

Active Localization of Gas Leaks using Fluid Simulation

Martin Asenov¹, Marius Rutkauskas², Derryck Reid², Kartic Subr¹, Subramanian Ramamoorthy¹

Abstract—Sensors are routinely mounted on robots to acquire various forms of measurements in spatio-temporal fields. Locating features within these fields and reconstruction (mapping) of the dense fields can be challenging in resource-constrained situations, such as when trying to locate the source of a gas leak from a small number of measurements. In such cases, a model of the underlying complex dynamics can be exploited to discover informative paths within the field. We use a fluid simulator as a model, to guide inference for the location of a gas leak. We perform localization via minimization of the discrepancy between observed measurements and gas concentrations predicted by the simulator. Our method is able to account for dynamically varying parameters of wind flow (e.g., direction and strength), and its effects on the observed distribution of gas. We develop algorithms for off-line inference as well as for on-line path discovery via active sensing. We demonstrate the efficiency, accuracy and versatility of our algorithm using experiments with a physical robot conducted in outdoor environments. We deploy an unmanned air vehicle (UAV) mounted with a CO₂ sensor to automatically seek out a gas cylinder emitting CO₂ via a nozzle. We evaluate the accuracy of our algorithm by measuring the error in the inferred location of the nozzle, based on which we show that our proposed approach is competitive with respect to state of the art baselines.

I. INTRODUCTION

We address the problem of using a robot-mounted sensor to actively search for features of a spatially extended field, e.g., source of a leaking gas, or to map such a field from point measurements. This is useful in numerous applications, such as disaster response [1], scientific data collection in difficult and inaccessible environments [2] and urban emissions mapping [3].

From a robotics perspective, the core problem is that of synthesizing paths with respect to an objective such as quality of reconstruction of the underlying spatial field, or the accuracy of localization of a spatio-temporal event (e.g., source of a gas leak). Traditional approaches for solving such problems include 3D surface reconstruction algorithms [4] and regression models such as Gaussian process [5].

What makes the practical problem challenging, and many of the traditional methods harder to apply, are resource constraints and lack of control over the experimental domain. When the sensor is mounted on an Unmanned Aerial Vehicle (UAV) and flown around from a more distant launch location, the number of samples that may be collected is limited (typically due to power constraints, but also due to other issues such as contamination risks). Moreover, the true underlying phenomena are typically varying, such as

when the field is created by dispersion under wind flows and one must reason about this (e.g., to locate the true source location despite wind). Lastly, the overall task is made difficult by the inability to obtain detailed ground truth to evaluate models - if the objective were to be defined as reconstruction of a spatial or temporal field, then we may not have a detailed description of the true underlying field. Some recent works that are noteworthy in this area, and shown to be successful in experimental settings include, e.g., [6] presents a method based on Gaussian process regression, and [7] which proposes another kernel based approximation model. The latter also accounts for the directionality of wind flows, by interpolating from past data with a modified kernel. In most such work, the treatment of wind as a cause of variation, and associated search strategies for actively locating environmental features (e.g., pinpointing the source location of a leak before wind flow disperses it) is either heuristic or considers wind to be constant over time (in which case reasoning about the flow may not be necessary).

In this paper, we propose the use of a fluid simulation as the model of the underlying phenomenon which allows us to directly address these causes of variation. We pose the problem of estimating the location of an environmental feature (e.g., source) as that of optimizing the fit between a fluid simulation model and the point-wise measurements of the resulting scalar field. This is a process of calibrating the simulation model, which we show can be done online as measurements arrive (indeed, in a way that can be implemented on resource constrained hardware within a standard commercial grade UAV) and used to drive search processes based on information-gain criteria.

The optimization problem can be solved in a number of ways. For the problem of estimating a spatially extended field given a sequence of measurements, one approach is to use Bayesian optimization, treating the true objective function implied by the environment as a random function with a Gaussian process prior. We develop an alternative approach, which exploits the shift-invariance of the phenomenon, that performs a larger simulation once at the beginning and poses location estimation as the problem of determining the optimal translation of a smaller field of interest. The use of the simulation model also allows us to address the active search version of this problem, where the sampling locations are determined online (and on-board), by iteratively computing a likelihood for the source location and flying to the point which maximizes this quantity. This proposed approach, which we call One-shot Grid Search (OGS), is efficient and requires relatively low computational resources for similar localization accuracy.

¹Institute of Perception, Action and Behaviour (IPAB), The University of Edinburgh

²Institute of Photonics and Quantum Sciences, Heriot-Watt University

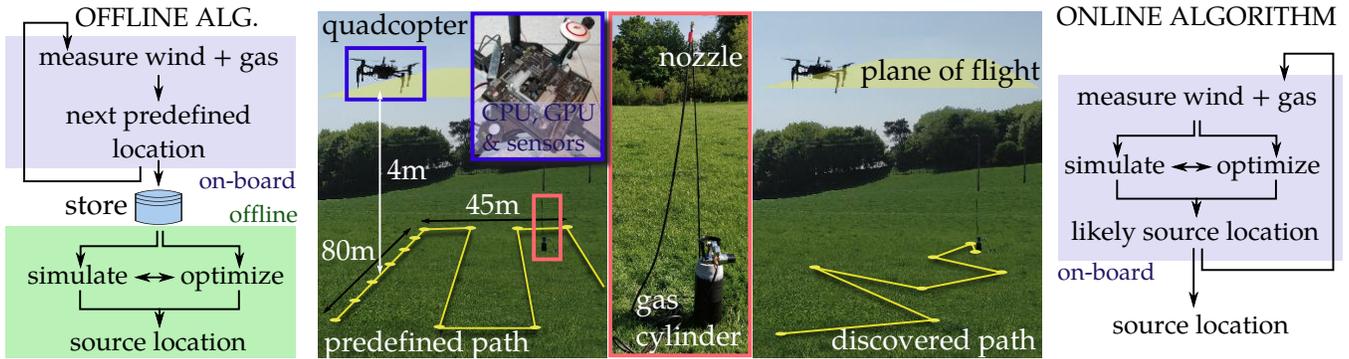


Fig. 1: **Gas-leak localization** We place a gas cylinder releasing CO_2 , via a nozzle, and measure gas concentrations, wind speeds and wind directions using a UAV. Our goal is to estimate the location of the nozzle from these point measurements. The UAV, which is equipped with sensors and on-board processors (CPU and GPU), flies on a fixed 2D plane 4m above the ground. We develop two algorithms for localization. Our *offline algorithm* (left) compares measurements, taken along a predefined path, with a fluid simulation (that considers measured wind). Our *online algorithm* (right) discovers a path to the source of the leak by repeatedly performing simulation and comparison steps on-board the UAV, and flying to the most likely location of the source given the measurements.

