Wind-Independent Estimation of Gas Source Distance From Transient Features of Metal Oxide Sensor Signals

JAVIER BURGUÉS[ID] **AND SANTIAGO MARCO, (Senior Member, IEEE)**

Institute for Bioengineering of Catalonia (IBEC), Barcelona Institute of Science and Technology, 08028 Barcelona, Spain
Department of Electronics and Biomedical Engineering, Universitat de Barcelona, 08028 Barcelona, Spain

Corresponding author: Javier Burgués (jburgues@ibecbarcelona.eu)

ABSTRACT The intermittency of the instantaneous concentration of a turbulent chemical plume is a fundamental cue for estimating the chemical source distance using chemical sensors. Such estimate is useful in applications such as environmental monitoring or localization of fugitive gas emissions by mobile robots or sensor networks. However, the inherent low-pass filtering of metal oxide (MOX) gas sensors—typically used in odor-guided robots and dense sensor networks due to their low cost, weight and size—hinders the quantification of concentration intermittency. In this paper, we design a digital differentiator to invert the low-pass dynamics of the sensor response, thus obtaining a much faster signal from which the concentration intermittency can be effectively computed. Using a fast photo-ionization detector as a reference instrument, we demonstrate that the filtered signal is a good approximation of the instantaneous concentration in a real turbulent plume. We then extract transient features from the filtered signal—the so-called ''bouts''— to predict the chemical source distance, focusing on the optimization of the filter parameters and the noise threshold to make the predictions robust against changing wind conditions. This represents an advantage over previous bout-based models which require wind measurements—typically taken with expensive and bulky anemometers—to produce accurate predictions. The proposed methodology is demonstrated in a wind tunnel scenario where a MOX sensor is placed at various distances downwind of an emitting chemical source and the wind speed varies in the range 10-34 cm/s. The results demonstrate that models optimized with our methodology can provide accurate source distance predictions at different wind speeds.

INDEX TERMS Gas detectors, chemical sensors, signal processing, machine learning, time series analysis.

I. INTRODUCTION

The detection of intermittent gas patches is key for rapid gas source localization (GSL) in turbulent environments where the ''chemical plume'' is a collection of gas patches rather than a continuous trail [1], [2]. In turbulent conditions, the instantaneous concentration is highly fluctuating and chaotic, and smooth concentration gradients that could be exploited by a mobile agent to reach the gas source are only observed after long time-averaging (on the order of 5-10 minutes [3]). It is known that some flying insects such as the male moth— which are excellent plume trackers—rely on the frequency of ''odor hits'' to efficiently navigate along turbulent plumes to find distant mates or food [4]. An odor hit is declared when the measured concentration (typically of pheromone) exceeds the limit of detection of the olfactory receptor [1]. Inspired by these animals, researchers have developed reactive plume tracking algorithms for mobile robots to autonomously localize explosives, drugs or gas leaks [5]. The limited success of these algorithms is often attributed to the slow response time of the metal oxide (MOX) gas sensors these robots are typically equipped with due to their commercial availability, low cost, high sensitivity to many gases and ease of use [6], [7].

Detecting odor hits with a MOX sensor is very challenging due to the low bandwidth of this technology (<0.1 Hz [1]) as compared to the bandwidth of chemical plumes (several KHz [8]). Fig. 1 illustrates this problem. As can be

---

The associate editor coordinating the review of this manuscript and approving it for publication was Mouloud Denai[ID].
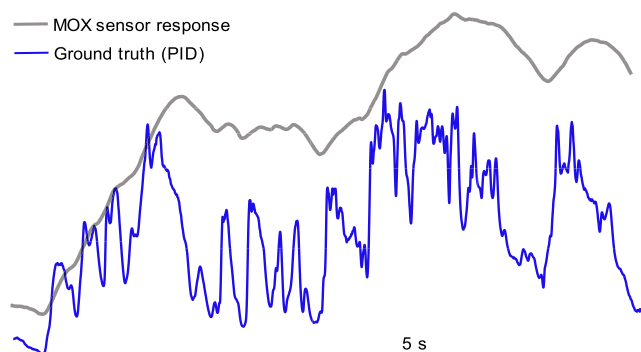
FIGURE 1. Illustration of the slow dynamics of a MOX gas sensor exposed to a turbulent plume. Data was captured in an open environment where a MOX sensor (gray trace) and a fast photo-ionization detector with a bandwidth of 330 Hz (blue trace) were placed close to each other in the centerline of a turbulent plume (See more details in Section III-A).

observed, the sensor output (gray trace) cannot follow the fast concentration fluctuations produced by strong turbulent wind (blue trace). It is clear that comparing the raw sensor output to a fixed threshold to detect individual odor hits [2], [9]–[13] will not be effective. The response time of a MOX sensor can be improved by different techniques: using novel coatings, modulating the sensor temperature, combining the response of multiple sensors hosted in independent gas chambers or post-processing the sensor signals. Signal processing methods are convenient because they can be directly applied to commercial sensors without changing the recommended operating mode and without adding extra components that would increase the cost and weight of the system. For example, if we assume that the dynamics of the MOX sensor can be modelled as ''leaky integration'' of the instantaneous concentration [14], a simple way to recover the (fast) stimulus from the slow response is to compute the smoothed derivative of the signal. Smoothing is necessary because differentiation degrades the signal-to-noise ratio (SNR) [15]. According to Schmuker *et al.* [16], the rising edges of the smoothed derivative—the so-called ''bouts''—are features that can be caused by individual odor filaments of the plume, so they could be potentially used to detect odor hits.

However, it is hard to validate that the detected bouts correspond to true odor hits because these are very difficult to observe or quantify. A proxy for such validation is the prediction error obtained when the detected bouts are used to predict the distance to the chemical source. If the detected bouts correspond to true odor hits, the source distance shall be estimated with high accuracy, according to previous experiments carried out with fast photo-ionization detectors [17]. Accordingly, Schmuker *et al.* [16] used this proxy to validate their algorithm in a wind tunnel scenario where a MOX sensor was placed in the centerline of a turbulent plume at different distances (range 25-145 cm) to a chemical source. Under a constant emission rate and wind speed, a linear regression model relating the bout count during three minutes and the source distance yielded a root mean squared error in cross-validation (RMSECV) of only 18 cm.

Based on these results, they claim that the number of bouts detected in a certain time interval is an accurate indicator of the source distance. However, these values were not benchmarked against traditional source proximity estimators, such as the mean or variance of the response [18], which may also perform well considering the long measurement windows (3 minutes). Additionally, when experiments were repeated at a higher wind speed, the parameters of the bout-based models changed considerably. This sensitivity to the wind speed implies that models trained at a certain wind speed are only accurate when the wind speed during future operation of the system is the same as in training, an unrealistic assumption in most scenarios. Although incorporating wind measurements into the predictive models could partially solve this issue, it limits the practical applicability of the algorithm because anemometers are expensive and heavy instruments that cannot be installed in robots with limited payload and are useless in indoor environments due to weak airflow.

