2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

Predicting against the Flow: Boosting Source Localization by Means of Field Belief Modeling using Upstream Source Proximity

Finn L Busch[1,2], Nathalie Bauschmann[1], Sami Haddadin[3], Robert Seifried[1], and Daniel A Duecker[3]

***Abstract*— Time-effective and accurate source localization with mobile robots is crucial in safety-critical scenarios, e. g. leakage detection. This becomes particular challenging in realistic cluttered scenarios, i. e. in the presence of complex current flows or wind. Traditional methods often fall short due to simplifications or limited onboard resources.**

We propose to combine source localization with a Gaussian Markov Random Field (GMRF). This allows to improve source localization hypotheses by building on the GMRF's concentration and flow field belief that are continuously updated by gathered measurements. We introduce the upstream source proximity (USP) as a natural metric that exploits the joint knowledge represented in the field belief's concentration and flow field, i. e. predicting sources upstream. As a result, our method yields a computationally efficient source localization and field belief module providing substantially more stable gradients than conventional concentration gradient-based methods.

We demonstrate the suitability of our approach in a series of numerical experiments covering complex source location scenarios. With regard to computational requirements, the method achieves update rates of 10 Hz on a RaspberryPi 4B.

I. Introduction

A. Motivation

Rapid and accurate localization of pollution sources constitutes a challenging and complex task to be performed by mobile robotic platforms [1]. Consequently, it has been an active area of research across all robotic domains ranging from air over ground to underwater. Their well-known source localization scenarios include gas leakage and oil spills.

While source localization in open spaces without obstacles is already hard, tracing sources within confined and cluttered scenarios heavily increases the complexity of this task. For instance, flows are deflected by wall and obstacle boundaries. This results in complex pollutant fields that have to be handled by the source localization algorithm. At the same time, within cluttered scenarios communication performance is usually weak and even worse for aquatic domains [2]. Hence, offloading of computationally heavy tasks from the mobile robot towards a base station cannot be guaranteed. Hence, suitable methods need to be lightweight enough to run on-board robot's hardware.

[1] Institute of Mechanics and Ocean Engineering, Hamburg University of Technology, Germany.
[2] Division of Robotics, Perception, and Learning (RPL), KTH Royal Institute of Technology, Sweden.
[3] Munich Institute of Robotics and Machine Intelligence (MIRMI), Technical University of Munich (TUM), Germany.
Contact: daniel.duecker@tum.de
This work was supported by the European Union's Horizon 2020 research and innovation programme as part of the project ReconCycle under grant no. 871352.

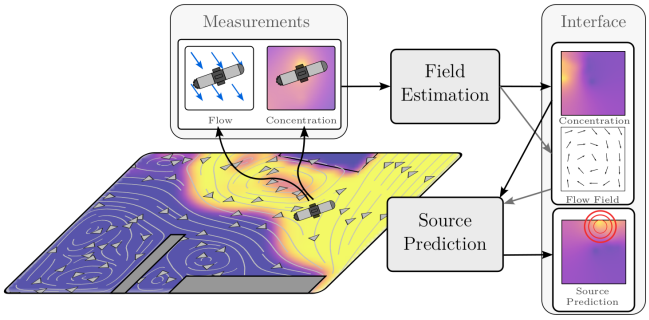

Fig. 1: A mobile robot aims to localize a pollution source within a confined and cluttered environment that is subject to flow. The environment is polluted by a contaminant (yellow). The mobile robot gathers flow and concentration measurements that are utilized by the field estimation model to construct a field belief. Using the field belief of both, pollution concentration and flow, the source localization algorithm predicts possible source locations.

B. Related Work

Source localization has been traditionally framed as a navigation problem, where the source is located by following simple, reactive rules, e. g. following airflow-following Anemotaxis [3], [4] or source search based on the quantity gradient [5]. For this, measurements on quantity and flow field are directly used to compute suitable control actions. We refer the reader to the works [6], [7] for a helpful overview on these approaches. However, as pointed out in [8], these simple control rules are usually limited to simple scenarios and fail in settings that include e. g. complex flows and obstacles, due to locality of the information used for navigation.

