On Gas Source Declaration Methods for Single-Robot Search

Pepe Ojeda*
University of Málaga
Email: ojedamorala@uma.es

Javier Monroy
University of Málaga

Javier Gonzalez-Jimenez
University of Málaga

***Abstract*—Source declaration, along with plume finding and plume tracking, is one of the needed processes for gas source localization (GSL). It is a fundamental part of the search, since it is responsible to decide whether the gas source has been found, and also to pinpoint its location. Despite its importance, source declaration is often ignored in most of the GSL research, the criteria for termination being selected in a seemingly arbitrary manner, or even not being discussed at all. A clear example of this is the large number of experiments in the literature that are declared concluded whenever the robot manages to physically reach the source, without formally declaring it. In this work, we seek to provide an overview of the most significant declaration methods that have been used in state-of-the-art research for single-robot GSL, analyzing their strengths and weaknesses. We also provide a preliminary experimental validation of these methods, focusing on how stable their performance is when their input parameters are modified.**

I. Introduction

The problem of gas source localization (GSL) with a single mobile robot is often divided into three main phases: plume finding, plume tracking, and source declaration.

Plume finding consists of devising an exploratory strategy for the robot to come into contact with the gas plume. This sub-problem is sometimes considered to not be part of the source localization, but instead a separate task, and GSL algorithms are often designed assuming that the robot starts inside the gas plume.

Plume tracking is the most studied part of source localization. During plume tracking, the robot uses sensory information (commonly gas and wind measurements) to navigate the environment towards the source of the gas. Even though the name plume-tracking implies the existence of a clear downwind gas plume, it is frequently used as an umbrella term to refer to any navigation based on gas measurements. Some counter-examples to the requirement of the existence of a plume include gradient-based methods [1] that assume the gas is being dispersed mostly through diffusion, without significant advection, and methods that are designed to work under heavy turbulence, which disrupts the formation of a plume [2].

A variety of strategies has been proposed for plume tracking, from purely reactive navigation, to probabilistic methods that use a model of gas dispersion to generate estimations of the source location and base its movements on this belief.

Lastly, source declaration is the process of concluding the search and producing a final decision on where the location of the source is. Thus, it involves establishing some termination criterion to be checked during the search, and a method for generating an estimated source location from the information that has been gathered up to that point. In the next sections, we discuss the most relevant methods of source declaration that have been reported in the literature, and also present experimental results of their performance.

II. Source Declaration Methods

The declaration of the source location is a challenging problem that is heavily dependent of the method employed for the search itself. In the case of algorithms that are purely navigational, without a probabilistic estimation process, it is particularly non-trivial to define a convergence criterion, since there is no belief about the source location that a measure of certainty could be attached to. Commonly, experiments carried out with these methods are declared immediately successful if the robot manages to get within a preset distance to the source location, skipping this step entirely [1].

With planned GSL strategies, typically based on probabilistic estimations [3], [4], [5], the problem of source declaration is better integrated into the search process. Since a belief about the source location is being maintained and updated, it is possible to exploit the uncertainty of this belief to declare the end of the search process. The details of how this declaration is carried out vary across research, but we outline the most common approaches below.

A. Fixed Probability Threshold

This method declares the end of the search when the estimated probability of a given candidate location being the source exceeds a fixed threshold. This strategy [5] is only applicable when the source location is modeled as a random variable with discrete range, as it would not be possible to talk about the probability of a single value in the continuous case. This declaration approach is simple and checking for its fulfillment is fast, but it has significant drawbacks. The main issue is the dependency on the spatial resolution considered by the search algorithm: if the resolution is set too fine, it may not be possible to pinpoint a specific cell in the grid as the location of the source. Instead, multiple neighboring cells might share similar probabilities of being the source of the gas release, and even though their combined probability may exceed the termination threshold, the source is never declared.

On the contrary, when a coarse discretization of the environment is being used, convergence may be easier, as the difference in position between neighboring cells is big enough that one can clearly discern which of the positions best fits the measurements acquired by the robot; but signaling a single cell as the source might not provide a precise location for the source, since each of them covers a larger area.

B. Stability of the Expected Value

This criterion establishes that the search is considered finished when the expected value of the source location has not changed significantly over a number of iterations [6], [7]. This approach is most often seen in methods that employ a particle filter to generate their estimations, being the final declared source position that of the expected value itself. The main drawback with this approach is its strong dependency on the values chosen for the maximum rate of change of the expected value and the number of iterations that must pass for it to be considered stable. Moreover, since this method does not directly analyze the spatial dispersion of the probability, but only its rate of change, it is possible for it to terminate the search early if the robot obtains multiple successive measurements that do not offer significant new information –even if the current belief still has high uncertainty.

C. Entropy

This method proposes using the entropy of the probability distribution as a measure of the uncertainty, and declaring the end of the search when it falls below a certain threshold [4], [8]. The entropy of the source location belief is often used by infotactic approaches to estimate the information gained by executing certain movements.