To obtain low prediction errors, low-amplitude bouts that may be produced by noise or by insufficient smoothing of the derivative must be filtered out. For that, Schmuker et al. used a noise threshold (b_{thr}) computed as the mean plus three standard deviations of the amplitude of bouts detected in the sensor baseline (i.e., in the absence of gas). This rule of thumb, known as the three-sigma rule [19], says that 99.73% of the data in a normal distribution lie within three standard deviations from the mean, so it is empirically used by Schmuker et al. to treat 99.73% as near certainty that all bouts detected in the absence of gas will be filtered out. In other words, it is a way to minimize the number of false positives. It should be noted that this way of estimating the noise threshold implicitly assumes that (i) the bout amplitudes in the sensor baseline are normally distributed, (ii) a false positive (FP) is more important than a false negative (FN), and (iii) the measurement noise is additive.

The first assumption (normality) is questionable because the bout amplitude is strictly positive by definition [16]. Indeed, in a previous work [20] we empirically observed that the amplitudes of bouts detected in the sensor baseline are not normally distributed. The second assumption (FPs are more important than FNs) arises when the cut-off threshold of a binary classifier is estimated using exclusively data from the null hypothesis. If we were able to observe the distribution of bout amplitudes at different distances of the chemical source (Fig. 2), we would observe that the overlapping between the alternative hypothesis (presence of gas) and the null hypothesis (absence of gas) increases as the sensor moves away from the source. Far from the source, when gas concentration is very low, the amplitudes of the detected bouts are similar to the amplitudes of bouts detected in clean air. In this situation, it is very important to set the threshold of the detector as low as possible to increase the sensitivity of the system (i.e., to correctly identify low-amplitude bouts truly produced by the plume). A threshold determined solely from the null hypothesis, such as the three-sigma threshold, can achieve very high specificity (low number of FPs) at the cost of
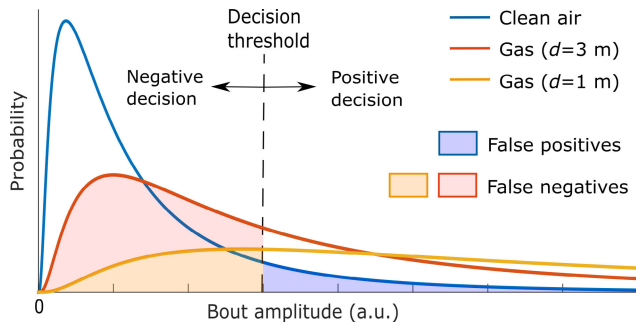
FIGURE 2. Simulated distributions of the bout amplitude in clean air (null hypothesis, blue curve) and at different distances to an emitting gas source (alternative hypothesis, red and yellow curves). The figure shows that the number of false negatives increase with distance to the source whereas the number of false positives remain constant. The distributions were generated using a log-normal function $ln(X) \sim N\left(\mu, \sigma^2\right)$ with $\mu = 0$ (blue curve), $\mu = 4$ (red curve), $\mu = 5$ (yellow curve) and $\sigma^2 = 1$ (all cases).


FIGURE 3. Flow diagram of the bout detection algorithm proposed by Schmuker *et al.* [16]. The meaning of each symbol is given in the text.

reduced sensitivity (high number of FNs). Regarding the third assumption (additive noise), we shall be aware that electronic noise can be additive but other sources of noise that also affect the sensor signal (e.g., chemo-transduction noise) usually depend on the chemical concentration [21]. This means that the noise level in clean air may not be representative of the noise level during gas exposure.

In this paper, we propose an optimization method for the parameters of the bout detection algorithm (smoothing factor and noise threshold) that instead of relying on unrealistic assumptions such as constant wind speed or normal distribution of the bout amplitudes, performs a multivariate grid search by varying all parameters simultaneously. The goal is to find the combination of parameters and predictive models (linear and non-linear) that provide the lowest prediction errors across different wind speeds. We then benchmark the results of the optimum bout-based models against the mean, maximum and variance of the response, and also against other bout-based features, such as the mean bout amplitude [29].

II. THE BOUT COMPUTATION ALGORITHM

To compute the bouts of a signal, Schmuker *et al.* [16] propose the signal processing pipeline (Equations 1-5 in [16]) illustrated in Fig. 3.

The sensor response x is first smoothed using a Gaussian low-pass filter [22] (with standard deviation σ_s s) to remove high-frequency noise. The smoothed response (x_s) goes through an EMA_α filter (with half-life time τ_h s) that produces the smoothed derivative (x_s') where the bouts are computed.

A. THE EMA_α DIGITAL FILTER

The EMA_α digital filter [14] is an approximation for a linear inverse filter of the dynamics of the MOX sensor response. At time $t = nT_s$, (being T_s the sampling interval) the output of the filter $y[n]$ is found by computing

$$y[n] = (1 - \alpha) \cdot y[n-1] + \alpha \cdot (x[n] - x[n-1]) \quad (1)$$

where $x[n]$ is the sensor response at time $t = nT_s$, $y[n-1]$ is the previous output of the filter, and α is the smoothing factor ($0 < \alpha \leq 1$). The EMA_α filter can be seen as the concatenation of two operations: first, taking the derivative of the input signal, i.e. $(x[n] - x[n-1])$; second, smoothing the derivative using an exponentially weighted moving average (EWMA) filter [23]. The smoothing factor α governs a trade-off between response time and SNR. As α increases, the closer is the filtered signal to the derivative and the faster is the response. However, being close to the derivative means also higher noise. A meaningful way to specify α is by the half-life time (s), τ_h, of the exponential decay, which is the time at which the exponential weight $(1 - \alpha)^k$ decays by one half

$$\alpha = 1 - \left(\frac{1}{2}\right)^{\frac{1}{\tau_h \cdot f_s}} \quad (2)$$

where τ_h is the half-life time (s) and f_s is the sampling frequency of x (Hz).

B. THE DETECTION OF BOUTS

The bouts are the rising edges of the smoothed derivative, x_s', which are delimited by two consecutive zero-crossings of the positive derivative of x_s', i.e. $x_s'' > 0$. A bout is characterized by two parameters: amplitude and duration. The amplitude of a bout is defined as x_s' at the end of the respective bout segment minus x_s' at the start of the same bout segment. The duration of a bout is defined as t at the end of the respective bout segment minus t at the start of the same bout segment. To remove low-amplitude bouts produced by noise, Schmuker et al. define the amplitude threshold

$$b_{thr} = \mu + 3\sigma \quad (3)$$

where μ and σ are the mean and standard deviation, respectively, of the amplitudes of bouts detected in the sensor baseline (i.e. in clean air) since these are surely produced by noise. All bouts with amplitude lower than b_{thr} are thus removed. If the underlying distribution follows a Gaussian shape, this threshold guarantees a specificity 99.73%.

