A Rao-Blackwellisation Approach to GDM-SLAM – Integrating SLAM and Gas Distribution Mapping (GDM)

Achim J. Lilienthal* Amy Loutfi* Jose Luis Blanco† Cipriano Galindo† Javier Gonzalez†
*Dept. of Technology, AASS, Örebro University, S-70182 Örebro, Sweden
achim@lilienthals.de, amy.loutfi@tech.oru.se
†Dept. of System Engineering and Automation, University of Malaga, 29071 Malaga, Spain
{jlblanco|cipriano|jgonzalez}@ctima.uma.es

***Abstract*— In this paper we consider the problem of creating a two dimensional spatial representation of gas distribution with a mobile robot. In contrast to previous approaches to the problem of gas distribution mapping (GDM) we do not assume that the robot has perfect knowledge about its position. Instead we develop a probabilistic framework for simultaneous localisation and occupancy *and* gas distribution mapping (GDM-SLAM) that allows to account for the uncertainty about the robot's position when computing the gas distribution map. Considering the peculiarities of gas sensing in real-world environments, we show which dependencies in the posterior over occupancy and gas distribution maps can be neglected under certain practical assumptions. We develop a Rao-Blackwellised particle filter formulation of the GDM-SLAM problem that allows to plug in any algorithm to compute a gas distribution map from a sequence of gas sensor measurements and a known trajectory. In this paper we use the Kernel Based Gas Distribution Mapping (Kernel-GDM) method. As a first step towards outdoor gas distribution mapping we present results obtained in a large, uncontrolled, partly open indoor environment.**

I. Introduction

Creating a spatial representation of gas distribution is an important and challenging subproblem within the field of mobile olfaction. Gas distribution mapping (GDM) could be used to determine the exact location of gas sources or perhaps even more importantly, to determine areas of high concentration of a harmful gas. Hindered by the temporally fluctuating character of turbulent gas transport, and the fact that chemical gas sensors provide information only about the small volume their surface interacts with, it is virtually impossible to measure the instantaneous concentration field without using a dense grid of sensors. However, it is sufficient to know the time-constant structure of a gas distribution for many applications such as air quality monitoring and surveillance of industrial sites. Furthermore, by using mobile robots to map the gas distribution, a contaminated area could be examined in rescue missions in order to provide incident planning staff with information to prevent rescue workers from being harmed or killed due to explosions, asphyxiation or toxication.

The main contribution of this paper is the integration of gas distribution mapping into a probabilistic framework for SLAM, essentially introducing pose uncertainty into gas distribution mapping. All previous approaches to the problem of gas distribution mapping assume that the robot has perfect knowledge about its position at all times. However, a perfect estimate of the robot pose will not be available in real world applications. It is therefore necessary to consider the uncertainty of the robot pose estimates in order to represent the available information in a gas distribution map. Instead of restricting ourselves to pure localization, we consider the full SLAM problem to address the case where a real robot is moving in a previously unknown environment.

In probabilistic estimation theory applied to the SLAM problem, Bayesian filtering provides a grounded framework for estimating unobserved variables given only noisy observations. Popular approaches for implementation of Bayes filtering include Extended Kalman Filters, and Particle Filters, which both have been extensively used in robotics. In a Particle Filter approach each particle represents a hypothesis of the variables being estimated, in our case the robot path and the maps. Our approach consists of a Rao-Blackwellised particle filter where a motion model is used to predict a prior distribution of the robot pose. The observation model from the range scanner is then used to update the pose estimation. Finally, the range readings together with the gas sensor measurements are used to update the maps. In this way, we obtain a GDM which is consistent with the estimation of both the occupancy map and the robot path.

The rest of this paper is organized as follows. After a brief review on related works (Section II), we begin with a discussion of the particular issues of gas distribution mapping and a description of the gas distribution algorithm used in this work (Section III). In Section IV, we outline the method used for integrating simultaneous localization and occupancy mapping with gas distribution mapping. Next, the experimental setup including the olfactory robot is detailed and an example of a gas distribution map applied in a challenging uncontrolled environment is shown (Section V). Finally, we conclude with a discussion and suggestions for future work (Section VI).