This approach overcomes some of the main limitations of the previously discussed methods: if a few cells in the grid accumulate most of the probability of being the source location, the entropy will be low, even if none of them are a clear winner; but if the probability is spread out over a large group of cells, the entropy will be high, regardless of how the probability distribution is changing over time.

The main problem with this solution is that entropy offers a measure of how much the probability is spread out over the possible values, but does not consider how similar those values are to each other. Therefore, the entropy of a distribution where two neighboring cells gather most of the probability will be equal to that of a distribution where these two high-probability cells are far apart from one another. However, when searching for a single source, it is trivial to see that the second case should not be considered to be on the same level of convergence as the first one.

D. Variance

This termination criterion is based on the variance of the probability distribution [2], [3]. Similarly to the entropy, variance serves as a measure of the spread of the probability over the possible values of the source location.

However, unlike entropy, variance can account for the similarity of the values that are being considered. In the

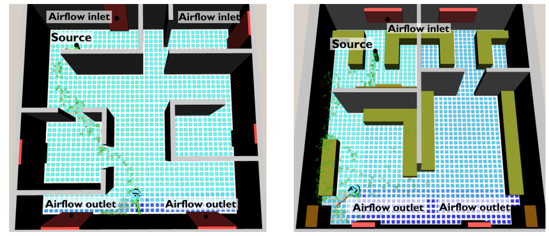

Fig. 1: Environments in which the experiments were carried out.

hypothetical case that was used as a counter-example for the entropy-based termination, where a few cells that are very distant from each other gather most of the probability, variance will appropriately be higher than when the same number of neighboring cells are the ones with high probability.

Since in most cases the source location random variable is defined to be a 2D or 3D vector, the actual measure of spread being used is the Generalized Variance [9] of the distribution, defined either as the determinant or the trace of the co-variance matrix.

III. Experimental Evaluation

This section presents a set of experiments aimed to evaluate and compare the source declaration criteria previously discussed. The experiments consist on performing a standard GSL search with a single robot and employing always the same algorithm [10] for controlling the navigation and the generation of probabilistic estimates. Only the source declaration strategy, and their respective parameters, are modified from one experiment to another, attempting to evaluate the impact on the final source declaration. The robot is equipped with a photoionization detector (*PID*) and a 2D anemometer.

Two indoor environments, with different dimensions but equal grid resolution, have been considered (see Figure 1). Experiment 1 was carried out in a 10x10m environment with multiple walls, leaving a total of 1058 free cells (*i.e.* not occupied by walls or furniture) in the grid that could be considered source candidates. Experiment 2 took place in a 11x12 environment with walls and furniture, leaving 485 possible source positions.

Figure 2 shows the results for each declaration method considered and for a range of their most significant input parameter. Each configuration was repeated 30 times. Specifically, we show and compare the proportion of runs that failed to terminate before a time limit of 10 minutes (%failed), the proportion of runs that terminated too early –*i.e.* before the robot reached the source– (%early), and the average error in the final estimated source position (error). It must be noted that the final estimated position is taken to be the expected value of the distribution, except in the case of the fixed probability threshold, where, to be congruent with the termination criterion, the estimation is the mode.

It can be observed that there is significant difference between the results of the studied methods. The probability threshold method shows to have a small range of values for which the proportion of failures is low, and setting the threshold outside that range results in either failing to declare the