C. OPTIMIZATION OF THE ALGORITHM PARAMETERS

The bout algorithm has three parameters to optimize (σ_s, τ_h and b_{thr}). The first two parameters (σ_s and τ_h) are coupled and they control the smoothness of the filtered signal. Schmuker et al. studied the influence of these two parameters on the relationship between the bout count and the source distance assuming that b_{thr} defined by (3) is an optimum threshold. They found that small values of σ_s and τ_h (i.e.,
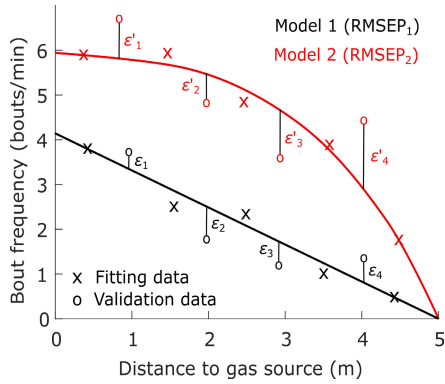
FIGURE 4. Sketch of the proposed optimization method. The bout frequency is computed at several downwind distances using different algorithm parameters (red and black crosses) leading to different functional relationships between the bout frequency and the source distance (red and black curves). The best model is the one that minimizes the prediction error in external validation samples (open circles).

FIGURE 5. Generator of turbulent plumes in an open environment. A pressurized air outlet (6.3 mm radius, 20 L/min) placed 50 cm behind a beaker (5 cm radius) filled with 200 mL of ethanol created the plume. A custom sensing board containing a naked MOX sensor and a fast photo-ionization detector was placed in the centerline of the plume, at 105 cm of distance to the gas source.

high bandwidth of the derivative filter) resolve better the short bouts encountered close to the source but fail to resolve the long bouts that occur further away, leading to a convex relationship between bout count and source distance. On the other hand, high smoothing (i.e. low-pass filtering) produces a concave behavior in which the sensitivity, i.e. the slope of the regression, is high far from the source and small near the source. They concluded that $\sigma_s = 0.3$s and $\tau_h = 0.4$ s are optimum values in their experimental dataset but no specific methodology to compute these values was provided. After the values of these two parameters are fixed, the value of the third parameter (b_{thr}) is computed through (3), which is not optimal due to the reasons mentioned in the introductory section.

To overcome these problems, we propose a specific methodology to optimize the bout algorithm parameters, sketched in Fig. 4. First, the bout frequency is computed at different distances to the source using different combinations of σ_s, τ_h and b_{thr}. The bout frequency (bouts/minute) is simply the number of bouts detected in a certain time interval divided by the length of such interval. For each parameter combination, the model (linear or non-linear) that bests fit the functional relationship between bout frequency (dependent variable) and source distance (independent variable) is selected via cross-validation (CV). In CV, a portion of the data is used to fit the models and the remaining data is used to evaluate the performance of the model in predicting samples not seen during model training. The goodness of the models is determined by the root mean squared error (RMSE)

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{n}\left(\hat{y}_i - y_i\right)^2}{n}} \qquad (4)$$

where $\hat{y}_i$ is the distance (m) predicted by the model, y_i is the true distance (m) and n is the number of validation samples. The performance of the model for different combinations of parameters can be assessed by computing the RMSE in external validation samples (RMSEP), which are samples reserved exclusively for validation purposes (i.e., excluded

from model fitting and optimization). Thus, the classification problem is converted into a regression problem that only requires knowing the distance at which measurement were taken.

III. EXPERIMENTAL

We used two experimental datasets in this work: (i) the open environment dataset [24] and (ii) the wind tunnel dataset [25]. The first one contains a small set of measurements performed with a MOX sensor and a fast Photo-Ionization Detector (PID) in an open environment. We use it to verify that the smoothed derivative of the MOX sensor response can be a good approximation of ground truth (PID signal). The second dataset contains a large collection of measurements performed with multiple MOX sensors (no PID in this case) inside a small wind tunnel. This comprehensive dataset, which is the same used by Schmuker et al. in their experiments, is used to test the proposed optimization of the bout detection algorithm. Data analysis was performed using Python (version 3.6.2) and the scikit-learn package (version 0.19.1) [26]. Some graphics were generated using MATLAB R2019a (The Mathworks, Nattick, USA).

A. OPEN ENVIRONMENT DATASET

A turbulent plume is generated in an open environment by passing an air flow stream over a beaker filled with ethanol (Fig. 5). A custom sensing board containing a naked (i.e., without cap) MOX sensor (TGS 2602, Figaro Engineering, Japan) and a fast PID (miniPID 201A, Aurora Scientific, Canada) with a bandwidth of 330 Hz were placed in the centerline of the plume at 105 cm distance to the gas source.

To ensure that both sensors are exposed to the same gas flow, the MOX sensor is hosted inside a miniaturized gas chamber (volume of 0.325 mL) and the inlet of the PID is connected to the exhaust port of the gas chamber (see more details in [24]). The PID is used both as reference device and to draw the sample inside the gas chamber (it has a pump with a flow rate of 1 L/min). The heater of the MOX sensor is powered at 5 V and the sensor output is acquired during 5 minutes at a sampling frequency of 1 kHz using a voltage divider ($RL = 10$ kΩ) and a USB-6002 datalogger

FIGURE 6. Schematic representation of the wind tunnel (2.5 m × 1.25 m). The five downwind measuring locations used in this work are indicated by black squares. A chemical plume has been outlined for illustration purposes (the actual plume might differ from that).

(National Instruments, USA). The PID output was sampled also at 1 kHz.

The smoothed derivative is computed through the convolution of the sensor response and a low-pass differentiator filter. For simplicity, this filter is defined as the convolution of a derivative kernel of the form $[1, -1]$ and a moving average filter with a window size of 20 ms.

B. WIND TUNNEL DATASET
The wind tunnel dataset [25] contains recordings from nine gas sensor arrays exposed to turbulent gas mixtures inside a small wind tunnel. We used a portion of this dataset corresponding to the measurements acquired by the gas sensor array #5 (placed in the plume centerline). Figure 6 shows a simplified schematic of the wind tunnel. Turbulent gas plumes of different gases were created by injecting pressurized gas into one end of the tunnel and dragging it with an exhaust fan from the other end of the tunnel. Different wind speeds (range 10-34 cm/s) were created by varying the rotational speed of the exhaust fan. The sensor array–which contains 8 MOX sensors–was sequentially placed at five downwind distances to the gas source (range 25 - 145 cm). Each sensor array integrates eight MOX sensors (Several TGS models: 2600 (2x), 2602 (1x), 2610 (1x), 2611 (1x), 2612 (1x) and 2620 (2x), from Figaro Engineering Inc.) operated at the same, constant heater voltage (range 4.0-6.0 V). The sensor response was measured with a voltage divider (10 kΩ load resistor) and sampled with a 12-bit ADC.

