Joint Estimation of Gas & Wind Maps
for Fast-Response Applications[☆]

Andres Gongora[1,2], Javier Monroy[2], Javier Gonzalez-Jimenez[2]

University of Malaga, Boulevard Louis Pasteur 35, Lab. 2.3.6, Malaga 29071, Spain

Abstract

This work addresses 2D gas and wind distribution mapping with a mobile robot for real-time applications. Our proposal seeks to estimate how gases released in the environment are distributed from a set of sparse and uncertain gas-concentration and wind-flow measurements; such that by exploiting the high correlation between these two magnitudes we may extrapolate their value for unexplored areas. Furthermore, because the air currents are completely conditioned by the environment, we assume a priori knowledge of static elements such as walls and obstacles when estimating both distribution maps. In particular, this joint estimation problem is modeled as a multivariate Gaussian Markov random field (GMRF), combining gas and wind observations under a common maximum a posteriori estimation problem. It considers two lattices of cells (a scalar gas-concentration field and a wind vector field) which are correlated following the physical laws of gas dispersal and fluid dynamics. Finally, we report various experiments in which our proposal is compared to other stochastic gas and gas-wind modeling methods under simulation, to evaluate their performance against a computer fluid-dynamics generated ground-truth, as well as under real and uncontrolled conditions.

Keywords: Gas distribution mapping, Wind flow modeling, Fluid mechanics, Machine olfaction, Gaussian Markov random field, Mobile robotics

1. Introduction

Gas distribution modeling (GDM) is the process by which the shape of a gas distribution in a given environment is estimated from a set of sparse measurements [1]. These are usually composed of gas observations (e.g. gathered by e-noses [2, 3]) and, in some cases, may also include wind data from an anemometer [4]. So far, multiple GDM approaches have been proposed for sensor networks [5, 6], mobile robots [7, 8] or even swarm of robots [9, 10], intended to accelerate the sampling process and in turn provide an accurate estimation of the gas distribution in the shortest possible time.

Being able to provide a fast estimate is paramount for applications where decision making is heavily reliant on information about the gas distribution. For example, when dealing with toxic substances in human environments, such as smoke during indoor fire-fighting, or accidental gas-leaks in industrial facilities. Yet, the main constraint of GDM in this regard is the need for multiple measurements that are conveniently distributed along the work area, in order to accurately predict the gas distribution. This requirement notably hinders the applicability of GDM for use-cases that are under time constrains or where not all areas are accessible.

In this work, we propose a novel approach that seeks to reduce the time and effort needed to provide an accurate estimation of the gas distribution by exploiting the correlation between gas dispersion and wind flow. To do so, we propose to model the gas concentration and wind flow in a joint distribution. Thus, we won't limit our approach to exploit wind measurements locally (as in previous approaches [11]), but present a method that estimates the wind currents in the whole work environment as a means of providing reasonable assumptions about the gas concentration at unexplored locations. This brings two important advantages: on the one hand, it boosts the estimation of the GDM itself, reducing the number of measurements when compared to previous methods, and on the other hand, it provides additional data for the posterior decision making process, as a wind map of the environment can be of great value when assessing the situation status.

Concretely, given an indoor environment where a gas is leaking at some unknown location, our goal is to provide an accurate estimate of the gas distribution in the whole environment in the shortest possible time. For this we consider the following assumptions:

---

[1]Corresponding author

[2]Authors are with the Machine Perception and Intelligent Robotics (MAPIR) research group, and the Biomedical Research Institute of Malaga (IBIMA), at the University of Malaga, Spain.

- The gas concentration and the 2D wind vector can be measured at multiple locations in the environment, for example with a fixed network of sensors, or with one or multiple mobile robots. In this work we will consider, without lack of generality, a single mobile robot equipped with an electronic nose and a 2D anemometer [12].

- A floor-plan of the work environment is available. It is exploited by our method to extract physical constrains on the wind currents, and in turn, the gas distribution.

- The gas distribution to be estimated is composed of a single target gas, i.e. we do not consider here mixtures of gases nor the simultaneous mapping of multiple odors. This could however be easily achieved by first introducing a gas classification step and then performing GDM over the resulting gas classes [13].

We model this estimation problem as a Gaussian Markov random field (GMRF) that combines gas and wind in joint probability distribution over two 2D lattices of cells, namely a gas-concentration map and wind vector map, inter-correlated attending to the physical laws of gas dispersal and fluid dynamics explained in Section 3. Thus, given a set of gas and wind observations, this formulation leads to a maximum a posteriori (MAP) problem whose solution is the most likely value of both magnitudes in the environment. Furthermore, because we assume all variables to be Gaussian, we can further derive the problem into a sparse least-squares formulation as detailed in Section 4. This has not only the advantage of being computable if real-time, but also allows us to retrieve the uncertainty of the estimated maps which, as will be discussed for the experiments in Section 5, provides useful information about their reliability.

2. Related Work

Distribution maps, understood as maps that display the spatial distribution of an attribute or phenomenon, are an active research topic in different disciplines. By integrating geo-spacial data they provide valuable information for analysis and a convenient way to determine physical locations of interest for the subject at hand. In this section, we review the most relevant works in the areas of wind flow and gas distribution mapping, with special interest in those works that combine both magnitudes to improve their estimations in a synergistic way.

2.1. Wind Flow Maps

Estimation of the wind flow maps over large geographical areas has been deeply studied in the fields of meteorology and oceanography where this topic is referred as *wind field retrieval*. Proposed methods commonly address the problem with satellite radar-data [14], sometimes combined with fixed weather stations and buoys [15], to invert the geophysical model function that relates the radar-backscatter measurements with the wind speed and direction on the planet's surface [16]. These approaches rely on atmospheric models to estimate the wind field [17, 18], which cannot be applied to indoor environments.

In a smaller scale domain, diverse works have tackled the problem of wind modeling by using trace gases [19], or the more complex and nowadays extended use of sophisticated computational fluid dynamics (CFD) tools [20]. Nevertheless, CFD-based techniques requires a complete knowledge of the boundary conditions of the environment (rarely available) as well as very high computational resources (beyond the capabilities of mobile platforms).

In this sense, data-driven approaches represent an interesting option for mobile robotics. The goal is to model the average wind flow from a set of measurements that the robot acquires while exploring the environment, or part of it. We can highlight works like [21], where a simple rule-based algorithm was applied to compute the wind flow patterns from a set of wind measurements, or [22] where the authors presented a wind field estimation algorithm for flying robots that considered a statistical distribution to fit their data (only wind speed). Later, in [23] the authors addressed the problem of building a spatial model of turbulent air flow in a joint orientation-speed space. A novel extrapolation method was presented to model the air flow as a linear combination of laminar and turbulent components. Two are the main differences with the proposal presented in this work: (i) we do model the effect that obstacles (e.g. walls, doors, etc) have over the wind map, while the former does not, and (ii) we can provide estimations of the wind flow from a low number of measurements, while the former requires large amount of data points to properly model the probabilistic distributions used for extrapolation.

Finally, we must highlight the work presented in [4] where a probabilistic framework based on GMRF was proposed to fast estimation of the 2D average wind flow at any location in an indoor environment. This work can be considered a simplified version of the current proposal, since only wind flow data was estimated and a detailed knowledge of the environment (i.e. boundary conditions) was assumed.

2.2. Gas Distribution Modeling

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 [24]. It is the task of a GDM algorithm to extrapolate sparse measurements in order to obtain an estimation of the gas dispersal at non-visited locations.

As was the case for wind flow modeling (WFM), many gas distribution models were developed back in the early 90s to tackle the problems of atmospheric dispersion [25]. Yet, such models are neither suitable for local scales, nor

are designed to capture all the relevant aspects of gas propagation with a sufficient level of detail. In this sense, the robotics community has proposed several alternatives to face this problem.

Works over the last decade include [26], where a Gaussian process mixture model was proposed, treating GDM as a regression problem using Expectation Maximization, or [27], where authors treated the gas distribution modeling as a density estimation problem, carrying out two parallel estimation processes: one for the mean and another for the variance of the map. Following this line, Blanco et al. [24] proposed a different approach to also estimate the average and variance of the gas distribution by employing a sparsified Kalman filter. More recently, exploiting the availability of other sensors onboard the robot (e.g. lasers and RGB-D cameras), the authors in [1] considered, for the first time, the presence of obstacles during the estimation of the gas distribution. The work presented here builds upon this work, aiming to improve the results in the estimation of the gas distribution by also considering wind measurements provided by an anemometer.

Special mention requires those works that account for the wind conditions when modeling the gas dispersion. For instance, in [28] the authors illustrated how the wind vector, measured with an ultrasonic anemometer onboard the robot, notably improved the posterior gas distribution mapping. Similarly, and although not directly applicable to indoor environments, the work of Neumann et al. [8] employed a micro drone equipped with gas sensor and anemometer to estimate the gas distribution outdoor. Also interesting is the work of Li et al. [29] where the authors considered an historical record of wind flow measurements to predict the gas distribution. However, applying this work to indoor scenarios is not straightforward as the assumption of homogeneous wind flow does not hold in most cases.

Finally, it must be stressed that all proposed GDM approaches that take into consideration wind information during the estimation process, only account for local measurements. Given the important impact that obstacles in the environment yields over the wind flow and its turbulent behavior, the spatial propagation of the wind measurements over the map might be quite unrealistic. Therefore, as proposed in this work, it becomes mandatory to not only account for the local wind measurements, but to enhance the GDM estimation capabilities by including a comprehensive wind model over the work area.

3. Formulation of the Problem

We introduce in this section the basis for the joint estimation of the wind and gas-distribution maps over a 2D lattice of cells using GMRFs. We will first elaborate on our problem to formulate it as a Markov random field (MRF) that describes all relevant state variables as Gaussian (hence the name GMRF). Afterward, we will introduce its factor-graph representation to help us define the

relationships and interactions between neighboring cells, and finally, we will derive an algebraic formulation of the underlying probability distribution from which a maximum a posteriori (MAP) estimate is developed.

3.1. Formulation as a Markov Random Field

The complexity of the problem at hand makes analytical solutions intractable. Thus, as in many other domains, we simplify the problem by considering a grid map. A noticeable particularity of such representation is that neighbor cells are not independent among them, and then they can be modeled as a discrete MRF. This way we preserve the underlying probabilistic nature of GDM and WFM, yet attain a computationally efficient solution for real-time applications. Naturally, this approach comes at the expense of accuracy and generality.

All the variables of interest that comprise the ensuing MRF can be grouped into two categories: unobserved and observed (i.e. unknowns and input data), although we will further divide the latter for convenience. Thus we end up with (i) the hidden variables representing the unknown gas and wind values (denoted as $\mathbf{m}$ for map), (ii) the measurements that the robot takes along the sampling process ($\mathbf{z}$), and (iii) prior knowledge about the environment and obstacles within it ($\mathbf{o}$). Then, we define the joint probability distribution of our problem as $p(\mathbf{m}, \mathbf{z}, \mathbf{o})$, and the posterior as $p(\mathbf{m} \mid \mathbf{z}, \mathbf{o})$, which we will maximize in Section 4 to obtain the most likely gas and wind fields given all the available data.

Next we analyze in detail each of those variables, providing a formal formulation of the problem.

3.1.1. Hidden Variables

We treat hidden variables (i.e gas concentration and wind vector) as a random field over a rectangular lattice $\mathbf{m} = \{\mathbf{m}_i\}_{i=1}^N$, where $\mathbf{m}_i = (g_i, \mathbf{w}_i)$ is a multivariate random variable that stands for the gas concentration g_i and wind vector $\mathbf{w}_i = (w_x, w_y)$ at each cell i. Fig. 1a shows an illustration of this grid-map, where each cell is represented by its coordinates $(x, y)_i$ and where $N = N_X \times N_Y$ denotes the dimensions of the map in terms of the number of cells, for a total of $3N$ unknowns as depicted in Fig. 1b. It must be stressed out that the random variables within a given cell are considered independent, but correlated with its neighbor cells.

