Informative Path Planning for Gas Distribution Mapping in Cluttered Environments*

Callum Rhodes, Cunjia Liu and Wen-Hua Chen

Abstract-Mobile robotic gas distribution mapping (GDM) is a useful tool for hazardous scene assessment where a quick and accurate representation of gas concentration levels is required throughout a staging area. However, research in robotic path planning for GDM has primarily focused on mapping in open spaces or estimating the source term in dispersion models. Whilst this may be appropriate for environment monitoring in general, the vast majority of GDM applications involve obstacles, and path planning for autonomous robots must account for this. This paper aims to tackle this challenge by integrating a GDM function with an informative path planning framework. Several GDM methods are explored for their suitability in cluttered environments and the GMRF method is chosen due to its ability to account for obstacle interactions within the plume. Based on the outputs of the GMRF, several reward functions are proposed for the informative path planner. These functions are compared to a lawnmower sweep in a high fidelity simulation, where the RMSE of the modelled gas distribution is recorded over time. It is found that informing the robot with uncertainty, normalised concentration and time cost, significantly reduces the time required for a single robot to achieve an accurate map in a large-scale, urban environment. In the context of a hazardous gas release scenario, this time reduction could save lives as well as further gas ingress.

I. INTRODUCTION

In response to Chemical, Biological, Radiological and Nuclear (CBRN) incidents, the quick acquisition of accurate information on the state of the environment can save lives. In this situational awareness problem, key media such as concentration maps and contextual information of contaminations would provide first responders with a greater understanding of the environment they are operating in. Gathering this information by manual sampling using a handheld device is risky and time consuming, therefore, a remote and automatic solution is the obvious step in mitigating the danger to human operators and reducing the time required to collect the data. One promising solution to this problem is to use autonomous vehicles such as Unmanned Aerial Vehicles (UAVs) or Unmanned Ground Vehicles (UGVs). Onboard these vehicles, there must be suitable algorithms that are capable of using an entire suite of perceptual sensors in order to enable the vehicle to operate autonomously. In this paper, both the Gas Distribution Mapping (GDM) method and the informative path planner are investigated as to achieve the goal of efficient and autonomous GDM within a cluttered environment.

For GDM within cluttered environments, to the best of the authors knowledge there is currently little research for Informative Path Planning (IPP). In this paper we propose a solution to this gap in the literature that convolves around a reward function informed by the outputs a GDM algorithm to deliver efficient and accurate GDM using an in-situ gas sensing mobile robot.

A. Gas Distribution Mapping

GDM attempts to answer the question of How is the gas distributed across an area? Whilst in certain scenarios it may only be of interest to either determine if a gas is present, or given the fact gas is present, where the source of the release is, the distribution across the whole environment can be useful for first responders as well as for path planning [1]. A combination of small scale fluctuations in the dispersal pattern and sparse measurements associated with mobile sensing makes gas mapping an exceptionally difficult task compared to radiation mapping which leverages dense continual data. Because of the sparse and transient nature of gas concentration readings, extra care must be taken when choosing a mapping model. For example, simple linear interpolation between readings is not suitable. There are three popular methods currently used for GDM (mobile robot specific applications) which are Kernel DM (and its variants) [2]-[5], Gaussian Markov Random Field (GMRF) [6], [7] and Gaussian Process (GP) [8], [9]. Of these methods, Kernel has the advantage of being online and is computationally lightweight but its simplistic nature shows a lack of accuracy in cluttered environments due to its inability to account for obstacle interactions. GP can be difficult to successfully implement due to its need to be trained appropriately with adequate data. Moreover, current formulations of GP do not encode obstacle information into the model. Although Jadidi et al. [10] showed that obstacle information can be encoded into GP for Simultaneous Localisation and Mapping (SLAM), it is an open question whether the same can be applied to GDM. GMRF is the most recent attempt at GDM and (whilst also being online and giving model uncertainties) the major advantage compared to others is the ability to model obstacle interactions with the plume and also model how obstacles affect the uncertainty of the models prediction. This modelling capability allows GMRF to be more suitable for cluttered environments, which realistically make up the vast majority of CBRN scenarios that require the use of a mobile robot. Based on these insights, this paper advocates GMRF as the GDM method for a Mobile Robot Olfaction (MRO) application and therefore it is used as the basis in

^{*}This work was supported in part by EPSRC DTP under the project No. 2126619

¹The authors are with the Department of Aeronautical and Automotive Engineering, Loughborough University, LE11 3TU, UK. {C.Rhodes, C.Liu5, W.Chen}@lboro.ac.uk

developing IPP for robotic GDM. It should be noted that GMRF computationally scales poorly when applied to very large maps, which needs to be further investigated.