II. RELATED WORK AND CONTRIBUTIONS

Robots used as sensing platforms Early work around autonomous sensing of physical phenomena involved ground-based mobile robots [6], [8], [9]. More recently, with the emergence of reasonably robust Unmanned Aerial Vehicle (UAV) platforms, often referred to as drones, they are being used as sensing platforms with benefits in terms of speed, manoeuvrability and ability to deal with hostile terrains, unobstructed by objects on the ground [7], [10]. UAVs bring their own challenges, such as reduced on-board power and the difficulty of finding sensors that fit within the form factor. Sensing technology has also continued to develop, e.g., making it possible to use spectrometers on UAVs [11]. We conduct experiments with a commercial off-the-shelf CO_2 sensor, but note that the computational methods presented here are sensor agnostic, assuming only that the sensor obtains point measurements from a scalar field.

Using simulations as models Models provide numerous advantages in machine learning [12], enabling inferences from limited data, and in planning [13], enabling counterfactual reasoning [14] and guided search. However, defining the structure of models in a way that leads to efficient inference while maintaining fidelity to complex arrangements of physical causes tends to be non-trivial.

The phenomena we consider in this paper involve gas flows. There is a long tradition of modelling such flows, including efficient computational methods aimed at graphics and animation applications [15]. The development of efficient solvers is also driven in the engineering community by the need to simulate phenomena such as fluid-structure interaction, yielding fast and approximate solvers through position-based dynamics methods [16], [17]. Simulation frameworks have also been developed aimed specifically at easing the development and testing of GDM and GSL algorithms [18].

Development of advanced simulation tools has led to new milestones in learning challenging robotics tasks. In [19],

the authors show that approximate simulators can enable the synthesis of complex behaviours such as pouring, that would have been hard to achieve through conventional means. This draws on earlier observations from human psychophysics, e.g., [20], that people seem to be able to reason about the flow of liquids in situations where the available data is necessarily sparse. These papers are situated within the broader topic of ‘intuitive physics’, which refers to the ways in which cognitive models of real-world physical phenomena seem to only require relatively simple representations of the true underlying phenomenon [21], [22]. In restricted settings, such representations have also been used for efficient neural network based model learning [23], [24] and calibration [25]. Such ideas have been explored within the problem of odour localization, by devising naive fluid models and search algorithms [26]; pre-computing dispersion maps using computational fluid dynamics and probabilistically weighting them at test time [27]; updating Gaussian analytical model using evolutionary strategies [28]; using matrix of static sensors [18] but so far haven’t been scaled to realistic outdoor environments, where usually limited samples are available.

In this paper, we utilise a reasonably accurate simulation of the phenomenon [29] but exploit simplifications inherent to the problem, such as that the dispersion process can be modelled on the 2-d plane¹ along which the point measurements are also being taken. Moreover, the process of dispersion is shift invariant [30], so that a single large simulation can be performed online, from which the flow patterns for different locations can be easily computed.

Active and adaptive measurements Specific problems such as the localization of gas sources have been approached using a variety of different algorithmic means. Bio-inspired

¹We observe that our approach is invariant to some degree of (small) noise, i.e., the situation of plain fields and gently rolling hills. Many realistic applications are indeed sited in such terrain, e.g., a petroleum refinery in the periphery of which one might wish to perform emissions monitoring.

approaches have been proposed devising heuristics to follow the wind gradient towards the source [10], [31]. In practice, such heuristics depend on the presence of specific environmental conditions, including constant wind speeds across the field of interest and the UAV being placed within the path of the gas. The source localization estimation has also been addressed using Bayesian methods, using particle filters in outdoor environment [32], [33]; Infotaxis which aim to maximize information gain by reducing the entropy [34]. Gaussian Markov random fields have also been used to address the problem of obstacles in indoor scenarios [35]. Another approach is to formulate the problem as one of regression from sparse measurements. Representative examples of this approach include [6], who use Gaussian process mixtures, and [9]. Kalman filter based estimation algorithms also work similarly [36]. A weakness of these methods has been that they do not explicitly consider the structure of the phenomenon in terms of a source and dispersion through wind flows, although refinements of the above procedures do indirectly account for these effects, e.g., [7], [8], [37].

Another aspect of active sensing is the method for collecting samples so as to maximise a notion of information gain. While the underlying exploration-exploitation tradeoffs can be posed formally in decision theoretic terms, most practical techniques tend to be myopic in their operation. Gaussian processes [38], [39] and the Kernel DM+V/W algorithm [7] address this question. One could also formulate this as optimal design of sequential experiments [40]. However, this requires access to analytically defined dynamics models which may be hard to construct for the specific scenario at hand. Notably, source term estimation was recently addressed with Bayesian estimation implemented using sequential Monte Carlo [41]. By using parameterized Gaussian plume dispersion model recursive Bayesian updates can be performed accounting for uncertainty in wind, dispersion, etc. This was additionally implemented on a UAV [42], performing outdoor localization of gas leaks using predefined flying pattern and ground station for performing computations. We formulate active sensing with a fluid simulation in the loop and devise an efficient algorithm for simulation alignment.