3.1.2. Gas and Wind Measurements

To represent the sensor observations, our MRF accounts for the set of measurements $\mathbf{z} = \{\mathbf{z}_g, \ \mathbf{z}_w\}$, as can be seen in Fig. 1c. This set is comprised of M_g gas measurements $\mathbf{z}_g = \{z_{g_k}\}_{k=1}^{M_g}$ and M_w wind measurements $\mathbf{z}_w = \{z_{w_k}\}_{k=1}^{M_w}$, each one taken at a specific position $(x, y)_k$ and uniquely associated to the closest cell in the map. Also, each measurement has a time-stamp t_k setting the moment when it was acquired, so that we may
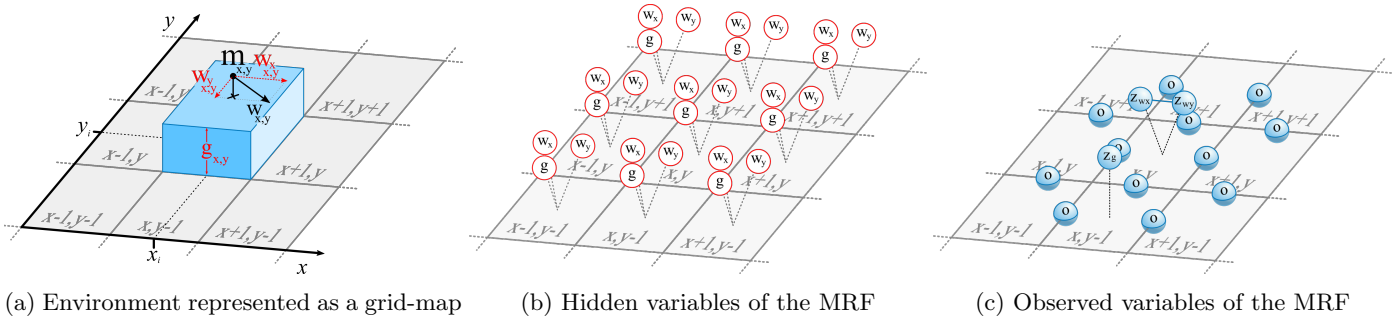
3

(a) Environment represented as a grid-map   (b) Hidden variables of the MRF   (c) Observed variables of the MRF

Figure 1: For any given location (x_i, y_i) at time t, figure (a) shows the gas intensity and wind at the discretized cell (x, y). By continuing this idea, we may assign the three random variables $g_{x,y}$, $w_{x_{x,y}}$ and $w_{y_{x,y}}$ to each cell as shown in (b). The observed variables, such information about the presence of obstacles between cells o or gas and wind observations z (assigned to the nearest cell) are depicted in (c).

also record how the environment changes over time [1]. Notice that this formulation allows for the possibility of more than one measurement falling onto the same cell $\mathbf{m}_i$, while also permitting the most general case of cells having no associated measurement.

3.1.3. Obstacle Information

To account for the physical distribution of obstacles in the work environment, the MRF also includes prior information capturing the correlation between neighboring cells. Assuming that there is some sort of map of the environment, either a deterministic floor plan or a occupancy grid map generated by the robot (i.e. each cell stores the probability of being occupied by an obstacle), we derive a *cell-connectivity* map [30] that indicates whether adjacent cells are separated by an obstacle that blocks the air-flow (see Fig. 1c). To keep the notation as clear as possible, we will refer to the latter as $\mathbf{o} = \{o_{i,j}\}$, where a value of $o_{i,j} = 0.0$ denotes that there is no obstacle between the adjacent cells i and j, $o_{i,j} = 1.0$ stands for the presence of an obstacle (e.g. a wall), and intermediate values indicate different degrees of certainty between both extremes (not applicable for a deterministic floor plan).

3.2. Formulation as a Factor Graph

The Hammersley–Clifford theorem [31] shows that some probabilistic models, such as Markov networks, can be represented as factor graphs, which are frequently used when performing inference of belief propagation in the networks. For the problem addressed here, this entails that the joint probability distribution $p(\mathbf{m}, \mathbf{z}, \mathbf{o})$ we are interested in can be factored as the product of potential functions $\varphi(\cdot)$ along the set of all its maximal cliques ($\mathcal{C}_m$) [32]:

$$p(\mathbf{m}, \mathbf{z}, \mathbf{o}) = \frac{1}{\zeta} \prod_{\mathcal{C} \in \mathcal{C}_m} \varphi_{\mathcal{C}}(\mathbf{v}_{\mathcal{C}}) \qquad (1)$$

where the proportionality constant ζ (called the partition function) is not relevant in our maximization problem, $\mathcal{C}$ denotes the different cliques, and $\mathbf{v}_{\mathcal{C}}$ the set of variables $(\mathbf{m}, \mathbf{z}, \mathbf{o})$ in each of these cliques. Each potential function

is thus an equation that relates a subset of variables from the MRF, defined over a single clique.

Potentials are restricted to be strictly positive by definition ($\varphi(\cdot) > 0$), therefore, for convenience in future calculations we express them as exponential functions[3], that is, $\varphi_C = e^{-E_{\mathcal{C}}(\mathbf{v}_{\mathcal{C}})}$, where we define $E(\cdot)$ as an energy function over each of the maximal cliques. Following this notation, the joint probability becomes:

$$p(\mathbf{m}, \mathbf{z}, \mathbf{o}) \propto \prod_{\mathcal{C} \in \mathbf{C}_m} \exp\{-E_{\mathcal{C}}(\mathbf{v}_{\mathcal{C}})\} = \exp\left\{-\sum_{C \in \mathbf{C}_m} E_{\mathcal{C}}(\mathbf{v}_{\mathcal{C}})\right\} \qquad (2)$$

Moreover, we will assume that all the potentials involved in our problem can be reasonably modeled as Gaussians functions (becoming the underlying graphical model a Gaussian Markov Random Field (GMRF)). This assumption works well in practice as demonstrated experimentally, and has the advantage of leading to a convenient least squares formulation of the problem. We can therefore establish a direct analogy between the energy functions $E_{\mathcal{C}}(\mathbf{v}_{\mathcal{C}})$ and the normal distribution $\mathcal{N}(\mu, \sigma) = k \; e^{r^2/2\sigma^2}$; where k is a normalization constant that is not relevant for our optimization problem, r is the residual of the distribution with regard to its mean μ, and σ is the standard deviation. Taking this assumption into consideration, the joint probability becomes a weighted sum of squared residuals:

$$p(\mathbf{m}, \mathbf{z}, \mathbf{o}) \propto \exp\left\{-\sum_{C \in \mathbf{C}_m} E_{\mathcal{C}}(\mathbf{v}_{\mathcal{C}})\right\} = \exp\left\{-\sum_{C \in \mathbf{C}_m} \frac{r_{\mathcal{C}}^2}{2\sigma_{\mathcal{C}}^2}\right\} \qquad (3)$$

At this point, all that is left to mathematically define our problem is to formulate each of the energy functions $E_{\mathcal{C}}(\mathbf{v}_{\mathcal{C}})$ that encode the relationship between the hidden variables (the map m) and the observable ones (observations z and obstacles o).

3.3. Formulation of the Energy Functions

The energy functions for all the factors and their associated cliques can be grouped according to their nature into

---

[3]Exponentials are monotonous growing functions and therefore do not change the locus of the maxima in our optimization problem.

(a) Gas observation   (b) Gas regularization   (c) Wind observation   (d) Wind regularization

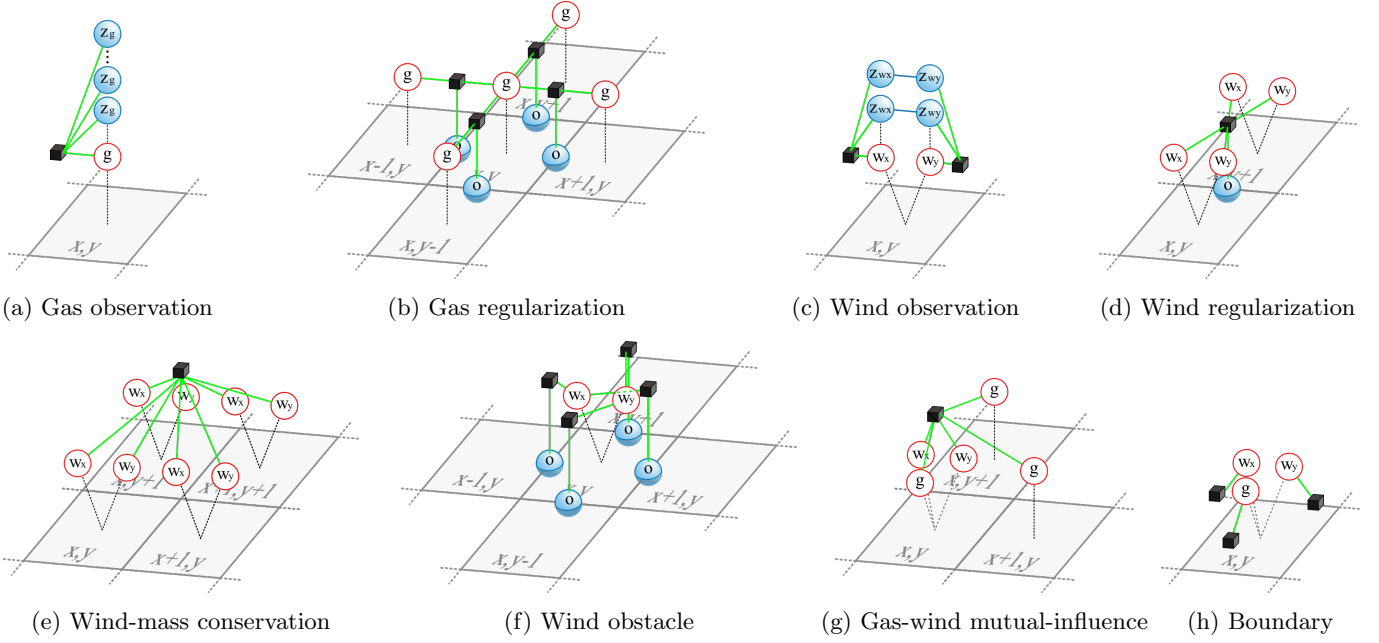(e) Wind-mass conservation   (f) Wind obstacle   (g) Gas-wind mutual-influence   (h) Boundary

Figure 2: Factor graph representation of the involved energy equations, showing which random variables from the GMRF formulation are affected by them. The hidden variables (gas and wind) are shown in red, the observed variables (observations and obstacle information) in blue, and the factors in black. Note that the green line denote that the connected variables are part of a given factor type, all of which repeat are repeated for the whole lattice.

eight different types, namely:

- Energies derived from the gas and wind observations, E_{gz} and E_{wz} respectively.

- Energies related to how these observations extrapolate to neighbor cells in the grid, referred to as regularization terms E_{gr} and E_{wr}.

- The energies involved in modeling the wind flow, accounting for both the presence of obstacles E_{wo} and the law of mass conservation E_{wc}.

- The energy that accounts for the interaction between gas and wind at any given cell, E_{gw}.

- The energy term that defines the default distribution when no observations has been gathered, E_{gw}.

Next, we derive the exact formulation of these energy terms individually. For each, we briefly reason how their physical meaning can expressed as a Gaussian, then provide an illustrative example of the underlying clique as a factor graph, and finally, formulate the energies as weighed squared residuals.

3.3.1. Energy of the Gas Observations (E_{gz})

This energy encompasses the cliques that establish the relationship between the observed gas intensity $\mathbf{z}_{\mathrm{g}}$ and the estimated gas concentration g in the environment. To model this energy, we assume that each observation is corrupted by two additive Gaussian errors: one from the inherent sensor noise ($\omega_{\mathrm{gz}_k} \sim \mathcal{N}(0, \sigma_{\mathrm{gz}}^2)$), and another, time-

dependent, that models the loss of information as the observation gets old ($\zeta_{\mathrm{gz}_k} \sim \mathcal{N}(0, \sigma_{\zeta_{\mathrm{gz}}}^2 \Delta t_{\mathrm{gz}_k})$). The latter models the potential changes that may have occurred in the gas concentration since the sensing at time t_{gz_k} by considering an increasing uncertainty over time [1]. Adding these terms to the ideal (noiseless) sensor model $h(\mathbf{g})$, we get:

$$z_{\mathrm{g}_k} = h(\mathbf{g}) + \omega_{\mathrm{gz}_k} + \zeta_{\mathrm{gz}_k} = g_{i_k} + \omega_{\mathrm{gz}_k} + \zeta_{\mathrm{gz}_k}, \quad (4)$$

where in the last step we considered that every measurement is associated to one single cell i, the one on which the robot e-nose is sniffing.

Under a probabilistic point of view, each gas-observation factor in the graphical model stands for the conditional probability density function (PDF):

$$p(\mathbf{z}_{\mathrm{g}_k}|\mathbf{g}) = p(\mathbf{z}_{\mathrm{g}_k}|g_{i_k}) = \mathcal{N}(g_{i_k}, \sigma_{\mathrm{gz}}^2 + \sigma_{\zeta_{\mathrm{gz}}}^2 \Delta t_{\mathrm{gz}_k}) \quad (5)$$

where, following the local Markov property[4], z_{g_k} is conditionally independent of all other variables given the concentration g_{i_k} (i.e. a gas-measurement only depends on the instantaneous intensity of the gas at the cell it is measured on). This leads to the gas observation factors (see Fig. 2a) to be expressed as energy functions in the form of:

$$E_{\mathrm{gz}} = \frac{1}{2} \sum_{k=1}^{M_{\mathrm{gz}}} \frac{(g_{i_k} - z_{\mathrm{g}_k})^2}{\sigma_{\mathrm{gz}}^2 + \sigma_{\zeta_{\mathrm{gz}}}^2 \Delta t_{\mathrm{gz}_k}} \quad (6)$$

where M_{gz} represents the total number of gas observations gathered since the beginning of the mapping process.

