Mapping Multiple Gas/Odor Sources in an Uncontrolled Indoor Environment using a Bayesian Occupancy Grid Mapping Based Method

Gabriele Ferri [1], Michael V. Jakuba [2], Alessio Mondini [1],
Virgilio Mattoli [3], Barbara Mazzolai [3], Dana R. Yoerger [4], Paolo Dario [1]

e-mail : g.ferri@sssup.it

[1] *CRIM Laboratory*	[2] *Australian Center for Field Robotics*	3 Center of Micro-BioRobotics@SSSA	[4] *Deep Submergence Laboratory*
Scuola Superiore Sant'Anna	*University of Sydney*	*Italian Institute of Technology*	*Woods Hole Oceanographic Institution*
Viale Rinaldo Piaggio 34, 56025 Pontedera (Pisa), Italy	*Rose Street Bldg, J04 Sydney, NSW, 2006 Australia*	*Viale Rinaldo Piaggio 34, 56025 Pontedera (Pisa), Italy*	*Woods Hole, MA 02543, USA*

Abstract - In this paper we address the problem of autonomously localizing multiple gas/odor sources in an indoor environment without a strong airflow. To do this, a robot iteratively creates an occupancy grid map. The produced map shows the probability each discrete cell contains a source. Our approach is based on a recent adaptation [15] to traditional Bayesian occupancy grid mapping for chemical source localization problems. The approach is less sensitive, in the considered scenario, to the choice of the algorithm parameters. We present experimental results with a robot in an indoor uncontrolled corridor in the presence of different ejecting sources proving the method is able to build reliable maps quickly (5.5 minutes in a 6 m x 2.1 m area) and in real time.

Keywords — indoor monitoring, gas source localization, gas source mapping, occupancy grid mapping.

1. Introduction

This paper addresses the issue of localizing/mapping multiple gas/odor sources in an indoor environment with no strong airflow. Gas source localization aims at solving the problem of rapidly localizing an emitter source ejecting a chemical agent. Improvements in this localization system could open up new horizons in robotics: we could use autonomous agents able to localize a gas/odor source in order to minimize the exposure of human agents to dangerous pollutants. These robotic automatic systems could find many applications in different scenarios: for environmental monitoring, for localizing hidden explosives or narcotics, for automatic humanitarian demining, for rescue tasks after accidents or natural disasters (survivors are seen as carbon dioxide sources), for the detection of dangerous chemicals leakages in a domestic or industrial environment and for the prevention of carbon monoxide poisoning caused by fire or inadequate ventilation of obstructed stoves.

The gas source localization task is not trivial. At medium and high Reynolds number gas dispersion is dominated by turbulent mixing [28]. Instantaneous gradients are poorly defined, time varying and frequently do not point towards the source [23]. Odor is concentrated in "packets", often with extremely low concentration measured between immersions in an odor packet.

Indoor environments present as a defining characteristic a lack of strong or persistent mean flow. The patchiness and intermittency of odor distribution typical of turbulent flows in the presence of strong mean flows is exacerbated by the lack of a persistent mean flow. In this case, the low energy turbulent mixing produces a patchier gas distribution. Local maxima in concentration have been observed at some distance from the source if the source has been active for some time [18].

The difference between an environment with the presence and the absence of a strong wind can be seen in Fig. 1 where the results of our source characterization experiments with one source ejecting ethanol are showed. In the figure the results of five sensors located at different distances from a source are reported.
These results indicate:

- without wind there are high fluctuations of gas concentration around the average value;
- with a strong and constant wind in an area near the source, the magnitude of fluctuations around the average gas value are reduced and gradient in concentration is relatively well defined. The energy carried by the strong wind leads to a more rapid mixing. The airflow can also be used by a robotic agent to infer the direction to the source.

To address the source localization issue, most researchers have proposed and implemented on robots gradient ascent based methods [12], [20], [27] and mixed behaviors of plume acquisition and plume following [7], [9] assuming or creating by fans a persistent mean flow. A strong and constant unidirectional airflow produces a plume of a well defined

shape to be followed up to its source.

A few works have addressed indoor environments with no artificially created strong ventilation, despite the fact that this condition is more realistic and common in indoor environments and represents real environments in which robots could work in the future [11]. In this setting, some researchers have proposed methods that either not rely or only partially rely on wind measurements [6], [10], [16], [17].

Multiple-source scenarios present additional problems to robotic searchers. Most reactive or behavior-based plume tracing algorithms implicitly assume that only a single source is present in the environment. While many of these algorithms would likely succeed in finding one source even in a multi-source environment, they offer no guidance on how to continue searching for other sources once a source has been found, or how to avoid re-finding the same source. To solve this problem we looked at map building techniques. Even if it may not be the most efficient way to locate a source, this approach can give important information on the features of gas dispersion or can be used in particularly challenging environments (not well defined plume, presence of multiple sources) to localize chemical sources where the robot might encounter difficulties in tracking a plume up to its source. Particularly in multi-source scenarios, map-building can supplant or complement reactive and behavior-based strategies by providing high-level guidance on the regions within a search area that have not been adequately searched.

In this work we demonstrate a map building technique that does not explicitly localize sources like reactive/behavior-based methods do, but creates maps of the environment from which areas containing sources are easily recognizable. In the approach pursued here a robot builds iteratively an occupancy grid map [22], [3] that indicates grid cells likely to contain sources.

We present a method based on an extension to Bayesian occupancy grid mapping standard algorithm proposed by Jakuba [15] that in this application creates maps with a reduced sensitivity to changes in algorithm parameters. In particular, it produces maps more consistent with the true number of sources in the case the probabilities with which the map is initialized (priors) underestimate the true number of present sources. In our setting the well known artifacts produced by traditional occupancy grid mapping techniques caused by independently estimating the state of each cell [31], [32] are emphasized because of the low prior with which the cells are initialized [15]: unlike the classical mapping problems using sonar sensors, only a few cells of the map are expected to be occupied containing a source.

A particular model for the sensor taking into account the sensed gas intensities to use all the information carried by the chemical measurements is presented.

The paper is organized as follows. Section 2 presents an introduction to techniques for mapping a gas concentration. In section 3 we introduce the classical occupancy grid mapping algorithm and we present our proposed method along with the procedure employed to select the algorithm's parameters. In sections 4 and 5 we describe the experimental setup and the results from experiments with a real robot using our occupancy grid algorithm and the standard one. We conclude with a discussion of our results and recommendations for future research.

2. Map building for chemical source localization

Some researchers have considered map-based approaches to chemical plume source localization. A map of the average gas concentrations can be produced from the measurements acquired by a grid of fixed sensors [8], [25]. Building a map with a grid of static sensors offers the advantage of reduced time consumption but much effort in sensors calibration is required. Furthermore, the system lacks in scalability because, increasing the area, the number of sensors needs to be incremented implying problems of cost and maintenance. Furthermore, a large number of metal oxide sensors can disturb the environment because the heaters present in the sensors can produce convective flows.

The use of robotic agents to produce gas distribution maps gives more flexibility. Hayes et al. [7] proposed to build a two-dimensional histogram over the number of "odor hits" (sensed values above a certain threshold) received in the corresponding area. The histogram bins contain the accumulated number of "odor hits" received by all robots (six robots were used) in the corresponding area during a random walk movement. However, the robots have to cover completely all the area of interest to produce a reliable map and the "odor hits", given their binary nature, strongly depend on the chosen threshold. During the mapping the robots have to scan the area and use a number of sampling locations relatively low not to make the task too much time-expensive. In order to produce a reliable gas distribution map of the whole area an extrapolation process is therefore needed. The proposed methods in literature have either built maps of concentration using different kinds of extrapolation [18], [21], [26], [29] rather than explicitly of source location or require the powerful assumption of a single source in the environment [5], [24].

In [18] the sensed gas intensities are convolved with a Gaussian kernel and experiments were carried out in an uncontrolled indoor environment. The method relies on the assumption that stabilised concentration maps represent the true average distribution, possibly with a small distortion due to the memory effect of sensors. This assumption is difficult to prove because there is no direct way to verify the observed concentration field (the "ground truth") for the mapping experiments. However some indications of the consistency of the produced maps comes from the applicability of the

produced concentration mapping method for gas source localization. In [21] gas measurements taken by 5 robots in a ventilated room are fused together with wind velocity in an advection-diffusion model (a non-turbulent transport model was assumed) using a Kalman Filter. The experimental area was divided into a grid and the gas concentration in each cell is represented by one state variable of the filter. In [26] experiments with a mobile robot driven on a zig-zag trajectory in a ventilated environment are reported. Since the measurement locations were not equally distributed, triangle-based cubic filtering was used for interpolation. The more jagged distribution with respect to that produced in the same environment from the averaged values sensed by a stationary sensor [26] shows a flaw of the interpolation. There is no means of averaging out the fluctuations of the gas concentration sensed at the sampling locations. A different approach aiming at addressing these problems was that proposed by Stachniss and co-workers [29] in which a sparse Gaussian process mixture model is used to represent the smooth background signal and to model sharp boundaries around patches of high concentrations.

Techniques to estimate the probability a source is in a certain position have been proposed: hidden Markov methods [5] and a Bayesian method [24].