Contributions In this paper, we:

- 1) formulate gas leak localization as model-based inference using a fluid simulator as the model,
- 2) develop a practical optimization algorithm based on a single simulation per iteration,
- 3) develop an online algorithm that locates gas leaks using active sensing,
- 4) demonstrate that our algorithm results in acceptable localization error in real experiments.

III. ONE-SHOT FLUID SIMULATION FOR LOCALIZATION OF GAS LEAKS

We localize the source of a leaking gas based on a discrete set of measurements of gas concentrations, wind speed and wind direction. We make the following assumptions: 1) there is a single source of gas within the domain of interest; 2) the ground plane is relatively flat; 3) gas and wind measurements

are made on a plane parallel to the ground, and above the source of leakage; 4) wind flow is time-varying but spatially constant within the domain (i.e., obstacle-free domain, not large enough for local variation to be significant).

Problem formulation We define a spatio-temporal domain $S \times T$ over which fields of interest, such as gas concentration and wind flow, are defined. Without loss of generalization, we define $S \equiv [0, 1] \times [0, 1]$ and $T \equiv [0, 1]$. We use a UAV to measure gas concentrations $g : S \times T \rightarrow \mathbb{R}$ and wind (speed and 2D direction) $\mathbf{w} : S \times T \rightarrow \mathbb{R}^2$ over this field. The sequence of measurement points are $s_i \in S$ and $t_i \in T$. We abbreviate measurements $g(s_i, t_i)$ and $\mathbf{w}(s_i, t_i)$ as g_i and \mathbf{w}_i respectively. These measurements are dependent on the location of the source of the gas $l \in S$ and hidden parameters θ , such as properties of the gas, which we assume remain fixed throughout an experiment.

We infer the most likely location of the source using a 2D Eulerian fluid simulator as our model. The simulator can be seen as a mapping from l to an approximate gas concentration field given the hidden parameters θ and a dense wind field. We rewrite this as a mapping from leak locations and spatio-temporal sites to gas concentrations $\tilde{g} : S \times S \times T \rightarrow \mathbb{R}$, given θ and $\tilde{\mathbf{w}}$. Here $\tilde{\mathbf{w}} : T \rightarrow \mathbb{R}^2$ is the spatially constant but temporally dense wind field reconstructed from sparse spatio-temporal measurements \mathbf{w} taken by the UAV. The gas concentration predicted by the model at s_i and t_i (where measurements were taken) are abbreviated as $\tilde{g}_i \equiv \tilde{g}(l, s_i, t_i | \theta, \tilde{\mathbf{w}})$. Our goal is to find a location l^* where measurements agree best with model predictions, i.e.,

$$l^* \equiv \underset{l \in S}{\operatorname{argmin}} \sum_{i=0}^k d(g_i, \tilde{g}_i) = \underset{l \in S}{\operatorname{argmin}} \sum_{i=0}^k d(g(s_i, t_i), \tilde{g}(l, s_i, t_i | \theta, \tilde{\mathbf{w}})) \quad (1)$$

where $d(\cdot, \cdot)$ is an appropriate distance metric and k is the number of measurements taken by the UAV. Fig. 2a summarises our notation.

One-shot Grid Search (OGS) The optimization in Eq. 1 can be performed naïvely by computing the objective function for many values of l , obtained by densely sampling S using a regular grid of $m \times n$ samples, and selecting the location l^* that yields the minimum value. This approach would be inefficient since it requires mn simulations to be performed. Our key observation is that the *model is shift invariant* if the parameters (θ and $\tilde{\mathbf{w}}$) are spatially stationary. Concentrations \tilde{g} produced by mn simulations with different source locations are identical up to a translation:

$$\tilde{g}(l + \Delta l, s_i, t_i | \theta, \tilde{\mathbf{w}}) = \tilde{g}(l, s_i + \Delta l, t_i | \theta, \tilde{\mathbf{w}}) \quad \forall (l + \Delta l) \in S, (s_i + \Delta l) \in S. \quad (2)$$

Fig. 2b illustrates the central idea. In addition, if we use an Eulerian simulator, a single run of the simulator \tilde{g} with a source located at l_0^* can be evaluated at several $(s_i + \Delta l, t_i)$ cheaply. Rather than repeating the simulation with different

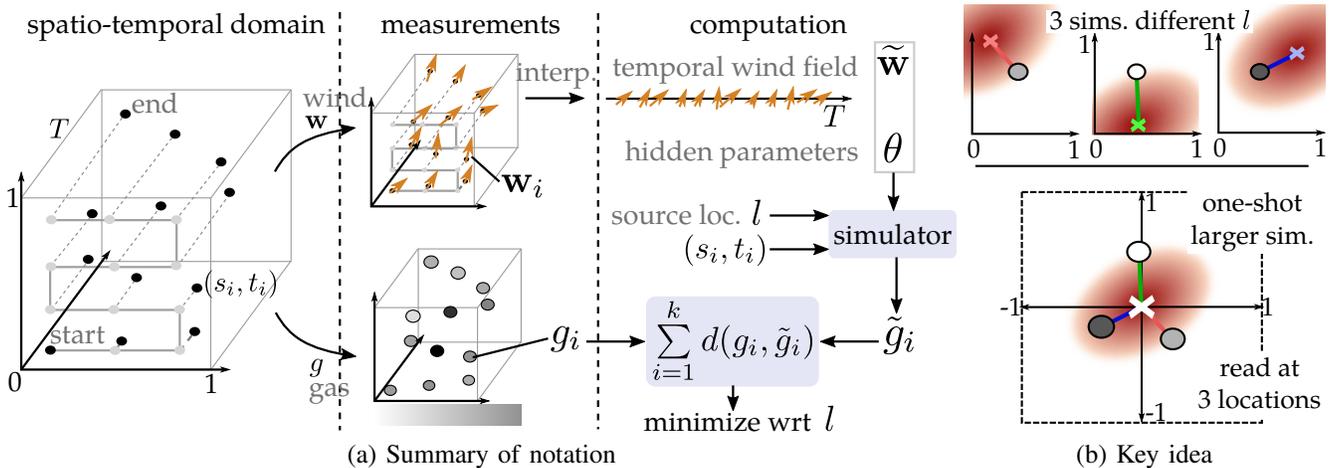


Fig. 2: **Overview** We locate the source of the leak by comparing the output of our physically-based model, \tilde{g}_i , with measurements at corresponding locations g_i . Rather than running multiple simulations with different values of l , we run one larger ($4\times$ domain) simulation and read multiple values from appropriate relative locations. We assume that \tilde{w} is only a function of time (spatially invariant), and so simulation output is spatially shift-invariant. Although the simulation output is a 3D gas concentration field, we visualize the spatial distribution (blurry red splats) at a slice in time.