---

[4]A variable in a Markov random field is conditionally independent of all other variables given all its neighbors.

3.3.2. Energy of the Gas Regularization (E_{gr})

The regularization energy is meant to represent how gases spread spatially. That is, this term models the correlation between the gas concentrations of neighboring cells, enabling the estimation of the gas concentration at cells that were not subject to direct sampling. It represents an important term in our GDM approach, as it allows us to reconstruct a coarse approximation of the whole gas distribution from even a reduced set of sparse measurements.

Particularly, we model this spatial correlation by means of a function that penalizes the difference of gas concentration ($d_{i,i'}$) between pairs of adjacent cells:

$$d_{i,j} = g_i - g_j \qquad (7)$$

where g_i and g_j are the gas concentrations at adjacent cells with lattice indices i and j, respectively. In terms of probability, this correlation can be modeled with a Gaussian probability distribution such as:

$$p(d_{i,j}) = \mathcal{N}(0, \sigma_{\mathrm{gr}}^2) \qquad (8)$$

with σ_{gr}^2 a fixed tolerance parameter that controls the *strength* of such correlation. Low values of σ_{gr}^2 will enforce neighbor cells to have a very similar gas concentration, while a more relaxed value of it will allow a greater tolerance among the gas concentration values.

It must be noticed that this formulation does not yet take into account the presence of obstacles in the environment, setting the same correlation between all neighbors cells in the grid. This can be easily solved by conditioning the latter probability to the cell connectivity map (see Fig. 3) as:

$$p(d_{i,j}) \approx \mathcal{N}\left(0, \ \frac{\sigma_{\mathrm{gr}}^2}{(1 - o_{i,j})^2}\right) \qquad (9)$$

For instance, if we are sure there is no obstacle between cells i and j (i.e. $o_{i,j} = 0$), then the gas correlation remains active, that is $p(d_{i,j} \mid o_{i,j} = 0) = \mathcal{N}(0, \sigma_{\mathrm{gr}}^2)$, whereas if there is an obstacle that obstructs the gas dispersion between these cells (i.e. $o_{i,j} \to 1$), it is reasonable to state that the difference of gas concentration can take any value with equal probability, that is, that $p(d_{i,j} \mid o_{i,j} = 1)$ follows a uniform distribution. Since mixing normal and uniform distributions would prevent the formulation of the estimator as a least-squares problem, we have approximated the latter as a Gaussian distribution with an infinite variance which not only matches the above assumptions perfectly, but also provides a smooth transition for unexplored areas where $o_{i,j}$ is close to 0.5.

Finally, we derive the energy function related to the gas regularization factors (see Fig. 2b) as:

$$E_{gr} = \frac{1}{2} \sum_{k=1}^{N_{\mathrm{r}}} \frac{(g_{i_k} - g_{j_k})^2}{\sigma_{\mathrm{gr}}^2 / (1 - o_{i_k, j_k})^2} \qquad (10)$$

being N_{r} the number of pairwise cliques of cell nodes in the GMRF, and i_k, j_k the adjacent cells for each of such cliques.
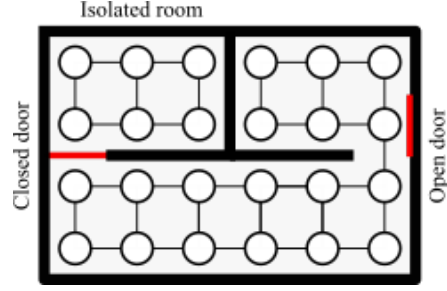

Figure 3: Illustration of how the physical disposition of obstacles, such as doors or walls, affect the correlation between hidden variables in the GMRF. In this particular example the left door is presumed to be closed, and thus, the gas and wind values inside the room are independent of the rest of the environment. Walls are denoted in black, doors in red, and the GMRF over the underlying grid-map is represented by the nodes and edges.

3.3.3. Energy of the Wind Observations (E_{wz})

This energy term establishes the relationship between the estimated value of the wind vector at any cell $\mathbf{w}_i$ and its measured values $\mathbf{z}_w$ with an anemometer. The logic behind it is analogous to that of the gas measurements: to account for sensor noise and age-based uncertainty of the gathered samples. Thus, the wind observation factors shown in Fig. 2c can be modeled by the energy function:

$$E_{\mathrm{wz}} = \frac{1}{2} \sum_{k=1}^{M_{\mathrm{wz}}} \frac{\| \mathbf{w}_{i_k} - \mathbf{z}_{\mathrm{w}_k} \|^2}{\sigma_{\mathrm{wz}}^2 + \sigma_{\zeta_{\mathrm{wz}}}^2 \Delta t_{\mathrm{gz}_k}}, \qquad (11)$$

with M_{wz} the total number of wind samples gathered since the beginning of the sampling process, and accounting for the two-dimensional nature of the wind vector.

3.3.4. Energy of the Wind Regularization (E_{wr})

Employing the same reasoning for the wind flow as we did for the regularization of the gas concentration, we can condition the difference of two (vertically or horizontally) adjacent cells of $\mathbf{w}$ to a Gaussian function that depends on the cell connectivity map $\mathbf{o}$. This leads to factors in to be formulated with energy functions:

$$E_{\mathrm{wr}} = \frac{1}{2} \sum_{k=1}^{N_{\mathrm{r}}} \frac{\| \mathbf{w}_{i_k} - \mathbf{w}_{j_k} \|^2}{\sigma_{\mathrm{wr}}^2 / (1 - o_{i_k, j_k})^2} \qquad (12)$$

where σ_{wr}^2 controls the likelihood that two adjacent cells have different wind vectors (w_x, w_y), and $k \in N_{\mathrm{r}}$ stands for all the pairs of adjacent cells.

3.3.5. Energy Derived from the Conservation of Wind Mass (E_{wc})

The wind flow is subject to additional constrains that have not equivalence in the gas concentration model. Such is the case of the mass-conservation law which states that the total mass of air in the environment should not change assuming the air is treated as an in-compressible gas. This is particularly applicable to indoor environments where

low wind speeds are to be expected (ambient air can be treated as an in-compressible fluid at slow wind-speeds (below 100 m/s) [33, 34]).

This constraint imposes that the divergence of the 2D wind-vector at any cell in the grid should be zero [35], that is, the amount of air entering a cell should be equal to the amount of air leaving it, which can be expressed as:

$$\Delta\mu = \sum_{i=1}^{N} \nabla \cdot \mathbf{w}_i = 0 \tag{13}$$

where $\Delta\mu$ stands for the total change of air mass in the environment and $\mathbf{w}_i$ the wind vector at the i'th lattice cell.

To account for real world phenomena such as changes in the ambient pressure, or the presence of obstacles that may be modeled as solid but in reality can let the air to pass through, we model this constraint as Gaussian random variable, then allowing deviations from the zero divergence criteria: $\Delta\mu_k \sim \mathcal{N}(0, \sigma_{wc}^2)$. Thus, the corresponding energy term becomes:

$$E_{\mathrm{wc}} = \frac{1}{2} \sum_{i=1}^{N} \frac{(\nabla \cdot \mathbf{w}_i)^2}{\sigma_{\mathrm{wc}}^2} \tag{14}$$

The divergence on a discrete vector field might is implemented here as the variation of wind flow among a set of four adjacent cells in a 2x2 pattern as depicted in Fig. 2e. Its formula reads:

$$\begin{aligned}
\nabla \cdot \mathbf{w}_i = \ & w_{x,(x_i,y_i)} && - w_{x,(x_i+1,y_i)} \\
& + w_{x,(x_i,y_i+1)} && - w_{x,(x_i+1,y_i+1)} \\
& + w_{y,(x_i,y_i)} && + w_{y,(x_i+1,y_i)} \\
& - w_{y,(x_i,y_i+1)} && - w_{y,(x_i+1,y_i+1)}
\end{aligned} \tag{15}$$

3.3.6. Energy Related to the Influence of Obstacles in the Wind Vector (E_{wo})

With this energy term we further restrict the wind vector by conditioning it on the surrounding objects. Concretely, we penalize any wind component that is perpendicular to an obstacle (i.e. an occupied cell), allowing tangential wind vectors but not perpendicular ones. For any cell i, this constraint is expressed as:

$$\mathbf{w}_i \perp \mathbf{o} = \mathbf{w}_i^\mathsf{T} \cdot \mathbf{n}_{\mathbf{o}_i} = \mathbf{0} \tag{16}$$

where $\mathbf{n}_{\mathbf{o}_i}$ denotes the normal vector to the obstacle surface encoded in the obstacle map $\mathbf{o}$ around the cell perimeter.

Similar to the wind-mass conservation, this constraint may not always hold in practice as the robot might misjudge the obstacles around a cell. This energy term is thus modeled as four Gaussian distributions, one for each side of the cell i as shown in Fig. 2f, to allow for some degree

of tolerance handled by σ_{wo}^2:

$$\begin{aligned}
E_{\mathrm{wo}} = \frac{1}{2} \sum_{i=1}^{N} \Bigg[\ & \frac{(w_{x,(x_i,y_i)} \, o_{(x_i+1,y_i)})^2}{\sigma_{\mathrm{wo}}^2} + \\
& \frac{(w_{x,(x_i,y_i)} \, o_{(x_i-1,y_i)})^2}{\sigma_{\mathrm{wo}}^2} + \\
& \frac{(w_{y,(x_i,y_i)} \, o_{(x_i,y_i+1)})^2}{\sigma_{\mathrm{wo}}^2} + \\
& \frac{(w_{y,(x_i,y_i)} \, o_{(x_i,y_i-1)})^2}{\sigma_{\mathrm{wo}}^2} \Bigg]
\end{aligned} \tag{17}$$

3.3.7. Energy Derived from the Correlation Between Gas Concentration and Wind Flow (E_{gw})

All previous factors and their associated energies have been defined on either the gas or the wind maps separately, with no flow of information from one to the other. With this energy term we want to model the mutual influence among these two magnitudes standing for how the gas distribution gets shaped into plumes by the wind flow, or equitably, how the presence of a gas plume reveals information about the general direction of the wind in the area.

Modeling the correlation that exists between gas and wind flow is not a trivial task as it depends on many parameters such as the turbulence index, temperature, or the nature of the gas being released, among others. Notwithstanding, from the analysis of several and varied experiments related to gas dispersion (see for example those of Meroney [37], Chov et al. [38] or Heist et al. [39], we have extracted a simple, yet relevant, connection between both magnitudes. Because gas plumes get stretched by wind, their gas concentration gradient is, on average, perpendicular to the main wind direction as illustrated in Fig. 4.

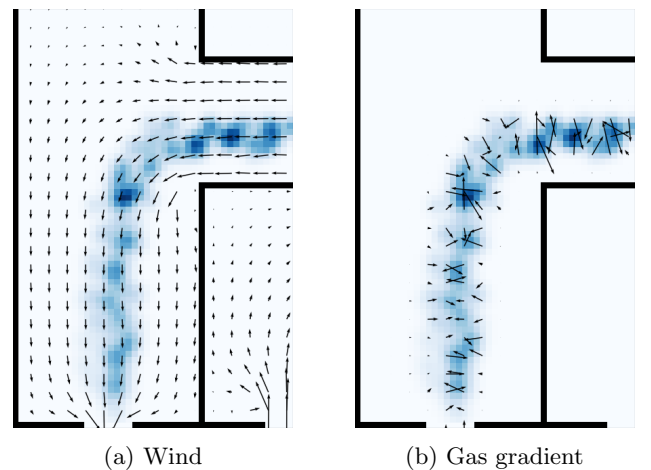(a) Wind   (b) Gas gradient

Figure 4: Closeup of an indoor environment near the wind exit point showing a CFD wind simulation with Open Foam [36] (black arrows), and a GADEN simulation [20] of a released gas distribution (blue intensity cloud). As can be observed, the gradient of the gas distribution in (b) is mostly perpendicular to the wind direction (a) along the whole plume excluding some noise due to concentration fluctuations.

7

But gas gradients may also be the result of diffusion, so we also need to consider the case of absence of wind flow. Put together, we get that for any given location, either the wind speed is zero, or the angle θ between the wind vector and the gas concentration gradient tends to 90^o. Hence, we derive an energy function to account for this correlation as:

$$E_{\text{gw}} = \frac{1}{2} \sum_{i=1}^{N} \frac{(\nabla g_i \cdot \mathbf{w}_i)^2}{\sigma_{\text{gw}}^2} \qquad (18)$$

where the gradient ∇g_i is computed by comparing the gas concentration of cell i with the ones that are above and to the right of it (see Fig. 2g), and σ_{gw}^2 is the tolerance parameter that controls the impact of this energy term.

