

Gas Classification in Motion: An Experimental Analysis

Javier G. Monroy*, Javier Gonzalez-Jimenez

MAPIR-UMA Group, Dept. System Engineering and Automation. University of Málaga, 29071 Málaga, Spain

Abstract

This work deals with the problem of volatile chemical classification with an electronic nose (e-nose), and particularly focuses on the case where the e-nose is not collecting samples in a stationary fashion but is carried by a moving platform (mobile robot, car, bike, etc). We bring to light that, under these specific circumstances, substantial changes in the transient response of the gas sensors arise (something that has not been considered until now). We experimentally demonstrate that these changes in the sensor's response have an important impact on the classification accuracy if not properly considered, resulting in a decrease of up to 30% in some configurations. We back our conclusions with an extensive experimental evaluation consisting of a mobile robot inspecting a long indoor corridor with two chemical volatiles sources (ethanol and acetone) more than 240 times, at four different motion speeds. The paper also reveals the relevance of training the classifiers with data collected in motion, and proposes different training schemes suitable to this problem.

Keywords: Volatile Chemical Classification, Gas Recognition, Robotics Olfaction, Motion, Electronic Nose, Pollution Monitoring

*Corresponding author

Email address: jgmonroy@uma.es (Javier G. Monroy)

1. Introduction

The classification of volatile chemical substances is an important task for many applications including the detection and diagnosis in medicine, quality control in food processing chains, finding drugs and explosives, or the monitoring of pollution levels in air. Among the different systems available to carry out such classification (gas chromatography, mass spectrometry, dynamic olfactometry, etc.) in this work we deal with electronic noses. Electronic noses (e-nose) are sensing devices capable of classifying chemical volatiles according to the readings of an array of non-selective gas sensors and some pattern recognition algorithm. Given their high versatility to host multiple sensors while still being compact and lightweight, e-noses have demonstrated to be a promising technology to real-world gas recognition, which is our main concern in this work.

From the variety of potential applications that may benefit from gas classification, some of them require to measure the environment continuously and at different locations. They include, among others, city odor mapping [1, 2, 3], pollution monitoring [4, 5] or leak detection [6, 7]. Two alternatives exist for this type of task: to deploy a network of fixed gas sensing devices (with the inherent problem of scalability to large areas), or to employ few mobile e-noses, which can be transported by vehicles, such as public buses and cars, drones, mobile robots, bikes, etc. We are interested in the latter, mobile measurement systems, which trade off temporal coverage against spatial coverage, enabling a high spatial resolution across large areas without the need for a large number of fixed nodes.

This work focuses on that specific problem, the classification of volatile compounds by an electronic nose carried by a mobile platform, that is, performing the classification in motion. The specific challenges here stem from both the inherent difficulties of performing chemical classification in real environments, and the lack of a clear understanding of what are the effects of the movement of the olfactory device on the gas sensor response. More specifically, our interest lies in two questions: (i) What is the relationship (if any) between the motion speed of an e-nose which is

continuously sampling the environment and the classification accuracy?, and (ii) How must we train the classifier to get the best possible performance?, in other words, should the classifier be trained using data gathered also in motion, or can data collected in a static setup (more feasible) be used for training?.

These questions arise as consequence of a previous work [2], where an odor monitoring campaign was carried out employing an e-nose. The objective of such campaign was to detect and broadly classify the odors present in a city located at southern Spain, by carrying the e-nose both by a person and on a bike. By analyzing the readings recorded over a total of three days, we came up with a series of difficulties to perform the classification, basically related to different responses of the e-nose depending, apparently, on the way it was carried. These findings suggested that the changes in the gas sensor response would have a non negligible impact on the chemical classification, and this is what we investigate in this work.

Our main contribution is an experimental study from which meaningful conclusions related to the two questions raised above are drawn. We also review in Section 2 different works that employ e-noses in non-static configurations, and then revise (in Section 3) the challenges of in-motion chemical recognition in real environments, which is a challenging open issue. In Section 4 we describe the specific setup used to gather the experimental data (e-nose, robot, classifiers, etc.), and then we present the results and learning related to the two questions in Section 5 and Section 6, respectively. Finally, Section 7 summarizes the presented work and proposes lines for future research.

2. Related Work

This section is devoted to review works that have dealt with volatile chemical recognition with an e-nose that changes its location during the sampling process. Particularly, we focus on works that analyze how the classification is conditioned by the movement of the gas sensing device.

One of the first works dealing with the movement of the e-nose for gas recognition was presented by Trincavelli and coauthors [8]. They used a robot trans-

porting an e-nose and conducted a preliminary investigation on the most suitable path to optimize the classification accuracy taking into account the possible effect of environmental variables on the signals collected along that path. All the experiments they carried out were restricted to a single motion speed of $0.05m/s$, which, not only is very low, but it remains fixed, so no conclusion about this parameter was drawn.

Later, in [9] an experimental setup is presented to study the changes in the signal properties of the gas sensors under similar environmental conditions to those encountered when the sensing device is in motion. Specifically, the authors analyzed how the classification performance is affected by turbulence and by the switching of the gas class, reaching the conclusion that some type of memory about previous state of the sensors may benefit the overall recognition accuracy. Yet, this work does not account for different motion speeds of the e-nose, neither different wind velocities that could lead to results relevant to the problem addressed here.

An interesting research work is that presented by Vergara *et al.* [10], which does not employ an e-nose in motion to gather the gas samples, but reaches to some meaningful conclusions that can be extrapolated to our problem. The authors compiled an extensive dataset where nine static e-noses were placed within a wind tunnel to evaluate the performance of several sensor arrays working in open sampling settings. Different locations of the e-nose, heater voltages, wind speeds, and chemical volatiles compose the list of variables that make this dataset one of the most complete currently available. Particularly, two important conclusions were highlighted by the authors with regard to the wind speed (which may be seen as the counterpart of moving the sensing device): the classification performance was affected by the wind speed used during training, and that, in order to increase the robustness of the system against air flow variations, one may want to train the system at all the expected system conditions. Following this findings, in this work we investigate if a similar effect shows up when the sensing device is in motion, and further explore the impact of the training settings.

In [5] a gas discrimination approach is proposed

which accounts for the amplitude of the sensor's response, as an estimation of the gas concentration, to refine the classification. Two different scenarios (indoor and outdoor) are considered to validate the proposal, employing an e-nose carried by a mobile robot that continuously samples the environment. To train the classifiers, they performed different trials releasing one chemical at a time. As in previous cases, the influence of the motion speed is not analyzed. Yet, the specific configuration employed (maintaining a similar motion speed during all the experiments), is, as we will show later, a favorable way to perform the classification.

Finally, related to environmental monitoring applications, we can find works where a gas sensing device is carried by a person [11], a bike [12], public transport vehicles [13] or even drones [14], while sensing the air quality. Despite sampling the environment in motion, they do not carry out a classification phase to discriminate the type of gas, but rather employ an array of gas sensors with disjoint selectivity (i.e. one sensor for each analyte to monitor, and usually discarding the cross-selectivity among classes). An interesting remaining question in these works is whether the concentration measurements of the different pollutants can also be improved by taking into account the motion speed of the sensing device.

