Preprint

This is the submitted version of a paper published in *Robotics and Autonomous Systems*.

Access to the published version may require subscription.

N.B. When citing this work, cite the original published paper.

Model-based Gas Source Localization Strategy for a Cooperative Multi-Robot System - A Probabilistic Approach and Experimental Validation Incorporating Physical Knowledge and Model Uncertainties

Thomas Wiedemann[a,*], Dmitriy Shutin[a], Achim J. Lilienthal[b]

[a]*German Aerospace Center, Oberpfaffenhofen, Germany*
[b]*Centre for Applied Autonomous Sensor Systems, Örebro University, Örebro, Sweden*

Abstract

Sampling gas distributions by robotic platforms in order to find gas sources is an appealing approach to alleviate threats for a human operator. Different sampling strategies for robotic gas exploration exist. In this paper we investigate the benefit that could be obtained by incorporating physical knowledge about the gas dispersion. By exploring a gas diffusion process using a multi-robot system. The physical behavior of the diffusion process is modeled using a Partial Differential Equation (PDE) which is integrated into the exploration strategy. It is assumed that the diffusion process is driven by only a few spatial sources at unknown locations with unknown intensity. The objective of the exploration strategy is to guide the robots to informative measurements location and by means of concentration measurements estimate the source parameters, in particular, their number, locations and magnitudes. To this end we propose a probabilistic approach towards PDE identification under sparsity constraints using factor graphs and a message passing algorithm. Moreover, message passing schemes permit efficient distributed implementation of the algorithm, which makes it suitable for a multi-robot system. We designed a experimental setup that allows us to evaluate the performance of the exploration strategy in hardware-in-the-loop experiments as well as in experiments with real ethanol gas under laboratory conditions. The results indicate that the proposed exploration approach accelerates the identification of the source parameters and outperforms systematic sampling.

Keywords: robotic exploration, gas source localization, multi-agent-system, partial differential equation, mobile robot olfaction, sparse Bayesian learning, factor graph, message passing

---

*Corresponding author

Email addresses: `thomas.wiedemann@dlr.de` (Thomas Wiedemann),
`dmitriy.shutin@dlr.de` (Dmitriy Shutin), `achim.lilienthal@oru.se` (Achim J. Lilienthal)

1. Introduction

In this paper we consider the task of finding gas sources, e.g. gas leaks, by exploring the resulting gas dispersion with autonomous robots. The proposed approach may be applied in technogenic accidents or during disaster response, where toxic or explosive material is leaking. In such cases localizing the sources is of high interest due to safety concerns. However, for civil protection agencies searching for toxic gas leaks in an already contaminated environment implies threats for human operators. Thus, employing robotic platforms in those scenarios might be beneficial with respect to safety aspects. Moreover, robots with a certain level of autonomy simplify the work of a human operator, as compared to just teleoperated platforms. For example a robot can instantaneously interpret the collected data and make appropriate decisions, which otherwise would require a trained operator or specialist. Following this motivation, we propose to make use of robotic platforms which – equipped with gas sensors – can sample the gas concentration in the environment and process the data in an automated fashion in order to localize gas sources.

1.1. Background and Related Work

A multi-robot system is capable of taking measurements at different locations at the same time. In contrast, a single robot is only able to take measurements serially. To observe the dynamic nature of gas dispersion simultaneous but spatially distributed measurements are necessary [1]. This is not possible with a single robot. Besides, a multi-robot system has additional advantages: i) multiple robots can perform the exploration task faster; ii) they are more robust due to natural redundancies; and, finally, iii) individual robots can make use of synergies, e.g. share the computational costs of algorithms.

However, in order to guide the robots to informative measurement locations, an exploration strategy is needed. The use of robotic platforms for gas source localization, plume tracking or gas distribution mapping in general is an active research field and a lot of different strategies can be found in the literature. For an overview of different approaches and terminology we refer the reader to [2], [3] and [4]. It is possible to categorize exploration strategies as (i) proactive (or predefined) and (ii) reactive (sometimes also called adaptive) strategies. While a predefined strategy follows a plan or trajectory that was defined beforehand, a reactive strategy reacts to current measurements and adjusts its plans accordingly. Both concepts are investigated in [5]. In our work presented here we design a reactive strategy. For gas distribution mapping with the goal to measure everywhere, a proactive strategy that ensures full coverage of the area may be the method of choice. However, for gas source localization reactive strategies are faster, as will be shown in this paper.

Since finding gas sources is also a challenge for animals like moths, beetles or bacteria, exploration strategies are often bio inspired, e.g. [6],[7]. These exploration strategies rely on local concentration gradients to move towards a gas

2

source. Those approaches can be group together under the term chemotaxis. Unfortunately turbulent airflow in the environment results in strong local gradients which do not necessarily point towards the source. Therefore in this paper we have chosen a different approach.

More advanced approaches, sometimes referred to as infotaxis, do not purely rely on concentration gradients, but are more information or entropy driven [8],[9],[10]. Moreover, additional knowledge about the environment or constraints could be utilized for the exploration strategy. For example the search for a gas source could be supported by observations of the current wind field [11] (anemotaxis), or information about obstacles could be used, as done in [12]. Another interesting research question is what you gain from sampling in 3D compared to pure ground base 2D measurements [13]. In general, exploiting model assumptions about the gas dispersion enables the use of higher level strategies. Here high level strategy means path planning and way point navigation to control the robot, in contrast to a low level controller coupling the sensor signals directly to the motors, as in Braitenberg style [3]. In this paper we also propose to use a model-based approach to design an adaptive exploration strategy.

A common way to model the process of a gas dispersion or the plume dynamics is to use PDEs [14],[15]. In this context the exploration or sampling problem is closely related to optimal sensor placement techniques. In the literature, some approaches consider this as an observer design problem. In those cases an observer performance is optimized by adapting the sensor location, e.g. in [16] or [14] to estimate a distributed process described by a PDE. From another perspective, the sampling problem could be treated as an optimal experimental design problem [17]. More details on these topics can be found in [18], where the author gives an introduction to optimal sensor location and experimental design problems. To sum up, from our point of view a PDE seems a very useful way to put physical knowledge about the gas dispersion into an exploration strategy.

Unfortunately, model-based approaches for designing high level sampling strategies for gas exploration are mostly evaluated in simulation and not in real world experiments. So it is unclear if those approaches also work in real world scenarios where the model assumption may not hold. Therefore, we propose to use a probabilistic framework. Our motivation for using a probabilistic formulation of the gas diffusion model is twofold. First, the probabilistic formulation naturally accounts for a possible model mismatch between the actual gas dispersion process and the used model. Second, it allows to quantify the uncertainty of the explored process, which we use to design an exploration strategy. In general, a probabilistic or statistical view is also a contemporary and relevant research topic in the context of gas distribution mapping [19, 20, 10]. Moreover, a probabilistic treatment enables to quantify informativeness of possible measurement locations [21]. Based on that we follow the idea of an uncertainty or entropy driven [22] exploration strategy.

1.2. Contribution

The main focus of the presented approach is the gas source localization task. Our approach does not rely on any assumption about the sources' strengths, the sources' positions or even the exact number of sources. Thus, in contrast to other work, e.g. [15, 8], we can handle cases with multiple sources where their number is unknown. We only assume a sparse distribution of sources. In other words, we do not know the exact number of sources, but we expect that there are only a few of them. For the gas source localization task, the presented adaptive exploration strategy for the multi-robot system is based on a mathematical model of the gas dispersion. This PDE is described in Section 1.2. By means of the model we can infer the location of the gas sources based on gas concentration measurements taken by multiple robots. Further, the model is transferred into a probabilistic formulation as discussed in Section 3. The probabilistic framework provides us with tools (i) to introduce the sparsity assumption about the source distribution and (ii) to quantify the spatial uncertainty about the gas and source distribution. To make use of synergies in the multi-robot-system Section 4 presents a Message Passing (MP) algorithm to perform all calculations in a distributed fashion. The actual exploration procedure grounded on the uncertainty quantification in our probabilistic model is explained in Section 5. This strategy guides the robots to their measurement locations. In contrast to the greedy algorithm in [23] where the robots only consider their direct neighborhood for a new measurement, in this paper robots possess a global view of the whole environment.

The methods and theory for modeling the PDE in a probabilistic framework and the algorithms for a distributed implementation were firstly introduced in [23]. This article is a substantial extension of a previous conference paper [24]. It provides a comprehensive, more detailed description of the approach described in [24]. Further, compared to our previous work the presented approach is evaluated in real world experiments. In general, it is difficult to evaluate gas distribution exploration in realistic experiments due to the difficulty of measuring ground truth gas concentrations. In addition, interesting gas distributions for realistic applications may be of toxic nature and dangerous to handle. To overcome those issues we evaluated the approach in two steps: First, we tested the exploration strategy in hardware-in-the-loop experiments, where we employed a real multi-robot system but only virtually simulated the gas dispersion. This simulation provides us with ground truth data and facilitates reproducibility of the experiments. Compared to pure simulations, in the hardware-in-the-loop setup we can test our system, how it is effected by other real world constraints. Beside constraints given by the robot dynamics, we are particular interested in how the distributed algorithm copes with limitations in a realistic communication system. Parameters like required data rate, communication latency etc. could be studied more easily on a real distributed system. Second, to ensure that the approach is able to cope with model uncertainties we tested the approach in a real world experiment under lab conditions with ethanol as a toy gas.

4

2. Gas Dispersion Process Model

In general, true physical mechanisms behind gas propagation are quite complex. Nonetheless, models exists that can "approximate" gas dynamics sufficiently well. Here we make use of such a model that describes spatial gas dynamics using a PDE of a diffusion equation. Hereby, we neglect other effects such as advection or turbulence. Further, we consider the problem as 2-dimensional, although the true gas dispersion is taking place in 3D. These approximations are necessary to keep the mathematical model feasible for online calculations on our multi-robot system. Besides, in this paper we restrict ourselves to a scenario with ground based robots and gas heavier than air. Such a scenario justifies the 2D approximation.

Consider now an exploration of a gas diffusion process in a bounded spatial region $\Omega \subset \mathbb{R}^2$ over some time interval $T \in \mathbb{R}_+$. A 2-dimensional gas diffusion process can be modeled with a linear parabolic PDE as follows

$$\frac{\partial f(\boldsymbol{x},t)}{\partial t} - \kappa \Delta f(\boldsymbol{x},t) = u(\boldsymbol{x},t),\ \boldsymbol{x} \in \Omega, t \in T, \tag{1}$$

where $f(\boldsymbol{x},t) : \Omega \times T \mapsto \mathbb{R}$ is a space- and time-variant function that represents gas concentration at location $\boldsymbol{x}$ and at time t. Parameter κ in (1) is a gas diffusion coefficient. The right-hand side $u(\boldsymbol{x},t) : \Omega \times T \mapsto \mathbb{R}$ of (1) is interpreted as a gas source distribution. In particular, it describes the source strength (or intensity) of inflow at location $\boldsymbol{x}$ and time t. Let us re-iterate that (1) is used as an approximation to the macroscopic gas dispersion caused by different complex physical mechanism.

