Time-Variant Gas Distribution Mapping with Obstacle Information

Javier G. Monroy · Jose-Luis Blanco · Javier Gonzalez-Jimenez

Abstract **This paper addresses the problem of estimating the spatial distribution of volatile substances using a mobile robot equipped with an electronic nose (e-nose). Our work contributes an effective solution to two important problems that have been disregarded so far: First, obstacles in the environment (walls, furniture, ...) do affect the gas spatial distribution. Second, when combining odor measurements taken at different instants of time, their 'ages' must be taken into account to model the ephemeral nature of gas distributions. In order to incorporate these two characteristics into the mapping process we propose modeling the spatial distribution of gases as a Gaussian Markov random field (GMRF). This mathematical framework allows us to consider both: (i) the vanishing information of gas readings by means of a time-increasing uncertainty in sensor measurements, and (ii) the influence of objects in the environment by means of correlations among the different areas. Experimental validation is provided with both, simulated and real-world datasets, demonstrating the out-performance of our method when compared to previous standard techniques in gas mapping.**

J. G. Monroy
Department of System Engineering and Automation, University of Málaga, Tel.: +34-952-13-2747, E-mail: jgmonroy@uma.es

J.L. Blanco
Department of Engineering, University of Almería

J. Gonzalez-Jimenez
Department of System Engineering and Automation, University of Málaga

1 Introduction

Gas distribution mapping (GDM) is the process of creating a representation of how gases spread in an environment from a set of spatially and temporally distributed measurements of relevant variables [Blanco et al, 2013; Lilienthal et al, 2009]. Foremost, these measurements include the gas concentration itself, but may also comprise wind, pressure or temperature.

In the last decade, GDM is gaining attention in the mobile robotics community because of the advantages a mobile robot offers when compared with the traditional approach based on networks of static gas sensors [Tsujita et al, 2005; Fenger, 1999]. To start with, a mobile robot usually carries only one but more expensive and powerful gas sensing device (e-nose), which can analyze more complex compounds [Airsense Analytics, 2014; Sanchez-Garrido et al, 2014; Sensigent Intelligent Sensing Solutions, 2014], and detect faster changes in the gas concentration [Gonzalez-Jimenez et al, 2011; G. Monroy et al, 2012]. Also, the robot can sample at a higher (and adaptive) resolution, while still providing the required accurate localization of each measurement. Moreover, the gas distribution map is created by a robot in an online fashion, allowing decision making to occur depending on such a map, e.g. for exploration tasks. Finally, a mobile robot can leverage environmental information provided by other sensors on board (cameras, laser scanners, etc.) to both enhance the GDM process itself, for example by detecting obstacles, as proposed in this paper, and to help in any other odor-related task, as

Javier G. Monroy, Jose-Luis Blanco, Javier Gonzalez-Jimenez
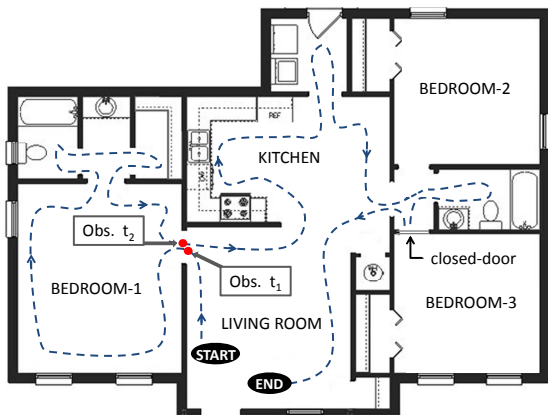
Fig. 1 An illustrative example where a robot is commanded to inspect an indoor environment by following the predefined path (the blue-dashed line). The robot gathers odor observations as it moves and builds a gas distribution map. Crucial aspects to be considered for such map building include how gas concentrations observed at different instants of time are combined (as for example, t_1 and t_2), and the influence of obstacles such as walls or furniture.

can be the identification of potential gas sources [Ishida et al, 2005].

Building a gas distribution model with a mobile robot turns out to be a tough problem for a number of reasons. First of all, and in contrast to most exteroceptive sensors employed in mobile robotics, an e-nose is a *point* sampling device[1], that is, it only samples the very near air around it. Furthermore, the dispersion of gases is strongly conditioned by the presence of large obstacles (i.e. objects) in the environment, such as walls and furniture. Consequently, when building a GDM, they should indeed be taken into consideration to yield accurate estimations. Finally, but not less important, odors are ephemeral due to the mechanisms that rule gas dispersion [Shraiman and Siggia, 2000] (with a strong predominance of advection and turbulence over molecular diffusion). Thus, we can say that the information conveyed by a given measurement quickly vanishes as time goes by.

The two latest points are pivotal characteristics that have been overlooked by previous works on GDM [Loufti et al, 2009; Blanco et al, 2013; Lilienthal et al, 2009; Stachniss et al, 2009; Turduev et al, 2014]. Traditionally, the influence of obstacles in the gas distribution has been only considered explicitly for the simulation of the gas dispersal [Pashami et al, 2010; Tauseef et al, 2011], and in some works on plume tracking [Marjovi and Mar-

ques, 2013]. However, when facing the GDM problem the information related to obstacles is neglected, leading to maps where nearby areas are always correlated, even when physical obstacles separate them. Moreover, existing approaches to GDM provide the same confidence to all gas measurements regardless of *when* they were taken. As a result, the estimated gas distribution averages out measurements taken at very different moments in time, something that strongly contradicts the vanishing nature of gases (odors).

The example in Figure 1, which illustrates the discussion above, shows a robot which is commanded to inspect the different rooms of a house to determine, for example, the possible existence of bad odors. As it moves, the robot collects new gas observations that are incorporated to the GDM. Occasionally the same place is revisited, thus samples from the same location must be somehow combined. Since time-separated observations are gathered at close locations (e.g. samples at t_1 and t_2), the GDM method has to deal with these questions: (i) are observations taken at t_1 and t_2 equally relevant?, and (ii) if not, how do their significances evolve over time? Furthermore, the GDM method has to settle how to account for the obstacles whose presence significantly alter the gas distribution. That is, to provide an estimation of the gas distribution based only on the observations collected along the covered path, or to additionally account for objects such as walls and furniture to model the correlation between the areas they separate (e.g. the presence of bad odors in the kitchen do not imply the same at bedroom-3, since its gas concentration can be considered 'independent' given the presence of walls and the closed door).

We can summarize the two contributions of the present work as follows. First, we propose accounting for the obstacles in the environment, obtaining maps which are more compliant with the mechanisms of gas dispersion. Secondly, we claim that the "age" of a measurement is of relevance in the GDM process. In particular, we propose to associate a time-decreasing weight to each gas measurement, modeling the fact that recent measurements represent the current gas distribution more faithfully than older ones. Thus, observations taken at the same locations and separated in time will be combined according to their respective weights.

As estimation tool for the gas mapping process, we propose to employ a Gaussian Markov random field (GMRF), which perfectly suits the characteristics of GDM by accounting not only for the information carried by the gas observations, but also for any *prior* knowledge which, in our case, includes both the obstacles in the environment (detected by the robot sensors), and the physics of how gases spatially distribute.

---

[1] An exception to this is the tunable laser absorption spectroscopy (TDLAS) technology [Trincavelli et al, 2012; Frish et al, 2005], which provides integral gas concentration measurements over the path of a laser beam. However, this is still an emerging technology with important drawbacks for robotic applications, including cost, weight and power consumption.

The rest of this paper is organized as follows. We first discuss the related literature on GDM with mobile robots in Section 2, to continue with the introduction of the proposed probabilistic model for GDM in Section 3. Then, we show how the maximum a posteriori (MAP) estimation becomes a sparse least squared problem in Section 4, and finally, in Section 5, we report simulated and real experimental results.

2 Related Work

We are interested in statistical modeling of gas distributions without making strong assumptions about the environmental conditions (temperature, pressure or airflows). Given that analytical solutions are intractable, it is common practice to divide the space into a regular lattice of cells (gridmap), and then estimate a probability density function (pdf) of the gas concentration at each cell of the grid. Under these circumstances, only a few gas distribution modeling methods have been proposed in the literature.

The most remarkable works in this field have been reported by Lilienthal and colleagues. In the pioneer work [Lilienthal and Duckett, 2004] they proposed the kernel-based method, which consists of convolving sensor readings with a Gaussian kernel [Bishop, 2007], thus providing a representation of the gas map without assuming any predefined parametric form for the distribution. This method was later extended for the case of multiple odor sources [Loufti et al, 2009] and to the three-dimensional case [Reggente and Lilienthal, 2009]. It was further shown 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 [Lilienthal et al, 2007].

More recently we find approaches that, in addition to providing the most-likely value for the gas distribution, also estimate the associate uncertainty (via a variance value). In [Stachniss et al, 2009], Stachniss *et al.* proposed an approach using Gaussian process mixture models (GPM), treating gas distribution modelling as a regression problem. The components of the mixture model and the gating function, that decides to which component a data point belongs, were learned using Expectation Maximization (EM). Later, in [Lilienthal et al, 2009], Lilienthal *et al.* carried out two parallel estimation processes, one for the mean and another for the variance, understanding the latter as the variability of the gas readings, not the uncertainty in the estimation which is the standard in probabilistic estimators. Results demonstrated that although providing similar maps to the obtained in [Stachniss et al, 2009], this method had the advantage of scaling better to larger

training datasets and to possess a simpler learning procedure. In [Blanco et al, 2013], Blanco *et al.* proposed another approach, in this case based on a Bayesian interpretation of the problem, which also obtains the variance of each map cell employing a sparsified Kalman filter.