3. The Challenge of in-Motion Chemical Recognition under Real Environmental Conditions

The recognition of volatile compounds in real, uncontrolled environments brings up complications beyond those encountered in the classification of chemical substances under well-controlled laboratory conditions. The two main causes of such additional difficulties are:

- **Absence of ground truth:** In real environments, both indoor and outdoor, the dispersion of gases is dominated by turbulent flows. A turbulent flow is that in which fluid particles move in a random and chaotic way within the flow field [15]. This entails a great difficulty to

track the dispersion of an odor release, and consequently, in most cases, the absence of ground-truth labels (i.e. unawareness of the nature of the chemical volatiles). Unlike in a standard classification scenario where a dataset with ground-truth labels is available and consequently divided into training, validation and testing samples, in real applications we can not be one-hundred percent sure about the identity of the gases that the e-nose is being exposed to at each time instant (particularly when multiple gases may be present). This undoubtedly represents one of the main drawbacks for carrying real experimentation, and usually entails the consideration of semi-controlled scenarios (as the one presented in this work) from which ground-truth labels can be reasonable assumed.

- **Sensor dynamics:** The fact that most gas sensors present slow responses and recovery times is a well-known and documented limitation (see [16, 17]). For discrimination purposes, this drawback has been usually palliated by allowing the sensors to reach steady state values before collecting the samples (by means of measuring chambers). However, as already reported in [18], the transient phase of the signal encloses valuable information for discrimination tasks. Furthermore, since e-noses are typically composed of non-selective gas sensors with different dynamics, the acquired signals (multivariate time series) present different dynamic regimes that consequently increase the already complex process of signal segmentation [19].

Besides the aforementioned problems in real scenarios, for the particular case where the sensing device is sampling the environment in motion, we need to additionally consider a series of changes in the dynamic properties of the collected signals derived from that motion. For example, by increasing the motion speed we also raise the chance of hitting multiple gas patches in a short period of time, which can lead to an accumulation of the responses (if the gas patches are not separated by a sufficiently large distance), and thus even to the possibility of mixing patches of

different gas classes, which, naturally, will be harder to differentiate.

Similarly, another effect that is commonly noticeable when increasing the motion speed of the gas sensing device is a decrease in the response amplitude of the gas sensors. This effect is observable, for example, when analyzing the readings of the monitoring campaign presented in [2], and illustrated in Fig. 1. An e-nose, carried both by a person and on a bike, is used to inspect a fixed path in a city center along three consecutive days. As can be seen, the sensors' response amplitudes are much lower when using the bike than on foot. This effect is due to the shorter exposure time to the volatile compounds (presumably in the form of gas patches), and the consequent lack of time to reach the steady state corresponding to the actual gas concentration. Recall that the temporal response of a MOX sensor to a gas concentration pulse can be reasonably approximated by two first-order models [20, 16], which are characterized by rise and recovery periods as well as a steady state value

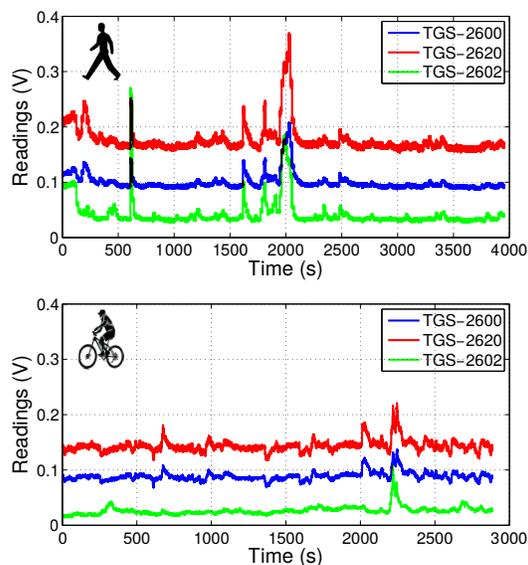


Figure 1: Illustration of the different signal amplitudes when varying the motion speed of the sensing device. The same e-nose, comprising 3 MOX sensors, gathers data from the same path: (top) when carried by a person on foot, and (bottom) when carried on a bicycle.

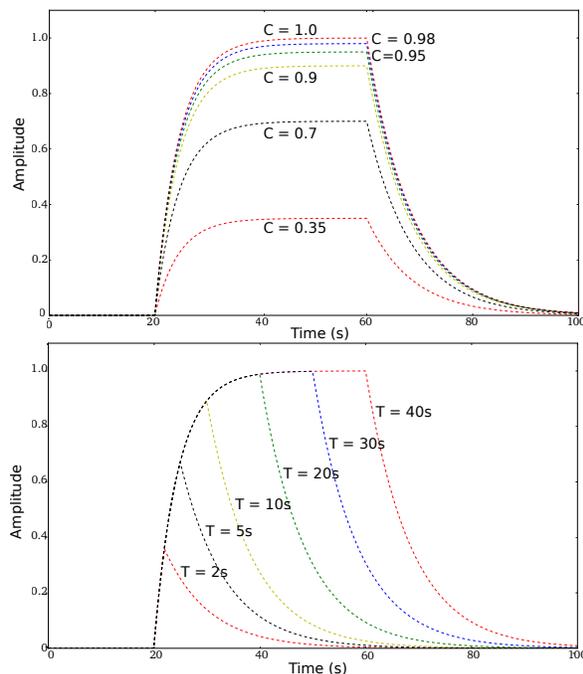


Figure 2: Simulation of the response of a MOX gas sensor to: (left) pulses of different concentration but identical exposure time ($T=40s$), and (right) pulses with identical concentration ($C=1.0$) but different exposure times. It can be seen that although the final amplitude of the sensor response is similar in both cases, the dynamic response is completely different.

which depends on both the sensor dynamics and the actual chemical concentration. At first glance, this decrease in the response amplitude may seem similar to that of exposing a still sensor to a lower chemical concentration. However, when sensing in motion, besides this decrease in the response amplitude, the transient response is also affected. Fig. 2 illustrates the differences in the sensors' response between the two cases, using the aforementioned model.

All these changes in the gas sensor response suggest that chemical classification is likely to be affected by the motion of the sensing device, and this is what we empirically investigate next.

4. Experimental Setup

The dynamics of gases in uncontrolled environments are characterized by high Reynolds numbers,

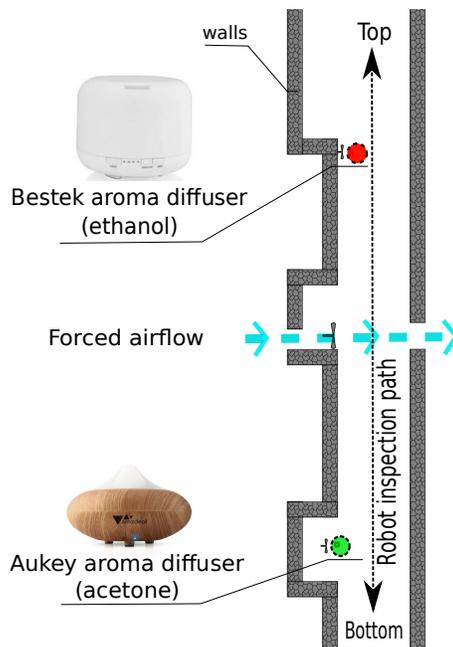


Figure 3: Schematic of the experimental setup (long indoor corridor), and pictures of the two gas sources. A barrier between the two gases at the middle of the corridor is created by a forced airflow generated by two fans.

which implies turbulent airflows and the chaotic dispersal of chemical volatiles. Artificial olfaction in such scenarios has to cope with a considerable number of environmental variables which are difficult (almost impossible) to monitor and control. A way to get around this issue is to gather data statistically representative of the phenomena under study.