In order to overcome this short-coming, elaborated observation models have been introduced that allow to update the predicted source location belief based on incoming measurements and available prior knowledge on the environment. However, this comes with numerous challenges. For instance, source location cannot be observed directly and the relation of the flow measurement on the source location belief is hard to express analytically [9]. Hence, various simplifying assumptions are required, e. g. homogeneous flow and absence of obstacles. As a consequence, observation model approaches typically fall short in complex scenarios. Models based on computational fluid dynamics (CFD) simulation are usually computational intractable for embedded deployment. However, recent works [10], [11] propose using pre-

computed libraries of CFD-simulations with candidate source locations which lowers the computational burden during runtime.

For source localization, models that effectively incorporate concentration and flow remains a bottleneck. Limitations range from the simplifying assumptions to the high computational effortThis motivates to review the domain of field belief models which aim to construct meaningful field representations based on gathered measurements.

Traditional field models are usually physics-based and, thus, require considerable computational resources. However, recent progress in data-driven methods boosted the usage of probabilistic methods. As a plus, these direct provide an uncertainty measure along with the field values of interest, e. g. concentration and flow. While Gaussian process (GP) regression constitutes a common first choice, it becomes computational heavy with increasing number of measurements [12]. Sparse formulations such as [13] and [14] have been proposed to alleviate this flaw. The Kernel DM+V method [15] constitute an effective alternative and was applied in [16] for modeling a gas distribution that is subject to simple wind flow. Gaussian Markov random field (GMRFs) reduce computational burden associated with GPs by exploiting the Markov property [17]. Thus, they constitute a promising option for deployment on mobile robots with limited computational resources, e. g. small-scale underwater robots [18], [19]. Moreover, GMRFs have been successfully extended to incorporate obstacles [20] and wind flows [21].

When viewing source localization and field modeling methods together, very few approaches utilize field modeling methods to predict source locations. One recent promising approach is [1] which combines the field beliefs of concentration and wind flow. The key idea is to forward simulate candidate source locations using a simplified model. The resulting map is then compared with the beliefs of concentration and wind flow to obtain a likelihood. However, forward simulating every candidate location is computationally heavy.

Overall, we observe that most works on source localization predict locations directly from measurements. For this, often strong simplifications are required or the method renders computationally heavy. In contrast, works that combine concentration and flow estimation models are rare. However, recently proposed approaches such as Gongora et al. [22] provide promising results and motivate further research in this direction. In particular, it seems promising to investigate whether source localization efficiency can be leveraged by a simultaneously constructed field belief model from the gathered field measurements.

C. Contribution

The contribution of this work is three-fold. First, we propose a novel approach to source localization that builds on a simultaneously constructed field belief estimate. This approach is computationally light since it leverages knowledge from the explored flow field and available prior knowledge on boundaries and obstacles. Second, we introduce the *upstream source proximity* (USP) metric as a natural objective to effectively find source locations in scenarios with complex fluid flows. Third, our approach is evaluated in a series of numerical experiments ranging from simple to complex field scenarios.

Throughout the paper, we investigate three claims: (i) field belief models can be used as an effective alternative to observation models for source localization, (ii) the proposed USP metric provides meaningful guidance to an autonomous system, as opposed to unguided exploration, (iii) our combined approach is computationally light enough to run online on embedded hardware.

II. Problem Statement

Consider the exploration scenario of a confined cluttered fluid environment subject to flow currents that is polluted from a leakage source. At start, the underlying flow field is unknown. The goal is to efficiently predict the leakage source location from a series of spatially distributed field measurements (pollution concentration and flow), e. g. gathered by a mobile robot, as depicted in Fig. 1.

While existing literature mostly focuses on open spaces with homogeneous flows, cluttered scenarios demand for methods handling complex flow fields. Hence, it renders appealing to incorporate flow information into the source prediction strategy.

We propose to combine the source localization algorithm with a stochastic field belief representation that jointly estimates the concentration and flow field. This is promising since it allows to profit from estimated current flows when predicting source locations. For this, we assume that the field does not change significantly over time, i.e. the flow field is developed. While our approach is capable of predicting multiple sources, this proceeding's scope lies on the case of a single pollution source.

