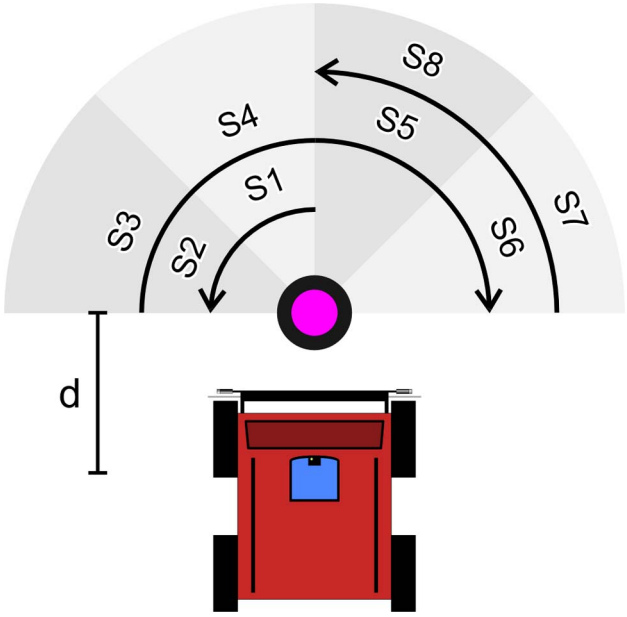

Fig. 4. Rotation manoeuvre performed to collect sensor data for gas source declaration. Indicated are the initial robot position, the gas source, the three successive rotations (given by arrows starting with the innermost one) and the sectors for which the mean and standard deviation is calculated as a feature.

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. 4). 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 $d \geq d_0 + \Delta d_{min}^{ns}$ were considered as being "not in the immediate vicinity of a source".

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 Fig. 4. Depending on the number M of gas sensors utilised, a maximum of $M \times 16$ features was extracted. Either the full $M \times 16$-dimensional input vector was utilised for training and testing, or only the $M \times 8$ mean or standard deviation values. Examples of the obtained feature vectors are depicted in [26].

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} ($j \in [1, M \times 8]$ or $j \in [1, M \times 16]$ 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

$$f_{ij}^{v} = \frac{F_{ij} - min\{F_{\bullet j}\}}{max\{F_{\bullet j}\} - min\{F_{\bullet j}\}}. \quad (1)$$

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 *horizontal* normalisation could be used:

$$f_{ij}^{h} = \frac{F_{ij} - min\{F_{i\bullet}\}}{max\{F_{i\bullet}\} - min\{F_{i\bullet}\}}. \quad (2)$$

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. The results obtained with horizontal and vertical normalisation provide therefore a lower and upper boundary of the classification performance that can be achieved in 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. 3, four experiments were carried out at four different directions (north, east, south, west) and eleven different distances Δd, alternating with four experiments in the direct vicinity of the gas source (see Fig. 3). Thus, a total of $N = 3 \times 4 \times 4 \times 11 \times 2 = 1056$ declaration trials were performed including 528 trials (50%) in the immediate vicinity of the source ($d = d_0$) and 528 experiments (50%) at a larger distance of $d = d_0 + \Delta d$ with $\Delta d \geq 5$ cm.

Based on the obtained data set, the classification performance that could be achieved with support vector machines [27] was evaluated by means of 5-fold cross-validation ("SVM classifier") for five different sensor combinations. These sensor combinations are indicated in the upper box in Fig. 5 by five symbols that show an iconic front view of the robot as in Fig. 2. 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. As a popular kernel function the radial basis function was used here as

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

In order to find suitable learning parameters, a grid search was carried out in the two-dimensional search space spanned

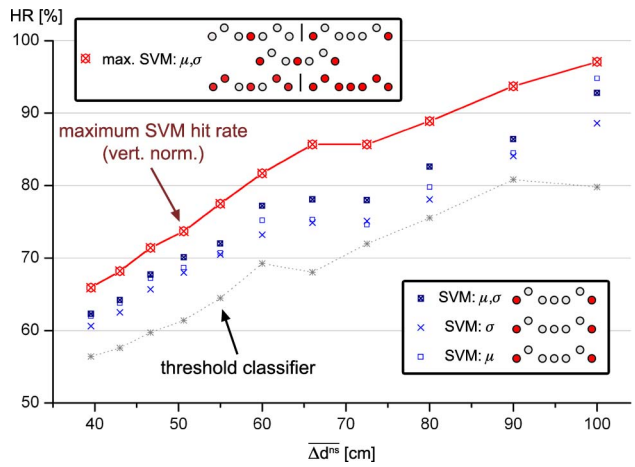

Fig. 5. Comparison of the classification performance obtained with the threshold classifier and the support vector machine using vertical normalisation. The achieved hit rate is plotted against the average distance of negative examples from the gas source. Different feature vectors were considered as indicated in the figure.