Both methods, based on flow measurements and on binary detections and non detections of chemical effluent, update the probability that some cells on a grid contain a source. Farrell et al. [5] propose a hidden Markov method for plume mapping and source localization to estimate the likelihood of odor detection versus position, the likelihood of source position versus position, the most likely path taken by the odor to a given location, and the path between two points most likely to result in odor detection. This approach makes sense in strongly advective environments when the width of the plume is small relative to the search area.
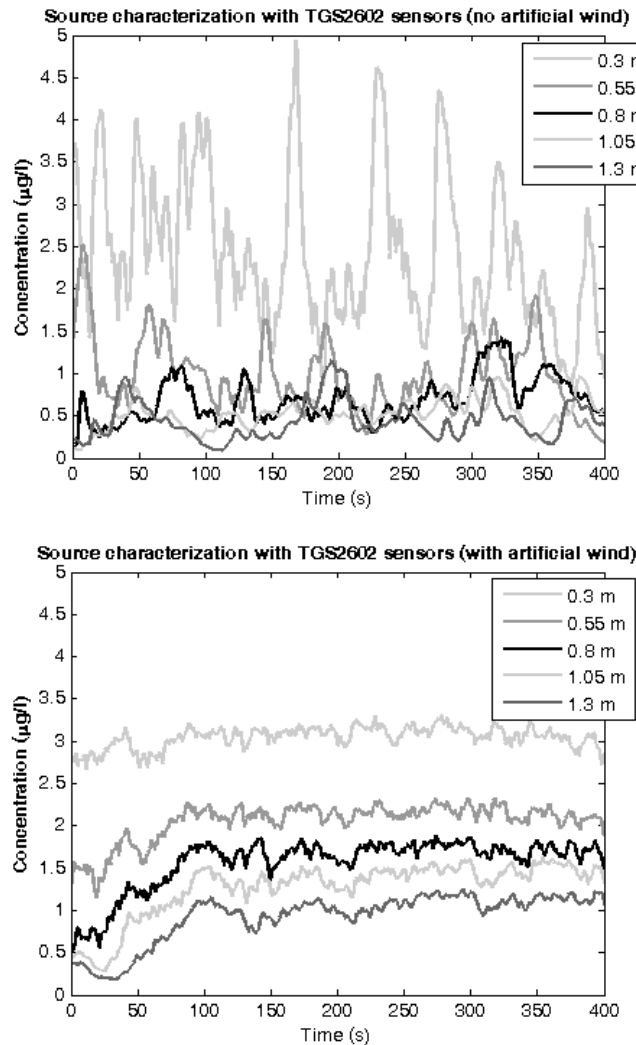

Fig. 1. Source (a circular 8.5 cm diameter dish ejecting ethanol) characterization. Sensor 0 is placed at a distance of 30 cm from the source. All sensors are spaced 25 cm, so that the farthest sensor (Sensor 4) is at a distance of 130 cm. Characterization without a predominant airflow (top). The distribution appears to be very patchy and an instantaneous concentration gradient is not present. Characterization with the presence of airflow (about 50 cm/s) generated by a fan placed behind the source (bottom) are presented. Here a gradient in concentration is relatively well defined.

In [24] a Bayesian methodology is proposed updating the state of a random field of possible locations for a source given binary detections/non-detections of effluent and a record of advective currents. This method is explicitly developed for short time-scale plumes suitable to be modelled as a collection of independently diffusing "filaments" [4]: each filament moves according to a random walk due to turbulent motion with a superimposed advection. The key assumption of this method is that a single source is present in the search area. This assumption makes this method different from occupancy grid mapping even if the grid-based representation of possible source location is similar to the idea of occupancy grid mapping. The assumption of one single source results in a very powerful one: this leads to an update rule that enables cells in portions of the grid not actually observed at a given time to have their posterior updated. The detection of the chemical at one location means all potential source locations not upwind to that position are unlikely to contain a source, while the upwind locations are likely to contain the source. Information about all cells of the map is extracted by a single measurement. A similar argument applies to the case where the effluent is not detected. Even if this method works in single-source scenario its extension to multiple sources is not straightforward.

For these reasons we considered an occupancy grid mapping based method that enables to build maps of the probability each discrete cell contains a source without considering the constraint of considering a single active source present in the search area.

3. Proposed algorithm

Occupancy grids were developed in 1980s as a finite-dimensional means of representing uncertainty in a robot's environment [3], [22]. The method requires discretization of the environment into a collection of small cells, each of which is approximated as being either completely occupied or else completely empty.

Occupancy grid mapping has been applied to many robotic applications: mapping of indoor environment [31], navigation, path planning [14], [30] and collision avoidance [2]. The maps have usually been created using long-range sensors as laser range finders [30], stereo vision [1] and sonar sensors [32].

Occupancy grid mapping is the process of fusing successive noisy and uncertain measurements into estimates of the likelihood of each cell being either occupied or empty. Recently, one of us has adapted occupancy grids to the mapping of likely plume source locations [15]. In that work, the binary state of grid cells is redefined to denote the presence or absence of an active plume source. A defining characteristic of chemical plume source localization considered within the context of occupancy grid mapping is the small number of cells expected to be occupied by active plume sources. From a Bayesian perspective, the prior probability of occupancy associated with each cell is much lower than is typical in occupancy grid maps of, for instance, the layout of an office or a building constructed from laser or sonar range-finder data [31], [32].

In this situation, the well known conflicts generated by standard grid mapping algorithm [13], [31] appear to be amplified [15]. These conflicts are caused by the sensors cone that sweeps on multiple cells inducing dependencies among the cells that cannot be taken into account by the standard occupancy grid mapping algorithm [31] that estimates independently the state of each cell. While in traditional applications these conflicts can generate maps where, for example, open doors appear closed, in source localization they generate maps with high posteriors respect to the true number of ejecting sources. These maps are, therefore, hardly usable to localize sources.

The map update rule employed in this work [15] is shown experimentally to produce better maps in our application than the standard Bayesian update rule. In particular, the resulting maps are more consistent with the true number of occupied cells, and less sensitive to the correct choice of prior.

We have also to underline the use of a gas sensor is quite different with respect to long-range sensors usually used in grid mapping because it provides punctual measurements. These punctual measurements (above all taking into account the gas intensities) carry information about a larger area: a high gas intensity means likely a source is near. A mechanism to retrieve the information carried by gas measurements is therefore needed to be able to build maps (see section 3.4).

3.1. Bayesian standard occupancy grid mapping algorithm

Let C denote the number of cells in an occupancy grid map, $m_{1:C}$ denote the collection of all binary cell states, and $z_{1:t}$ denote the collection of all measurements. Then posed as a Bayesian C-dimensional binary estimation problem, all information about the state of the map would be contained in the posterior:

$$p(m_{1:C} \mid z_{1:t}) \tag{1}$$

Unfortunately, the number of possible maps is 2^C, and so it is usually prohibitive to store or compute this full posterior. Consequently, occupancy grid mapping methods typically seek the marginal posteriors for the occupancy of each cell,

$$p(m_s \mid z_{1:t}) \tag{2}$$

However, it is not possible in general to compute these marginals exactly without first computing the full posterior and then marginalizing. To avoid the computational burden implied by this task, occupancy grid mapping methods employ map update rules that attempt to generate the marginal posteriors directly from measurements and the current state of the map.

The standard Bayesian map update rule produces inexact marginal posteriors, but it is simple to implement and popular [32]. For numerical reasons, the standard algorithm computes the so-called log odds of occupancy for each cell at each time step,

$$l_{t,i} = \log \frac{p(m_i \mid z_{1:t})}{1 - p(m_i \mid z_{1:t})} \tag{3}$$

from which it is trivial to recover the associated marginal posterior. At time t = 0, this quantity is initialized to the log-odds prior for each cell,

$$l_0 = \log \frac{p(m_i)}{1 - p(m_i)} \tag{4}$$

Besides the prior probability of occupancy $p(m_i)$ for each cell, the only other quantity required is an inverse sensor model to provide the probability that the cell being updated is occupied given the current measurement,

$$p(m_i \mid z_t) \tag{5}$$

Pseudo-code for standard algorithm is provided in Table 1. A critical assumption is required in the derivation of the algorithm of Table 1 [32]:

$$p(z_t \mid z_1, \dots z_{t-1}, m_i) = p(z_t \mid m_i) \tag{6}$$

This equation states that the current measurement is independent of all other measurements predicated on knowledge of the cell currently being updated. Because sensors sweep over multiple cells, this assumption is usually incorrect and consequently the marginal posteriors computed by the standard algorithm are inexact. However, assumption (6) greatly simplifies the problem of computing the marginal posteriors by enabling the mapping problem to be treated as a collection of C independent binary estimation problems [31]. It remains to specify a sensor model suited to indoor plume source localization.

