A Kalman Filter Based Approach to Probabilistic Gas Distribution Mapping

Jose Luis Blanco*, Javier G. Monroy†, Javier González-Jimenez†, and Achim Lilienthal‡
*Dept. of Engineering, University of Almería, Spain
†Dept. of System Engineering and Automation, University of Málaga, Spain
‡AASS Research Centre, Örebro University, Sweden

Abstract—Building a model of gas concentrations has important industrial and environmental applications, and mobile robots on their own or in cooperation with stationary sensors play an important role in this task. Since an exact analytical description of turbulent flow remains an intractable problem, we propose an approximate approach which not only estimates the concentrations but also their variances for each location. Our point of view is that of sequential Bayesian estimation given a lattice of 2D cells treated as hidden variables. We first discuss how a simple Kalman filter provides a solution to the estimation problem. To overcome the quadratic computational complexity with the mapped area exhibited by a straighforward application of Kalman filtering, we introduce a sparse implementation which runs in constant time. Experimental results for a real robot validate the proposed method.

Index Terms—alman Filter, Gas Distribution Mapping, Mobile Olfactionalman Filter, Gas Distribution Mapping, Mobile OlfactionK

I. Introduction

Modeling the gas distribution of an environment implies deriving a truthful representation of the observed gas distribution from a set of spatially and temporally distributed measurements of relevant variables, foremost gas concentration, but also wind, pressure, and temperature, for example. Building gas distribution models (GDM) is a challenging task, mainly because in many realistic scenarios gas is dispersed by turbulent advection, which creates packets of gas that follow chaotic trajectories [16].

While an exact description of turbulent flow remains an intractable problem, it is possible to approach the problem by aiming at a representation of the *average* gas distribution. A gas distribution model should therefore represent an estimate of the time-averaged concentration and the statistics of the expected fluctuations. In this sense, a gas distribution model is *truthful* if it explains new observations well and allows to identify hidden parameters such as the location of the source of gas, for example.

Instead of trying to solve the physical equations governing gas distribution, we create a statistical model of the observed gas distribution from the sparse set of measurements, treating gas sensor readings as random variables. Under the assumption of a *static* gas distribution and given a sufficient number of measurements, such a description will provide a truthful representation.

Gas distribution modeling constitutes an ideal application area for mobile robots since they can carry out the required repetitive measurement procedure without suffering from fatigue, can provide a higher (and adaptive) resolution of the distribution model than a stationary sensor network and offer the required accurate localization, the capabilities to create the gas distribution model on-line and to decide based on this model which locations need to be observed next. Mobile robots equipped with gas sensors, as our robot shown in Fig. 1, have a great potential
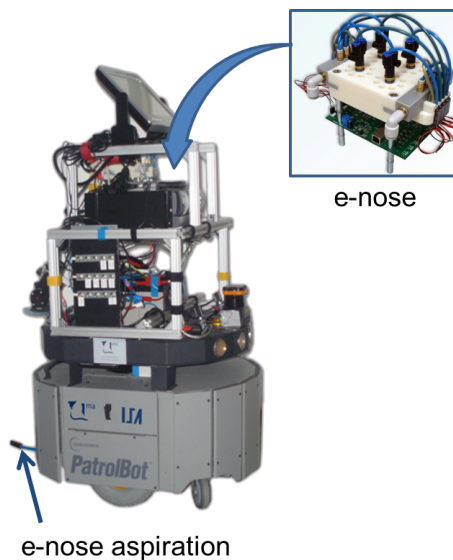

Fig. 1. A picture of the e-nose mounted on the mobile robot employed in the experiments.

for pollution monitoring in public areas [1] or inspection of hazardous industrial facilities.