values of l_k^* , we run the simulator once with $l_0^* \equiv (0, 0)$ and read several values of the gas concentration field $M_{ij} \equiv \tilde{g}(l_0^*, s_i + \Delta l_j, t_i, |\theta, \tilde{w}|)$, $j = 1, \dots, mn$. We construct the entire $k \times mn$ matrix \mathbf{M} using only one simulation. Each column of \mathbf{M} contains \tilde{g}_i corresponding to a source location. We solve the optimization problem in Eq. 1 by identifying p as follows:

$$l^* = \Delta l_p \text{ where } p = \underset{j \in \{1, \dots, mn\}}{\operatorname{argmin}} \sum_{i=0}^k d(g_i, M_{ij}). \quad (3)$$

Since this OGS approach requires the shifted source (and samples) to be within the domain, we run the simulation on a larger domain $S \equiv [-1, 1] \times [-1, 1]$, always place the source at the origin and adjust the relative locations read appropriately (see Fig. 2b). Thus, rather than running mn simulations, we run one simulation with four times as many Eulerian grid cells. We define q as the vector of distances of each column of \mathbf{M} to g_i and use this to derive the likelihood for different source locations on a grid.

Wind estimation The above formulated approach relies on having access to wind estimates \mathbf{w} . In order to acquire such measurements on a UAV, we use the pitch, yaw and rotation provided by the IMU which through a series of transformations can estimate the current wind [7], [43]. The rationale is that the wind speed and direction are directly related (through the transformation implied by a flight dynamics model) to the control signal that must be applied within the UAV, when hovering in place. Depending on the direction of the wind, the UAV will lean in a different way; the stronger the wind, the more it will lean. Using this approach the direction of the wind can be directly estimated, however the inclination angle of the UAV with respect to the ground has to be calibrated with respect to the strength of the wind speed. We use an off-the-shelf wind simulator provided with commercial UAV we

use, although more precise methods such as a wind tunnel can be used [7]. Importantly, we show that the proposed method is robust to imprecise wind speed measurements.

Fluid simulator Another requirement for our method is having access to a simulator oracle \tilde{g} . There are many ways to express a physical model of fluids, but we opt to use a stable Navier-Stokes solver due to its efficiency and ease of implementation [29]. The solver is realized by dividing the space into voxels and iteratively updating the velocity and density. The differential equation for solving for the density is linear with respect to the density term and thus easier to solve. For solving for the velocity a semi-Lagrangian technique is used, producing stable result like the density solver². The simulator depends on different parameters - we assume that we have prior knowledge of the diffusion properties of the gas, *accurate* wind direction and *approximate* wind speed can be estimated as described above, and we normalize the gas readings to be invariant of the quantity of gas released. The rest of the parameters, number of cells (fidelity of the simulation), number of wind locations, simulation timestep, solver iterations, tend to be a trade-off between the accuracy and speed of the simulation.

IV. EXPERIMENTS

First, we perform a series of offline experiments, by collecting data with a UAV³ flying at predefined waypoints. We compare our method to existing gas localization and mapping baselines found in literature. We then benchmark our optimization method against standard approaches for

²An online demo of the simulator and its behaviour can be found at <https://gas-drone-simulation.neocities.org>

³We do not model the turbulent effect of the propellers - using a smoke flare, we visually inspected the effect of the propellers and found that it has little impact on the larger scale gas dynamics.

solving the proposed simulation alignment problem. Finally, we conduct sensitivity analysis of the different hyperparameters of the fluid simulation used.

Secondly, we carry out a set of online experiments, dynamically selecting new waypoints as part of the optimization procedure. We use readings from a noisy simulator as data to evaluate the performance of our method against other approaches. Finally, we perform active sensing experiments on a UAV using our algorithm.

A. Offline algorithm

In order to evaluate the proposed approach we conducted a total of 13 flights. Each flight, taking approximately 10 minutes, visits 16 waypoints. We use a DJI M100 UAV, integrated with a TX2 for data logging and processing and CozIR-A CO₂ sensor. As a gas source we use a compressed CO₂ cylinder. The UAV flies at a constant height of 4 m, covering an area of 80 m × 45 m, with the bottle being placed at an unknown location somewhere within that area (see Fig. 1). As we do not have access to the ground truth gas distribution, we evaluate our algorithm using *localization error* – the distance between the location of the maximum gas concentration and the true location of the cylinder (determined through GPS measurement).

Comparison with related work In the first set of experiments, we compare our algorithm against standard approaches from the literature, such as one using Gaussian Process regression [6] with a Radial Basis Function kernel, variance 15 and lengthscale 7, as well as the TD Kernel DM+V/W algorithm [8], with cell size 0.2, kernel size 10, evaluation radius 10, time scale 1 and wind scale 0.001 (we perform grid search to find optimal parameters for the baselines). The collected air samples are used to fit a 2-d concentration map, with DM+V/W additionally using wind samples for reshaping the kernel function and scaling the readings based on the timestep taken. As previous approaches do not explicitly model the source location, we use the peak of the posterior as a proxy for this quantity [7]. We compare this against the computed l^* from OGS (Sim-Likelihood). For the simulation parameters, we use gas release 25, simulation fidelity 1/1 and diffusion $1e-4$ (see Fig. 5). We show examples of the posterior mean (GP and DM+V/W) and likelihood of the source of the gas (ours), together with the overall error between the methods in Fig. 3.