3.2. A forward sensor model for plume source localization

We conjecture a characteristic common to all plume source localization scenarios is that measurements consist fundamentally of plume presence (detection) and absence (non-detection). Of course, auxiliary measurements of sensed concentration or wind direction are often also available, but these latter quantities only influence how plume detections and non-detections should be interpreted in terms of likely source locations. If one accepts the binary nature of plume detection (z_t=1) and non-detection (z_t=0), then a simple model for the probability of registering a detection given complete knowledge of the locations of all sources in the map (i.e. the true state of each grid cell) is [15]:

$$p(z_t = 1 \mid m_{1:C}) = 1 - (1 - P_f^t) \prod_{s \in O} (1 - P_s^t) \tag{7}$$

where P_c^t (s and c are used as a subscript) denotes the probability that sufficient signal from cell c arrives at time t to trigger a detection, P_f^t denotes the probability that a false alarm occurs at time t, and the product is over all occupied cells (O) in the map. The probability of non-detection is simply one minus this quantity.

To understand how (7) is derived, we start to show the probability of non detection z_t=0 given none or one occupied cell,

$$p(z_t\text{=}0 \mid O = \varnothing) = 1 - P_f^t$$

$$p(z_t\text{=}0 \mid O = s) = (1 - P_f^t)(1 - P_s^t) \tag{8}$$

The first relation of (8) states that the probability of a non detection given no cell is occupied is equal to the probability that a false alarm did not occur. The second states that if only one cell in the entire map is occupied, the probability of a

non-detection is equal to the probability that a false alarm does not occur and that the single occupied cell does not present a sufficient signal to trigger a detection at the location of measurement. Assuming that all the occupied cells of the map have independent probability to trigger a detection it follows that,

$$p(z_t=0 \mid m_{1:C}) = (1-P_f^t)\prod_{s\in O}(1-P_s^t) \tag{9}$$

From (9) it is trivial to derive (7). This model admits the possibility of multiple cells presenting simultaneously sufficient signal to the detector to cause a detection, but it does not admit multiple occupied cells reinforcing one another's signals to trigger a detection when the signal present at the detector from any single occupied cell is insufficient on its own. The sensor model described by (7) is a forward model, that is, it describes the probability of a measurement given the map. The parameters of forward models (P_c^t and P_f^t in the above) are usually easy to specify from calibration data. Section 3.4 contains an example of this procedure. However the standard Bayesian occupancy grid mapping algorithm requires the inverse marginals of (7), i.e. $p(m_i \mid z_t)$. It can be shown that (7) is analytically both invertible and marginalizable [15]. The results given a detection and non-detection are, respectively,

$$p(m_c \mid z_t = 1) = \frac{1-(1-P_f^t)(1-P_c^t)\prod\limits_{s\neq c}(1-P_s^t\, p(m_s))}{1-(1-P_f^t)\prod\limits_{s}(1-P_s^t\, p(m_s))}\, p(m_c) \tag{10a}$$

$$p(m_c \mid z_t = 0) = \frac{1-P_c^t}{1-P_c^t\, p(m_c)}\, p(m_c) \tag{10b}$$

where $p(m_c)$ is the prior for cell c.

3.3. An alternative map update algorithm

In plume source localization problems where the expected number of sources (occupied cells) is small, an alternative map update rule often produces better results than the standard algorithm given in Table 1. The marginal posteriors given by the inverse sensor model (10a) are exact. Therefore, if the set of all measurements consists of a single measurement, i.e. if $Z^t = \{z_1\}$, then (10a) can be used to compute the exact marginal posteriors. Now suppose these one-step marginal posteriors are regarded as independent:

$$p(m_1, m_2, ..., m_c \mid z_1) = \prod_{i} p(m_i / z_1) \tag{11}$$

This same requirement was placed on the prior $p(m_1, m_2, ..., m_C \mid z_1)$, thus posteriors satisfying this property can be regarded as specifying a new, "revised," prior. Furthermore, if at $t = 2$ the priors $p(m_c)$ in (10a) are replaced by the single-step marginal posteriors $p(m_c \mid z_1)$, and if these are independent, then the resultant marginal posteriors $p(m_c \mid z_2)$ produced will remain exact. Replacing the priors in (10a) with the revised priors $\tilde{P}_c^t = p(m_c \mid z_{t-1})$ at each timestep defines an alternative map update rule [15]. The algorithm pseudo-code is showed in Table 2.

Because OG algorithms intrinsically involve conditional independence assumptions, this algorithm was named after its particular assumption: Independence of Posteriors (IP). In principle, independence of the posteriors could be assumed in conjunction with any sensor model; however, the forward model (7) is special because the priors enter explicitly into the inverse sensor model generated from it, thereby providing a mechanism for interpreting the current measurement in terms of the present belief in the states of the cells in the map. Like other "context-sensitive" OG methods specific to sonar range finders [33], the IP algorithm interprets new measurements in light of the present state of the map. In fact, independence needs to hold only following detections for the marginal posteriors computed according to Table 2 to be exact; the marginal posteriors remain independent following any number of non-detections [15]. This asymmetry is apparent in Table 2 in the case of non-detection: the posterior odds do not depend on the state of any cells other than the cell in question i. Its existence can also be understood on intuitive grounds: correct determination of cell occupancy following a detection is strongly influenced by the state of nearby cells because the detected plume may have originated from one of the other cells rather than the cell in question. In contrast, the state of nearby cells is irrelevant following a non-detection. Since no detection has occurred, by definition no cell has presented signal to the detector, and consequently no ambiguity arises as to which cell is responsible for generating the measurement (none are responsible). Uncertainty in the state of the cells in

the perceptual field still remains because the combination of a plume and chemical sensor represents an imperfect composite sensor with probability of detection less than one.

Like the standard algorithm (showed in Table 1), the IP algorithm is inexact: however, the central assumption of independent posteriors required to derive it produces an algorithm that exhibits markedly different behavior. By revising the prior at each iteration a weak dependence is maintained between measurements within the map itself. Because measurements are interpreted differently depending on the state of the map, the results produced by IP algorithm depend on the order in which measurements are collected. By contrast, the classical algorithm produces identical maps regardless of the order in which measurements are added. This aspect of the IP algorithm represents a potential weakness, since bad measurements (false alarms or measurements that do not conform the assumed model) collected early during the survey can alter the interpretation of further measurements.

However, for environments with low priors, it does tend to produce maps that are more consistent with the true number of occupied cells and to be less sensitive to the choice of the algorithm parameters. Furthermore, its computational and storage costs are equivalent to the standard algorithm.

3.4. Algorithm implementation

It remains to specify the parameters of the sensor model, a prior probability of occupancy, $p(m_i)$, and to choose a grid cell size. In this section, we demonstrate how to specify the parameters, P_c^t and P_f^t, of the forward sensor model (7) from readily attained calibration data. The invertibility of (7) avoids the often more difficult task of specifying the parameters of an inverse model directly. Empirical expressions for generating the P_c^t were derived from the sensor and source characterization data shown in Fig. 1.

3.4.1. Sensors calibration and source characterization

As odor sensors we opted for TGS 2602 commercial sensors from Figaro Engineering Inc. They are metal oxide semiconductor gas sensors and in the presence of alcohol vapors, their internal resistance changes logarithmically. The choice of this kind of sensors was due to their high sensitivity, their usable long life span and their comparatively high robustness to changing environmental conditions [19]. They are small, cheap and can easily be integrated in the measurement circuit [19]. Their main drawbacks are a needed high operating temperature that results in a high power consumption, a slow recovery time after the gas is removed and a significant variability in response characteristics between individual sensors. For this reason, at first, some calibration sessions were carried out. The sensors (6 sensors: 5 for source characterization, 1 mounted on the robot) were tested in a hermetically sealed box, where a known and increasing amount of alcohol was periodically injected. A fan was also included to accelerate alcohol vaporization and rapid mixing. The signal coming from the sensors was acquired using a DAQCard-6024E™, from National Instruments. Data were sampled every 1 ms; we only stored the mean value computed on 500 measurements. In order to calibrate our sensors, the collected data were fitted with the characteristic bi-logarithmic function given by the sensors datasheet. Best fit parameters were extrapolated. A third degree polynomial was used to fit the data. The found coefficients of the polynomial were used to convert sensor voltage output in concentrations [6].

Once the sensors were calibrated, some preliminary experiments were carried out with environmental conditions similar to those present in the survey area in order to characterize the gas dispersion. We used an 8.5 cm diameter circular dish containing ethanol free to evaporate as a source for the calibration. The dimension of this source was similar to those used during our experiments with robot (see section 4). We used 5 sensors placed in a row, spaced 25 cm. The first one is at a distance of 30 cm from the source. Data acquisitions started after a source had been free to eject for 10 minutes. Two different kinds of trials were carried out: one without the introduction of an external artificial airflow and the other with a fan, placed 40 cm behind the source, generating an airflow of about 50 cm/s. In the trials where wind was present, the sensors were placed inside the generated plume. The results are shown in Fig. 1.