B. Robotic Path Planning for Autonomous GDM

Path planning for GDM is of crucial importance when considering the scenarios it is likely be applied to. Fast deployment and acquisition of gas distribution characteristics can mean the difference between safe containment and potential loss of life. Whilst path planning within mobile robotics has been extensively researched for a wide variety of generic robotics tasks, very little research has been performed for the specificity of the GDM domain.

Most research for GDM has employed some form of sweep pattern wherein the robot navigates along a predetermined trajectory with the majority of efforts focusing on path planning to resolve the single point source [2], [11], [12], rather than optimising for full GDM. Environmental monitoring tasks that do focus on the mapping process treat path planning as a coverage maximisation problem, such as those surveyed by Galceran et al. in [13]. Whilst this may be suitable for environmental monitoring tasks such as aerial surveying [14] and aquatic monitoring [15], GDM robots often operate in obstacle rich environments that interact with the distribution being modelled. This requires further consideration as being explored in this paper.

II. GAUSSIAN MARKOV RANDOM FIELD

To spatially model the gas distribution across a map, the GMRF method is employed. Markov random fields are a nodal representation of an environment and therefore, a given map is discretised into interconnected cells, with each cell being represented by a node in the Markov network. In Markov random fields, the construction of the nodal network (see Fig. 1) in conjunction with measurement data can be entirely defined in a Jacobian matrix J. An example can be found in Eq (1).

$$J = \begin{bmatrix} m_1 & m_2 & \cdots & m_i & \cdots & m_j & \cdots & m_N \\ 1 & 1 & -1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ 2 & 1 & 0 & \cdots & -1 & \cdots & 0 & \cdots & 0 \\ \vdots & & \vdots & & & & \\ 0 & 0 & \cdots & 1 & \cdots & -1 & \cdots & 0 \\ 0 & 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & & & \vdots & & & \\ 1 & 0 & 0 & \cdots & 1 & \cdots & 0 & \cdots & 0 \end{bmatrix}$$
(1)

where m_j , j = 1, ..., N, represents all map nodes in the field, and all nodal connections (1 : L) are pairwise defined with 1 and -1 (this is predefined and does not change assuming the map formulation does not change). Observations are also accounted for in the Jacobian, with the location of each observation z_k , k = 1, ..., M, being iteratively added to the Jacobian with a value of 1 at the corresponding cell, m_i , i = 1, ..., N.

The construction of the Markov random field has the benefit of the nodal map being entirely flexible to the specific

scenario. Monroy et. al [6] propose a formulation for GDM wherein map nodes are connected to other nodes that are spatially orthogonal, i.e., a node in free space is connected to four cells above, below, left and right. To account for obstacles, if the connection between two nodes is blocked by an obstacle, then the connection is deleted leaving the node with three (or fewer) connections, as shown in Fig. 1.



Fig. 1. An example of a GDM GMRF formulation showing the concentration cells, m, the gas observations, z, the prior factors, F_p , the observation factors, F_o , and a potential obstacle boundary. The diagram shows how the obstacle boundary isolates the cell $m_{x,y-2}$, to the current cell being sampled, $m_{x,y}$.

Once the construction of the network is completed, the maximum likelihood estimation for the distribution can be attained. However, for the sake of simplicity, only key calculation steps are included in this paper. For detailed derivation of the GDM specific formulation of the GMRF, please refer to [6].

For GDM, the solution to the modelling problem is to iteratively update prior concentration values of each node, given new observations. Therefore, at each calculation of the model, a change in the concentration value of each cell, denoted by the matrix Δm_i , is calculated to update the GMRF, such that

$$m_i = m_{prior_i} + \Delta m_i \tag{2}$$

where m_i is the posterior mean concentration of the cell and m_{prior_i} is the prior mean concentration of a cell.