II. Related Works

The problem of creating a gridmap that represents the distribution of an analyte gas is still a relatively new field to mobile olfaction. Some works have attempted to create a spatial representation of the gas concentration without the use of a mobile robot by taking simultaneous measurements with stationary sensors. In Ishida et al. [7], the time-averaged response of metal-oxide gas sensors over 5 minutes at 33 grid points distributed over an area of $2 \times 1m^2$ was used to characterise the experimental environment. With an increasing

area, however, establishing a dense grid of gas sensors would involve problems such as cost and a lack of flexibility. Furthermore, an array of metal oxide sensors would cause a severe disturbance to the gas distribution due to the convective flow created by the heaters built into these sensors [6].

In Hayes et al. [4] concentration measurements were acquired with a mobile robot and a representation of the gas distribution was created by two dimensional histograms whose bins contained the number of odour hits received as the robot performed a random walk behavior. An odour hit was registered when the sensed concentration exceeded a predefined threshold. This method requires perfectly even coverage of the environment and it is doubtful whether it could be performed in large unknown environments. Also only binary information is used to create the map and therefore much of the fine gradations of the concentration measurements is discarded.

All approaches to gas distribution mapping so far assume perfect knowledge about the robot pose. This is also true for the original version of the Kernel-GDM method by Lilienthal and Duckett [8], described in Section III, which is extended in this paper to address the GDM-SLAM problem. To our best knowledge the GDM-SLAM problem has not been considered by other authors so far.

III. Kernel Based Gas Distribution Mapping

The general gas distribution mapping problem given the robot trajectory x^t and the gas sensor measurements z^t_{gas} is to estimate the posterior probability distribution over gas distribution maps m_{gas}

$$p(m_{gas}|x^t, z^t_{gas}). \tag{1}$$

Due to fundamental differences between range sensing with a laser scanner and gas sensing with metal oxide sensors (which are the most widely used gas sensors in mobile robotic applications, see Section V-B) Bayesian estimation cannot be applied to the gas distribution mapping problem in the same way as to estimate an occupancy grid map.

The main differences are, first, that the sensor readings do not allow to derive the instantaneous concentration levels directly. Metal oxide gas sensors are known to recover slowly after the target gas is removed (15 to 70 seconds [1]) and therefore perform temporal integration implicitly. Sensor readings can be comparatively high although the instantaneous concentration level is actually close to zero if a high gas concentration was sensed previously. Second, a snapshot of the gas distribution at a given instant contains little information about the distribution at another time due to the chaotic nature of turbulent gas transport. Turbulence generally dominates the dispersal of gas. As a consequence, the instantaneous concentration field of a target gas released from a small static source is a chaotic distribution of intermittent patches with peak concentration values an order of magnitude higher compared to the time-averaged values [12]. Third, in contrast to a typical range-finder sensor, a single measurement from a gas sensor provides information about a very small area because it represents only the reactions at the sensor's surface (≈ 1 cm^2).

Altogether, it is futile to attempt to create a map of the instantaneous gas distribution with a mobile robot. Therefore, we instead consider the problem of estimating the posterior over *time-averaged* gas distribution maps m^{av}_{gas}

$$p(m^{av}_{gas}|x^t, z^t_{gas}). \tag{2}$$

Another consequence of the peculiarities of gas transport and gas sensing is that little information about the geometrical location of the robot can be obtained from gas sensor measurements. Compared to the observation likelihood for laser range scans, it is therefore possible to approximate the observation likelihood for gas sensor measurements by a constant value, provided that the laser scanner observes appropriate features. This approximation is used in Eq. 16 and visualised in Fig. 3.

In order to estimate a grid map that represents the time-averaged relative concentration of a detected gas, we use the Kernel-GDM method introduced by Lilienthal and Duckett [8]. The main idea is to interpret gas sensor measurements z^t_{gas} as noisy samples from a time-constant distribution. This implies that the gas distribution in fact exhibits time-constant structures, an assumption that is often fulfilled in unventilated and unpopulated indoor environments [15]. It is important to note that the noise is caused by the large fluctuations of the instantaneous gas distribution while the electronic noise on individual gas sensor readings is negligible [5].