3.4.2. Determining the parameters of the sensor model

The percentage of measurements above a given reference concentration is shown in Fig. 2 for each gas sensor in the array. Based on these data, we fit analytical functions to the experimentally observed probability that measured concentration q exceeds a given reference concentration $p(q{\geq}q_{ref}; r, q_{ref})$ as a function of range r from the source and the reference concentration q_{ref}. First we fit a Gaussian curve to the data from each sensor to yield the probability of the measured concentration exceeding a reference concentration at the range r_i from the source (r_i=0.3, 0.55, 0.8, 1.05, 1.3 m) associated with each sensor and as a function of reference concentration:

$$p(q{\geq}q_{ref}; r_i, q_{ref}) = e^{-\frac{1}{2}\left(\frac{q_{ref}}{\sigma(r_i)}\right)^2} \tag{12}$$

where $\sigma(r_i)$ denotes a range-dependent scale parameter determined by the fit. Next we fit an exponential to $\sigma(r_i)$ to permit computation of $p(q \geq q_{ref};\ r,\ q_{ref})$ for all ranges:

$$\sigma(r) = e^{-5.849(r-0.325)} + 0.1323 \tag{13}$$

Fig. 2 shows the $\sigma(r_i)$ and the interpolating exponential function. Fig. 3 shows the final analytical approximation to $p(q \geq q_{ref}; r,\ q_{ref})$.

Let r_c denote the range between the robot and a cell c. Then the probability $p(q \geq q_{ref}\ ;\ r_c,\ q_{ref})$ determined thus far corresponds to the P_c^t we seek if the reference concentration q_{ref} is set to the measured concentration q as long as q exceeds some static master threshold q_0 discriminating a detection ($z_t=1$) and non-detection ($z_t=0$). For detections, the equivalence between $p(q \geq q_{ref}\ ;\ r_c,\ q_{ref})$ and P_c^t holds for a dynamically altered reference concentration ($q_{ref}= q$); however, non-detections require setting $q_{ref} = q_0$ in order to compute the correct P_c^t. In fact P_c^t denotes the probability that sufficient signal from cell c arrives at time t to trigger a detection; since a detection would have been declared only for $q \geq q_0$, in a non-detection we fix $q_{ref}= q_0$. Thus after each measurement the P_c^t are computed according to:

$$P_c^t = \begin{cases} p(q \geq q_{ref}; r_c, q_{ref}) & q \geq q_0 \quad \text{(detection)} \\[2mm] p(q \geq q_0; r_c, q_0) & q < q_0 \quad \text{(non detection)} \end{cases} \tag{14}$$

In this way, the measured concentration can be used to help infer range to source without violating the imposed form of the forward sensor model (7). In multiple-source environments, effluent from two or more sources can mix to produce, on average, higher measured concentrations than would be present at the same range from a single source. The forward sensor model prohibits inferring the number of sources contributing to the current measurement, and therefore high measured concentrations will cause the mapping algorithms employing this model to build maps that favor the explanation of a single nearby source rather than multiple sources farther away.

On the basis of Fig. 3, the probability of measuring a concentration above about 0.7 μl/l (1 μl/l of ethanol is equal to 383.6249 ppm or 0.789 g/m^3) is essentially independent of range for ranges greater than about 0.5 m from the source. This suggests a natural static master threshold for detection q_0 of about 0.7 μl/l. With a threshold near this value, the P_c^t associated with cells at least 1.5 m from the robot are less than about 0.02. On the basis of this fact, we chose to use a static perceptual radius of 1.5 m, outside of which all P_c^t are effectively regarded as zero. Computing the P_c^t at each iteration for all cells in the small map used in our experiments (~1250 cells) presents only a modest computational load; however, by employing a static perceptual radius, the algorithms presented here scale without modification to map sizes limited only by available memory.

Our experiments all employed a probability of false alarm, $P_f^t =0$. In our experience, the probability of measuring a concentration above about 0.7 μl/l in the absence of a nearby source is negligible. Finally, we chose a basic cell size for our experiments of 10 cm x 10 cm, on the order of the size of the sources used. We also present results from maps with 30 cm x 30 cm grid cells.

4. Experimental set-up

The experiments were performed using one robot of MOMO platform [6]. The PC tracks robot position (via web-cams), communicates with the robot, checks its state (through an FM communication channel – an AUR EL XTR-434H from AUREL spa Inc. transceiver is used either with the PC or the robots), collects the data produced by the robot actions and allows a human operator to tele-operate the robotic agent. One single robot (see Fig. 4) (17 cm long and 17 cm wide) has two encoders on its wheels for rough odometry, a main micro controller that manages movement (PIC 18F452 from Microchip Technology Inc.), two sets of high capacity rechargeable Ni-MH batteries (assuring at least two hours of durability), two bumper whiskers as touch sensors, and a gas sensor.

In the experiments described in this paper, the robot was tele-operated using the FM wireless connection by a human operator. The positions were acquired using two Logitech Quickcam Express II webcams (640x480 resolution) fixed to the ceiling. The robot was equipped with one TGS 2602 sensor from Figaro Engineering Inc. for alcohol detection.

The experiments were performed in an area of one corridor of CRIM Lab building. The dimensions of the area were 6 m x 2.1 m (see Figs. 5 and 6). The data processing was akin to that applied to our calibration data (see section 3.4.1). Once a sampling was commanded, the robot stopped and computed the average on a temporal window of 3 s. The robot followed a trajectory as that showed in Fig. 6. The samplings are spaced 40 cm: one trial consisted in 64 samplings.
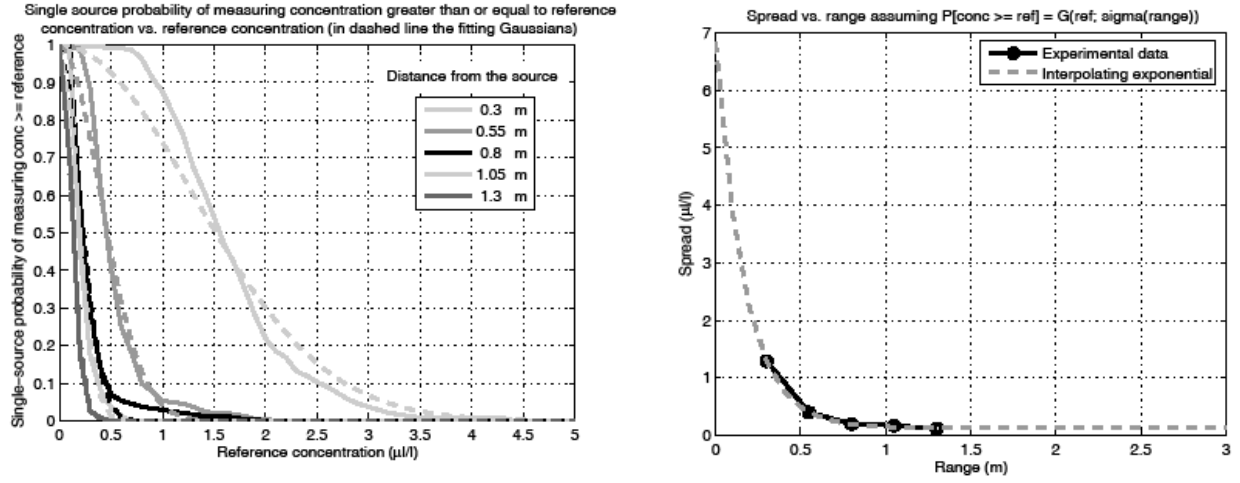
Fig. 2. (Left) Single source probability of measuring a concentration greater than or equal to reference concentration vs. reference concentration. The probabilities for the different sensors are shown. (Right) Standard deviation of the fitted Gaussian for the different sensors and the interpolating exponential.
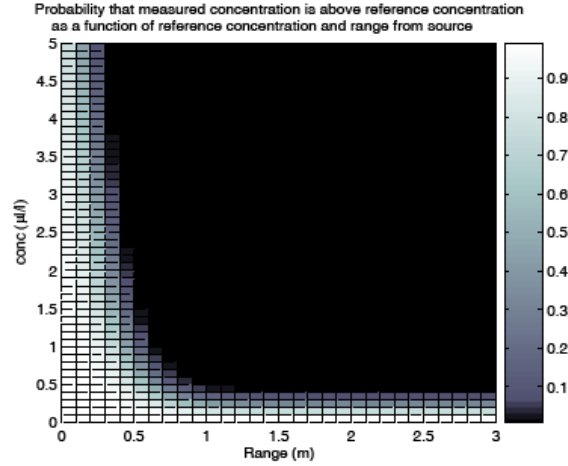

Fig. 3. Probability that measured concentration is above reference concentration as a function of the reference concentration and range from source.

Several trials were carried out in the presence of three sources ejecting ethanol. The sources used in the trials (labelled S1, S2 and S3) are different in shape and so eject alcohol at different rates. S1 was a circular 8.5 cm diameter dish, S2 was formed by two circular 8.5 cm diameter dishes and S3 was composed by a single circular 15 cm diameter dish. In all the trials the odour sources were kept at the same positions. The reported experiments were carried out in 5 experimental sessions carried out in different days during the month of February (second half of the month). In each session the first trial started after the sources had been ejecting ethanol for 10 minutes. The next trial started about 10 minutes after the end of the previous one. The temperature was between 18 and 20 Celsius degrees and there were no drastic changes of environmental conditions during the experiments. No artificial airflow was introduced. The heating system of the building was not turned off in order to create a realistic indoor scenario.

5. Experimental results

In this section we report the results of 15 experiments with the robot. Experimental results show that the proposed method can provide suitable maps for localizing multiple gas ejecting sources. A good map provides an indication both where sources are and are not likely to be, thus indicating which areas have been explored and which still need to be further investigated.