To solve the change in mean concentration across the field given an observation, Eq (3) is solved for Δm using the Cholesky decomposition of the Hessian (formed from the Jacobian) $H = R^{\top}R$.

$$H\Delta m = -g$$

($J^{\top}\Lambda J$) $\Delta m = -J^{\top}\Lambda(f(m) - y)$ (3)

$$\Delta m = R/(R^{\top}/-g) \tag{4}$$

where H refers to the $N \times N$ Hessian, Λ refers to the $(L + M) \times (L+M)$ information matrix consisting of diagonal prior and observation factor values (i.e., F_p and F_o), and g is the gradient vector. The gradient vector accounts for the current concentration value, m_i , of cell i, and calculates a gradient based on the values of surrounding cells (weighted by the prior factor, F_p) and observations at the node (weighted by the observation factor, F_o) shown in Eq (5).

$$g_i = \sum_k F_{o_k}(m_i - z_k) + \sum_j F_{p_j}(m_i - m_j)$$
 (5)

where z is the value for all observations, k, associated with the node i. The gradient on priors is summed over all neighbour cells, j, of the node i.

Moreover, to calculate modelling uncertainty at a node, σ_i , the covariance matrix is recovered by performing the inverse of the Hessian and taking the diagonal values, Eq (6).

$$\sigma_i^2 = H_{i,i}^{-1}$$
 (6)

Since the Hessian diagonal values are proportional to the number of connections and observations (depending on the prior/observation information matrix) then cells that have few connections exhibit a higher uncertainty. This has interesting implications for uncertainty driven path planning since cells that are in cluttered regions have a high uncertainty compared to cells in more open space.

It should be noted that the original formulation by Monroy et al. included a reductive term on the observations based on the time that an observation is taken. Due to a quasisteady state simulation scenario, this additional factor is not required but would be necessary for time varying mapping applications.

III. IPP REWARD FUNCTIONS

For IPP, a reward function needs to be defined based on some measure of information to find the best goal location for the robot to travel to. The reward function used in a path planner should consider all metrics that are available to it whilst maximising the information gain [16]. The output of the GMRF provides a predicted mean concentration map and the associated uncertainty of all cells in the map (the information to be maximised). Since the connection relationship between nodes are defined in the Jacobian, a distance metric, d, between a cell and all other cells can easily be established using a simple search method such as A* search along these connections (other obstacle avoidance algorithms could surrogate but the focus of this paper is on goal selection). As shown in the comprehensive gas sensor technology survey by Liu et al. [17], gas sensors are subject to a response time that varies depending on the type of sensor deployed. This response time often dictates that the robot has to remain stationary in order to get a representative concentration value and provide accurate mapping [9] (however more advanced sensors such as photo-ionisation detectors, can allow the robot to collect readings in transit). Based on this inherent issue with gas sensing technology, the sampling time must also be reflected in the cost function. The sampling time cost, C, is shown in Eq (7), and refers to a simple case UGV (more complex dynamic models could be included at this stage). Since point sampling is assumed, full trajectory information reward functions commonly used in environmental monitoring applications [18], cannot be used

as the concentration data along a trajectory does not well represent the actual concentration.

$$C_i = d_i / v_r + t_{sample} \tag{7}$$

where d_i refers to the distance from the robots current location to a cell, *i*. This is then converted to a traversal time using an assumed constant robot velocity, v_r , in order to incorporate the time needed to sample, t_{sample} .

The two GMRF outputs and the A* search, gives three possible variables (mean, uncertainty and traversal cost) to be considered when selecting candidate locations for sampling local concentrations.

As explained previously, GMRF requires a relatively coarse resolution to be computationally tractable. This vastly reduces the number of candidates for traversal to being, at a maximum, the number of nodes in the Markov field. It is also likely this number is reduced as nodes that are within obstacles do not need to be considered.

In [19], similar variables are applied to a global artificial potential field and then maximised. The reward functions proposed are constructed in a similar manner but in the following, each of the variables are also tested in isolation to examine their individual contribution to final proposed function. Five proposed reward functions are tested, which contain some permutations of the cost and information criteria as follows:

• Maximising uncertainty reduction - The candidate with the highest variance, σ^2 , is chosen

$$r = \operatorname*{argmax}_{i \in N} \left(\sigma_i^2 \right) \tag{8}$$

where N refers to the number of nodes within the GMRF

• Time weighted uncertainty reduction - The candidate with the highest variance weighted with the cost in traversing to and sampling the location is selected

$$r_i = \operatorname*{argmax}_{i \in N} \left(\sigma_i^2 / C_i \right) \tag{9}$$