A total of 900 distinct experiments were performed by varying the distance to source (6 possible), gas (10 possible), sensor temperature (5 possible) and wind speed (3 possible). Each experiment was repeated 20 times. In each experiment, the following procedure was performed: (i) Measure the baseline response of the sensors for 20 s in the absence of gas, (ii) Release the selected gas for 3 minutes, (iii) Circulate clean air for 1 minute to record the sensor recovery and (iv) Purge the wind tunnel by setting the fan at maximum speed. Even though the gas release started at time $t = 20$ s, the recorded signals show a transient behavior between $t = 30$s and $t = 90$ s because the gas requires some time to travel through the tunnel. The time to reach a stable mean response depends on the distance to the source (sensors closer to the source stabilize faster) and the wind speed (signals stabilize faster in experiments performed at high wind speeds). A common


FIGURE 7. Frequency response of the low pass differentiator filter used to process the MOX sensor signal in the open environment dataset. The response of an ideal differentiator is displayed as a reference.

time frame where all signals are stable is $t = [100, 200]$ s. We denote this period as the "stable gas release".

We use the following data from the wind tunnel dataset: board #5 (located in the plume centerline), sensor #4 (TGS 2600), heater voltage of 6.0 V and Acetaldehyde gas. We selected the same board used by Schmuker et al. in their experiments to compare our results to their results. The choice of sensor, heater voltage and gas were arbitrary, and analyzing the influence of these parameters in the results is out of the scope of this study. For the selected configuration, sensor recordings are available at 6 distances, 3 wind speeds and 20 trials per distance/wind combination (360 experiments). To find the optimum parameters of the algorithm, the bouts of the 360 experiments are extracted using different values of σ_s, τ_h and b_{thr}. We perform a grid search: $\sigma_s = \{0.05, 0.3, 1.6\}$ s, $\tau_h = \{0.05, 0.4, 1.6\}$ s and $b_{thr} = \{\mu + \sigma, \mu + 3\sigma, \mu + 5\sigma\}$, where the center values of each parameter match those used by Schmuker et al.

For each combination of these parameters (27 possible) and wind speeds (3 possible), we select the model that best fits the relationship between bout frequency and distance to the source, using a cross-validation scheme on the bouts extracted from the stable gas release period of the first 10 trials. Due to the limited number of values of the independent variable (6 distances), we use simple models such as linear, polynomial (up to order 3) and single-term exponential. The optimum models are then used to predict the distance to the source from the bouts extracted in the last 10 trials (unseen during model fitting/optimization).

IV. RESULTS AND DISCUSSION
We first visualize the signals in the open environment dataset and then proceed to analyze in detail the wind tunnel dataset.

A. OPEN ENVIRONMENT DATASET
Fig. 7 shows the frequency response of the low-pass differentiator filter described in Section III-A. As can be observed, the derivative is computed only at low frequencies (cut-off frequency of approximately 20 Hz) while high-frequency noise is attenuated. The result of filtering the MOX signal with this filter is shown in Fig. 8. The smoothed derivative
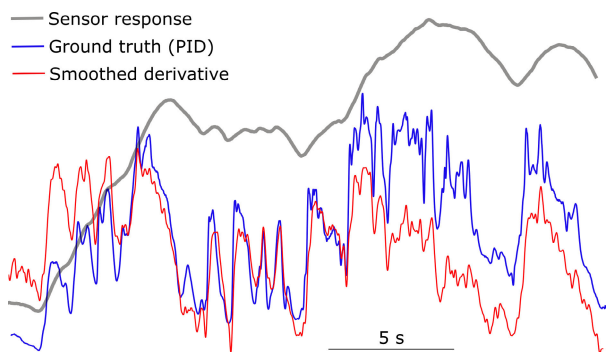
FIGURE 8. Comparison between the dynamics of the MOX sensor, the PID and the smoothed derivative of the MOX signal in the open environment dataset. The sensor and the PID were placed at 105 cm distance of the gas source. The figure illustrates that the smoothed derivative can follow the fluctuations of the instantaneous concentration.


FIGURE 9. Raw signals captured by sensor #4 (TGS 2600) of board #5 (plume centerline) in trial #1 at different downwind distances (coded by line transparency) and wind speed (coded by line color). The figure shows how the m ean amplitude and variability of the signals decrease with increasing downwind distance and wind speed.

is a much better reconstruction of ground truth, effectively improving the response time of the sensor.

B. WIND TUNNEL DATASET

The signals recorded in the plume centerline (Fig. 9) are characterized by a stable *plateau* concentration, due to gas accumulation within the wind tunnel, modulated by some fluctuations due to turbulence. The wind speed has a strong influence on the mean intensity and fluctuations of the signals, with lower wind speeds yielding signals with higher intensity and variability (remember that the wind speed is related to the ventilation efficiency of the tunnel).

The mean and standard deviation of the sensor signals are two straightforward features that can be used to predict the source distance. However, predictive models built using these features exhibit undesirable properties such as non-monotonicity and low sensitivity far from the source (Fig. 10), leading to relatively high RMSEP values of 21-28 cm. In the case of the mean response, these errors are partly produced by the non-monotonic behavior observed in the mean response at $d = 50$cm. This anomalous behaviour could be related to the high variance of the wind speed at that location (see Fig. 3b in [25]) probably owing to the geometry of the wind tunnel and the location of the exhaust fan. In the case of the standard deviation, the high scattering among
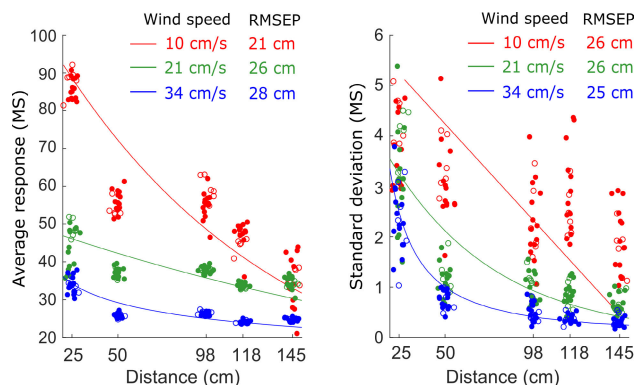

FIGURE 10. Mean and standard deviation of the response of sensor #4 (board #5) at different downwind distances and wind speeds (color-coded). The solid lines represent the optimum model for the fitting samples (trials 1-10, solid circles), whereas the RMSEP (see legend) is computed on the external validation samples (trials 11-20, open circles).