In this work we introduce a probabilistic method to learn a gas distribution model of planar environments given a sequence of localized gas sensor readings, that is, we assume that robot localization is either solved or decoupled from gas mapping as in [11], a common practice in the mobile robotics olfaction literature. The space is divided into a two-dimensional lattice where cells are treated as hidden variables to be estimated through sequential Bayesian estimation. As discussed later on, the simple sensor model we propose allows the derivation of an efficient implementation of a Kalman Filter, such that updates can be performed in constant time. The estimated model, also referred as the *map* of gas concentrations, keeps a density distribution of the expected concentration at each cell, including its uncertainty. We present experimental results with a mobile robot equipped with an electronic nose that validate our proposal.

This paper is organized as follows. We first discuss related works in Section II, then we introduce the map model in Section III and derive the probabilistic method in Section IV. Finally, experimental results for a dataset gathered by a real mobile robot are discussed.

II. Related research

This section gives an overview of the work in the area of gas distribution mapping with a particular focus on methods that have been developed for mobile robots.

According to the assumptions about the nature of the gas distribution, these methods can be classified as model-based approaches or statistical approaches. Model-based approaches as in Ishida *et al.* [7] infer the parameters of an analytical gas distribution model from the measurements. They crucially depend on the underlying model. As discussed above, the application of complex numerical models based on fluid dynamics simulations is not feasible in practical situations. Simpler analytical models, on the other hand, often rest on rather unrealistic assumptions and are of course only applicable for situations in which the model assumptions hold.

Among the statistical approaches, histogram methods take the spatial correlation of concentration measurements into account because of the implicit extrapolation on the measurements by the quantization into histogram bins. Hayes *et al.* [6] suggest a two-dimensional histogram where the bins contain the accumulated number of "odor hits" received in the corresponding area. Odor hits are counted whenever the response level of a gas sensor exceeds a defined threshold. In addition to the dependency of the gas distribution map on the selected threshold, a problem with using only binary information from the gas sensors is that much useful information about fine gradations in the average concentration is discarded. A further disadvantage of histogram methods for gas distribution modeling is their dependency on the bin size and that they require perfectly even coverage of the inspected area.

Kernel extrapolation gas distribution mapping, which can be seen as an extension of histogram methods, was introduced by Lilienthal and Duckett [9]. Spatial integration is carried out by convolving sensor readings and modeling the information content of the point measurements with a Gaussian kernel. The Kernel extrapolation method was extended for the case of multiple odor sources [14] and it was demonstrated how a post-processing step, in which the obtained map is interpreted by an analytical physical model, allows to locate the gas source with a higher certainty and accuracy [13]. It was further shown on the basis of the Kernel extrapolation method how gas distribution mapping methods can be embedded into a Blackwellized particle filter approach to account for the uncertainty about the position of the robot [11].

All the methods mentioned so far model the average or the peak gas concentration but not the concentration fluctuations, or *variance*. The probabilistic model introduced in this paper estimates both parameters for each location (concretely, for each cell of the grid). Estimating the predictive variance is important for techniques that suggest new measurement locations based on the current model (sensor planning), for evaluating the model quality in terms of the data likelihood and for integrating the predictions into probabilistic localization methods [2]. Additionally, the Bayesian estimation of the variance proposed in this paper allows taking into consideration a *transition model* of the system, providing a promising tool to model certain instances of the GDM problem in the presence of wind.

Another method which predicts the mean concentration and its uncertainty using Gaussian process mixture models ("GPMM") was presented by Stachnis *et al.* [17]. The proposed method treats gas distribution modeling as a regression problem. In contrast to the approach introduced here, the model is represented directly using the training data. Since it requires the inversion of matrices growing with the number of training samples n, the computational complexity for learning the GPMM is $\mathcal{O}(n^3)$, while the sparse Kalman filter implementation introduced later on achieves a constant update time per observed measurement.

More recently, Lilienthal *et al.* [12] proposed the Kernel DM+V algorithm to estimate in addition to the distribution mean, the predictive variance per grid cell. They carried out two parallel estimation process, one for the mean and another for the variance, with the aim to adapt to the real variability of gas readings. The method proposed in this paper is based on a Bayesian interpretation, providing the covariance of the mean gas concentration as an estimate of the variance at each grid cell. As mentioned above, a remarkable advantage of the Kalman Filter-based mapping with respect to previous proposals is its potential for integrating in the gas mapping process a *transition model* that accounts for environmental information such as wind. This *transition model* for the gas concentration map is not addressed here and remains as future work.

