SmokeBot

Mobile Robots with Novel Environmental Sensors for Inspection of Disaster Sites with Low Visibility

Project start: January 1, 2015

Duration: 3.5 years

Deliverable 2.6

Gas Detection and Identification

Due date: month 16 (April 2016)

Lead beneficiary: ORU

Dissemination Level: Public

Main Authors:

Victor Hernandez Bennetts (ORU)

Han Fan (ORU)

Version History:

0.1: initial version, vh, April, 2016

0.2: modified version, vh, May, 2016

Contents

A Introduction and purpose of this document

The SmokeBot project will develop a new tele-operated robotic platform for supporting fire brigades for search and rescue missions with a special focus on operation environments with limited visibility. In disaster scenarios, where SmokeBot operates, the presence of hazardous chemicals is a risk factor to consider and thus, it is critical to have sensing modalities that allow to detect, identify and quantify gaseous substances present in the environment. Gas identification can be carried out with laboratory or hand-held equipment based on spectroscopy, optics and analytical chemistry. However, these devices are expensive and often, their bulky size prevents them from being used in field inspections. In order to provide gas sensing capabilities to the SmokeBot platform, an array of high bandwidth (Task T2.1) sensors are used together with a set of pattern recognition algorithms. In complex disaster scenarios, such as the ones targeted by SmokeBot, it is expected that the gas sensors will be exposed to intermittent gas concentration levels, often of unknown type. Therefore SmokeBot develops a set of Gas Identification algorithms (T2.3) to be used in such uncontrolled environments and that can provide an estimate of the identity of the gases to which the sensors are being exposed.

This document describes the proposed gas identification system and the achieved results in the period from M1-M16 and as well provides an overview of future developments on this task. The rest of this report is structure as follows: Section B provides a general overview of the problem of gas identification in uncontrolled environments as well as an exploratory data analysis carried out to evaluate the feasibility of gas identification in uncontrolled environments. Section C presents the technical description and evaluation of a supervised gas discrimination algorithm, tailored according to the observations presented in Section B. Section D presents the technical description and evaluation of a non supervised approach developed for gas identification in open environments. Section E presents a summary and outlook towards future developments.

B General Overview

Gas identification in controlled environments such as laboratories is commonly carried out under constant environmental conditions (i.e. temperature, humidity, airflow) and measurements are acquired in sampling chambers in three phases [1]. First, the sensors are exposed to a reference gas (e.g. clean air) in order drive the sensors to a known state or a baseline response level. Then, a gas sample of constant concentration is transported towards the sensor array. When the sensors start interacting with the gas sample a transient response pattern is produced. After a few seconds to a few minutes, the sensors reach a steady response where typically data analysis is carried out.

However, having bulky three phase samplings systems on-board the SmokeBot platform would drastically reduce the payload capabilities and mobility of the overall system. Moreover in a disaster scenario, such as the ones targeted by SmokeBot, environmental conditions are constantly changing and thus, the use of three-phase sampling processes in not possible. Instead, gas identification has to be carried out with devices commonly referred to as Open Sampling Systems (OSS) [2]. In OSS, data samples are directly taken from the environment. This means that constant concentration values cannot be expected and instead, intermittent concentrations of unknown gas identities are recorded under changing airflow, humidity and temperature conditions. Figure 1(b) shows the response of an OSS mounted on a mobile robot that was commanded to explore an indoor location where a gas source was present. Notice that compared to Figure 1(a), the sensor response does not show a clear three phase profile and that a steady response is never reached.

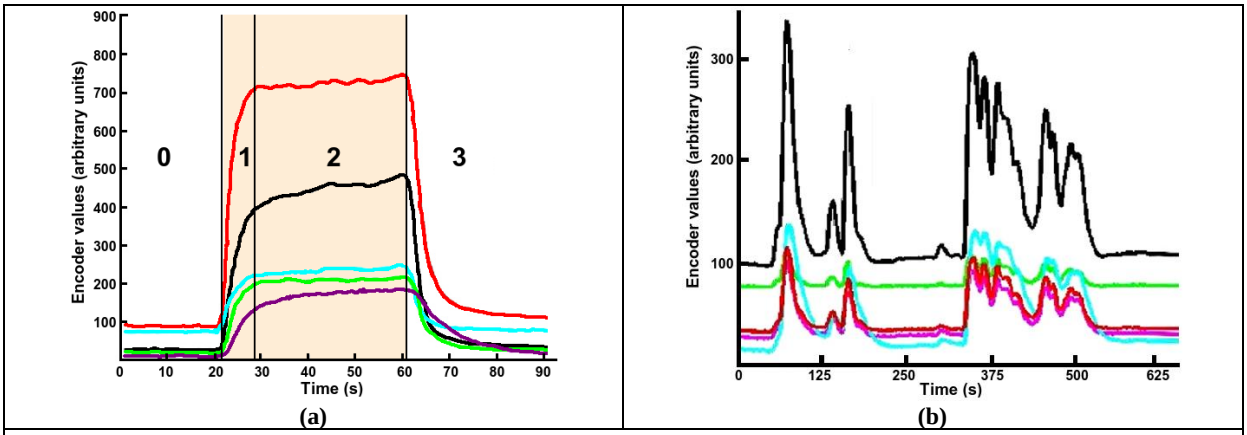

Figure 1: Response patterns acquired with arrays of metal oxide sensors using different sampling processes. (a) Three phase sampling process. The numbers on the figure indicate the different stages of the sensor response: 0 - Baseline response, 1 - transient response (rising edge), 2 - steady state, 3 - transient response (recovery edge). The shaded area denotes the time interval during which the sensors were exposed to the gas sample. (b) Response pattern acquired with an OSS mounted on a mobile robot. A controlled three phase sampling process is not possible due to the chaotic nature of the environment. Both images are adapted from [2].

B.1 Data Exploration

Considering the above mentioned challenges, we used datasets collected with mobile robots with OSS in order evaluate the feasibility of conducting gas identification in open environments. Two different datasets were used namely, a dataset corresponding to experiments conducted indoors [3] and a second dataset that corresponds to outdoor experiments [4].

The indoor experiments were conducted in a 5m x 5m x 2m closed room (Figure 2(a)) in which Ethanol and propanol were released in six separate trials at a constant rate of 0.2 l/min from a tube placed on the floor. With no artificial airflows (e.g fans), a weak circulating air flow field (0.01 – 0.03 m/s) was formed by natural convention. In this scenario, a mobile robot was programmed to move in a pre-defined trajectory to collect data with the OSS. Regarding the outdoor experiments, acetone and propanol sources were placed inside plastic containers and were allowed to evaporate. A bubbler (i.e. a pump) was used to facilitate the evaporation process. A couple of fans were placed near the gas source to spread the gas patches away (See Figure 2(b)). A mobile robot equipped with an OSS was used to collect data at random positions in the environment.

