



<http://www.diva-portal.org>

Preprint

This is the submitted version of a paper presented at *IEEE International Conference on Autonomous Systems (ICAS)*.

Citation for the original published paper:

Wiedemann, T., Shutin, D., Lilienthal, A J. (2021)

Experimental Validation of Domain Knowledge Assisted Robotic Exploration and Source Localization

In: *2021 IEEE International Conference on Autonomous Systems (ICAS)* Montreal, Canada

<https://doi.org/10.1109/ICAS49788.2021.9551145>

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

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:oru:diva-97011>

EXPERIMENTAL VALIDATION OF DOMAIN KNOWLEDGE ASSISTED ROBOTIC EXPLORATION AND SOURCE LOCALIZATION

Thomas Wiedemann, Dmitriy Shutin

Institute of Communications and Navigation
German Aerospace Center
Oberpfaffenhofen, Germany

Achim J. Lilienthal

Center of Applied Autonomous Sensor Systems
Örebro University
Örebro, Sweden

ABSTRACT

In situations where toxic or dangerous airborne material is leaking, mobile robots equipped with gas sensors are a safe alternative to human reconnaissance. This work presents the Domain Knowledge Assisted Robotic Exploration and Source Localization (DARES) approach. It allows a multi-robot system to localize multiple sources or leaks autonomously and independently of a human operator. The probabilistic approach builds upon domain knowledge in the form of a physical model of gas dispersion and the *a priori* assumption that the dispersion process is driven by multiple but sparsely distributed sources. A formal criterion is used to guide the robots to informative measurement locations and enables inference of the source distribution based on gas concentration measurements. Small-scale indoor experiments under controlled conditions are presented to validate the approach. In all three experiments, three rovers successfully localized two ethanol sources.

Index Terms— mobile robot olfaction, gas source localization, Bayesian inference, swarm exploration.

1. INTRODUCTION

Mobile robotic platforms, like rovers and Unmanned Aerial Vehicles (UAVs), are the means of choice when it comes to exploration missions in hazardous environments. For example, in disaster relief scenarios or Chemical, Biological, Radiological and Nuclear (CBRN) events, mobile robots can be dispatched to survey an area of interest and provide an overview of the current situation. There it is important to explore the environment as fast as possible and provide reliable information early enough to civil protection agencies and first responders. It is obvious that while a single robot would need a certain time to explore a region of interest, multiple robots can accomplish the same task faster.

However, the deployment of multiple robots brings along several challenges. Whereas a single robot can be easily teleoperated, steering and coordinating many robots in real-time is a too complex task for a single operator or even for a team of coordinators. In addition, in disaster scenarios, very few




Fig. 1: Robotic gas source localization scenario

human resources are available for controlling robots. Autonomy, on the other hand, allows to address these challenges: through cooperation, the robots can coordinate themselves and accomplish the exploration task independently of an operator. One of the key elements required to implement such an autonomous multi-robot system is an exploration strategy – an algorithm that allows the robots to decide where to collect information or measurements.

In this paper, we consider the task of exploring the dispersion of a toxic or dangerous airborne trace substance (referred to as “gas” in the latter) leaking from an unknown number of sources (see Fig. 1). Our goal is to localize the sources using gas concentration measurements taken by in-situ sensors mounted on the robots. This paper shortly presents our Domain Knowledge Assisted Robotic Exploration and Source Localization (DARES) strategy developed in [1]. As the main contribution, the paper presents results of an evaluation of the DARES approach in experiments under laboratory conditions. In contrast, our previous work studied the approach only in simulations.

2. DOMAIN KNOWLEDGE ASSISTED ROBOTIC EXPLORATION AND SOURCE LOCALIZATION

In the past, many gas source localization strategies for robotic applications were based on the idea that the gas concentra-

tion rises monotonously with proximity to a source. These approaches are often referred to as *chemotaxis* [2, 3]. Using *chemotaxis* a robot tries to follow the gradient of the gas concentration. Such approaches are often supported by additional information like airflow or wind [4]. However, the monotonicity of the concentration distribution does not hold in many real-world environments, since gas dispersion is disturbed by turbulence [5]. Over time, more sophisticated gas source localization strategies have emerged. They take into account more complex mathematical models of the gas dispersion process [6, 7, 8]. These strategies aim at maximizing the information gain obtained with collected measurement data. Consequently, these methods are termed *infotaxis* [9]. The proposed DARES approach likewise follows an infotactic concept.