III. Methodology

A. Concept

The key idea of this paper is to leverage source localization efficiency by incorporating information on the present complex flow field. This can be broken down to the simple idea: If we measure pollution, its source can be expected upstream the estimated flow (and vice versa if no pollution is detected). This motivates combining a stochastic field belief, in form of a GMRF, that predicts the pollution concentration *and* the flow field based on gathered field measurements. Instead of using incoming measurements directly for source localization we use these to simultaneously build-up a field model that effectively provides insights in the surrounding field, e. g. flow effects originating from boundary conditions. Figure 2 illustrates this advantageous effect using only a small number of measurements. This belief of the concentration and flow field is then used as an highly informative input for the source localization problem, see Fig. 1.

It is intuitive to expect the pollution source upstream a flow. This motivates the definition of a new metric to which we refer to as *upstream source proximity* (USP) distance d_{USP}. It incorporates both, the pollution concentration level

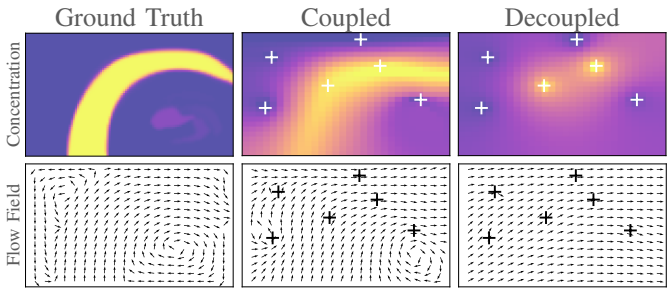

Fig. 2: Effect of *coupling* the flow and concentration field within the belief. The coupled estimate (center) based on six randomly sampled measurements is more effective in capturing the ground truth (left) compared to the decoupled case (right). Furthermore, exploiting basic prior knowledge of the flow field, e. g. mass-conservation and low flow velocities perpendicular obstacles, allows to incorporate knowledge on scenario structure. This allows constructing a comprehensive flow field belief from very few measurements (center bottom).

and the predicted flow field direction. Hence, by fusing both information into one metric we expect to improve prediction efficiency of the source location hypotheses. This is because we receive a more stable gradient for source tracking and identification in comparison to standard concentration gradients, see Fig. 3.

B. Field Estimation

Our field estimation module is based on the method proposed by Gongara, Monroy, and Gonzalez-Jimenez [22]. The method employs a GMRF over the concentration $\boldsymbol{q}$ and the flow velocities $\boldsymbol{f}^x$, $\boldsymbol{f}^y$ in x and y-direction, respectively. However, for the sake of brevity, we refer the reader to the detailed derivation in [22] and rather point out adaptations relevant to the scope of this work.

The GMRF is formulated as a factor graph, where factors are specified by an energy function reading

$$E_i = \frac{1}{2}\frac{\boldsymbol{r}_i^2}{\sigma_i^2}, \tag{1}$$

where $\boldsymbol{r}_i$ is the residual of factor i and is a function of the estimation variable. Moreover, σ_i^2 is the factor's variance which scales the relative importance of the factors. The edge factors E_e connect neighboring nodes and the observation factors E_y connect observations to the observed node. They read

$$E_\mathrm{e}(e_i) = \frac{1}{2}\frac{(q_0(e_i) - q_1(e_i))^2}{\sigma_\mathrm{e}^2} \quad \text{and} \tag{2}$$

$$E_\mathrm{y}(y_i) = \frac{1}{2}\frac{(y_i - q(y_i))^2}{\sigma_\mathrm{y}^2}, \tag{3}$$

where e_i denotes the edge formed by a neighboring node pair and $q(y_i)$ yields the concentration of the node related to the observation y_i. In order to encourage mass conservation within the flow field we define

$$E_\mathrm{mc}(i,j) = \frac{1}{2}\frac{(\nabla \cdot \boldsymbol{f}_{i,j})^2}{\sigma_\mathrm{mc}^2}. \tag{4}$$