For both scenarios, two different types of experiments were carried namely single source and dual source experiments. In the first type of experiments, a single gas source was placed in the environment at different independent trials. For the second experiment type, two gas sources of different kind were placed at the same time in the environment separated at least by 50cm between each other.

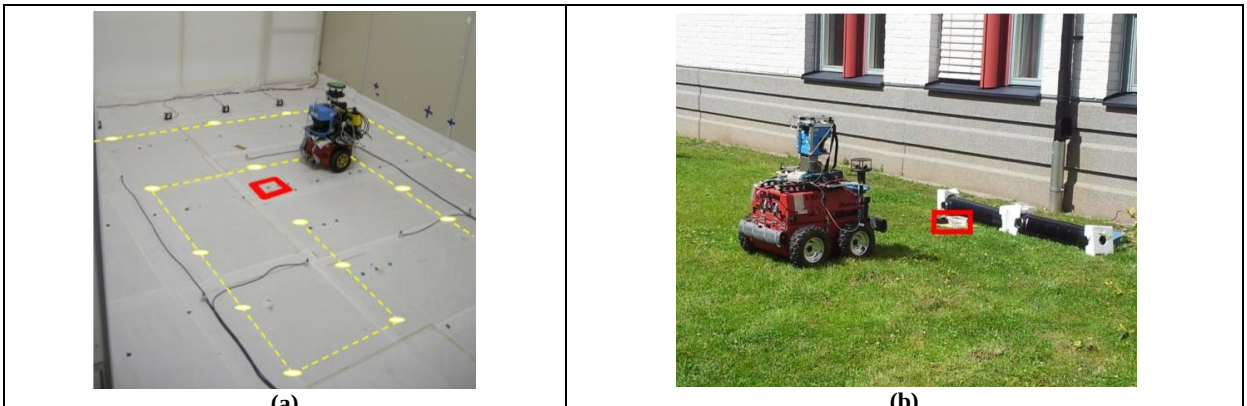

Figure 2: Experimental scenarios considered in the data exploration process. (a) Indoor environment [3]. (b) Outdoor environment [4]

From acquired data, a set of informative features (i.e. descriptors) can be extracted in order to identify the gases present in the environments. This process is commonly referred to as feature extraction. Feature extraction methods for OSS should consider that no steady response state is reached due to the dynamics of the environment and thus, descriptors have to be extracted from e.g. the rising/decaying edges of the sensor responses (see Figure 1(a)). Common methods for transient state feature extraction include the use of the Fourier transform, Wavelet transform, phase space descriptors, etc. The work of Trincavelli [2] is arguably the most comprehensive study related to gas identification with OSS.

We performed transient feature extraction over the data acquired in the indoor environment. Corresponding to the results of Trincavelli [2], we achieve a significantly improved classification success rate by performing gas discrimination only on the rising edges of the sensors response. Figures 3(a) and 3(b) illustrate the analysis we perform over the rising edges of the sensor responses. In this particular case, features were extracted using the Discrete Wavelet Transform approximation coefficients. Notice that in Figure 3(b), a sparse feature space is observed. This means that useful information, for example time intervals when the sensors are exposed to clean air, is discarded. While class overlapping occurs in a relatively small region in the feature space, it is difficult to determine at which point the overlap of the sensor response exactly ends.

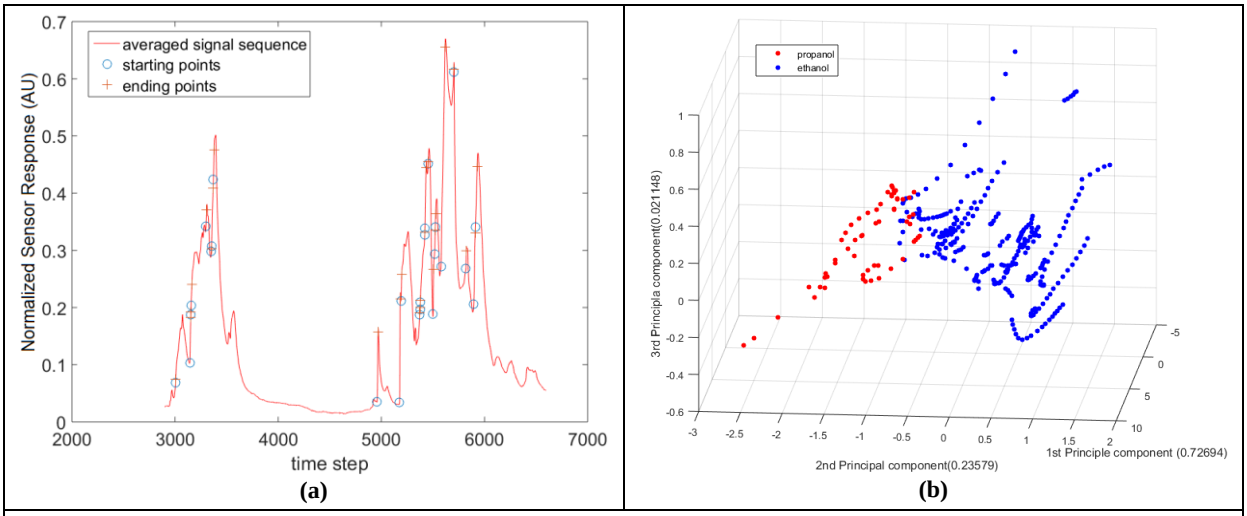

Figure 3: Feature selection using wavelet transform. (a) A segmented sensor response, where the start and the end of the rising edges are detected. (b) Feature space plot obtained from the selected response segments.

Alternatively, we used the instantaneous sensor responses as features as illustrated in Figures 4(a) and 4(b) for the indoor and outdoor locations respectively. In the figures, brighter color shades are assigned to higher concentration measurements, while low concentrations are plotted in gray tones. From the figures, the correlation between class separability and gas concentration can be seen. Gas identification at high concentration regions is relatively trivial due to the relative distance among the two classes. On the other hand, at lower concentrations, there is a clear overlapping between classes, which complicates gas discrimination.