• Uncertainty weighted concentration - The candidate with the highest normalised concentration, V_i , weighted by how certain the model is at said location. Normalised concentration is defined as the ratio between a cells concentration and the average map concentration (which will usually converge to the background concentration). This is employed so that scaling based on the type of gas measured is not required.

$$V_i = m_i / \left(\frac{\sum_{j=1}^N m_j}{N}\right) \tag{10}$$

$$r = \operatorname*{argmax}_{i \in N} \left(V_i \times \sigma_i^2 \right) \tag{11}$$

where m_i refers to the mean concentration value of a cell as modelled by the GMRF

• Time weighted concentration - The candidate with the highest normalised concentration weighted with the cost of traversal. It is also employed that no location shall be visited twice, to avoid local minima.

$$r = \operatorname*{argmax}_{i \in N} \left(V_i / C_i \right) \tag{12}$$

• Time weighted joint uncertainty and concentration -The candidate with the highest normalised concentration weighted by how certain the model is at said location, and also taking into account the cost of traversal to the candidate.

$$r = \operatorname*{argmax}_{i \in N} \left((V_i \times \sigma_i^2) / C_i \right)$$
(13)

The above functions are tested against a manually selected lawnmower sweep that runs for an equivalent length of time (Figure 2). It should be noted that a myopic approach has been implemented due to the computational limitations of using GMRF in the following large scale scenario. For small scale scenarios, an approach which uses the model to predict several steps into the future should be used.

IV. SIMULATION EXPERIMENTS

To test the performance of the proposed reward functions, a large scale, outdoor, cluttered scenario is used so that the obstacle inclusion of GMRF can be fully exploited. Ground truth concentration data is obtained from the DAPPLE dispersion experiment [20] and consists of a field test validated Computational Fluid Dynamics (CFD) simulation of a source release under steady wind conditions. As stated in [21], one of the greatest challenges in odour sensitive robotics is the availability of ground truth datasets, so this dataset provides a good platform to base validated conclusions from. A temporal snapshot of the CFD simulation is shown in Figure 2, along with the predefined lawnmower sweep (not accounting for obstacle avoidance).

The GMRF is set up with default parameters of $F_o = 10$ and $F_p = 0.5$ and initialised with a cell resolution of 13m x 13m. This resolution is very large but the CFD simulation is of a large scale chemical release which shows ingress over the environment of several hundreds of metres, therefore, this resolution gives a tractable solution and suitable resolution to model the plume. This scale leads to an average GMRF update time of 2.5s and a cell count of N = 5038, with a matrix dimension of 68 x 74 cells. Since the source release power is constant throughout the simulation (fluctuates temporally but with long term plume stability as shown in Figure 4) and the fact that the simulation starts at full gas ingress, the time decay functionality on observations is not required. GMRF parameters are set based on a brief accuracy study of various parameter values not shown here.

In the experiments, 10 starting locations are randomly chosen that span areas both in and out of the plume, as well as starting from inside and outside the cluttered regions. This is to ensure that the robustness of the proposed algorithms can be demonstrated.



Fig. 2. Temporal snapshot of the CFD simulation concentration data along with the predefined lawnmower sweep. Sweep trajectory shown is symbolic of the lawnmower pattern, as the actual trajectory accounts for obstacle avoidance so that the correct time cost is applied to the simulation. Each of the 10 starting locations for IPP are shown with green circles.

Alg	orithm 1 IPP using GMRF				
1:	procedure INITIALISATION				
2:	Create $GMRF() \leftarrow occmap$				
3:	$Obs \leftarrow [$]				
4:	$m_{prior} \leftarrow 0$				
5:	$time \leftarrow 1$				
6:	set $CFD(time)$				
7:	$pos \leftarrow [-450, -380]$				
8: procedure MAIN					
9:	while $time \leq 10000s$				
10:	$Obs \leftarrow [Obs, CFD(pos)]$				
11:	$[m, \sigma] \leftarrow GMRF(pos, Obs, m_{prior})$				
12:	$m_{prior} \leftarrow m$				
13:	$C \leftarrow d/vel + t_{sample}$				
14:	$r \leftarrow RewardFunction(m, \sigma, C)$				
15:	$pos \leftarrow r$				
16:	$time \leftarrow time + C(r)$				
17:	set $CFD(time)$				