FIGURE 11. Frequency response of the low pass differentiator filters used to process the MOX sensor signal in the wind tunnel dataset. The filter with default parameters is highlighted. The response of an ideal differentiator is displayed as a reference.

different trials seems to indicate that a 3-minute measurement window may be too short to obtain a reliable estimate of this feature in the type of plume generated in the wind tunnel.

Based on this first analysis, it seems that the mean and the standard deviation of the response are not reliable features for predicting the distance to the source and this motivates studying whether bout-based features can improve the results. The frequency response of the low-pass differentiator filters obtained with different combinations of σ_s and τ_h are shown in Fig. 11. As can be seen, the shape of the filter is mostly governed by σ_s whereas the effect of τ_h is to change the cut-off frequency. For example, for a fixed value of τ_h filters with higher σ_s attenuate faster the signal (narrower transition band). The impact of τ_h in the frequency response increases with decreasing values of σ_s. For instance, at $\sigma_s = 1.6$s filters with different values of τ_h are quite similar whereas the differences are much larger at $\sigma_s = 0.05$s. The combination $\sigma_s = 0.3$ s and $\tau_h = 0.4$s (optimum according to Schmuker et al.) produces a filter with a moderate transition band and a cut-off frequency of approximately 0.5 Hz.
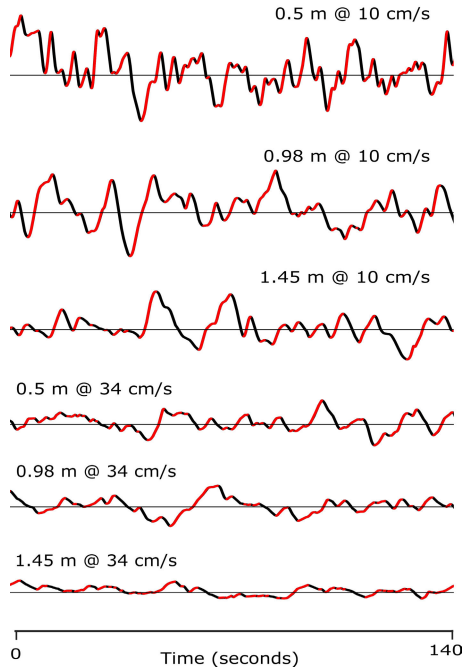
FIGURE 12. EMA_α-filtered signals ($\sigma_s = 0.3$ s, $\tau_h = 0.4$ s) during the stable gas exposure of the first trial of various distance/wind combinations (indicated above each subplot). All detected bouts ($b_{thr} = 0$) are marked in red.

FIGURE 13. Bout frequency and mean bout amplitude computed on sensor #4 (board #5) at different downwind distances and wind speeds (color-coded). The bout algorithm parameters were $\sigma_s = 0.3$ s, $\tau_h = 0.4$ s and $b_{thr} = \mu + 3\sigma$. The solid lines represent the optimum model for the fitting samples (trials 1-10), whereas the RMSEP (see legend) is computed on external validation samples (trials 11-20).


FIGURE 14. Histogram of bout amplitudes (sensor #4, board #5, $\sigma_s = 0.3$ s, $\tau_h = 0.4$ s) in clean air (red) and at various downwind distances (blue, green, yellow) of an active chemical source. Wind speed is 34 cm/s. The figure illustrates the high overlapping between the histograms obtained at different distances. The y-axis is in logarithmic units. The $\mu + 3\sigma$ and 99.7% thresholds correspond to bout amplitudes of 0.06 and 0.11 *MS*, respectively.

Examples of filtered signals at different distances to the source using this combination of parameters are shown in Fig. 12. The first thing that can be observed is that differentiation removed the plateau previously observed in the raw signals of Fig. 9. The amplitude and frequency of the detected bouts decrease with increasing distance to the source or increasing wind speed. As the sensor moves away from the source, the signals become weaker and low-amplitude bouts comparable to those detected in the sensor baseline become more frequent. At a certain distance, it will be hard to distinguish bouts induced by the plume from those produced by noise.

The bout frequency (BF) computed with $b_{thr} = \mu + 3\sigma$ produces a quasi-linear behaviour across the studied distance range and wind speeds (Fig. 13). The RMSEP decreases with increasing wind speed and is always lower than the one obtained for the mean and standard deviation of the response (Fig. 10). This is probably related to the fact that the variance of the wind speed in the wind tunnel decreases with increasing rotational speed of the exhaust fan (see Fig. 3 in [25]). Although the obtained RMSEP values may seem sufficiently small, it should be remained that they only represent the performance of the system when training and test wind speed is the same. From the graphic it is obvious that the prediction error will increase if, for example, the model is trained at 10 cm/s (red line) and tested at 34 cm/s (blue samples).

On the other hand, the mean bout amplitude (MBA) shows an exponentially increasing behavior at every wind speed (Fig. 13). The RMSEP increases with increasing wind speed due to the very low sensitivity of the exponential models far
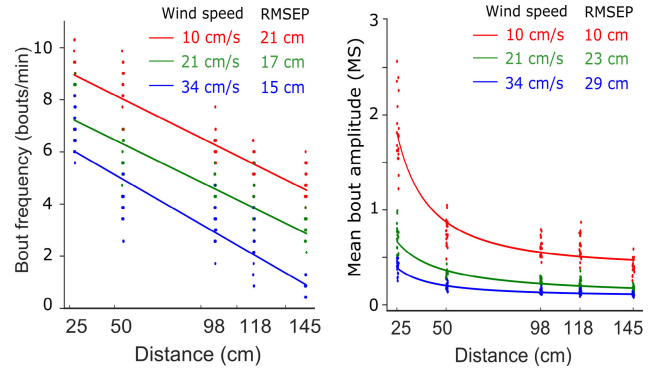
from the source. Interestingly, at the lowest wind speed the MBA provides the lowest RMSEP across all studied features (10 cm). We also found that the MBA is less sensitive to the value of b_{thr} than the BF, which degrades its behaviour if b_{thr} is either too low or too high. It should be noted that the BF counts bouts in a similar manner regardless of their amplitude, so that a high-amplitude bout produced by the plume is as important as a low-amplitude bout produced by noise and not properly filtered out. On the other hand, the MBA increases in proportion to the amplitude of each detected bout, so the negative effect of low-amplitude bouts is diluted.

1) IS THREE-SIGMA AN OPTIMUM THRESHOLD?
The $\mu + 3\sigma$ threshold implicitly assumes that the amplitudes of the bouts detected in the sensor baseline follow a Gaussian distribution. This is obviously not correct because the bout amplitude is a random variable bounded from zero (Fig. 14). Resultingly, the $\mu + 3\sigma$ threshold represents a 98.1% percentile in this dataset, instead of the 99.87% corresponding to a normal distribution. This is, approximately 1.9% (instead of

FIGURE 15. Bout frequency vs. downwind distance for different algorithm parameters. Wind speed is 21 cm/s. Each subplot represents a different value of b_{thr} (indicated above each subplot). The line color indicates the value of σ_s (red: 0.05 s, green: 0.30 s, blue: 1.60 s) and the line style indicates the value of τ_h: (-) 0.05 s, (--) 0.40 s, (:) 1.60 s.