In addition, it can be noticed that high concentration measurements are sparse. This means that most of the measurements belong to the low to the mid concentration regions in the feature space (see the histogram plots at the top right corners on the figures). A density based classification algorithm would tend to assign higher class posterior probabilities to measurement points that lie on densely populated regions, and low posterior probabilities for sparsely represented concentrations. Thus, in order to consider the particular characteristics of gas sensing in uncontrolled environments, it is required to incorporate gas concentration information into the algorithm.

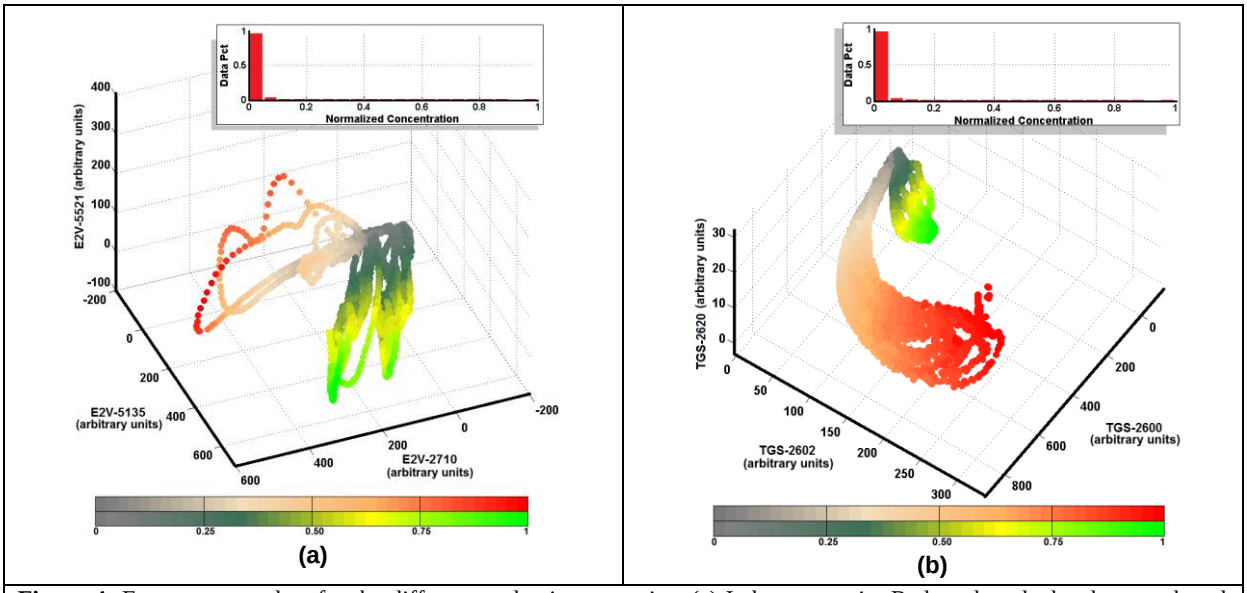

Figure 4: Feature space plots for the different evaluation scenarios. (a) Indoor scenario. Red marker shades denote ethanol samples and green marker shades denote propanol samples. (b) Outdoor scenario. Red marker shades denote acetone, green marker shades denote propanol. Most of the data lies in low concentration regions as shown in the histogram plots in both figures.

The above observations can be used to design a gas identification algorithm for uncontrolled environments. This means that the gas identification algorithm has to consider class separability with respect to the gas concentration and the misrepresentation of the high concentration measurements when making predictions of the gas identity. In the following section we describe an implementation and validation of a supervised gas identification algorithm that is tailored according to the conclusions reached in this section.

C A Supervised Approach for Gas Identification with OSS

As stated above, the goal of the algorithm is to incorporate the observed correlation between gas identification and concentration level. Figure 5 shows a block diagram of the proposed supervised algorithm for gas identification. The input to the algorithm corresponds to the response set acquired with a sensor array (i.e. "odour fingerprint") and the output corresponds to the posterior estimation of the identity of the acquired samples. Considering the specifications of the gas sensor box from task T2.1, it is assumed that only uncalibrated, partially selective gas sensors are used.

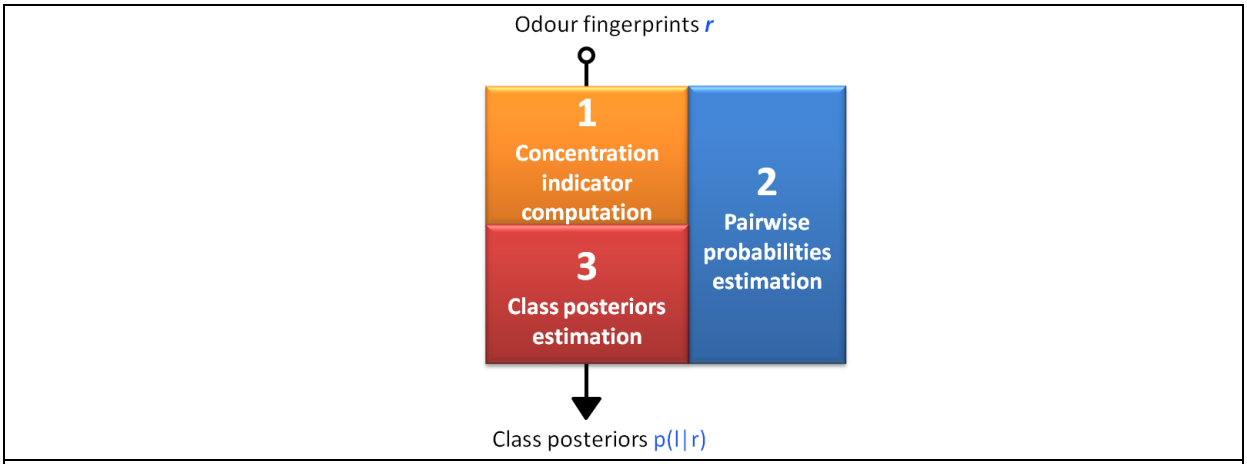

Figure 5: Block diagram of the proposed gas identification algorithm.

C.1 Concentration indicator computation

This stage computes a non calibrated concentration indicator using the instantaneous responses of the gas sensor array. The instantaneous sensor responses are good indicators of the concentration level since the sensor responses from e.g. a MOX sensor are linearly proportional to the logarithm of the gas concentration [5]. For an array of D sensors, we computed a concentration indicator I_c as follows:

$$I_c = \mathbf{max}_{r_i^{(j)} \in \mathbf{r_i}} \left\{ r_i^{(1)}, r_i^{(2)}, \cdots, r_i^{(D)} \right\} \qquad (1)$$

The above equation considers that the response matrix $\mathbf{r_i}$ is composed of the conductance readings of the sensor array. The conductance values in $\mathbf{r_i}$ increases according to the concentration level.