3.3.8. Energy to Set a Default Value (E_d)

To ensure the convergence of our approach in areas where observations have not yet been gathered (preventing the problem from being ill-defined), it is necessary to set a *default* value for the gas concentration and the wind vector. This may be the case of physically isolated areas (i.e. completely surrounded by obstacles), or cells too far away from the locations where observations have been taken and therefore not under the influence of other energies.

Assuming we have no prior information about the values of the gas concentration or the wind speed, the most common approach is to resort to uniform distributions: $[0, \infty)$ for g_i and $(-\infty, \infty)$ for both w_{x_i} and w_{y_i}, respectively. But because we want all energy functions to be Gaussian to ensure a least-squares solution, we assume that both can be modeled as factors (Fig. 2h) with energy function:

$$E_{\text{d}} = \frac{1}{2} \sum_{i=1}^{N} \left[\frac{g_i^2}{\sigma_{\text{d}}^2} + \frac{\| \mathbf{w}_i \|^2}{\sigma_{\text{d}}^2} \right] \qquad (19)$$

where the value of σ_{d}^2 controls how close to zero should be the default value. In practice, this parameter will be very large to indicate that we are not enforcing any particular default value.

4. Maximum a Posteriori Estimation

In this section we describe how to derive the MAP estimate from the joint probability distribution $p(\mathbf{m}, \mathbf{z}, \mathbf{o})$, and show how this MAP becomes a least-squares problem given the proposed GMRF formulation. Moreover, present how to retrieve the uncertainty associated to the estimation, and then discuss the selection of the most important parameters and their impact on the estimate.

4.1. Derivation

Obtaining the MAP estimate $\hat{\mathbf{m}}$ for our problem is equivalent to say that we seek to maximize the posterior probability $p(\mathbf{m}|\mathbf{z}, \mathbf{o})$, that is, to obtain the most likely value of the gas and wind maps $\mathbf{m} = \{\mathbf{g}, \mathbf{w}\}$ conditioned to all available data, i.e. the measurements $\mathbf{z}$ and the occupancy information $\mathbf{o}$.

Given we are only interested in the location of this posterior's maxima and not in its value, and since $p(\mathbf{m}|\mathbf{z}, \mathbf{o}) \propto p(\mathbf{m}, \mathbf{z}, \mathbf{o})$ we can establish a direct relation between the MAP and the energy functions previously introduced in Section 3.3.1. By taking the negative logarithm over the joint, the complete energy function $E(\mathbf{m})$ becomes the well known least-squares form of a GMRF inference problem [40], which in our case reads:

$$\hat{\mathbf{m}} = \arg \max_{\mathbf{m}} p(\mathbf{m} \mid \mathbf{z}, \mathbf{o})$$
$$= \arg \min_{\mathbf{m}} E(\mathbf{m}), \qquad (20)$$

where the total energy is defined as:

$$E(\mathbf{m}) = E_{\text{gz}} + E_{\text{gr}} + E_{\text{wz}} + E_{\text{wr}}$$
$$+ E_{\text{wc}} + E_{\text{wo}} + E_{\text{gw}} + E_{\text{d}}. \qquad (21)$$

Because all the energy terms in Eq. (21) are defined as weighted quadratic errors (i.e. $E_k = \Lambda_k r_k^2$), we can express the above equation as the sum of squared residuals weighted by an information matrix,

$$E(\mathbf{m}) = \frac{1}{2} \sum_{k=1}^{M} \Lambda_k r_k^2 = \frac{1}{2} \mathbf{r}^{\mathsf{T}} \mathbf{\Lambda} \mathbf{r}, \qquad (22)$$

where the vector $\mathbf{r}$ corresponds to the residuals of all the M energies (i.e. their deviation from the expected values of the underlying Gaussian distributions), and where $\mathbf{\Lambda}$ is a diagonal matrix of size $M \times M$.

The minimum of the least-squares formulation in Eq. (22) can then be found by solving the *Gauss-Newton* method equations [41]:

$$\underbrace{(\mathbf{J}^{\mathsf{T}} \mathbf{\Lambda} \mathbf{J})}_{\text{Hessian } \mathbf{H}} \mathbf{\Delta m} = - \underbrace{\mathbf{J}^{\mathsf{T}} \mathbf{\Lambda} \mathbf{r}}_{\text{Gradient } \mathbf{d}} \qquad (23)$$

where $\mathbf{J} = \delta \mathbf{r} / \delta \mathbf{m}$ is the Jacobian of the residual $\mathbf{r}$ with respect to $\mathbf{m}$.

As opposed to previous approaches that only accounted for GDM[1] or WFM[4] and where a closed solution could be found, introducing the energy term that correlates both magnitudes E_{gw} makes the problem non-linear with respect to the map $\mathbf{m}$. The latter entails that a numerical approximation must be found by successive iterations using the *Newton Method* [42],

$$\mathbf{m}_i \leftarrow \mathbf{m}_{i-1} + \mathbf{\Delta m}_i, \qquad (24)$$

$$\mathbf{\Delta m}_i = -\mathbf{H}_i^{-1} \mathbf{d}_i \qquad (25)$$

until the halting criterion $\|\mathbf{\Delta m}_i\| < R$ is met for some arbitrary chosen error bound R and starting condition $\mathbf{m}_0$. Notice that this implies that the Hessian $\mathbf{H}$ and the gradient $\mathbf{d}$ must also be recomputed at each iteration, as they are both functions of $\mathbf{m}$.
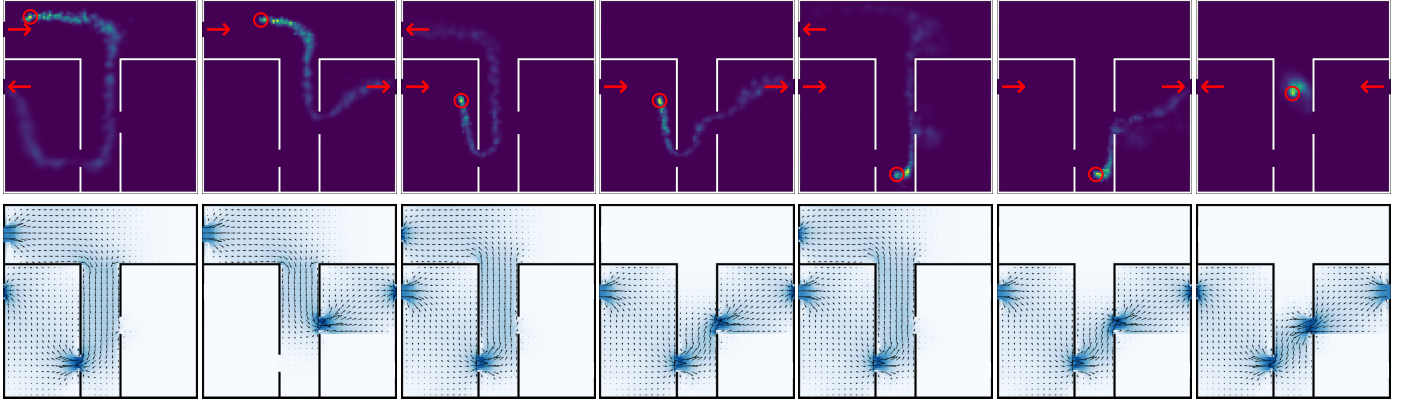
Figure 5: Virtual test scenarios employed to train the parameters for the GMRF estimator. All scenarios were generated in an indoor-like environment, using computer fluid dynamics (CFD) to simulate the wind conditions (bottom images), and the filament-based particle simulator GADEN for the gas distributions (top images). These account for different initial conditions (e.g. wind speed, gas-source location, amount of released gas, etc.) to emulate realistic situations. Please note that this figures show an instantaneous snapshots, as the environments were continuous-time simulations. The red circles denote the location of the simulated source, and the arrows the wind inlets and outlets.

4.2. Recovering the Uncertainty

To conveniently interpret the results of the estimated map $\hat{\mathbf{m}}$ at any given time, it is necessary to pair the values estimated for the gas concentration and wind vector with a confidence value or uncertainty measure. Obtaining such measure on a GMRF model is achieved by computing the inverse of the Hessian [43] and retrieving the diagonal values, as each diagonal element in $\mathbf{H}^{-1}(i,i)$ corresponds to the *standard error of the mean* [41] of the variables in the estimations of $\hat{\mathbf{m}}$, that is, $\hat{\mathbf{g}}$, $\hat{\mathbf{w}}_x$ and $\hat{\mathbf{w}}_y$.

Because the size of $\mathbf{H}$ increases quadratically with the number of cells in the environment, it becomes very computationally expensive to invert the whole matrix. If the application only requires the uncertainty of the gas estimate, but not the wind, it is possible to offset this cost by carefully grouping the variables in the Jacobian from Eq. (23) during implementation such that only the pertaining submatrix in $\mathbf{H}$ has to be inverted.

4.3. Parameters of the Estimator

The resolution of the least-squares problem derived in our proposal depends on eleven parameters, namely: the cell size c of the grid map $\mathbf{m}$, the time-dependant information-loss parameters $\sigma_{\zeta_{gz}}$ and $\sigma_{\zeta_{wz}}$ that codify how observations lose importance as they get older, and the eight Gaussian variances corresponding to the different energies introduced in Section 3.3). Next, we analyze each of these parameters, focusing on their impact in the estimation process.

- **c** : The cell size of the grid used to encode the map $\mathbf{m}$ determines the resolution at which the different predictions are made. Its mayor impact is on the computational cost, rising considerably for large maps with high resolution. Its optimal value is therefore a trade-off between the size of the environment, the desired resolution, and the computational power available.

- $\sigma_{\zeta_{gz}}$ and $\sigma_{\zeta_{wz}}$: These two parameters control the weight of gas and wind measurements over time. Low values translate into observations being trusted for long periods of time, while high values will quickly disregard them in favor of newer ones. The selection of these parameters is highly dependent on the desired application. For example, if the goal is to obtain an estimation that represents the average gas and wind distributions, low values should be chosen to include all past observations. Whereas, if the application requires a dynamic estimation that represents the latest state of the environment, high values should be chosen instead to only account for the most recent data. For the experiments presented in the next section, we consider a linear increasing uncertainty limiting the influence of gathered observations to 10% of their original weight after 10 minutes because the environments are not prone to fast changes.

- σ_i $\forall i$ *in* E : Each of the energy functions defined in Eq. (21) is weighted by a variance parameter that controls the contribution of its energy term to the global estimation. The optimal value of these parameters is not straightforward as many of them have not a direct physical interpretation. For example, the weights for the gas-sensor and anemometer measurement noise can be set from the tolerance provided by the manufacturer of the corresponding transducers, but the weight of the regularization energies between neighbor cells, or that that controls how much the wind direction may deviate from a perfect tangent to obstacles, are not so straightforward.

We present here the selection of these weights by resorting on an iterative optimization process over simulated data. Concretely, we employ GADEN [20], a finite-element computer fluids simulator (refer to Section 5.2 for more information) to generate the training environment shown in Fig. 5. This figure illustrates

a generic indoor-like setting over which we considered different setups and environmental conditions, ranging from well formed gas-plumes to turbulent and uneven gas distributions. Over the simulated data we performed a balanced [44] stochastic gradient-descent optimization that minimizes the mean squared error (MSE) between prediction and ground-truth. The values obtained from this optimization are depicted in Table 1.

It must be stressed that, although this approach is extremely slow given the high number of parameters to optimize, in general, the results are quite stable to some changes in the environmental set-up. This allows applying the values presented in Table 1 in other scenarios that are different from the training set. An exception to this involves changes in the cell size as this parameter has a notable influence in the others, being advisable to re-optimize.

5. Experiments and Discussion

This section covers multiple experiments, both simulated and real, aimed to evaluate and compare the proposed approach under different set-ups. We consider a mobile robot equipped with an e-nose and an anemometer [12] whose mission is to explore the environment and provide an estimation of the gas distribution in the shortest possible time.

To provide quantitative results, as well as to allow a stringent comparison to other GDM methods, we first present a set of simulated experiments where a ground-truth is available to assess the exact accuracy of the estimated maps. Afterwards, we show the results of a real experiment focusing on the consistency of the predictions and the time employed to provide them. Finally, we alsom provide a performance analysis that accounts for the the sensors noise level, the computational cost and the influence of the number of observations.

Parameter	gmrf-g	kdm-vw	gmrf-gw
cell size	0.1 m	0.1 m	0.1 m
σ_{gz}	0.1	-	0.1
σ_{gr}	1.128	-	1.128
σ_d	10000	-	10000
$\sigma_{\zeta_{gz}}$	0.012	-	0.012
σ_{wz}	-	-	0.1
σ_{wr}	-	-	0.825
σ_{wc}	-	-	0.048
σ_{wo}	-	-	0.22
σ_{gw}	-	-	0.12
$\sigma_{\zeta_{wz}}$	0.012	-	0.012
σ_0	-	0.38	-
σ_Ω	-	4.956	-
γ	-	0.33	-

Table 1: Parameter values for the three tested GDM methods.