The Kernel-GDM method compensates for the small overlap between single measurements by convolving the sensor readings with a two-dimensional Gaussian kernel. It has a notable analogy with the problem of estimating density functions using a Parzen window approach [11] with a Gaussian kernel. However, when creating the gas distribution map, we do not sample from the gas distribution directly. It is therefore necessary to make the assumption that the trajectory of the robot (respectively, the trajectory of the sensors) roughly covers the available space. As detailed in Section III, the Kernel-GDM method maintains two temporary grid maps obtained from spatial integration of the points of measurement convolved with the Gaussian kernel. One temporary grid map $M_{xz_{gas}}$ integrates the points of measurement weighted by the sensor measurements and the second temporary grid map M_x integrates the points of measurement without a weight assigned. The gas distribution m^{av}_{gas} is estimated from the grid map $M_{xz_{gas}}$ normalised to M_x, which corresponds to sampling from the (normalised) gas distribution if the sensor readings are considered as a measure of how many samples were drawn from the particular grid cell. Because of the normalisation to M_x, a perfectly even coverage of the inspected area is not required so that the robot trajectory not necessarily has to be customised for gas distribution mapping.

The Kernel-GDM method can cope to a certain degree with the temporal and spatial integration of successive readings that metal-oxide gas sensors perform implicitly due to their slow response and long recovery time [9]. In order to obtain a faithful representation of gas distribution despite the slow sensor dynamics ("memory effect"), the robot's path needs to fulfill the requirement that the directional component of the distortion due to the memory effect is averaged out. This can either be achieved approximately by random exploration or

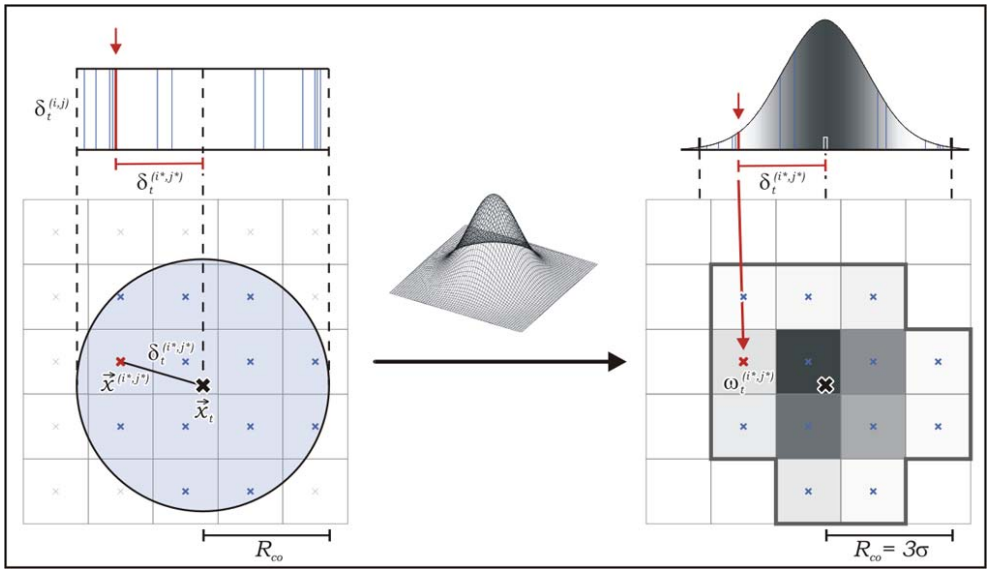

Fig. 1. Discretisation of the Gaussian weighting function onto the grid. Left side: for each grid cell within a cutoff radius R_{co} (represented by a circle) around the point of measurement $\vec{x}_t$, the displacement $\vec{\delta}_t^{(i,j)}$ is calculated. The corresponding distances are indicated for the 13 affected cells by the vertical lines drawn in the upper part of the figure. Right side: the weights $w_t^{(i,j)}$ are determined by evaluating the Gaussian function for the displacement values of the affected cells, using $\sigma = 1/3R_{co}$ in this example. The resulting weights are indicated by shadings of grey (dark shadings correspond to high weights).