C.2 Pairwise Probability Estimation

In this stage, the pairwise probabilities between the target compounds and the pairwise probabilities between the target compounds and air are computed. For example, for a two class gas identification problem, the pairwise probabilities P_{1v2} and P_{2v1} are computed. In addition P_{1va} and P_{2va}, which represent the pairwise probabilities between target compounds and air are estimated. There is no specific requirement related to the selection of the classification method to compute P_{1v2} and P_{2v1} as long as $P_{1v2} = 1 - P_{2v1}$.

The concentration information is integrated in the gas identification problem by using I_c to compute the pairwise posteriors P_{1va} and P_{2va}. To illustrate the process of computing P_{kva} (where k is the gas identity), Figure 6 shows a plot of I_c vs P_{2v1}. Notice that there is a high uncertainty in the classification (given by fluctuations in the class probabilities) at low concentration values. It is thus desired to model P_{kva} in such a way that the uncertainty in the class probability estimation decreases proportionally to the concentration level.

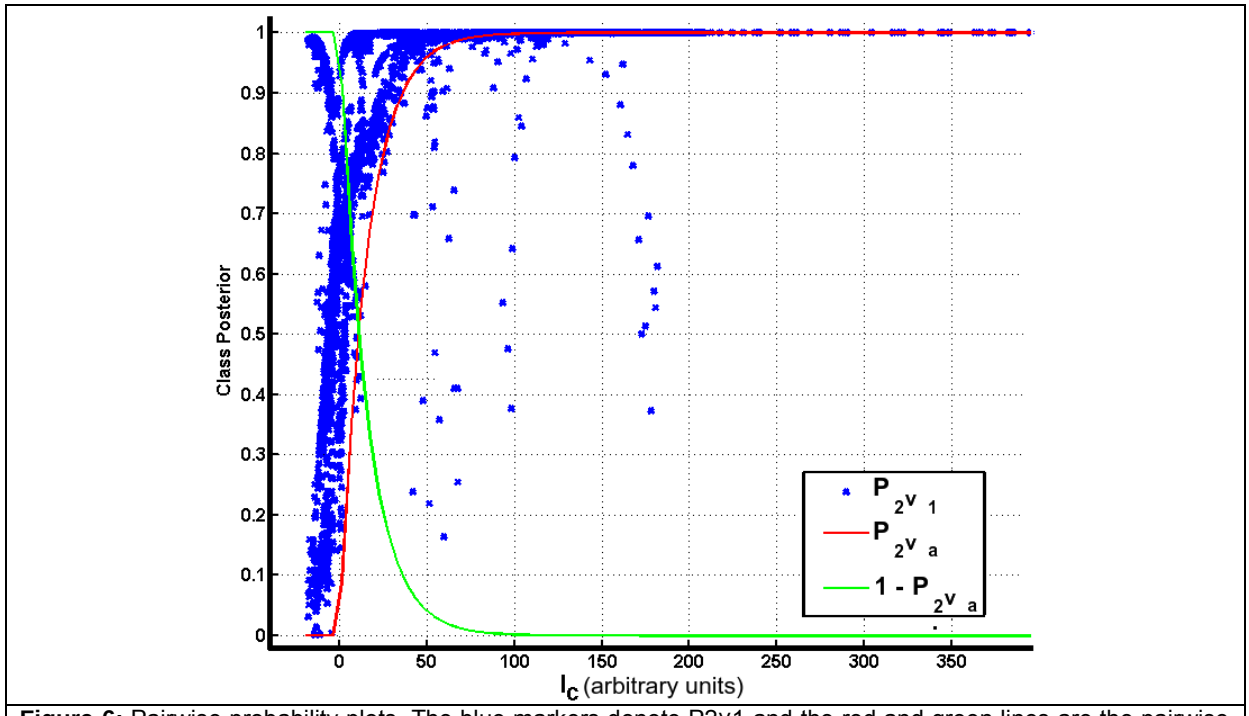

Figure 6: Pairwise probability plots. The blue markers denote P2v1 and the red and green lines are the pairwise probabilities P_{2va} and 1 - P_{2va} respectively.

Based on the above observations, the pairwise probabilities $P_{k\vee a}$ between the target compounds and air can be computed as follows:

$$P_{k\vee a}(I_c) = 1 - e^{-\beta_k I_c}\big|_{k=1,2} \qquad (2)$$

In the above equation, the functional parameter β_k, determines the rate of change in the class probability predictions. The functional parameters β_k can be individually learned from the data by dividing the training dataset according to their labels l_k and using the pairwise probabilities between the compounds (e.g. $P_{2\vee 1}$) as target variables.

C.3 Class Posterior Estimation

The final computation of the class posteriors p(l|r) is carried out by coupling the pairwise probability estimations. For the implemented gas identification system, we used the algorithm proposed by Hastie and Wu in [6] [7], which addresses the estimation of the posteriors as a minimization of the Kullback-Leibler (KL) divergence between the pairwise estimates and the true distributions.

C.4 Experimental Validation

Using the datasets described in the Section C we performed two different validation processes namely, quantitative and qualitative evaluations. The quantitative evaluation is carried out with data from the single source experiments where ground truth is available. As a metric for the performance of the classifier we use the success rate of correctly classified samples. Those measurements that are labeled as "air" are not considered in the quantitative evaluation. Moreover, we considered two different conventional classifiers namely Mixture of Gaussians Classifer (MoGC) and K Nearest Neighbours Classifier (KNN-C) to evaluate the parameter selection sensitivity for the classifiers used to compute the pairwise probabilities between the target compounds.

Figures 7(a) and 7(b) show the performance plots vs different parameters for both the indoor and outdoor environments using 5 fold cross validation. The performance criterion used for evaluation was the percentage of correctly classified points. It can be seen in the figures that the performance of the algorithm improves with fewer (for K-NNC) and larger numbers of Gaussians (for MoGC). While better performances can be achieved when increasing the number of Gaussians, the difference between the best and the worst performance of the trained classifiers is not substantial (2%outdoors and around 1% in the robot arena). Overall, the best performances were higher than 99% for the indoor experiment and higher than 97% for the experiments conducted outdoors.