(a) Scenario I, 15m x 5m (simulation)

(b) Scenario II, 11m x 10m (simulation)

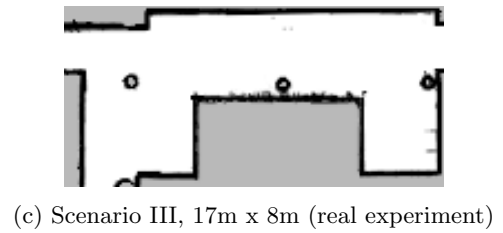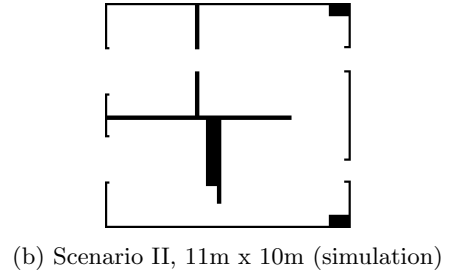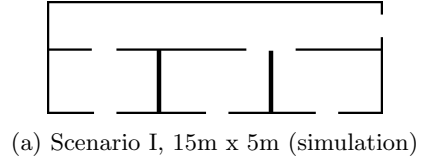(c) Scenario III, 17m x 8m (real experiment)

Figure 6: Floor plan of the test scenarios. Scenario I (a) was simulated in 2D, Scenario II (b) in 3D, and Scenario III (c) is the robot's navigation map of where we conducted the real experiment.

5.1. Experimental Setup

Both real and simulated experiments have been conducted in the indoor scenarios shown in Fig. 6 where a mobile robot is tasked to explore the environment, following a predefined fixed path, to provide an estimation of the gas distribution. We assume the robot has access to the occupancy grid-map of the environment, which we consider static (i.e. the occupancy probability is not updated along the experiment), and where inlets and outlets at the boundaries of the experimental area have been marked as open ($o_i = 0$). The latter allows the estimated wind to enter or exit from any direction. In other words, we assume all inlets/outlets could be open, and leave it to our GDM method to figure out which wind currents make the most sense given the input data.

For the environmental conditions of all three scenarios, we consider wind currents of around 1 m/s (typical of indoor scenarios [45]) and a single gas source releasing ethanol at a fixed location (dependent on the scenario). Our selection of ethanol is justified by it being heavier than air and accumulating close to the floor, which in turn makes it easier to detect for a wheeled ground robot. Also, ethanol is safe to handle without special equipment and can be easily removed by venting the environment for a few minutes, which was particularly convenient for the real experiments.

Lastly, in order to evaluate the advantages and disadvantages of the proposed method, from now on referred as gas-wind GMRF distribution modeling (GW-GMRF), we present a detailed comparison with two other state-of-
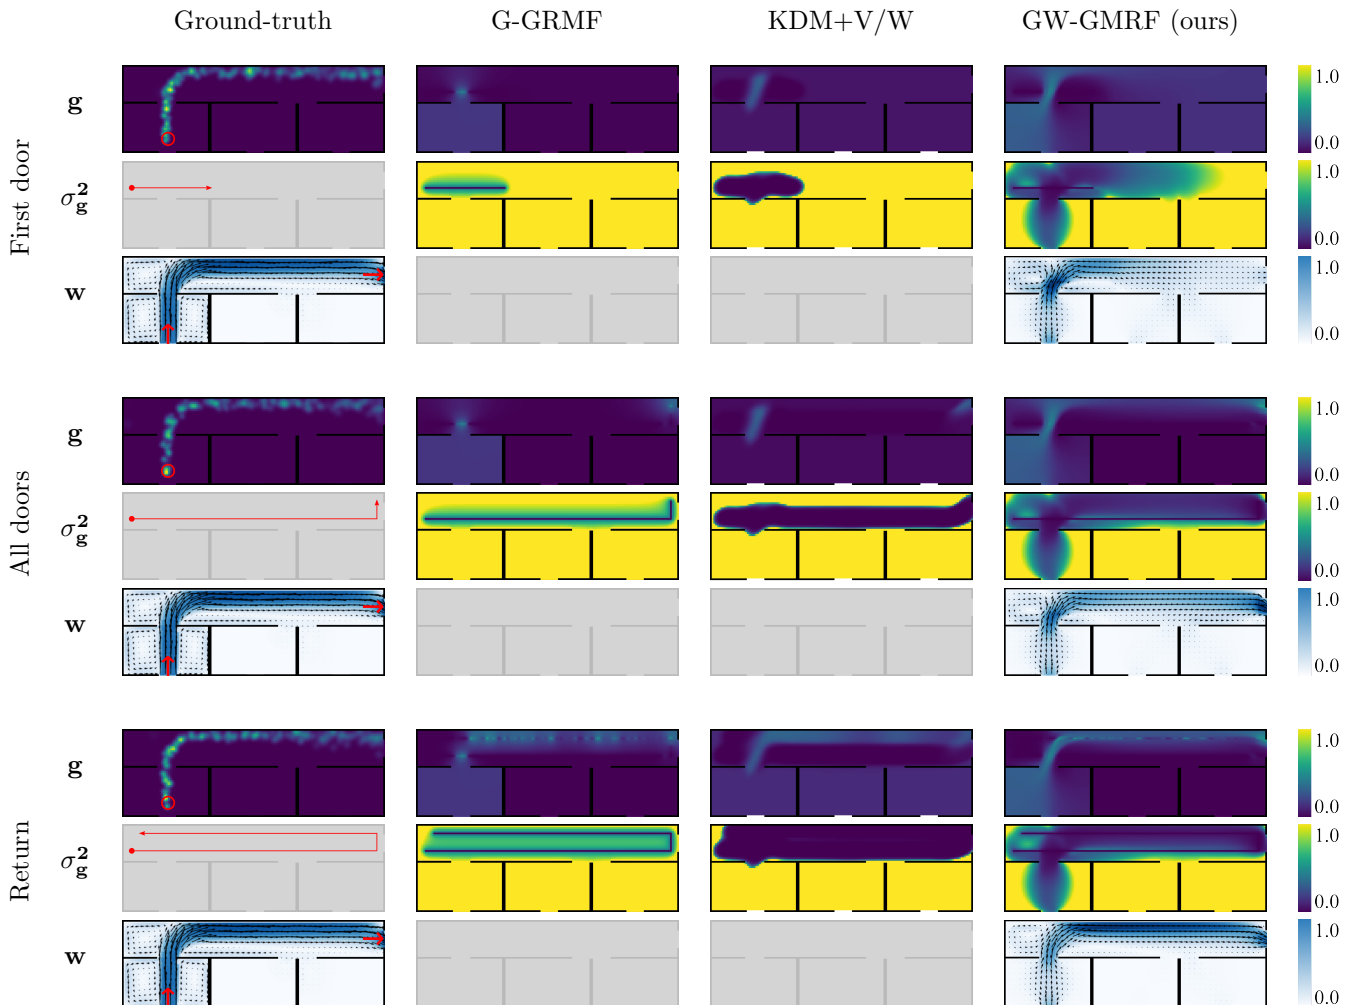
Figure 7: Simulation of Scenario I for a quick exploration path along the main corridor, illustrating the GDM estimates at the beginning, middle, and end of the exploration. The instantaneous ground-truths show in red the location of the gas source, the exploration path, and the wind inlets and outlets. The scale for the gas and gas-uncertainty maps are limited to 1 *gas unit* and 1 *gas unit*$^{-2}$ respectively, and 1 m/s for the wind map.

the-art GDM estimators, namely gas distribution modeling based on GMRF (G-GMRF) [1] and KDM+V/W [28]. The former stems from a previous work and can be regarded as a much simple version of our current approach since it does not consider any wind information to estimate the gas distribution, though it relies on a similar GMRF formulation. The latter estimator, KDM+V/W, is a kernel-based method that uses a multivariate Gaussian weighting function to model the information provided by the e-nose and the anemometer. Wind data is applied to regulate the shape of the kernel, which in turn controls the amount of extrapolation.

To quantify the performance of each GDM method we compute the RMSE of the gas concentration estimates as well as the negative logarithm of the maximum likelihood (NLML), a more relevant magnitude that also accounts for their associated uncertainty [46]. In this context, we must note that KDM+V/W does not provide a statistically significant uncertainty margin for its predictions, but a confidence interval α which ranges from 0 (not confident

at all) to 1 (highly confident). This entails a problem for comparison since these magnitudes are not equivalent and then, cannot be compared. This prevents us from computing the NLML for KDM+V/W. However, to enable a qualitative comparison in subsequent illustrations, we use the formula $\Sigma_{KDM} = \frac{1}{\alpha} - 1$ to map the provided confidence value to the range $[0, \inf]$, and treat it akin to a gas-estimate uncertainty.

5.2. Simulation Experiments

Two different scenarios have been considered for running the simulated experiments (see Fig. 6a and Fig. 6b). Scenario-I is a long corridor with adjacent rooms intended to study gas-plumes formed by strong wind currents in 2D, whereas Scenario-II is an office-like environment meant to test the GDM methods under generally slower but more turbulent conditions in 3D. In both cases, we first compute a ground-truth with OpenFOAM [47] for a reliable CFD simulation of the wind, and then simulate with GADEN [20] a continuous gas-filament representation of
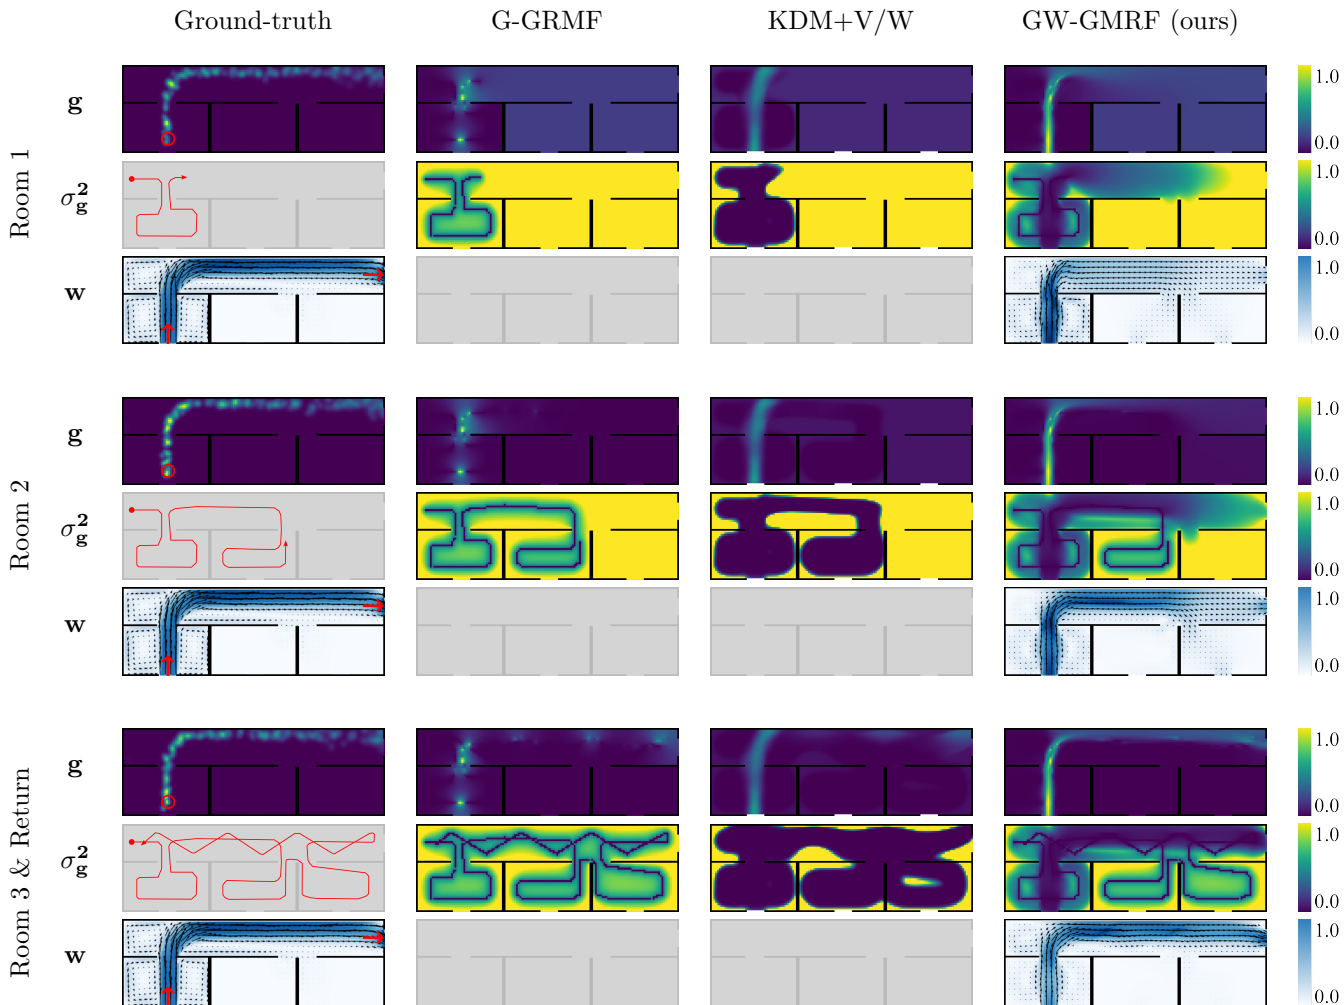
11

Figure 8: Simulation of Scenario I at the beginning, middle, and end of a complete exploration of all rooms. The instantaneous 2D ground-truths show in red the location of the gas source, the robot's path, and the wind inlets and outlets. The scale for the gas and gas-uncertainty maps are limited to 1 *gas unit* and 1 *gas unit^{-2}* respectively, and 1 m/s for the wind map.

the gas distribution [48]. The wind conditions were determined by closing all doors and windows except two: one that acts as a wind inlet at 1 m/s and another as isobaric outlet. It must be stressed that simulations are conducted in continuous time, that is, both gas and wind may contain dynamic eddies or *puffs*. For this reason we do not start the GDM experiment until the initial transient-response has settled (usually 30 seconds after simulation start) to ensure impartial conditions with respect to the robot's starting position, and normalize the concentration for convenience. Lastly, we do not corrupt the sensor measurements at this point, but discuss in Section 5.4 its effect on the gas estimates. Gas and wind samples were taken at a rate of 10 Hz and the robot moved at 0.5 m/s.