III. A STOCHASTIC MODEL FOR GDM

As in most previous works, we simplify the problem of estimating the gas concentration in an environment by estimating a two-dimensional map. A map m is modeled as a random field where m_{xy} are scalar variables representing the gas concentration at coordinates (x, y).

In this work we propose a very simple probabilistic model for gas measurements: an observation z_t taken by the robot at time t is simply the actual value of the gas concentration at that point of space (denoted here as m_c), corrupted by additive Gaussian noise of variance σ_n^2, that is:

$$z_t = m_c + n_t \quad , \quad n_t \sim \mathcal{N}\left(0, \sigma_n^2\right) \tag{1}$$

The model is backed up by the physical principle of gas sensors which indeed are limited by a few square millimeters of sensing surface. Nevertheless, in practice the slow reaction time of sensors leads to an "averaging" effect over time. This effect can be reduced by forcing the robot to move very slowly or by using a specific e-nose configuration [5], [15].

Going back to the map model, and given analytical solutions are intractable, we divide the space into a regular lattice of *cells*. Since our aim is a probabilistic model of gas concentrations, the probability density function (pdf) of the concentration will be estimated at each cell within this gas concentration grid (GCG). In principle, this map model resembles occupancy grid (OG) maps [3] used for sonar or laser mapping. However, two fundamental properties set GCGs apart from OGs:

- In an OG, each cell is uniquely characterized by the *discrete* property of occupancy, thus each cell is modelled through a Bernoulli distribution. In contrast, the property we are modeling in a gas map is concentration, a *continuous* variable. Thus, we propose to model the density of cells as Gaussians.
- Many common sensors provide information about a much larger portion of space in comparison to gas sensors. This is the reason why assuming independence between cells is a common and plausible approach to building OGs [3] (a notable exception is [18]): several cells are observed simultaneously, while a gas sensor takes just one reading of a point. Motivated by this observation, the present approach does not assume independence between cells which would lead to an severe lack of information about locations the robot has not visited yet. Moreover, assuming certain correlations between neighbor cells has a clear foundation in the way gasses spread through an environment, thus the assumption of cell dependency arises naturally in gas mapping.

To summarize our model, we represent the map of gas concentrations, m, as a multidimensional Gaussian distribution,