Figs. 7 and 8 show the posteriors (also in log odds form) of maps generated using the samplings acquired in two different experiments (they are named in the following "experiment 1" and "experiment 2"). The parameters used to generate both the maps are those presented in section 3.4. The priors for the cells were chosen to be consistent with 5

expected sources (0.004). The grid cell size for these maps was selected as 10 cm x 10 cm. Equations (10a) and (10b) were used in standard occupancy grid mapping algorithm as inverse sensor model.

Multiple sources of different sizes are present, diffusing ethanol in a turbulent environment. Moreover, the plumes generated by the different sources can coalesce. Nevertheless, our algorithm succeeds in building maps that are easily interpretable in terms of location and number of likely sources. There are clearly three separated candidate areas, one for each source. The areas are near the sources and are a bit to the "north-east" with respect to the true locations. This suggests the presence of a weak wind (further evidence appears in the raw data of Fig. 6, where the samplings of "experiment 1" are reported). This little gap between declared positions and true ones derives from the fact we did not use the wind information in our model given the difficulty in having a reliable measurement.
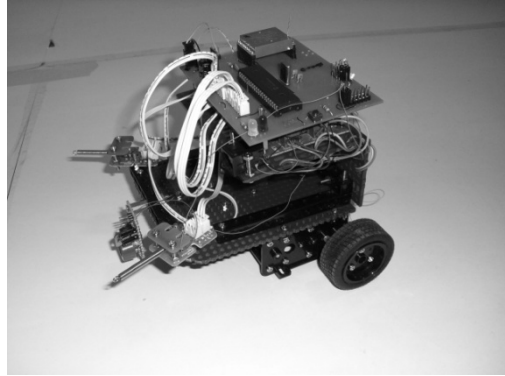

Fig. 4. One robot of MOMO platform.

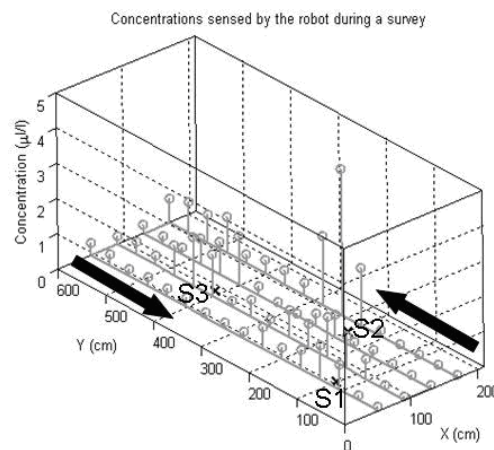Fig. 5. Experimental area in which experiments were carried out.

Fig. 6. Trajectory of the robot with gas samplings and sources locations (black crosses).

One merit factor that we consider in this work is the expected number of sources given the measurements, $E(\#sources \mid z_{1:t})$. It is computed as the sum over all the cells of the map of the posteriors, that is,

$$E(\# sources \mid z_{1:t}) = \sum_i p(m_i \mid z_{1:t}) \tag{15}$$

By comparing the posterior expected number of sources to the actual number of sources we can gauge the accuracy of the posteriors in the entire map. Accurate posteriors are necessary if decisions would be made on the basis of the map.
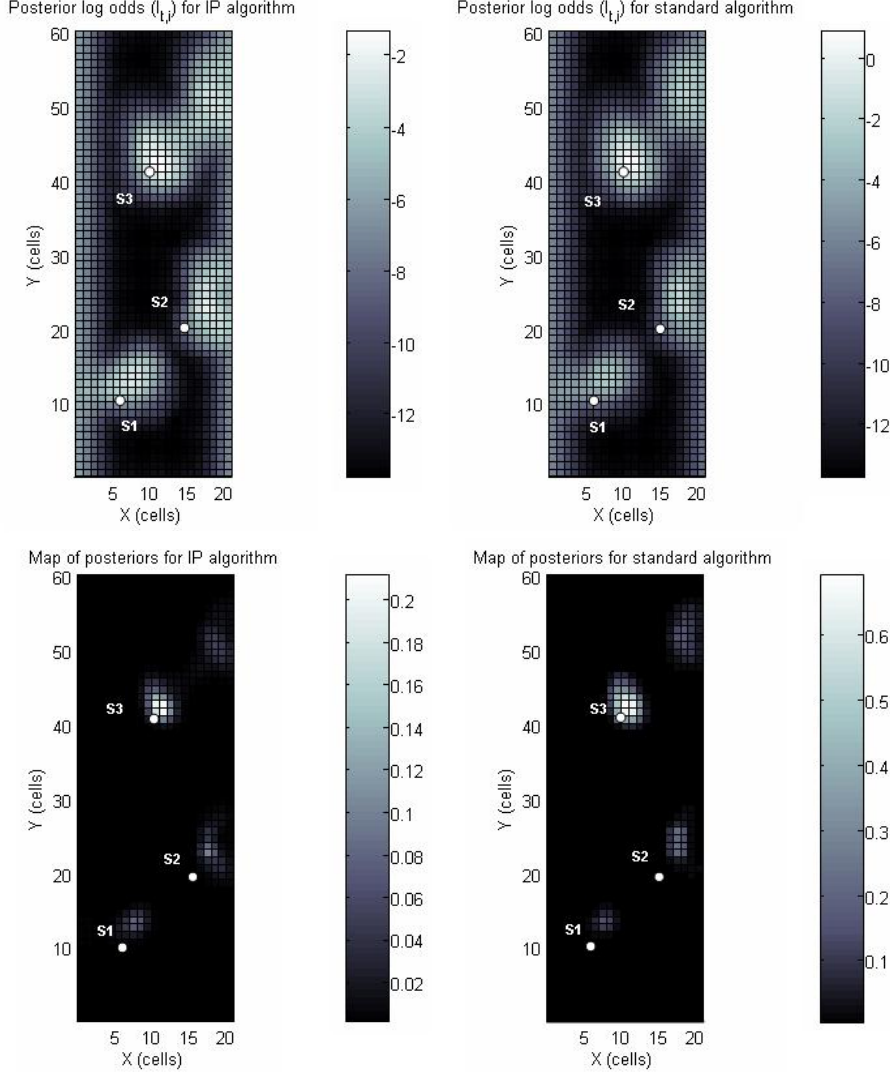

Fig. 7. Posteriors in log odds form (top) and posteriors (bottom) for one trial using IP and standard algorithm. The priors are equal to 4 10^{-3} (consistent with 5 expected sources). The maps are created from "experiment 1" dataset. The standard algorithm presents slight overestimated posteriors with respect to IP and the peaks related to the three different sources are less uniform as value than in maps produced with IP (*Note the different grey scales used*).

Figs. 7 and 8 compare the maps produced by the IP and the standard algorithm applied to two sets of experimental data: while these maps appear to be qualitatively similar there are differences. The first difference is in the estimated number of expected sources given the measurements: IP gives a result of 5.84 and 5.44 and it is an essentially correct number, standard algorithm produces 13.62 and 13.78 (overestimated numbers). Secondly, the IP produces a map with peaks of increased probability more uniformly distributed if compared to those produced by the standard algorithm. For example, in Fig. 8, we can notice that for standard algorithm the spot related to S3 contains cells with probabilities higher if compared to the spots related to S1 and S2.

Even if the grid cell size of 10 cm x 10 cm gives good results producing smooth maps, an interesting point to investigate is what happens to the produced maps if the size of the grid cells is incremented. This might be necessary when a large indoor environment has to be mapped and the robot is not provided with enough memory resources. The data of the "experiment 1" were used to create a map with the grid cell size of 30 cm x 30 cm (see Fig.9). The used prior is consistent with an expectation of 5 sources in the experimental area. Even if the produced maps are less smooth they are useful to localize the sources. The $E(\# sources \mid z_{1:t})$ in this case are 5.97 for the IP and 7.53 for the standard. So far, it appears from the showed maps that standard algorithm overestimates the number of expected sources slightly with respect to the IP