To this aim, we have designed an experimental setup which allows us to repeat the trials under similar environmental conditions. Concretely, the setup consists of a long indoor corridor (30m) where two different gas sources, at fixed positions, are continuously releasing volatiles by means of two ultrasonic diffusers. The first source is composed of an "Aukey aroma diffuser" filled with an acetone dilution in water at 40% concentration, while the second source is based on a "Bestek aroma diffuser" filled with an ethanol dilution in water at a 35% concentration. A schematic of the setup and pictures of the ultrasonic diffusers are shown in Fig. 3.

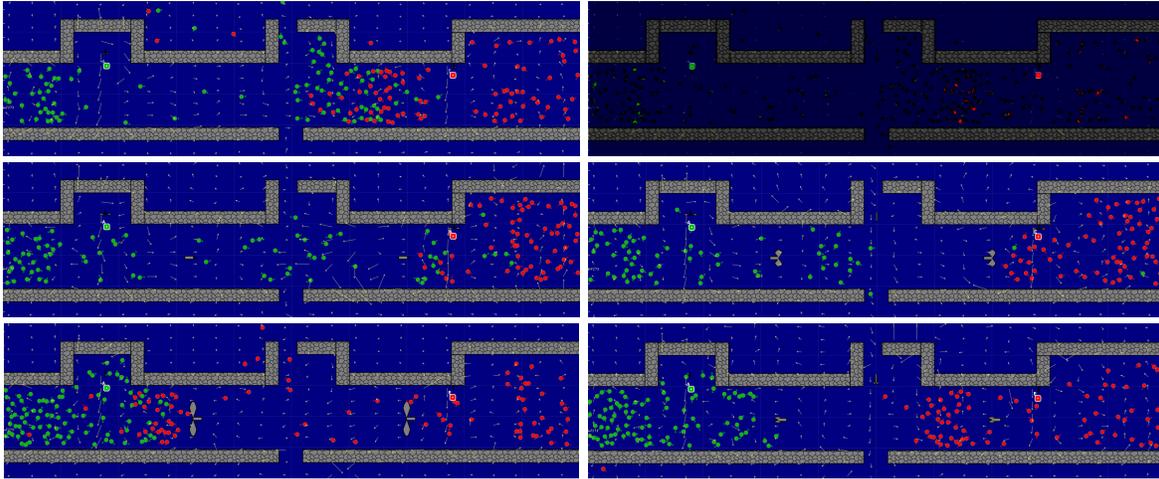


Figure 4: Wind flow and particle dispersion simulation of two gas sources in our experimental setup for different environmental conditions. Left column images display the uncontrolled scenario case, while right column images plot similar environmental conditions for the case of forcing a perpendicular airflow at the middle of the corridor to reduce the mixture of gases: (a,d) no mainstream flow, (b,e) mainstream flow towards ethanol source, and (c,f) mainstream flow towards acetone source (indicated with an arrow). It can be noticed that the forced airflow at the middle of the corridor reduces considerably the mixture of gases.

A mobile platform carrying an e-nose is then commanded to traverse the corridor back and forth at four nominal speeds: *low* $\approx 0.2m/s$, *medium* $\approx 0.4m/s$, *high* $\approx 0.5m/s$ and *very-high* $\approx 0.6m/s$. These motion speeds are only guaranteed when passing close to the gas sources but not at the end of the corridor, where the robot has to stop and turn around. Consequently, for the experimental evaluation we will discard samples taken at the corridor ends, since they are not of interest here. In order to cancel out random phenomena, at each trial the robot inspects the corridor 20 times (i.e. 10 round-trip runs along the corridor), summing up a total distance of about 600m. Furthermore, to keep the influence of environmental changes low, for each motion speed we repeated the trials three times, with a period of at least one hour between different trials, opening the windows during this interval to allow a rapid cleaning of the air in the corridor. The sequence in which the 12 trials (4 speeds x 3 repetitions) are carried out was set randomly, avoiding trials of the same speed to be tested consecutively. Thus, we gathered a rich dataset composed of four motion speeds, with three

inspection trials each, and 20 iterations each trial, having a total of 240 inspections of the corridor.

In this work we do not consider gas mixtures, only transitions from one volatile to another. Thus, although there is a considerable distance between both gas sources, we have introduced a forced airflow as a barrier to reduce the mixture of gases in the corridor. This airflow is generated by two fans, set at different heights and perpendicular to the robot path at the middle of the corridor and in front of a doorway, so that they drive away the volatiles present in this area to the outdoor while introducing clean air from the room behind. To better illustrate this, Fig. 4 plots a 2D simulation of the wind flows and the distribution of particles generated by both gas sources for different environmental conditions (see [21]). As can be noticed, for the uncontrolled scenario (left column images), the mixture of gases is always present, changing its location according to the environmental conditions. However, by considering the additional perpendicular airflow, the mixture of gases is, though not completely avoided, highly reduced. This configuration allows us to have a ground-truth (GT) of the

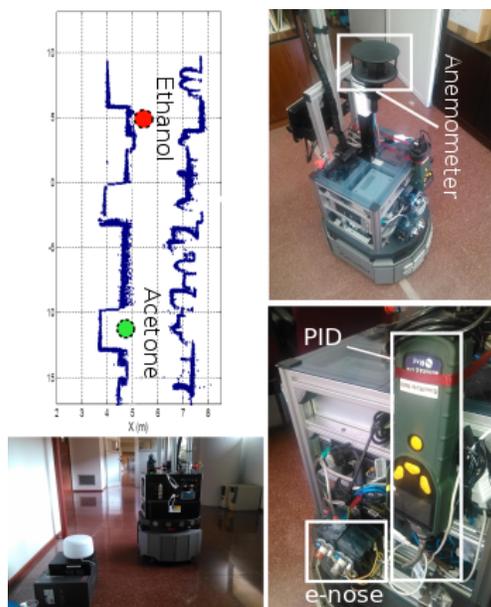


Figure 5: Points-map of the environment where the experiments have been carried out (refer to Fig. 3), and pictures of the mobile robot and its on-board sensors.

class labels based on location, that is, for each sample taken at point (x,y) the GT label is defined as that of the closest gas source. Additionally, we obviate samples taken at the central area, thus preventing to a great extent using gas mixtures by the recognition system.

