

Integrating SLAM into Gas Distribution Mapping

Achim Lilienthal, Amy Loutfi, Jose Luis Blanco, Cipriano Galindo and Javier Gonzalez

Abstract—In this paper we consider the problem of creating a spatial representation of a gas distribution in an environment using a mobile robot equipped with gas sensors. The gas distribution mapping method used models the information content of a given measurement about the average concentration distribution with respect to the point of measurement. In this paper, we present an extension which can consider the uncertainty about the robot’s position in the gas distribution mapping. We present a preliminary result where a mobile robot equipped with gas sensors creates a map of a large indoor environment, using both spatial and olfactory information.

I. INTRODUCTION

Creating a spatial representation of a gas distribution in an environment is an important and challenging subproblem within the field of mobile olfaction. Gas distribution mapping (GDM) could be used as a means to determine the exact location of gas sources or perhaps even more importantly, determine areas of high concentrations of a harmful gas (that may not always be present at the source location). 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 probably impossible to measure the instantaneous concentration field without using a dense grid of sensors. Nonetheless, it is often 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, 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 contribution of this paper is a description of a gas distribution mapping algorithm which is able to take into account the uncertainty in the pose estimate. This is an important aspect to consider that is inherent to a real robot moving in an unknown environment. In probabilistic estimation theory applied to the SLAM problem, Bayesian filtering

A. Lilienthal and A. Loutfi are with the Center for Applied Autonomous Sensor Systems Örebro University, 701-82 Örebro, Sweden (web: www.aass.oru.se)

J. Blanco, C. Galindo and J.Gonzalez are with the System Engineering and Automation Department, University of Malaga, 29071 Malaga, Spain (email: {jlblanco,...}@ctima.uma.es).

This work was partly supported by the Spanish Government under research contract DPI2005-01391 and the Swedish National Science Agency, Vetenskapsradet.

provides a grounded framework for estimating unobserved variables given only noisy observations. Popular approaches for implementation of Bayes filtering include Extended Kalman Filter, and Particle Filters. Both of these approaches have been extensively used in robotics. For particle based representation, 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 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, which together with the gas sensor measurements is 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. We begin with a brief description of related works in the field of creating gas concentration gridmaps (Section 2). An overview of the gas distribution algorithm used in this work is then described (Section 3). In Section 4, we outline the method used for integrating simultaneous localization and mapping with gas distribution mapping. Next, the experimental setup of an olfactory robot is given and an example of a gas concentration map applied in an uncontrolled environment is shown. Finally, we conclude with a discussion of future works.

II. RELATED WORKS

The problem of creating a gridmap which represents the distribution of a gas source is still a relatively new field to mobile olfaction. Some works have attempted to create spatial representation of the gas concentration in an environment without the use of a mobile robot but rather by taking simultaneous measurements by multiple stationary sensors. In Ishida et al. [6], the time-averaged gas sensor response over 5 minutes at 33 grid points distributed over an area of $21m^2$ was used to characterise the experimental environment. With an increasing area, however, establishing a dense grid of gas sensors would involve an arbitrarily high number of fixed gas sensors, which poses 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 [5].

In Hayes et al. [3] gas measurements were acquired with a mobile robot and a representation of the gas distribution was

created by two dimensional histogram 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 and unknown environments. Furthermore, only binary information is used to create the map and therefore much of the fine gradations in the average concentration is discarded.

III. KERNEL BASED GAS DISTRIBUTION MAPPING

The general gas distribution mapping problem given the robot trajectory is

$$p(m_{gas}|x^t, z_{gas}^t). \quad (1)$$

Due to fundamental differences between range sensing with a laser scanner and gas sensing with metal oxide sensors 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. As a consequence, sensor readings can be comparatively high although the instantaneous concentration level is actually close to zero if a high gas concentration was sensed previously. In order to estimate the instantaneous concentration level from a sequence of sensor readings, one would need to model the interaction between the gas sensors and the (unobserved) times when the sensors were “hit” by patches of gas and also their respective (unobserved) strength. 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 that are generally an order of magnitude higher compared to the time-averaged values [11]. 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 ($\approx 1 \text{ cm}^2$).