Both functions $f(\boldsymbol{x},t)$ and $u(\boldsymbol{x},t)$ are unknown. Our exploration aims at estimating both functions from measurements performed by individual robots. We will assume that measurements are performed according to the following model

$$y(\boldsymbol{x},t) = f(\boldsymbol{x},t) + \epsilon(\boldsymbol{x},t), \tag{2}$$

i.e., we measure a value of a gas concentration perturbed by additive Gaussian noise $\epsilon(\boldsymbol{x},t)$ at location $\boldsymbol{x}$ and time t. In the sequel, we will assume that $\epsilon(\boldsymbol{x},t)$ is spatially and temporally white and normally distributed. Its statistical properties will be specified later in more details.

This model reflects many common gas sensors (e.g., Metal Oxide (MOX) sensors or Photoionization detectors (PIDs)). As we see, the information about $f(\boldsymbol{x},t)$ is acquired directly by means of noisy gas concentration measurements. In contrast, the source distribution $u(\boldsymbol{x},t)$ is hidden and must be inferred from measurements indirectly using some inference procedure.

To enable numerical treatment of the exploration problem, we approximate the space- and time-continuous system (1) with a discrete equivalent. Specifically, we use Finite Difference Method (FDM) [25] for this purpose. Although other, more advanced methods exist for this approximation, the chosen discretization method is simple and illustrates well the proposed methodology. The investigations on the choice of discretization or the use of finite elements

instead of finite differences are left outside the scope of this paper, despite the fact that they do impact the numerical approximation quality (see for example [26]).

To discretize the time, we consider the system at discrete time intervals $t = nT_s$, with $n \in \mathbb{N}_0$ and a sampling period T_s. The discretization in space is done as follows. The exploration domain Ω is divided into C smaller subdomains Ω_c, $c = 1, \ldots, C$, which form a grid with I rows and J columns, so that $C = IJ$. For simplicity we use quadratic cells. This choice, however, affects neither our inference approach nor the exploration strategy. The functions $f(\boldsymbol{x}, t)$ and $u(\boldsymbol{x}, t)$ can then be represented by discrete vectors $\boldsymbol{f}[n] \in \mathbb{R}^C$ and $\boldsymbol{u}[n] \in \mathbb{R}^C$ containing the concentration values and source strengths for each grid cell at time instance n. Note that during a time instance and within a grid cell the concentration values and source strengths are considered as constant.

Based on the discretization, the FDM simply replaces the differential operators in (1) by appropriate finite differences. Thereby, we obtain a system of linear equations as a numerical approximation of the PDE (1). More precisely, we get one equation for each sub-domain Ω_c, i.e., for each grid cell $c = 1, \ldots, C$. These equations define residuals r_c that are set to zero. They represents a relation between the concentration $f_c[n]$ of the considered cell c and its four spatial neighbors: $f_{c-1}[n]$, $f_{c+1}[n]$, $f_{c-J}[n]$, and $f_{c+J}[n]$ as well as the source strength $u_c[n]$ (inflow) in the cell and the concentration $f_c[n-1]$ of the previous time stamp. These define an auxiliary state vector $\boldsymbol{s}_c[n] = [f_c[n], f_{c-1}[n], f_{c+1}[n], f_{c-J}[n], f_{c+J}[n], f_c[n-1], u_c[n]]^T$. The residuals depend on this state vector:

$$r_c(\boldsymbol{s}_c[n]) = 0, \quad c = 1, \ldots, C. \tag{3}$$

In particular for our problem and a cell c this can be rewritten as

$$\left[\frac{4\kappa T_s + d^2}{-T_s d^2}, \frac{\kappa}{d^2}, \frac{\kappa}{d^2}, \frac{\kappa}{d^2}, \frac{\kappa}{d^2}, \frac{1}{T_s}, 1 \right] \boldsymbol{s}_c[n] = 0 \tag{4}$$

with d denoting the cell width. We would also like to stress that equations in (4) are specific for the used PDE and the FDM approximation. For other PDEs, these equations need to be appropriately modified following similar discretization steps.

Let us remark that after the discretization, the source distribution $u(\boldsymbol{x}, t)$ is represented by a vector $\boldsymbol{u}[n]$, with each element corresponding to the source strength (or intensity) in a particular cell. Thus our assumption that sources are sparsely distributed in the domain Ω implies that the vector $\boldsymbol{u}[n]$ is sparse in a sense that most of its entries are zero. The number and the index of nonzero entries in $\boldsymbol{u}[n]$ reflect the number and locations of the active gas sources, which we are interested to infer from measurements.

Now consider K robots which explore Ω by means of sampling the gas concentration process $f(\boldsymbol{x}, t)$ according to (2). By collecting the measurements $y_k[n]$

6

of individual robots in a vector $\boldsymbol{y}[n] = [y_1[n], \ldots, y_K[n]]^T$, we can represent the measurement performed by the robots as

$$\boldsymbol{y}[n] = \boldsymbol{M}[n]\boldsymbol{f}[n] + \boldsymbol{\epsilon}[n]. \tag{5}$$

Here $\boldsymbol{M}[n] = [\boldsymbol{m}_1[n], \ldots, \boldsymbol{m}_K[n]]^T$ is a spatial $K \times C$ sampling matrix. Each vector $\boldsymbol{m}_k[n] \in \mathbb{R}^C$, $k = 1, \ldots, K$, is defined as a zero vector except one element that equals 1 at an index that corresponds to the spatial cell c measured by the robot k at time n. Finally, we state the properties of the additive noise process $\boldsymbol{\epsilon}[n]$ in (5). Specifically, we assume $\boldsymbol{\epsilon}[n]$ to be spatially and temporally white, normally distributed random vector with zero mean and precision τ_m, i.e., $\boldsymbol{\epsilon}[n] \sim N(\boldsymbol{0}, \tau_m^{-1}\boldsymbol{I})$.

We would like to avoid a synchronization of the robots that ensures that at each time stamp there is a measurement of each robot. Instead we prefer a system where we can plug in measurements whenever it is available. This means there may be a time stamp n without measurements of robot k. In order to account for this "no measurement available case", we just set the precision of the corresponding element $\boldsymbol{\epsilon}_k[n]$ to zero. This causes the line k of the equation system (5) to be ignored in the probabilistic formulation, as will become obvious from the next section section.

3. Probabilistic Inference Framework

Our motivation for using a probabilistic formulation of the gas diffusion model is twofold. First, the probabilistic formulation naturally permits us to account for a possible model mismatch between the actual gas dispersion process and the used PDE model (1). Second, it will allow us to quantify the uncertainty of the explored process in different regions of the domain Ω. This property forms the basis for our exploration strategy, as we will show later. In particular, we use information-theoretic tools to quantify uncertainties (see Section 3.2 for more details).

In the probabilistic setting the gas concentration vectors $\boldsymbol{f}[n]$ and source strength vectors $\boldsymbol{u}[n]$ are modeled as random vectors. The underlying probabilistic structure will be explained in the following.

3.1. Bayesian Formulation

As a first step, we relax equation (3). In particular, we assume that C equations in (3) are allowed to deviate from zero. This deviation – the residual of a grid cell– is a normally distributed random variable with zero mean and precision τ_s. These residuals are assumed to be conditionally independent and identically distributed for all grid cells in Ω. The latter assumption is reasonable since equations in (3) essentially "act" locally on neighboring cells. The role of τ_s is to regulate our trust in the model: with $\tau_s \to \infty$ we recover the deterministic case, which encodes the assumption that the dispersion process is accurately represented by (1). For small τ_s we allow the dispersion process to deviate from the model (1) and thus "tolerate" other dynamical effects that are not captured

7

with a pure diffusion. This relaxation is an important feature of the presented approach, since we are able to parametrize our trust in the model by a single scalar value with a physical interpretation.

Under the above assumptions and using (3), we can now define the conditional Probability Density Function (PDF) of the gas concentration distribution $\boldsymbol{f}[n]$ at time n as

$$p(\boldsymbol{f}[n]|\boldsymbol{f}[n-1], \boldsymbol{u}[n]) \propto \prod_{c=1}^{C} \mathrm{e}^{-\frac{\tau_s}{2}\left(r_c(f_c[n], f_{c-1}[n], f_{c+1}[n], f_{c-J}[n], f_{c+J}[n], f_c[n-1], u_c[n])\right)^2}.$$

(6)

In a similar way the measurement model is casted in a probabilistic setting. Based on (5) and the assumed Gaussian noise characteristics we can formulate the gas concentration likelihood function as follows:

$$p(\boldsymbol{y}[n]|\boldsymbol{f}[n]) \propto \mathrm{e}^{-\frac{\tau_m}{2}\|\boldsymbol{M}[n]\boldsymbol{f}[n]-\boldsymbol{y}[n]\|^2} \propto \prod_{k=1}^{K} \mathrm{e}^{-\frac{\tau_m}{2}\left(\boldsymbol{m}_k[n]^T \boldsymbol{f}[n]-y_k[n]\right)^2}. \quad (7)$$

To complete the probabilistic formulation of our model, we also need to specify two prior PDFs: an initial gas concentration $\boldsymbol{f}[0]$ and a source prior PDF $p(\boldsymbol{u}[n])$.

The source prior $p(\boldsymbol{u}[n])$ is a mechanism that we use to incorporate our assumptions about the source sparsity. To do so we appeal to Sparse Bayesian Learning (SBL) techniques [27]. This approach we first proposed in [23]. In the following we only give a short summary of the approach and refer the reader to [23] for more details. We introduce a new hyper-parameters $\gamma_c[n]$ for each cell. The hyper-parameters $\gamma_c[n]$ represent the parameters of the prior for $u_c[n]$; they are treated as random variables and are estimated along with the other model parameters. Sparsity is introduced through construction of a hierarchical prior $p(\boldsymbol{u}[n], \boldsymbol{\gamma}[n])$, which is parameterized with hyper-parameters $\boldsymbol{\gamma}[n]$ as follows:

$$\begin{aligned} p(u_c[n]|\gamma_c[n]) &= N(u_c[n]|0, \gamma_c^{-1}[n]) \\ p(\gamma_c[n]) &= Ga(\gamma_c[n]|a_\gamma, b_\gamma), c = 1, \ldots, C. \end{aligned} \quad (8)$$

where $Ga(\cdot|a, b)$ is a gamma PDF with parameters a and b. The product $p(\boldsymbol{u}[n]|\boldsymbol{\gamma}[n])p(\boldsymbol{\gamma}[n])$ defines a so called Gaussian scale mixture [28].[1]

Here we will make use of a popular version of SBL that uses non-informative hyper-prior $p(\gamma_c[n]) \propto \gamma_c[n]^{-1}$ obtained when $a_\gamma \to 0$ and $b_\gamma \to 0$ [29, 30, 27]. This prior can be recognized as a non-informative Jeffrey's prior. The motivation for this choice is twofold. First, the resulting inference schemes typically demonstrate superior (or similar) performance as compared to schemes derived based on other hyper-prior selections [28]. Second, very efficient inference algorithms can be constructed and studied [31, 32, 33, 34, 35].