4.1. The Robot and the E-nose

All the experiments in this work have been conducted with a mobile robot carrying the e-nose. The reason to employ a mobile robot instead of any other mobile platform is because it offers the possibility to automate the data collection process, to program in advance the sweeping strategy, as well as the easy control of the motion speed while taking advantage of high level operations such as obstacle avoidance or self-localization. Concretely, we have employed Rhodon, a lab-robot built upon a commercial platform, on which different devices are integrated. Apart from the e-nose which will be described next, Rhodon is equipped with two laser scanners: SICK PLS and Hokuyo URG (used here for localization

and obstacle avoidance), a Gill WindSonic ultrasonic anemometer for measuring the wind flows along the corridor, and a miniRAE Lite photo Ionization detector (PID) as an alternative gas detector. Fig. 5 shows pictures of the mobile robot and its on-board sensors, as well as the points-map of the experimental scenario as generated by an ICP-based method fed with the readings of the on-board 2D laser scanners. This points-map is used for the robot localization along the experiments.

For the detection of the chemical volatiles, Rhodon has been equipped with an easily configurable, modular e-nose [22] composed of an array of 10 MOX gas sensors, which provides gas readings at a rate of 5Hz. Table 1 lists the sensors as well as their target gases and detection ranges. As can be noticed, the sensors have different and overlapping selectivity and detection ranges, since in real scenarios there is little knowledge about the gases the e-nose is going to be exposed to. It is worth mentioning that most sensors provides sensitivity to alcohols (e.g. ethanol) and, though not usually mentioned in their specification sheets, they are responsive also to acetone in a lower degree. Fig. 6 shows a detailed picture of the e-nose as mounted on the mobile robot.



Figure 6: Picture of the e-nose used in the experiments for the detection (and posterior classification) of acetone and ethanol. The e-nose is composed of an array of 10 MOX gas sensors with overlapping selectivity and detection range.

Manufacturer	Sensor Model	Target gases	Detection Range
Figaro	TGS 2600(1)	Hydrogen, Ethanol, CO	1 - 30 ppm
Figaro	TGS 2600(2)	Hydrogen, Ethanol, CO	1 - 30 ppm
Figaro	TGS 2602	Toluene, H_2S , Ethanol, Ammonia	1 - 30 ppm
Figaro	TGS 2620(1)	Ethanol, Hydrogen, Iso-butane	50 - 5000 ppm
Figaro	TGS 2620(2)	Ethanol, Hydrogen, Iso-butane	50 - 5000 ppm
Figaro	TGS 2611	Methane, Iso-butane, Ethanol	500 - 10000 ppm
Hanwei	MQ2	LPG, Propane, Methane, Ethanol	100 - 2000 ppm
Hanwei	MQ3	Ethanol, Benzine (Breathalyzer)	0.05 - 10 ppm
Hanwei	MQ5	LPG, Natural Gas (not ethanol)	200 - 10000 ppm
Hanwei	MQ9	Methane, Propane, CO, LPG.	100 - 10000 ppm

Table 1: List of gas sensors mounted on the e-nose, their target gases and range of detection. When more than one detection range is specified by the manufacturer, here we show the one related to alcohols.

4.2. Classifier and Feature Extraction

Although not directly related to the scope of this work, the selection of a suitable classifier is an important factor when evaluating the classification accuracy, but this is not our focus here. Instead, we are interested in, given a certain classifier, analyzing the changes in its performance for different motion speeds. Thus, in this work we have employed a batch of nine different classifiers, namely: "K Nearest Neighbors", "Linear SVM", "RBF SVM", "Decision Tree", "Random Forest", "AdaBoost", "Naive Bayes", "LDA" and "QDA" from the *scikit-learn* machine learning python library. For the sake of clarity, in this article we restrict the results to the two classifiers which have delivered the best overall performance: RBF SVM and Naive Bayes.

Likewise, the features extracted from the e-nose signals, which are used to feed the classifier, though crucial to achieve a good classification rate, are not the main concern in this work. Thus, we have extracted a number of the most common ones (see Table 2) and, to reduce this large (correlated) set, we have applied PCA over the full feature set to retain the two first components.

4.3. The Collected Dataset

While sweeping the experiment area, the robot records the following data: accurate localization, linear and angular motion speeds, wind measurements, and chemical responses from both, the e-nose and

id	Feature description
0-9	Instantaneous response
10	Average (among sensors) of the instantaneous response
11	STD (among sensors) of the instantaneous response
12-21	Temporal average for a 1s sliding window
22	Average (among sensors) of the temporal average
23	STD (among sensors) of the temporal average

Table 2: List of features extracted from the response of the 10-MOX gas sensors array based e-nose.

the PID. Figures 7 and 8 plot the readings of the on-board sensors for a run at low motion speed, with respect the sensing location and time, respectively. As can be observed, the chemical substances are detected by both, the e-nose and the PID, and their maximum readings are obtained close to the real source's locations. As expected, the long recovery of MOX gas sensors "invents" high gas concentration values at positions far from the sources, which may wrongly suggest that the gases are heavily distributed in the corridor. A more precise estimation of the gas distribution is provided by the PID, which being a much faster sensor than the e-nose does not suffer for such long recovery periods.

To avoid considering samples where the motion speed is not the predefined, or where some mixture of gases may exist, we restrict the collection of samples to the shaded areas displayed in Fig. 7(b) and Fig. 7(c), which corresponds to locations less than 2m

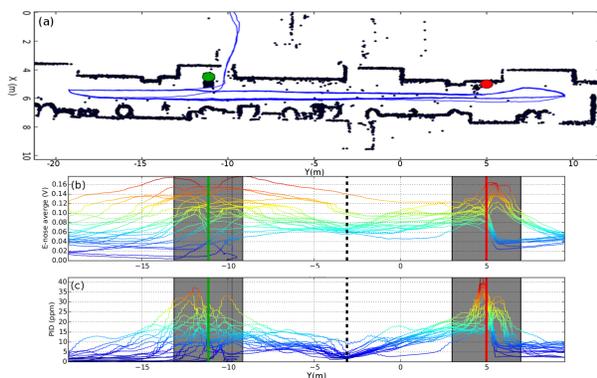


Figure 7: (a) Points-map of the experimental area (black points) and path followed by the robot during a chemical inspection at low speed (blue-solid line). (b) Average response of the e-nose (among all sensors) vs y-axis location, and (c) PID response in real concentration units (ppm) vs y-axis location. The positions of the gas sources are marked as vertical green/red solid lines (acetone/ethanol, respectively), while the location of the artificial airflow is set as a dashed black line. The shaded regions determine the areas where samples are collected.

apart from the gas sources. Furthermore, to ensure a correct labeling of the samples collected at these areas, we set an additional filter to discard samples where the PID response is below a predefined threshold ($PID_{th} = 5ppm$). To illustrate the need of such filter, Fig. 9 depicts two different scenarios where the location-based labeling fails to provide the correct labels. By applying the PID filter, we can get rid of these wrong labeled samples.

Related to the perpendicular airflow which was set to reduce the mixture of gases, it can be noticed how the response of the PID is minimum at this area (because of the clean air injected), and how this artificial airflow is several times stronger than the mainstream in the corridor (see wind measurements shown in Fig. 8). This justifies our assumption of not having mixture of gases among the collected samples.