algorithm. Thus far we have utilized the correct prior: for "correct" we mean that the prior is compatible with the true number of sources present in the search area. However, the problem with the standard algorithm is its strong dependence on priors and thresholds.
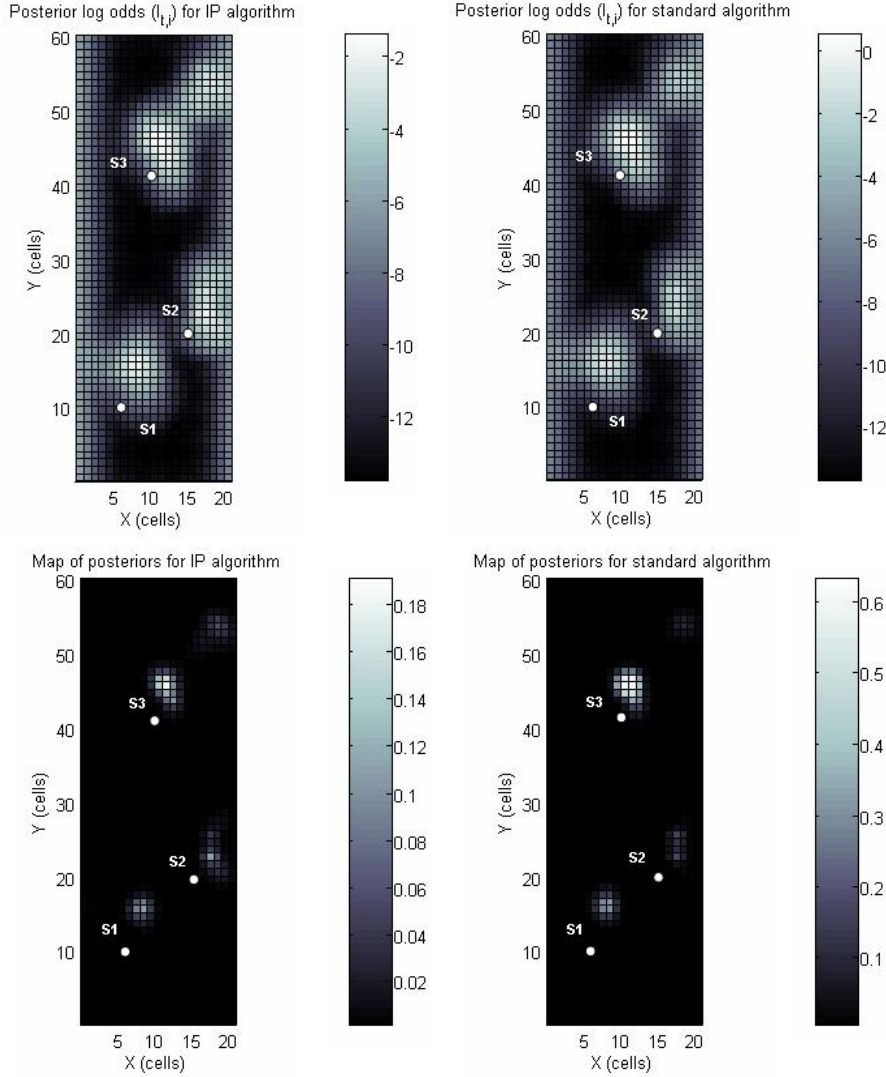

Fig. 8. Posteriors in log odds form (top) and posteriors (bottom) for one trial using IP and standard algorithm. The priors are equal to $4 \cdot 10^{-3}$ (consistent with 5 expected sources). The maps are created from "experiment 2" dataset. The standard algorithm presents slight overestimated posteriors with respect to IP and the peaks related to the three different sources are less uniform as value than in maps produced with IP (*Note the different grey scales used*).

This is evident from Fig. 10 where the posteriors produced by the algorithms with priors equal to $4 \cdot 10^{-5}$ are reported (this is the case where the prior underestimates the true number of sources). The data from which the maps are created are once again those of "experiment 1". It is evident that while the IP algorithm continues to produce a good map, the standard one totally overestimates the posteriors. This is evident in the number of expected sources given the measurements. While for IP algorithm $E(\#sources \mid z_{1:t}) = 7$ and this is still a good estimate, the standard one results in 195 estimated sources. It is clearly a completely unrealistic number. When viewed in log odds form the peaks produced by the standard algorithm are still in the right position; however it is unclear how to base decisions on these values.

The dependence of the two algorithms on the priors and three different thresholds is shown in Fig. 11 where the expected number of sources given the posteriors versus different used priors is reported. On the graph, the results of the two algorithms with the thresholds of 0.6 μl/l, 0.8 μl/l and 1 μl/l are reported. The averages of the results of the datasets from 15 experiments are reported. While IP algorithm proves very robust to the priors decreasing, the standard one starts to overestimate the number of sources as priors decrease. The two algorithms tend to have similar behaviors only when the priors grossly overestimate the true number of sources present in the searching area. When the priors grossly overestimate the true number of sources neither algorithm produces posteriors consistent with the true number of sources.
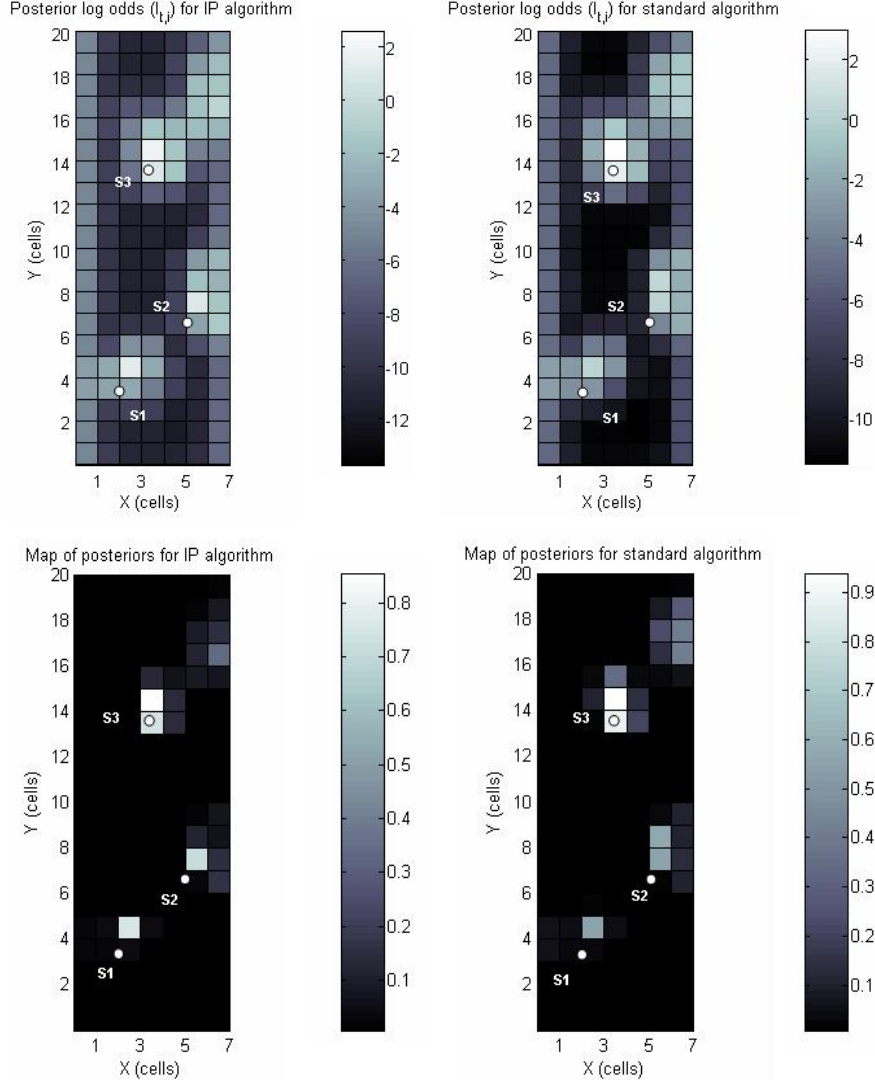
12

Fig. 9. Posteriors in log odds form (top) and posteriors (bottom) created from "experiment 1" dataset using IP and standard algorithm. The priors are equal consistent with 5 expected sources. The cell size is 30 cm x 30 cm. The maps appear less smooth, but always usable to localize the sources. The standard algorithm presents slightly overestimated posteriors.

One delicate issue for map creation in either algorithm is the decision of the threshold to discriminate if a measurement is a detection or not: a too low threshold could consider too many measurements to be detections while they could be due to the saturation of the environment not bringing useful information about sources position, while a too high one might generate maps that miss some sources. The threshold depends clearly on the quantity of odor present in the air and that is function of different factors: wind, sources dimensions/ejecting rates and time sources were free to eject before the robot starts its survey. However, to choose the threshold, firstly, we relied on the fact that sources had being ejecting for at least 10 minutes before the survey started giving a certain time to the gas to diffuse; secondly, on the fact that even if the odor sources are different in shapes, they are all of the same order of magnitude. The robustness of the chosen threshold has been investigated using also 0.6 μl/l and 1 μl/l as thresholds, and, once again while for IP the maps do not change significantly, for standard algorithm the posteriors of the produced maps are more influenced by the used threshold.