$$m \sim \mathcal{N}\left(\mu, \Sigma\right) \tag{2}$$

where the mean vector $\mu = \{\mu_i\}_{i=1}^N$ keeps the average concentration for each of the N cells, and the $N \times N$ matrix Σ is the full covariance matrix. Thus, each cell m_{xy} is individually
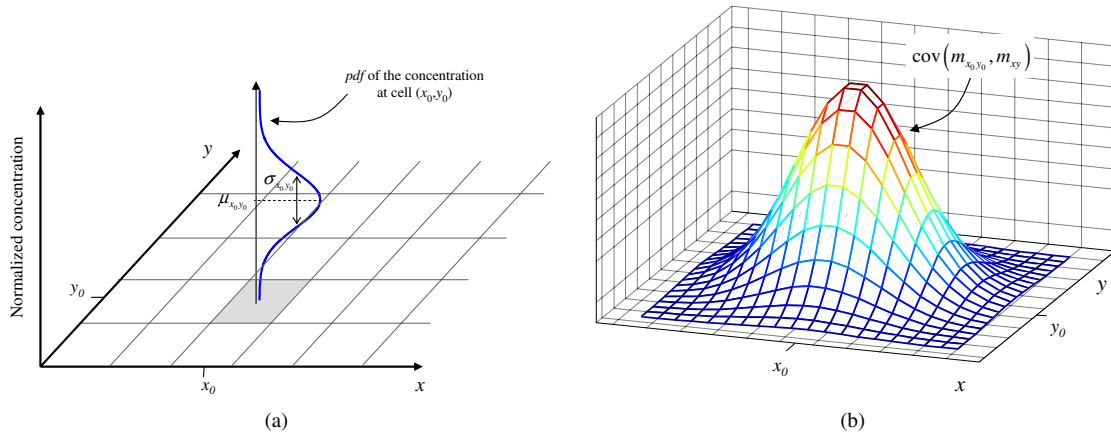
Fig. 2. (a) The 2D map is represented by a lattice where each cell keeps the estimate of gas concentration by means of a Gaussian density, represented here in the vertical axis. (b) We also estimate the covariance between cells and their neighbors. This plot shows the initial value of this covariance, representing the assumption that closer cells have more similar concentrations.
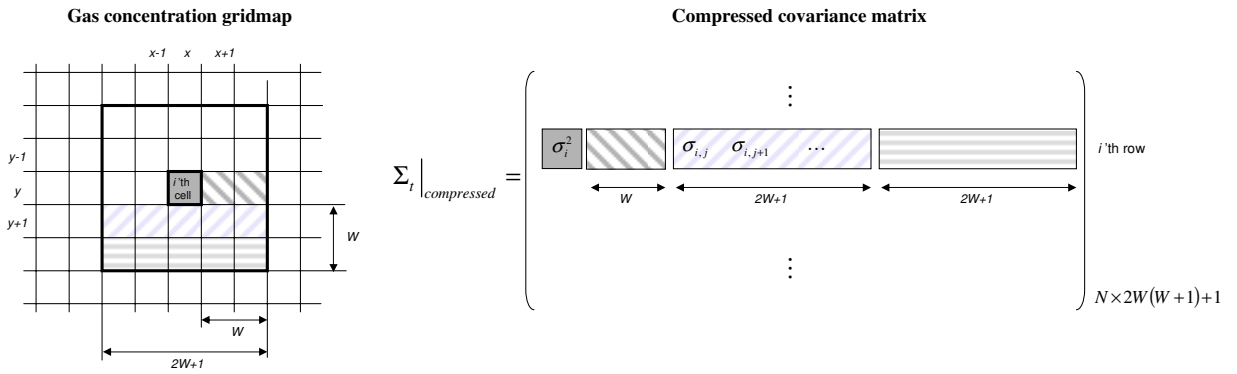

Fig. 3. Our sparse Kalman filter approach takes into account correlations between one cell and those within a given range W only (left). By assuming that the rest of covariance values are null, only a subset of the full $N \times N$ covariance matrix must be kept, the "compressed" covariance, its structure being depicted at the right. Note that only half of the covariances ($\sigma_{i,j}$ iff $i \leq j$) must be kept due to the symmetry of Σ_t.

modelled by $\mathcal{N}\left(\mu_{xy}, \sigma_{xy}^2\right)$, as depicted in Fig. 2(a), in addition to the covariances with the other cells. For convenience, mean values are *normalized gas concentrations* in the range $[0, 1]$.

The initialization of the covariance matrix Σ is the only point in our method where a physical model of gas distribution enters. Inspired by the kernel-based gas mapping algorithm [10], closer cells are assigned higher correlations which is modeled by an isotropic 2D gaussian as depicted in Fig. 2(b). The initial variance of each cell is set to a value larger than the range of normalized concentrations, e.g. $\sigma_{xy}^2 = 3$, such as the Gaussian approximates a uniform distribution for each unobserved cell.

The notation used in the rest of the paper deserves further discussion. Referring to cells by their 2-d coordinates (e.g. μ_{xy}) is useful for visualizing the spatial arrangement of cells. Nevertheless, when dealing with the state vector-covariance matrix representation it becomes more convenient to denote individual cells by a single index, e.g. μ_c, thus the variance of a given cell c is denoted by σ_c^2 and cross-covariances by $\sigma_{i,j}$.

IV. DERIVATION

In this section we first introduce Bayesian estimation of gas concentration grids using Kalman filtering, then we describe how a sparse representation of covariances leads to a dramatic improvement in efficiency resulting in a constant time complexity (i.e. independent of the map size).