Moreover, with the aim to better understand the results presented in the following sections, we now analyze the collected data, highlighting three important parameters which are key to reach trustable conclusions, namely: the average response of the e-nose, the mean gas concentration measured by the PID, and

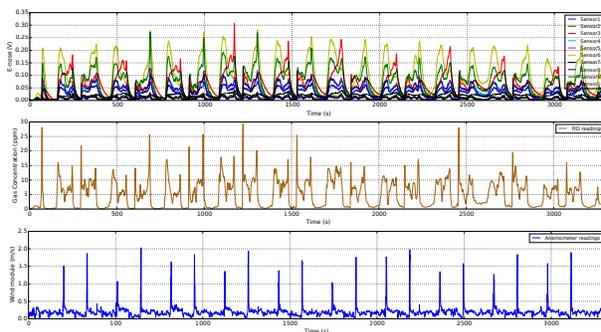


Figure 8: E-nose, PID and anemometer readings over time during a chemical inspection at low speed. The different inspections of the corridor (20 in total) are clearly distinguishable by the successive peaks in the response of both the e-nose and the PID. Also, it can be seen how the artificial airflow set at the middle of the corridor (see Section 4) is correctly detected by the on-board anemometer (strong peaks).

the average wind velocity detected by the anemometer. Fig. 10 displays these parameters for the different motion speeds and trials that compose the full dataset. Likewise, since we will also consider the case of training the classifiers with data collected in a static configuration, we include as well the corresponding values for this scenario to allow a proper comparison.

The data for the static configuration was collected by placing the e-nose at a fixed distance to the gas source, and employing a small fan to disperse the

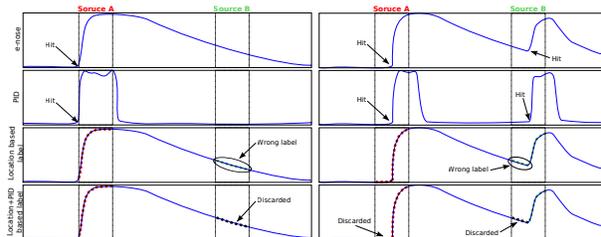


Figure 9: Illustration of the labeling problem due to the long recovery of MOX gas sensors. (left) Scenario where only one gas is detected, (right) scenario where both gases are detected. In both cases, red-dots represent samples labeled as gas A, green-dots are samples labelled as gas B, while black-dots are samples discarded by the PID filter. As can be seen, the location-based labeling incurs sometimes in errors which are overcame by doing use of the PID filter.

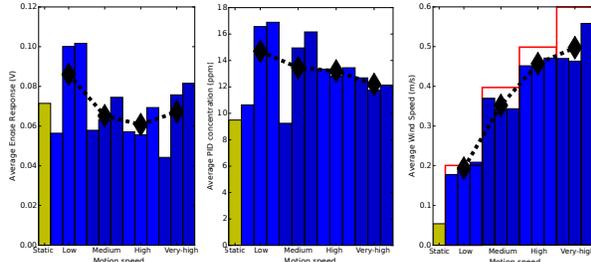


Figure 10: (Bar-plots) Average sensor’s responses for the static and in motion configurations: (left) average e-nose response, (center) average PID sensed concentration, and (right) average wind velocities together with the nominal motion speeds (red solid line). (Line-plots) Averages over the three trials that compose each motion speed. These values are calculated using only the selected samples (PID readings over 5ppm, and collected at less than 2m from the gas sources).

gases towards the sensing device. The objective of this simple setup is to collect data which account for the natural turbulence phenomena of gas dispersion, similar to the case of a wall-mounted e-nose. The distance between the gas source and the e-nose was set empirically to keep the average response of the e-nose similar to that of the motion experiments (approximately 1.2 m). This is fundamental to discard changes in the classification accuracy due to different e-nose amplitudes during the training and test phases.

From the average measurements of the PID we can see that the gas concentration is almost constant along the different experiments and motion speeds ([12 – 15]ppm). Yet, it must be stressed that it is practically impossible to maintain a constant and stable gas concentration along the different trials, mainly because of the turbulent and chaotic mechanisms of gas dispersion. Nonetheless, a small decrease is noticed when increasing the motion speed, possibly due to the turbulence generated around the robot.

Similarly, the average e-nose response is also affected when comparing different velocities ([0.06 – 0.09]V). In general, a decrease is noticed when the motion speed grows, something that goes in accordance to the effect illustrated in Fig. 1. An exception to this trend is the case of "Very-high" speed, where

two out of the three trials present average concentrations higher than those obtained with slower velocities. The exact cause of such change of tendency cannot be precisely determined with the current setup, yet we suspect that it is related to the long recovery of the e-nose and the accumulation in the gas sensor response when re-visiting a gas source in a short time-span. There are two main reasons that lead us to that belief: first, the PID readings for these experiments do not show this contradictory increase, which suggests that this effect is something inherent to the e-nose, and second, by increasing the motion speed we shorten the time between exposures to the gas sources, which very likely provokes accumulation in the gas sensor response as already described in Section 3.

Another interesting point that should be mentioned is the fact that the static data was collected so that the average e-nose response falls within the range of e-nose values during the motion experiments, but, as seen on the PID readings shown in Fig. 10, this corresponds to employing lower gas concentrations than in the motion trials. This is related to the response and recovery times of gas sensors, and to the shorter exposure times when increasing the motion speed (as explained in Section 3).

Finally, Fig. 10(right) plots the average wind speeds measured by the on-board anemometer for the different trials, together with the nominal motion speeds (red solid line). Here we can see how the average wind velocities increase with the motion speeds used in this work, but they do not match them. This difference between sensed wind velocity and nominal motion speed is due mainly to the automatic control performed by the reactive navigator of the robot (which reduces the motion speed in the presence of obstacles, and prevents us to test higher velocities), as well as to the turbulences generated by the robot’s own movement and the natural wind-flows present in the environment. For this work, we are more interested in the "sensed" motion speeds than the nominal values, since those are the ones that really influence the e-nose measurements. Thus, hereinafter, we will refer to them as low, medium, fast and very-fast motion speeds. Lastly, for the static configuration case, it must be noticed that the small wind speed

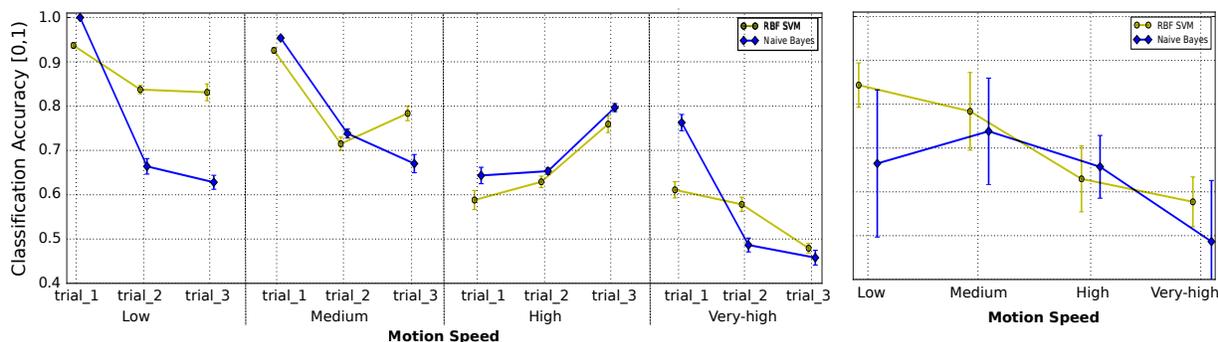


Figure 11: (Left) Classification accuracy for Naive Bayes and RBF SVM (average among the 20 repetitions of each trial) for different motion speeds when training with static data samples. (Right) Average over the three trials of each motion speed. A decrease in the classification rate is noticeable when increasing the motion speed.

measured (around 0.05 m/s) is produced by the fan employed to disperse the gases towards the sensing device.