---

[1] A recent work [28] extends the framework by generalizing $p(\boldsymbol{u}[n]|\boldsymbol{\gamma}[n])$ to be the PDF of a power exponential distribution, which makes the hierarchical prior a power exponential scale mixture distribution.

The final ingredient is the prior gas concentration $p(\boldsymbol{f}[0])$. Since we have no information about the initial gas concentration distribution, we choose $p(\boldsymbol{f}[0])$ with zero mean and a very high variance: $p(\boldsymbol{f}[0]) = \prod_{c=1}^{C} N(f_c[0]|0, 10^3)$.

Now, using Bayes theorem we can construct the desired posterior PDF for the processes of interest by combining (6), (7), (8) and the prior $p(\boldsymbol{f}[0])$ as follows:

$$
\begin{aligned}
p(\boldsymbol{f}[0]...\boldsymbol{f}[N], \boldsymbol{u}[1]...\boldsymbol{u}[N], \boldsymbol{\gamma}[1]...\boldsymbol{\gamma}[N]|\boldsymbol{y}[1]...\boldsymbol{y}[N]) = \\
p(\boldsymbol{f}[0]) \prod_{n=1}^{N} p(\boldsymbol{y}[n]|\boldsymbol{f}[n]) p(\boldsymbol{f}[n]|\boldsymbol{f}[n-1], \boldsymbol{u}[n]) \\
\prod_{c=1}^{C} p(u_c[n]\big|\gamma_c[n]) \prod_{c=1}^{C} p(\gamma_c[n]).
\end{aligned}
\tag{9}
$$

Let us point out that the prior $p(\boldsymbol{u}[n], \boldsymbol{\gamma}[n])$ not only reflects the fact that there are only a few "active" cells with sources, but also plays the role of a regularization term in a classical deterministic setting. To be more precise, inspecting the log of a posterior (9) would reveal that the prior (8) introduces a penalty term $\sum_c \gamma_c[n] \left(u_c[n]\right)^2$, which is a weighted ℓ_2-norm of the vector $\boldsymbol{u}[n]$. From a perspective of exploration the regularization is an important part of the observation procedure. In fact, solving equations arising from the FDM for $\boldsymbol{u}[n]$ without the regularization would not be possible at early phases of the exploration, since there would be too few measurements available.

3.2. Uncertainty Quantification

The key idea of our exploration strategy is to direct robots to regions, where our knowledge about the explored process is currently low. This requires a mechanism for quantifying the information we have about the process itself or about process parameters. If the uncertainty about the gas concentration value or source strength in a grid cell is high, then this grid cell is a good candidate for the next measurement location. In the following we present how the uncertainty could be quantified based on the developed probabilistic model.

Indeed, the PDFs of variables associated with a grid cell – gas concentration or source intensity – can be used to compute the amount of information we have about these variables. In order to evaluate only a single grid cell, we calculate the marginal PDF for each cell based on the joint posterior (9) by integrating over all other variables and parameters. While in simple cases like a Gaussian distribution, the marginal PDF could be calculated in closed form, the introduction of the Gamma distribution in the hierarchical prior prevents an analytical calculation in our case. Yet as we will show, the marginalization can be performed efficiently by representing (9) using a factor graph [36] and performing inference on this graph using message passing algorithms (see also [23]). This algorithm and its implementation is described in more detail in the next section. As a result we obtain the marginal distribution of the source strength $u_c[n]$ and gas concentration $f_u[n]$.

9

It has been shown in [23] that the marginal distributions of both gas concen-
tration and source intensity in a grid cell can be approximated with Gaussian
PDF. As such, the second order moments of these variables can be used as a
gauge of the information content[2]. Specifically, the variance of a cell is used
to quantify the corresponding cell uncertainty; cells with a higher variance are
more uncertain and therefore are more "interesting" as potential location for a
new measurement. In this context "interesting" means that a measurement at
this location will help to reduce the error in the parameter estimation more, as
compared to other locations.

4. Distributed Inference Algorithm

As described in the previous section, marginal PDFs of the variables can be
used to quantify of the uncertainty of the latter. This is the key to the proposed
exploration strategy. As closed form solutions are rarely available, a numerical
estimation of the quantities of interest is needed. In this chapter we discuss
how to calculate these PDFs in a distributed fashion based on the graphical
formulation of the posterior (9). This graphical representation is used to derive
a MP algorithm that calculates the marginal PDFs of the required variables in a
distributed fashion. Further we show how this algorithm is implemented on our
multi-robot system in a distributed fashion. The derivation of the algorithm is
detailed in [23], to which we refer the interested reader for further details. In
the following we will merely summarize the key inference steps and expressions.

4.1. Factor Graph Representation

Our inference algorithms is based on representation of the posterior PDF (9)
using Factor Graphs (FGs) [36, 37]. A FG is an undirected bipartite Bayesian
network being composed of value nodes, which represent random variables, and
factor nodes, which model functional dependencies between variables according
to the factorized representation of a joint PDF. In our case, the factorized rep-
resentation of the latter is obtained by inserting equation (6) into (9). Figure 1
depicts the FG for our posterior PDF (9) for a single cell c and a time instance n.
The variable nodes (depicted as spheres) correspond to the random variables of
the concentration $f_c[n]$, source strength $u_c[n]$ and hyper-parameter $\gamma_c[n]$. These
variables are related to each other via factor nodes, depicted by cubes in the
figure and represented by capital letters. The factor nodes model the depen-
dencies and constraints according to the factorized posterior formulation. Four
types of factors can be identified in the graph. A factor node Y_c represents a
measurement taken by an agent at the cell c. This factor formally represents the
likelihood function (7). Clearly, this factor is only present whenever a measure-
ment at the cell c has been performed; it is absent when no measurement has
been taken at this location. The next factor node R_c represents the PDE model

---

[2]Recall that for a Gaussian random variable the entropy is related to the square root of
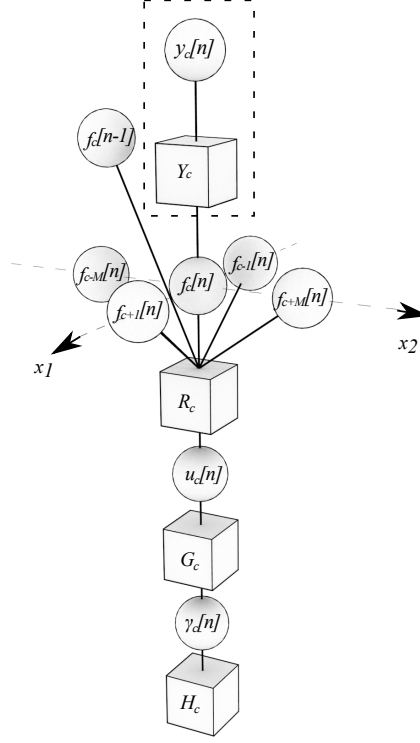the variable's variance.

Figure 1: Factor Graph: This graph represents the part of the posterior PDF associated with a single grid cell. It models the relations between variable nodes (spheres) by factor nodes (cubes).

in (1) and the used discretization the model. According to the used discretization, the R_c relates the concentration value $f_c[n]$ to that of the neighboring grid cells together with the source strength $u_c[n]$ in this cell according to (6). The other two factor nodes G_c and H_c represent the parametric prior $p(u_c[n]|\gamma_c[n])$ and the hyper-prior $p(\gamma_c[n])$ in (9), respectively.

4.2. Message Passing Algorithm

Inference on FGs can be efficiently implemented using MP algorithms [36, 37]. MP is a powerful tool to calculate marginal distributions of random variables given a FG. For the considered problem two algorithms are used: the sum product algorithm [38] (also called loopy belief propagation) and variational MP [39]. The latter is particularly handy when approximate inference is desired, as will be discussed later.

In both cases, the MP algorithms work as follows. Certain messages are exchanged between nodes of the FG along the edges: from factor nodes to variable

11

nodes and from variables to factors. The messages represent probabilistic densities functions, or beliefs. By iteratively exchanging messages between nodes, the outgoing messages of variable nodes converge to the marginal distribution of the corresponding variables. In the following we will use the following notation to denote messages. For a message outgoing from a node A towards a node B we will use notation $m_{A \to B}$.

The challenge in application of MP algorithms is the computation of the messages. In our case for the graph Fig. 1 it can be verified that all messages connected to factors R_c and Y_c could be calculated according to the sum product algorithm. Moreover, the messages are Gaussian PDFs. As such they can be presented as $m_{A \to B} = N(x|\mu_{A \to B}, \tau_{A \to B}^{-1})$ where $\mu_{A \to B}$ is the message mean and $\tau_{A \to B}$ is the message precision. For our model these can be computed analytically in closed form. Only these two values have to be communicated along the edges of the graph to fully communicate the belief. The rules for calculating the messages are summarized in Table 1 (see also [23] for a detailed derivation of the message expressions). Let us also remark, that the message $\mu_{Y_c \to f_c[n]}$ corresponds to the actual measured value and $\mathcal{N}(c)$ denotes the set of all neighbored cells of c excluding c itself.

to simplify further notation let us also introduce two auxiliary vectors

$$\boldsymbol{s}_c[n] = [f_c[n], f_{c-1}[n], f_{c+1}[n], f_{c-J}[n], f_{c+J}[n], f_c[n-1]]^T$$

and

$$\boldsymbol{z}_c[n] = [f_{c-1}[n], f_{c+1}[n], f_{c-J}[n], f_{c+J}[n], f_c[n-1], u_c[n]]^T,$$

which aggregate all concentrations $f_{c' \in \mathcal{N}(c)}$ together with f_c and the source strength u_c, respectively.