A. Kalman filtering

Let μ_t and Σ_t denote the mean and covariance matrix of the map estimate at time step t. Kalman filtering (KF) [8] allows us to incrementally update this estimate given new evidences, or observations z_t. Without loss of generality, we assume that only one gas sensor is read at each time. In the case of robots with several sensors, the following equations are just applied more than once sequentially.

Since the sensor model proposed in this work is linear, see Eq. (1), the original KF algorithm suffices to our problem. Here the updated mean vector is computed as:

$$\mu_t = \mu_{t-1} + \mathbf{K}_t \left(z_t - \mathbf{H}_t \mu_{t-1} \right) \tag{3}$$

where $\mathbf{H}_t$ is the $1 \times N$ matrix of the observation model. In our problem, this matrix has a special configuration which will lead to important simplifications:

$$\mathbf{H}_t = \begin{pmatrix} 0 & \cdots & 1 & \cdots & 0 \end{pmatrix} \tag{4}$$
$$\mathbf{H}_t[k] = \begin{cases} 1 & , & k = c \quad \text{(the current cell)} \\ 0 & , & k \neq c \end{cases}$$

that is, all entries but one are zero due to the simple sensor model in Eq. (1) which states that the sensor measures the gas concentration at the current cell c. It follows then that the Kalman gain, $\mathbf{K}_t$ in Eq. (3), becomes:

$$\mathbf{K}_t = \mathbf{\Sigma}_{t-1}\mathbf{H}_t^\top \left(\mathbf{H}_t\mathbf{\Sigma}_{t-1}\mathbf{H}_t^\top + \sigma_n^2\right)^{-1} \qquad (5)$$

$$= \mathbf{\Sigma}_{t-1}\mathbf{H}_t^\top \left(\sigma_c^2 + \sigma_n^2\right)^{-1}$$

$$= \begin{bmatrix} \sigma_{1,c} \\ \vdots \\ \sigma_{N,c} \end{bmatrix} \left(\sigma_c^2 + \sigma_n^2\right)^{-1}$$

which leads to the following update rules for each i'th cell's mean and cross-covariances:

$$\mu_i' = \mu_i + (z_t - \mu_c)\frac{\sigma_{i,c}}{\sigma_c^2 + \sigma_n^2} \qquad (6)$$

$$\sigma_{i,j}' = \sigma_{i,j} - \frac{\sigma_{i,c}\ \sigma_{c,j}}{\sigma_c^2 + \sigma_n^2} \qquad (7)$$

where for clarity prime variables refer to time t and unprimed variables to time $t-1$. This formulation provides us with an exact solution to gas mapping given our sensor model. However, its computation demands $\mathcal{O}(N)$ and $\mathcal{O}(N^2)$ operations for the mean and the covariance matrix, respectively, being N the number of cells. This computational burden is better revealed by noting that the number of cells N grows linearly with the mapped *area*, thus the method has an overall complexity of $\mathcal{O}(A^2)$ for A being the mapped area. Storage is another drawback since keeping all the covariances also demands quadratic space with respect to the map area. Therefore, the method above can be directly applied only to small maps. We develop in the following an optimized version of the algorithm which dramatically reduces the computational and storage complexities.

B. Sparse implementation

While in landmark-based SLAM sparse filters are well-known and exploit real independencies between far-off landmarks [19], in a gas grid any cell has some degree of correlation with its neighbor, vanishing quickly with distance as illustrated in Fig. 2(b).

Our proposed implementation of Kalman filtering for GDM hence relies on the truncation of covariances between any cell and those ones out of a window centered at the current cell, as depicted in Fig. 3. The window size, W, determines the range of cells affected by a gas reading at some given location. Thus, for each new measurement, the mean of all cells is updated using (6) and only some covariances are updated through (7).