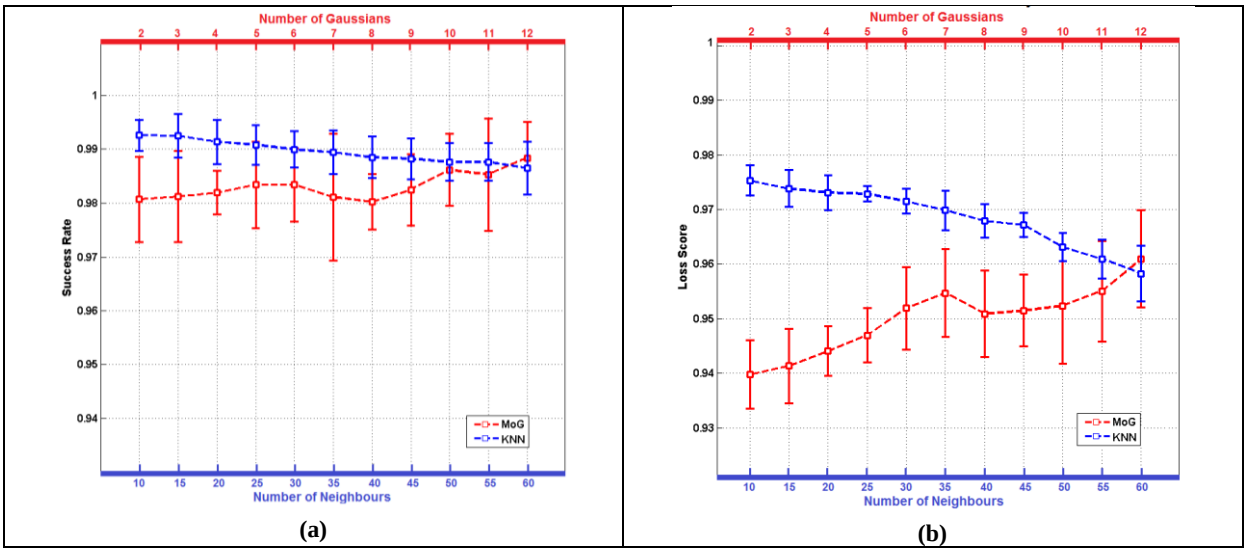

Figure 7: Performance plots for the proposed supervised classifier. (a) Indoor scenario. (b) Outdoor scenario.

Figures 8(a) and 8(b) show segments of the concentration indicator I_c coloured according to their predicted posteriors for indoors and outdoors respectively. The measurements are coloured according to the class posteriors computed with the proposed gas identification algorithm and square markers are used to highlight misclassified measurements. Notice that predictions are made with higher confidence for measurements were the concentration is higher while at lower concentrations, predictions are made with lower confidence. Classification errors occur close to the baseline response level. In addition, the proposed approach does not assign high confidences to the erroneous predictions. Misclassified predictions were made with an average confidence of 49% in the robot arena and 59% in the outdoor courtyard.

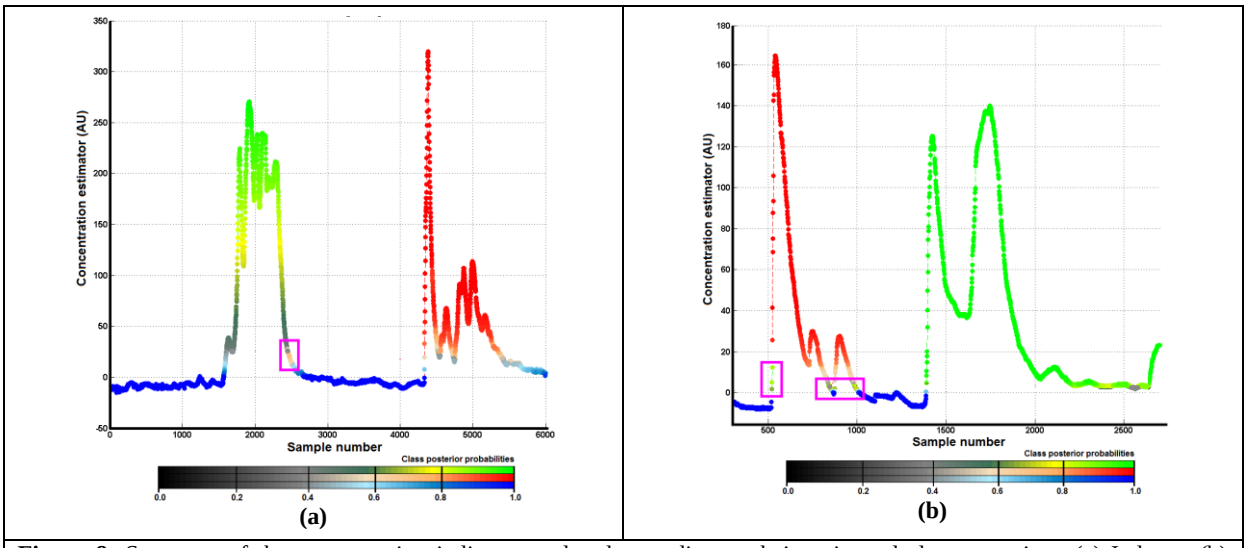

Figure 8: Segments of the concentration indicators colored according to their estimated class posteriors. (a) Indoors. (b) Outdoors. In both figures the magenta markers denote the location of misclassified measurements.

In addition to the single source experiments, further evaluation was carried out using the dual source experiments. Due to the lack of ground truth (i.e. class labels for each measurement), the evaluation was conducted only in a qualitative way. Figure 9 shows an example of such evaluation. The figure presents the spatial distribution of the measurements acquired with the robot in the indoor environment. The Z axis corresponds to the normalized concentration level. Square markers indicate the position of the gas sources (red propanol, green acetone). The classifier was trained by using the single source experiment as training data. Notice that the concentration level indicates that gas plumes moves from the gas sources to the upper left corner in the environment. In a similar way, the gas identification algorithm shows clean air (blue markers) near the bottom right corner and makes confident predictions in areas where the concentration is higher. It can be observed that the class predictions are correlated with the actual locations of the gas sources.