None of these works take into account the constraints imposed by the obstacles of the environment when estimating the gas distribution, neither the physical fact that the information provided by a gas sensor vanishes with time. The latter, however, was pointed out by Asadi *et.al* [Asadi et al, 2011], proposing an extension to the kernel method which incorporates an exponential time decay to the formulation. Nonetheless, to the best knowledge of the authors, no map estimator taking this into account was ever reported. Therefore, the GMRF-based approach proposed in this work exploits, for the first time, both concepts.

3 Modeling GDM as a Markov Random Field

In this section we introduce the basis for the estimation of the gas distribution over a 2D lattice of cells using GMRFs. We also describe the highly-sparse structure of the problem which leads to efficient estimates of the problem and, finally, we present our model for observation time-varying uncertainty.

3.1 Probabilistic Model for GDM

The proposed approach aims at estimating the probability density function of the gas concentration in an environment. As in most previous works on GDM we simplify the problem by estimating a discrete two-dimensional map, dividing the space into a rectangular lattice of *cells*. A map $\mathbf{m} = \{m_i\}_{i=1}^{N}$ is then modeled as a random field where m_i are scalar variables standing for the gas concentration inside the i'th cell with coordinates (x_i, y_i). Let N be the overall number of variables in the map, such that if the map is $N_x \times N_y$ cells, $N = N_x N_y$. Notice that this model resembles occupancy grids in robotics, with the difference of not holding a discrete distribution (occupied vs. free) but a continuous magnitude (see Figure 2) .

Our goal is to obtain the maximum a posteriori (MAP) estimation of $\mathbf{m}$, along with its uncertainty, given the gas concentrations measured by the robot e-nose (random variables $\mathbf{z}$) and some prior knowledge that includes (i) how the gas spreads over the environment, and (ii) how the perceived obstacles affect the propagation of gases between nearby cells. Given the
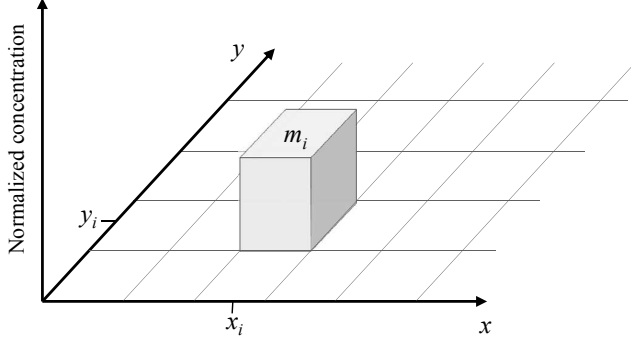
Fig. 2 The 2D map is represented by a lattice where each cell keeps the estimate of gas concentration, represented here along the vertical axis.

small space sampled by an e-nose (even when employing pumps or fans to aspire the air), this prior is extremely important for inferring the gas concentration at distant locations not subject to direct sensing.

Our proposal is to use a Markov random field (MRF), a tool widely employed in other estimation problems on grids. For example, in image processing, where statistical models are defined for the intensity of image pixels [Winkler, 2003]. Notice the strong analogy between problems such as image de-noising or image restoration and the GDM stated here, where gridmap cells play the role of pixels.

According to the *Hammersley-Clifford* theorem [Clifford, 1990], the joint probability distribution $p(\mathbf{m}, \mathbf{z})$ that we want to maximize can be expressed as a *Gibbs* distribution, that is, it can be factored as the product of the potential functions $\varphi(\cdot)$ for the set of all its maximal cliques ($\boldsymbol{\mathcal{C}}_m$) [Bishop, 2007]:

$$p(\mathbf{m}, \mathbf{z}) = \frac{1}{Z} \prod_{\mathcal{C} \in \boldsymbol{\mathcal{C}}_m} \varphi_\mathcal{C}(\mathbf{n}_\mathcal{C}) \qquad (1)$$

where the proportionality constant Z (called the *partition function*) is not relevant in our problem, $\mathcal{C}$ denotes the different cliques and $\mathbf{n}_\mathcal{C}$ the set of variables ($\mathbf{m}, \mathbf{z}$) in that clique.

Since we are restricted to potential functions which are strictly positive ($\varphi(\cdot) > 0$), it is convenient to express them as exponentials:

$$p(\mathbf{m}, \mathbf{z}) \propto \prod_{\mathcal{C} \in \boldsymbol{\mathcal{C}}_m} \exp\{-E(\mathbf{n}_\mathcal{C})\} \qquad (2)$$

$$= \exp\left\{-\sum_{\mathcal{C} \in \boldsymbol{\mathcal{C}}_m} E(\mathbf{n}_\mathcal{C})\right\}$$

where E is an energy function, obtained by adding up the energies of each of the maximal cliques.

An intuitive and convenient way of dealing with the dependencies encoded in a MRF is to consider its *factor graph* [Dellaert and Kaess, 2006; Loeliger, 2004] as the graphical model from which to derive the optimization equations. In this graphical model, each potential functions $\varphi(\cdot)$ over a maximal clique becomes a factor $\mathcal{F}$. As shown in Figure 3, this model comprises two kinds of nodes: (i) gas concentrations at cells (unknowns to be estimated), and (ii) gas observations (known data). We also define two distinct sets of factors between nodes: *observation factors* which represent sensor observations and constrain the concentration value of a cell i according to all sensor measurements taken by the robot at that cell, and *prior factors* which, being independent of observations, capture any a priori knowledge on how the gas distribution behaves over space.

Attending to the two different set of factors, the joint probability distribution can then be expressed as:

$$p(\mathbf{m}, \mathbf{z}) \propto \exp\left\{-\sum_{\mathcal{C}_o} E_o(\mathbf{n}_{\mathcal{C}_o}) - \sum_{\mathcal{C}_p} E_p(\mathbf{n}_{\mathcal{C}_p})\right\} \qquad (3)$$

3.2 Factor Parameters

We assume that all the conditional distributions involved in the problem can be reasonably modeled as Gaussians, thus the underlying graphical model becomes a Gaussian MRF (GMRF). This assumption works well in practice, as demonstrated experimentally. Therefore, we need to provide the parameters of each Gaussian distribution that appears in our graphical model in order to have it completely defined.
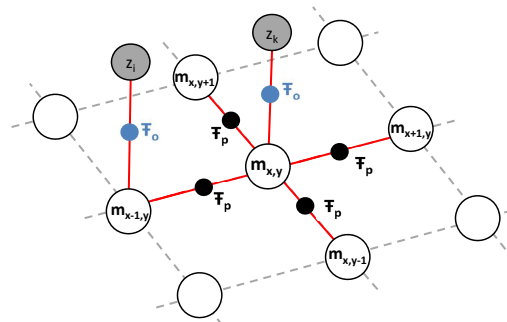
Fig. 3 Factor graph derived from the MRF employed in our approach. There are two types of nodes: gas concentrations at cells (white circles), and gas observations (grey shaded circles), and two kind of factors: prior factors ($\mathcal{F}_p$), and observation factors ($\mathcal{F}_o$).

3.2.1 *Observation Factors*

They encode the observation model, that is, the relationship between an e-nose reading and the true gas concentration of the cell at which it was taken.

Let M be the number of e-nose observations collected by the robot. Each observation consists of a gas concentration value z_k taken at a particular cell i_k at a given instant of time t_k, with $k = 1 \dots M$. Each such observation is assumed to be corrupted with two additive Gaussian errors: one from the inherent sensor noise ($\omega_k \sim \mathcal{N}(0, \sigma_s^2)$) and another time-dependent term that models the potential changes that may have occurred since the sensing time ($\zeta_k \sim \mathcal{N}(0, \sigma_\zeta^2(t - t_k))$). With this last noise, we model the loss of information of a measurement as an increase of uncertainty (variance). Denoting the ideal (noiseless) sensor model as $h(\mathbf{m})$, we have:

$$z_k = h(\mathbf{m}) + \omega_k + \zeta_k \qquad (4)$$

Given the moderately large size of grid cells in GDM (typically in the range of decimeters) it becomes reasonable to assume that every measurement is associated to one single cell, the one which the robot e-nose is sniffing at, which takes us to the minimalistic sensor model $h(\mathbf{m}) = h(m_{i_k}) = m_{i_k}$.

Under a probabilistic point of view, each observation factor in the graphical model then stands for the conditional pdf:

$$p(z_k|\mathbf{m}) = p(z_k|m_{i_k}) = \mathcal{N}(m_{i_k}, \sigma_s^2 + \sigma_\zeta^2(t - t_k)) \qquad (5)$$

where we have applied the conditional independence between z_k and the rest of the cells given m_{i_k}.

Then, the energy function associated to the observation factors can be expressed as:

$$E_o = \sum_{\mathcal{C}_o} E_o(\mathbf{n}_{\mathcal{C}_o}) = \sum_{k=1}^{M} \frac{(m_{i_k} - z_k)^2}{\sigma_s^2 + \sigma_\zeta^2(t - t_k)} \qquad (6)$$

The time-increasing variance $\sigma_\zeta^2(t - t_k)$ above means that if two observations from the same cell are combined to estimate its gas concentration, more weight is given to the most recent one. Eventually, during a GDM process the variance of older measurements will become large enough as to be neglected. Thus, in practice, only a finite set of M observations will account for the estimation, bounding the overall computational complexity of our method for a fixed-size map.

3.2.2 *Prior Factors*

These factors capture the knowledge about how gases distribute spatially. Particularly, we want to model the correlation between gas concentrations of neighboring cells. We insist in the necessity and relevance of this term in GDM, since gas observations from e-noses provide us very localized information, i.e. only for the cell at which the sample was taken.