The advantages of this sparse representation are twofold. Firstly, the complexity of each update is reduced from $\mathcal{O}(N^2)$ (determined by the update of covariances) to $\mathcal{O}(W^4)$, which represents a great improvement given that N is typically several orders of magnitude larger than W. Note also that the complexity becomes *independent* of the actual mapped area. Secondly, storage requirements for the covariances also greatly decrease from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \cdot W^2)$. One possible arrangement is the "compressed" covariance layout pictured in Fig. 3. As an example, a full covariance matrix for a real gas gridmap is shown in Fig. 4(a). It can be clearly observed how all elements but those in a diagonal band are zero, hence they do not contribute useful information to the filter. In the compressed matrix in Fig. 4(b), only the band diagonal elements are kept, thus all the information is preserved while requiring a fraction of the memory. The exact gain in memory depends on the map size and the value of W, but the improvement grows with the size of maps.

V. EXPERIMENTAL RESULTS

To validate our proposal, we have carried out the following experiments[1]. A mobile robot equipped with an "electronic nose"

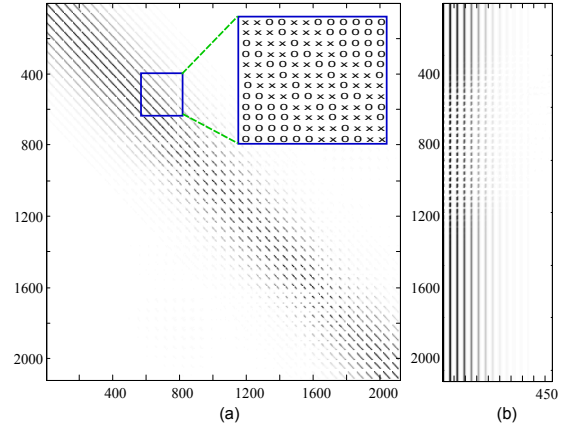[1] A video is available at YouTube in www.*hidden for blind revision*

Fig. 4. (a) A covariance matrix for a gas gridmap, where it becomes obvious that most of the correlations occur between very close cells, thus virtually all the information is kept in a band diagonal (darker colors mean stronger correlations). (b) The compressed covariance matrix proposed in this work, where one row exists for each cell and contains the relevant covariances only.

(e-nose) (see Fig. 1) was guided through an office room with an alcohol source (a cup) placed on the floor in the middle of the room. All doors and windows in the room were kept shut during the experiment to prevent strong air flows. Inside the e-nose, four different Figaro sensors [4] provide us in parallel the gas concentration of different chemicals.

Both the robot path and the occupancy grid map obtained as the robot performed SLAM to localize itself are represented in Fig. 5(a). The gas readings collected as the robot moves are plotted in Fig. 5(b). After applying the straighforward implementation of Kalman filtering (with the full $N \times N$ covariance matrix), we obtain an estimated map where the peak roughly coincides with the actual location of the gas source. The mean and standard deviation of each cell in the gas grid can be seen in Fig. 5(c)-(d), respectively. The level of uncertainty associated to each cell quickly increases with the distance to the actual robot path, as can be seen in Fig. 5(d).

As could be expected, the mean of cells is modified far beyond the robot path, thus the approach is successful in interpolating the gas readings to locations not visited by the robot. The parameter σ_d, which controls the influence area of measurements, has been set to 30cm in this experiment – a value comparable to those in Kernel-based GDM. This value of σ_d has been determined by optimizing the observations likelihood for the present dataset, although values approximately in the range $25 - 50cm$ leads to sensible maps. It must be noted that the optimal parameters are determined manually for each dataset, thus a more concise analysis of optimal configurations across several environments remains being a future work.