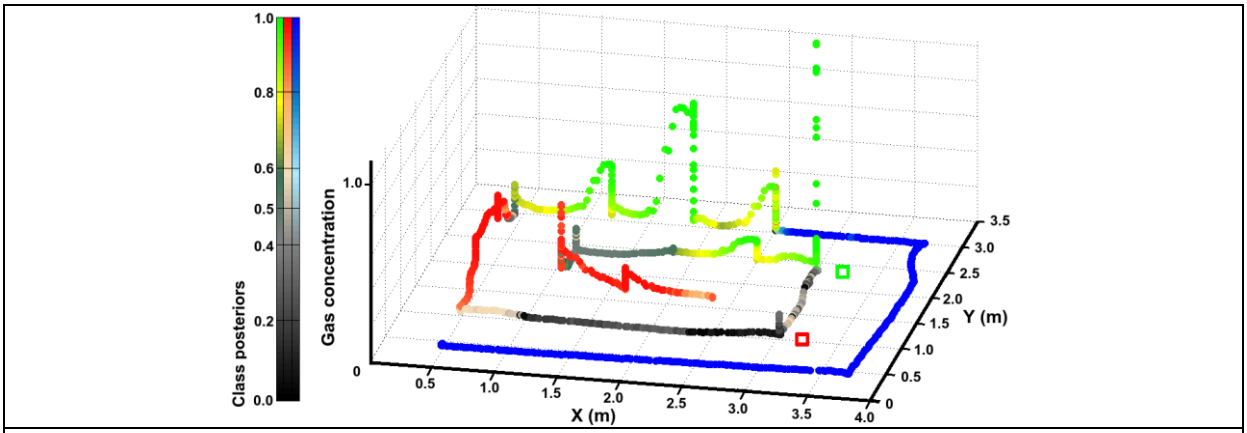

Figure 9: Spatial distribution of the predicted class posteriors in an indoor dual source experimental run. The color shades denote the estimated class posteriors.

D Towards unsupervised Gas Identification with OSS

A key challenge when training gas identification algorithms that use partially selective sensors (such as the ones developed in T2.1) is that ground truth datasets have to be acquired for all possible chemicals that might be present in the target application. Experiments are commonly conducted with only one or two known target compounds. This results in the gas identification system being able to only discriminate among the target compounds. Moreover, unknown chemicals are erroneously identified as one of the target compounds that were used as training datasets.

In SmokeBot's application scenario where a gas leak might have occurred, it can be expected that unknown, by-product gases might be present in the environment and thus, the gas identification system should be able to discriminate between unknown chemicals. Therefore, we presented an approach to perform gas identification with OSS using an unsupervised approach. The proposed gas identification approach is based on the clustering algorithm originally presented in [8] by Rodriguez et al. and is designed to address the particular characteristics of gas identification with OSS.

The algorithm presented by Rodriguez et al. [8] assumes that cluster centers can be identified by considering two indicators namely ρ_i and δ_i. The indicator ρ_i denotes the local density for measurement *i* in the feature space. The indicator δ_i, also computed in the feature space, finds the similarity value between measurement i and the most similar measurement j that has a higher local density than *i*. The equations to compute δ_i and ρ_i are as follows:

$$\rho_i = \sum_{i=1, i\neq j}^{N} e^{-\left(\frac{d_{ij}}{d_c}\right)^2} \quad (3)$$

$$\delta_i = \mathbf{min}_{i,:\rho_j>\rho_i}\left(d_{ij}\right) \quad (4)$$

In the above equations, the parameter d_{ij} corresponds to the similarity measurement between data points *i* and *j*. By considering the data distribution in the feature space, we opted for a cosine similarity measurement. From several typical feature spaces including various types of gas, we have an intuitive observation that the data points belonging to the same class may have large geometrical distances (e.g., Euclidean distance) between each other, however they do not differ so much when measured by the angles of the feature vectors, which implies cosine similarity is a more suitable metric for the case of gas discrimination. Given two measurement points I and J in a feature space of dimension D, the cosine similarity measurement d_{ij} is given as follows:

$$d_{ij} = \frac{\sum_{r=1}^{D} I_r J_r}{\sqrt{\sum_{r=1}^{D} I_r^2}\sqrt{\sum_{r=1}^{D} J_r^2}} \quad (5)$$

The parameter d_c is referred to as the kernel size and according to [8] it can be heuristically set by decreasingly sorting a vector Δ of all the pairwise similarity measurements d_{ij}. Then, d_c is set equal to the value of the second percentile of Δ.

In the current implementation, data clusters are determined by e.g. ranking the acquired measurements according to the indicator $\gamma_i = \delta_i\rho_i$ or by setting a decision boundary in a δ_i vs ρ_i plot in a systematic way. Once the cluster centers have been identified, the next step is to assign the remaining data points to the most similar cluster center. This is done by propagating the labels from the cluster centers towards the data points in areas in the feature space that exhibits lower local densities. This means that a given data point *i* will get the same cluster label as the most similar data point that has a higher local density value.

D.1 Experimental Validation

The proposed unsupervised algorithm was evaluated with the datasets described in Section B. It was assumed that no previous information existed about the possible target compounds in the environment and therefore, the algorithm starts without previous knowledge on the gas identities and ideally, it should identify all the gaseous compounds present in the environment.

Under the above assumption that no information exist about target compounds, a quantitative evaluation was carried out with the single source experiments. Figure 10 shows the feature space plot for the indoor environment. Notice that both gases are clustered as only one class. This is due to the fact that the most likely cluster centers are detected close to the baseline level (e.g. close to concentration levels near to zero), where most of the data lies and where the class overlapping occurs.

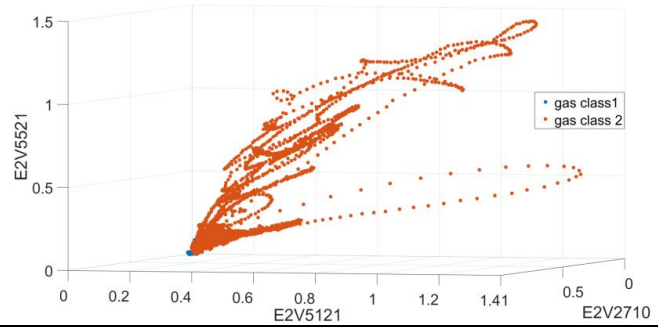

Figure 10: Experimental results using the unsupervised approach without considering a class for low concentration measurements (i.e. air). These plots were computed using the single source indoor experiments. Notice that the algorithm fails to identify the gaseous compounds. Instead, low concentration measurements, close to zero are identified as "class 1" class while the rest of the data is labeled as "class 2".