The simulation process for collecting the data (outlined in Algorithm 1) is as follows. The GMRF Jacobian formulation is created based on a predetermined occupancy map (*occmap*), and an empty observation vector, *Obs*, initialised. The CFD simulation is then started from t = 1 and the initial position of the robot, *pos*, is set. After sampling the concentration of a location and inserting the reading into the GMRF algorithm at the corresponding node, a 'next best view' location, *r*, is assessed from the proposed reward functions. The robot then sets this as its destination and the CFD simulation is stepped forward in time, according to the estimated time it should take the robot to reach its destination and take a sample (which is set to 10s). The robot then takes a sample at the proposed location by averaging noise added CFD values that are within a 5m radius of the robot (to



Fig. 3. The RMSE of the modelled plume against the ground truth for the starting location [-450, -380], plotted over time for each strategy. The time the source location is sampled is marked with a star for added context.

simulate the transient nature of sampling over the 10s sample time). This new reading is then inserted into the GMRF and the updated maps are used to inform the next stage of the path planning process. This process is repeated until the end of the CFD data available, at t = 10000s. A video example of a simulation run is available at *youtu.be/GWiZQaayapQ*.

V. RESULTS

In order to assess the performance of each function, the Root Mean Square Error (RMSE) of the modelled plume to that of the time averaged CFD data (ground truth) are recorded at each time interval. This allows the evaluation of not only the final level of accuracy of each strategy, but also how quickly it took each algorithm to suitably resolve the plume. RMSE is chosen as it puts a higher error penalty on data that are far away from the ground truth (i.e. poorly resolving the distribution close to the source where a concentration spike is seen). From a practical point of view, inaccurately modelling high concentration areas (false negatives) could be highly dangerous to a first responder.

$$RMSE = \sqrt{\frac{\sum_{i=1}^{i \in N} (m_{c_i} - m_i)^2}{N}}$$
(14)

where m_c is the ground truth mean concentration attained from the CFD simulation.

The time to resolve, T_r score (Table I) refers to how many times faster, on average across all experiments, the algorithm resolves the plume compared to the resolve time of the sweep. An algorithm is deemed to adequately resolve the plume when the RMSE drops below 0.0095 (this has been determined as the appropriate limit of accuracy of the GMRF to model the plume from extended testing). This metric is key for evaluating how quickly each algorithm attains an accurate map. A N/A score is shown if an algorithm did not adequately resolve the plume within the simulation time.

The lawnmower sweep is the simplest method and covers the entire map ensuring full coverage. Reviewing an example

Strategy	Final RMSE x10 ⁻³	Min RMSE $x10^{-3}$, (s)	T_r score	No. samples		
Sweep	9.5	9.2	1	110		
σ^2	11.3	10.8	N/A	55		
σ^2/C	9.5	9.3	0.56	195		
$V \times \sigma^2$	9.4	8.8	1.06	108		
V/C	9.2	8.9	1.15	211		
$(\dot{V} \times \sigma^2)/C$	8.7	8.5	2.05	209		
TABLE I						

SIMULATION RESULTS FOR THE LAWNMOWER SWEEP AND 5 PROPOSED ALGORITHMS AVERAGED ACROSS ALL STARTING LOCATIONS.

run, Fig. 3, initial behaviour is similar to the proposed methods but then plateaus at 1000s-2000s. This is caused by the robot sweeping into an area that is not in the plume, despite the fact that parts of the plume have already started to be resolved on the first pass. This is the main drawback of an uninformed sweep. The sweep then revisits the plume on a second pass and eventually samples the source locations at t = 3024s. To ensure full coverage, the sweep has to traverse a greater distance between samples leading to a relatively low average number of samples taken over the simulation (110). The average minimum RMSE of the map is comparable to the other methods but significantly worse due to the lack of sampling resolution within the plume.

Navigating based greedily on maximum uncertainty is the worst strategy. Large traversal times are ignored and therefore the average number of samples of this strategy is the lowest (55). Since the mean concentration map is not accounted for, the plume is poorly resolved. This proves the need for extra considerations when performing uncertainty driven path planning for GDM.

Adding the traversal cost to the maximum uncertainty criterion allows a 4 times increase in the number of samples. This significantly increases RMSE reduction as the algorithm progressively investigates the cluttered region. However, efforts are still spent resolving areas that are modelled to be of very low concentration, which leads to long periods of time where the plume is not being resolved.

Trying to maximise all information (mean and uncertainty) without accounting for cost proves to be a useful (if naive) solution. The robot successfully locates the plume, however less efficiently due to the lack of constraint on traversal time. When the plume is located, t = 2507s, the algorithm resolves the plume efficiently since the V term, keeps the robot softly constrained to the plume.