Regarding the sparse implementation described in section IV-B, we have observed that there exists a minimum window size (W) which leads to an acceptable approximation of the full covariance implementation, though for greater W values the approximated maps converge very quickly to the exact one at a fraction of the computational cost. To quantify the improvement we have applied the sparse KF with a range of different W values. The results are summarized in Fig. 5(e)-(f), where the average errors in both the mean and the variance of cells, respectively, are plotted against increasing sizes of the window for values in the range $14 \leq W \leq 26$. An accurate approximation is obtained for values of $W \geq 20$, approximately, while the time required to build the entire map remains around 20 seconds – refer to Fig. 5(g) – in contrast to the consideration of the full covariance

matrix which requires more than 130 seconds. It must be noted that the robot took an overall time of 768 seconds to collect the dataset, thus both methods are capable of real-time mapping.

Naturally, these performance results are related to the cell size, set in this experiment to $c = 10cm$. In general, a finer grid will provide more accurate results, at the expense of a greater computational time. This burden is derived from the need to enlarge the window size W to compensate the smaller size of cells. As a rule of thumb from the results in Fig. 5(e)-(f), the window size should be $K \gtrsim 7\sigma_d/c$. Hence the convenience of keeping the grid as coarse as possible. Typically, good results can be obtained with cell sizes c in the range $5 - 50cm$.

VI. Conclusions

In this paper we have revisited the problem of map building for the case of gas concentrations. We have approached the problem from a Bayesian perspective and employed an optimized version of Kalman filtering to generate a model of the gas distribution in a planar environment. The main contribution of this work is the introduction of a fast, probabilistic algorithm which considers uncertainty in gas maps, and provides the mathematical background for integrating in the gas mapping process a *transition model* that accounts for environmental information such as wind. This *transition model* is not addressed here and remains as future work. The method has been validated with a real dataset and despite the noisy measurements, the obtained map correctly reflects a peak in the concentration at the approximate location of the source. Due to its probabilistic nature, the proposed approach is compatible with localization and SLAM methods relying, uniquely or partly, on gas sensors. Future works will explore these possibilities.

VII. Acknowledgments

References

[1] DustBot - Networked and Cooperating Robots for Urban Hygiene. http://www.dustbot.org.

[2] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte Carlo localization for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 2, 1999.

[3] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989.

[4] Figaro. Figaro corporate website: http://www.figarosensor.com/.

[5] J. Gonzalez-Jimenez, J. G. Monroy, and J. L. Blanco. The multi-chamber electronic nose, an improved olfaction sensor for mobile robotics. *Sensors*, 11(6):6145–6164, 2011.

[6] A. Hayes, A. Martinoli, and R. Goodman. Distributed Odor Source Localization. *IEEE Sensors Journal, Special Issue on Electronic Nose Technologies*, 2(3):260–273, 2002. June.

[7] H. Ishida, T. Nakamoto, and T. Moriizumi. Remote Sensing of Gas/Odor Source Location and Concentration Distribution Using Mobile System. *Sensors and Actuators B*, 49:52–57, 1998.

[8] R. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

[9] A. Lilienthal and T. Duckett. Building Gas Concentration Gridmaps with a Mobile Robot. *Robotics and Autonomous Systems*, 48(1):3–16, August 2004.

[10] A. J. Lilienthal and T. Duckett. Building gas concentration gridmaps with a mobile robot. *Robotics and Autonomous Systems*, 48(1):3–16, August 31 2004.

[11] A. J. Lilienthal, A. Loutfi, J. L. Blanco, C. Galindo, and J. Gonzalez. A rao-blackwellisation approach to gdm-slam. integrating slam and gas distribution mapping. In *Proceedings of the European Conference on Mobile Robots (ECMR)*, pages 126–131, September 19–21 2007.

[12] A. J. Lilienthal, M. Reggente, M. Trincavelli, J. L. Blanco, and J. Gonzalez. A statistical approach to gas distribution modelling with mobile robots, the kernel dm+v algorithm. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 570–576, October 11 – October 15 2009.

[13] A. J. Lilienthal, F. Streichert, and A. Zell. Model-based Shape Analysis of Gas Concentration Gridmaps for Improved Gas Source Localisation. pages 3575 – 3580, Barcelona, Spain, 2005.