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 *time-averaged* gas distribution map given the robot trajectory

$$p(m_{gas}^{av}|x^t, z_{gas}^t). \quad (2)$$

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 in particularly when considering a time-averaged gas distribution. The observation likelihood for the gas sensor measurements can thus be approximated by a constant value compared to the observation likelihood for the laser range scans. This property is used in Eq. (16) and visualised in Fig. 4.

In order to estimate a grid map that represents the time-averaged relative concentration of a detected gas, we use the kernel extrapolation gas distribution mapping method introduced by Lilienthal and Duckett [7]. The main idea is to interpret the gas sensor measurements z_t 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 [14]. 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 [4].

The gas distribution mapping method compensates for the small overlap between single measurements by convolving the sensor readings with a two-dimensional Gaussian kernel. The kernel can be seen as modelling the information content of a given measurement about the average concentration distribution with respect to the point of measurement. This information content decreases with increasing distance to the point of measurement.

There is also a notable analogy of the kernel extrapolation gas distribution mapping method with the problem of estimating density functions using a Parzen window approach [10] 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. The kernel extrapolation gas distribution mapping 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_{gas}^{av} 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 extrapolation gas distribution mapping method can cope to a certain degree with the temporal and spatial integration of successive readings that metal-oxide gas sen-

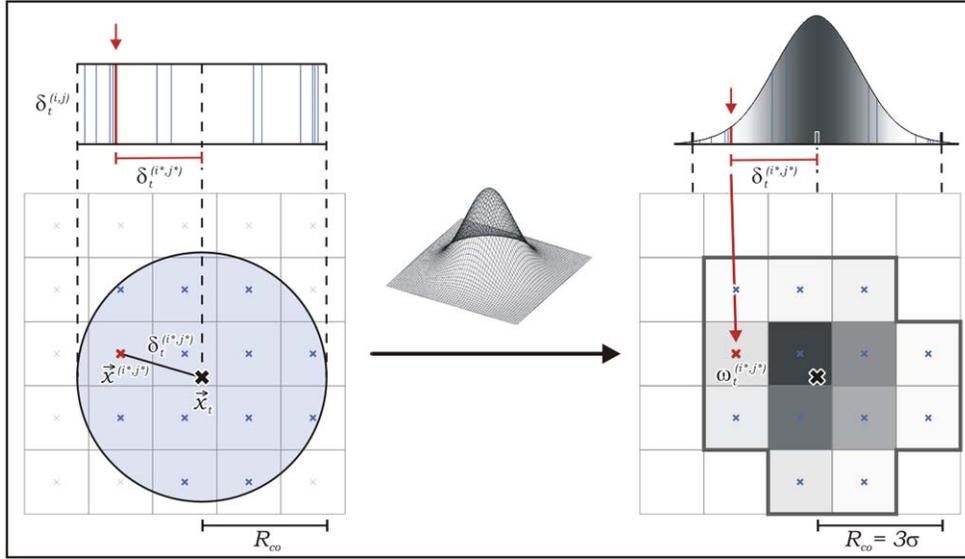


Fig. 1. Discretisation of the Gaussian weighting function onto the grid. Left side: For each grid cell within a cutoff radius R_{co} (represented by a circle) around the point of measurement \vec{x}_t , the displacement $\delta_t^{(i,j)}$ is calculated. The corresponding distances are indicated for the 13 affected cells by the vertical lines drawn in the upper part. Right side: According to the displacement, the weights $w_t^{(i,j)}$ are determined for all these grid cells (surrounded by a strong border) by applying a Gaussian function. As an example, a Gaussian with $\sigma = 1/3R_{co}$ is used. The weights are indicated by shadings of grey (dark shadings correspond to high weights).

