

Building Gas Concentration Gridmaps with a Mobile Robot

Achim Lilienthal^a Tom Duckett^b

^a*University of Tübingen, WSI, D-72076 Tübingen, Germany.
lilien@informatik.uni-tuebingen.de*

^b*Center for Applied Autonomous Sensor Systems, Dept. of Technology, University
of Örebro, S-70182 Örebro, Sweden. Tom.Duckett@tech.oru.se*

Abstract

This paper addresses the problem of mapping the structure of a gas distribution by creating concentration gridmaps from the data collected by a mobile robot equipped with gas sensors. By contrast to metric gridmaps extracted from sonar or laser range scans, a single measurement from a gas sensor provides information about a comparatively small area. To overcome this problem, a mapping technique is introduced that uses a Gaussian weighting function to model the decreasing likelihood that a particular reading represents the true concentration with respect to the distance from the point of measurement. This method is evaluated in terms of its suitability regarding the slow response and recovery of the gas sensors, and experimental comparisons of different exploration strategies are presented. The stability of the mapped structures and the capability to use concentration gridmaps to locate a gas source are also discussed.

Key words: mobile nose, gas distribution mapping, gas source localisation.

1 Introduction

This paper addresses the problem of representing gas distribution in indoor environments by a mobile robot equipped with a gas sensitive system, comprising an on-board array of gas sensors. A new algorithm is presented for creating concentration gridmaps by combining the recorded gas sensor readings of the robot with location estimates. Intended applications include chemical mapping of hazardous waste sites and localisation of a gas source, especially in environments where it is impractical or uneconomical to install a fixed array of gas sensors. The proposed method does not require artificial ventilation of the environment, e.g., by imposing a strong, unidirectional airflow as in previous approaches for gas source localisation [12, 25, 7].

Gridmaps were originally introduced to mobile robotics in the 1980s as a means of creating maps using wide-angle measurements from ultrasonic range-finder sensors [4]. The basic idea is to represent the robot's environment by a grid of small cells. In a conventional gridmap, each cell contains a certainty value representing the belief that the corresponding area is occupied by an object. In a gas concentration gridmap, each cell contains an estimate of the relative concentration of a detected gas in that particular area of the environment. There are several problems in creating such a representation that are specific to mobile robots equipped with gas sensors.

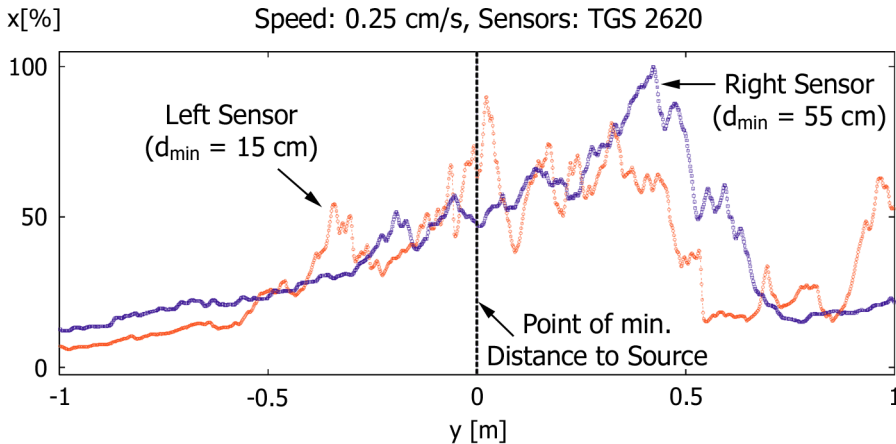


Fig. 1. Example of gas sensor readings recorded while the robot passed a gas source (ethanol) along a straight line at a low speed of 0.25 cm/s. The curve displays relative conductance values of two metal oxide sensors mounted on the left and right side of the robot with a separation of 40 cm.

A major problem is that the distribution of gas molecules in an environment that is not strongly ventilated tends to be dominated by turbulence and convection flow rather than diffusion, which is known to be a considerably slower transport mechanism for gases in general [20]. This typically results in a jagged

pattern of temporally fluctuating eddies [18, 22]. These effects are illustrated in Fig. 1, which shows typical sensor readings in the vicinity of a gas source (evaporating liquid ethanol). In this experiment, the robot passed the source along a straight line at low speed in order to measure the distribution of the analyte accurately. The curve in Fig. 1 indicates that the turbulent gas distribution creates many local concentration maxima, and that the absolute maximum is often located some distance from the actual location of the gas source if this source has been active for some time. In addition, the gas distribution varies over time.