In contrast to the message listed in the Table 1, the messages related to the sparsity-inducing source priors for the factor nodes G_c and H_c are not analytical tractable by the sum-product algorithm. Therefore, we use the Variational Message Passing (VMP) techniques to compute analytical approximation of the messages. Moreover, as we have shown in [23], it is possible to compute fixed point expressions for the VMP inference expressions and thus accelerate convergence. Combining this with the other messages computed using the sum-product algorithm leads to simple update rules for the message $m_{u_c[n] \to R_c}$ as summarized in Table 1. We would like to stress that the FG in Fig. 1 represents only a single cell. The overall graph requires computing the messages for all cells and factors. Nonetheless, those messages could be calculated in closed form and in random order or all at once in parallel. This flexibility makes the algorithm particularly suitable for a distributed implementation on a multi-robot system. To this end, we split the overall graph into different parts that correspond to different 2D regions of our environment. A simplified version of the overall FG and the partitioning of this graph are exemplarily shown in Fig. 2. For the sake of clarity, the connections to nodes of the previous time stamp are not displayed. Further, the nodes Y_c are only exemplarily shown for some cells, where it is assumed that a measurement was taken. Each region (i.e. subgraph) is assigned to one robot of our multi-robot system and each robot is able

12

Message	Equation
$m_{Y_c[n] \to f_c[n]}$	$N(f | \mu_{Y_c \to f_c[n]}, \tau_m^{-1})$
$m_{f_c[n] \to R_c}$	$m_{Y_c \to f_c[n]} \displaystyle\prod_{c' \in \mathcal{N}(c)} m_{R_{c'} \to f_c[n]} \propto N(f | \mu_{f_c[n] \to R_c}, \tau_{f_c[n] \to R_c}^{-1})$ $\tau_{f_c[n] \to R_c} = \tau_m + \displaystyle\sum_{c' \in \mathcal{N}(c)} \tau_{R_{c'} \to f_c[n]}; \quad \mu_{f_c[n] \to R_c} = \frac{1}{\tau_{f_c[n] \to R_c}} \left[\displaystyle\sum_{c' \in \mathcal{N}(c)} \mu_{R_{c'} \to f_c[n]} \tau_{R_{c'} \to f_c[n]} + \mu_{Y_c \to f_c[n]} \tau_m \right]$
$m_{R_c \to u_c[n]}$	$\displaystyle\int e^{-\frac{\tau_s}{2}(r_c(s_c[n], u_c[n]))^2} \times m_{f_c[n-1] \to R_c} \times m_{f_c[n] \to R_c} \times \prod_{c' \in \mathcal{N}(c)} m_{f_{c'}[n] \to R'_c} \, ds_c[n] \propto N(u | \mu_{R_c \to u_c[n]}, \tau_{R_c \to u_c[n]}^{-1})$ $\tau_{R_c \to u_c[n]} = \left(\tau_s^{-1} + a^T \mathrm{diag}\{\tau_{f_c[n] \to R_c}, \tau_{f_{c-1}[n] \to R_c}, \tau_{f_{c+1}[n] \to R_c}, \tau_{f_{c-J}[n] \to R_c}, \tau_{f_{c+J}[n] \to R_c}, \tau_{f_c[n-1] \to R_c}\}^{-1} a \right)^{-1}$ $\mu_{R_c \to u_c[n]} = [\mu_{f_c[n] \to R_c}, \mu_{f_{c-1}[n] \to R_c}, \mu_{f_{c+1}[n] \to R_c}, \mu_{f_{c-J}[n] \to R_c}, \mu_{f_{c+J}[n] \to R_c}, \mu_{f_c[n-1] \to R_c}] a$ $a = - \left[\dfrac{4\kappa T_s + d^2}{-T_s d^2}, \dfrac{\kappa}{d^2}, \dfrac{\kappa}{d^2}, \dfrac{\kappa}{d^2}, \dfrac{\kappa}{d^2}, \dfrac{1}{T_s} \right]^T$
$m_{R_c \to f_c[n]}$	$\displaystyle\int e^{-\frac{\tau_s}{2}(r_c(s_c[n], u_c[n]))^2} \times m_{f_c[n-1] \to R_c} \times m_{u_c[n] \to R_c} \times \prod_{c' \in \mathcal{N}(c)} m_{f_{c'}[n] \to R_c} \, dz_c[n]$ $\tau_{R_c \to f_c[n]} = a_0^2 \left(\tau_s^{-1} + b^T \mathrm{diag}\{\tau_{u_c[n] \to R_c}, \tau_{f_{c-1}[n] \to R_c}, \tau_{f_{c+1}[n] \to R_c}, \tau_{f_{c-J}[n] \to R_c}, \tau_{f_{c+J}[n] \to R_c}, \tau_{f_c[n-1] \to R_c}\}^{-1} b \right)^{-1}$ $\mu_{R_c \to f_c[n]} = [\mu_{u_c[n] \to R_c}, \mu_{f_{c-1}[n] \to R_c}, \mu_{f_{c+1}[n] \to R_c}, \mu_{f_{c-J}[n] \to R_c}, \mu_{f_{c+J}[n] \to R_c}, \mu_{f_c[n-1] \to R_c}] b$ $b = \left[-1, \dfrac{\kappa}{d^2}, \dfrac{\kappa}{d^2}, \dfrac{\kappa}{d^2}, \dfrac{\kappa}{d^2}, \dfrac{1}{T_s} \right]$
$m_{u_c[n] \to R_c}$	$\begin{cases} N\!\left(u_c \,\big|\, \dfrac{\mu_c^2 \tau_c - 1}{\mu_c \tau_c}, \left(\tau_c + \dfrac{\tau_c}{\tau_c \mu_c^2 - 1}\right)^{-1}\right); & \tau_c \mu_c^2 - 1 > 0 \\[2mm] N(u_c | 0, \infty^{-1}) = \delta(u_c); & \tau_c \mu_c^2 - 1 \le 0, \end{cases}$ with $\mu_c = \mu_{R_c \to u_c[n]}$ and $\tau_c = \tau_{R_c \to u_c[n]}$

Table 1: Update rules for the message passing algorithm (see [23] for detailed derivation)

to calculate all messages related to its own part. The edges crossing the border of a region are marked as red lines in Fig. 2. These messages communicated over these edges have to be exchanged between the robots via an communication link. At the current status of our work we divide the region in rectangular sub-regions with equal number of nodes and assign those randomly to the robots. Of course one can think of more intelligent ways to assign the sub-regions to robots in order to minimize communication load or allow communications only between neighboring robots.
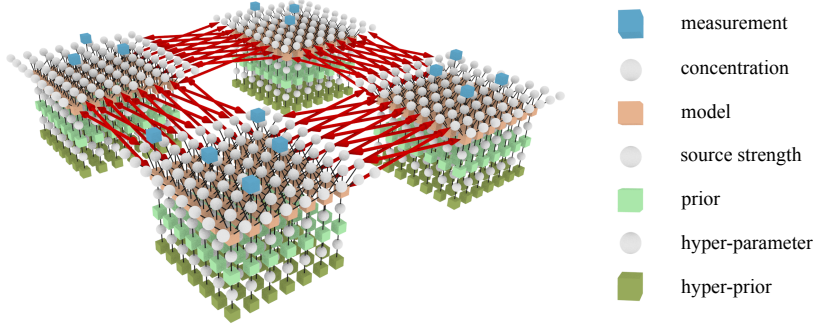

Figure 2: Distributed Factor Graph: This figure illustrates the overall factor graph. This is a simplified version without time dependencies. The graph is spatially split into four regions. Big arrows represent messages that are communicated between different agents.

5. Exploration Procedure

After convergence of the messages, we obtain marginal PDFs of all the variables in the graph. Our interest is in the posterior variances of source variables $u_c[n]$, $c \in \{1, \ldots, C\}$, as it can be used to quantify the uncertainty about the sources at the corresponding cells (see Section 3.2). The latter is used by each robot to propose potential cells of interest for making new concentration measurements, as explained in the following.

First, the cells are rated according to the inverse variance, i.e., precision $\tau_c[n]$, of the source strength marginal $p(u_c[n]) \propto N(\hat{u}_c[n], \tau_c[n])$. Then, a pre-defined number P of cells with the lowest precision is selected as a proposal for a new measurement locations. These P cells are computed as follows: each robot selects K cells with the lowest precision using part of the graph it is responsible for. These cells are then combined with proposals of the other robots, such that a total of P cells is obtained. Let us note that although each robot generates possible way points using only its own part of the factor graph, it uses a global set of possible measurement positions to make measurements. Once an assignment is made, the selected measurement location is removed from the list of potential way points to avoid inter-agent conflicts. Clearly, this requires a corresponding coordination protocol to communicate decision to other members in the swarm.
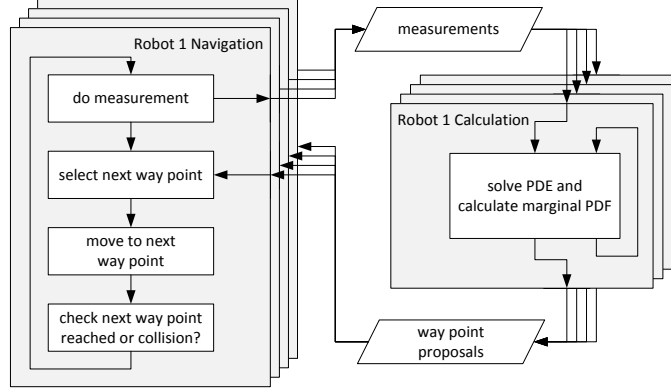
14

Figure 3: Exploration Procedure: The exploration is implemented with two main loops. The right loop solves the PDE and produces new way point proposals, the left one controls the individual robots.

After the assignment of the measurement location to a robot in the swarm is made, the robot moves to the selected way point. Despite robots are assigned different way points, a collision avoidance mechanism is needed to avoid inter-robot collisions.

In our work we used a reactive collision avoidance strategy. On its way, the robot monitors the current distance to all other robots based on received position information. If the distance reduces below a pre-defined safety threshold, i.e., when two robots are getting too close, the robot stops and selects another way point from the set of available proposals such as to increase the distance itself and the other robots on a collision course. Finally, when a robot reaches its goal, it makes a measurement which is then incorporated in the algorithm, after which the whole exploration cycle is repeated.

The overall procedure is depicted in Fig. 3. As can be seen, the robotic navigation and solution of the PDE actually run in two separate loops. These loops run independent of each other and asynchronously. The loop responsible for solving the PDE and updating the way point proposals runs with a constant frequency of $1Hz$ $(T_s = 1s)$. The navigation loop of the robots is not synchronized with the calculation loop. It just considers the last available generated list of way point proposals, but does not wait for an up-to-date list. Further the navigation loop does not run with a constant frequency, since the time to reach a new measurement location always differs on the distance to the new location. We would like to stress that the navigation loop of each robot is also not synchronized with other robots. Thus no robot has to wait for results of the other robots. Whenever a measurement becomes available at a time stamp n, it is directly inserted into the PDE solver. Hence, the two loops are only connected by the exchange of proposed way points and measurements.