Moreover the influence of obstacles on the flow,e. g. flow $\boldsymbol{f}_i$ perpendicular to obstacles $\boldsymbol{n}_o$ is close to 0, is captured by

$$E_\mathrm{o}(i,j) = \frac{1}{2}\frac{(\boldsymbol{f}_i^\top \boldsymbol{n}_o)^2}{\sigma_\mathrm{fo}^2}, \tag{5}$$

where $\boldsymbol{n}_\mathrm{o}$ denotes the normal vector of an adjacent obstacle. Finally, coupling between flow and concentration $E_\mathrm{xf}(i,j)$ reads

$$E_\mathrm{xf}(i,j) = \frac{1}{2}\frac{(\nabla q(i,j) \cdot \boldsymbol{f}(i,j))^2}{\sigma_\mathrm{fc}^2} \tag{6}$$

resulting into a small concentration gradient along the flow direction, Given this model, a field belief is constructed. For a given node s, the belief is then characterized by the normal distributions

$$\boldsymbol{q}(s) \sim \mathcal{N}\left(\mu_q(s), \sigma_q^2(s)\right), \tag{7}$$

$$\boldsymbol{f}(s) \sim \mathcal{N}\left(\begin{bmatrix}\mu_f^x(s)\\ \mu_f^y(s)\end{bmatrix}, \begin{bmatrix}\sigma_f^{x^2} & 0\\ 0 & \sigma_f^{y^2}(s)\end{bmatrix}\right). \tag{8}$$

C. Source Prediction using Upstream Source Proximity

In particular in the beginning of an exploratory mission, the available information is insufficient to predict source locations. Moreover, the information is likely limited to explored domains, e. g. the room where the exploration starts.

Thus, it is promising to extrapolate from information available in vicinity. For instance, the surrounding fluid flow is part of the overall flow system and may provide insights on source search directions. Hence, instead of predicting the actual source location, the revised objective is to predict the points closest to the source within the stream of the pollutant. Specifically, tracing the stream of the pollutant upstream eventually leads to the actual source location and thus provides guidance early on. For this, we introduce the *Upstream Source Proximity* (USP) d_USP. It is defined as the proximity of a given point to the source when following the stream of the pollutant. This means that $d_\mathrm{USP} \to \infty$ for approaching the actual source and $d_\mathrm{USP} \to 0$ for points very far from the source. It is important to note, that the USP does not correspond to any measurable distance. Rather it describes the *proximity* of a point to the source, relative to other points in the domain. Consider Fig. 3, where an exemplary field is depicted together with four candidate points. The USP for the four given points can be computed by comparing them to one another. Following the flow field, it becomes apparent that the blue point is situated upstream of the red point and downstream of the green point, reading

$$d_\mathrm{USP}(\textcolor{red}{\bullet}) < d_\mathrm{USP}(\textcolor{blue}{\bullet}) < d_\mathrm{USP}(\textcolor{green}{\bullet}). \tag{9}$$

To fulfill this inequality expression, it is sufficient to compare each point to the points downstream of it and add to their USPs. For instance, the blue point's USP can be computed by considering the red point only. The orange point is not situated within the stream of pollutants and consequently has a USP of 0. In a similar fashion, a grid of points can be defined and the USP can be computed for all of them, yielding a dense map of USP values for the domain. The

dependence of each point's USP on downstream neighbors only renders each point independent of all other points, given its direct neighbors. This allows formulating the problem in a local fashion. In this work, we adopt the structure of the GMRF formulation, placing a source candidate point at each position. Furthermore, we utilize the concentration estimate $\boldsymbol{q}$ to formulate the likelihood of a point being situated within the stream of the pollutants and the flow field estimates $\boldsymbol{f}^x, \boldsymbol{f}^y$ to determine the points downstream of a given point in a probabilistic fashion. We define the USP of a single point s as

$$d_{\mathrm{USP}}(s) = \sum_{s \in n(s)} p(s, s')[1 + d_{\mathrm{USP}}(s')], \tag{10}$$

where $n(s)$ denotes the set of the 8 neighboring points of s. Hereby, we define the transition function $p(s, s')$ as

$$p(s, s') = p_{\mathrm{stream}}(s) p_{\mathrm{migration}}(s, s'), \tag{11}$$

where

$$p_{\mathrm{stream}}(s) = \mathrm{CDF}(\boldsymbol{q}(s) > c), \tag{12}$$

denotes the likelihood of s being situated within the stream of the concentration. We model this as the likelihood of s exhibiting a larger concentration estimate than a threshold value c, given by the cumulative distribution function of the field belief CDF. Hereby, c remains a tuning parameter. The migration function $p_{\mathrm{migration}}(s, s')$ incorporates the relative proportions of concentration in s flowing to s', according to the flow field estimates. We approximate the proportions by sampling 200 transitions from the normal distribution of the flow estimates at point s, see Eq. 8. The formulation proposed in Eq. 10 can be interpreted as a Markov Decision Process (MDP). To solve the MDP, we use Value Iteration which, as discussed in [23], is an efficient method with guaranteed convergence.