The comparison between IP and standard shows that IP method is much less sensitive than standard algorithm to both priors and thresholds changing.
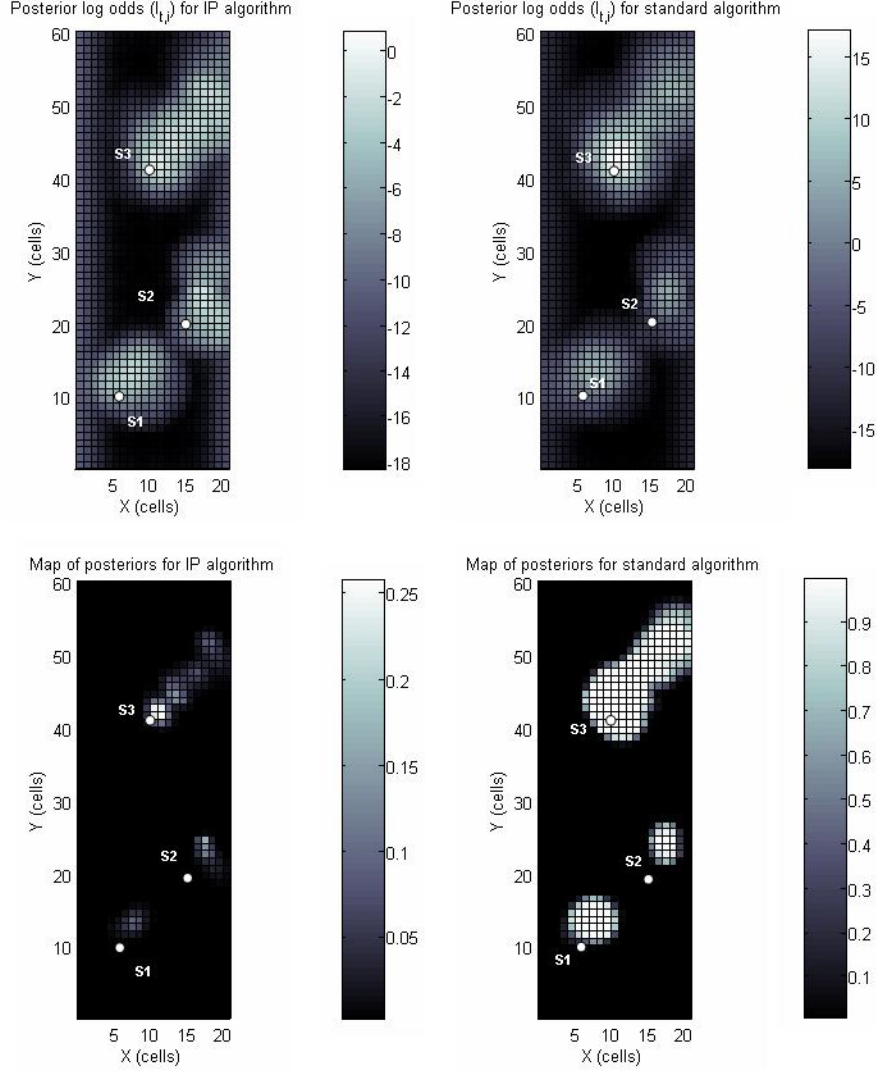
Fig. 10. Posteriors in log odds form (top) and posteriors (bottom) from "experiment 1" dataset using IP and standard algorithm. The priors are equal to 4 10^{-5} (prior underestimates the true number of sources). While IP continues to produce maps coherent with the true number of sources, the standard algorithm overestimates the posterior (*Note the different grey scales used*).

However, when the priors grossly overestimated the true number of sources the two algorithms are not able to build a map consistent with the true number of sources present in the searching area.

The maps created by IP algorithm appear to be suitable to map different sources ejecting a chemical. The probable positions produced by the maps, as told, are a bit to the "north-east" of the source due to a light airflow blowing towards that direction. To quantify the errors between true source position and estimated one, we report in Fig. 12 the averages on the experimental trials of the distances between real and estimated position of the sources. The probable source position for the different sources was selected as the center of mass of the probabilities of the three spots on the maps that present an increased probability to contain an occupied cell. The average of the errors was in all three cases lower than 50 cm and the S3 is very well localized (average error of about 18 cm). The relatively low standard deviations suggest that map creation with our method is quite insensitive to the gas distributions sensed in different experiments (this is also deducible from Figs. 7 and 8 where the maps are very similar). The method filters the information contained in the different gas acquisitions to estimate the most likely sources' positions.
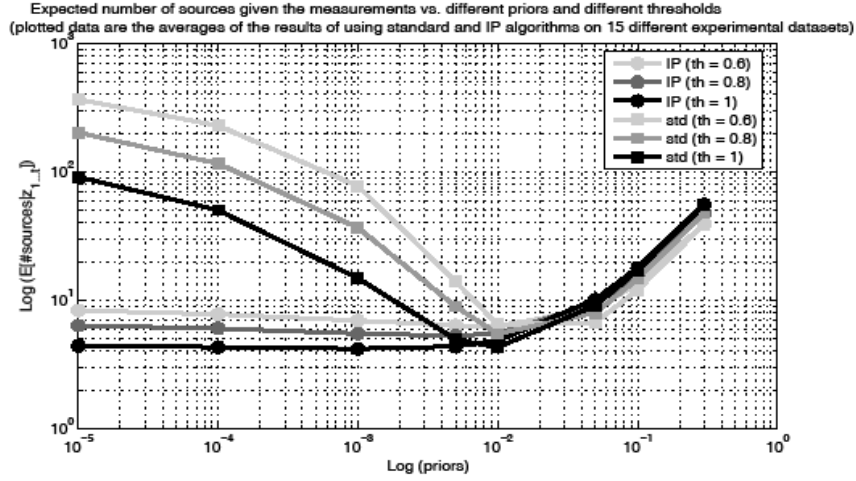
Fig. 11. Expected number of sources given the posteriors versus different used priors and different used detection threshold. The values are the averages of the results of applying IP and standard algorithm on 15 different experimental datasets. IP results are much less sensitive to both thresholds and priors variations, though both algorithms are incapable to produce a map consistent with the true number of sources when the priors grossly overestimate the true number of sources.
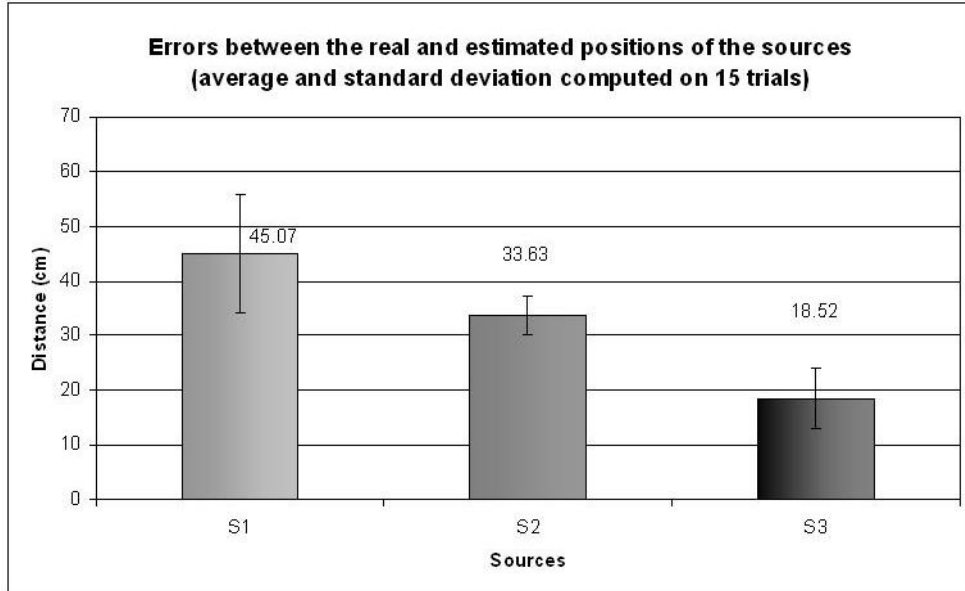

Fig. 12. Errors between estimated and real sources position. The reported values are the averages and the standard deviations of the 15 experimental performed trials.

It is interesting to remark that the proposed method can build reliable maps in a relatively short time and in an online manner. In our experiments an area of 6 m x 2.1 m was mapped using 64 measurements (each gas acquisition lasted 3 s). The final time for one trial is about 5.5 minutes. With non-detections, moreover, the explored areas that are not likely to contain a source are explicitly showed. This allows a hypothetical human operator not only to know in real-time where the sources likely are, but also to know areas that are candidates to be empty, so to make decisions while the robot is mapping the area without waiting for the end of the mapping process.

6. Conclusions and future works

In this paper, we address the problem to localize, with a mobile robot, multiple sources ejecting a chemical in an indoor environment with no strong airflow. These conditions represent a realistic domestic or industrial scenario. Without a strong airflow the task to localize a source gets harder because no information about the airflow can be used to infer the source direction, and because the low energy turbulent mixing produces a patchy and intermittent gas concentration. In our setting the presence of multiple sources ejecting a chemical is considered. The presence of multiple sources increases the difficulties of the problem: reactive/behaviour based algorithms offer no clear extension to environments with multiple sources. To overcome these problems we have described a map building based method aiming at creating a map of the environment: the survey area is divided into cells and each cell presents the probability to contain a source. Even if our

method does not explicitly localize the source (the robot moves on a predefined path), it can give useful information to a human operator: areas where sources are likely to be, areas that are probably empty and areas not explored are indicated.

To create this kind of map we use occupancy grid mapping methods. We show that in our particular setting (low prior with which the cells are initialized) the well known artifacts of traditional occupancy grid mapping techniques [32] caused by independently estimating the state of each cell, are amplified [15] producing maps with posteriors overestimating the true number of sources and hardly usable for our needs. To solve this, we adopt an algorithm (IP) that is an extension [15] to the traditional Bayesian grid algorithm. This IP algorithm is essentially based on a "revised prior" in order to build a linkage between previous measurements that resulted in detections of the gas and the more recent ones. A particular sensor model is assumed taking into account the intensities of the measured gas. Our method has been tested in a corridor of CRIM Lab and has been showed to create reliable maps quite insensitive to the specific dataset used.