Correlation between cells: Previous approaches based on kernel methods [Lilienthal et al, 2009] or Kalman filtering [Blanco et al, 2013] have modeled the correlations between observations and cells as Gaussian functions. Similarly, we model the correlation between cells by penalizing the difference in the gas concentration (l_{ij}) between pairs of (vertically and horizontally) adjacent cells:

$$l_{i,j} = m_i - m_j \qquad (7)$$

where m_i, m_j are the gas concentrations at adjacent cells with lattice indices i and j, respectively. Each prior factor then stands for the following pdf:

$$p(l_{i,j}|\neg o_{i,j}) = \mathcal{N}(0, \sigma_p^2) \qquad (8)$$

meaning that adjacent cells are forced to have the same gas concentration with a tolerance stated by σ_p^2. The meaning of the conditioning on $\neg o_{i,j}$ is explained next.

Obstacles: The influence in the gas distribution of the objects in the environment must be accounted for while modeling the expected difference between adjacent cells in Eq. (7). We assume that the probability of their intermediary space to be occupied, $P(o_{i,j}) \in [0, 1]$, is readily available in the form of an occupancy grid representation of the environment. Note that the cell size of this grid is not required to match that of the gas map and, in practice, will often be finer.

Denoting as $o_{i,j}$ the fact that an obstacle exists between cells i and j, we can then apply the law of total probability over the two only possibilities (either $o_{i,j}$ or $\neg o_{i,j}$) to obtain:

$$p(l_{i,j}) = p(l_{i,j}|o_{i,j})P(o_{i,j}) + p(l_{i,j}|\neg o_{i,j})P(\neg o_{i,j}) \qquad (9)$$

where $P(\neg o_{i,j}) = 1 - P(o_{i,j})$ and $p(l_{i,j}|\neg o_{i,j})$ was already given in Eq. (8). Regarding our a priori distribution for $l_{i,j}$ in the case of an obstacle blocking the way between two cells, it seems reasonable to assume no correlation at all. Then, a good candidate for $p(l_{i,j}|o_{i,j})$ is a uniform distribution over a sufficiently large interval. Since mixing Gaussians and uniform densities
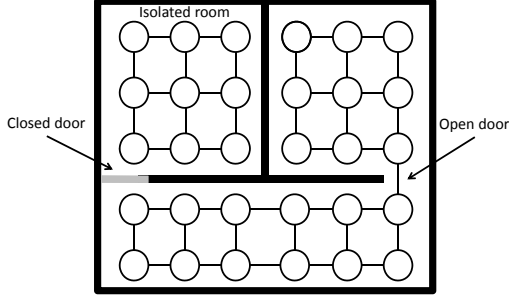
Fig. 4 Example MRF which considers the physical obstacles (objects) in the environment. Nodes represent the gas concentration at cells and edges encode the a priori assumption of correlation between some adjacent nodes. Factors and observations have been omitted here for clarity.

would prevent the formulation of the estimator as a least-squares problem, the following approximation is conveniently proposed:

$$
p(l_{i,j}) \approx \mathcal{N}\left(0, \frac{\sigma_p^2}{(1 - P(o_{i,j}))^2}\right) \tag{10}
$$

which exactly matches the ideal model in Eq. (9) for $P(o_{i,j})$ equals to 0 or 1, while providing a smoothly changing Gaussian model for intermediary values. Since the relevant parts of the map will have occupancy probabilities close to these extremes it becomes a minor issue that Eq. (10) only poorly models Eq. (9) in unexplored areas where $P(o_{i,j})$ is close to 0.5. As depicted in Figure 4, when two adjacent cells are physically separated by an obstacle ($P(o_{i,j}) \to 1$), no correlation is assumed between these cells. Note that this leads to infinite variance in *Eq. (10)*, which is not problematic since the estimator ultimately handles inverse variances instead.

Finally, the energy function capturing the prior constraints reads:

$$
E_p = \sum_{\mathcal{C}_p} E_p(\mathbf{n}_{\mathcal{C}_p}) = \sum_{k=1}^{L} \frac{(m_{i_k} - m_{j_k})^2}{\sigma_p^2/(1 - P(o_{i,j}))^2} \tag{11}
$$

with L the number of pairwise cliques of cell nodes in the GMRF, and i_k, j_k the adjacent cells for each such pairwise clique k.

4 Maximum a Posteriori Estimation of the GDM

We show next how the MAP estimation of the GDM becomes a least-squares problem for the proposed GMRF model. We also describe that the uncertainty of the estimated map can be retrieved from its graphical model.

4.1 Derivation

We start by conditioning the gas concentration map $\mathbf{m}$ to all available data, that is, to all gas observations $\mathbf{z} = \{z_1, \dots, z_M\}$. Then, we seek to maximize the posterior $p(\mathbf{m}|z_{1:M}) \propto p(\mathbf{m}, z_{1:M})$ which gives us the MAP estimate $\hat{\mathbf{m}}$.

By taking the negative logarithm over such posterior, the complete energy function becomes the well-known least-squares form of a GMRF inference problem [Madsen et al, 2004], which in our case reads:

$$
E(\mathbf{n}) = \sum_{k=1}^{L} \frac{(m_{i_k} - m_{j_k})^2}{\sigma_p^2/(1 - P(o_{i,j}))^2} \tag{12}
$$
$$
+ \sum_{k=1}^{M} \frac{(m_{i_k} - z_k)^2}{\sigma_s^2 + \sigma_\zeta^2(t - t_k)}
$$

We can rearrange the terms of the energy function $E(\cdot)$ as a sum of quadratic errors $\mathbf{r}$ weighted by an information matrix $\boldsymbol{\Lambda}$, i.e. $E(\mathbf{n}) = \mathbf{r}^\top \boldsymbol{\Lambda} \mathbf{r}$. Errors can be conveniently defined in terms of a prediction function $\mathbf{f}(\mathbf{m})$ such that $\mathbf{r} = \mathbf{f}(\mathbf{m}) - \mathbf{y}$ for some vector of known data $\mathbf{y}$. From the assumed statistical independence between variables and model noises, it follows that $\boldsymbol{\Lambda}$ is diagonal, leading to:

$$
E(\mathbf{n}) = \mathbf{r}^\top \boldsymbol{\Lambda} \mathbf{r} = \sum_{k=1}^{L+M} \Lambda_k (f_k(\mathbf{m}) - y_k)^2 \tag{13}
$$

where the k subscript denotes the corresponding scalar entry in each matrix or vector. The minimum of the quadratic expression in Eq. (13) can be found by solving the *Gauss-Newton method* equations [Fernandez-Madrigal and Blanco, 2013; Dennis and Schnabel, 1983]:

$$
\underbrace{(\mathbf{J}^\top \boldsymbol{\Lambda} \mathbf{J})}_{\text{Hessian } \mathbf{H}} \boldsymbol{\Delta}\mathbf{m} = -\underbrace{\mathbf{J}^\top \boldsymbol{\Lambda}(\mathbf{f}(\mathbf{m}) - \mathbf{y})}_{\text{Gradient } \mathbf{g}} \tag{14}
$$

where $\mathbf{J} = \frac{d\mathbf{r}}{d\mathbf{m}}$ is the Jacobian of the error function $\mathbf{r}$.

It should be emphasized that unlike other mapping problems in mobile robotics, both factor types ($\mathcal{F}_p, \mathcal{F}_o$) are linear with the map $\mathbf{m}$, which implies that $\hat{\mathbf{m}}$ can be solved in closed form, without iterating.

Matching the generic Eq. (13) to our particular case in Eq. (12) we have:

$$
\begin{aligned}
\mathbf{f}(\mathbf{m}) &= [l_1(\cdot) \cdots l_L(\cdot) \mid m_{i_1} \cdots m_{i_M}]^\top \\
\mathbf{y} &= [0 \quad \cdots \quad 0 \mid z_1 \cdots z_M]^\top
\end{aligned} \tag{15}
$$

from which in section 4.2 we will derive the sparse Jacobian expressions.

Regarding the $(L + M) \times (L + M)$ information matrix $\boldsymbol{\Lambda}$, its first L diagonal entries correspond to the prior factors, that is, to the correlation between adjacent cells, and is $\Lambda_{p_k} = (1 - P(o_{i,j}))^2/\sigma_p^2$. The rest M diagonal entries are the weights of e-nose observations, which decrease over time according to their "age", that is, $\Lambda_{o_k} = 1/(\sigma_s^2 + \sigma_\zeta^2(t - t_k))$.

As a important remark for an efficient implementation, the $\mathbf{Ax} = \mathbf{b}$-like system of equations to be solved in Eq. (14) is highly sparse due to the strongly local structure of the constraints, leading to a symmetric banded Hessian matrix with a bandwidth of $W = \max(N_x, N_y)$. Then, a sparse $LL^\top$ (Cholesky) decomposition can be used for efficiently factoring and solving the system in $\mathcal{O}(N_x N_y W^2)$ and $\mathcal{O}(N_x N_y W)$ time, respectively [Bjorck, 1996, p.220]. This means that, for example, estimating a square map with $N = N_x N_y$ cells has an overall cost that grows with $\mathcal{O}(N^{1.5})$.

Furthermore, symbolic factorization of the linear system does not need to be updated between successive time steps, as long as the size of the gridmap does not vary. Notice that general methods for sparse matrix factoring, such as AMD [Davis, 2004], will automatically group connected cells into clusters according to the structure of the mapped environment, e.g. following its division into individual rooms.