in a strict manner by using a predefined path where the robot passes each point in the trajectory equally often from opposite directions. If the trajectory of the robot fulfills this requirement and sufficient time is given for the map to converge, the time-constant structures of the gas distribution will be represented faithfully in the gridmap, being slightly expanded and blurred but not shifted. The validity of the gridmaps produced by the Kernel-GDM algorithm therefore degrades gracefully with respect to the ratio between the time constant of the sensor dynamics and the speed of the robot. The algorithm introduces the kernel width σ as a selectable parameter, corresponding to the size of the region of extrapolation around each measurement. The value of σ has to be set large enough to obtain sufficient coverage according to the path of the robot. Conversely, this means that for a larger kernel width a faster convergence can be achieved while preserving less detail of the gas distribution in the map. Consequently, the selected value of the kernel width σ represents a trade-off between the need for sufficient coverage and the aim to preserve fine details of the mapped structures. Parameter selection and the impact of sensor dynamics are discussed in more detail in [9].

The Kernel-GDM Algorithm

The sensor readings are convolved using the univariate two dimensional Gaussian function

$$f(\vec{x}) = \frac{1}{2\pi\sigma^2} e^{-\frac{\vec{x}^2}{2\sigma^2}}. \tag{3}$$

Then, the following steps are performed:

- In the first step the normalised readings r_t are determined from the raw sensor readings R_t as

$$r_t = \frac{R_t - R_{min}}{R_{max} - R_{min}}, \tag{4}$$

using the minimum and maximum (R_{min}, R_{max}) value of a given sensor.

- Then, for each grid cell (i, j) within a cutoff radius R_{co}, around the point $\vec{x}_t$ where the measurement was taken at time t, the displacement $\vec{\delta}_t^{(i,j)}$ from the grid cell's centre $\vec{x}^{(i,j)}$ is calculated as

$$\vec{\delta}_t^{(i,j)} = \vec{x}^{(i,j)} - \vec{x}_t. \tag{5}$$

- Now the weighting $w_t^{(i,j)}$ for all the grid cells (i, j) is determined as

$$w_t^{(i,j)} = \begin{cases} f(\vec{\delta}_t^{(i,j)}) & : \quad \delta_t^{(i,j)} \leq R_{co} \\ 0 & : \quad \delta_t^{(i,j)} > R_{co} \end{cases} \tag{6}$$

- Next, two temporary values maintained per grid cell are updated with this weighting: the total sum of the weights

$$M_x : W_t^{(i,j)} = \sum_{t'}^{t} w_{t'}^{(i,j)}, \tag{7}$$

and the total sum of weighted readings

$$M_{xz_{gas}} : WR_t^{(i,j)} = \sum_{t'}^{t} r_{t'} w_{t'}^{(i,j)}. \tag{8}$$

- Finally, if the total sum of the weights $W_t^{(i,j)}$ exceeds the threshold value W_{min}, the value of the grid cell is set to

$$c_t^{(i,j)} = WR_t^{(i,j)} / W_t^{(i,j)} \quad : \quad W_t^{(i,j)} \geq W_{min}. \tag{9}$$

An example that shows how a single reading is convolved onto a 5 × 5 gridmap is given in Fig. 1.

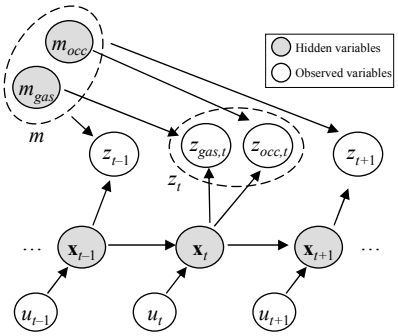

Fig. 2. The Dynamic Bayesian Network (DBN) of the SLAM problem for odour and occupancy grid mapping. Dependencies between variables are represented as directed arcs. Please note that this graphical model implies that the map m can be estimated from the observations z^t given a known robot path hypothesis x^t. Observations from the range scanner and the gas sensors are modelled as dependent on their respective maps only.