Bayesian optimization vs one-shot grid search A standard approach to solving the optimization problem in Eq. 1 would be to treat the objective function as a random function with a Gaussian process (GP) prior over it. Then, based on each measurement, the prior could be updated to form a posterior distribution over the objective function. Based on this posterior, an *acquisition function* can be constructed to determine the next sample location. We perform experiments using this approach, in order to then compare the results against our proposed optimization procedure. We start with a GP prior using the Matern52 kernel, sampling from it twice to obtain a first (random) estimate l_0^* . Then, we estimate

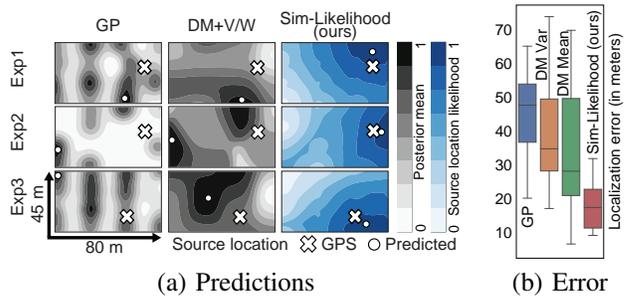


Fig. 3: **Predictions** (a) Posterior mean/variance of gas concentrations obtained using state of the art methods (column 1 [6] and column 2 [8]) against the likelihood computed using our method (column 3). Rows represent results of different experiments. The source location was the same for the first two experiments, but environmental conditions (wind) were different. (b) Overall error across all experiments.

$d(g_i, \tilde{g}(l_0^*, s_i, t_i | \theta, \tilde{w}))$ and use it to update the posterior (which only contains two samples). Based on this posterior, and using an acquisition function (Lower Confidence Bound, with alpha parameters 0.5, 1, 2 and 3, Expected Maximization or Maximum Probability of Improvement), we determine the next sampled source location l_1^* . We then continue sampling for new leak locations, as prescribed by the corresponding acquisition function. For each of these locations, we run simulations and update the GP posterior. We evaluate these results from Bayesian optimization against the corresponding values for OGS, measuring running times and precision for the two optimization methods over a limited number of location samples l as shown in Fig. 4.

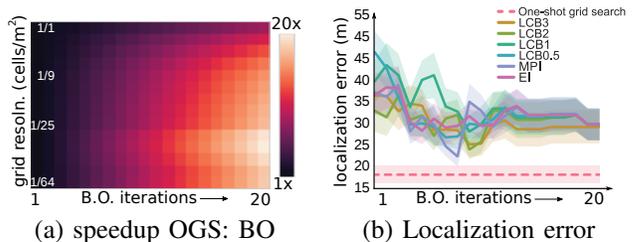


Fig. 4: **One-shot grid search (OGS) vs Bayesian Optimization (BO)** (a) **Compute time** Speed up in the computational time for different resolution of the simulations. (b) **Localization error** We use all the experimental data collected from Sec. IV-A to evaluate the localization error of different acquisition functions of BO with OGS.

Sensitivity to inaccuracy in parameters In order to determine the robustness of our approach to errors in the setup of the simulations, we perform a sensitivity analysis. We analyze the effect of different parameters on localization error by artificially perturbing the underlying values. We study the effects of inaccuracies in the quantity of gas released, diffusion coefficient, wind speed, wind direction and simulation grid resolution. We start with default values for the parameters (see caption of Fig. 5), and then assess

localization error when each parameter is individually modified to one of six different values. The resulting errors and the perturbed values of each parameter are shown in Fig. 5.

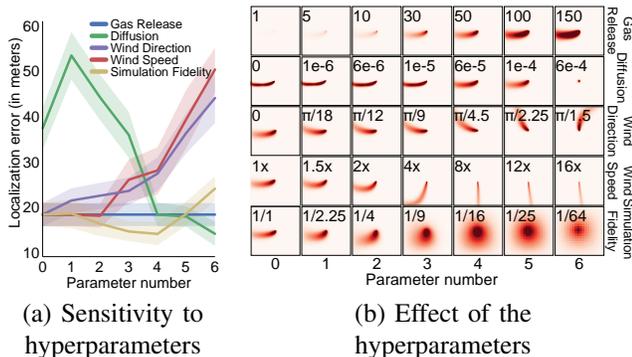


Fig. 5: **Sensitivity and effect of the hyperparameters** (a) **Sensitivity** Mean and variance localization error for all offline experiments for different parameters’ variations. We start with the default parameters (gas release:50, diffusion:0.0001, wind speed:1x, wind direction:0, simulation fidelity:1 cell/1m²) and individually vary each of them. (b) **Effect** Visualization of different values of the hyperparameters in a 160×160 cell simulation. We use the wind flow from one of the experiments from Sec. IV-A to drive the simulation, hence some of the observed dynamics.

B. Online algorithm

Noisy simulation We use sparse noisy readings taken every 20 seconds from a simulator (80 × 80 cells) with an arbitrary location for the gas source. We assume that the hyperparameters of the simulation are known and that wind flow is spatially constant, and we add up to 10% multiplicative noise to the simulated readings. We experiment with multiple start locations, both on and off the path of the gas as shown in Fig. 6. We use the following parameters - gas release:20, diffusion:6e − 4, wind speed:25 (simulator metric), wind direction:π ± π/2 (primarily coming from the right, but uniformly changing its direction every 30 seconds).

Comparison with related work Using the noisy simulator, we also evaluate the convergence rate of different algorithms as shown in Fig.7. We perform three sets of experiments - no wind in the simulation, constant wind and variable wind. An acquisition function is used for each of the algorithms to select each consecutive measurement point. For GP and DM+V/W we use Lower Confidence Bound with alpha parameter 3 and for our approach we use the suggested likely region. In addition to the baseline regression approaches, we also generate example trajectories for different wind conditions using Infotaxis with parametric plume model and parameters as in [34] as shown in Fig. 8.