Let us also point out that the presented approach models the source strengths

as a time dependent variable, i.e. $\boldsymbol{u}[n]$ is a function of time n. Thus the algorithm can handle appearing and disappearing sources (see also [40]). Since, the algorithm is not only supposed to find the source but also to monitor the sources afterwards, it is not possible nor desired to define a clear termination criterion. However, in the following experimental evaluation, we kept the source strengths constant over time. The experiments were stopped by the operator as soon as no apparent change in the estimates was observable anymore.
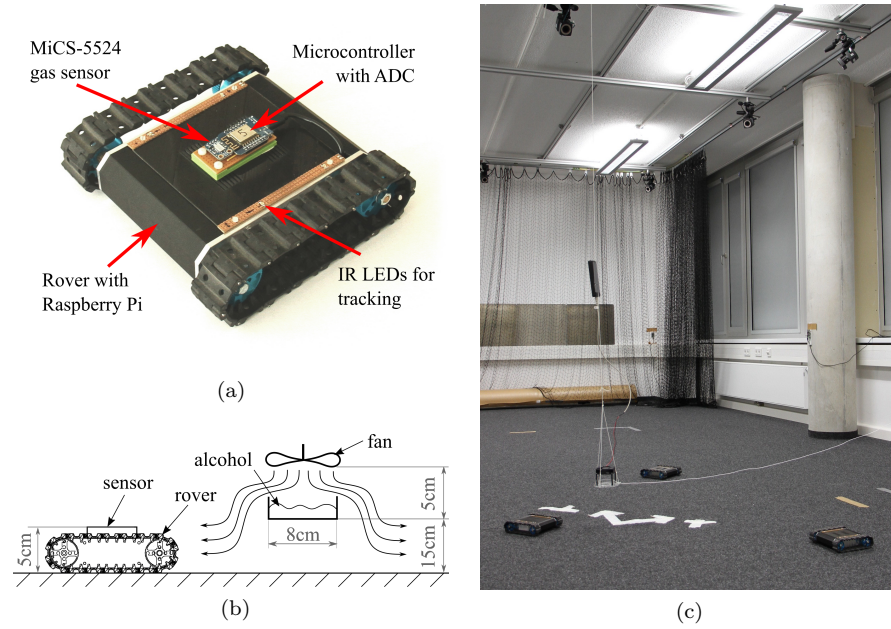
6. Experimental Setup

Figure 4: Experimental setup: Picture (a) depicts the robotic platform carrying the gas sensor. The schematic in (b) shows the design of the artificial gas source and (c) a picture of the actual realization, however, on this picture the rovers are not equipped with gas sensors.

Evaluation of gas mapping or gas source localization strategies in realistic real-world scenarios is in general quite hard [41]. While ground truth data may be available for source locations, this is not the case for the gas concentration distribution. Even though in this work we mainly focus on the gas source localization, for better understanding of the used approach, an accurate estimation of the gas concentration is also of interest. To address these issues we therefore analyze our exploration strategy in two different types of experiments.

First, we carry out hardware-in-the-loop experiments with synthetic gas dispersion. This implies that a real robotic system is used, yet the gas dispersion process is simulated. In this way we obtain ground truth data of the source

16

distribution and gas concentration from the simulation. However, the simulation uses the same mathematical model as the model-based exploration strategy. This perfect model match is not usually observed in reality. For this reason, we also carried out experiments with real gas, more specifically with ethanol vapor and a robotic swarm in lab conditions.

For our experiments, small, custom-built ground robots are used. Each robot is equipped with a Raspberry Pi 2 – a low power single-board computer with Linux OS (900MHz quad-core ARM Cortex-A7 CPU, 1GB RAM). From this computer it is possible to send control commands to a micro-controller that implements a velocity controller for two motors driving the tracks of the robot (see Fig. 4a). The experiments are done in a laboratory with a commercial optical tracking system. Using active infrared LEDs mounted on the robot the tracking system is able to compute accurate robot location and orientation in realtime with an accuracy of $\approx$ 1cm. This can be considered as an almost perfect localization with respect to the discretization of the environment with a cell size larger than 10cm. We used from 3 to 5 robots in the experiments. Robot are able to exchange data through a WiFi communication link. Also, through this link the robots receive their current positions. The interconnection of all involved components is implemented using the Robot Operating System (ROS) [3]. We now outline the setting for the two experiments.

In the hardware-in-the-loop experiments, the gas diffusion is simulated for a two dimensional case. The data are generated according to equation (1). Whenever a measurement is demanded by the exploration procedure for a robot, the equation is evaluated at the current position of the robot. Additionally, we perturb the measurement with an additive white noise ξ. For the evaluation of the PDE a Finite Volume Method solver is used [42]. The used spatial discretization grid has 12×30 cells. For our example scenario we have chosen Dirichlet boundary conditions $f(x,t) = 0$, except for the right border, where we use a Neumann boundary condition $\frac{\partial f(x,t)}{\partial x} = 0$. This would correspond to an open field scenario, where material can flow off at all boarders except for the right one because of e.g. a wall or a similar obstacle. Note that in this setting the system will not reach a steady state. Thus, we can be sure to observe a dynamic process. For the virtual gas simulation we considered the concentration and source strengths as unit-less. The discrete grid size, the time difference between two discrete time stamps and the diffusion coefficient κ are set to 1 in the simulation. However, later the concentration field is fitted to match our laboratory environment with a scale of $6x2.4$m.

For the second experiment with the ethanol gas, the robots are equipped with a MiCS 5524 metal oxide gas sensor (SGX Sensortech Ltd, Switzerland). The sensor voltage is measured with an AD-converter of an ESP8266 micro-controller, which provides the sensor reading to the whole system via wireless LAN. In order to convert the binary values of the ADC to concentrations, we used the data sheet of the sensor [43], which specifies a linear dependency in the

---

[3]http://www.ros.org/

17

log-log domain between the sensor's resistor R_s and the ethanol concentration. Therefore, we model this dependency by an exponential function [44] as follows:

$$y = \beta \cdot R_s^{\alpha}, \tag{10}$$

where y is the measured gas concentration plugged into our Bayesian approach (i.e. $y = \mu_{Y_c \to f_c}[n]$). The parameter α in (10) corresponds to the slope of the sensor sensitivity in the log-log domain. According to the data sheet it is set to $\alpha = -1.6$. The parameter β was chosen in such a way that y is normalized to be unit-less and roughly in the range of $[0, 1])$. In this way, we can use the exact same parametrization of the algorithm as in the hardware-in-the-loop case, where we considered unit-less concentration and source values in the same range. Further, the parameter β is individually adjusted for each sensor (or each robot in our case) so that the sensor responses are roughly the same across the swarm when exposed to the same constant concentration. In our case β was chosen to be approximately 10. The resistor R_s was calculated according to the measured voltage drop U on a load resistance R_L of a voltage divider consisting of R_s and R_L as follows:

$$R_s = R_L \left(5V/U - 1\right) \propto \left(5V/U - 1\right). \tag{11}$$

We would like to remark that we incorporate R_L in the constant β when inserting (11) into (10).

As a toy gas source in the experiments we used ethanol evaporating from a culture dish ($8cm$ diameter) filled with approximately 5g of 94% ethanol assay. Above the culture dish we mounted a small fan (see Fig. 4c). The airflow caused by this fan avoids a saturation of ethanol concentration in the layer above the liquid. Thus, it accelerates the evaporation. Moreover, the air flow facilitates a radial dispersion of the ethanol gas. The whole structure hangs down from the ceiling, so that the robots are able to drive below the source without any collision with the dish containing the alcohol solution. Finally, we would like to note that the culture dish is relatively small compared to the cell size of $20cm$ and its height of $15cm$ over ground. In this way we make sure that no cell is shadowed by the dish, and the highest concentration is still below the source.

7. Evaluation

In this section we summarize the results of our experimental studies. We begin with the analysis of the results for hardware-in-the loop experiment with the proposed exploration strategy. Then, the second experiment with real ethanol gas dispersion is discussed.

7.1. Hardware-in-the-loop exploration experiment

As we mentioned, in this experiment we intend to demonstrate the performance of the numerical Bayesian solver and benchmark the proposed exploration strategy. Specifically, we compared the proposed exploration strategy against
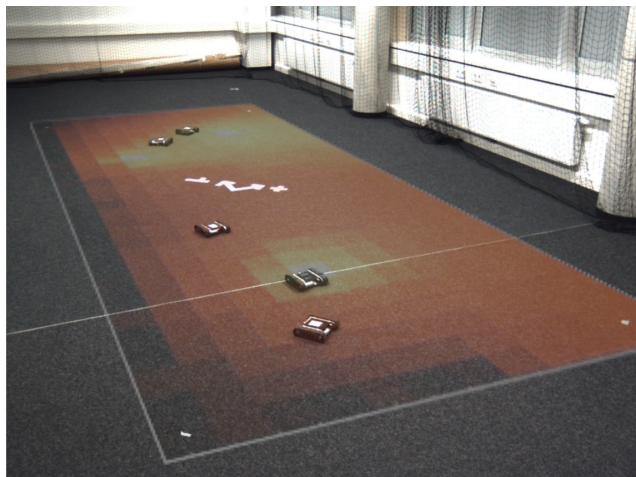
18

Figure 5: Lab Environment: The picture shows our lab during an experiment. The simulated concentration field is projected to the ground in a post-processing step.

exploration with a predefined sweeping trajectory for the case of 5 robots. The sweeping trajectories are generated by simply dividing the environment into five equal regions and generating a predefined "meander" trajectory for each of these regions. In literature this kind of trajectories are also referred to as "lawn mower" path. In these trajectories the measurements will fully cover the whole environment after a certain time, i.e. each grid cell is measured at least once. This strategy is reasonable if no prior knowledge or model assumptions are available. We compare the performance of the strategies by means of efficiency (in terms of the required number of measurement samples needed to achieve convergence of the source signals) and quality of the estimates in terms of the achievable estimation error. Note that using the number of measurement instead of using the actual time for the exploration is of an advantage in our case, since the actual time highly depends on the time spent on making a measurement; these are different for hardware-in-the-loop and real gas measurements. Also, in order to see how well the spatially distributed sparse sources $\boldsymbol{u}[n]$ are identified, we use a so-called Earth Mover's Distance (EMD) measure, which is analogous to a Wasserstein metric for discrete distributions [45]. The use of such distance metric is motivated by the fact that classical Mean Squared Error (MSE) does not adequately represent distances between sparse signals. Instead, sparse signals are treated as distributions, and corresponding metrics – Wasserstein metric in a continuous case and Earth Mover's Distance (EMD) metric in the discrete case – better reflect distances between estimated sparse signals and the "ground truth" signal. In particular, EMD measures the effort needed to "displace" one distribution onto another one. In our case we use EMD to compare the estimated vector $\boldsymbol{u}[n]$ with the ground truth vector $\hat{\boldsymbol{u}}$ with all elements set to zero except for the three cells containing a source.
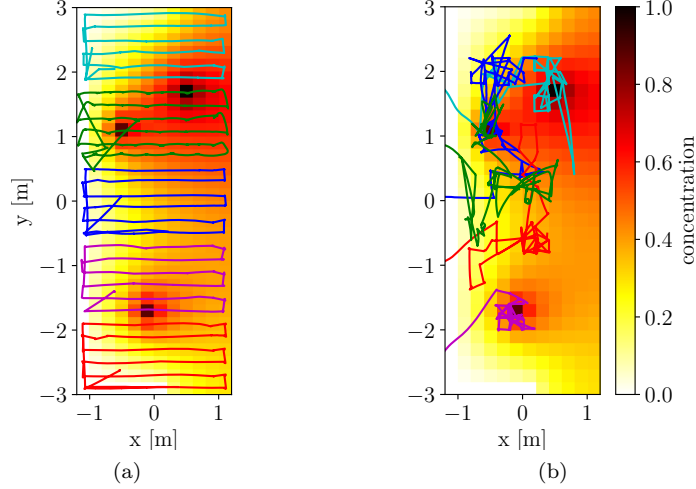
19

Figure 6: Hardware-in-the-loop experiments: The figure compares the meander trajectory (a) and the proposed exploration strategy (b). The trajectories superimpose the simulated concentration field. Note that the shown concentration field is only a snapshot at the corresponding time stamp of the trajectories. This field is time variant during the exploration process.