Our key idea is to assist the robots by *a priori* available domain knowledge about the gas dispersion process. With this additional information, the robots can localize the sources faster, i.e. with fewer measurements, as shown in [10]. In particular, we assist the robots by providing a physical description of gas dispersion in terms of a Partial Differential Equation (PDE). Additionally, since the exact number of sources is assumed as unknown, we endow the model with an assumption that the sources are sparsely distributed. This weak assumption turned out to be very beneficial in order to localize the sources [1].

In what follows we explain how to encode our knowledge and assumptions in a probabilistic gas dispersion model suitable for estimating the sources from concentration measurements. Afterward, the model is used to design an exploration strategy that guides the robots to informative measurement locations with the objective to reduce the uncertainty of the estimates.

2.1. Probabilistic Gas Dispersion Model

From physics it is known that the gas dispersion process over some domain of interest Ω can be approximated by the (stationary) advection-diffusion PDE [11]:

$$\begin{aligned} -\nabla^2 f(\mathbf{x}) + \mathbf{v}(\mathbf{x}) \nabla f(\mathbf{x}) &= u(\mathbf{x}), & \mathbf{x} \in \Omega & \quad (1) \\ \text{s.t. } f(\mathbf{x}) &= 0, & \mathbf{x} \in \partial\Omega & \quad (2) \end{aligned}$$

where $f(\mathbf{x}) = 0$ is a boundary condition on the boundary $\partial\Omega$.

Here, we restrict ourselves to the static two dimensional case ($\Omega \subset \mathbb{R}^2$), where function $f(\mathbf{x})$ denotes the gas concentration at location \mathbf{x} . The right-hand side of (1) models the source distribution. More precisely, the function $u(\mathbf{x})$ represents the source strength or amount of material inflow at a location \mathbf{x} . Furthermore, the vector-valued functions $\mathbf{v}(\mathbf{x}) \in \mathbb{R}^2$ describes the two components of the airflow field at a location \mathbf{x} .

However, estimation of a function $u(\mathbf{x})$, which in general requires application of calculus of variations, cannot be solved analytically. Instead we approximate (1) numerically

using Finite Element Method (FEM) [1]. To this end, the continuous environment Ω is discretized using a finite number of N nodes spanning a mesh. The continuous functions f, u, \mathbf{v} are approximated by a finite number of linear shape functions – finite elements – that linearly interpolate between the mesh nodes.¹ Thus the numerical approximations of the continuous functions can be fully parameterized by the values at the mesh nodes. The values at the mesh nodes are aggregated in vectors $\mathbf{f}, \mathbf{u}, \hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2 \in \mathbb{R}^N$ corresponding to discretization of the continuous functions f, u , and \mathbf{v} , respectively. Thus, the variational problem (1) can be equivalently represented with a system of $N + B$ algebraic equations:

$$\begin{cases} r_i(\mathbf{f}, \mathbf{u}, \hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2) = 0, & i = 1, \dots, N \\ r_i(\mathbf{f}) = 0, & i = N + 1, \dots, N + B \end{cases} \quad (3)$$

where $r_i(\mathbf{f}) = 0, i = N + 1, \dots, N + B$, represent B equations obtained after discretization of the boundary condition (2) at $\partial\Omega$. Essentially, B is the number for of nodes in the mesh that are located at the border of Ω .

We now cast the deterministic model (3) into a probabilistic setting. Instead of demanding r_i to be exactly zero, we assume that the equations only hold with a certain precision τ_s . The derivations (or residuals) of individual equations are assumed to be spatially and temporally white zero-mean Gaussian samples. This assumption results in the following conditional probability density function for the gas concentration given the source distribution and airflow field:

$$p(\mathbf{f} | \mathbf{u}, \hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2) \propto \prod_{i=1}^{N+B} e^{-\frac{\tau_s}{2} (\mathbf{f}, \mathbf{u}, \hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2)^2}. \quad (4)$$

Furthermore, we assume the L robots to be equipped with in-situ gas sensors. Each robot takes K_l noisy concentration measurements at different locations. Thus, the likelihood function of our model can be defined as:


$$p(\mathbf{y}_1, \dots, \mathbf{y}_L | \mathbf{f}) \propto \prod_{l=1}^L \exp\left(-\frac{\tau_m}{2} \|\mathbf{M}_l^T \mathbf{f} - \mathbf{y}_l\|^2\right), \quad (5)$$

where $\mathbf{y}_l \in \mathbb{R}^{K_l}, l = 1, \dots, L$, are measurements taken by the robot l , $\mathbf{M}_l \in \{0, 1\}^{K_l \times N}$ is a binary selection matrix that “picks” elements in \mathbf{f} corresponding to the location of the robot when the measurement was taken. Besides, τ_m is the sensor measurement noise precision.