sors perform implicitly due to their slow response and long recovery time [8]. 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 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-based extrapolation algorithm therefore degrades gracefully with respect to the ratio between the time constant of the sensor dynamics and the speed of the robot (i.e. the slow sensor dynamics). The algorithm introduces the kernel width σ as a selectable parameter, corresponding to the size of the region of extrapolation around each measurement. This parameter allows the user to decide between a faster or more accurate map building process. Its value 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

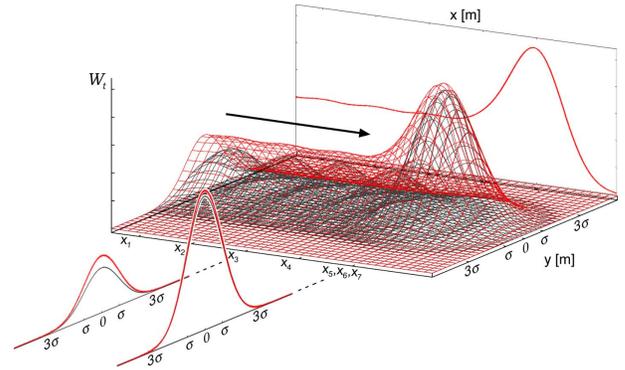


Fig. 2. Example of the information content regarding the average concentration distribution as modelled when calculating a kernel-based extrapolation gridmap for a hypothetical series of measurements. A sensor trajectory is assumed consisting of a constant velocity movement along a straight path and an immediate stop after the fifth time step. (i.e., measurements x_5 , x_6 and x_7 were all taken at the same physical location).

to preserve fine details of the mapped structures. Parameter selection and the impact of sensor dynamics are discussed in more detail in [8].

An example of a sum of Gaussian kernels for a set of measurement points is shown in Fig. 2, assuming a hypothetical sensor trajectory that consists of a constant velocity movement along a straight path and an immediate stop after the fifth time step. Referring to the interpretation of the kernel as a model of the information content, large values of the

sum of kernels correspond to locations where the certainty about the average concentration is high. Fig. 2 reveals that, assuming a suitable kernel width, a high certainty is assigned along the sensor trajectory and that the information content of the sensor readings is modelled as being approximately constant if the robot was driven with a constant, not too fast speed. This can be seen in the projection along the sensor path ($y = 0$) in Fig. 2. The roughly constant value reflects the fact that the sensor readings contain information about the average concentration along the path, which is approximately independent of the actual points of measurements due to the integration of successive readings. While the information content of the sensor readings regarding the average concentration is modelled as being high along the trajectory, it decreases quickly orthogonal to the path (and also to the front end of the path) where concentration values can be estimated by extrapolation only. Generally, the information content is modelled as being higher if the robot was driven more slowly, especially in cases where a number of successive measurements were performed at a particular spot. This can be seen in the projections orthogonal to the sensor path ($x = x_3$ and $x = x_5, x_6, x_7$) in Fig. 2. In the case where several measurements were performed on the spot, the concentration value calculated by averaging over the subsequent measurements represents a temporally integrated quantity that naturally contains more information about the average concentration at this particular location than a single measurement. Further, the temporally averaged gas sensor response value has a higher information content at adjacent places because in addition to a higher certainty about the average concentration the temporal mean also carries out some spatial integration due to the spatial fluctuation of gas. This appears as a slightly enlarged peak in the model of information content.

Step-by-Step Explanation of Kernel Based Gas Distribution Mapping

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

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

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

$$r_t = \frac{R_t - R_{min}}{R_{max} - R_{min}}, \quad (4)$$

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

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

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

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

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

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

and the total sum of weighted readings

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

- 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

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

An example that shows how a single reading is convolved onto a 5×5 gridmap is given in Fig. 1. First, thirteen cells are found to have a distance of less than the cutoff radius from the point of measurement (Fig. 1, left). These cells are indicated in the right side of Fig. 1 by a surrounding strong border. The weightings for these cells are then determined by evaluating the Gaussian function for the displacement values. In this example, the cutoff radius was chosen to be three times the width σ . The weights are represented by shadings of grey. Darker shadings indicate higher weights, which correspond to a stronger contribution of the measurement value r_t in the calculation of the average concentration value for a particular cell.