We placed three sources in the environment with the following source amplitudes and locations: source 1 with $\hat{u} = 1.0$ at $x = -0.2m, y = -1.8m$, source 2 with $\hat{u} = 1.0$ at $x = -0.6m, y = 1.0m$, and source 3 with $\hat{u}_3 = 0.8$ at $x = 0.4m, y = 1.6m$. The sources are placed so as to disclose different aspects of the source localization problem. One source is isolated in the lower region, where the other two are more close to each other (see Fig. 6); the latter incurs some strong spatial correlation between them. Further one of those is placed in an area with a generally higher concentration level. For a detailed study on the effect of the number of sources, we refer the reader to [40]. Besides, we use the Normalized Mean Square Error (NMSE) to quantify the gas concentration estimation error, defined as $||\boldsymbol{f}[n] - \hat{\boldsymbol{f}}[n]||^2/||\hat{\boldsymbol{f}}[n]||^2$, where $\hat{\boldsymbol{f}}[n]$ is a "ground truth" gas concentration obtained by directly solving the PDE equation, Note that since $\boldsymbol{f}[n]$ is not sparse, the use of NMSE is justified. The corresponding results are shown in Fig. 6 and Fig. 7 for the considered experiment.

Fig. 6a and Fig. 6b present the trajectories of the meander and the proposed exploration strategy. Let us re-iterate that the very nature of the proposed exploration strategy is adaptive, i.e., the algorithm will react to the made measurements. As such, the trajectory is not deterministic. The generated trajectories are overlaid with the simulated concentration field computed at the time when the sources were correctly identified. It can be seen that the gas distribution is driven by three sources located at the concentration peaks.
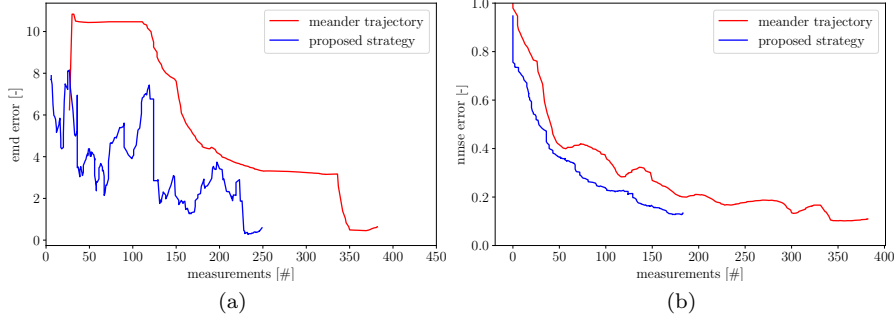
20

Figure 7: Hardware-in-the-loop results: The two plots compare the performance of the meander trajectory and the proposed exploration strategy. In (a) the error is plotted measured regarding the estimated source distributions by means of the Earth Mover's Distance (EMD). In (b) the Normalized Mean Square Error (NMSE) of the estimated concentration field compared to the ground truth is shown.

₅₄₅ Fig. 7 depicts the estimation performance for both source and concentration signals for the evaluated exploration strategies. The curves in Fig. 7a show the EMD error between the estimated source distribution $\boldsymbol{u}[n]$ and the true source distribution $\hat{\boldsymbol{u}}$ in relation to the number of collected measurements. The curves in Fig. 7b depict the NMSE between the estimated concentration field ₅₅₀ and the ground truth. As we see, using the meander trajectory the multi-robot system is able to identify the source distribution after approximately 340 measurements. This is the location where the EMD drops effectively towards zero. Our results show that measurements very close to the source are needed to successfully identify it. Unfortunately, this means that the performance of ₅₅₅ the meander trajectory highly depends on the position of the sources. If they are already covered at the beginning of the trajectory, fewer measurements are needed. However, to be conservative the worst case has to be considered and this means a full coverage of the region. In our example, 360 measurements would be necessary for that. These findings question the meander trajectory as a good ₅₆₀ reference strategy. For future work we would recommend other benchmark algorithms that do not suffer from this property.

In contrast to the meander, the curve for the proposed exploration strategy converges after only 230 measurements in Fig. 7a. This indicates that robots were able to identify the sources with fewer measurements. As can be seen ₅₆₅ from the trajectory in Fig. 6b the measurements are concentrated around the source locations. Obviously, the corresponding measurements contain more information about the sources, which is the reason for a better performance of the proposed exploration strategy. Based on Fig. 7b, it can be seen that the estimated concentration field for both strategies reaches a low NMSE. Here ₅₇₀ the performance of the proposed strategy is better, too. However subjectively speaking the difference is relatively small according to Fig. 7b.
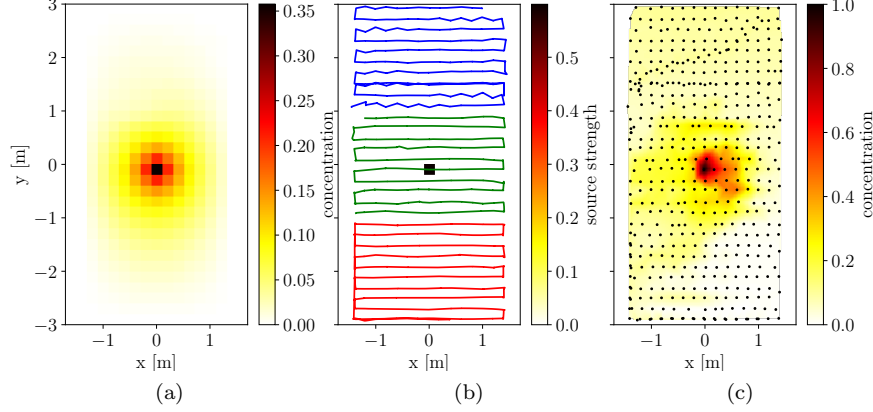
21

Figure 8: Meander trajectory: The figure shows the results of sampling the real ethanol gas distribution by means of a meander trajectory. Where (a) depicts the estimated gas concentration field based on the PDE model, (b) respectively shows the estimated source strength distribution. Here a single peak in x=0, y=0 was detected. These two distributions correspond to the time stamp, when full coverage by the trajectories was reached. Further, (c) illustrates the raw measurements taken at the location marked by a dot and linearly interpolated in between. Additionally, the trajectories of the robots are shown in (b).

Additionally, the hardware-in-the-loop experiments enable us to analyze other performance indicators of our algorithm and system properties. For example we can measure the gross data rates containing all overheads caused by ROS, OS, TCP etc. Specifically, in our case a data rate of less than 70kBytes/s is required for the communication link between two robots. This allows us to define specifications for a communication system required for future real-world experiments. Similarly, we can investigate the processor load of the on-board computers caused by the algorithm. In particular, the on-board computers were able to generate way point proposals with an update-rate of 1.0Hz. This is fast enough for typical applications, especially when the measuring and recovering time of MOX sensors are considered. These may be in the range of $1 - 20sec$ for sensor's response and in worst case up to more than $1min$ for the recovery time [46].

7.2. Real Gas Experiment

Now, we consider the experiment with a single real ethanol gas source. Here we employ only 3 robots. In a first step we repeat the exploration with a predefined sweeping trajectory. In this experiment the ethanol source as described in section 6 was placed in the middle of our $3m \times 6m$ lab environment. This field was discretized so as to create a 15×30 spatial sampling grid. In the experiment we assume Dirichlet boundary condition $f(x,t) = 0$ at the borders of the exploration environment. Note that this assumption is only chosen due to a lack

of the actual knowledge about the conditions at the border of the exploration environment. Clearly, this also causes a mismatch between the reality and the used model. Also, in order to take the inertia of the gas sensor into account (response time of the sensor is < 2s [47]) a specific measurement procedure is designed as follows:

1. once the robot arrives to the target position, it stops
2. it then waits for 1s to allow the sensor to reach a steady state.
3. then, a measurement starts over the time period of $5s$
4. finally, the robot moves to the next measurement position.

Thus, the actual measurement at one position is averaged over 5s. Such a long measurement time is needed to reduce the measurement noise caused by the ADC, or to smooth out short time-scale turbulence in the airflow. However, as the performance is measured with respect to the number of measurements, the actual time needed to take a measurement plays a minor role in the evaluation of the results. We would also like to remark that the overall performance of both considered strategies depends on the number of grid cells. Their number scales with (i) the size of the considered environment and (ii) with the selected spatial resolution, i.e. the size of the cell. The impact of the discretization of the exploration environment in this experiment we leave outside the scope for this paper. For the direct comparison of the two experiments, we made sure that the resolution and the number of cells in the environment is the same.

The results of the meander exploration are shown in Fig 8. In Fig 8b the trajectories are plotted as an overlay above the estimated source strength distribution. The black peak at $x = 0$, $y = 0$ which indicates that a single source at this location with an approximated strength of 0.6 was found. Fig 8a shows the estimated gas concentration based on the PDE model. Further, Fig 8c depicts the raw concentration measurements collected at the locations marked with a dot. [4]

As can be seen, the model-based estimated gas concentration in Fig 8a differs from the raw observations in Fig 8c. Although the estimated gas distribution exhibits a radial shape[5], the real concentration shows an asymmetric pattern. This asymmetric pattern with increased concentration in the lower left area was observed during multiple sweeping experiments. A possible explanation for this phenomenon would be a room specific airflow, caused by a none-airtight door located in the lower left corner of the lab, or by an aligned temperature-gradient in the room. Also, the air flow introduced by the fan placed over the source may cause an asymmetric concentration distribution. Even though, the observations of the real gas dispersion do not perfectly match our model assumption, the source position estimated based on the model is correct within the accuracy of the discretization. Thus, the diffusion PDE approximates the

---

[4]Between the measurement locations the concentration was linearly interpolated to better visualize the raw data.