IV. Probabilistic GDM-SLAM

The general SLAM problem is stated as to simultaneously estimate the map m and the robot path $x^t = \{x_1, ..., x_t\}$, where each x_t represents the robot pose at time step t. Set out as a Bayesian filtering problem conditioned on the sequence of robot actions $u^t = \{u_1, ..., u_t\}$ and observations $z^t = \{z_1, ..., z_t\}$, the probability distribution to be estimated is:

$$p(x^t, m|u^t, z^t). \tag{10}$$

The graphical model for this problem is shown in Fig. 2 as a Dynamic Bayesian Network (DBN), where the hidden variables (represented by shaded circles) are to be estimated from the only known data, i.e. the sequence of actions and observations. The directed arcs in this graph represent statistical dependence between variables. Notice that the estimation of the map m is related to the inverse sensor model of the observations z_t, which in turn depend on the estimation of the robot path x^t. The inverse sensor model is used for estimating the map from observations since this implies to traverse the arrows in the DBN in the opposite direction to the actual dependence (Fig. 2). However, provided that maps can be analytically estimated given a robot path hypothesis, the complexity of estimating the distribution in (10) can be highly reduced by considering the factorization

$$p(x^t, m|u^t, z^t) = p(x^t|u^t, z^t)p(m|x^t, u^t, z^t) \tag{11}$$

and subsequently performing estimation of the first term only (the robot path) whereas analytically computing the second one (the maps). This technique to reduce the dimensionality of the estimation problem by exploiting the structure of the variables is called Rao-Blackwellised Particle Filter (RBPF) in Estimation Theory [2]. To estimate the robot path we represent its distribution by a set of M weighted particles,

$$\left\{x^{t,[i]}\right\}_{i=1...M} \sim p(x^t|u^t, z^t) \tag{12}$$

where associated weights $\omega_t^{[i]}$ account for the fact that the particles $x^{t,[i]}$ are not exactly distributed according to the density being estimated. Particle filtering for robotics and RBPFs are extensively discussed elsewhere [14, 2].

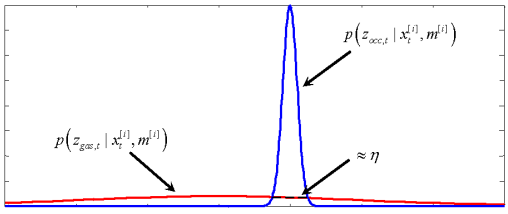

Fig. 3. One-dimensional example which illustrates that the likelihood of the gas sensor observation can be approximated by a constant value η under practical assumptions (discussed in the text). Consequently, the information provided by the gas sensors can be neglected for localization purposes.

In this work we consider that a map m comprises two different grid maps: the occupancy map m_{occ} and the gas distribution map m_{gas}. Assuming independency between them, we can estimate the map hypotheses $m_{occ}^{[i]}$ and $m_{gas}^{[i]}$ for each particle i separately. In a similar way, we define the observations z_t as the pair of observations $z_{occ,t}$ and $z_{gas,t}$ for the range scanner and the gas sensors, respectively. Notice that both observations are also conditionally independent given a robot path hypothesis, as illustrated in Fig. 2. If we consider the sequential Bayesian estimation of the robot path distribution in (12) under the Markov assumption we obtain the recursive formula

$$p(x_t|u^t, z^t) \propto \tag{13}$$
$$p(z_t|x_t, m) \int p(x_t|x_{t-1}, u_t)p(x_{t-1}|u^{t-1}, z^{t-1})dx_{t-1}.$$

Here two stochastic models are required: the observation model $p(z_t|x_t, m)$, and the robot motion model $p(x_t|x_{t-1}, u_t)$ which we obtain from odometry. In a RBPF, the latter distribution is not necessary in closed form, since we need only a mechanism to randomly draw samples from it. Assuming the standard proposal distribution [2], particles for each time step t are generated directly by sampling from the motion model:

$$x_t^{[i]} \sim p(x_t|x_{t-1}^{[i]}, u_t). \tag{14}$$

Accordingly, weights are updated through the observation likelihood function:

$$\omega_t^{[i]} \propto \omega_{t-1}^{[i]} p(z_t|x_t^{[i]}, m^{[i]}). \tag{15}$$

Intuitively, this means that those particles that better explain the current observations are assigned higher weights. If we take into account now the conditional independence between the pair of observations, we obtain

$$\begin{aligned} p(z_t|x_t^{[i]}, m^{[i]}) &= p(z_{gas,t}, z_{occ,t}|x_t^{[i]}, m_{gas}^{[i]}, m_{occ}^{[i]}) \\ &= p(z_{gas,t}|x_t^{[i]}, m_{gas}^{[i]})p(z_{occ,t}|x_t^{[i]}, m_{occ}^{[i]}) \\ &\approx \eta p(z_{occ,t}|x_t^{[i]}, m_{occ}^{[i]}). \end{aligned} \tag{16}$$