4.2 Jacobian, Hessian and Gradient

Based on Eq. (14) and the particular structure denoted in Eq. (15), next we devise the structure of the system matrices:

– Jacobian $\mathbf{J}$: The $\mathbf{J}$ matrix contains the $\frac{d\mathbf{r}}{d\mathbf{m}}$ for every factor in the graph. Rows of the Jacobian derived from *prior factors* contain zeros but for the i'th and j'th columns, corresponding to adjacent cells in the map (see Eq. (11)), which have values 1 and -1, respectively. Rows for *observation factors* are all zeros except at the column of the observed cell.

$$
\mathbf{J} = \begin{array}{c} \\ 1 \\ 2 \\ \vdots \\ L \\ L+1 \\ L+2 \\ \vdots \\ L+M \end{array}
\begin{array}{c} \begin{matrix} m_1 & m_2 & \cdots & m_i & \cdots & m_j & \cdots & m_N \end{matrix} \\
\left[\begin{matrix}
1 & -1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\
1 & 0 & \cdots & -1 & \cdots & 0 & \cdots & 0 \\
& & & \vdots & & & & \\
0 & 0 & \cdots & 1 & \cdots & -1 & \cdots & 0 \\
0 & 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\
1 & 0 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\
& & & \vdots & & & & \\
0 & 0 & \cdots & 1 & \cdots & 0 & \cdots & 0
\end{matrix} \right]
\end{array}
$$

– Hessian $\mathbf{H}$: Since all functions in our problem are linear, the Hessian is exactly $\mathbf{J}^\top \boldsymbol{\Lambda} \mathbf{J}$. The existence of two blocks in the Jacobian matrix, with its upper block corresponding to the prior factors which typically will not change over time, advices us to decompose the Hessian into the sum of two components:

$$\mathbf{H} = \mathbf{H}_p + \mathbf{H}_o \qquad (16)$$

The first part, $\mathbf{H}_p$, only contains the following nonzero entries:
  – Each off-diagonal entry $H_p(i, j)$ is the sum $\sum_k J(k, j)\Lambda_{p_k}$ for each *prior factor* k relating cells (i, j). Following the graphical model in Figure 3, and attending to the sparse structure of the Jacobian, in our case $H_p(i, j) = -\Lambda_{p_k}$ if cell i is adjacent to cell j, zero otherwise.
  – Each diagonal element $H_p(i, i)$ becomes the sum $\sum_k J(k, i)\Lambda_{p_k}J(k, i) = \sum_k \Lambda_{p_k}$, for each *prior factor* k defined over cell i, that is, for each neighbor of cell i.

The second part, $\mathbf{H}_o$, is exactly diagonal and depends only on the observations. The $i'th$ element of its diagonal $H_o(i, i)$ amounts to $\sum_k \Lambda_{o_k}$, being k the index of each observation taken at cell i.

– Gradient vector $\mathbf{g}$: The gradient vector $\mathbf{g} = \mathbf{J}^\top \boldsymbol{\Lambda} \mathbf{r}$, with length the number of grid cells N, simply becomes:

$$g_i = \sum_k \Lambda_{o_k}(m_i - z_k) + \sum_j \Lambda_{p_k}(m_i - m_j) \qquad (17)$$

for all the observations k taken at cell i, and all the neighbor cells j of cell i.

4.3 Recovering the Uncertainty

To obtain the uncertainty of the estimation at each cell, we have to compute the diagonal of $\mathbf{H}^{-1}$. Each diagonal element $\mathbf{H}^{-1}(i, i)$ corresponds to the variance of cell i (σ_{ii}^2).

Given that matrix inversion is a computationally expensive operation, and since we are only interested in recovering the diagonal of the inverse Hessian, efficient approximations as the one presented by Golub and Plemmons [Golub and Plemmons, 1980] can be employed. They proposed a method for recovering only the entries σ_{ij} of the covariance matrix $\boldsymbol{\Sigma}$ that coincide with nonzero entries in the factor matrix $\mathbf{R}$, being $\mathbf{R}$ the upper triangular matrix that results from the Cholesky decomposition of $\mathbf{H}$.

4.4 Parameter Selection

The proposed GDM method depends on four parameters: the cell size c, the Gaussian errors σ_s^2 and $\sigma_\zeta^2(t-t_k)$, and the tolerance parameter σ_p^2.

The cell size c, which determines the resolution at which different predictions can be made, must be selected attending to the compromise between map resolution and computational cost. In general, for mobile robotics olfaction, given the actual limitations of the gas sensor technology (usually point sampling devices) it makes no sense to have cell sizes below 0.2m, since to gather samples under such resolution will require a tedious and costly inspection of the scenario.

Two are the parameters that control the additive Gaussian errors in the proposed method: the sensor noise parameter σ_s^2, related to the physical characteristics of the sensor employed, and the parameter $\sigma_\zeta^2(t - t_k)$, which models the aging of observations. The latter is conditioned by the dynamics of the gas concentration (presence of sources, airflows, etc.) and the structure of the environment, and it is typically initialized to zero $(\sigma_\zeta^2(0) = 0)$.

These two parameters determine the time-variant information associated to each observation (Λ_{o_k}), as described in Section 4.1. The initial, and maximum, information provided by an observation is given by the sensor noise $(\Lambda_{o_k}(\Delta t = 0) = \frac{1}{\sigma_s^2})$, while the loss of information is controlled by the time-increasing rate of $\sigma_\zeta^2(t - t_k)$. Thus, in environments where the gas concentration changes rapidly (e.g. due to the presence of strong turbulent airflows), the value of Λ_{o_k} should decrease quickly to manifest that gas samples can only be trusted for a short period of time.

Selecting the appropriate value and decreasing rate of Λ_{o_k} is, however, not a simple task. They can be modeled considering not only previously gathered gas observations, but also taking into account information from other sensing modalities (e.g. slower rates for cells where potential gas sources are detected by vision, or even variable rates according to wind measurements). In this article we do not claim any aging model, but the introduction of the concept within GDM methods, thus, a simple linear model for decreasing Λ_{o_k} will be considered in the experimental section.

Finally, the selection of the tolerance parameter σ_p^2 is related to how gas observations extrapolate to neighbor cells. A high tolerance means that neighbor cells are not forced to have similar values, and so extrapolation is reduced. The influence of this parameter can be clearly seen in Figure 5, where the uncertainty maps of a simple experiment employing different values of the parameter σ_p^2 are shown. To better illustrate the effect

(a) $\sigma_p^2 = 5$ (b) $\sigma_p^2 = 2$

(c) $\sigma_p^2 = 1.25$ (d) $\sigma_p^2 = 0.5$

(e) Occupancy map
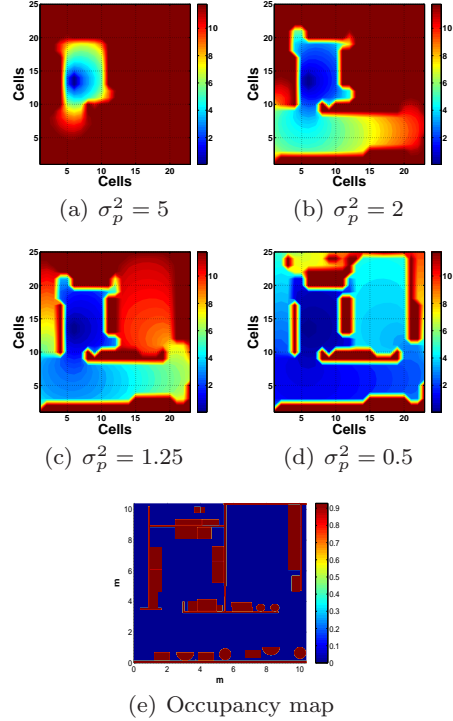
Fig. 5 (a-d) GMRF uncertainty maps of an illustrative experiment consisting of a single gas observation taken at cell(6,14)), for different values of the tolerance parameter σ_p^2. Lower values of the tolerance parameter correspond to higher extrapolations. (e) Occupancy map of the experimental area depicting the probability $P(o_{i,j}) \in [0, 1]$ of a cell to be occupied.

of σ_p^2, only one gas observation is collected in this experiment (at cell (6,14)). As can be noticed, high values of σ_p^2 lead to reduced extrapolations, since, for example, an observation taken inside a room has low influence (reduction of uncertainty) in cells outside it (see Figure 5(a)). However, for small tolerance values, high extrapolation is achieved, and so even only one sample can reduce considerably the uncertainty of far away cells.

5 Experiments

This section presents experiments aimed at validating the performance of the proposed method when estimating the GDM of a time-variant gas distribution in an indoor scenario. The experiments consist of a mobile robot which is patrolling an area with several rooms while building a GDM of the inspected environment. The localization of the robot is handled by an efficient particle filter implementation [Blanco et al, 2010], upon the information provided by the robot odometry and laser scans, over a known map.

To provide quantitative results of the accuracy of the estimated gas maps, as well as to allow for a stringent comparison between different GDM methods, we first perform an experiment in simulation in order to have a ground truth (GT) of the gas distribution. Then, we carry out a real test in very similar conditions.

5.1 Experiment Setup

Both the real and simulated tests are conducted in an indoor scenario composed of two adjacent rooms connected through a corridor (see Figure 6). Each room is inspected through a complete loop by a mobile robot equipped with an e-nose which collects samples along the path (red-dashed line). To go from one room to the other, the robot traverses the corridor (path marked as blue-dashed line).

To study how the proposed method behaves under time-variant gas distributions, both rooms are alternatively inspected by repeating the pattern "left(x3), right(x1)", under the following three cases: (i) when both rooms are clean of odors, (ii) when an ethanol leak appears in the leftmost room, and (iii) when the ethanol odor is removed and both rooms become clean again.