Other problems relate to the gas sensors. In contrast to range-finder sensors such as sonar or laser, a single measurement from an electronic gas sensor provides information about a very small area. This problem is further complicated by the fact that the metal-oxide sensors typically used for this purpose do not provide an instantaneous measurement of the gas concentration. Rather, these sensors are affected by a long response time and an even longer recovery time. The time constants of rise and decay for the complete gas sensitive system (mobile nose) used here were estimated as $\tau_r \approx 1.8$ s and $\tau_d \approx 11.1$ s respectively [15]. Thus, considerable integration of successive measurements is carried out by the sensors themselves. The impact of this memory effect on concentration mapping is discussed in Section 3.3.

To overcome these problems, a mapping technique is introduced that integrates many gas measurements over an extended period of time. Spatial integration of the point measurements is carried out by using a Gaussian weighting function to extrapolate on the measurements, assuming a decreasing likelihood that a given measurement represents the true concentration with respect to the distance from the point of measurement. By integrating many measurements along the path of the robot, the underlying structure of the gas distribution can be separated from the transient variations due to turbulence. We show also that it is possible under certain limited conditions to use the grid cell with the maximum concentration value as an approximation to the location of the gas source, particularly when the shape of the distribution is roughly circular with a strong central peak.

In order to build complete concentration gridmaps, the path of the robot should roughly cover the entire space, although perfectly uniform exploration is not necessary. To increase spatial accuracy it is also advantageous to pass particular points from multiple directions. The method assumes that the pose of the mobile robot is known with high accuracy. In this paper, the location estimates required for map building were obtained from the external, vision-based absolute positioning system W-CAPS [16], which is briefly described in Section 2. However, the results are expected to apply to any mobile robot equipped with a suitably accurate positioning system, e.g., by carrying out simultaneous localisation and mapping with other sensor systems [3].

The rest of this paper is structured as follows. After a brief review of related work (Section 1.1), the experimental setup is described in Section 2. Next, the algorithm for creating gas concentration gridmaps is introduced (Section 3) and discussed in terms of parameter selection (Section 3.2) and its suitability regarding the slow response and recovery of the gas sensors (Section 3.3). Different data acquisition strategies are then discussed in Section 4 and an experimental comparison of the different exploration strategies is given in Section 5, followed by conclusions and suggestions for future work (Section 6).

1.1 Related Work

Most work on chemical sensing for mobile robots assumes an experimental setup that reduces the influence of turbulent transport by either minimising the source-to-sensor distance in *trail following* [26, 28, 27, 21] or assuming a strong airstream in the environment [9, 25, 23, 7]. A strong airstream means that additional information about the local wind speed and direction can be obtained from an anemometer. Thus strategies become feasible that utilise the instantaneous direction of flow as an estimate of the source direction [2] by combining gas searching behaviours with periods of upwind movement. Under the assumption of isotropic and homogeneous turbulence, and a unidirectional wind field with a constant average wind speed, it is further possible to model the time-averaged spread of gas [8]. The effect of turbulent air movement can be described in this case with a diffusion-like behaviour. Under these assumptions, the effect of turbulent air movement can be described with a diffusion-like behaviour ruled by an additional diffusion coefficient. The available wind measuring devices, however, are limited in their applicable range. With state-of-the-art anemometers based on the cooling of a heated wire [12], the bending of an artificial whisker [24] or the influence on the speed of a small rotating paddle [23], reliable readings can be obtained only for wind speeds in the order of at least 10 cm/s.

To the best knowledge of the authors, there have been only a few suggestions for creating spatial representations of gas distribution. A straightforward method to create a representation of the time-averaged concentration field is to measure the response over a prolonged time with a grid of gas sensors. This technique has been used on various occasions by Ishida and co-workers. The time-averaged gas sensor response over 5 minutes at 33 grid points distributed over an area of $2 \times 1 \text{ m}^2$ was used in [11], for example, 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 [13].