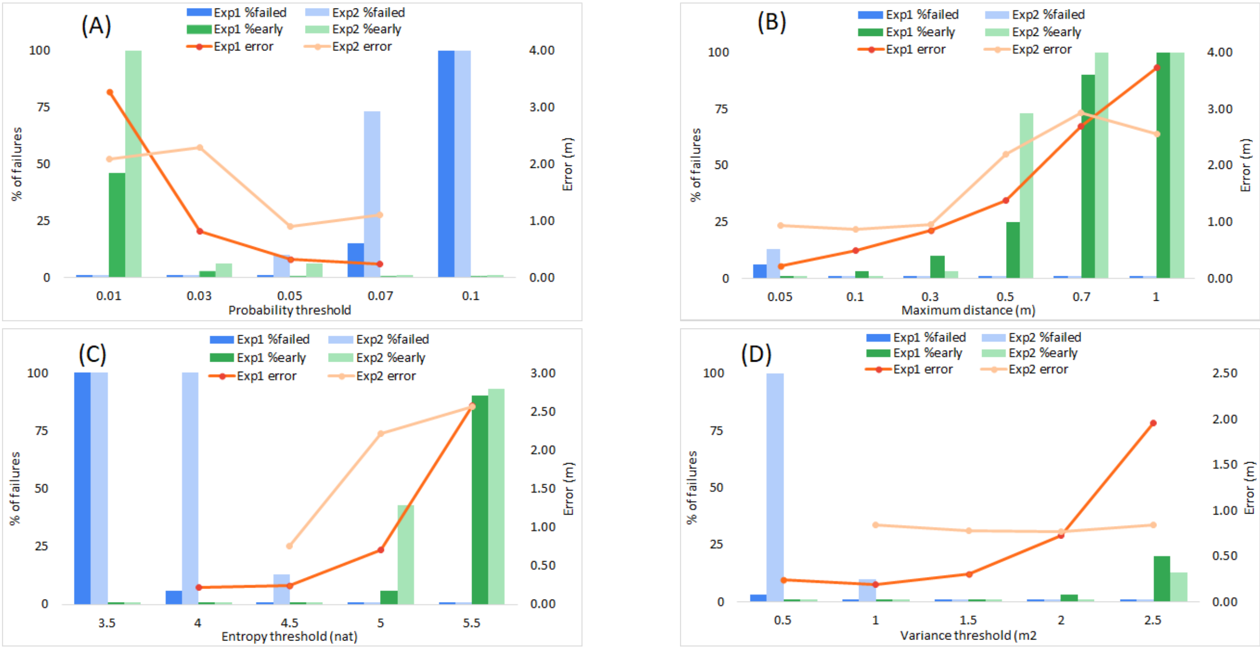

Fig. 2: Experimental results carried out in two indoor scenarios (Exp1 and Exp2) for a set of source declaration methods and a range of their input parameters. (a) Based on a single cell exceeding a fixed probability threshold. (b) Over-time stability of the expected value of the source. (c) Entropy of the probability distribution of the source location. (d) Variance of the probability distribution of the source location.

source withing the time-limit, or terminating before the source is found (with the consequent increase in the localization error. This termination criterion is, as can be seen, very sensitive to the threshold used and therefore must be set carefully.

The method that relies on the stabilization of the expected value is able to very reliably terminate inside of the time limit, although it is not rare for it to terminate early and produce a high error in the estimated location. An important advantage of this method is that it is fully agnostic of the total number of possible values (*i.e.* grid cells) considered, and thus the results are reasonably stable when changing environments.

The entropy method shows the greatest difference between the two considered environments, and a fairly high sensitivity to the chosen threshold. Although testing in more environments would be required for definitive conclusions, the results here indicate that this method might be too sensitive to changes in the number of spatial subdivisions.

Finally, the variance method shows acceptable results across a wide range of values, but, similarly to the entropy method, there is clear difference in the results for each environment, which suggests that the experimental conditions must be taken into account when setting the threshold.

IV. Conclusions

We have presented a brief review of state-of-the-art gas source declaration methods, focusing on the case of single-robot search. We have found some of them, particularly the variance method and the expected value method, to be less sensitive to changes in the input parameters or the number of possible source locations, making them more reliable and easier to configure correctly after some preliminary testing. We hope this article may be helpful to other researchers working on the field of source localization, and that more attention may be brought to the issue of source declaration, as it is an important but often overlooked part of GSL.

References

[1] J. Macedo, L. Marques, and E. Costa, "A comparative study of bio-inspired odour source localisation strategies from the state-action perspective," *Sensors*, vol. 19, no. 10, 2019.

[2] M. Hutchinson, H. Oh, and W.-H. Chen, "Entrotaxis as a strategy for autonomous search and source reconstruction in turbulent conditions," *Information Fusion*, vol. 42, pp. 179–189, 2018.

[3] 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*, vol. 27, no. 6, pp. 2388–2402, 2019.

[4] F. Rahbar, A. Marjovi, and A. Martinoli, "Design and performance evaluation of an algorithm based on source term estimation for odor source localization," *Sensors (Switzerland)*, vol. 19, no. 3, 2019.

[5] C. Sanchez-Garrido, J. Monroy, and J. Gonzalez-Jimenez, "Probabilistic estimation of the gas source location in indoor environments by combining gas and wind observations," *Frontiers in Artificial Intelligence and Applications*, vol. 310, pp. 110–121, 2018.

[6] 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.

[7] J. G. Li, J. Yang, J. Liu, G. D. Lu, and L. Yang, "Odor-source searching using a mobile robot in time-variant airflow environments with obstacles," *Proceedings of the 33rd Chinese Control Conference, CCC 2014*, pp. 8559–8564, 2014.

[8] F. Rahbar and A. Martinoli, "A Distributed Source Term Estimation Algorithm for Multi-Robot Systems," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 5604–5610, 2020.

[9] S. S. Wilks, "Certain generalizations in the analysis of variance," *Biometrika*, vol. 24, no. 3/4, pp. 471–494, 1932.

[10] P. Ojeda, J. Monroy, and J. Gonzalez-Jimenez, "Information-Driven Gas Source Localization Exploiting Gas and Wind Local Measurements for Autonomous Mobile Robots," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1320–1326, 2021.