To prove the need for uncertainty driven path planning, planning based solely on concentration and traversal cost is tested. Due to to the initial no concentration state, the algorithm defaults to an automated sweep due to not being able to sample the same location twice. This initially proves worse than the manually selected sweep. Once the robot detects the edge of the plume, the RMSE significantly drops but the lack of exploratory term dictates that the robot tediously resolves its current interpretation of the plume location (also shown with the highest total number of samples at 211).



Fig. 4. Mean concentration maps for the best performing strategy (the proposed algorithm) against the standard sweep at 1000s, 2000s, 4000s and 8000s. The mean concentration data from the CFD simulation (the ground truth) is shown for comparison.

This is shown with no RMSE drop from t = 2000s - 8500s, after which the source location is sampled and the RMSE begins to drop again. At this point, the plume is well resolved due to its attraction to the high concentration area, but hits the ground truth data limit of the simulation at 10000s. With starting locations that are initially near the source, this method resolves very quickly (hence a reasonable time score), but Fig. 3 shows how this method is not robust far away from the plume.

The final and most informed algorithm considers all the factors outlined. It shows the quickest initial reduction in RMSE and in 9 out of 10 starting locations is the quickest to resolve, averaging twice as fast as the sweep. Once the plume is located, the V term attracts the robot to keep resolving the plume whilst still navigating to areas that reduce the uncertainty of the model. Across all metrics, this proposed algorithm performed the best.

Figure 4 shows the mean maps for best strategy $(V \times \sigma/C)$, the lawnmower sweep and the CFD ground truth at 4 time steps across the simulation. The CFD maps show the large scale stability of the plume, but also how the simulation is locally time variant in regions and hence why the RMSE values in Fig. 3 change despite the algorithms converging. They also provide visual stimuli as to how well the model would be able to inform a first responder as to potential evacuation zones (a metric impossible to attain from the raw RMSE output).

VI. CONCLUSION

The proposed algorithm leverages the unique uncertainty output from the GMRF formulation to quickly and efficiently resolve a gas release in a cluttered large scale urban environment. The incorporation of normalised concentration and traversal cost leads to an accurate distribution map acquisition approximately twice as fast as a manual sweep and quicker than all other proposed reward functions that only use a subset of the variables utilised in the final algorithm. It is found that the uncertainty term dictates the robot to be more exploratory, biasing toward unsampled locations, whilst the normalised concentration term keeps the robot within the region of interest. Accounting for traversal and sampling cost maximises the number of samples that the robot takes within a given time frame.

The autonomous nature of the algorithm means there is no need for the robot to be first situated in zone with concentration data available and only requires an obstacle map of the surroundings. Another benefit compared to a manual sweep is that a sweep requires some form of luck to suitably sample an unknown area, which is not the case for the final algorithm which intelligently searches to find the plume.

Further work on IPP in GDM is however, still required. An experimental validation of the results seen in the simulation and testing on a different scenario (such as an indoor stage) will further validate the work shown here.

ACKNOWLEDGEMENTS

The author would like to thank Defence Science and Technology Laboratory (DSTL) for providing the DAPPLE experiment dataset used as the ground truth in this paper.