We aim at a Bayesian inference approach, and therefore we also need to define the prior distribution for the airflow $p(\mathbf{v}_1, \mathbf{v}_2)$, as well as the source prior distribution $p(\mathbf{u})$. What concerns the airflow prior, in this work we assume it to be a multivariate Gaussian:

$$p(\mathbf{v}_1, \mathbf{v}_2) = p(\mathbf{v}_1)p(\mathbf{v}_2) = N(\mathbf{v}_1 | \boldsymbol{\mu}_{v1}, \Sigma_{v1})N(\mathbf{v}_2 | \boldsymbol{\mu}_{v2}, \Sigma_{v2})$$

¹These are also known as Lagrange elements of first-order



(a)



(b)

Fig. 2: The figure shows in (a) the experimental setup in the laboratory. The robots are indicated with A. The two sources are hanging from the ceiling (indicated with B) so that the robots can drive below without collision. The artificial airflow is generated by fan C. As an overlay, the estimated gas concentration after 12min in one experimental run is shown. The error in the estimated source distribution is plotted in (b) over time and averaged over three experiments.

where μ_{vj} , Σ_{vj} , $j = 1, 2$ are fixed design parameters. Practically, these can be set from, e.g., weather forecast or determined using anemometer sensors in the field. The modeling of a source prior is a bit more involved. To incorporate the sparsity assumption we use Sparse Bayesian Learning and represent $p(\mathbf{u})$ using a hierarchical prior. The model is augmented with hyperparameters α , such that $p(\mathbf{u}, \alpha) = p(\mathbf{u}|\alpha)p(\alpha) = N(\mathbf{u}|\mathbf{0}, \text{diag}\{\alpha\}^{-1})\text{Ga}(\alpha)$, where $\text{Ga}(\alpha)$ is a Probability Density Function (PDF) of a Gamma distribution. The hyperparameters α are estimated alongside other parameters. Based on equation (4) and (5) our Bayesian inference approach gives us the posterior

$$p(\mathbf{f}, \mathbf{u}, \mathbf{v}_1, \mathbf{v}_2, \alpha | \mathbf{y}_1, \dots, \mathbf{y}_L) \propto p(\mathbf{y}_1, \dots, \mathbf{y}_L | \mathbf{f})p(\mathbf{f} | \mathbf{u}, \hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2)p(\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2)p(\mathbf{u} | \alpha)p(\alpha). \quad (6)$$

To maximize this posterior we represent it using a factor graph and perform inference using message passing, which can also be implemented in a distributed setting over the network of robots. Due to space constraints, we refer the reader to our works [12, 10] where the inference algorithm is described in more detail. In our setup, each robot calculates messages of a partition of the whole graph and shares the results with the other robots.

2.2. Exploration Strategy

The maximum of the posterior (6) provide us the source distribution \mathbf{u} based on gas concentration measurements \mathbf{y}_l , $l = 1, \dots, L$, taken by robots. In order to achieve a high level of autonomy, however, an intelligent sampling strategy is needed that will guide robots to new, informative sampling locations.

For this purpose, we propose an uncertainty-driven exploration strategy, where the new measurements are taken at locations that maximally reduce the uncertainty of the obtained estimates. To this end, a gauge for the spatial uncertainty is needed with respect to different locations in the environment.

Here the probabilistic inference approach becomes useful. Recall, that \mathbf{f} represents a concentration value at each discretized location of the environment Ω . We can calculate a marginal distribution for each entry f_i , $i = 1, \dots, N$, of \mathbf{f} :

$$p(f_i) \propto \int \dots \int p(\mathbf{f}, \mathbf{u}, \hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \alpha | \mathbf{y}_1, \dots, \mathbf{y}_L) d\mathbf{f} \sim f_i \approx N(f_i | \mu_{f_i}, \sigma_{f_i}^2); i = 1, \dots, N \quad (7)$$

where we use the notation $\int \dots \int d\mathbf{f} \sim f_i$ to indicate a marginalization operation over all variables except for f_i . The proposed idea approximates $p(f_i)$ with a Gaussian distribution and uses the variance $\sigma_{f_i}^2$ as a gauge of concentration uncertainty.