5. The influence of the E-nose Motion Speed on the Classification Performance

This section is dedicated to the first of the two questions addressed in this work: How does the e-nose speed affect the classification accuracy?. As already pointed out in the introduction, this knowledge turns to be useful for a number of important applications of artificial olfaction, including environmental odor and pollution mapping in cities, inspection and monitoring of contaminated areas, etc.

To analyze to which extent the motion of the gas sensing device may affect the classification accuracy, we train the classifiers with samples of each chemical volatile collected in the static setup, and then, we study the classification performance for the set of increasing motion velocities. As described in Section 4.3, the data collected in the static configuration accounts for the natural turbulence phenomena of gas dispersion, similar to the case of a wall-mounted e-nose, but lacks the effects of moving the sensing device.

Given the specific training and test configuration of this experiment, it is not possible to adopt the classical cross-validation procedure to obtain robust estimations. Yet, it is possible to perform a "sample

level" cross-validation by randomly selecting a portion of the training set and test set to estimate the classification accuracy (instead of using all the training and test data available), and repeating the process for a number of times. Concretely, we randomly select 80% of the training and testing data to estimate the accuracy, and repeat this partitioning for 10 times. This allow us to obtain not only a single value for the accuracy, but also the related variance.

Fig. 11 shows the mean plus-minus one standard deviation of the classification results for each experimental trial (left sub-figure), as well as the average for each motion speed (right sub-figure), for the selected classifiers and features (see Section 4.2). From these results, different conclusions can be drawn. First, a noticeable reduction in the classification accuracy is observed when increasing the motion speed. This seems to confirm our suspicions about the negative impact that the motion speed has over classification rate because of the different changes in the sensor signals which are not considered during the static training. Thus, it is reasonable that, for this type of training, higher velocities (which may incur in more prominent changes in the sensor signals) show lower classification accuracies.

Then, it can be seen how results for low and medium motion speeds are relatively high (close to 0.8). This suggests that this kind of static training may be satisfactory for scenarios where the speed of the gas sensing device is kept low. We will go deeper

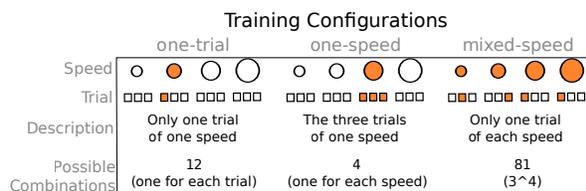


Figure 12: Different training configurations with samples gathered in motion.

into this matter in the next section, where the second question addressed in this work is analyzed in detail. Finally, notice that the estimator’s variance related to each trial is quite small, while there is a high variability among the different trials that compose the set of experiments of each motion speed (which is represented by the long error-bars in the right sub-figure). This high variance captures the reality of real applications, where an important number of uncontrolled parameters influences the collected data, and consequently, the classification outcomes.

In conclusion, these results support our suspicion that sampling chemical volatiles with an e-nose in motion carries a series of changes in the gas sensor’s response that do affect the posterior classification of the gases. Furthermore, a static training seems not to be appropriated for this scenario, even if it accounts for the turbulence phenomena of gas dispersion, and even if it presents similar e-nose’s amplitudes than those from the motion experiments.

6. The Relevance of Training

After corroborating in Section 5 that the classification accuracy is affected by the motion speed of the sensing device (when trained with data collected in a static configuration), in this section we analyze to what extent this drawback can be palliated by selecting a more appropriate training procedure.

Concretely, we now train the classifiers with data also collected in motion, proposing three different training configurations to analyze their impact on the classification accuracy (see Fig. 12 for a graphical representation of these three configurations).

(a) One-trial training: Here, we train the classifiers with data from one of the four speeds and only

one out of the three trials that compose the whole dataset. This guarantees that the training set is independent of the testing one because sample separation is performed at experiment level. The total of 12 train-test combinations are tested (one for each trial being used as training set), and for each combination we carried out the “sample level” cross-validation by randomly selecting 80% of the data and repeating the process 10 times to extract the average and variance of the estimator.

Fig. 13 shows the classification accuracy for each motion speed, averaging the three combinations that employ the same training speed. As can be seen, the overall classification accuracy has improved with respect to the static training configuration (SVM is in most cases above 0.9), something expected since now the training data is more similar to the testing data. Furthermore, the deterioration of the classification accuracy when varying the motion speed with respect to the one used at the training phase, is also reduced. Yet, in most cases the best results are obtained at the training speed, since the classifiers are tuned to the specific dynamics corresponding to that velocity.

An interesting point to remark is the high variability shown when using this training configuration, not only when estimating the accuracy at other motion speeds than the one used for training, but also for the same speed. This indicates that there are other factors that influence the classification which are not captured with only one trial.

(b) One-speed training: As in the previous case, we train the classifiers at one specific motion speed, but now we consider all the three trials that compose the set of data. This configuration may produce overfitting when testing at the same motion speed that in training, but yet, may be useful to extract conclusions for other motions since more data is available than in the above case (i.e more variability in the selected features at the training step).

Fig. 14 shows the averaged classification accuracy for this configuration (mean±std). In general, we can observe an important improvement with respect to the previous case, when only one trial was used for training (see Fig. 13). Not only the overall accuracy has improved, but the variability of the estimator has considerably been reduced. Related to the decreasing

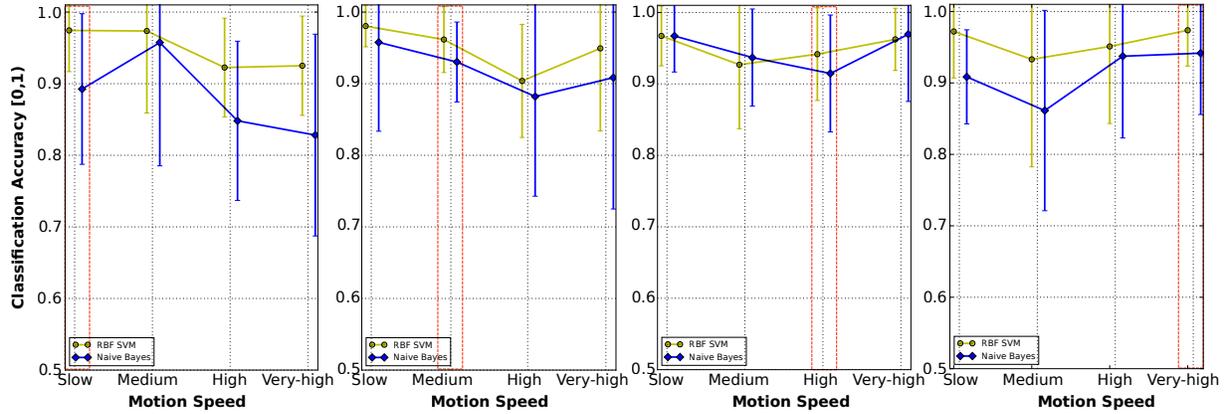


Figure 13: Classification accuracy when the classifiers are trained with in-motion data from only one trial at a time. The dashed-red rectangle indicates the motion speed of the training dataset.