[5]This is so since a diffusion PDE with a single source has a Gaussian function as a solution.
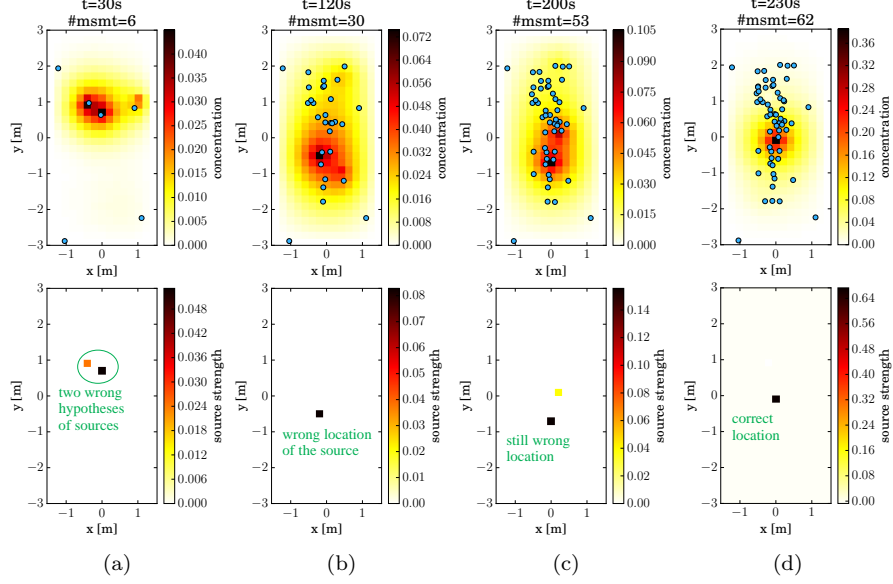
Figure 9: Proposed strategy: Here four snapshots of the estimated concentration and source distribution are shown. In (a) after 30s and 6 measurements, in (b) after 120s and 30 measurements, in (c) after 200s and 53 measurements and in (d) after 230s and 62 measurements. The blue dots indicate the locations of measurements taken until the corresponding time stamp.

real gas dispersion sufficiently well in the scenario considered here. It shows the advantage of the probabilistic approach that is able to deal with imperfect model assumptions.

635   As we can see from the experiments with a meander-based movement strategy, having a sufficient number of measurements allows a PDE-based diffusion model to approximate well the gas and source distribution. Now, in the second step we switch to the proposed model-based exploration strategy as described in Section 5. We also keep the same experimental setup. The results of the an experiment run are exemplarily shown in Fig. 9. In Fig. 9a to 9d snapshots of the estimated concentration and source distribution are shown.

    In Fig. 10 we compare the performance of the meander trajectory with that generated by the proposed exploration strategy using the same EMD and NMSE metrics as in first experiment. Let us also point out that for some time instances the estimated source distribution was exactly equal to zero for all grid cells. For these cases the EMD criterion could not be computed. Thus, we excluded these occurrences from the performance plot in Fig. 10a and linearly interpolated the error curve; these segments are indicated by dotted lines.

650   From the EMD plot it can be seen that the meander trajectory correctly identifies the source after approximately 225 measurements, as indicated by the convergence of the error curve. The 225 measurements correspond to exactly the half of all 450 cells. In other words the robot of the second trajectory (i.e. green

24

in Fig. 8b) is very close to the source. This shows once more that measurements
close to the source are very important for the reconstruction. Unfortunately, it
also shows that the performance of the meander highly depends on the position
of the source. This is a general drawback of the meander trajectory. In contrast
by the proposed strategy the robots were able to identify the source with less
than 100 measurements. At this point let us comment on the discretization
that was used. A perfect localization within the accuracy of the discretization
corresponds to an EMD error of less or equal to 1. This is due to the fact
that in our setup the actual source is located at the edge of two cells, but for
calculating the EMD we consider the ground truth source distribution as a single
peak located only in one cell.

Now, let us study the MSE performance. In contrast to the hardware-in-
the-loop experiments, in the real world experiments the NMSE in Fig.10b of the
estimated concentration is rather high, especially, in case of some runs with the
proposed strategy. Of course, this is not a very precise performance indicator,
since the concentration was not compared to real ground truth values but to
the collected raw data by the meander trajectory. However, it gives a first hint
on a certain mismatch between the model assumptions of the gas dispersion
and reality. Besides, the meander strategy achieves a lower NMSE in Fig.10b
compared to the proposed strategy. While the EMD error plot shows that
the proposed strategy is better for source localization since it achieves a low
estimation error for the source distribution faster, for concentration mapping
tasks a predefined sweeping trajectory may be better since the accuracy of the
estimated concentration field is better according to the NMSE when model
mismatches are present.

Still the NMSE is very high. Unfortunately in the real world experiment this
can be either caused by a wrong reconstruction (i.e. model mismatch) or the
way we defined our "ground truth". For example we consider only a static map
based on the raw measurements as ground truth, however the gas concentrations
is a dynamic process that changes over time.

8. Discussion and Conclusion

The results of the experiments have shown that a model-based exploration is
of advantage for sampling a gas diffusion process in order to localize gas sources.
Using an intelligent exploration strategy, the number of required measurements
can be reduced while the capability to locate gas sources and to some degree
also to map the gas distributions is preserved. This property is favorable for
applications, where a measurement is expensive or consumes a lot of time.

The potential of the presented approach arises from the uncertainty driven
strategy for taking new measurements in combination with the assumption that
the sources are sparsely distributed in the environment and their number is
small. This assumption is encoded with a prior PDF that assumes the proba-
bilistic source strength distribution to have zero mean and unknown variance.
Under this assumption evidence for a source with non-zero posterior mean effec-
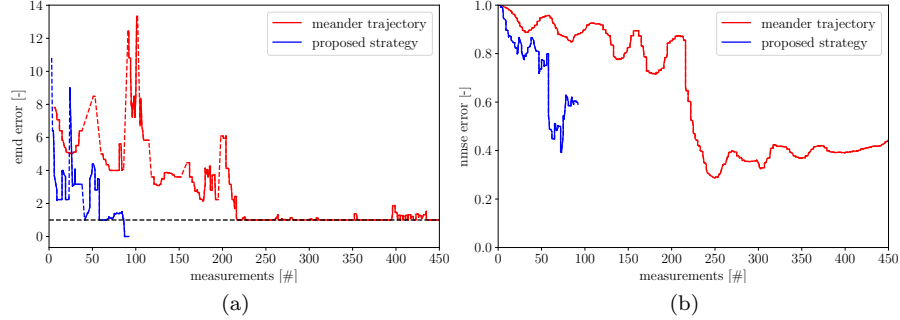tively "contradicts" this prior assumption; as such the uncertainties of the esti-

Figure 10: Performance in real experiments: The two plots compare the performance of the meander trajectory and the proposed exploration strategy. In (a) the error is plotted regarding the estimated source distributions by means of the Earth Mover's Distance (EMD). Here ground truth data were available. The plot (b) shows the Normalized Mean Square Error (NMSE) of the estimated concentration field compared to the raw data collected by the meander in Fig 8c.

mated sources in the corresponding regions grow. As a consequence, the robots concentrate their measurements on informative regions around the sources ac-
₇₀₀ cording to the proposed exploration strategy. It also implies that based on our approach robots tend to repeat measurements at the same location several times in contrast to classical coverage or sweeping algorithms where each point in space is only measured once. Intuitively, this feature is not counter-productive and seems to be important for monitoring dynamic processes. Repetitive measure-
₇₀₅ ments do provide information about time dependencies and are able to reduce the impact of noise. On the down side we observed that favoring the regions around the sources, may cause single robots to get stuck in the neighborhood of a source. This may prevent - under certain conditions - to discover all sources. However, we did not observe this in our experiments. Multiple robots, however,
₇₁₀ are able to effectively reduce the uncertainty around the sources by multiple simultaneous measurements, and individual robots can "escape" from the source location. In future work a mechanism to declare a source as being found, after its PDF has converged, may also help to explore other regions not visited yet.

In the hardware-in-the-loop experiments we have shown the potential of this
₇₁₅ model-based multi-robot exploration for gas source localization. While it seems obvious that prior information based on a model facilitates the exploration, it is unclear how the model-mismatch between the used dispersion model and real gas dispersion affects the exploration. Unfortunately a detail analytical study or quantification of the model mismatch is impossible in real world conditioned by
₇₂₀ the lack of ground truth data. However, it would be possible to setup simulations with a defined discrepancy in the model used by the exploration strategy and the one used by the simulator. This is part of our future work. Nevertheless, from the practical point of view in the experiments with ethanol gas we have

26

shown that our exploration approach is able to handle a certain degree of model-mismatch quite well. The high error between the reconstructed concentration field and gas measurements indicates clearly that the model does not fit perfectly to the observations. Nonetheless, the exploration strategy is able to identify the source, despite model mismatches.

We believe that key reasons for the model mismatch are turbulence and advection mechanisms that are not explicitly represented with (1), yet represent important phenomena in reality. Incorporation of these effects into the model is part of future work. For example, the advection-based diffusion can be incorporated directly into the PDE through a convection term.

A further important advantage of the presented approach is its robustness with respect to the sensor calibration. Correct concentration measurements with exact scaling are not needed. Instead normalized concentration values are sufficient as long as the outputs of the sensors on different robots are all in the same range. Moreover, we noticed that it is even possible to localize the source when a linear relation between the concentration and the sensor voltage is assumed.

Nonetheless, there are also some limitations of the presented approach. For instance in the experiments the source was placed in the middle of the region to be explored. In general it is easier to locate a source that is farther away from the border. Furthermore, only one source was considered in the real world experiments. Thus, additional experimental studies should be considered in future work.

In summary, let us mention that realistic gas dispersion problems are complex dynamical processes. From a practical perspective having a very complex model that adequately represents reality might lead to computationally very complex inverse problems. Instead, simpler probabilistic models, like the one used in this work, can be seen as a numerically feasible approximation, which can be estimated based on concentration measurements in an adequate (in the context of robotic exploration) amount of time.

Appendix A.

In Section 7.2 we have shown a single exemplary experiment run. Of course we performed several runs to evaluate our approach. However, in order to compare two different runs it would be necessary to have similar environmental conditions. The two runs (meander and proposed strategy) are carried out immediately after each other within less than $30min$. With an increasing time in between experiments the environmental condition changes. For example because of the evaporation ethanol the background concentration level is rising. Further because of different temperature gradient in the room at different times of the day, the air flow may look different. Nevertheless, we present four experiment runs in Figure A.11, where the fourth correspond to the experiment presented in Section 7.2. The other three were carried out at different days. A comparison based on the EMD is still valid since the source position has not changed.

However, the NMSE should be considered with caution because of possible changes in the environmental conditions.
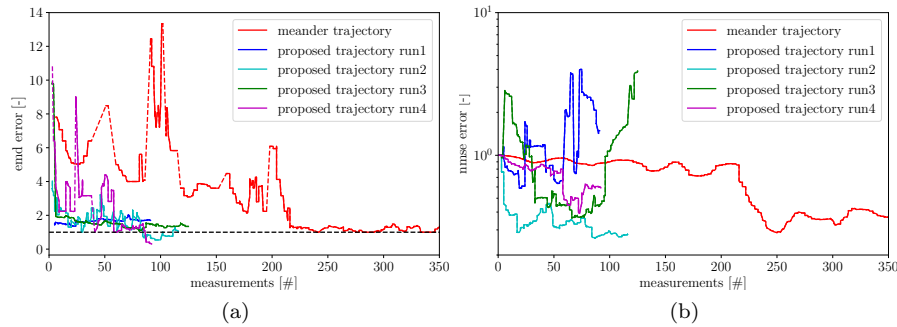

Figure A.11: Performance in real experiments: The two plots compare the performance of the meander trajectory and the proposed exploration strategy in four different experiment runs. In (a) the error is plotted regarding the estimated source distributions by means of the Earth Mover's Distance (EMD). Here ground truth data were available. The plot (b) shows the Normalized Mean Square Error (NMSE) of the estimated concentration field compared to the raw data collected by the meander.