IV. Analysis

We evaluate our proposed method responding to our claims (i-iii) from Sec. I-C For this purpose, we investigate the method in two studies in variations of the scenario depicted in Figs.4 and 5.

A. Evaluation Scenarios

Throughout our investigation we consider the scenario depicted in Fig. 4 consisting of four rooms with two flow inlets, one outlet, and the pollutant source at five different locations, referred to as *A* to *C*. The ground truth of the scenario is generated through a multiphase OpenFOAM CFD simulation [24] where water constitutes the fluid and the pollutant is gasoline. The resulting pollutant fields are depicted in Fig. 5.

For the guidance study, we employ a simple breadth-first-search (BFS) planner that operates on the GMRF grid and greedily searches for the shortest path from the robot's current position to the most probable source location. The agent can move one step at a time on the grid, in either a vertical or horizontal direction, with the constraint of not penetrating obstacles, and starts at the outlet in the bottom left. At each step, the robot collects a flow and concentration measurement. The robot concludes the mission if it has been stationary at the current source prediction for more than 5 steps.

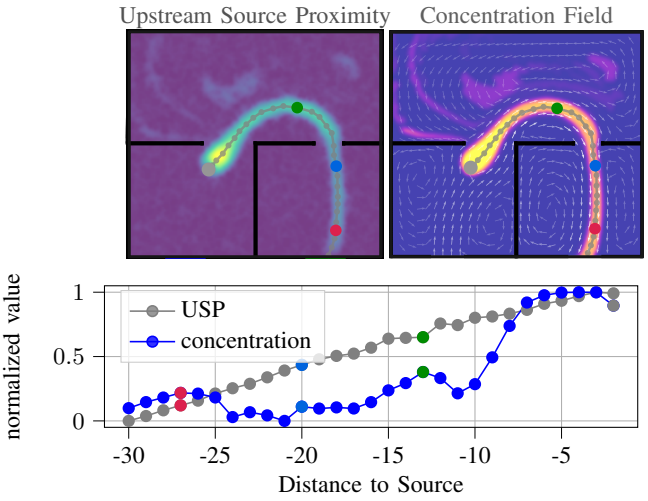

Fig. 3: Comparison of USP metric and the traditional concentration metric evaluated on an exemplary source location scenario. The USP provides a more stable gradient compared to the concentration level which dips, e. g. , in the vicinity of doors.

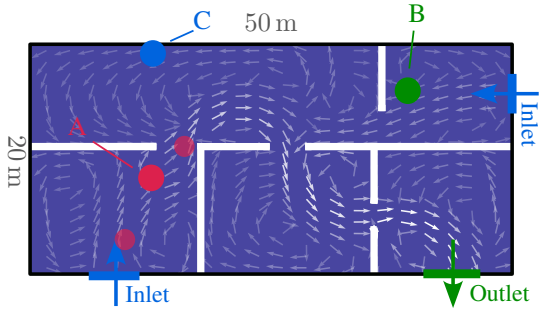

Fig. 4: Exploration scenario consisting of four rooms, two fluid inlets (*blue*), one outlet (*gree*). We investigate three source locations of varying complexity. Location A has two variations to investigate the sensitivity to small deviations in the source location.

B. Uniform Sampling

In this first study, we investigate the method's performance under uniformly sampled random noisy measurements. The goal is to evaluate the performance of the source localization method independent of a connected planning algorithm. Specifically, we respond to claim (i+iii). For this, the source location algorithm is executed in all 5 location variants for 100 independent runs to obtain meaningful statistics. The results are depicted in Fig.6. We observe that with increasing number of measurement the source prediction error drops rapidly. An interesting deviation is the error towards Location C which drop fastest but has the highest remaining error.