REFERENCES

- [1] R. R. Murphy, J. Peschel, C. Arnett, and D. Martin, "Projected needs for robot-assisted chemical, biological, radiological, or nuclear (CBRN) incidents," 2012 IEEE International Symposium on Safety, Security, and Rescue Robotics, SSRR 2012, vol. 00, no. c, pp. 1-4, 2012
- [2] A. Lilienthal and T. Duckett, "Gas Source Localisation by Constructing Concentration Gridmaps with a Mobile Robot," in Proceedings of the European Conference on Mobile Robots (ECMR 2003) 2003
- [3] A. J. Lilienthal, M. Reggente, M. Trinca, J. L. Blanco, and J. Gonzalez, "A statistical approach to gas distribution modelling with mobile robots - The Kernel DM+V algorithm," 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009, pp. 570-576, 2009.
- [4] M. Reggente and A. J. Lilienthal, "Using local wind information for gas distribution mapping in outdoor environments with a mobile robot," Proceedings of IEEE Sensors, pp. 1715-1720, 2009.
- S. Asadi, C. Plagemann, C. Stachniss, M. Reggente, and A. J. Lilien-[5] thal, "Statistical Gas Distribution Modeling Using Kernel Methods," Intelligent Systems for Machine Olfaction, vol. 1, no. March, pp. 153-179, 2011.
- [6] 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.
- [7] A. Gongora, J. Monroy, and J. Gonzalez-Jimenez, "Joint Estimation of Gas & Wind Maps for Fast-Response Applications," Applied Mathematical Modelling, vol. 87, pp. 655-674, 2020.
- [8] C. Stachniss, C. Plagemann, and A. J. Lilienthal, "Learning gas distribution models using sparse Gaussian process mixtures," Autonomous Robots, vol. 26, no. 2-3, pp. 187-202, 2009.
- [9] M. Hutchinson, P. Ladosz, C. Liu, and W. H. Chen, "Experimental assessment of plume mapping using point measurements from unmanned vehicles," in Proceedings - IEEE International Conference on Robotics and Automation, vol. 2019-May, 2019, pp. 7720-7726.
- [10] M. G. Jadidi, J. V. Miro, R. Valencia, and J. Andrade-Cetto, "Exploration on continuous Gaussian process frontier maps," Proceedings - IEEE International Conference on Robotics and Automation, pp. 6077-6082, 2014.
- [11] A. Marjovi and L. Marques, "Multi-robot olfactory search in structured environments," Robotics and Autonomous Systems, vol. 59, no. 11, pp. 867-881, 2011. [Online]. Available: http://dx.doi.org/10.1016/j.robot.2011.07.010
- [12] M. Hutchinson, C. Liu, and W. H. Chen, "Source term estimation of a hazardous airborne release using an unmanned aerial vehicle," Journal of Field Robotics, vol. 36, no. 4, pp. 797-817, 2019. [Online]. Available: http://dx.doi.org/10.1002/rob.21844
- [13] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," Robotics and Autonomous Systems, vol. 61, no. 12, pp. 1258-1276, 2013. [Online]. Available: http://dx.doi.org/10.1016/j.robot.2013.09.004
- [14] M. Coombes, T. Fletcher, W.-H. Chen, and C. Liu, "Optimal Polygon Decomposition for UAV Survey Coverage Path Planning in Wind," Sensors, vol. 18, no. 7, 2018. [Online]. Available: https://www.mdpi.com/1424-8220/18/7/2132
- [15] V. Hombal, A. Sanderson, and D. R. Blidberg, "Multiscale adaptive sampling in environmental robotics," IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems, pp. 80-87. 2010.
- [16] C. Rhodes, C. Liu, and W.-H. Chen, "An Information Theoretic Approach to Path Planning for Frontier Exploration," in Embedded Inteligence: Enabling & Supporting RAS Technologies. UK-RAS Network, jan 2019, pp. 40-43. [Online]. Available: https://www.ukras.org/wp-content/uploads/2019/03/UKRAS19-Proceedings-Final.pdf
- [17] X. Liu, S. Cheng, H. Liu, S. Hu, D. Zhang, and H. Ning, "A survey on gas sensing technology," Sensors (Switzerland), vol. 12, no. 7, pp. 9635-9665, 2012.
- [18] M. G. Jadidi, J. V. Miro, and G. Dissanayake, "Samplingbased Incremental Information Gathering with Applications to Robotic Exploration and Environmental Monitoring," The International Journal of Robotics Research, 2016. [Online]. Available: http://arxiv.org/abs/1607.01883
- [19] P. Neumann, S. Asadi, J. H. Schiller, A. J. Lilienthal, and M. Bartholmai, "An artificial potential field based sampling strategy for a gas-

sensitive micro-drone," IROS Workshop on Robotics for Environmental Monitoring (WREM), pp. 34-38, 2011

- [20] D. Martin, G. Nickless, C. S. Price, R. E. Britter, M. K. Neophytou, H. Cheng, A. G. Robins, A. Dobre, S. E. Belcher, J. F. Barlow, A. S. Tomlin, R. J. Smalley, J. E. Tate, R. N. Colvile, S. J. Arnold, and D. E. Shallcross, "Urban tracer dispersion experiment in London (DAPPLE) 2003: Field study and comparison with empirical prediction," Atmospheric Science Letters, vol. 11, no. 4, pp. 241–248, 2010. J. Monroy and J. Gonzalez-Jimenez, "Towards Odor-Sensitive Mobile
- [21] Robots," Rapid Automation, pp. 1491-1510, 2019.