IV. INTEGRATING SLAM INTO GAS DISTRIBUTION MAPPING

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

$$p(x^t, m | u^t, z^t) \quad (10)$$

The graphical model for this problem is shown in Fig. 3 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

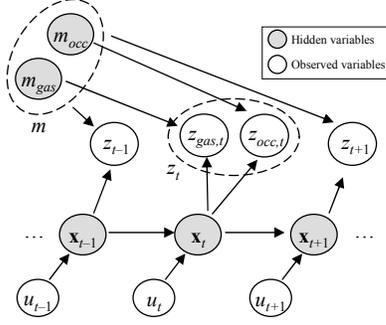


Fig. 3. The Dynamic Bayesian Network (DBN) of the SLAM problem for odour and occupancy grid mapping. Dependencies between variables are represented as directed arcs. From this graphical model it becomes clear that the map m can be estimated from the observations z^t given a known robot path hypothesis x^t . Observations from the range scanner and the gas sensors are modelled as dependent on their respective maps only.

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. 3). 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:

$$p(x^t, m | u^t, z^t) = p(x^t | u^t, z^t) p(m | x^t, u^t, z^t) \quad (11)$$

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,

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

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 ([13], [2]).

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. 3. If we consider the sequential Bayesian estimation of the robot path

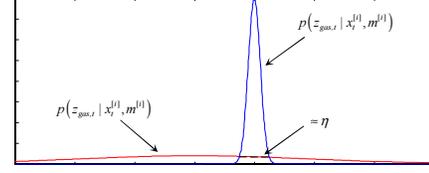


Fig. 4. Due to the precision of the range scanner, the information provided by the gas sensors can be safely neglected for localization purposes. As illustrated with this one-dimensional example, it is reasonable to approximate the likelihood of the gas sensor observation with a constant value η .

distribution in (12) under the Markov assumption we obtain the recursive formula:

$$p(x_t | u^t, z^t) \propto p(x_t | x_{t-1}, u_t) p(x_{t-1} | u^{t-1}, z^{t-1}) \quad (13)$$

Here two stochastic models are required: the observation model $p(z_t | x_t, m)$, and the robot motion model $p(x_t | x_{t-1}, u_t)$ (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:

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

Accordingly, weights are updated through the observation likelihood function:

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

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:

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

Notice that we approximate the observation likelihood for the gas sensors by a constant value η , which is a reasonable assumption if the precision of the range scanner sensor dominates the product in (16), as is clearly the case for laser range finders. This approximation is illustrated in Fig. 4 with a one-dimensional example. In practice, this approximation means that we disregard the information provided by the gas sensor for updating the robot pose estimation, although it will be used to update the gas concentration map.

Regarding the second term in (11), the distribution of maps, the occupancy grid $p(m_{occ}^{[i]} | x^{t,[i]}, z_{occ}^t)$ for the i 'th hypothesis is updated by well-known sensor integration methods ([9], [12]). For estimating the gas concentration map,

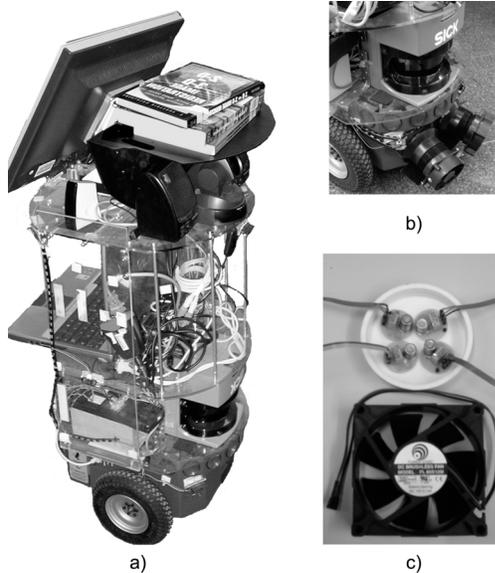


Fig. 5. Sancho, the service robot. a) The original version of Sancho for delivery applications. b) Partial view of the robot focusing on the two electronic noses mounted on Sancho for our experiments. c) Each e-nose is composed of four gas sensors, a fan that provides a constant air flow, and a retractable plastic tube (not shown in the picture) that directs the air flow to the sensors.