In the last step we approximate the observation likelihood for the gas sensors by a constant value η (similar to Grisetti et al. [3]), assuming that the observation likelihood of the range sensor dominates the product in (16). This is illustrated in Fig. 3 with a one-dimensional example. Given the precision of laser range scanners, this approximation is reasonable as long as appropriate features can be observed with the laser scanner. We do not consider situations where the range measurements

do not contain features, such as is in the middle of a wide open space, for example. This approximation means that we disregard the information provided by the gas sensors for updating the robot pose estimation. The gas sensor measurements are used, however, to update the gas distribution map.

Regarding the second term in (11), the distribution over maps, the occupancy grid $p(m_{occ}|x^{t,[i]}, z^t_{occ})$ for the i'th hypothesis is updated by well-known sensor integration methods ([10, 13]). The gas distribution map $p(m^{av}_{gas}|x^{t,[i]}, z^t_{gas})$ for the i'th particle is calculated by the Kernel-GDM algorithm detailed in Section III. Finally, the map estimates can be calculated as the marginal of the map distributions, taken over all the hypotheses of robot paths as

$$p(m|z^t) = \sum_i \omega^{[i]}_t p(m|x^{t,[i]}, z^t). \quad (17)$$

Since the expectation of a sum is the sum of the expectations,

$$E[p(...)] = E[\sum_i \omega^{[i]} p(...)] = \sum_i \omega^{[i]} E[p(...)], \quad (18)$$

we can compute the mean over map posteriors from the gas distribution maps obtained from the Kernel-GDM algorithm for each particle, which approximate the terms $E[p(m^{av}_{gas}|x^{t,[i]}, z^t_{gas})]$ in Eq. 18.

V. Experiments

A. Robot

Our experiments have been conducted using a service robot, called Sancho, which is intended to work within human environments as, for example, a conference or fair host (see Fig. 4a). It is constructed upon a pioneer 3DX mobile base whose structure has been devised to contain the sensorial system, including a radial laser scanner, a set of 10 infrared sensors, a colour motorized camera, and a pair of electronic noses, which are placed at a low position in the frontal part of the robot (see Fig. 4b). All devices of Sancho are managed by a Pentium IV laptop computer at 2.4GHz with wireless communication to remote servers or to the internet, enabling, for instance, remote users to command and to control the robot.

B. Gas Sensors

Located on the front of Sancho approximately 11 cm from the floor are two electronic noses based on TGS Figaro technology. Each e-nose consists of four TGS sensors (TGS 2600 (x2), 2620, 2602) placed in a circular formation on a plastic backing (see Fig. 4c). The sensors are fitted inside a retractable plastic tube sealed with a cpu fan that provides a constant airflow into the tube (see Fig. 4b). The two e-noses are separated at a distance of 14 cm (measured from the center of the circular backing).

Readings from the gas sensors are collected by an on-board data acquisition system with a sampling frequency of 1.25 Hz. Prior to experimentation, the sensor arrays for both e-noses were heated for approximately 30 minutes reaching temperatures between 300-500 °C, needed for proper operation. Metal oxide sensors exhibit some drawbacks worth noting. Namely the low selectivity, the comparatively high power consumption (caused by the heating device) and a weak durability. Furthermore, metal oxide sensors are subject to a long response time and an even longer decay time. However, this type of gas sensor is most often used for mobile noses because it is inexpensive, highly sensitive and relatively unaffected by changing environmental conditions like room temperature or humidity.