Real experiments We perform three real-world active sensing experiments with the setup described in Sec. IV-A (Supplementary material). The domain is 40m × 40m, and we use 1 grid cell per 4 m² in simulation. We perform optimization using OGS, for which the simulation and

comparison calculations are performed within an on-board processor. After each taken reading (gas and wind), the UAV updates its beliefs about the likelihood of the source, saves the current state of the optimization and flies to the suggested by the optimization most likely source location l^* - repeating until the optimization procedure converges to the same location. In our experiments, we had strong winds from the Southwesterly direction in the first two experiments and weak varying winds in the third experiment. The estimated likelihood and waypoints discovered, are shown in Fig. 9.

V. INTERPRETATION OF RESULTS AND DISCUSSION

More accurate than state-of-the-art In practice, the alternate methods based on Gaussian process regression and the DM+V/W algorithm provide similar results. DM+V/W performs somewhat better, as it includes reshaping of the smoothing kernel based on wind information. As we show in Fig. 3, these methods often predict the highest mean to be away from the true location, caused by the wind moving the released gas. Although DM+V/W indirectly accounts for those effects via wind kernel reshaping and time scaling, convergence is slower. By using wind measurements in simulation, we are able to capture these dynamics, allowing our proposed approach to achieve lower localization error (< 20m) than GP (45m) and DM+V/W (< 40m). Likewise, our approach achieved better performance than with standard Bayesian optimization approaches to compute the minimum (35m). Using fluid simulation yields a higher computational cost compared to traditional approaches. However, we show that our algorithm is fast (up to 20× faster than BO) and can be implemented within a commercial on-board computer.

Fidelity and sensitivity Our algorithm is robust to inaccurate parameters as seen in Fig. 5. The extent of released gas did not have any impact on localization, as both simulated and real-world readings are normalized. However, it is sensitive to the diffusion coefficient (green curve) used in simulation, as well as to errors in wind estimation. We found that we can tolerate 2× inaccuracy in wind speed and an offset of π/12 in direction from the calibrated mappings. We use a standard approach [7], [43] to estimate wind, which lies within the required tolerances. We observed, curiously, that a grid resolution of 1 cell per 16m² yields best results.

Active sensing Our greedy (exploitation) strategy consistently locates the source using very few (usually less than 10) samples, to similar accuracy as the offline algorithm. As with approaches that utilize gradient-search, we are able to follow the path of wind and gas readings. In contrast to those methods, we can also locate the source if we are not on the path of the gas based on expected concentration readings for different source locations given wind conditions. Importantly, we find that variable wind tends to be informative leading to faster convergence, unlike competing methods, where it’s often a limiting factor. As we normalize the readings, we indirectly account for different rates of gas release. Thus even when is on the path of the gas flow, the algorithm explores the neighborhoods of the path. Finally, at each step, we also update uncertainties over possible locations.

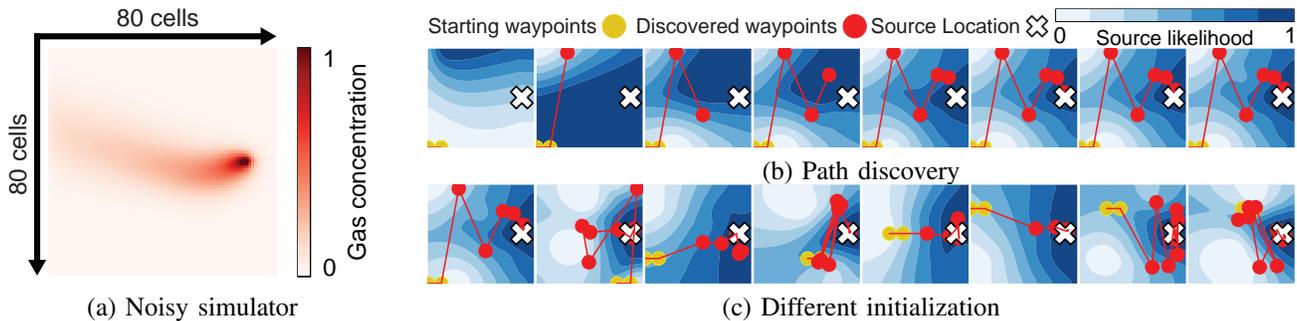


Fig. 6: **Active sensing using synthetic data** (a) **Noisy simulator** We use noisy readings from our model to evaluate the proposed approach – the source is placed in the middle-right, with the wind blowing East-West. (b) **Path discovery** Successive locations suggested during the discovery of one path are shown with red circles along with initialization locations (yellow circles). (c) **Different initialization** The convergence of the algorithm is illustrated for eight different starting locations.

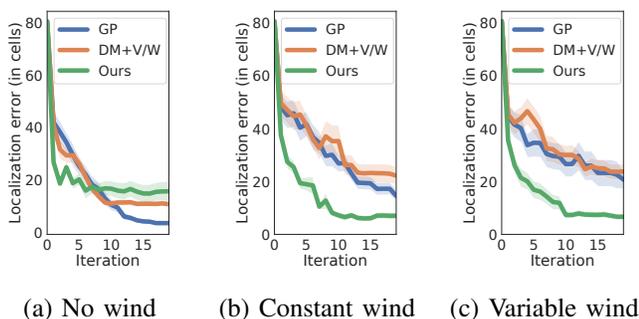


Fig. 7: **Convergence of active sensing** Convergence rate of different algorithms for different wind conditions - no wind (a), constant wind (b) and variable wind (c). In the experiments a noisy simulator was used.

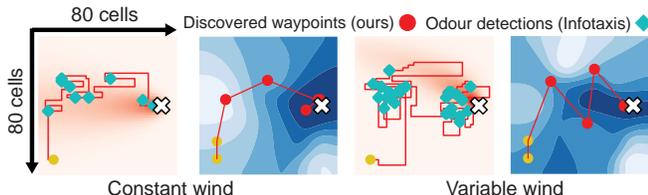


Fig. 8: **Trajectories** Generated trajectories for the here proposed approach and 'Infotaxis' with parametric plume model. Odour detections with gas flow (for Infotaxis) and discovered waypoints with source likelihood map (for ours) are shown.