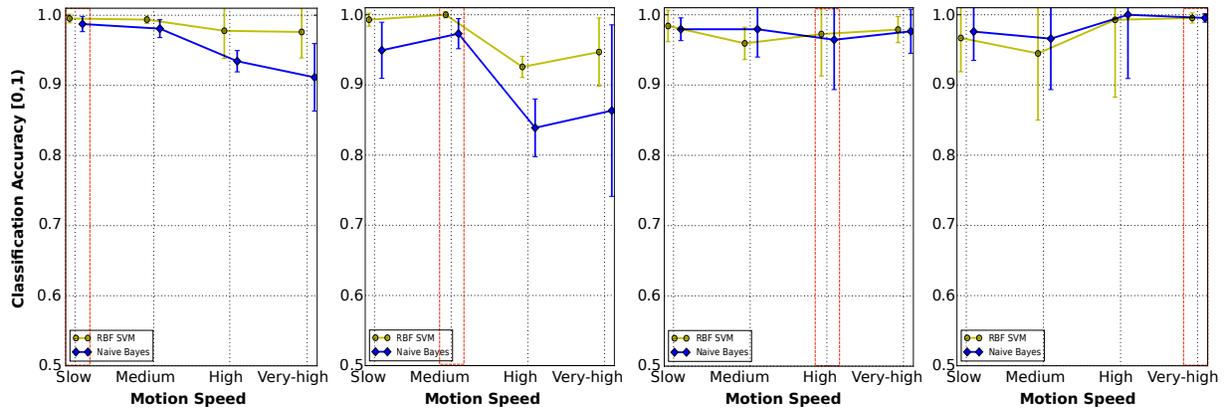


Figure 14: Classification accuracy when the classifiers are trained with in-motion data from the three trials of one motion speed. The dashed-red rectangle indicates the motion speed of the training dataset.

performance when testing at motion speeds different from the training one, we see that this effect is still noticeable but the accuracy drop is smaller.

(c) **Mixed-speed training:** Finally, we provide the classification results for the case of training the classifiers with a mixture of trials from the four motion speeds. Concretely, we select one trial for each speed, resulting in a training dataset composed of four trials. As in the previous training-configurations, we analyze all possible combinations for the training dataset, 81 in this case (three trials to the power of four motion speeds 3^4), and also carry

out the "sample level" cross-validation (randomly selecting 80% of the data for a total of ten repetitions).

Fig. 15 plots the average classification results from this training configuration. It can be observed that the mean accuracy values are almost independent of the motion speed, as well as the variance values are considerably small. Yet, the variance of the classification accuracy seems to increase with the motion speed, possibly due to the fact that higher velocities usually involve more and more significant changes in the sensor's response. Finally, although this training configuration may be the ideal one with respect to

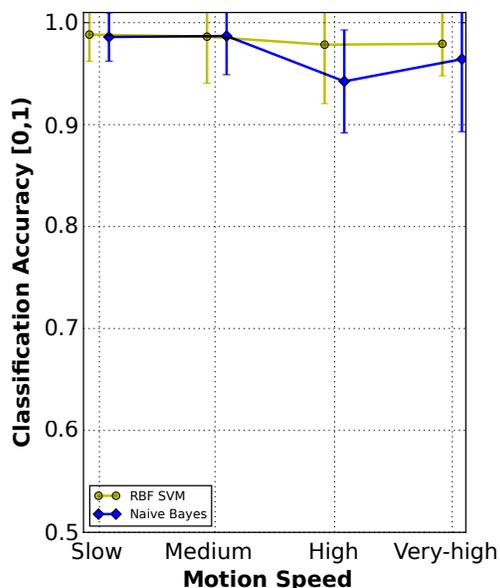


Figure 15: Classification accuracy for different motion speeds when the classifiers are trained with data from all tested speeds (one trial of each).

the classification accuracy, is the less practical in real applications, since it forces the acquisition of training samples at all possible motion speeds.

7. Conclusions and Future Research

In this work we have analyzed the problem of volatile chemical classification with an e-nose in motion, and specifically focused on the impact that the motion of the sensing device has over the classification accuracy.

We have reviewed the principal problems of real olfaction applications, and detailed some of the most common alterations that can be noticed in the gas sensor's response when considering motion. Then, we presented a dataset composed of 12 long trials, 2 chemical volatiles, 4 different motion speeds and a total of 240 inspections of an indoor scenario, to analyze and extract statistically representative conclusions of the phenomena under study.

Through different experiments, we have empirically corroborated that the changes in the sensor's

signal, induced by the movement of the sensing device, have an impact on the classification accuracy. Results shown a deterioration up to 30% when the motion speed of the data used for training highly differs from that of the testing.

Then, we have studied the influence of the training configuration, showing that training a classifier with data collected in motion yields, on average, more accurate outcomes than using a static setup. Moreover, we have found that it is not necessary to train the classifiers with data gathered at the exact same speed than the testing data to remove this negative correlation, but it suffices to capture the underlying dynamics. An extrapolation to other speeds not considered in the experiments is difficult to provide. Yet, since the performance's drop is caused by the different dynamics induced in the sensor signals, it seems reasonable that low motion speeds cause minor accuracy drops, while high speeds will be more problematic if not considered in the training. Nevertheless, we believe that the absolute speed is not a determinant parameter, but the gap between the speeds used to collect the training and testing datasets.

Therefore, for a real application, we can conclude that training the classifier with data collected at a few different motion speeds (from the range of possible velocities) should be enough to palliate this negative effect, and improve the overall classification performance. The optimal set of speeds to be considered during training remains still open, being necessary to perform an optimization phase that accounts for the particular classifier employed, the gas sensor types and features selected, and the dynamic characteristics of the environment to sample.

Finally, we have presented the case of training the classifier with data collected at all possible motion speeds, demonstrating that the selected classifiers were able to cope with the effects of the motion speed in the sensors response if presented at the training phase. The latter reinforces the independence of the results with the classifiers.

Our next steps along this line of research will be towards exploiting these results in real-world applications, extending our previous works on odor urban monitoring, and analyzing the impact on the classification and the map generation with mobile e-noses.

Acknowledgements

This work is partially supported by the Ministry of Economy and Competitiveness of Spain under grant DPI2014-55826-R, and by the Andalusia Regional Government and the European Regional Development Fund (ERDF) under project TEP2012-530.