Robots are then sent to locations with the highest variance, i.e. with the highest uncertainty, which in turn implies the highest Shannon entropy of the concentration value. Taking a measurement at this location would reduce the total concentration uncertainty the most.

3. EXPERIMENTAL EVALUATION

We evaluated our approach in experiments under laboratory conditions as shown in Figure 2. We placed two Petri dishes hanging from the ceiling and filled with ethanol modeling 2 sources. They acted as sources of ethanol vapor in our

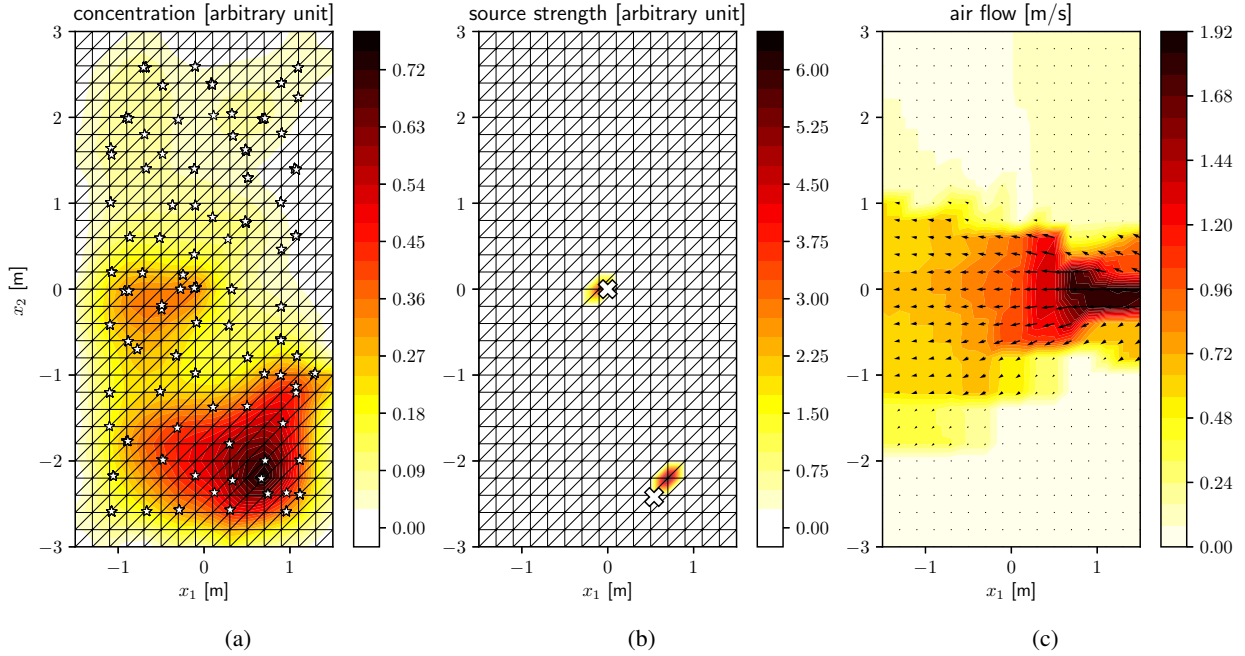


Fig. 3: The figure depicts a snapshot of an exploration experiment after 12min (corresponding to Figure 2). In (a) the estimated gas concentration is shown. Besides, the white stars indicate all measurement locations. In (b) the estimated source distribution (spatial source strength) is plotted, where the white crosses indicate the actual position of the ethanol sources. In (c) the artificial airflow field generated by fans and used in the inference approach is shown.

experiments. Above the culture dishes, fans were mounted blowing air downwards to accelerate evaporation and dispersion. Three small robots were deployed in the experiment. The robots were equipped with Photoionization detectors (alphasense PID-AH2) to measure the ethanol concentration in the room. Also, the robots make use of a camera tracking system for precise indoor localization and can move to desired waypoints enabled by a robotic path planner. We generated an artificial airflow in the room using multiple fans. The air flow turned out to be long-term stable and was sampled before the experiments using multiple anemometers. The measured spatial airflow field, as depicted in Figure 3c, was used as the airflow prior in the inference approach (6). Namely, based on the measurements the parameters μ_{vj} , Σ_{vj} , $j = 1, 2$, of the airflow prior were computed. To evaluate the exploration strategy, we compare the estimated source distribution \mathbf{u} to the ground truth source distribution \mathbf{u}_{gt} by means of the Earth Mover's Distance (EMD) [13]. Vector \mathbf{u}_{gt} is an all-zero vector, with an exception of locations corresponding to sources. Note that in reality, the actual source inflow rate is unknown. Therefore, we set the elements in \mathbf{u}_{gt} corresponding to sources to 1 and normalized the estimated source distribution to 1, too, before comparison. The resulting EMD performance is plotted in Figure 2b and averaged over three experimental runs. It can be seen that by intelligently sampling the gas concentration the robots successively reduce the er-