[14] A. Loutfi, S. Coradeschi, A. J. Lilienthal, and J. Gonzalez. Gas Distribution Mapping of Multiple Odour Sources using a Mobile
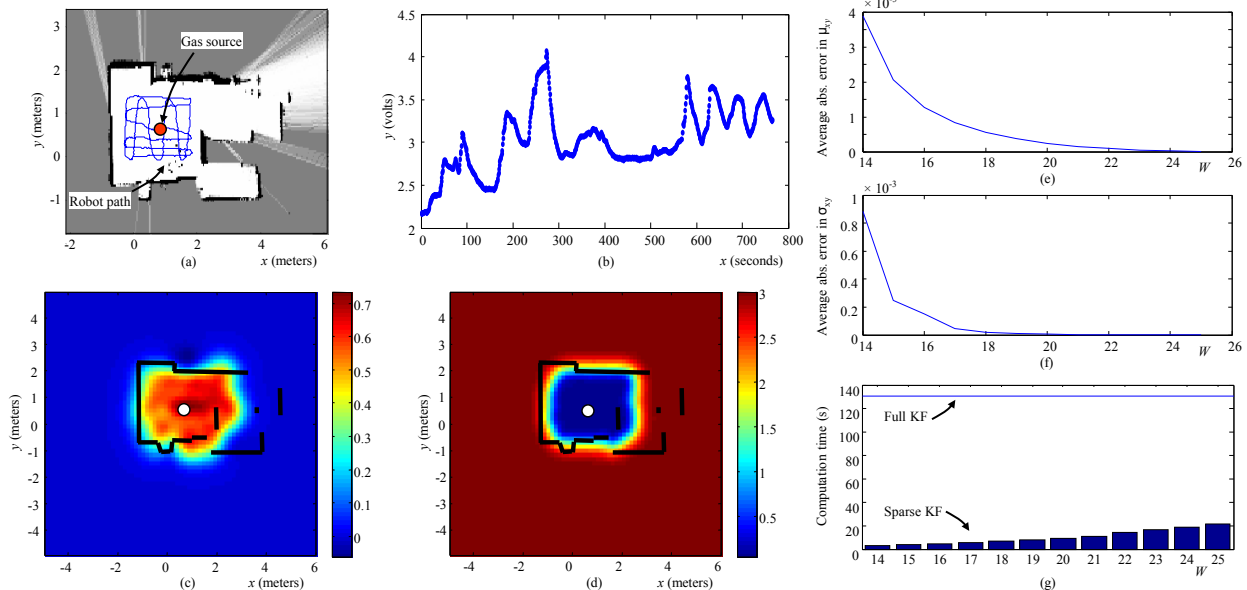
Fig. 5. Experimental results for gas readings gathered by a real mobile robot. (a) The robot path and the occupancy grid of the office where the experiment took place. (b) The gas readings over time. (c)-(d) The mean and standard deviation, respectively, of each cell in the gas concentration grid built with our method. (e)-(f) The errors between the exact KF solution and the sparse filter for increasing sizes of the window W, and (g) the corresponding computation times.

Robot. *Robotica*, June 4 2008. Published online by Cambridge University Press.

[15] J. G. Monroy, J. Gonzalez-Jimenez, and J. L. Blanco. Overcoming the slow recovery of mox gas sensors through a system modeling approach. *Sensors*, Manuscript submitted for publication.

[16] B. Shraiman and E. Siggia. Scalar Turbulence. *Nature*, 405:639–646, 8 June 2000. Review Article.

[17] C. Stachniss, C. Plagemann, A. Lilienthal, and W. Burgard. Gas Distribution Modeling Using Sparse Gaussian Process Mixture Models. In *Robotics: Science and Systems (RSS)*, Zurich, Switzerland, June 2008.

[18] S. Thrun. Learning Occupancy Grid Maps with Forward Sensor Models. *Autonomous Robots*, 15(2):111–127, 2003.

[19] M. Walter, R. Eustice, and J. Leonard. Exactly Sparse Extended Information Filters for Feature-based SLAM. *International Journal of Robotic Research*, 26(4):335–359, 2007.