In particular, we analyze the estimated gas distribution maps by taking snapshots at the following representative instants of time: (t_0) at the start of the experiment, (t_1) after inspecting both rooms when no odor has yet been released, (t_2) once the ethanol has spread over the leftmost room, and it has been inspected once, (t_3) after the three inspections (loops) in the leftmost room with odor presence, (t_4) after the first inspection of the leftmost room, again clean of odors, and finally (t_5) after the three inspections of the leftmost room with no odor presence.

For each instant of time t_i, the path followed by the robot will be displayed as a grey-thin line, highlighting the last seconds of it as a thick-white line. The start point is marked as a circumference, and the current robot localization as a triangle. Furthermore, for an easier interpretation of results we have highlighted obstacles in the map (walls and office furnishings) in white color, normalized the gas measurements in the range [0,1], and set an upper bound of 12 on the uncertainty values to be displayed (for graphical considerations).

Related to the selection of parameters, for this experiment, given the dimensions of the experimental area, we set $c = 0.5m$ and the robot average speed to 0.25 m/s. Similarly, and following the reasoning on Section 4.4, we set the Gaussian sensor noise to $\sigma_s^2 = 0.1$, or equiv-
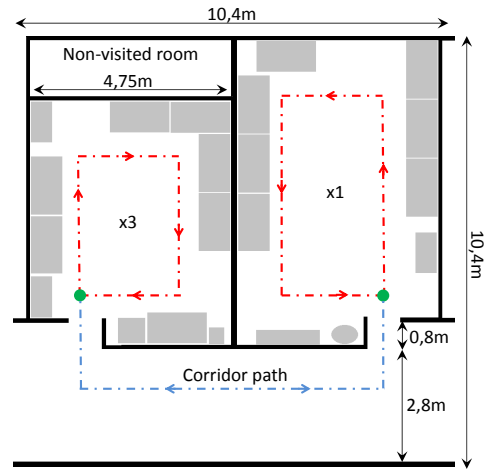
(a)
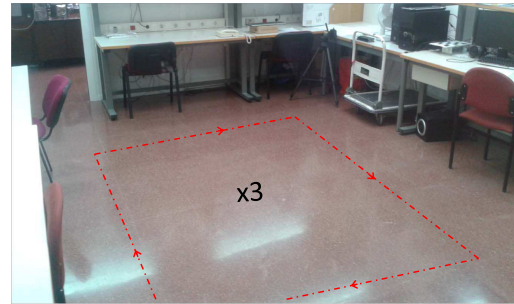
(b)

Fig. 6 (a) Sketch of the experimental scenario. White color represents free space, grey color obstacles as tables or cabinets, and black color walls. (b) Picture of the leftmost room of the above plan. The path followed by the robot is marked as a dashed line.

alently $\Lambda_{o_k}(\Delta t = 0) = 10$, and select a decreasing information rate to obtain an observation lifespan of $60s$. Finally, we set $\sigma_p^2 = 2$ to obtain moderate extrapolation results (see Figure 5).

In this section, given that the experimental map is known a priory, the occupancy probabilities are restricted to binary values ($P(o) = 0$ for free space, $P(o) = 1$ for obstacles), which gives a value for Λ_p of 0.5 and 0, respectively. Note however that this is not a limitation of our method, which can handle uncertainty in the occupancy of the environment (see Section 3.2.2).

5.2 Simulated Experiment

For the synthetic scenario we have made use of the simulation framework presented in [G.Monroy et al, 2013], which allows an easy comparison between the ground truth of the gas distribution and the maps estimated by different GDM methods.
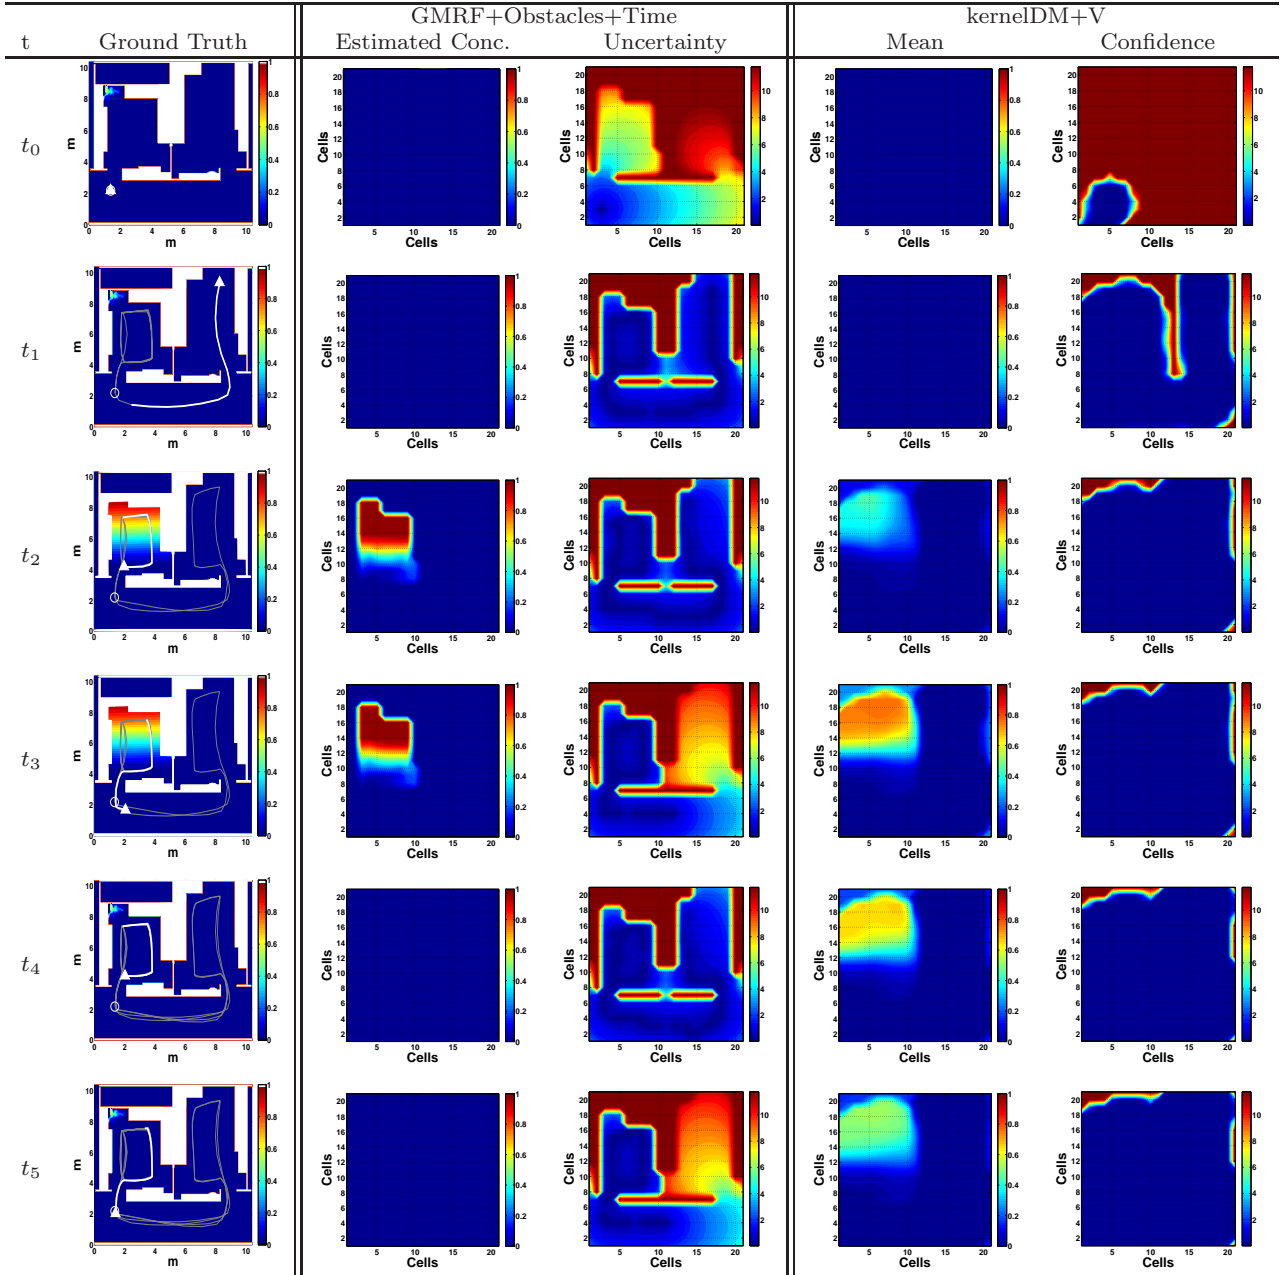
First, in Figure 7 we compare the estimated gas concentration and associated uncertainty maps at the
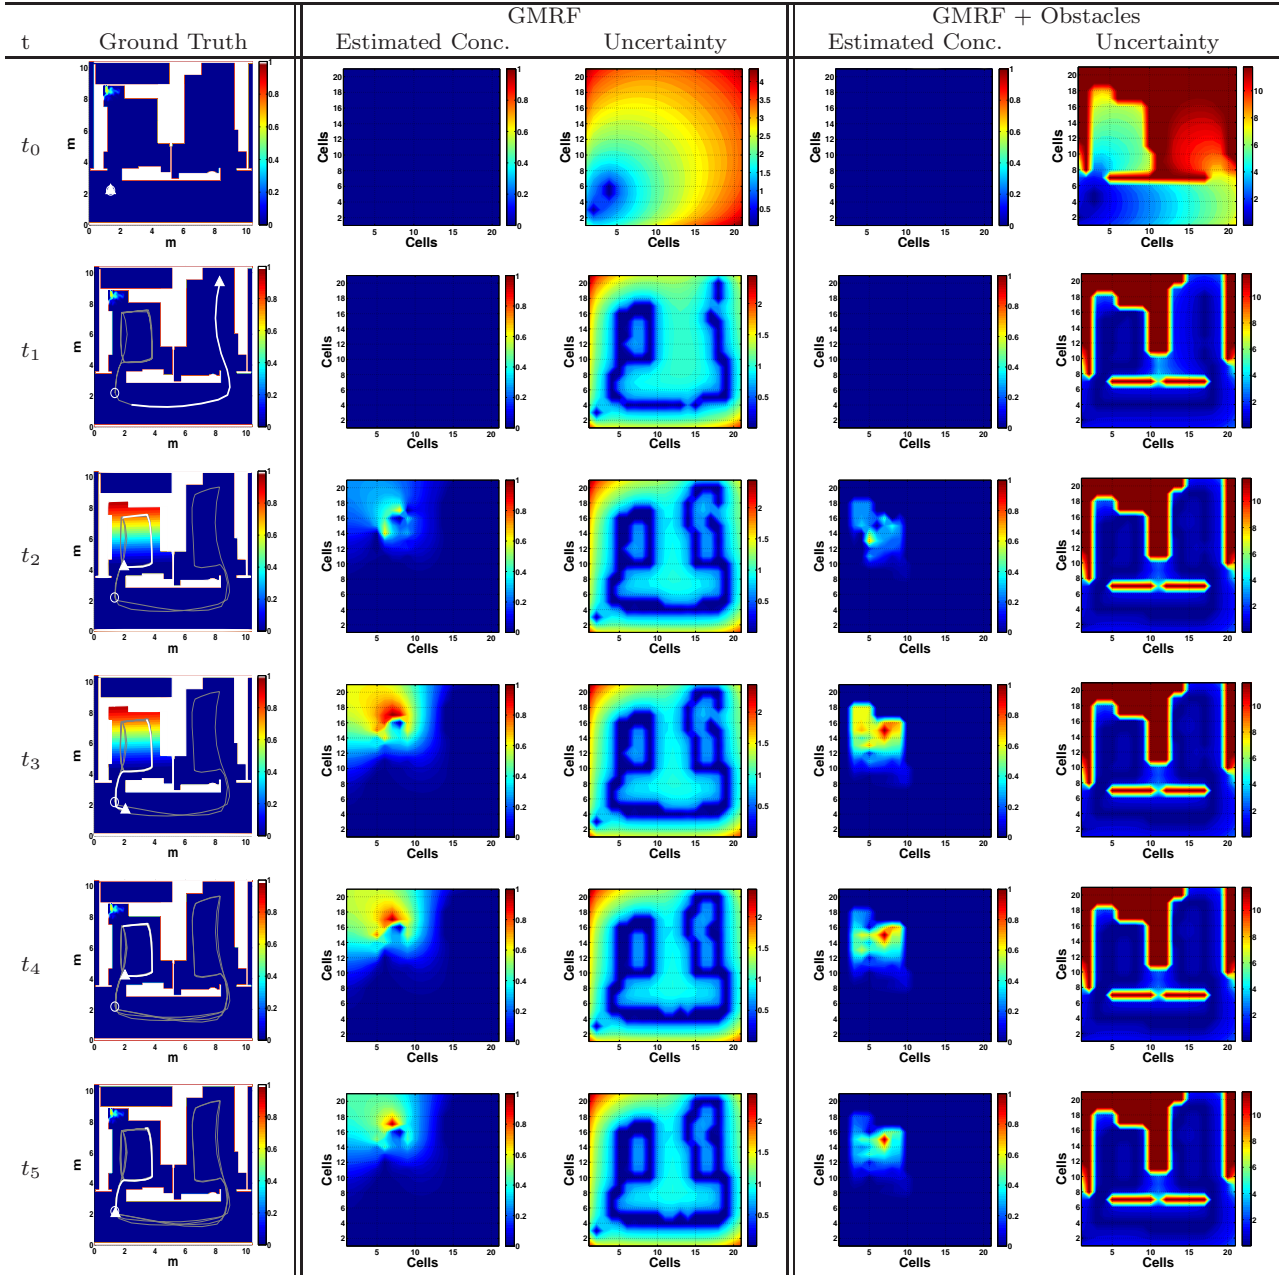
Fig. 7 Maps of the estimated gas concentration and associated uncertainty at the different time steps t_i of the simulated experiment, as estimated by: (a) the proposed GMRF method without accounting for obstacles or aging in the observations, and (b) when introducing obstacles.