Fig. 2. Koala Robot with the Örebro Mark III mobile nose and the gas source used in the experiments. The small image in the top left corner shows a Figaro gas sensor used in the Mark III mobile nose.

Gas measurements acquired with a mobile robot were used by Hayes et al. to create a representation of the gas distribution by means of a two-dimensional histogram [7]. The histogram bins contained the number of “odour hits” received in the corresponding area while a random walk behaviour was performed. “Odour hits” were counted whenever the sensed concentration exceeded a fixed threshold. In addition to the dependency of the gas distribution map on the selected threshold, the problem with using only binary information from the gas sensors is that much useful information about fine gradations in the average concentration is discarded. It would also take a very long time to obtain statistically reliable results, and there is no extrapolation on the measurements apart from the quantisation into the histogram bins. So it is doubtful whether this approach would scale well to larger environments. A further disadvantage of this method is that it requires perfectly even coverage of the inspected area by the mobile robot.

2 Experimental Setup

2.1 Robot and Gas Sensors

The experiments were performed with a Koala mobile robot equipped with the Mark III mobile nose [15], comprising 6 tin oxide sensors manufactured by Figaro (see Fig. 2). This type of chemical sensor shows a decreasing resistance in the presence of reducing volatile chemicals in the surrounding air. In

consequence of the measurement principle, metal oxide sensors exhibit some drawbacks, including low selectivity, comparatively high power consumption (caused by the heating device) and weak durability. In addition, metal oxide sensors are subject to a long response time and an even longer recovery time [15]. However, this type of gas sensor is most often used on mobile robots because it is inexpensive, highly sensitive and relatively unaffected by changing environmental conditions such as room temperature and humidity. The sensors were placed in sets of three (of type TGS 2600, TGS 2610 and TGS 2620) inside two separate tubes containing a suction fan each. Multiple, redundant sensor types were used only to increase the robustness of the system (there was no attempt to discriminate different odours). Papst Fans (405F) were used to generate an airflow of 8 m³/h. The distance between the two sets of sensors was 40 cm.

2.2 Environment, Gas Source and Absolute Positioning System

All experiments were performed in a rectangular laboratory room at Örebro University (size 10.6 × 4.5 m²). The robot’s movement was restricted so that its centre was always located inside the central region where precise and reliable position information is available. The air conditioning system in the room was deactivated in order to eliminate the possibility of a dominant constant airflow.

To emulate a typical task for an inspection robot, a gas source was chosen to imitate a leaking tank. This was realised by placing a paper cup filled with ethanol on a support in a bowl with a perimeter of 12 cm (see Fig. 2). The ethanol dripped through a hole in the cup into the bowl at a rate of approximately 50 ml/h. Ethanol was used because it is non-toxic and easily detectable by the tin oxide sensors.

To record the position of the robot, the vision-based absolute positioning system W-CAPS [16] was applied, which tracks a distinctly coloured object mounted on top of the robot (the cardboard “hat” shown in Fig. 2). The positioning system uses four Philips PCVC 740K web-cameras mounted at fixed positions with a resolution of 320 × 240 pixels to triangulate the (x,y) position of the centre of the colour blob. By combining up to 6 single position estimates, it provides centimeter level accuracy. Fig. 3 shows the camera positions and the respective fields of view. The graded shadings indicate the number of cameras that can sense each part of the environment. The heading ϑ of the robot can be estimated from the tangent to the robot’s path while the robot is moving at non-zero speed. By fusing these estimates with the values provided by odometry, an accurate estimate of the heading is obtained that is not subject to long term drift.




Fig. 3. Testbed area for the gas concentration mapping experiments. The figure shows a floor plan of the laboratory room, and the fields of view for each of the four cameras that were used to track the robot’s position, shaded according to the number of cameras that can sense a particular region.

3 Gas Concentration Gridmaps

This section presents the method for creating gas concentration gridmaps from a sequence of sensor measurements collected by a mobile robot.