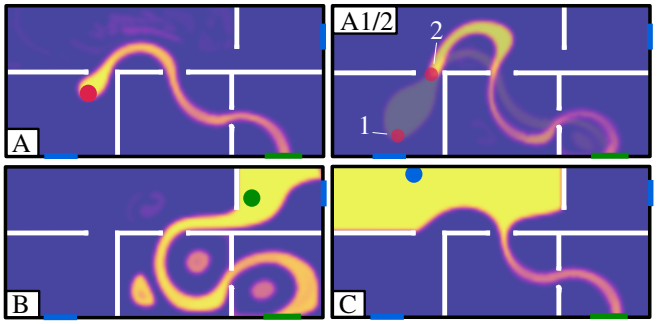

Fig. 5: Resulting simulated pollutant distribution for the different source locations. The distributions for variants A1 and A2 are depicted in top right.

TABLE I: Computation times under varied number of nodes evaluated on a Desktop PC with a Ryten 7 3700X CPU and a RaspberryPi 4B. In all cases, the source localization algorithm can be executed with an update frequency >10 Hz.

# of Nodes	Type	Desktop PC	RaspberryPi 4B
20 × 17	Source Loc.	6.1 ms	18.0 ms
	GMRF	38.9 ms	272.2 ms
30 × 25	Source Loc.	14.7 ms	41.1 ms
	GMRF	277.4 ms	2 120 ms
40 × 33	Source Loc.	26.7 ms	73.2 ms
	GMRF	1 101 ms	9 126 ms

This can be explained when looking into Fig. 5 where we observe that the whole room around C is covered with a high concentration of the pollutant. When comparing Location A and its variants A1 and A2 we see a considerable sensitivity on the resulting pollutant tail. However, in all cases, the source location is determined accurately.

Additionally, we evaluate the computational effort required for our method. For this, we run the algorithm on both a full-fledged desktop computer and a RaspberryPi 4B single-board computer as it is commonly used on mobile robotics platforms. Table I depicts the update times for the source localization and the GMRF field separately. We see that the efforts for the USP source localization are for considerably smaller than the GMRF computation times. It is worth pointing out that the proposed source localization method is light enough to run under all configurations with more than 10 Hz on the RaspberryPi computer. The strong increase in the computation time for the GMRF on the RaspberryPi is most likely due to Pi-specific bottlenecks in the Eigen-library. Overall, this study demonstrates the general suitability of the combined approach for source localization.

C. Combination with Baseline Planner

In this second study, we evaluate the method's capability of guiding a mobile agent toward the source location. Moreover, we respond to claim (ii). During the derivation of the method, we revised the goal of predicting source locations to predicting the USP, in order to define a measure that can be predicted in a meaningful way with heavily limited information, and directs the agent towards the source location. We simulate the agent controlled by the BFS planner 50 times

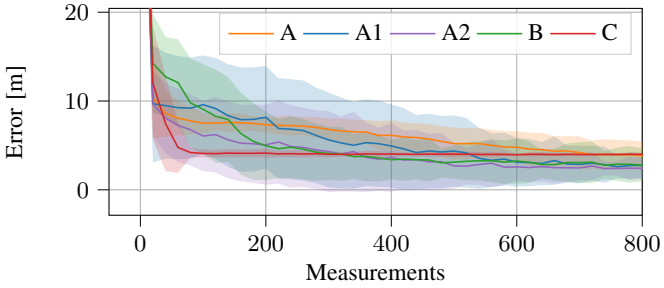

Fig. 6: Distance to the source with uniformly sampled measurements over 100 runs. The scenario's dimensions are 50 m × 20 m.

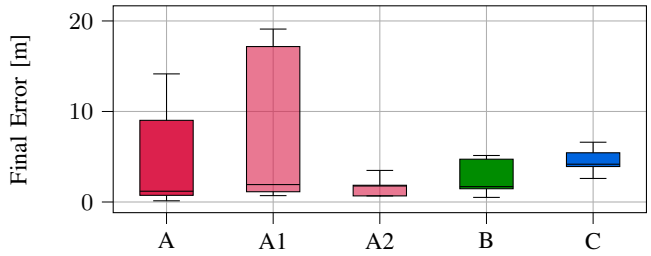

Fig. 7: Box Plot of the Final error using BFS planner. Whiskers at 1.5 × IQR.

for each setup as discussed in Sec. IV-A. The final errors for all setups are depicted as box plots in Fig. 7. We observe that for all setups, the median error remains far below 10 m, indicating successful source localization. Furthermore, as expected, scenario C yields the largest median error.