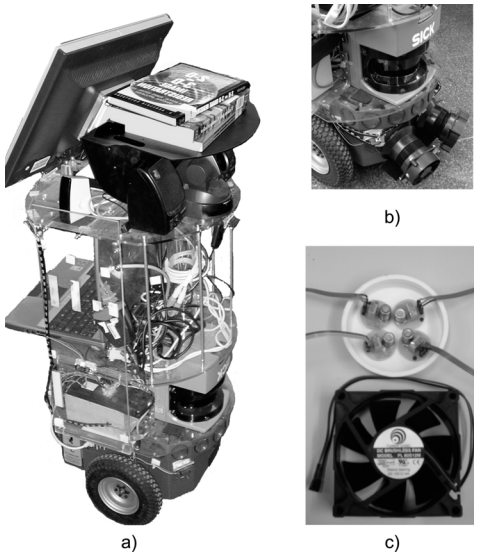

Fig. 4. Service robot Sancho. a) The original version of Sancho for delivery applications. b) Partial view of the robot focusing on the two electronic noses mounted on Sancho for our experiments. c) Each e-nose is composed of four gas sensors, a fan that provides a constant air flow, and a retractable plastic tube (not shown in the picture) that directs the air flow to the sensors.

C. Environment

Experiments were carried out within one of the wings of the Computer Science building at the University of Málaga (Spain). The testing scenario comprises two long corridors (one indoor and one outdoor) connected through two passages. Test results are presented from the runs conducted in the indoor portion of the corridor. The environment was in no way modified for the purpose of the experiment. Furthermore, people were occasionally present in the corridor, moving about and at times entering or closing doors.

An ethanol gas source was used and was contained in a cup approximately 6 cm in diameter and 5 cm high. The small size of the cup proved to be convenient since the robot was able to drive directly over the source. The source was prepared at a distance beyond the experimentation area, it was then covered and moved into position in the corridor approximately 30 minutes prior to experimentation. The cover was then removed just before an experimentation trial would begin.

D. Results

An implementation of the system was made by moving the robot at a speed of 5 cm/s in a spiral sweeping fashion. Fig. 5a shows the maximum likelihood path taken by the robot indicated by a solid black line. Fig. 5b shows the gas distribution map merged with the laser scan data. The source location in the figure is indicated by a circular ring. Here, different shadings of gray are used to indicate concentration values, where dark shading corresponds to low and light shading

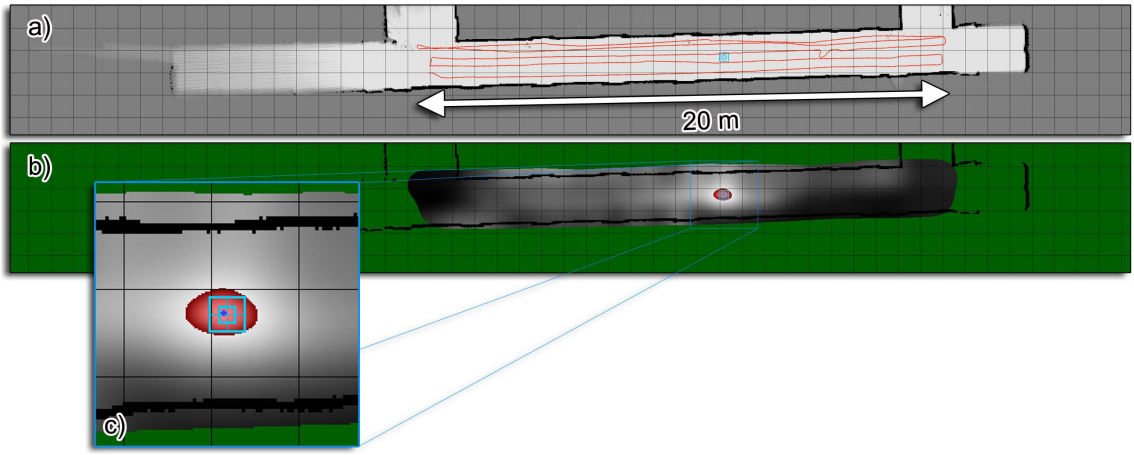

Fig. 5. (a) Laser Scan of the explored corridor and the most likely path taken by the robot. (b) The gas distribution map merged with the laser scan information. The actual source position is indicated by (light blue) concentric squares and a cross. Concentration values higher than 80% of the maximum are indicated with a second of range of dark to light shading in red. (c) Enlarged view of the gas distribution map at the source location.

to high relative concentration values. To better illustrate the variations in the measured concentration, a different shading color is used for cells containing concentration values higher than 80% of the maximum.