different instants of time (t_i) for two configurations of the GMRF method: when neither aging, nor obstacles are considered, and when accounting only for obstacles. Then, in Figure 8 comparison between the GMRF accounting for both obstacles and aging, is performed against the kernelDM+V method, which is probably the more accurate and effective existing method for GDM.

It is important to clarify that unlike the GMRF method, where the variance captures the uncertainty in the estimation in the interval $(0, \infty)$, the variance provided by the kernelDM+V represents the estimated variability of the observations, and so, variance maps cannot be directly compared (see Section 2). Instead, we decide to compare with the kernelDM+V confidence α, which has a similar meaning to that of the GMRF variance, though normalized to the interval $[0, 1)$. To

Fig. 8 Maps of the estimated gas concentration and associated uncertainty at the different time steps t_i of the simulated experiment, as estimated by: (a) the proposed GMRF method accounting for obstacles and observations's aging (compare the results with those in Figure 7), and (b) the kernelDM+V method. Notice that for the latter, the transformed confidence values are depicted following Eq. (18).

fairly compare both magnitudes, we employ the transformation:

$$\sigma_i^2 = \frac{1}{\alpha_i} - 1 \qquad (18)$$

where σ^2 represents the corresponding uncertainty values for the kernelDM+V. To avoid numerical problems we imposed $\alpha_i > 0.001$.

Although an optimization procedure based on the negative log predictive density (NLPD) can be used to select the KernelDM+V parameters when working offline as explained in [Lilienthal et al, 2009], in this work we need an online comparison, and so, such optimization cannot be followed. Furthermore, the suggested optimization procedure assumes that the gas concentration is static along the experiment duration, which completely contradicts our case of study. Instead, we select

its parameters based on the observations found in the cited paper together with some practical restrictions: we select the cell size as in our algorithm ($c = 0.5m$), and to allow a proper extrapolation of individual readings, we set the kernel width to $\sigma = 0.9m$ (complying with the restriction $\sigma > c$). Finally, the cutoff radius parameter is set to 4σ as proposed in [Lilienthal et al, 2009].

The following conclusions can be drawn from this experiment:

1. At time steps t_0 (initialization) and t_1 (after inspecting both rooms for the first time), all the compared methods estimate correctly the absence of volatiles in the environment as appreciated in their respective concentration maps. However, their uncertainty maps differ considerably. The most noticeable difference is appreciated when introducing the obstacles in the environment, influencing the correlation between map cells, and consequently the uncertainty maps. Under these circumstances the proposed GMRF method provides low uncertainties for the visited cells, while for those falling on obstacles or in non-visited rooms (as the isolated room at the top left of the map), still maintain high values, as desired. This contribution enables a better comprehension of the estimated maps, since the structure of the environment (rooms, corridor, walls, etc) is easily appreciated.

2. Once the gas is released, we take snapshots at two time instants: t_2 after only one loop through the contaminated area, and t_3 after three loops, as displayed in Figure 6. Here, since both the KernelDM+V as well as the limited configurations of GMRF and GMRF+Obstacles do not consider the time at which observations are taken, the estimated concentration maps are the result of averaging recent observations with older ones, that is, with observations gathered when no gas was still released. Notice how at such time instants, those methods fail to detect the high concentration values present in the ground truth maps. This, which represents one of the main limitations of existing GDM methods, is successfully overcome by the proposed approach, as appreciated in the GMRF+Obstacles+Time maps in Figure 8. By increasing the uncertainty of observations as they become older, together with the fact that we consider the presence of obstacles in the environment, our approach is able to detect and correctly localize the high gas emissions with only one lap over the contaminated zone. The main difference between time-instants t_2 and t_3, is that at t_3 the uncertainties at the right room start increasing as a consequence of the rise in the variance of observations,
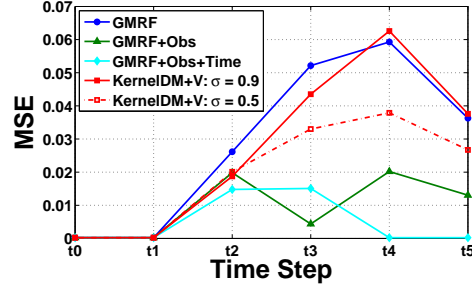
Fig. 9 Mean squared error (MSE) at the different time steps (t_i) of the simulated experiment, for the KernelDM+V and the different configurations of the GMRF method.

which allows the method to remove them from the set of observations.

3. Finally, at time steps t_4 and t_5, when the gas has been removed, the "average" effect of methods not considering the age of observations can still be appreciated as gas patches in the concentration maps. On the contrary, our approach adapts faster to changing gas concentrations, and so it correctly provides a zero concentration estimation at the left room, even after only one lap (t_4).