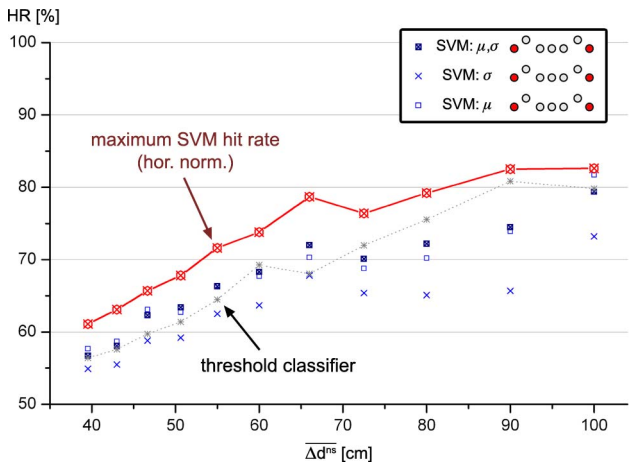

Fig. 6. Comparison of the classification performance obtained with the threshold classifier and the support vector machine using horizontal normalisation.

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 $\gamma = 2^{-3}, 2^{-2.75}, ..., 2^{8}$ and $C = 2^{-6}, 2^{-5.75}, ..., 2^{6}$, which covers the parameter range where all the optimal combinations were found in initial tests. More details and a comparison with the classification performance that was achieved with a feedforward neural network are given in [26].

In order to assign the training examples to different categories according to whether they are recorded in the immediate vicinity of the gas source or not, the examples were separated by the distance to the source at which the rotation manouevre was performed. Positive examples were assumed when the data were collected at the minimal distance d_0 where the trajectory of the gas sensors just avoids hitting the source at the point of closest approximation. Negative examples were assumed when the data were collected at larger distances of $d \geq d_0 + \Delta d^{ns}_{min}$. The correlation between the classification performance and the mean distance of negative examples to the source was investigated by disregarding those trials with $0 < \Delta d < \Delta d^{ns}_{min}$ for evaluation. In order to preserve an even proportion between positive and negative examples, the loss of negative examples was then compensated by omitting the same amount of randomly chosen positive examples.

As a reference that indicates the value of measuring the absolute intensity in order to determine proximity to a gas source, the classification performance that can be achieved by selecting an optimal threshold value regarding the mean sensor signal during the rotation manoeuvre was also tested as a possible method for gas source declaration ("threshold classifier"). However, apart from the problem of sensor drift due to changing environmental conditions or ageing of the sensors [25], the suitability of the threshold classifier is limited by the fact that a weak sensor response occurs in the case of non-zero wind speeds also if the robot is located upwind from the source. Moreover, the signal obtained from a real distribution is superimposed by strong local concentration variations due to turbulence (typically with high peak-to-mean ratios of 10 or more [15]) and depends also on the time since the source was uncovered because a non-stationary situation is considered.

Fig. 5 shows the maximum hit rate that could be obtained with the support vector machine using vertical normalisation and by selecting an optimal threshold on the raw data ("threshold classifier"). The optimal threshold was acquired with respect to the mean sensor signal during the rotation manoeuvre, meaning that the threshold classifier refers to a time-averaged concentration value. It should be noted that no cross-validation was carried out in order to determine the performance of the threshold classifier.

The general trend of the classification performance against the average distance of negative examples from the gas source can be seen in Fig. 5. This plot shows a comparison of the cross-validation hit rate obtained with the support vector machine using vertical normalisation and the performance of the threshold classifier. Apart from the maximum hit rate that was achieved with the support vector machine, considering all the feature vectors, individual results for three selected feature vectors are also given. In the same way, Fig. 6 shows the results obtained with horizontal normalisation.