5.2.1. Scenario-I

Scenario I is a 2D indoor environment meant to experiment with well formed and laminar gas plumes. We force a wind flow from the bottom-left room towards the right side of the corridor, and place the gas source in front of

the inlet window which leads to the plume seen in Fig. 7. Because we want to test how much information is needed to delimit the general shape of the distribution, we run this experiment twice: the first time the robot is tasked with a quick exploration of the main corridor (Fig. 7), and the second time with a thorough exploration of all rooms (Fig. 8). In this way we can test whether the additional wind information allows for an accurate estimate with less samples, and if so, how our method compares to the others.

For the exploration path that covers the corridor only, all three gas estimates reach a similar RMSE with regard to the ground-truth (Fig. 9, left column). However, GW-GMRF is slightly faster (i.e. needs less data) than the others to do so. Notice how in Fig. 7, after the robot crosses the first door, our method predicts various possible wind-paths that may carry gas towards all possible exits, but discards the wrong ones as soon as it passes in front of the other doors, allowing it to predict the general shape of the plume without having actually measured it yet. In contrast, G-GMRF and KDM+V/W do not model the plume
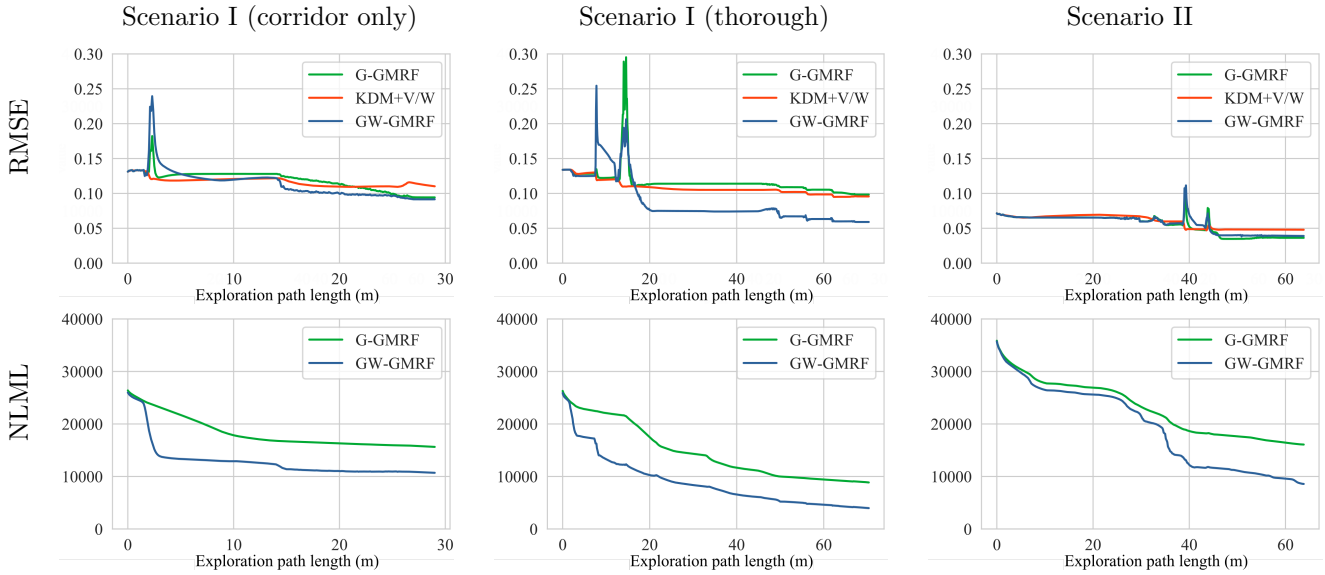
Figure 9: Comparison of the RMSE (top) and NLML (bottom) between the estimated gas maps and the ground-truth for the tested GDM methods depending on the length of the exploration path. Note that NLML is not computed for KDM+V/W because it provides no statistically meaningful uncertainty margin. For both metrics, lower values are better (i.e. better approximation of the ground-truth).
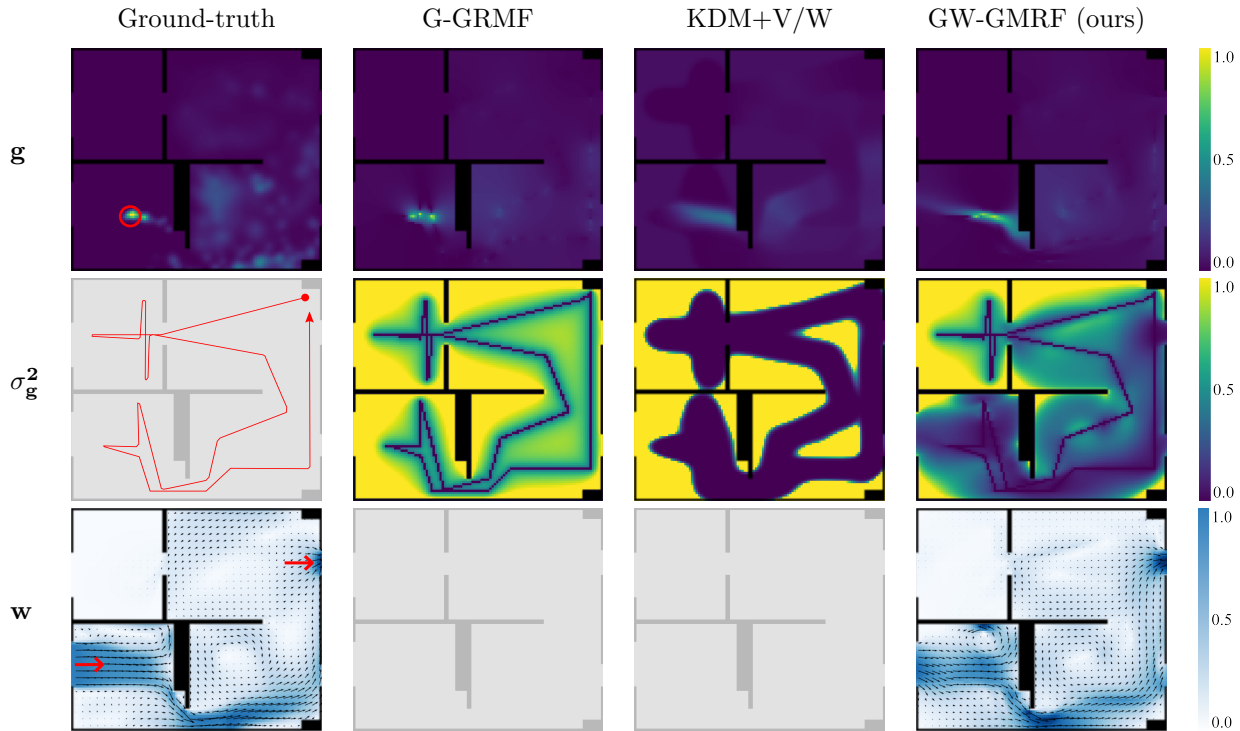

Figure 10: Simulated exploration of Scenario II. The ground-truth shows a slice of the 3D gas and wind distributions at the height at which the virtual robot carries its sensors, on top of which the location of the gas source, the robot's path, and the wind inlets and outlets are highlighted in red. The scale for the gas and uncertainty maps are limited to 1 *gas unit* and 1 *gas unit*$^{-2}$, and 1 m/s for the wind map.

until the end when the robot drives through it on its return path. Likewise, during the thorough exploration and after entering the first room in Fig. 8, KDM+V/W and GW-GMRF provide both good approximations of the head of the plume, but whilst KDM+V/W is very conservative and limits its prediction to the inspected area (notice its uncertainty map), GW-GMRF can already extrapolate to

the middle section of the corridor (the wind has nowhere else to go). The downside of this amount of extrapolation is that the RMSE of GW-GMRF goes up in both cases (Fig. 9) because it predicts gas around all possible exits the wind could take, which as shown in the ground truth, is wrong. But because it assigns very high uncertainty to these locations, the prediction's NLML does actually im-

13

(a) Photo of Scenario III
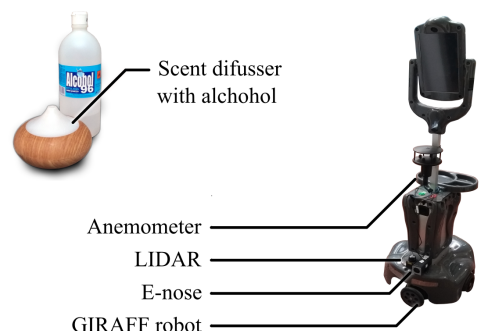
(b) Gas source and Giraff robot with sensors

Figure 11: Snapshot of the real experiment (a). Figure (b) shows the gas source, an ultrasonic scent-diffuser filled with ethanol, as well as the location of the sensors on the Giraff mobile robot employed to gather the gas-wind data.

prove faster in both cases than that of G-GMRF. In fact, when we compare the results between both paths, the additional wind information allows GW-GMRF to achieve a similar prediction in terms of RMSE and NLML after exploring only 3 m of the corridor as G-GMRF after following the thorough path for 30 m. Which makes it suited for quick response applications without sacrificing on accuracy for longer explorations.

5.2.2. Scenario-II

This scenario is intended to represent realistic 3D gas distributions as they may occur in indoor environments. Our goal is to test the performance of our GDM method when, as it is usually the case in practice, the robot carries its sensors at a fixed height and can therefore not detect the gases flowing under or over them. For this purpose we have programmed the virtual robot to sample at 0.7 m from the floor, and because the output of the three GDM methods are 2D only, compare the results against the cross-section of the ground-truth at this same height.

As shown in Fig. 10, an exploration of the environment provides very similar gas maps in all three cases. Their peculiarities are however more apparent during intermediate stages of the experiment, as was already the case for Scenario I. At the beginning and before any significant gas concentration is detected, there is virtually no difference between them in terms of their RMSE (Fig. 9, right column). But once the robot enters the lower half of Scenario II, where the gas is accumulating (after about 30 m of traveled path, GW-GMRF extrapolates with all the wind information it has gathered so far the gas concentration in the remainder of the map. As a result, its able to reduce its NLML distance to the ground truth much faster than G-GMRF (Fig. 9, right column), even though its RMSE starts to fluctuate while getting close to the source before settling at the same final value than G-GMRF. Hence, the advantage of GW-GMRF in this type of situations is not an improved accuracy of the gas maps, but (i) the ability to make good predictions at a much earlier stage of the exploration, and (ii) to do so for a much wider surface area as denoted by the uncertainties in Fig. 10.

5.3. Real Experiment

For the real experiment we chose Scenario III (Fig. 11a) because it offers slow but consistent wind currents in the range of 1 m/s. It is a short corridor that leads outside through the doors on both sides (see Fig. 6c) which were open at the same time. As shown in Fig. 11b, the gas source for the experiment was a scent diffuser filled with about 75% water and 25% alcohol, and the robot was an autonomous Giraff system [49] equipped with a portable e-nose [2] and an anemometer. As for the sensors, we placed the e-nose as close as possible to the bottom because ethanol is heavier than air, but were forced to mount the anemometer about 55 cm higher to avoid wind-shadows caused by the robot's shape. This mismatch should however have no mayor impact on the results, as we measured the wind-speed to be about the same at both heights. Lastly, we must remark that the software on the robot aided in a better response time for the e-nose [50] and also compensated for the wind component caused by the robot's own motion using LIDAR-based odometry [51]. The robot's speed was limited to 0.5 m/s, and samples were taken at a rate of 10 Hz.