Source of error, limitations and future work Although our method improves on state of the art in gas localization, the error is still about 15m. We identify two potential sources for this residual error. Firstly, our model is based on 2D space S while gas diffusion in the real world is in 3 dimensions. Thus, even though our algorithm finds the highest density of the gas on the plane of flight, it may not be exactly above the nozzle releasing the gas. A second source of error is that we model the wind as spatially invariant, which is a coarse approximation for increasingly larger domains. It

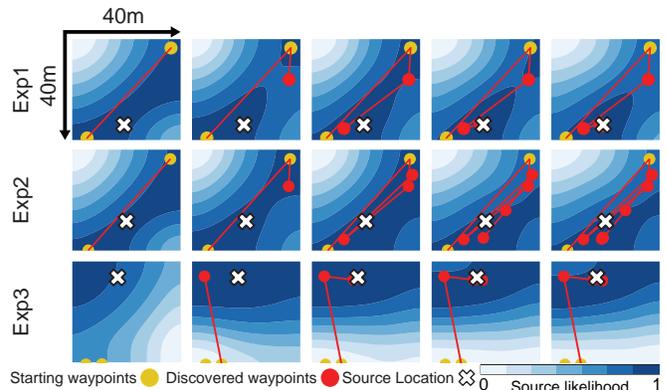


Fig. 9: **Active sensing by UAV** The waypoints visited by the UAV. The fluid simulator and OGS are executed at each step, as described in the text. We initialize the algorithm with two predefined waypoints (yellow circles), after which the UAV automatically seeks the source using our online algorithm for active sensing. Our algorithm converges in 5, 7 and 4 iterations, in those examples.

would be beneficial to extend this work by using simulations supporting obstacles and multiple sources [18]. This will require development of new methods able to encapsulate the variability of the more complex simulations from limited samples, while preserving fast inference at runtime.

VI. CONCLUSION

We formulate gas-leak localization as an optimization problem, minimizing the discrepancy between simulated gas flow and point-wise measurements of leaking gas. We propose a practical optimization algorithm that can be used offline as well as online for active sensing. We evaluate our algorithm by implementing it on a UAV equipped with sensors to detect CO_2 . Our algorithm is able to cope with dynamically varying wind, and efficiently localizes the source of leaks even if the UAV is not started on the path of the gas. We show through experiments that our proposed approach outperforms baselines from the literature in its ability to minimize localization error.