Thus, we defined a detection limit Λ, given by a percentage of the concentration indicator I_c (See Section C). A given measurement is considered as "air" if it falls below Λ. Figures 11 and 12 show the achieved results with Λ=0.0005% and Λ=0.4% for indoors and outdoors respectively. These thresholds were determined in a systematic way. Figures 11(a) and 12(a) correspond to the decision graphs, which show two highly likely cluster centers. Figures 11(b) and 12(b) correspond to the feature spaces. Notice that no class posterior is estimated and only a class label is provided by the algorithm. Similar to the case of the supervised algorithm, we defined the performance metric as the percentage of correctly classified measurements. With the above mentioned concentration thresholds, the achieved performance is approximately 94% for both scenarios.

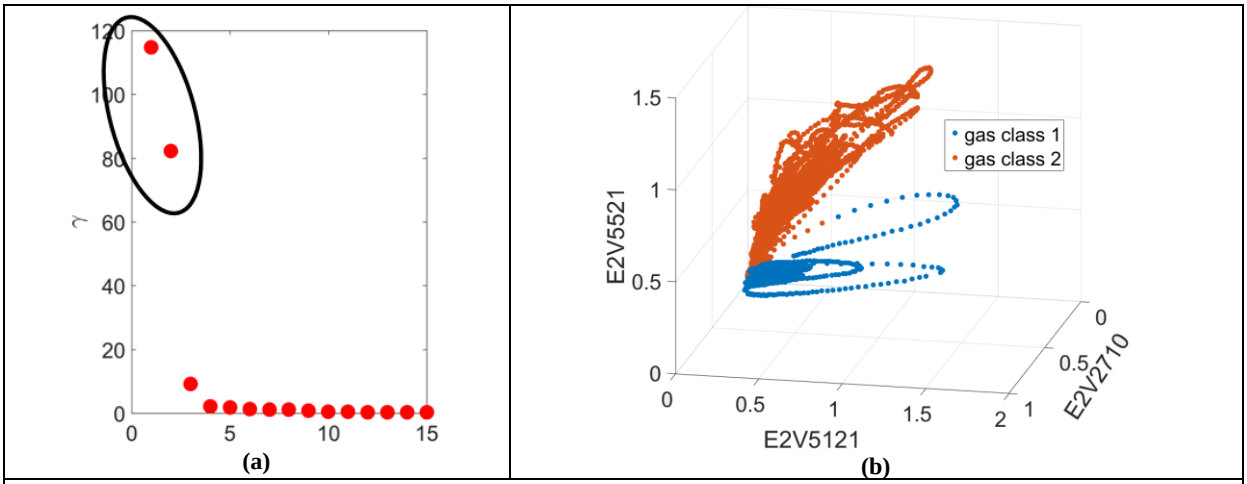

Figure 11: Results with the single source indoor experiments. The concentration threshold Λ was set to 0.00005%. (a) Decision graph (γ_i plot) with the two most likely cluster centers circled. (b) Feature space plot.

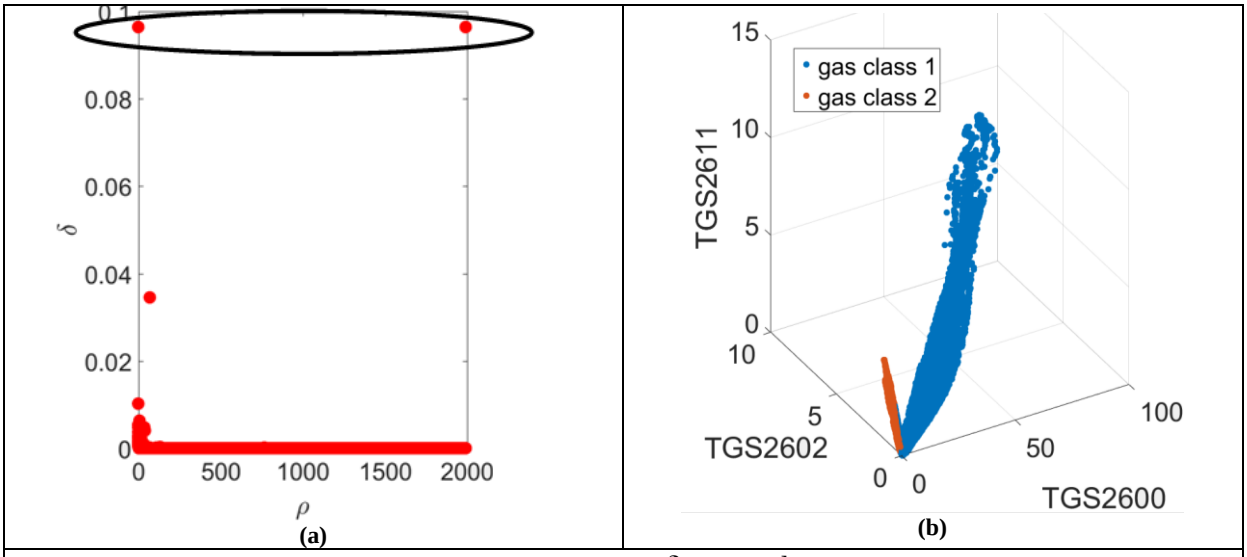

Figure 12: Results with the single source outdoor experiments (δ_i vs ρ_i plot) with the two most likely cluster centers circled. The concentration threshold Λ was set to 0.4%. (a) Decision graph. (b) Feature space.

Similarly to the case of the supervised algorithm, we conducted a qualitative evaluation using the dual source experiments. Figures 13(a) and 13(b) shows the spatial distribution of the measurements in the indoors and outdoors experiments respectively. Black colored markers denote the measurements that fell below Λ. Square markers indicate the actual position of the gas sources. Notice that the results are comparable to the case of the single source experiments, where the predicted labels correlate with the position of the two different gas sources.