In order to create gas concentration gridmaps, the cells have to be updated multiple times. Gas sensor readings represent only the concentration at the very small area of the sensor’s surface ($\approx 1 \text{ cm}^2$). Nevertheless these readings contain information about a larger area, for two reasons. First, despite the jagged, fluctuating nature of gas distribution [18, 22], it is reasonable to assume that the gas concentration in the vicinity of the point of measurement does not change drastically because of the smoothness of the time-constant structures. Second, the metal-oxide gas sensors perform temporal integration of successive readings implicitly due to their slow response and recovery time. Thus spatial information is integrated along the path driven by the robot.

Note that errors in measurement are not modelled by the mapping algorithm because the electronic noise on individual gas sensor readings is negligible [10]. Rather, the algorithm models the decreasing amount of information contained in single sensor readings about the average gas concentration in the vicinity of the point of measurement. By applying this model, the algorithm is able to determine the spatial structure of the average gas distribution with limited distortion (discussed in Section 3.3) from a series of readings, which do not represent the gas concentration directly due to the response characteristics of the sensors. Individual sensor readings, however, are assumed to represent the actual physical effect measured (the rate of combustion reactions) exactly.

3.1 The Algorithm

The sensor readings are convolved using the radially symmetric two dimensional Gaussian function

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

This weighting function indicates the likelihood that the measurement represents the average concentration at a given distance from the point of measurement. In detail the following steps are performed:

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

$$r_t = \frac{R_t - R_{min}}{R_{max} - R_{min}}, \quad (2)$$

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

$$\vec{\delta}_t^{(i,j)} = \vec{x}^{(i,j)} - \vec{x}_t. \quad (3)$$

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

$$w_t^{(i,j)} = \begin{cases} f(\vec{\delta}_t^{(i,j)}) & : \delta_t^{(i,j)} \leq R_{co} \\ 0 & : \delta_t^{(i,j)} > R_{co} \end{cases} \quad (4)$$

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

$$W_t^{(i,j)} = \sum_{t'}^t w_{t'}^{(i,j)}, \quad (5)$$

and the total sum of weighted readings

$$WR_t^{(i,j)} = \sum_{t'}^t r_{t'} w_{t'}^{(i,j)}. \quad (6)$$

- 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

$$c_t^{(i,j)} = WR_t^{(i,j)} / W_t^{(i,j)} \quad : \quad W_t^{(i,j)} \geq W_{min}. \quad (7)$$




Fig. 4. Discretisation of the Gaussian weighting function onto the grid.

An example that shows how a single reading is convolved onto a 5×5 gridmap is given in Figure 4. First, thirteen cells are found to have a distance of less than the cutoff radius from the point of measurement (Fig. 4, left). These cells are indicated in the right side of Fig. 4 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.

3.2 Parameter Selection

The concentration mapping algorithm requires three parameters: the width σ of the Gaussian function, the cutoff radius R_{co} and the threshold W_{min} . While the actual values of R_{co} and W_{min} do not have a strong influence on the resulting concentration map, the parameter σ is critical. A value of σ has to be chosen that is high enough to satisfy the requirement for sufficient extrapolation on the gas concentration measurements, but low enough to preserve the fine details of the mapped structures. In this paper, parameter values of $\sigma = 15$ cm, $R_{co} = 3\sigma$ and $W_{min} = 10.0 \times (\text{number of sensors})$ were chosen based on the results of an initial set of experiments.

Fig. 5 shows the sum of the Gaussian weighting functions for a sequence of steps along a straight line (with a step width of 2σ). The figure shows the last four steps plus the current one indicated by an arrow. It can be seen that the




Fig. 5. Sum of the Gaussian weighting functions for a sequence of steps along a straight line.

resulting sum W_t models the information content in a way that assigns the strongest weight along the sensor trajectory. Here, the certainty is particularly high because the gas concentration at the corresponding locations actually contributed to the measurements. Due to the memory effect of the metal oxide sensors, the sensor readings represent a low-pass filtered concentration value integrated along the path driven.