Acknowledgment

References

[1] A. Marjovi, L. Marques, Multi-robot odor distribution mapping in realistic time-variant conditions, 2014 IEEE International Conference on Robotics and Automation (ICRA) (2014) 3720–3727 doi:10.1109/ICRA.2014.6907398.

[2] A. J. Lilienthal, A. Loutfi, T. Duckett, Airborne chemical sensing with mobile robots, Sensors 6 (11) (2006) 1616–1678. doi:10.3390/s6111616.

[3] G. Kowadlo, R. A. Russell, Robot odor localization: A taxonomy and survey, International Journal of Robotics Research 27 (8) (2008) 869–894. doi:10.1177/0278364908095118.

[4] H. Ishida, Y. Wada, H. Matsukura, Chemical sensing in robotic applications: A review, IEEE Sensors Journal 12 (11) (2012) 3163–3173. doi:10.1109/JSEN.2012.2208740.

28

[5] A. Lilienthal, T. Duckett, Building gas concentration gridmaps with a mobile robot, Robotics and Autonomous Systems 48 (1) (2004) 3–16. `doi:10.1016/j.robot.2004.05.002`.

[6] R. A. Russell, A. Bab-Hadiashar, R. L. Shepherd, G. G. Wallace, A comparison of reactive robot chemotaxis algorithms, Robotics and Autonomous Systems 45 (2) (2003) 83–97. `doi:10.1016/S0921-8890(03)00120-9`.

[7] B. Gao, H. Li, F. Sun, 3D moth-inspired chemical plume tracking, 2015 IEEE International Conference on Robotics and Biomimetics, IEEE-ROBIO 2015 24 (1) (2015) 1786–1791. `doi:10.1109/ROBIO.2015.7419031`.

[8] H. Hajieghrary, M. A. Hsieh, I. B. Schwartz, Multi-agent search for source localization in a turbulent medium, Physics Letters, Section A: General, Atomic and Solid State Physics 380 (20) (2016) 1698–1705. `doi:10.1016/j.physleta.2016.03.013`.

[9] M. Vergassola, E. Villermaux, B. I. Shraiman, 'Infotaxis' as a strategy for searching without gradients, Nature 445 (7126) (2007) 406–409. `arXiv:t8jd4qr3m`, `doi:10.1038/nature05464`.

[10] J. D. Albertson, T. Harvey, G. Foderaro, P. Zhu, X. Zhou, S. Ferrari, M. S. Amin, M. Modrak, H. Brantley, E. D. Thoma, A Mobile Sensing Approach for Regional Surveillance of Fugitive Methane Emissions in Oil and Gas Production, Environmental Science & Technology 50 (5) (2016) 2487–2497. `doi:10.1021/acs.est.5b05059`.

[11] H. Ishida, Y. Kagawa, T. Nakamoto, T. Moriizumi, Odor-source Localization In Clean Room By Autonomous Mobile Sensing System, Proceedings of the International Solid-State Sensors and Actuators Conference - Transducers '95 1 (1995) 115–121. `doi:10.1109/SENSOR.1995.717349`.

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

[13] J. R. Gemerek, S. Ferrari, J. D. Albertson, Fugitive gas emission rate estimation using multiple heterogeneous mobile sensors, in: 2017 ISOCS/IEEE International Symposium on Olfaction and Electronic Nose (ISOEN), 2017, pp. 1–3. `doi:10.1109/ISOEN.2017.7968897`.

[14] S. Li, Y. Guo, B. Bingham, Multi-robot cooperative control for monitoring and tracking dynamic plumes, Proceedings - IEEE International Conference on Robotics and Automation (2014) 67–73`doi:10.1109/ICRA.2014.6906591`.

[15] M. A. Demetriou, T. Egorova, N. A. Gatsonis, Estimation of a gaseous release into the atmosphere using a formation of UAVs, IFAC-PapersOnLine 49 (18) (2016) 110–115. `doi:10.1016/j.ifacol.2016.10.148`.

29

[16] M. A. Demetriou, D. Ucinski, State Estimation of Spatially Distributed Processes Using Mobile Sensing Agents, American Control Conference (ACC) (2011) 1770–1776.

[17] F. Pukelsheim, Optimal Design of Experiments, John Wiley & Sons, Inc., 1993.

[18] D. Uciński, Optimal Measurement Methods for Distributed Parameter System Identification, CRC Press, 2004.

[19] 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, Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on (2009) 570–576`doi:10.1109/IROS.2009.5354304`.

[20] J. G. Monroy, A. J. Lilienthal, J. L. Blanco, J. Gonzalez-Jimenez, M. Trincavelli, Probabilistic gas quantification with MOX sensors in Open Sampling Systems - A Gaussian Process approach, Sensors and Actuators, B: Chemical 188 (2013) 298–312. `doi:10.1016/j.snb.2013.06.053`.

[21] D. J. C. MacKay, Information-Based Objective Functions for Active Data Selection, Neural Computation 4 (4) (1992) 590–604. `doi:10.1162/neco.1992.4.4.590`.

[22] P. Whaite, F. Ferrie, Autonomous exploration: driven by uncertainty, IEEE Transactions on Pattern Analysis and Machine Intelligence 19 (3) (1997) 193–205. `doi:10.1109/34.584097`.

[23] T. Wiedemann, C. Manss, D. Shutin, Multi-agent exploration of spatial dynamical processes under sparsity constraints, Autonomous Agents and Multi-Agent Systems 32 (1) (2018) 134–162. `doi:10.1007/s10458-017-9375-7`.

[24] T. Wiedemann, C. Manss, D. Shutin, A. J. Lilienthal, V. Karolj, A. Viseras, Probabilistic modeling of gas diffusion with partial differential equations for multi-robot exploration and gas source localization, in: 2017 European Conference on Mobile Robots (ECMR), 2017, pp. 1–7. `doi:10.1109/ECMR.2017.8098707`.

[25] J. C. Strikwerda, Finite Difference Schemes and Partial Differential Equations, 2nd Edition, Society for Industrial and Applied Mathematics, Philadelphia, 2004.

[26] W. Gong, Approximations of Parabolic Equations With Measure Data, Mathematics of computation 82 (281) (2013) 69–98. `doi:10.1090/S0025-5718-2012-02630-5`.

[27] M. E. Tipping, Sparse Bayesian Learning and the Relevance Vector Machine, Journal of Machine Learning Research.

[28] R. Giri, B. Rao, Type i and type ii bayesian methods for sparse signal recovery using scale mixtures, IEEE Transactions on Signal Processing 64 (13) (2016) 3418–3428. `doi:10.1109/TSP.2016.2546231`.

[29] D. Wipf, S. Nagarajan, A New View of Automatic Relevance Determination, Advances in Neural Information Processing Systems 20 (2008) 1625–1632.

[30] D. P. Wipf, B. D. Rao, Sparse Bayesian learning for basis selection, Signal Processing, IEEE Transactions on 52 (8) (2004) 2153–2164. `doi:10.1109/TSP.2004.831016`.

[31] M. E. Tipping, A. Faul, J. J. T. Avenue, J. J. T. Avenue, Fast marginal likelihood maximisation for sparse bayesian models, in: Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics, 2003, pp. 3–6.

[32] D. Shutin, T. Buchgraber, S. R. Kulkarni, H. V. Poor, Fast variational sparse bayesian learning with automatic relevance determination for superimposed signals, IEEE Transactions on Signal Processing 59 (12) (2011) 6257–6261. `doi:10.1109/TSP.2011.2168217`.

[33] D. Shutin, S. R. Kulkarni, H. V. Poor, Incremental reformulated automatic relevance determination, IEEE Transactions on Signal Processing 60 (9) (2012) 4977–4981. `doi:10.1109/TSP.2012.2200478`.

[34] D. Wipf, S. Nagarajan, A new view of automatic relevance determination, in: Proceedings of the 20th International Conference on Neural Information Processing Systems, NIPS'07, Curran Associates Inc., USA, 2007, pp. 1625–1632.

[35] T. L. Hansen, B. H. Fleury, B. D. Rao, Superfast Line Spectral Estimation, ArXiv e-prints`arXiv:1705.06073`.

[36] H. A. Loeliger, An Introduction to Factor Graphs, Signal Processing Magazine, IEEE 21 (1) (2004) 28 – 41.

[37] F. R. Kschischang, B. J. Frey, H. Loeliger, Factor graphs and the sum-product algorithm, IEEE Transactions on Information Theory 47 (2) (2001) 498–519. `doi:10.1109/18.910572`.

[38] J. Pearl, Probabilistic Reasoning in Intelligent Systems, Kaufmann, San Francisco, 1988.

[39] J. Winn, C. M. Bishop, Variational Message Passing, Journal of Machine Learning Research 6 (2005) 661–694.

[40] T. Wiedemann, C. Manss, D. Shutin, Multi-agent exploration of spatial dynamical processes under sparsity constraints, Autonomous Agents and Multi-Agent Systems 32 (1) (2018) 134–162. `doi:10.1007/s10458-017-9375-7`.

[41] V. Hernandez Bennetts, A. J. Lilienthal, E. Schaffernicht, S. Ferrari, J. Albertson, Integrated simulation of gas dispersion and mobile sensing systems, in: Workshop on Realistic, Rapid and Repeatable Robot Simulation, 2015.

[42] A. Logg, K.-A. Mardal, G. Wells, Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book, Springer Publishing Company, Incorporated, 2012.

[43] SGX Sensortech Limited, The MiCS-5524 is a compact MOS sensor, 1084 Rev. 6.

[44] K. Ihokura, J. J. Watson, The stannic oxide gas sensor : principles and applications, Boca Raton : CRC Press, 1994, includes bibliographical references and index.

[45] Y. Rubner, C. Tomasi, L. J. Guibas, A Metric for Distributions with Applications to Image Databases, Computer Vision, 1998. Sixth International Conference on (1998) 59–66doi:10.1109/ICCV.1998.710701.

[46] A. Lilienthal, T. Duckett, A stereo electronic nose for a mobile inspection robot, ROSE 2003 - 1st IEEE International Workshop on Robotic Sensing 2003: Sensing and Perception in 21st Century Roboticsdoi:10.1109/ROSE.2003.1218709.

[47] MiCS Application Note 4 Using MiCS Sensors for Alcohol Detection Application- Breathalyser, Tech. Rep. June, SGX sensortech (2009).