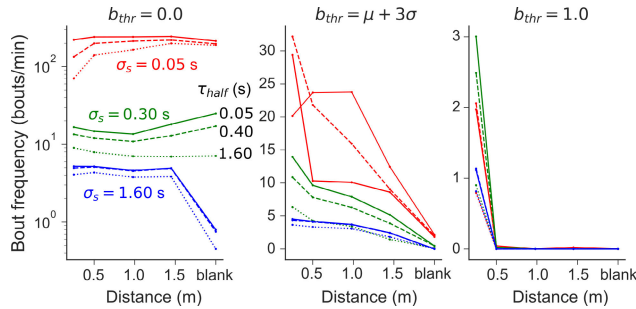**TABLE 1.** Prediction error (cm) of the bout frequency models under several combinations of smoothing factors (rows), amplitude thresholds (blocks) and test wind speeds (columns). The training wind speed is (A) 10 cm/s; (B) 21 cm/s; (c) 34 cm/s.

A

σ_s (s)	τ_h (s)	$b_{thr}=\mu+1\sigma$			$b_{thr}=\mu+3\sigma$			$b_{thr}=\mu+5\sigma$		
		10	21	34	10	21	34	10	21	34
0.05	0.05	30	57	76	24	51	51	28	45	52
	0.40	41	47	57	30	52	74	17	61	81
	1.60	33	41	57	45	41	41	40	43	75
0.30	0.05	24	30	45	21	40	72	19	47	82
	0.40	25	36	49	21	42	71	18	49	83
	1.60	23	36	58	20	41	77	19	44	84
1.60	0.05	27	43	66	28	58	111	24	57	104
	0.40	30	44	68	28	52	100	25	56	105
	1.60	29	42	62	28	49	82	29	56	91

B

σ_s (s)	τ_h (s)	$b_{thr}=\mu+1\sigma$			$b_{thr}=\mu+3\sigma$			$b_{thr}=\mu+5\sigma$		
		10	21	34	10	21	34	10	21	34
0.05	0.05	39	43	46	68	37	45	135	33	45
	0.40	42	44	45	65	25	46	61	23	37
	1.60	38	34	48	54	41	37	63	28	45
0.30	0.05	25	25	30	32	19	31	37	18	37
	0.40	31	23	29	40	17	37	48	17	34
	1.60	43	21	35	48	20	41	53	20	38
1.60	0.05	33	33	38	39	30	35	46	28	38
	0.40	34	36	41	39	32	36	46	31	38
	1.60	35	38	41	41	37	38	62	33	38

C

σ_s (s)	τ_h (s)	$b_{thr}=\mu+1\sigma$			$b_{thr}=\mu+3\sigma$			$b_{thr}=\mu+5\sigma$		
		10	21	34	10	21	34	10	21	34
0.05	0.05	47	45	43	97	44	39	463	97	36
	0.40	54	50	43	85	39	29	92	33	31
	1.60	43	43	44	90	67	20	87	50	21
0.30	0.05	40	31	20	60	35	15	66	34	16
	0.40	54	30	19	65	34	15	80	39	18
	1.60	68	38	21	90	50	20	90	45	22
1.60	0.05	48	37	30	75	49	21	81	53	19
	0.40	49	38	33	75	49	22	80	51	20
	1.60	60	44	32	81	51	25	89	54	21

The green rectangle highlights the bout frequency with optimum parameters ($\sigma_s = 0.3$ s, $\tau_h = 0.05$ s, $b_{thr}=\mu + \sigma$). The red rectangle highlights Schmuker's combination ($\sigma_s = 0.3$ s, $\tau_h = 0.4$ s, $b_{thr}=\mu + 3\sigma$).

0.13%) of the bouts detected in clean air will be incorrectly declared as "true" bouts. One could be tempted to increase the threshold to reach the 99.87 percentile and keep a low number of false alarms; however, this will reduce the sensitivity of the detector (increase the false negative rate) because the positive and negative classes are highly overlapped (especially at high wind speeds and distant downwind measuring locations).

2) GRID-SEARCH OPTIMIZATION OF ALGORITHM PARAMETERS

We performed a multivariate optimization of the three algorithm parameters (σ_s, τ_h and b_{thr}) by varying all of them simultaneously. We found that if the threshold is set too low (e.g. $b_{thr} = 0$) or too high (e.g. $b_{thr} = 1$), there is no combination of smoothing factors that produces a monotonically increasing relationship between bout frequency and source distance with enough sensitivity across the studied distance range (Fig. 15). The desired behaviors are only found for intermediate thresholds (e.g. $\mu + 3\sigma$) and a subset of smoothing parameters. For example, the combinations {$\sigma_s = 0.05$ s, $\tau_h = 0.4$ s} (red dashed line) and {$\sigma_s = 0.3$ s, $\tau_h = 0.05$ s} (green solid line) exhibit monotonic behaviors and high sensitivity.

However, the optimum configuration should not only achieve monotonicity and high sensitivity, but also small scattering of the different trials around the regression line, i.e. it must minimize the RMSEP. Computing the RMSEP at different wind speeds (Table 1) confirmed that moderate smoothing of the raw signal ($\sigma_s = 0.3$ s) combined with low smoothing of the derivative ($\tau_h = 0.05$ s) and low amplitude threshold ($b_{thr} = \mu + \sigma$) yields, in average, the best results. Higher thresholds, such as $\mu + 3\sigma$ or $\mu + 5\sigma$, can produce lower prediction errors when the wind conditions in training and test are the same; however, the proposed combination ($\sigma_s = 0.3$ s, $\tau_h = 0.05$ s, $b_{thr} = \mu + \sigma$) is the most robust when the wind conditions in training and test are different. For example, if the training and test wind speed are both 10 cm/s, Schmuker's combination achieves 12% lower errors than the proposed combination (21 cm versus 24 cm). However, if
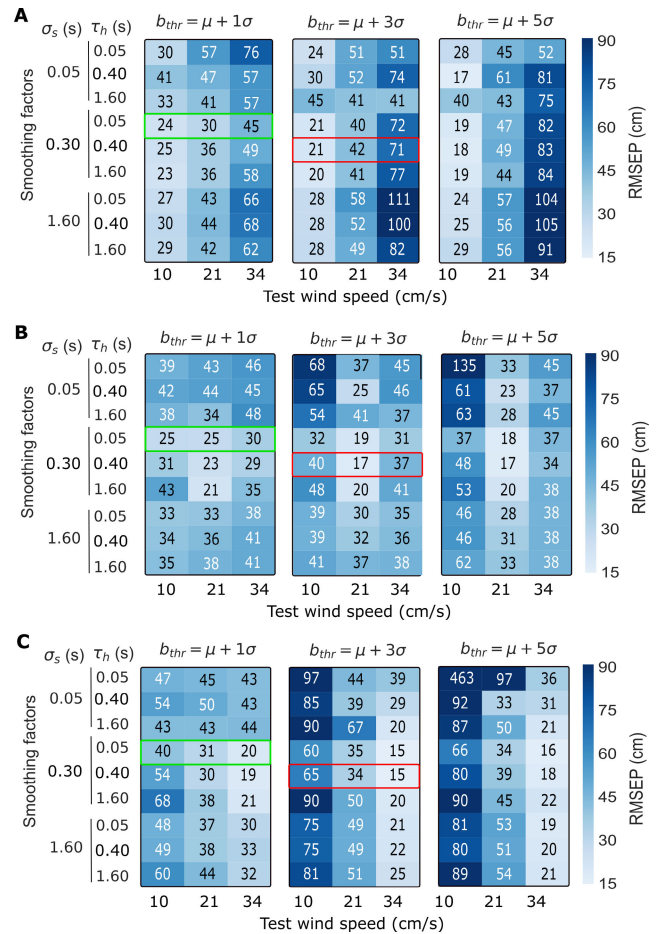
the test wind speed increases to 0.34 m/s while keeping the training speed at 0.10 m/s (worst case scenario in this dataset) the error in Schmuker's combination (71 cm) is 58% higher than in the proposed combination (45 cm).