represented by $p(m_{gas}^{[i]} | x^{t,[i]}, z_{gas}^t)$ in the above formulation, we will assume in this work that the GDM introduced in the previous section provides an approximation of the *average concentration* for each cell in the grid. However, one could devise other approaches where Bayesian estimation is performed to estimate the exact distribution of the average concentration at each cell. This has not been addressed so far.

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. 5a). It is constructed upon a pioneer 3DX mobile base whose structure has been devised to contain the sensorial system. The sensorial systems includes a radial laser scanner, a set of 10 infrared sensors, a colour motorized camera, and a pair of electronic noses placed at a low position in the frontal part of the robot (see Fig. 5b). All devices of Sancho are managed by a Pentium IV laptop computer at 2.4GHz with wireless communication that connect Sancho 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). Four sensors are placed in a circular formation on a plastic backing (see Fig. 5 (c)). The sensors are then placed inside a retractable plastic tube sealed with a cpu fan that provides a constant airflow into the tube (see Fig. 5 (b)). 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 located on the frame of the robot and a sampling frequency of 1.25 Hz was used. Prior to experimentation, the sensor array 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.

C. Environment

Experiments were carried out within one of the wings of the Computer Science building at the University of Málaga (Spain). Fig. 6 depicts the testing scenario which 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, during an experimental trial, 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 in the indoor wing of the region indicated in Fig. 6. In Fig. 7 the path most likely taken by the robot was determined and is indicated by a solid black line. Fig. 8 shows the gas concentration 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

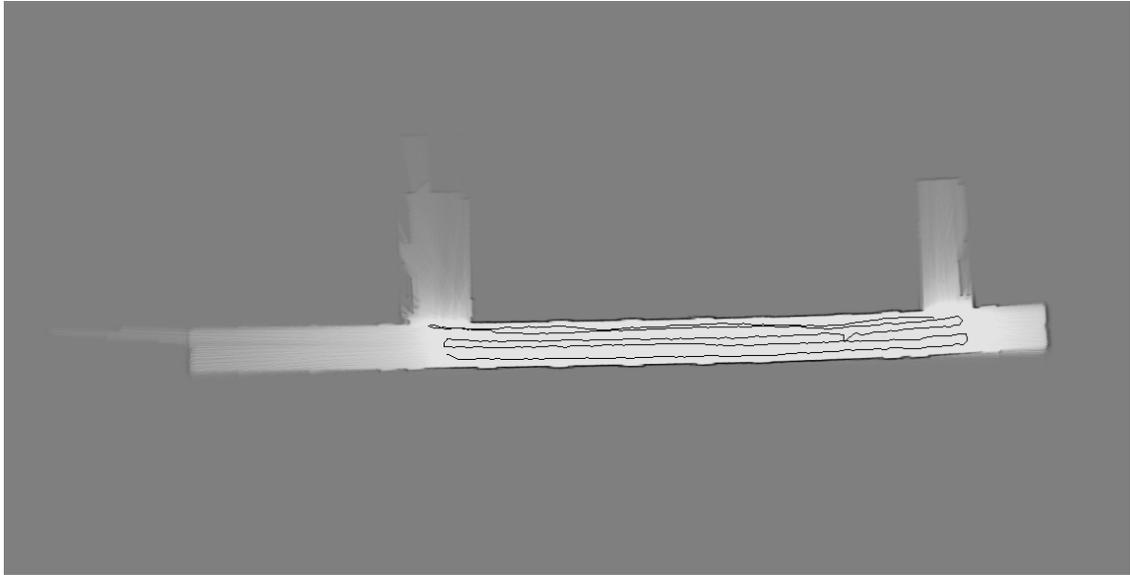


Fig. 7. Laser Scan of the explored corridor and the most likely path taken by the robot.