In Fig. 8, we display a heatmap over visited areas of the domain for case A. Hereby, it becomes apparent that in all runs, the agents do not deviate far from the stream of the pollutant, effectively tracing the stream to the source. This indicates that the guidance provided by the algorithm is effectively guiding the agents, which in turn avoids unnecessary exploration into regions that do not lead to the source.

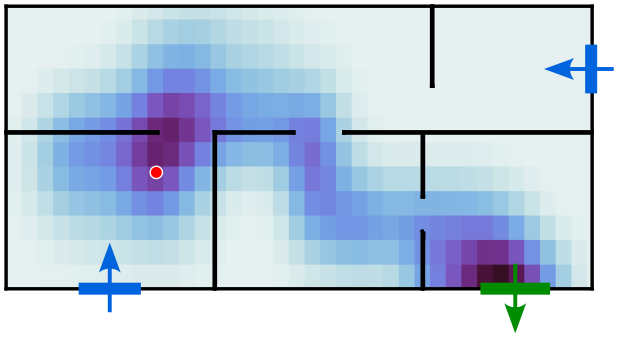

Fig. 8: Heatmap of visited areas in Scenario A together with the ground truth source location (red). The robot traces the stream to the source by following the subsequent source predictions, without exploring areas outside of the stream.

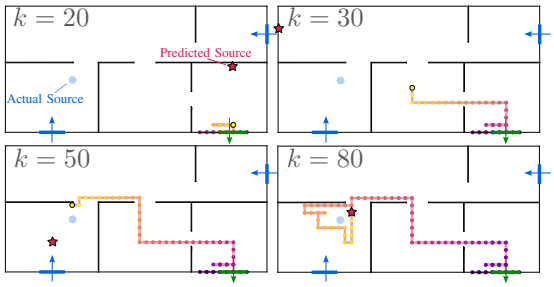

Fig. 9: Snapshot sequence from one localization run using the of a Breadth-First-Search-Planner. The source predictions are subsequently refined as more information is gathered.

A sequence of time points for one exemplary simulation together with the corresponding source predictions is shown in Fig. 9. Here, we observe that the predictions directly lead the agent from the initial position to the source. While some initial predictions might be inaccurate in terms of source locations, they nevertheless lead to the right direction, allowing the agent to gather more meaningful information and refine the predictions.

Overall, we conclude that the proposed method is capable of accurately predicting source locations in complex, realistic scenarios. By providing meaningful guidance to autonomous systems, our approach enables efficient and fast source localization, which is essential for rapid responses in critical situations. Moreover, the method's ability to circumvent local minima by offering stable gradients represents a substantial advantage over conventional concentration-gradient based methods. This, coupled with its comparably low computational demands, makes the method well-suited for deployment in practical applications.

V. Conclusion

In this work, we study problem of source localization in cluttered fluid environments that are subject to complex flow fields. For this, we propose to combine source localization with a probabilistic field belief model that simultaneously constructs beliefs on the concentration and the surrounding flow field. Furthermore, we propose the *upstream source proximity* (USP) metric to exploit the information flow field provided by the belief model. This allows to effectively predict the source location upstream the current flow. We demonstrate the performance in a series of numerical simulation and its computational effort on a RaspberrryPi 4 single board computer achieving a update rate of 10 Hz.

Future work will include the extension to time-varying flow fields and multi-source scenarios will be investigated Moreover, the approach will be validated onboard a small-scale mobile ground robot.

References

[1] P. Ojeda, J. Monroy, and J. Gonzalez-Jimenez, "Robotic gas source localization with probabilistic mapping and online dispersion simulation," 2023.

[2] S. Watson, D. A. Duecker, and K. Groves, "Localisation of unmanned underwater vehicles (uuvs) in complex and confined environments: A review," *Sensors*, vol. 20, no. 21, 2020.