REFERENCES

- [1] R. R. Murphy, S. Tadokoro, D. Nardi, A. Jacoff, P. Fiorini, H. Choset, and A. M. Erkmen, "Search and rescue robotics," in *Springer Handbook of Robotics*. Springer, 2008, pp. 1151–1173.
- [2] D. Pieri and J. A. Diaz, "In situ sampling of volcanic emissions with a uav sensorweb: Progress and plans," in *Dynamic Data-Driven Environmental Systems Science*. Springer, 2015, pp. 16–27.
- [3] E. L. Glaeser and M. E. Kahn, "The greenness of cities: carbon dioxide emissions and urban development," *Journal of urban economics*, vol. 67, no. 3, pp. 404–418, 2010.
- [4] T. K. Dey, *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis (Cambridge Monographs on Applied and Computational Mathematics)*. New York, NY, USA: Cambridge University Press, 2006.
- [5] P. B. Jónsson, J. Wang, and J. Kim, "Scalar field reconstruction based on the gaussian process and adaptive sampling," in *Ubiquitous Robots and Ambient Intelligence (URAI), 2017 14th International Conference on*. IEEE, 2017, pp. 442–445.
- [6] C. Stachniss, C. Plagemann, A. J. Lilienthal, and W. Burgard, "Gas distribution modeling using sparse gaussian process mixture models," in *Robotics: Science and Systems*, vol. 3, 2008.
- [7] P. P. Neumann, S. Asadi, A. J. Lilienthal, M. Bartholmai, and J. H. Schiller, "Autonomous gas-sensitive microdrone: Wind vector estimation and gas distribution mapping," *IEEE robotics & automation magazine*, vol. 19, no. 1, pp. 50–61, 2012.
- [8] M. Reggente and A. J. Lilienthal, "Using local wind information for gas distribution mapping in outdoor environments with a mobile robot," in *Sensors, 2009 IEEE*. IEEE, 2009, pp. 1715–1720.
- [9] A. J. Lilienthal, M. Reggente, M. Trincavelli, J. L. Blanco, and J. Gonzalez, "A statistical approach to gas distribution modelling with mobile robots-the kernel dm+ v algorithm," in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*. IEEE, 2009, pp. 570–576.
- [10] M. Hutchinson, H. Oh, and W.-H. Chen, "A review of source term estimation methods for atmospheric dispersion events using static or mobile sensors," *Information Fusion*, vol. 36, pp. 130–148, 2017.
- [11] J. R. Demers, F. Garet, and J.-L. Coutaz, "Atmospheric water vapor absorption recorded ten meters above the ground with a drone mounted frequency domain thz spectrometer," *IEEE sensors letters*, vol. 1, no. 3, pp. 1–3, 2017.
- [12] C. M. Bishop, "Model-based machine learning," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 371, no. 1984, 2013.
- [13] M. Ghallab, D. Nau, and P. Traverso, *Automated Planning and Acting*. Cambridge University Press, 2016.
- [14] A. Bordallo, F. Previtali, N. Nardelli, and S. Ramamoorthy, "Counterfactual reasoning about intent for interactive navigation in dynamic environments," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 2943–2950.
- [15] J. Stam, *The Art of Fluid Animation*. CRC Press, 2015.
- [16] M. Macklin and M. Müller, "Position based fluids," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 4, p. 104, 2013.
- [17] M. Macklin, M. Müller, N. Chentanez, and T.-Y. Kim, "Unified particle physics for real-time applications," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, p. 153, 2014.
- [18] J. Monroy, V. Hernandez-Bennetts, H. Fan, A. Lilienthal, and J. Gonzalez-Jimenez, "Gaden: A 3d gas dispersion simulator for mobile robot olfaction in realistic environments," *Sensors*, vol. 17, no. 7, p. 1479, 2017.
- [19] T. Lopez-Guevara, N. K. Taylor, M. U. Gutmann, S. Ramamoorthy, and K. Subr, "Adaptable pouring: Teaching robots not to spill using fast but approximate fluid simulation," in *Conference on Robot Learning*, 2017, pp. 77–86.
- [20] C. Bates, P. Battaglia, I. Yildirim, and J. B. Tenenbaum, "Humans predict liquid dynamics using probabilistic simulation," in *CogSci*, 2015.
- [21] P. W. Battaglia, J. B. Hamrick, and J. B. Tenenbaum, "Simulation as an engine of physical scene understanding," *Proceedings of the National Academy of Sciences*, vol. 110, no. 45, pp. 18 327–18 332, 2013.
- [22] M. B. Chang, T. Ullman, A. Torralba, and J. B. Tenenbaum, "A compositional object-based approach to learning physical dynamics," *arXiv preprint arXiv:1612.00341*, 2016.
- [23] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine, "Learning to poke by poking: Experiential learning of intuitive physics," in *Advances in Neural Information Processing Systems*, 2016, pp. 5074–5082.
- [24] K. Fragkiadaki, P. Agrawal, S. Levine, and J. Malik, "Learning visual predictive models of physics for playing billiards," *arXiv preprint arXiv:1511.07404*, 2015.
- [25] J. Wu, I. Yildirim, J. J. Lim, B. Freeman, and J. Tenenbaum, "Galileo: Perceiving physical object properties by integrating a physics engine with deep learning," in *Advances in neural information processing systems*, 2015, pp. 127–135.
- [26] G. Kowadlo and R. A. Russell, "Naïve physics for effective odour localisation," in *Proceedings of the Australian Conference on Robotics and Automation*, 2003.
- [27] C. Sanchez-Garrido, J. Monroy, A. J. Gonzalez-Jimenez, *et al.*, "Probabilistic localization of gas emission areas with a mobile robot in indoor environments," 2018.
- [28] A. Lilienthal, F. Streichert, and A. Zell, "Model-based shape analysis of gas concentration gridmaps for improved gas source localisation," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*. IEEE, 2005, pp. 3564–3569.
- [29] J. Stam, "Real-time fluid dynamics for games," in *Proceedings of the game developer conference*, vol. 18, 2003, p. 25.
- [30] I. G. Damousis, M. C. Alexiadis, J. B. Theocharis, and P. S. Dokopoulos, "A fuzzy model for wind speed prediction and power generation in wind parks using spatial correlation," *IEEE Transactions on Energy Conversion*, vol. 19, no. 2, pp. 352–361, 2004.
- [31] R. Sykes, D. Henn, S. Parker, and R. Gabruk, "Scipuff-a generalized hazard dispersion model," American Meteorological Society, Boston, MA (United States), Tech. Rep., 1996.
- [32] P. P. Neumann, V. Hernandez Bennetts, A. J. Lilienthal, M. Bartholmai, and J. H. Schiller, "Gas source localization with a micro-drone using bio-inspired and particle filter-based algorithms," *Advanced Robotics*, vol. 27, no. 9, pp. 725–738, 2013.
- [33] J.-G. Li, Q.-H. Meng, Y. Wang, and M. Zeng, "Odor source localization using a mobile robot in outdoor airflow environments with a particle filter algorithm," *Autonomous Robots*, vol. 30, no. 3, pp. 281–292, 2011.
- [34] M. Vergassola, E. Villermaux, and B. I. Shraiman, "infotaxis as a strategy for searching without gradients," *Nature*, vol. 445, no. 7126, p. 406, 2007.
- [35] J. G. Monroy, J.-L. Blanco, and J. Gonzalez-Jimenez, "Time-variant gas distribution mapping with obstacle information," *Autonomous Robots*, vol. 40, no. 1, pp. 1–16, 2016.
- [36] J. L. Blanco, J. G. Monroy, A. Lilienthal, and J. Gonzalez-Jimenez, "A kalman filter based approach to probabilistic gas distribution mapping," in *Proceedings of the 28th Annual ACM Symposium on Applied Computing*. ACM, 2013, pp. 217–222.
- [37] S. Asadi, S. Pashami, A. Loutfi, and A. J. Lilienthal, "Td kernel dm+ v: Time-dependent statistical gas distribution modelling on simulated measurements," in *AIP Conference Proceedings*, vol. 1362, no. 1. AIP, 2011, pp. 281–282.
- [38] J. Le Ny and G. J. Pappas, "On trajectory optimization for active sensing in gaussian process models," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*. IEEE, 2009, pp. 6286–6292.
- [39] R. Ouyang, K. H. Low, J. Chen, and P. Jaillet, "Multi-robot active sensing of non-stationary gaussian process-based environmental phenomena," in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2014, pp. 573–580.
- [40] X. Huan and Y. M. Marzouk, "Sequential bayesian optimal experimental design via approximate dynamic programming," *arXiv preprint arXiv:1604.08320*, 2016.
- [41] M. Hutchinson, C. Liu, and W.-H. Chen, "Information-based search for an atmospheric release using a mobile robot: Algorithm and experiments," *IEEE Transactions on Control Systems Technology*, no. 99, pp. 1–15, 2018.
- [42] —, "Source term estimation of a hazardous airborne release using an unmanned aerial vehicle," *Journal of Field Robotics*, 2018.
- [43] P. P. Neumann and M. Bartholmai, "Real-time wind estimation on a micro unmanned aerial vehicle using its inertial measurement unit," *Sensors and Actuators A: Physical*, vol. 235, pp. 300–310, 2015.