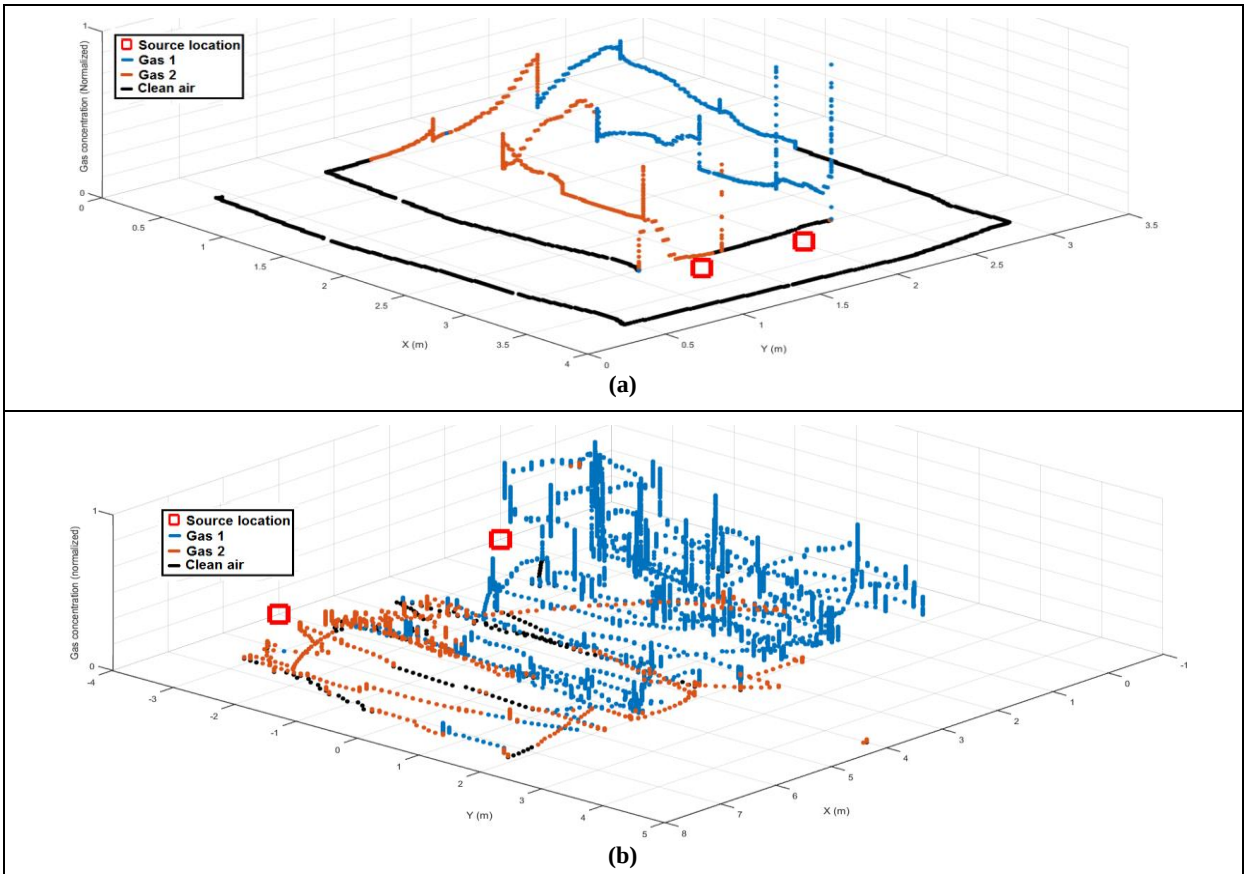

Figure 13: Spatial distribution of the predicted class labels in the dual source experiments. (a) Indoors. (b) outdoors.

E Summary and Outlook

In disaster scenarios, where SmokeBot operates, gas leaks and by-product chemical compounds might be present and therefore, the on-board sensing modalities should allow detecting and identifying the chemicals in the air. In SmokeBot, a set of high bandwidth partially selective sensors (task T2.1) and pattern recognition algorithms are used to perform gas identification. We presented two different approaches for gas identification namely, a supervised and an unsupervised approach. Regarding the supervised approach, consistent results were achieved in two different indoor and outdoor environments. The evaluation shows that gas identification can be conducted with high performance rates.

However, one of the key difficulties with supervised methods is the collection of training data. This is particularly difficult for Mobile Robotics Olfaction (MRO) related applications due to the time consuming process of preparing the experimental scenarios and the interference caused by environmental conditions. Moreover, for applications such as the one targeted by SmokeBot, not all the target chemicals are known.

Thus we presented a second gas identification approach based on unsupervised learning. We assume that no previous information is given about the target gases and the gas sensing system should be able to identify how many different gases might be present in the environment. The evaluation of our proposed unsupervised approach shows that the results are comparable to those achieved with the supervised method.

However, there are still open issues with the implementation that will be addressed in the upcoming months. First, the clustering algorithm will be extended in order to provide class posteriors instead of crisp labels. This is an important development for further gas sensing related tasks. For example, class posterior can be used in T.2.4 to create maps that indicate the likelihood of finding a given compound in the environment. An algorithm will be developed to automatically learn the concentration threshold that determines the gas identification capabilities of the system. In addition, the algorithm will be extended to consider already existing information about target gases (e.g. measurements with reference gases). This means that once a reference gas has been learned, the algorithm will identify those measurements that are highly unlikely to belong to the known compound.

Additional developments for this deliverable include dissemination efforts and the evaluation of the proposed algorithms with the high bandwidth gas sensing unit currently under development in T2.1. Tests with the gas sensing unit will be performed in cooperation with UWAR, once a working prototype is ready. The evaluation will be carried out with the sensing unit mounted on the robot and different testing conditions will be considered for example identifying specific target compounds, and detecting how many gas compounds are present in the environment.

References

[1] T.C. Pearce, S.S. Schiffman, H.T. Nagle, and J.W. Gardner. Handbook of machine olfaction. Wiley, Germany, 2003.

[2] Marco Trincavelli. Gas discrimination for mobile robots. PhD thesis, Örebro University, December 2010.

[3] Yuta Wada, Marco Trincavelli, Yuichiro Fukazawa, and Hiroshi Ishida. Collecting a database for studying gas distribution mapping and gas source localization with mobile robots. In Proceedings of the 5th International Conference on Advanced Mechatronic (ICAM), pages 183–188, 2010.

[4] Victor Pomareda, Victor Hernandez Bennetts, Ali Abdul Khaliq, Marco Trincavelli, Achim Lilienthal, and Santiago Marco. Chemical source localization in real environments integrating chemical concentrations in a probabilistic plume mapping approach. In Proceedings of the 15th International Symposium on Olfaction and Electronic Nose (ISOEN), 2013.

[5] K. Ihokura and J. Watson. The Stannic Oxide Gas SensorPrinciples and Applications. Taylor & Francis, 1994.

[6] Trevor Hastie and Robert Tibshirani. Classification by pairwise coupling. The Annals of Statistics, 26(2):451–471, 04 1998.

[7] Ting Fan Wu, Chih Jen Lin, and Ruby C. Weng. Probability estimates for multi-class classification by pairwise coupling. J. Mach. Learn. Res., 5:975–1005, December 2004.

[8] Alex Rodriguez and Alessandro Laio. Clustering by fast search and find of density peaks. Science 344, 1492 (2014)