C. OVERALL PERFORMANCE

A comparison of the prediction error of different signal features is presented in Table 2. The general trend is that most estimators achieve low RMSEP when the wind speeds in training and test are the same but degrade otherwise. In the case of matching wind speeds, the variance of the ema filtered signal and the BF with Schmuker's parameters achieve the lowest fitting errors (13-21 cm). The low end of this range corresponds to training and testing at 34 cm/s, whereas the highest errors are obtained when training and testing are both performed at 10 cm/s. In the case of non-matching wind speeds, the proposed method (outlined with a green box) seems to greatly improve the other estimators. For example, in the worst-case scenario (training at 10 cm/s and testing at

TABLE 2. Prediction error (cm) of several signal features, for various training wind speeds (indicated on the top of each block) and test wind speeds (indicated on the bottom of each column).

	w_{train} = 10 cm/s			w_{train} = 21 cm/s			w_{train} = 34 cm/s		
Mean	22	54	76	71	26	306	84	76	29
Variance	44	44	44	69	32	92	76	52	23
Maximum	22	55	79	66	23	297	84	71	25
Variance$_{EWMA}$	20	177	510	66	17	107	80	52	14
Bout frequency*	24	30	45	25	25	30	40	31	20
Bout frequency	21	42	71	40	17	37	65	34	15
Bout amplitude	10	57	82	47	23	33	69	34	29
	10	**21**	**34**	**10**	**21**	**34**	**10**	**21**	**34**

* Optimum   Test wind speed (cm/s)

The green rectangle highlights the bout frequency with optimum parameters (σ_s = 0.3 s, τ_h = 0.05 s, b_{thr}= $\mu + \sigma$).

34 cm/s) it achieves a RMSEP of 45 cm whereas the BF with Schmuker's parameters (second best estimator) achieves an RMSEP of 71 cm (58% difference). A similar thing happens when training at 34 cm/s and testing at 10 cm/s. In the case of training at medium wind speed (21 cm/s), the proposed method keeps the RMSEP in the range 25-30 cm regardless of the wind speed while the other estimators degrade the RMSEP by at least a factor of two.

V. CONCLUSION

We have experimentally demonstrated that multivariate optimization of the parameters of the bout detection algorithm can improve the robustness of the algorithm output under changing wind conditions. The optimum configuration that we found applies moderate smoothing to the raw signal, low smoothing to the derivative and low noise threshold. The latter fact suggests that high sensitivity is more important than high specificity for accurately predicting the distance to a gas source in this dataset, although this shall be further validated using other sensors and gases available in the dataset. Even if a certain application requires high specificity, applying the $\mu + 3\sigma$ threshold may produce unexpected results because the distribution of amplitudes of baseline bouts does not follow a Gaussian distribution. In this case, empirically estimating the threshold corresponding to the desired percentile of the negative class seems a more reliable approach.

An additional advantage of lower thresholds is the lower detection delay and the higher number of bouts detected per minute, which theoretically allows for shorter measurement windows. Indeed, the 2-minute measurement windows that we used in this work may be considered too long for a real GSL application. Future work may explore the influence of the measurement window in the results, by analyzing chunks of measurements of different size. In this context, it makes sense that the robot reacts in real-time to each detected bout instead of performing static measurements at fixed locations. A straightforward approach could be to use the bouts as a replacement of odor hits in GSL algorithms such as Infotaxis [2] or Pang and Farrell's algorithm [13]. The wind tunnel dataset already contains a 2D grid measurements in which

the viability of this approach could be simulated prior to performing experiments with a robot in a real scenario.

A second question addressed by this work is whether the bout frequency outperforms other estimators of source proximity such as the mean and variance of the response or the bout amplitude. Although most estimators worked reasonably well in the wind tunnel dataset if the training and test wind speed were equal, the bout frequency with optimized parameters was the best estimator when considering all wind conditions. Specifically, the optimized models achieved a maximum prediction error of 45 cm (over a distance range of 1.45 m) in the worst-case scenario where the models are fit at wind speed of 10 cm/s and tested at 34 cm/s. In similar conditions, the optimum model reported by Schmuker et al. and the variance of the response (a typical measure of intermittency) yielded maximum errors of 71 cm (58% difference) and 92 cm (104% difference), respectively. The reason for the good performance of the bout frequency may be the linear relationship with the distance to the source, as compared to the polynomial and exponential behaviours observed in other signal features. The response variance, which some authors consider a reliable estimator of source proximity, did not work well in our experiments. In the likely case that the test wind conditions are unknown, our results suggest that lower prediction errors will be obtained if the models are trained at medium wind speed.

In follow-up works we plan to exploit the relationship between bout amplitude and bout duration to improve the classification of bouts. One problem of fixed amplitude thresholds is that bouts with high amplitude and high duration that appear in the baseline are often classified as true bouts. Linear classifiers such as support vector machines (SVM) or non-linear ones such as random forests or k-nearest neighbours (KNN) could generate decision regions that overcome this issue.

REFERENCES

[1] D. Martinez, L. Arhidi, E. Demondion, J.-B. Masson, and P. Lucas, "Using insect electroantennogram sensors on autonomous robots for olfactory searches," *J. Vis. Exp.*, vol. 90, no. 90, 2014, Art. no. e51704.

[2] M. Vergassola, E. Villermaux, and B. I. Shraiman, "'Infotaxis' as a strategy for searching without gradients," *Nature*, vol. 445, no. 7126, pp. 406–409, 2007.