References

- [1] G. Onkal-Engin, I. Demir, S. N. Engin, Determination of the relationship between sewage odour and BOD by neural networks, *Environmental Modelling & Software* 20 (7) (2005) 843–850. doi:<http://dx.doi.org/10.1016/j.envsoft.2004.04.012>.
- [2] J. G. Monroy, J. Gonzalez-Jimenez, C. Sanchez-Garrido, Monitoring household garbage odors in urban areas through distribution maps, in: *IEEE Sensors*, 2014, pp. 1364–1367. doi:[10.1109/ICSENS.2014.6985265](https://doi.org/10.1109/ICSENS.2014.6985265).
- [3] F.-M. Schleif, B. Hammer, J. G. Monroy, J. Gonzalez-Jimenez, J.-L. Blanco, M. Biehl, N. Petkov, Odor recognition in robotics applications by discriminative time-series modeling, *Pattern Analysis and Applications* (2015) 1–14 doi:[10.1007/s10044-014-0442-2](https://doi.org/10.1007/s10044-014-0442-2).
- [4] D. Hasenfratz, O. Saukh, C. Walser, C. Hueglin, M. Fierz, T. Arn, J. Beutel, L. Thiele, Deriving high-resolution urban air pollution maps using mobile sensor nodes, *Pervasive and Mobile Computing* 16, Part B (0) (2015) 268–285. doi:<http://dx.doi.org/10.1016/j.pmcj.2014.11.008>.
- [5] V. Hernandez Bennetts, E. Schaffernicht, V. Pomareda Sese, A. Lilienthal, M. Trincavelli, A novel approach for gas discrimination in natural environments with open sampling systems, in: *IEEE Sensors*, 2014, pp. 2046–2049. doi:[10.1109/ICSENS.2014.6985437](https://doi.org/10.1109/ICSENS.2014.6985437).
- [6] J.-G. Li, Q.-H. Meng, Y. Wang, M. Zeng, Odor source localization using a mobile robot in outdoor airflow environments with a particle filter algorithm, *Autonomous Robots* 30 (3) (2011) 281–292. doi:[10.1007/s10514-011-9219-2](https://doi.org/10.1007/s10514-011-9219-2).
- [7] T. Lochmatter, X. Raemy, A. Martinoli, Odor Source Localization with Mobile Robots, *Bulletin of the Swiss Society for Automatic Control* (46) (2007) 11–14.
- [8] M. Trincavelli, S. Coradeschi, A. Loutfi, Classification of odours for mobile robots using an ensemble of linear classifiers, in: *Proceedings of the 13th International Symposium on Olfaction and Electronic Nose (ISOEN)*, 2009, pp. 475 – 478. doi:[10.1063/1.3156587](https://doi.org/10.1063/1.3156587).
- [9] M. Trincavelli, A. Loutfi, An inspection of signal dynamics using an open sampling system for gas identification, in: *Workshop in Networked and Mobile Robot Olfaction in Natural, Dynamic Environments*, ICRA, 2010.
- [10] A. Vergara, J. Fonollosa, J. Mahiques, M. Trincavelli, N. Rulkov, R. Huerta, On the performance of gas sensor arrays in open sampling systems using inhibitory support vector machines, *Sensors and Actuators B: Chemical* 185 (2013) 462–477. doi:<http://dx.doi.org/10.1016/j.snb.2013.05.027>.
- [11] E. Bales, N. Nikzad, N. Quick, C. Ziftci, K. Patrick, W. Griswold, Citisense: Mobile air quality sensing for individuals and communities design and deployment of the citisense mobile air-quality system, in: *6th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2012, pp. 155–158. doi:[10.4108/icst.pervasivehealth.2012.248724](https://doi.org/10.4108/icst.pervasivehealth.2012.248724).
- [12] B. Elen, J. Peters, M. V. Poppel, N. Bleux, J. Theunis, M. Reggente, A. Standaert, The aeroflex: A bicycle for mobile air quality measurements, *Sensors* 13 (1) (2013) 221–240. doi:[10.3390/s130100221](https://doi.org/10.3390/s130100221).
- [13] S. Devarakonda, P. Sevusu, H. Liu, R. Liu, L. Iftode, B. Nath, Real-time air quality monitoring through mobile sensing in metropolitan areas, in: *Proceedings of the 2Nd ACM*

- SIGKDD International Workshop on Urban Computing, UrbComp '13, 2013, pp. 1–8. doi:
[10.1145/2505821.2505834](https://doi.org/10.1145/2505821.2505834).
- [14] P. Neumann, M. Bartholmai, J. H. Schiller, B. Wiggerich, M. Manolov, Micro-drone for the characterization and self-optimizing search of hazardous gaseous substance sources: A new approach to determine wind speed and direction, in: IEEE International Workshop on Robotic and Sensors Environments (ROSE), 2010, pp. 1–6. doi:[10.1109/ROSE.2010.5675265](https://doi.org/10.1109/ROSE.2010.5675265).
- [15] S. Sklavounos, F. Rigas, Validation of turbulence models in heavy gas dispersion over obstacles, Journal of Hazardous Materials 108 (12) (2004) 9–20. doi:[http://dx.doi.org/10.1016/j.jhazmat.2004.01.005](https://doi.org/http://dx.doi.org/10.1016/j.jhazmat.2004.01.005).
- [16] J. G. Monroy, J. Gonzalez-Jimenez, J.-L. Blanco, Overcoming the slow recovery of MOX gas sensors through a system modeling approach, Sensors 12 (10) (2012) 13664–13680. doi:[10.3390/s121013664](https://doi.org/10.3390/s121013664).
- [17] J. Gonzalez-Jimenez, J. G. Monroy, J.-L. Blanco, The multi-chamber electronic nose—an improved olfaction sensor for mobile robotics, Sensors 11 (6) (2011) 6145–6164. doi:[10.3390/s110606145](https://doi.org/10.3390/s110606145).
- [18] A. Hierlemann, R. Gutierrez-Osuna, Higher-order chemical sensing, Chemical Reviews 108 (2) (2008) 563–613. doi:[10.1021/cr068116m](https://doi.org/10.1021/cr068116m).
- [19] D. Graves, W. Pedrycz, Multivariate segmentation of time series with differential evolution, in: Conference of the European Society for Fuzzy Logic and Technology (IFSA/EUSFLAT), 2009, pp. 1108–1113.
- [20] E. Di Lello, M. Trincavelli, H. Bruyninckx, T. De Laet, Augmented switching linear dynamical system model for gas concentration estimation with mox sensors in an open sampling system, Sensors 14 (7) (2014) 12533–12559. doi:
[10.3390/s140712533](https://doi.org/10.3390/s140712533).
- [21] C. Xie, Interactive heat transfer simulations for everyone, The Physics Teacher 50 (4) (2012) 237–240. doi:[http://dx.doi.org/10.1119/1.3694080](https://doi.org/http://dx.doi.org/10.1119/1.3694080).
- [22] C. Sanchez-Garrido, J. G. Monroy, J. Gonzalez-Jimenez, A configurable smart e-nose for spatio-temporal olfactory analysis, in: IEEE Sensors, 2014, pp. 1968–1971. doi:[10.1109/ICSENS.2014.6985418](https://doi.org/10.1109/ICSENS.2014.6985418).