Our method presents the following characteristics:

- it does not need to assume the presence of only one source to work (multiple sources were present in the corridor);
- it does not rely on information about the airflow to work;
- it can provide suitable maps after a short time (a 6 m x 2.1 m area mapped in about 5 minutes and a half);
- it can work in an online fashion.

All these features make it suitable to be used in real indoor monitoring task.

In future works we will study possible solutions in combining the use of this map building method with reactive algorithms. The maps should be useful to tell the robot where to explore for new sources and avoid to refind the previously localized sources.

REFERENCES

[1] J. Buhmann, W. Burgard, A.B. Cremers, D. Fox, T. Hofmann, F. Schneider, J. Strikos and S. Thrun, The mobile robot Rhino, AI Magazine, 16(1) (1995).

[2] J. Borenstein and Y. Koren, The vector field histogram – fast obstacle avoidance for mobile robots, IEEE Journal of Robotics and Automation, 7(3) (1991) 278-288.

[3] A. Elfes, Occupancy grid: a probabilistic framework for robot perception and navigation, PhD Thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, 1989.

[4] J.A. Farrell and J. Murlis and X. Long and W. Li and R.T. Cardé, Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes, Environmental Fluid Mechanics, Vol. 12 (2002) 143-169.

[5] J.A. Farrell, S. Pang and W. Li, Plume mapping using hidden Markov methods, IEEE Transactions on Systems, Man and Cybernetics Part B: Cybernetics, Vol. 33, No.6 (2003) 850-863.

[6] G. Ferri, E. Caselli, V. Mattoli, A. Mondini, B. Mazzolai and P. Dario, SPIRAL: a Novel Biologically-Inspired Algorithm for Gas/Odor Source Localization in an Indoor Environment with no Strong Airflow, Robotics and Autonomous Systems, Vol. 57, Issue 4, 30 (April 2009), 393-402.

[7] A.T. Hayes, A. Martinoli and R.M. Goodman, Distributed odor source localization, IEEE Sensors Journal, Vol.2, No.3 (2002) 260-273.

[8] H. Ishida, T. Nakamoto, T. Moriizumi, Remote sensing of gas/odor source location and concentration distribution using mobile systems, Sensors and Actuators B, 49 (1998), 52-57.

[9] H. Ishida, G. Nakayama, T. Nakamoto and T. Moriizumi, Controlling a gas/odor plume-tracking robot on transient responses of gas sensors. IEEE Sensors Journal, Vol.5, No.3 (2005) 537-545.

[10] H. Ishida, H. Tanaka, H. Taniguchi and T. Moriizumi, Mobile robot navigation using vision and olfaction to search for a gas/odor source, Autonomous Robots 20 (2006) 231-238.

[11] H. Ishida, Robotic systems for gas/odor source localization: gap between experiments and real-life situations, in: Proceedings of Workshop on Robotic Olfaction of IEEE International Conference on Robotics and Automation (ICRA 2007), Rome, Italy, 2007.

[12] S. Kazadi, R. Goodman, D. Tsikata, D. Green and H. Lin, An autonomous water vapor plume tracking robot using passive resistive polymer sensors, Autonomous Robots, 9 (2000) 175-188.

[13] K. Konolige, Improved occupancy grids for map building, Autonomous Robots, 4 (1997) 351-367.

[14] D. Kortenkamp, R.P. Bonasso and R. Murphy, AI-based mobile robots: case studies of successful robot systems, MIT Press: Cambridge, Massachusetts, 1998.

[15] M.V. Jakuba, Stochastic mapping for chemical plume source localization with application to autonomous hydrothermal vent discovery, PhD Thesis, WHOI-MIT, 2007.

[16] A. Lilienthal, D. Reiman and A. Zell, Gas source tracing with a mobile robot using an adapted moth strategy, Autonome Mobile Systeme (AMS) Karlsruhe, December 4-5 (2003) 150-160.

[17] A. Lilienthal and T. Duckett, Experimental analysis of smelling Braitenberg vehicles, in: Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003), Coimbra, Portugal, (2003) pp. 375-380.

[18] A. Lilienthal and T. Ducket, Building gas concentration gridmaps with a mobile robot, Robotics and Autonomous Systems, 48 (2004) 3-16.

[19] A. Lilienthal, A. Loutfi and T. Duckett, Airborne chemical sensing with mobile robots, Sensors 6 (2006) 1616-1678.

[20] C. Lytridis, E.E. Kadar and G.S. Virk, A systematic approach to the problem of odour source localization, Autonomous Robots, 20 (2006) 261-276.

[21] L. Marques, A. Martins and A. de Almeida, Environmental monitoring with mobile robots, in: Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2005), Edmonton, Canada, (2005) pp. 3624 -3629.

[22] H.P. Moravec, Sensor fusion in certainty grids for mobile robots, AI Magazine 9 (1988) 61-74.

[23] J. Murlis, J.S. Elkinton and R.T. Cardè, Odor plumes and how insects use them, Annu. Rev. Entomol., 37 (1992) 505-532.

[24] S. Pang and J.A. Farrell, Chemical plume source localization, IEEE Transactions on Systems, Man and Cybernetics Part B: Cybernetics, Vol.36, No.5 (2006) 1068-1080.

[25] A. H. Purnamadjaja and R.A. Russell, Congregation behaviour in a robot swarm using pheromone communication, in: Proceedings of the Australasian Conference on Robotics and Automation, Sydney, Australia (2005), pp. 1-7.

[26] P. Pyk, S. Bermudez i Badia, U. Bernardet, P. Knusel, M. Carlsson, J. Gu, E. Chanie, B.S. Hansson, T.C. Pearce and P.F.M.J. Verschure, An artificial moth: chemical source localization using a robot based neuronal model of moth optomotor anemotactic search, Autonomous Robots, 20 (2006) 197-213.

[27] R.A. Russell, A. Bab-Hadiashar, R.L. Scheffer and G. Wallace, A comparison of reactive chemotaxis algorithms, Robotics and Autonomous Systems, 45 (2003) 83- 97.

[28] B.I. Shraiman and B.I. Siggia, Scalar turbulence, Nature, Vol.405, No. 6787 (2000) 639-646.

[29] C. Stachniss, C. Plagemann and A. Lilienthal, Learning gas distribution models using sparse Gaussian process mixtures, Autonomous Robots, Volume 26 , Issue 2-3 (April 2009), 187-202.

[30] S. Thrun, Learning metric-topological maps for indoor mobile robot navigation, Artificial Intelligence, 99(1) (1998) 21-71.

[31] S. Thrun, Learning occupancy grid maps with forward sensor models, Autonomous Robots 15 (2003) 111-127.

[32] S. Thrun, W. Burgard and D. Fox, Probabilistic robotics, MIT Press, Cambridge, Massachusetts, 2005.

[33] H. P. Moravec and D. W. Cho, "A bayesian method for certainty grids," in Working notes of AAAI 1998 Spring Symposium on Robot Navigation, Stanford, CA, 1989, pp. 57–60.

Table 1. Pseudo code of standard Bayesian occupancy grid mapping algorithm.

Standard occupancy grid mapping $(\{l_{t-1,i}\}, z_t)$
for all cells m_i do
    if m_i is in the perceptual field of z_t then
        $l_{t,i} = l_{t-1,i} + \textbf{log\_inverse\_sensor\_model}\ (m_i, z_t) - l_0$
    else
        $l_{t,i} = l_{t-1,i}$
    endif
endfor
return $\{l_{t,i}\}$

where $\textbf{log\_inverse\_sensor\_model}(m_i, z_t)$ is the log odds form of the inverse sensor model, $\log \dfrac{p(m_i \mid z_t)}{1 - p(m_i \mid z_t)}$

Table 2. Pseudo code of IP algorithm.

IP algorihm $(\{l_{t-1,i}\}, \{P_i^t\}, \{z_t\})$
Time starts at instant $t = 1$

for all cells m_i do
    if m_i is in the perceptual field of z_t then
        if $z_t = 1$ (z_t results in a detection)

$$\tilde{P}_i^t = \frac{e^{l_{t-1,i}}}{1 + e^{l_{t-1,i}}}$$

$$l_{t,i} = \log \frac{1 - (1 - P_f^t)(1 - P_i^t)\prod_{s \neq i}(1 - P_s^t \tilde{P}_s^t)}{1 - (1 - P_f^t)\prod_{s \neq i}(1 - P_s^t \tilde{P}_s^t)} + l_{t-1,i}$$

        else ($z_t = 0$, z_t results in a non detection)

$$l_{t,i} = \log (1 - P_i^t) + l_{t-1,i}$$

        endif
    else
        $l_{t,i} = l_{t-1,i}$
    endif
endfor
return $\{l_{t,i}\}$

where $l_{0,i} = \log \dfrac{p(m_i)}{1 - p(m_i)}$. It is trivial from $\{l_{t,i}\}$ to recover $p(m_i \mid z_{1:t})$. In fact $p(m_i \mid z_{1:t}) = \dfrac{e^{l_{t,i}}}{1 + e^{l_{t,i}}}$.