Because there is no ground-truth for his scenario, we chose a reasonably simple setup. We placed the gas source in the middle of the corridor such that the wind caused the gas to accumulate only on one half of the map. Accordingly, the experiment in Scenario III was meant to test if our method could correctly determine that, if the robot stars on the side with gas, there should be none once it drives next to the source.

The exact path followed by the robot is shown on the left-most column of Fig. 12. It starts on the bottom-left, a location with little to no gas, then passes through the gas plume and next to the source, and finally stops after one meter after it. As shown in the same figure, the estimated gas maps for all three methods are very similar in shape and concentration. Their only major difference lies in their associated uncertainties, where GW-GMRF stands out because it covers a much wider surface area. In fact, if we use the uncertainties as a mask and only plot the gas estimates as in the last row of Fig. 12, we can appreciate
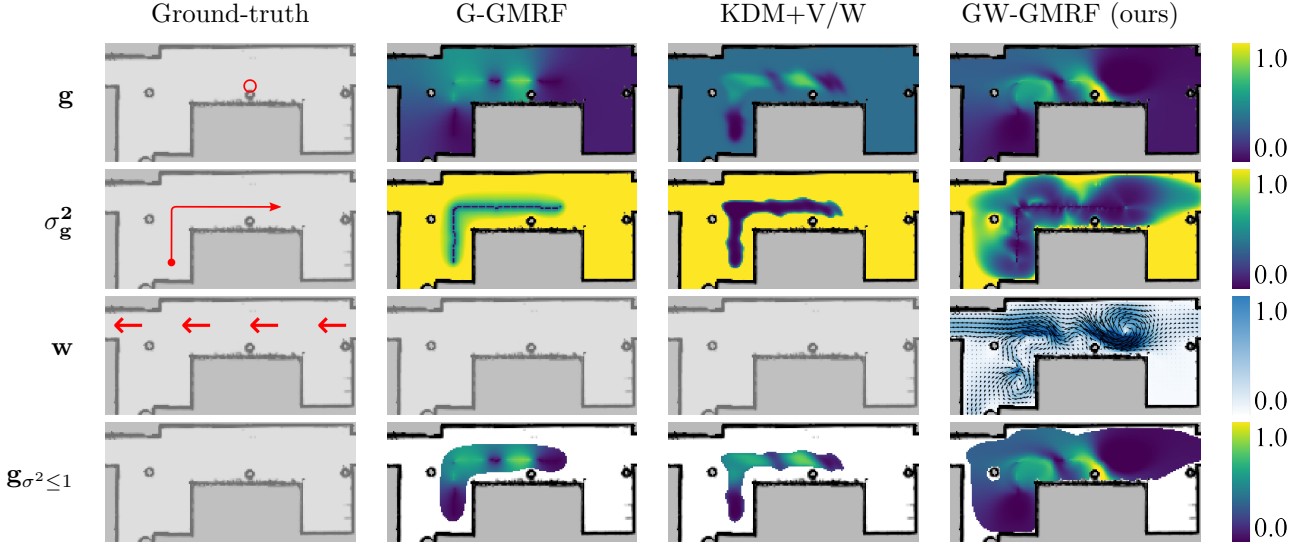
14

Figure 12: Estimated gas maps for the real experiment. Although no exact ground truth is available, the location of the gas source (red circle) and the dominant wind direction (red arrows) should cause the gas to accumulate only on the left side of the environment. The robot's explorations goes from left to right and ends after passing next to the source to test whether the estimated gas map can predict that there should be no gas on the unexplored right side. The last row shows the GDM estimates using the uncertainty as a cell-wise mask with a threshold of 1, $\mathbf{g}_{\sigma^2 \leq 1} = \mathbf{g} \odot (\sigma_{\mathbf{g}}^2 \leq 1)$.

how our method extrapolates considerably more than the other two for the same exploration path; which appears to be correct given the direction of the wind and the position of the gas source.

5.4. Performance and Resilience Experiments

Lastly, we have performed several experiments to analyze GW-GMRF in terms resilience to sensor noise as well as overall behavior, and again, compare it against the other GDM methods for reference. We have reused for this purpose the simulation setup from Section 5.2 for its ground-truth, but to keep this section short only show the results of Scenario II (which we deem more complex) because the results for Scenario I followed the same trend and provided no additional information.

5.4.1. Tolerance to Sensor Noise

Given that GW-GMRF method relies heavily on extrapolation to predict the gas and wind fields within unvisited rooms, sensor noise could potentially lead to unrealistic estimations. To test for this possibility we have repeated the experiment from Scneario II, but added increasing amounts of Gaussian noise to the gas-sensor and the anemometer. First, we tried with the the original parameter selection from Table 1, but later also repeated the experiment for a configuration meant to compensate for the increased noise (e.g. decreasing weight on σ_{gz} and σ_{wz}) As depicted in Fig. 13, the RMSE of all three methods degrades by a similar amount for small input noise, but becomes more noticeable for GW-GMRF for higher noise values. Likewise, the NLML does also degrade faster for GW-GMRF than for G-GMRF, and even becomes worse in extreme cases where the standard deviation of the noise
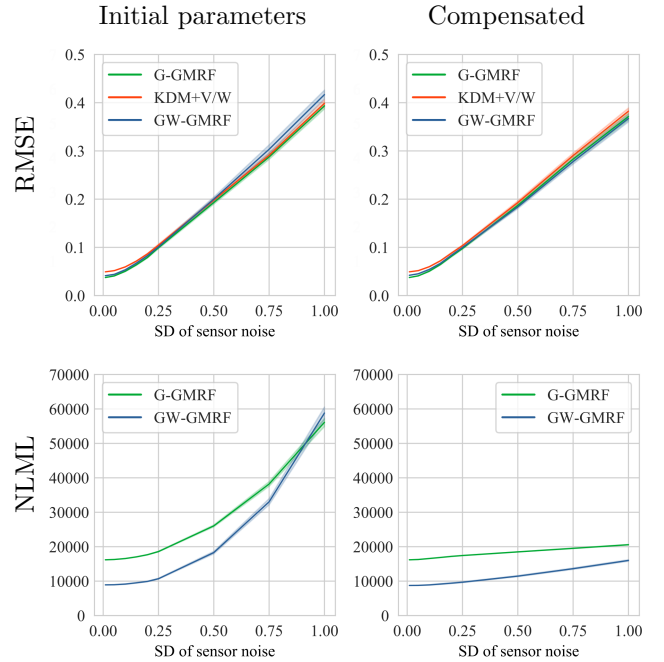

Figure 13: Effect of different sensor-noise levels on the GDM estimates of Scenario II, plotted for the initial parameter selection (left column) and for a configuration meant to compensate for the noise (right column). The shaded area around each line represents the standard deviation over 25 iterations. Lower values are better.

is as high as the ground truth. As suspected, our new method is less resilient to noise than the alternatives. Although it is only of concern for applications with extreme noise levels, and even then, it can be counteracted by readjusting the parameters of the methods as shown for estimation when the parameters are adjusted to compensate.
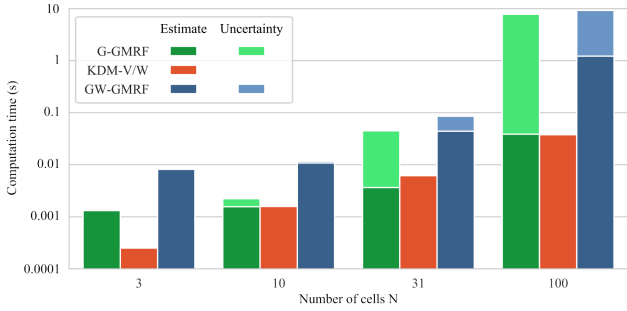
15

Figure 14: Execution time of the GDM methods depending on the total number of cells in the environment, separated into estimation of the gas distribution, and computation of its associated uncertainty. Please note that the vertical scale is logarithmic. The test was implemented in Python 3 and run on an Intel Core i7-8700 with 16 GB of RAM. Lower is better.

5.4.2. Computational Complexity

Although GW-GMRF can achieve a better estimate than the other methods during the early stages of the robot's exploration, it is considerably slower to compute. To compare their execution times, we run a series of experiments with an environment of increasing size. As a generic case, we used a square environment with obstacles around the borders only and sampled along its diagonal to ensure that the number of observations is linearly proportional to its size. As expected and shown in Fig. 14, GW-GMRF is a magnitude slower to compute than G-GMRF, and increases with exponential complexity regarding the number of cells rather than being linear like KDM+V/W. Depending on the application it might be possible to offset this cost by reducing the rate at which the map is recalculated or by lowering the map's resolution, but GW-GMRF remains prohibitive for small robots with limited on-board computation-capabilities.

5.4.3. Accuracy Dependence on Available Data

Fig. 15 shows how the amount of sensor data affects the accuracy of the predictions. For this test we have uniformly sampled Scenario II at random location in a *Monte Carlo*-like approach, starting with 10 samples and gradually increasing to 5000. As can be observed in Fig. 15a, the RMSE of all three methods follow a similar trend at the beginning, but the prediction of KDM+V/W does not improve once the environment is over-sampled (around 1000 to 1500 observations, as Scenario II has 1100 cells) whereas the GMRF-based approaches benefit from additional exploration to obtain more data as the environment changes over time. Both G-GMRF and GW-GMRF have a similar RMSE throughout, but GW-GMRF's achieves a better NLML with a fraction of the samples (e.g. 200 vs 1000). This difference is most noticeable for little observations where the ability of extrapolate with wind information is most useful.
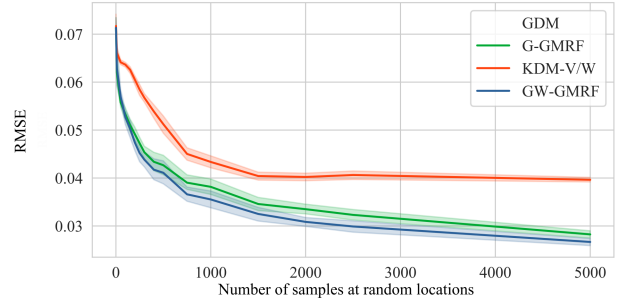
(a) RMSE vs. number of random samples
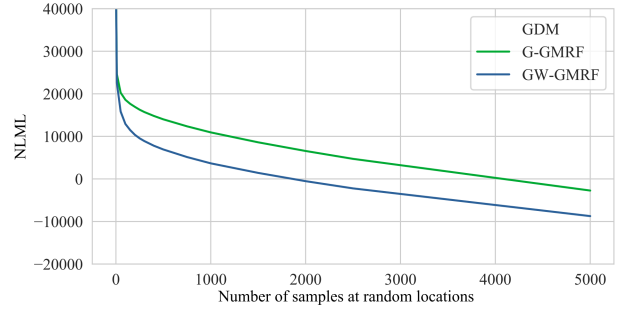
(b) NLML vs. number of random samples

Figure 15: RMSE and NLML in the estimation of the gas distribution in Scenario II for increasing number of samples at random locations. The experiment has been repeated 25 times to obtain a standard deviation (indicated by the shaded area), but because was uniform manner the variability is very low. Lower values are better.

6. Conclusions and Outlook

In this paper we presented GW-GMRF, a new GMRF-based gas distribution model for real-time applications that accounts for wind-flow modeling in the environment to boost the estimation. That is, it models the law of gas dispersion with simple but sensible rules that account for the effect of obstacles and wind currents to extrapolate the gas concentration at remote locations that were not subject to direct sampling. With this method, GW-GMRF not only provides a gas map of the environment, but also a statistically significant uncertainty for each position as well as a wind estimation map.

Experimental results have demonstrated that GW-GMRF works remarkably well in indoor environments with the presence of wind flows, and outperforms other similar methods in terms accuracy and reliability of the prediction. However, our approach is considerably more costly in terms of computational power when compared to methods like G-GMRF or KDM+V/W. This becomes especially relevant for scenarios where there is little to no wind, and consequently the ability of our method to extrapolate is hindered, not providing substantial benefits over its predecessor G-GMRF.

Moreover, GW-GMRF offers additional advantages that were not further explored in this article, including:

- Adaptation to the dynamic nature of the occupancy grid-map. The obstacle connectivity map **o** does not

16

need to be static. Our formulation is meant to deal with situation where the robot has not fully explore the environment and the cells have an occupation probability other than 0.0 or 1.0. Every time new a prediction is made, GW-GMRF uses the latest sensor data and obstacle information, for which error in previous estimation are not carried into the next. For example, the estimation can be updated with new information about the state of a door or window with other sensors (e.g. computer vision) in favor of more accurate extrapolations.