Note that the area indicating highest concentration values in the gas distribution map, corresponds to the source location, seen best in the close-up shown in Fig. 5c. While this correspondence provides a good indication that the gas distribution map corresponds to the true gas distribution it is not possible to conclude from this result that the suggested approach is generally able to represent the average gas distribution accurately. A thorough ground truth evaluation of the obtained gas distribution map is, however, very difficult due to the difficulty of measuring the gas distribution at the same time and the same height as with the sensors mounted on the robot. It is nonetheless remarkable that the gas distribution mapping algorithm was able to cope with the dimension of the environment (corridor of approx. 20 m x 2 m) and with the uncontrolled environment (an indoor wing connected to outdoor junctions).

VI. Conclusion

In this paper we present a probabilistic framework that integrates SLAM (with "M" = occupancy mapping) and gas distribution mapping. In contrast to all previous approaches, which assume that the robot has perfect knowledge about its position at all times, we propose an algorithm to create a two dimensional gas distribution gridmap that takes the uncertainty of the pose estimates into account. We develop a Rao-Blackwellised particle filter formulation of the GDM-SLAM problem and show which dependencies in the posterior can be neglected under certain practical assumptions.

Finally, we present results with the suggested approach obtained with a mobile robot equipped with gas sensors and a laser range scanner. To our knowledge, the experimental testbed used for evaluation is the largest environment for which a gas distribution map was created by a mobile robot so far and due to its partly open character this experiment can also be considered as a first step towards outdoor gas distribution mapping. This is an important contribution as mobile olfaction platforms move towards real application domains. Future work will primarily focus of establishing the validity of the proposed algorithm through experimentation with the current mobile platform.

References

[1] K. J. Albert and N. S. Lewis, "Cross Reactive Chemical Sensor Arrays," *Chem. Rev.*, vol. 100, pp. 2595–2626, 2000.

[2] A. Doucet, N. de Freitas, K. Murphy, and S. Russell, "Rao-Blackwellised particle filtering for dynamic Bayesian networks," *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pp. 176–183, 2000.

[3] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters," *IEEE Transactions on Robotics*, vol. 32, pp. 16–25, 2006.

[4] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Distributed Odor Source Localization," *IEEE Sensors*, vol. 2, no. 3, pp. 260–273, 2002.

[5] H. Ishida, A. Kobayashi, T. Nakamoto, and T. Moriizumi, "Three-Dimensional Odor Compass," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 2, pp. 251–257, April 1999.

[6] H. Ishida, M. Tsuruno, K. Yoshikawa, and T. Moriizumi, "Spherical Gas-Sensor Array for Three-Dimensional Plume Tracking," in *Proc. IEEE ICAR*, 2003, pp. 369–374.

[7] H. Ishida, T. Yamanaka, N. Kushida, T. Nakamoto, and T. Moriizumi, "Study of Real-Time Visualization of Gas/Odor Flow Images Using Gas Sensor Array," *Sensors and Actuators B*, vol. 65, pp. 14–16, 2000.

[8] A. J. Lilienthal and T. Duckett, "Creating Gas Concentration Gridmaps with a Mobile Robot," in *Proc. IEEE IROS*, 2003, pp. 118–123.

[9] ——, "Building Gas Concentration Gridmaps with a Mobile Robot," *Robotics and Autonomous Systems*, vol. 48, no. 1, pp. 3–16, August 2004.

[10] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," *Robotics and Automation. Proceedings. 1985 IEEE International Conference on*, vol. 2, 1985.

[11] E. Parzen, "On the estimation of a probability density function and mode," *Annals of Mathematical Statistics*, vol. 33, pp. 1065–1076, 1962. [Online]. Available: citeseer.ist.psu.edu/parzen62estimation.html

[12] P. J. W. Roberts and D. R. Webster, "Turbulent Diffusion," in *Environmental Fluid Mechanics - Theories and Application*, H. Shen, A. Cheng, K.-H. Wang, M. Teng, and C. Liu, Eds. ASCE Press, Reston, Virginia, 2002.

[13] S. Thrun, "Learning occupancy grids with forward models," *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, vol. 3, 2001.

[14] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. The MIT Press, September 2005.

[15] M. R. Wandel, A. J. Lilienthal, T. Duckett, U. Weimar, and A. Zell, "Gas Distribution in Unventilated Indoor Environments Inspected by a Mobile Robot," in *Proc. IEEE ICAR*, 2003, pp. 507–512.