ror in the estimated source distribution. Figure 3b also shows the estimated source distribution after 12min for one experiment. There, the two peaks in the estimated source distribution nearly perfectly match the actual position of the ethanol sources indicated by the two crosses.

4. CONCLUSION

The paper shortly summarizes the DARES approach towards an exploration of gas sources. Based on a physical model of gas dispersion and the assumption that the dispersion process is driven by multiple sparsely distributed sources, a probabilistic model was formulated. The model is the foundation of a Bayesian inference approach to estimate the source distribution based on gas concentration measurements taken by robots. Further, the model is used to derive an exploration strategy that autonomously guides the robots to informative measurement locations. The DARES approach has been evaluated in small-scale experiments in an indoor environment under controlled conditions. There it has been shown that the robots are able to successfully localize ethanol vapor sources. In the future, the DARES approach needs to be evaluated also in a more challenging environment, for example in outdoor scenarios.

5. REFERENCES

- [1] Thomas Wiedemann, Achim J. Lilienthal, and Dmitriy Shutin, "Analysis of model mismatch effects for a model-based gas source localization strategy incorporating advection knowledge," *Sensors*, vol. 19, no. 3, 2019.
- [2] R. Andrew Russell, Alireza Bab-Hadiashar, Rod L. Shepherd, and Gordon G. Wallace, "A comparison of reactive robot chemotaxis algorithms," *Robotics and Autonomous Systems*, vol. 45, pp. 83–97, 2003.
- [3] Gideon Kowadlo and R. Andrew Russell, "Robot odor localization: A taxonomy and survey," *International Journal of Robotics Research*, vol. 27, no. 8, pp. 869–894, 2008.
- [4] David J. Harvey, Tien Fu Lu, and Michael A. Keller, "Comparing insect-inspired chemical plume tracking algorithms using a mobile robot," *IEEE Transactions on Robotics*, vol. 24, no. 2, pp. 307–317, 2008.
- [5] Ali Marjovi and Lino Marques, "Multi-robot odor distribution mapping in realistic time-variant conditions," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 3720–3727.
- [6] Julian Ruddick, Ali Marjovi, Faezeh Rahbar, and Alcherio Martinoli, "Design and Performance Evaluation of an Infotaxis-Based Three-Dimensional Algorithm for Odor Source Localization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 1413–1420.
- [7] Qing Hao Meng, Wei Xing Yang, Yang Wang, and Ming Zeng, "Collective odor source estimation and search in time-variant airflow environments using mobile robots," *Sensors*, vol. 11, pp. 10415–10443, 2011.
- [8] Hadi Hajieghrary, M. Ani Hsieh, and Ira B. Schwartz, "Multi-agent search for source localization in a turbulent medium," *Physics Letters, Section A*, vol. 380, no. 20, pp. 1698–1705, 2016.
- [9] Massimo Vergassola, Emmanuel Villermaux, and Boris I. Shraiman, "'Infotaxis' as a strategy for searching without gradients," *Nature*, vol. 445, no. 7126, pp. 406–409, 2007.
- [10] Thomas Wiedemann, Dmitriy Shutin, and Achim J. Lilienthal, "Model-based gas source localization strategy for a cooperative multi-robot system - A probabilistic approach and experimental validation incorporating physical knowledge and model uncertainties," *Journal of Robotics and Autonomous Systems*, vol. 118, pp. 66–79, 2019.
- [11] John Crank, *The Mathematics Of Diffusion*, Clarendon Press, Oxford, second edition, 1975.
- [12] Thomas Wiedemann, Christoph Manss, and Dmitriy Shutin, "Multi-agent exploration of spatial dynamical processes under sparsity constraints," *Autonomous Agents and Multi-Agent Systems*, vol. 32, no. 1, pp. 134–162, 2018.
- [13] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas, "A Metric for Distributions with Applications to Image Databases," in *IEEE International Conference on Computer Vision*, 1998, pp. 59–66.