- GW-GMRF is compatible with multiple data input sources as long as each individual gas or wind sample can be uniquely assigned to a single position and carries it's own time-stamp. This can be exploited to further accelerate the map generation by using multiple sensors on separate robots, or to improve the accuracy by monitoring key locations with dedicated sensors while the robot explores the rest.

- If there a prior gas distribution of the work environment which could provide a good starting point for G-GMRF, it can be entered as the mean value of the Gaussian that comprise the default energy E_d. As a result, the output of our estimator will tend towards said prior while there is no data or it does not contradict the observations, while retaining high uncertainty until real evidence of its veracity becomes available.

Notwithstanding, we must remark that our method is not designed to achieve accuracy comparable to the slower but more precise CFD simulations. Being formulated as a GMRF that, in turn, becomes an easy to solve least-squares problem, all variables must necessarily be *Gaussian*. This is a debatable assumption that favors computation speed, but also prevents GW-GMRF from including other distributions that might be better suited to model the environment. In fact, there is a caveat in the gas estimate because of this very reason: the physical gas concentration of a cell can only be zero (i.e. perfectly clean air) or positive, yet a *Gaussian*'s tails imply that there is always a chance for negative values as well. Arguably, a Gamma or a Chi-squared distribution would have been better choices for the gas variables if they could have been implemented without subverting GW-GMRF's efficiency. Nevertheless, this assumption works exceedingly well in practice, as evidenced in the experimental sections of this work, and converges toward the ground-truth as more sensor data becomes available.

Lastly, we would like to highlight that, at this point, GW-GMRF is completely passive. It depends on the robot's exploration to acquire new data and, in turn, it provides accurate estimates, but it does not tell the robot where to explore. One possible solution would simply involve sending the robot to the most uncertain location as it should provide the most valuable data. Yet this naive approach does not account for the energy and time costs of the exploration, nor does it plan a path that goes purposefully through other uncertain locations to avoid having to explore them in the future. Hereof, one of our next steps will focus on closing the GDM-robot exploration loop for GW-GMRF in the most effective possible way.

References

[1] J. Monroy, J.-L. Blanco, J. Gonzalez-Jimenez, Time-variant gas distribution mapping with obstacle information, Autonomous Robots 40 (1) (2016) 1–16. doi:10.1007/s10514-015-9437-0.

[2] A. Gongora, J. Monroy, J. Gonzalez-Jimenez, An electronic architecture for multipurpose artificial noses, Journal of Sensors 2018.

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

[4] J. Monroy, M. Jaimez, J. Gonzalez-Jimenez, Online estimation of 2d wind maps for olfactory robots, in: International Symposium on Olfaction and Electronic Nose (ISOEN), 2017, pp. 1–3. doi:10.1109/ISOEN.2017.7968883.

[5] W. Tsujita, A. Yoshino, H. Ishida, T. Moriizumi, Gas sensor network for air-pollution monitoring, Sensors and Actuators B: Chemical 110 (2) (2005) 304–311.

[6] A. Marjovi, A. Arfire, A. Martinoli, High resolution air pollution maps in urban environments using mobile sensor networks, in: 2015 International Conference on Distributed Computing in Sensor Systems, IEEE, 2015, pp. 11–20.

[7] A. Loutfi, S. Coradeschi, A. J. Lilienthal, J. Gonzalez, Gas distribution mapping of multiple odour sources using a mobile robot, Robotica 27 (2) (2009) 311–319.

[8] P. P. Neumann, S. Asadi, A. J. Lilienthal, M. Bartholmai, J. H. Schiller, Autonomous gas-sensitive microdrone: Wind vector estimation and gas distribution mapping, IEEE robotics & automation magazine 19 (1) (2012) 50–61.

[9] M. Turduev, Y. Ataş, P. Sousa, V. Gazi, L. Marques, Cooperative chemical concentration map building using decentralized asynchronous particle swarm optimization based search by mobile robots, in: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, 2010, pp. 4175–4180.

[10] A. Marjovi, L. Marques, Multi-robot odor distribution mapping in realistic time-variant conditions, in: 2014 IEEE Int. Conference on Robotics and Automation (ICRA), IEEE, 2014, pp. 3720–3727.

[11] M. Reggente, A. J. Lilienthal, The 3d-kernel dm+v/w algorithm: Using wind information in three dimensional gas distribution modelling with a mobile robot, in: IEEE Sensors, 2010, pp. 999–1004.

[12] J. Monroy, F. Melendez-Fernandez, A. Gongora, J. Gonzalez-Jimenez, Integrating olfaction in a robotic telepresence loop, in: IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), 2017, pp. 1012–1017. doi:10.1109/ROMAN.2017.8172427.

[13] A. Gongora, D. Chaves, A. Jaenal, J. G. Monroy, J. G. Jiménez, Toward the generation of smell maps: Matching electro-chemical sensor information with human odor perception., in: APPIS, 2018, pp. 134–145.

[14] M. Portabella Arnús, Wind field retrieval from satellite radar systems, Universitat de Barcelona, 2002.

[15] P. Doubrawa, R. J. Barthelmie, S. C. Pryor, C. B. Hasager, M. Badger, I. Karagali, Satellite winds as a tool for offshore wind resource assessment: The great lakes wind atlas, Remote Sensing of Environment 168 (2015) 349 – 359. doi:10.1016/j.rse.2015.07.008.

[16] L. Landberg, L. Myllerup, O. Rathmann, E. L. Petersen, B. H. Jørgensen, J. Badger, N. G. Mortensen, Wind resource estimation—an overview, Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology 6 (3) (2003) 261–271.

[17] G. A. Grell, J. Dudhia, D. R. Stauffer, A description of the fifth-generation penn state/ncar mesoscale model (mm5). ncar technical note, National Center for Atmospheric Research, Boulder.

[18] R. Laprise, D. Caya, G. Bergeron, M. Giguère, The formulation of the andré robert mc2 (mesoscale compressible community) model, Atmosphere-Ocean 35 (sup1) (1997) 195–220.

[19] S. Kato, S. Murakami, New ventilation efficiency scales based on spatial distribution of contaminant concentration aided by numerical simulation, American Society of Heating, Refrigerating and Air-conditioning Engineers (ASHRAE) Transactions (1988) 09–330.

[20] J. Monroy, V. Hernandez-Bennetts, H. Fan, A. Lilienthal, J. Gonzalez-Jimenez, Gaden: A 3d gas dispersion simulator for mobile robot olfaction in realistic environments, MDPI Sensors 17 (7: 1479) (2017) 1–16. doi:10.3390/s17071479.

[21] G. Kowadlo, R. A. Russell, Improving the robustness of naïve physics airflow mapping, using bayesian reasoning on a multiple hypothesis tree, Robotics and Autonomous Systems 57 (6-7) (2009) 723 – 737.

[22] L. Rodriguez, J. A. Cobano, A. Ollero, Wind field estimation and identification having shear wind and discrete gusts features with a small uas, in: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2016, pp. 5638–5644.

[23] V. H. Bennetts, T. P. Kucner, E. Schaffernicht, P. P. Neumann, H. Fan, A. J. Lilienthal, Probabilistic air flow modelling using turbulent and laminar characteristics for ground and aerial robots, IEEE Robotics and Automation Letters.

[24] J.-L. Blanco, J. Monroy, J. Gonzalez-Jimenez, A. Lilienthal, A kalman filter based approach to probabilistic gas distribution mapping, in: Proceedings of the 28th Annual ACM Symposium on Applied Computing, 2013, pp. 217–222. doi:10.1145/2480362.2480409.

[25] J. Monroy, J. Gonzalez-Jimenez, Towards odor-sensitive mobile robots, in: Rapid Automation: Concepts, Methodologies, Tools, and Applications, IGI Global, 2019, pp. 1491–1510.

[26] C. Stachniss, C. Plagemann, A. J. Lilienthal, Learning gas distribution models using sparse gaussian process mixtures, Autonomous Robots 26 (2-3) (2009) 187–202.

[27] A. J. Lilienthal, M. Reggente, M. Trincavelli, J. L. Blanco, J. Gonzalez, A statistical approach to gas distribution modelling with mobile robots-the kernel dm+ v algorithm, in: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE, 2009, pp. 570–576.

[28] M. Reggente, A. J. Lilienthal, Using local wind information for gas distribution mapping in outdoor environments with a mobile robot, in: SENSORS, 2009 IEEE, IEEE, 2009, pp. 1715–1720.

[29] J.-G. Li, B. Sun, F.-L. Zeng, J. Liu, J. Yang, L. Yang, Experimental study on multiple odor sources mapping by a mobile robot in time-varying airflow environment, in: 2016 35th Chinese Control Conference (CCC), IEEE, 2016, pp. 6032–6037.

[30] D. Portugal, R. P. Rocha, Extracting topological information from grid maps for robot navigation., in: ICAART (1), 2012, pp. 137–143.

[31] P. Clifford, Markov random fields in statistics, Disorder in physical systems: A volume in honour of John M. Hammersley 19.

[32] N. M. Nasrabadi, Pattern recognition and machine learning, Journal of electronic imaging 16 (4) (2007) 049901.

[33] J. Katz, A. Plotkin, Low-speed aerodynamics, Vol. 13, Cambridge university press, 2001.

[34] D. A. Kaminski, M. K. Jensen, Introduction to thermal and fluids engineering, Wiley New York, 2005.

[35] J. M. Hyman, M. Shashkov, Natural discretizations for the divergence, gradient, and curl on logically rectangular grids, Computers & Mathematics with Applications 33 (4) (1997) 81–104.

[36] H. Jasak, A. Jemcov, Z. Tukovic, et al., Openfoam: A c++ library for complex physics simulations, in: International workshop on coupled methods in numerical dynamics, Vol. 1000, IUC Dubrovnik Croatia, 2007, pp. 1–20.

[37] R. N. Meroney, Wind-tunnel experiments on dense gas dispersion, Journal of hazardous materials 6 (1-2) (1982) 85–106.

[38] F. K. Chow, P. W. Granvold, C. M. Oldenburg, Modeling the effects of topography and wind on atmospheric dispersion of co2 surface leakage at geologic carbon sequestration sites, Energy Procedia 1 (1) (2009) 1925–1932.

[39] D. Heist, S. Perry, L. Brixey, A wind tunnel study of the effect of roadway configurations on the dispersion of traffic-related pollution, Atmospheric Environment 43 (32) (2009) 5101–5111.

[40] K. Madsen, H. B. Nielsen, O. Tingleff, Methods for non-linear least squares problems (1999).

[41] J.-A. Fernández-Madrigal, Simultaneous Localization and Mapping for Mobile Robots: Introduction and Methods: Introduction and Methods, IGI global, 2012.

[42] J. E. Dennis Jr, R. B. Schnabel, Numerical methods for unconstrained optimization and nonlinear equations, Vol. 16, Siam, 1996.

[43] Y. Pawitan, In all likelihood: statistical modelling and inference using likelihood, Oxford University Press, 2001.

[44] D. Eigen, R. Fergus, Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture, in: Proceedings of the IEEE int. conf. on computer vision, 2015, pp. 2650–2658.

[45] L. J. Lo, D. Banks, A. Novoselac, Combined wind tunnel and cfd analysis for indoor airflow prediction of wind-driven cross ventilation, Building and environment 60 (2013) 12–23.

[46] J.-L. Blanco, J. González-Jiménez, J.-A. Fernández-Madrigal, An alternative to the mahalanobis distance for determining optimal correspondences in data association, IEEE transactions on robotics 28 (4) (2012) 980–986.

[47] H. G. Weller, G. Tabor, H. Jasak, C. Fureby, A tensorial approach to computational continuum mechanics using object-oriented techniques, Computers in physics 12 (6) (1998) 620–631.

[48] J. A. Farrell, J. Murlis, X. Long, W. Li, R. T. Cardé, Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes, Environmental fluid mechanics 2 (1-2) (2002) 143–169.

[49] M. Luperto, J. Monroy, F.-A. Moreno, J. R. Ruiz-Sarmiento, N. Basilico, J. Gonzalez-Jimenez, N. A. Borghese, A multi-actor framework centered around an assististe mobile robot for elderly people living alone, in: IEEE International Conference on Intelligent Robots - Workshop on Robots for Assisted Living (IROS), 2018.

[50] J. G. Monroy, J. González-Jiménez, J. L. Blanco, Overcoming the slow recovery of mox gas sensors through a system modeling approach, Sensors 12 (10) (2012) 13664–13680.

[51] M. Jaimez, J. Monroy, M. Lopez-Antequera, J. Gonzalez-Jimenez, Robust planar odometry based on symmetric range flow and multi-scan alignment, IEEE Trans. on Robotics (2018) 1623–1635doi:10.1109/TRO.2018.2861911.