To quantitatively compare the different configurations of the proposed GMRF method, among themselves, and with other methods previously proposed in the literature, we employ the mean squared error performance measure (MSE). The MSE is calculated as the difference between the ground-truth concentration c_i and the estimated MAP value m_i, or the mean value μ_i for the kernelDM+V method.

$$MSE = \frac{1}{N} \sum_{i=1}^{N} (c_i - m_i))^2 \qquad (19)$$

Figure 9 shows the MSE for the different time steps (t_i) of the simulation. Since the kernelDM+V performance is dependent of the kernel size, we also include the performance measures of or a reduced kernel size of $\sigma = 0.5$, which corresponds to the smaller kernel size given the selected cell size in the experiment. As can be appreciated, our approach introduces a notable improvement, being remarkable the inclusion of obstacles in the GDM process.

A particular case is the comparison between the simplest GMRF configuration and the kernelDM+V, showing both similar performance measures (even advantageous for the kernelDM+V when considering the reduced kernel size). This unveils the importance of considering obstacles and observation ageing in the GDM process, which not only reduce considerably the MSE

on the estimated gas concentration, but also provide more realistic uncertainty maps.

Other performance measures which not only account for the estimated concentration values, but also consider the uncertainty on the estimation (e.g the NLPD), have not been considered in this work because of the inherent differences in the variances of the compared methods. While for the kernelDM+V the variance represents the estimated variability of the gas observations, for the proposed GMRF method the variance represents the uncertainty about the estimated gas concentration, and thus, to the best knowledge of the authors, they cannot be compared. Similarly, kernelDM+V confidence values α, which have been used in the graphical comparison because of the likeness with the uncertainty of the estimations, can neither be used to compute performance measures such as the NLPD because they are not strictly variances, and so, a quantitatively comparison would be meaningless.

5.3 Real Experiment

Following the same setup than in the simulated experiment, we show next the results of a real experience. In this case the gas observations are collected with a photo ionization detector (PID[2]), mounted in a *Patrol-Bot* [MobileRobots Inc., 2014] mobile platform. To allow comparison with the results of the simulated experiment, the *ppm* readings provided by the PID are normalized to the range [0,1].

To generate the gas leak in the left room, we place an ethanol bottle in front of a fan to boost the gas dispersion. The ethanol bottle was timely opened or closed to match the setup described in Section 5.1. This configuration generates a gas plume towards the door heading to the corridor when the ethanol bottle is open, while it helps cleaning the room from gases when the bottle is closed.

In this case, only the full configuration of the proposed GMRF is compared to the kernelDM+V. For a fair comparison among them, the experimental data is collected and saved to a log file for off-line processing. In this way, differences from both estimations will only obey to differences in the methods and not to the data.

Figure 10 shows the estimated concentration and uncertainty maps generated by the GMRF and the KernelDM+V. As in the simulation experiment, snapshots of the maps at the different time steps (t_i) are depicted. Since the ground truth of the gas distribution is not

available, we plot instead the obstacles map together with the robot localization.

As can be noticed, the results are quite similar to those obtained in simulation, which corroborates the already highlighted advantages of our proposal. The main difference resides in the presence of low concentration gases in the left room at time steps t_4 and t_5, after the ethanol bottle is closed. This can be attributed to the fact that the time elapsed since the ethanol bottle was closed till the robot inspected back the room, was not enough for the gas to disappear.

5.4 Execution Time

Finally, we compare the execution time of the different configurations of the GMRF method and the kernelDM+V for different map sizes. This is done by considering the same experimental setup as described in Section 5.1, but reducing the cell size to increase the overall number of cells in the map. Figure 11 summarizes this comparison, plotting the execution time corresponding to the resolution of the least-squares problem together with the recovery of the associated uncertainty (see Section 4.3). As can be appreciated, the computational cost of the kernelDM+V method is almost independent of the map size, and in the order of few milliseconds. This is because the number of cells to be updated for each new gas observation is trunked by the cutoff radius parameter, keeping almost constant the execution time. Conversely, the GMRF method is strongly dependent of the map size, updating all map cells at each time step in order to reflect the effect of obstacles and observation aging.

An interesting property shared by both methods is the fact that they allow batch of observations. This is specially important for the GMRF method, meaning that we are not forced to execute the algorithm for each new gas observation, but we can stock observations for few milliseconds or even seconds and then update the map maintaining the same execution time as in the case of considering only one observation at a time. An example of this property can be observed when comparing the execution time of *GMRF+Obs* and *GMRF+Obs+Time*. The main difference between these configurations is that the latter works only with a finite set of M observations from the set of all observations (the most recent ones), yet their execution times are almost identical. This property allows the proposed GMRF method to be used online even for big map sizes, although reducing the updating frequency.

[2] Model *ppbRAE2000* from RAESystem, with a 10.6 eV UV lamp.
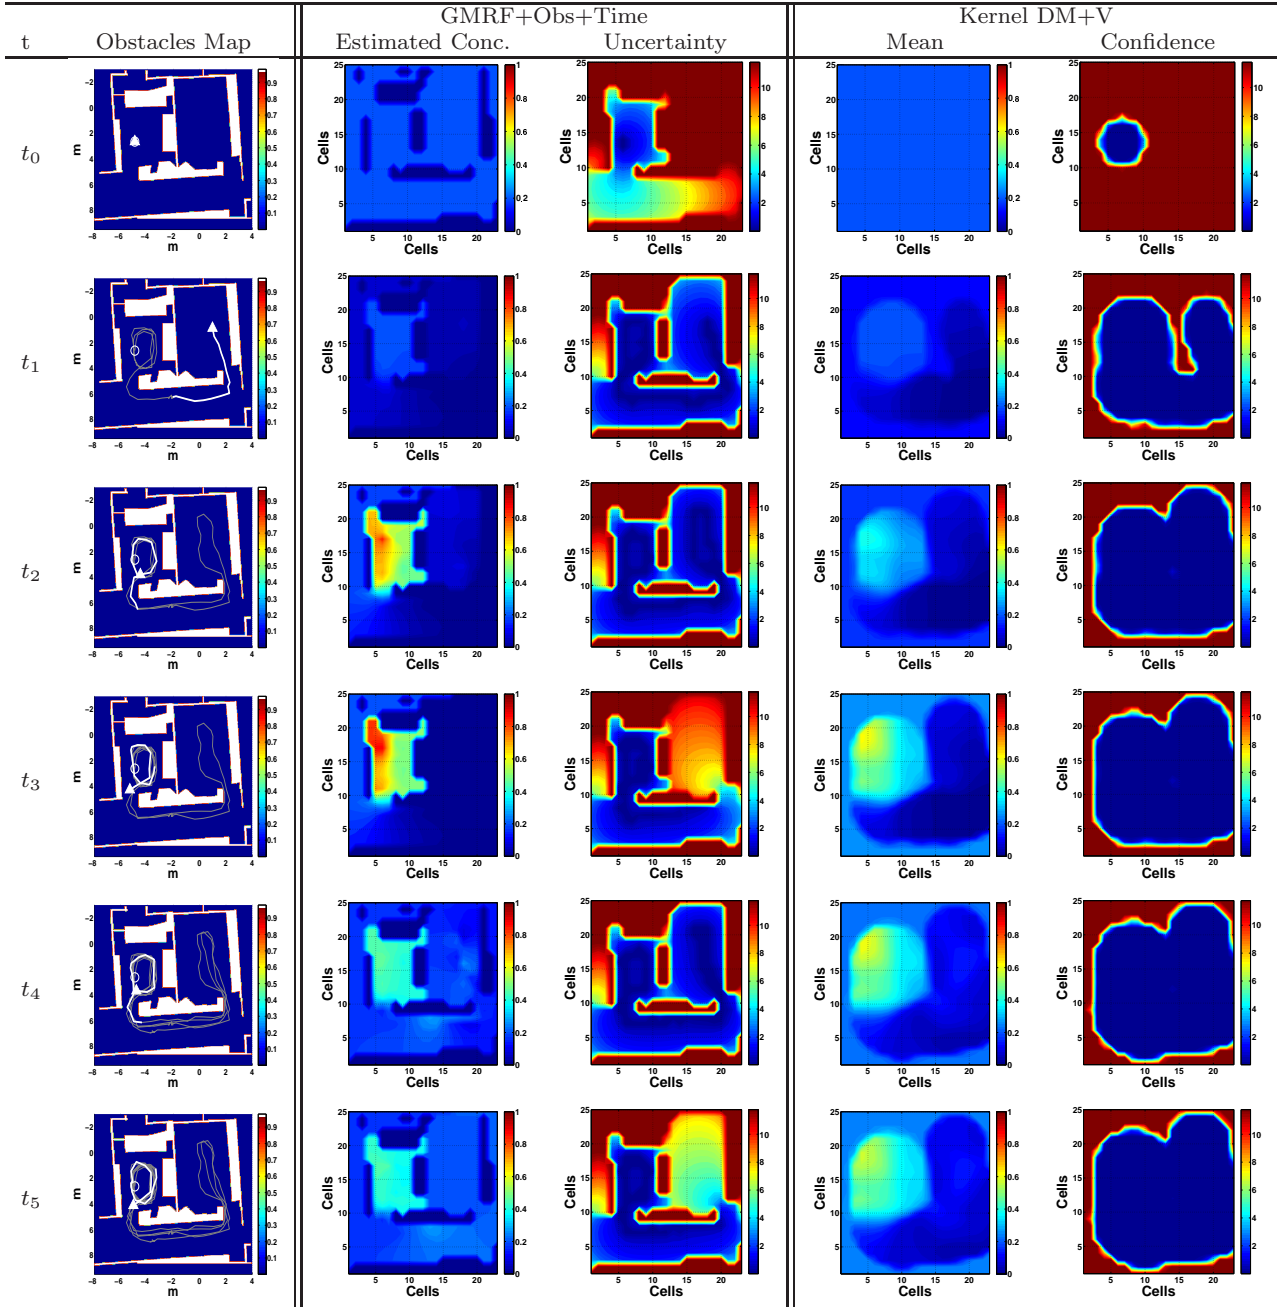
Fig. 10 Maps of the estimated gas concentration and associated uncertainty at the different time steps t_i of the real experiment, as estimated by: (a) the proposed GMRF method accounting for obstacles and observations's aging, and (b) the kernelDM+V method. Notice that for the latter, the transformed confidence values are depicted following Eq. (18).