As expected, the classification performance decreases with decreasing distance to the source. The observed course, however, is not linear and three different regions can be distinguished. An approximately linear descent was found when the average distance of negative examples was above 72.5 cm (corresponding to $\Delta d^{ns} \geq 50$ cm) or below 60 cm ($\Delta d^{ns} \leq 30$ cm) with a roughly constant plateau in between. A similar profile was observed for all of the classifiers tested, probably indicating a transition between domains where proximity to the gas source can be detected using different properties of the concentration field.

The classification performance that was achieved with the support vector machine using vertical normalisation was generally higher compared to the performance of the threshold

classifier. This result corresponds to the fact that information about the relative course of the sensor signal can be exploited by the SVM in addition to information about the absolute intensity, which is only used by the threshold classifier. In contrast to vertical normalisation, horizontal normalisation does not preserve the absolute intensity. Although the achieved classification performance was generally lower using horizontally normalised feature vectors, the maximum hit rate achieved with the SVM was nevertheless higher than the performance of the threshold classifier (see Fig. 6).

While using only the mean values of each sector or only the standard deviation yielded similar hit rates in the case of vertical normalisation, the relative course of the standard deviation was found to be of little value with horizontal normalisation. Consequently, the maximum hit rate with vertical normalisation was most often found using both the mean values and the standard deviations, while the additional benefit of using also the standard deviations was rather small in the case of horizontal normalisation. This tendency can be seen with the unconnected symbols in Figs. 5 and 6, which indicate the results that were obtained using only the two outermost sensors.

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.

The results of this ongoing work demonstrate the feasibility of the approach and show that high classification rates can be achieved using support vector machines. An analysis of the classification rate depending on the desired declaration accuracy is also presented in this paper, and a comparison of the performance of the SVM classifier with the classification rate that can be achieved by selecting an optimal threshold on the raw data (threshold classifier) is given.

Future work will address the question, which are the important features for classification (feature selection), the prediction capability of the distance to the source (regression) and the suitability of Bayesian learning techniques.

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] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] C. G. Lomas, *Fundamentals of Hot Wire Anemometry*. Cambridge University Press, 1986.

[12] M. Pluijm, G. Sars, and C. H. Massen, "Calibration Unit for Micro-Anemometers at Very Low Air Velocities," *Appl. Sci. Res.*, vol. 43, pp. 227 – 234, 1986.

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

[14] 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.

[15] P. J. W. Roberts and D. R. Webster, "Turbulent Diffusion," in *Environmental Fluid Mechanics - Theories and Application*, K. T. John Steele, Steve Thorpe, Ed. Academic Press, 2002.

[16] 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.

[17] R. A. Russell, A. Bab-Hadiashar, R. L. Shepherd, and G. G. Wallace, "A Comparison of Reactive Chemotaxis Algorithms," *Robotics and Autonomous Systems*, vol. 45, pp. 83 – 97, 2003.

[18] H. Ishida, T. Yamanaka, N. Kushida, T. Nakamoto, and T. Moriizumi, "Study of Real-Time Visualization of Gas/Odor Flow Images Using Gas Sensor Array," *Sensors and Actuators B*, no. 65, pp. 14–16, 2000.

[19] 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.

[20] J. Atema, "Eddy Chemotaxis and Odor Landscapes: Exploration of Nature With Animal Sensors," *Biological Bulletin*, vol. 191, no. 1, pp. 129–138, 1996.

[21] H. Ishida, T. Nakamoto, and T. Moriizumi, "Remote Sensing of Gas/Odor Source Location and Concentration Distribution Using Mobile System," *Sensors and Actuators B*, no. 49, pp. 52–57, 1998.

[22] G. Kowadlo and R. A. Russell, "Naive Physics for Effective Odour Localisation," in *Proceedings of the Australian Conference on Robotics and Automation*, 2003.

[23] A. Lilienthal and T. Duckett, "Gas Source Localisation by Constructing Concentration Gridmaps with a Mobile Robot," in *Proceedings of the European Conference on Mobile Robots (ECMR 2003)*, Warsawa, Poland, 2003, pp. 159 – 164.

[24] 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.

[25] 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.

[26] A. Lilienthal, H. Ulmer, H. Fröhlich, A. Stützle, F. Werner, and A. Zell, "Gas Source Declaration with a Mobile Robot," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2004)*, 2004, pp. 1430 – 1435.

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