[3] E. A. Bakkum and N. J. Duijm, "Vapour cloud dispersion," in *Methods for the Calculation of Physical Effects 'Yellow Book'*, C. J. H. van den Bosch and R. A. P. M. Weterings, Eds., 3rd ed. The Hague, The Netherlands: Ministerie van Verkeer en Waterstaat, 2005, ch. 4, pp. 1–139.

[4] D. R. Webster and M. J. Weissburg, "Chemosensory guidance cues in a turbulent chemical odor plume," *Limnol. Oceanogr.*, vol. 46, no. 5, pp. 1034–1047, Jul. 2001.

[5] G. Kowadlo and R. A. Russell, "Robot odor localization: A taxonomy and survey," *Int. J. Robot. Res.*, vol. 27, no. 8, pp. 869–894, Aug. 2008.

[6] V. H. Bennetts, A. J. Lilienthal, P. P. Neumann, and M. Trincavelli, "Mobile robots for localizing gas emission sources on landfill sites: Is bio-inspiration the way to go?" *Frontiers Neuroeng.*, vol. 4, Jan. 2012, p. 20.

[7] A. J. Lilienthal, A. Loutfi, and T. Duckett, "Airborne chemical sensing with mobile robots," *Sensors*, vol. 6, no. 11, pp. 1616–1678, Nov. 2006.

[8] A. R. Jones and D. J. Thomson, "Simulation of time series of concentration fluctuations in atmospheric dispersion using a correlation-distortion technique," *Boundary-Layer Meteorol.*, vol. 118, no. 1, pp. 25–54, Jan. 2006.

[9] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Distributed odor source localization," *IEEE Sensors J.*, vol. 2, no. 3, pp. 260–271, Jun. 2002.

[10] S. Pashami, A. J. Lilienthal, and M. Trincavelli, "Detecting changes of a distant gas source with an array of MOX gas sensors," *Sensors*, vol. 12, no. 12, pp. 16404–16419, 2012.

[11] H. Ishida, T. Ushiku, S. Toyama, H. Taniguchi, and T. Moriizumi, "Mobile robot path planning using vision and olfaction to search for a gas source," in *Proc. IEEE Sensors*, Oct./Nov. 2005, p. 4.

[12] H. Ishida, G. Nakayama, T. Nakamoto, and T. Moriizumi, "Controlling a gas/odor plume-tracking robot based on transient responses of gas sensors," *IEEE Sensors J.*, vol. 5, no. 3, pp. 537–545, Jun. 2005.

[13] S. Pang and J. A. Farrell, "Chemical plume source localization," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 36, no. 5, pp. 1068–1080, Oct. 2006.

[14] M. K. Muezzinoglu, A. Vergara, R. Huerta, N. Rulkov, M. I. Rabinovich, A. Selverston, and H. D. I. Abarbanel, "Acceleration of chemo-sensory information processing using transient features," *Sens. Actuators B, Chem.*, vol. 137, no. 2, pp. 507–512, Apr. 2009.

[15] S. M. Kay, *Fundamentals of Statistical Signal Processing*. Upper Saddle River, NJ, USA: Prentice-Hall, 1993.

[16] M. Schmuker, V. Bahr, and R. Huerta, "Exploiting plume structure to decode gas source distance using metal-oxide gas sensors," *Sens. Actuators B, Chem.*, vol. 235, pp. 636–646, Nov. 2016.

[17] K. A. Justus, J. Murlis, C. Jones, and R. T. Cardé, "Measurement of odor-plume structure in a wind tunnel using a photoionization detector and a tracer gas," *Environ. Fluid Mech.*, vol. 2, nos. 1–2, pp. 115–142, Jun. 2002.

[18] D. Martinez and E. M. Moraud, "Reactive and cognitive search strategies for olfactory robots," *Neuromorphic Olfaction*, vol. 5, pp. 153–172, Apr. 2013.

[19] D. J. Wheeler, *Understanding Statistical Process Control*. Knoxville, TN, USA: SPC Press, 1992.

[20] J. Burgués, V. Hernández, A. J. Lilienthal, and S. Marco, "Smelling nano aerial vehicle for gas source localization and mapping," *Sensors*, vol. 19, no. 3, p. 478, 2019.

[21] G. Korotcenkov and B. K. Cho, "Instability of metal oxide-based conductometric gas sensors and approaches to stability improvement (short survey)," *Sens. Actuators B, Chem.*, vol. 156, no. 2, pp. 527–538, 2011.

[22] I. T. Young and L. J. van Vliet, "Recursive implementation of the Gaussian filter," *Signal Process.*, vol. 44, no. 2, pp. 139–151, 1995.

[23] S. W. Roberts, "Control chart tests based on geometric moving averages," *Technometrics*, vol. 1, no. 3, pp. 239–250, 1959.

[24] J. Burgués, L. F. Valdez, and S. Marco, "High-bandwidth e-nose for rapid tracking of turbulent plumes," in *Proc. IEEE Int. Symp. Olfaction Electron. Nose (ISOEN)*, May 2019, pp. 1–3.

[25] A. Vergara, J. Fonollosa, J. Mahiques, M. Trincavelli, N. Rulkov, and R. Huerta, "On the performance of gas sensor arrays in open sampling systems using inhibitory support vector machines," *Sens. Actuators B, Chem.*, vol. 185, pp. 462–477, Aug. 2013.

[26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Oct. 2011.

JAVIER BURGUÉS received the bachelor's degree in telecommunication engineering from the University Autónoma of Madrid, in 2010, the master's degree in computer science from the University of Southern California, in 2013, and the Ph.D. degree in engineering and applied sciences from the University of Barcelona, in 2019. He is currently pursuing the Ph.D. degree in smart chemical instrumentation with the Institute for Bioengineering of Catalonia (IBEC), where he has been working under the advisement of Dr. S. Marco, since 2014. His main research interests include the applications of signal processing and pattern recognition techniques to chemical sensor data, integration of chemical sensors into robotic platforms, and the development of bioinspired flight algorithms for localization and mapping of chemical sources.

SANTIAGO MARCO received the University degree in applied physics, in 1988, and the Ph.D. degree in microsystem technology from the University of Barcelona, in 1993. He held a European Human Capital Mobility Grant for a postdoctoral position at the Department of Electronic Engineering, University of Rome Tor Vergata. Since 1995, he has been an Associate Professor of electronic instrumentation with the Department of Electronics, University of Barcelona. In 2004, he had a sabbatical leave at EADS Corporate Research, Munich, working on ion mobility spectrometry. In 2008, he was appointed as the Leader of the Signal and Information Processing for Sensing Systems Lab, Institute of Bioengineering of Catalonia. His research interest includes the development of signal/data processing algorithmic solutions for smart chemical sensing based in sensor arrays or microspectrometers integrated typically using microsystem technologies.

• • •