6 Conclusions and Future Work

In this paper we have revised the problem of creating a map of the gas distribution and proposed a new approach that accounts for two important issues that had been ignored in previous works: the aging of the observations and the presence of obstacles in the environment. The former is achieved by introducing a time-decreasing weighting factor to each gas observa-

tion. Thus, when combining observations taken at close locations, their ages are used to determine their relevances, avoiding the detrimental "average" effect of previous approaches. On the other hand, obstacles such as walls or furniture are now taken into account by modeling the correlation between the map cells they separate. We have addressed the problem in a probabilistic manner modeling the problem as a MAP estimation over a Gaussian Markov random field (GMRF).
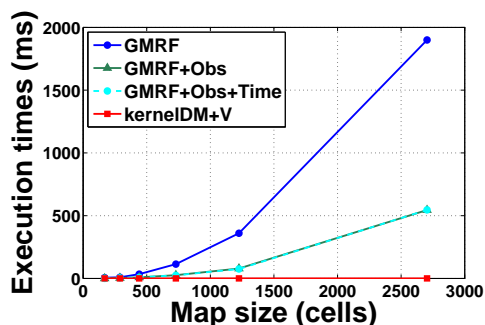
Fig. 11 Execution time for the different configurations of the GMRF method and the kernelDM+V, for different map sizes.

Our approach has been validated with simulated and real experiments, providing qualitative and quantitative comparison with classical methods and demonstrating the advantages in the estimation of the gas distribution.

Future work includes the consideration of additional sources of information such as vision or semantics to improve the way cell correlations are modeled, exploiting properties of the obstacles that influence the gas dispersion (shape, height, etc). Furthermore, efficient alternatives to obtain the GDM of non rectangular scenarios will be studied, for example, by only estimating the gas concentration at desired areas instead of the complete rectangular lattice.

References

Airsense Analytics (2014) PEN3. http://www.airsense.com/en/products/pen3/

Asadi S, Pashami S, Loutfi A, Lilienthal AJ (2011) Td kernel dm+v: Time-dependent statistical gas distribution modelling on simulated measurements. In: Proceedings of the 14th International Symposium on Olfaction and Electronic Nose (ISOEN), vol 1362, pp 281–283

Bishop CM (2007) Pattern Recognition and Machine Learning. Springer

Bjorck A (1996) Numerical methods for least squares problems. Society for Industrial and Applied Mathematics

Blanco JL, Gonzalez-Jimenez J, Fernandez-Madrigal JA (2010) Optimal filtering for non-parametric observation models: Applications to localization and SLAM. The International Journal of Robotics Research 29(14):1726–1742, DOI 10.1177/0278364910364165

Blanco JL, G Monroy J, Gonzalez-Jimenez J, Lilienthal AJ (2013) A kalman filter based approach to probabilistic gas distribution mapping. In: 28th Symposium On Applied Computing (SAC), pp 217–222, DOI 10.1145/2480362.2480409

Clifford P (1990) Markov random fields in statistics. In: Grimmett G, Welsh D (eds) Disorder in Physical Systems: A Volume in Honour of John M. Hammersley, Oxford University Press, Oxford, pp 19–32

Davis TA (2004) A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. ACM Trans Math Softw 30(2):165–195, DOI 10.1145/992200.992205

Dellaert F, Kaess M (2006) Square root sam: Simultaneous localization and mapping via square root information smoothing. The International Journal of Robotics Research 25(12):1181–1203

Dennis J, Schnabel R (1983) Numerical Methods for Unconstrained Optimization and Nonlinear Equations. Classics in Applied Mathematics, Society for Industrial and Applied Mathematics

Fenger J (1999) Urban air quality - their physical and chemical characteristics. Atmospheric Environment 33(29):4877–4900, DOI doi:10.1016/S1352-2310(99)00290-3

Fernandez-Madrigal JA, Blanco JL (2013) Simultaneous Localization and Mapping for Mobile Robots: Introduction and Methods. Information Science Reference

Frish MB, Wainner RT, Green BD, Laderer MC, Allen MG (2005) Standoff gas leak detectors based on tunable diode laser absorption spectroscopy. In: Proc. SPIE 6010, Infrared to Terahertz Technologies for Health and the Environment, DOI 10.1117/12.630599

G Monroy J, Gonzalez-Jimenez J, Blanco JL (2012) Overcoming the slow recovery of mox gas sensors through a system modeling approach. Sensors 12(10):13,664–13,680, DOI 10.3390/s121013664

GMonroy J, Blanco JL, González-Jiménez J (2013) An open source framework for simulating mobile robotics olfaction. In: Proceedings of the 15th International Symposium On Olfaction and Electronic Nose (ISOEN)

Golub GH, Plemmons RJ (1980) Large-scale geodetic least-squares adjustment by dissection and orthogonal decomposition. Linear Algebra and its Applications 34(0):3 – 28, DOI 10.1016/0024-3795(80)90156-1

Gonzalez-Jimenez J, G Monroy J, Blanco JL (2011) The multi-chamber electronic nose. an improved olfaction sensor for mobile robotics. Sensors 11(6):6145–6164, DOI 10.3390/s110606145

Ishida H, Ushiku T, Toyama S, Taniguchi H, Moriizumi T (2005) Mobile robot path planning using vision and olfaction to search for a gas source. In: IEEE Sensors, DOI 10.1109/ICSENS.2005.1597899

Lilienthal AJ, Duckett T (2004) Building gas concentration gridmaps with a mobile robot. Robotics and Autonomous Systems 48(1):3–16

Lilienthal AJ, Loutfi A, Blanco JL, Galindo C, Gonzalez-Jimenez J (2007) A rao-blackwellisation approach to GDM-SLAM: Integrating SLAM and gas distribution mapping (GDM). In: 3rd European Conference on Mobile Robots (ECMR)

Lilienthal AJ, Reggente M, Trincavelli M, Blanco JL, Gonzalez-Jimenez J (2009) A statistical approach to gas distribution modelling with mobile robots - the kernel dm+v algorithm. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS., pp 570 –576, DOI 10.1109/IROS.2009.5354304

Loeliger HA (2004) An introduction to factor graphs. IEEE Signal Processing Magazine 21(1):28–41, DOI 10.1109/MSP.2004.1267047

Loufti A, Coradeschi S, Lilienthal AJ, Gonzalez-Jimenez J (2009) Gas distribution mapping of multiple odour sources using a mobile robot. Robotica 27:311–319, DOI 10.1017/S0263574708004694

Madsen K, Nielsen HB, Tingleff O (2004) Methods for non-linear least squares problems (2nd ed.)

Marjovi A, Marques L (2013) Optimal spatial formation of swarm robotic gas sensors in odor plume finding. Autonomous Robots 35(2-3):93–109, DOI 10.1007/s10514-013-9336-1

MobileRobots Inc (2014) Corporate website. http://www.mobilerobots.com

Pashami S, Asadi S, Lilienthal AJ (2010) Integration of openfoam flow simulation and filament-based gas propagation models for gas dispersion simulation. In: Proceedings of the Open Source CFD International Conference

Reggente M, Lilienthal AJ (2009) Three-dimensional statistical gas distribution mapping in an uncontrolled indoor environment. In: Proceedings of the 13th International Symposium on Olfaction and Electronic Nose (ISOEN), vol 1137, pp 109–112

Sanchez-Garrido C, G Monroy J, Gonzalez-Jimenez J (2014) A configurable smart e-nose for spatio-temporal olfactory analysis. In: IEEE Sensors, pp 1968–1971, DOI 10.1109/ICSENS.2014.6985418

Sensigent Intelligent Sensing Solutions (2014) Cyranose 320. http://www.sensigent.com/products/cyranose.html

Shraiman BI, Siggia ED (2000) Scalar turbulence. Nature 405:639–646, DOI 10.1038/35015000

Stachniss C, Plagemann C, Lilienthal AJ (2009) Gas distribution modeling using sparse gaussian process mixtures. Autonomous Robots 26(2-3):187–202

Tauseef S, Rashtchian D, Abbasi S (2011) Cfd-based simulation of dense gas dispersion in presence of obstacles. Journal of Loss Prevention in the Process Industries 24(4):371 – 376, DOI 10.1016/j.jlp.2011.01.014

Trincavelli M, Hernandez Bennetts V, Lilienthal AJ (2012) A least squares approach for learning gas distribution maps from a set of integral gas concentration measurements obtained with a tdlas sensor. In: IEEE Sensors, pp 1–4, DOI 10.1109/ICSENS.2012.6411118

Tsujita W, Yoshino A, Ishida H, Moriizumi T (2005) Gas sensor network for air-pollution monitoring. Sensors and Actuators B: Chemical 110(2):304 – 311, DOI 10.1016/j.snb.2005.02.008

Turduev M, Cabrita G, Kırtay M, Gazi V, Marques L (2014) Experimental studies on chemical concentration map building by a multi-robot system using bio-inspired algorithms. Autonomous Agents and Multi-Agent Systems 28(1):72–100, DOI 10.1007/s10458-012-9213-x

Winkler G (2003) Image analysis, random fields and Markov chain Monte Carlo methods: a mathematical introduction, vol 27. Springer Verlag