[3] Y. Kuwana, I. Shimoyama, Y. Sayama, and H. Miura, "Synthesis of pheromone-oriented emergent behavior of a silkworm moth," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS '96*, vol. 3, 1996, pp. 1722–1729 vol.3.

[4] A. Rutkowski, S. Edwards, M. Willis, R. Quinn, and G. Causey, "A robotic platform for testing moth-inspired plume tracking strategies," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*, vol. 4, 2004, pp. 3319–3324 Vol.4.

[5] R. Russell, A. Bab-Hadiashar, R. L. Shepherd, and G. G. Wallace, "A comparison of reactive robot chemotaxis algorithms," *Robotics and Autonomous Systems*, vol. 45, no. 2, pp. 83–97, 2003.

[6] Z. Li, Z. F. Tian, T.-F. Lu, and H. Wang, "Assessment of different plume-tracing algorithms for indoor plumes," *Building and Environment*, vol. 173, p. 106746, 2020.

[7] X.-X. Chen and J. Huang, "Odor source localization algorithms on mobile robots: A review and future outlook," *Robotics and Autonomous Systems*, vol. 112, pp. 123–136, 2019.

[8] P. Ojeda, J. Monroy, and J. Gonzalez-Jimenez, "An evaluation of gas source localization algorithms for mobile robots," in *Proceedings of the 3rd International Conference on Applications of Intelligent Systems*, ser. APPIS 2020. New York, NY, USA: Association for Computing Machinery, 2020.

[9] H. Magalhães, R. Baptista, J. Macedo, and L. Marques, "Towards fast plume source estimation with a mobile robot," *Sensors*, vol. 20, no. 24, 2020.

[10] M. Asenov, M. Rutkauskas, D. Reid, K. Subr, and S. Ramamoorthy, "Active localization of gas leaks using fluid simulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1776–1783, 2019.

[11] C. Sánchez-Garrido, J. G. Monroy, and J. G. Jiménez, "Probabilistic estimation of the gas source location in indoor environments by combining gas and wind observations." in *APPIS*, 2018, pp. 110–121.

[12] D. Barber, *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.

[13] A. Smola and P. Bartlett, "Sparse greedy gaussian process regression," in *Advances in Neural Information Processing Systems*, T. Leen, T. Dietterich, and V. Tresp, Eds., vol. 13. MIT Press, 2000.

[14] E. Snelson and Z. Ghahramani, "Sparse gaussian processes using pseudo-inputs," in *Advances in Neural Information Processing Systems*, Y. Weiss, B. Schölkopf, and J. Platt, Eds., vol. 18. MIT Press, 2005.

[15] 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 *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009, pp. 570–576.

[16] 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*, 2009, pp. 1715–1720.

[17] H. Rue and L. Held, *Gaussian Markov Random Fields: Theory and Applications*, ser. Monographs on Statistics and Applied Probability. London: Chapman & Hall, 2005, vol. 104.

[18] D. A. Duecker, A. R. Geist, E. Kreuzer, and E. Solowjow, "Learning environmental field exploration with computationally constrained underwater robots: Gaussian processes meet stochastic optimal control," *Sensors*, vol. 19, no. 9, 2019.

[19] D. A. Duecker, B. Mersch, R. C. Hochdahl, and E. Kreuzer, "Embedded stochastic field exploration with micro diving agents using bayesian optimization-guided tree-search and gmrfs," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8649–8656.

[20] 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, Jan 2016.

[21] A. Gongora, J. Monroy, and J. Gonzalez-Jimenez, "Joint estimation of gas and wind maps for fast-response applications," *Applied Mathematical Modelling*, vol. 87, pp. 655–674, 2020.

[22] ——, "Joint estimation of gas and wind maps for fast-response applications," *Applied Mathematical Modelling*, vol. 87, pp. 655–674, 2020.

[23] M. van Otterlo and M. Wiering, *Reinforcement Learning and Markov Decision Processes*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 3–42.

[24] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, "A Tensorial Approach to Computational Continuum Mechanics using Object-Oriented Techniques," *Computers in Physics*, vol. 12, no. 6, pp. 620–631, 1998.