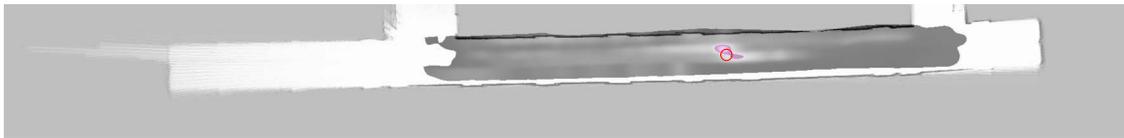


Fig. 8. The gas concentration map merged with the laser scan information. The actual source position is indicated by a (red) circular ring. Concentration values higher than 80% of the maximum are indicated with a second of range of dark to light shading in purple.

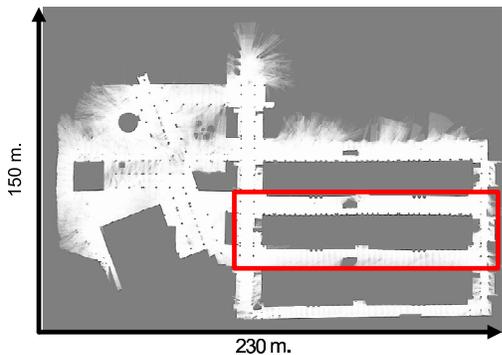


Fig. 6. Map of the Computer Science building at the University of Malaga (Spain). The region of the test scenario which contains an indoor corridor and an outdoor corridor is marked with a rectangle.



Fig. 9. Enlarged view of the gas concentration map at the source location.

and light shading 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. A few points are worth noting with regards to the figure. First, in the gas concentration map, the area indicating highest concentration values, corresponds to the source location, seen best in Fig. 9. This is a good indication of a reasonable performance of the

gas concentration mapping, although, it can not be said with certainty that areas of high concentration will necessarily correspond to source location due to the complex mechanism behind the gas transport (particularly in a complex environment). Nonetheless, having this correspondence implies that the gas concentration mapping was able to cope with both the large region of a corridor (approx. 20 m x 2 m) and in the uncontrolled environment (an indoor wing connected to outdoor junctions).

VI. CONCLUSION

In this work, we have presented a conceptual framework whereby a gas distribution mapping algorithm was integrated in a mobile robotic system. We have also presented an example of this framework in practice using a mobile robot equipped with a number of heterogeneous sensing modalities. In our work, the observation models from the range scanner and the gas sensors are used to obtain a gas distribution map

which is consistent with the estimation of the occupancy map and thus the robot path. To the authors knowledge, the experimental testbed used for evaluation is not only the largest used so far but also uncontrolled where no explicit effort was made to regulate ventilation or airflow. 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 experimental validation 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] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Distributed Odor Source Localization," *IEEE Sensors*, vol. 2, no. 3, pp. 260–273, 2002.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] A. J. Lilienthal and T. Duckett, "Creating Gas Concentration Gridmaps with a Mobile Robot," in *Proc. IEEE IROS*, 2003, pp. 118–123.
- [8] —, "Building Gas Concentration Gridmaps with a Mobile Robot," *Robotics and Autonomous Systems*, vol. 48, no. 1, pp. 3–16, August 2004.
- [9] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," *Robotics and Automation. Proceedings. 1985 IEEE International Conference on*, vol. 2, 1985.
- [10] 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
- [11] 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.
- [12] S. Thrun, "Learning occupancy grids with forward models," *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, vol. 3, 2001.
- [13] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. The MIT Press, September 2005.
- [14] 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.