3.3 Impact of Sensor Dynamics

As a consequence of the memory effect of the gas sensors, the mapped values show asymmetrically blurred edges and a slightly shifted centre of the area of maximum concentration compared to the real distribution. This effect can be seen in Fig. 6, which shows how a rectangular step pulse would be mapped by the gas concentration mapping algorithm. In Fig. 6(a) the real distribution can be seen, which is a step pulse with an assumed duration of $\Delta t = 10$ s. The response of the gas-sensitive system is shown in Fig. 6(b). This curve was calculated using a first order sensor model and the parameters $\tau_r \approx 1.8$ s and $\tau_d \approx 11.1$ s for the Örebro Mark III mobile nose, as calculated in [15]. Fig. 6(c) shows the Gaussian weighting functions multiplied by the corresponding sensor readings. The samples were assumed to be recorded at a rate of 2 Hz and a width of $\sigma = 1$ s was used for the Gaussian weighting function (see Eq. 1). This corresponds to a distance of $\sigma = 5$ cm, if a situation is considered where a robot drives with a constant velocity of 5 cm/s through a 50 cm wide area of constant concentration. Finally, the normalised curve of the mapped values is depicted in Fig. 6(d), calculated according to Eq. 7.

Comparing the real distribution with the mapped values in Fig. 6(d), the asymmetrical shift as well as the blurring effect can be seen. This corruption is, however, not critical. Due to the low speed of the robot, which never exceeded 5 cm/s during the experiments presented here, the expected shift would be

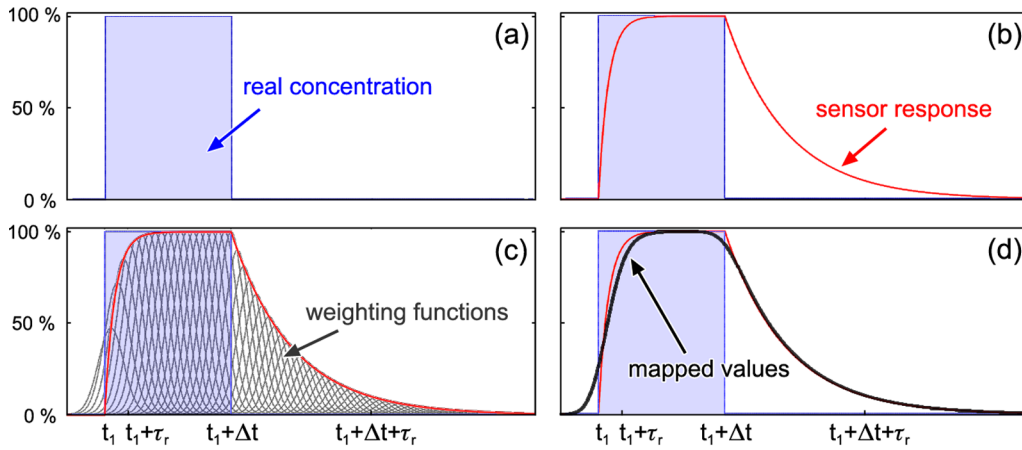


Fig. 6. Mapping of a rectangular step pulse. The figure shows (a) a simulated step stimulus varying between the minimum and maximum concentration levels, (b) the sensor response as calculated for the Örebro Mark III mobile nose, (c) the Gaussian weighting functions multiplied by the corresponding sensor readings, and (d) the resulting curve of the mapped values.

in the order of 10 cm at most. This effect would be even smaller for smooth distributions, which the metal-oxide sensors can follow more closely than a step-like one. In addition, the directional component of both effects would be averaged out if the robot passed a given point from different directions. If this condition is fulfilled, then the positions of concentration maxima should be represented faithfully by the described mapping process. The remaining effect of the blurred edges should be tolerable as long as the time-constant structures can be distinguished in the maps created.

4 Data Acquisition Strategy

This section describes the different exploration strategies used by the mobile robot to collect the sensor data for concentration mapping in our experiments.

The actual trajectory along which the sensor readings were recorded is expected to have a minor influence on the resulting concentration map. While uniform exploration is not necessary, it is required that the trajectory roughly covers the available space. In order to obtain better accuracy, it is further advantageous to pass particular points from multiple directions (see Section 3.3). This can be achieved in a straightforward way by driving the robot along a *pre-defined path* that complies with these requirements. Two different predefined trajectories, a rectangular spiral and a sweeping movement (see Sections 4.1 and 4.2) were tested here. On the other hand, it might be not desirable to specify a fixed trajectory if a robot is not exclusively used for gas distribution mapping. Instead, it might be very useful for some applications (e.g., search

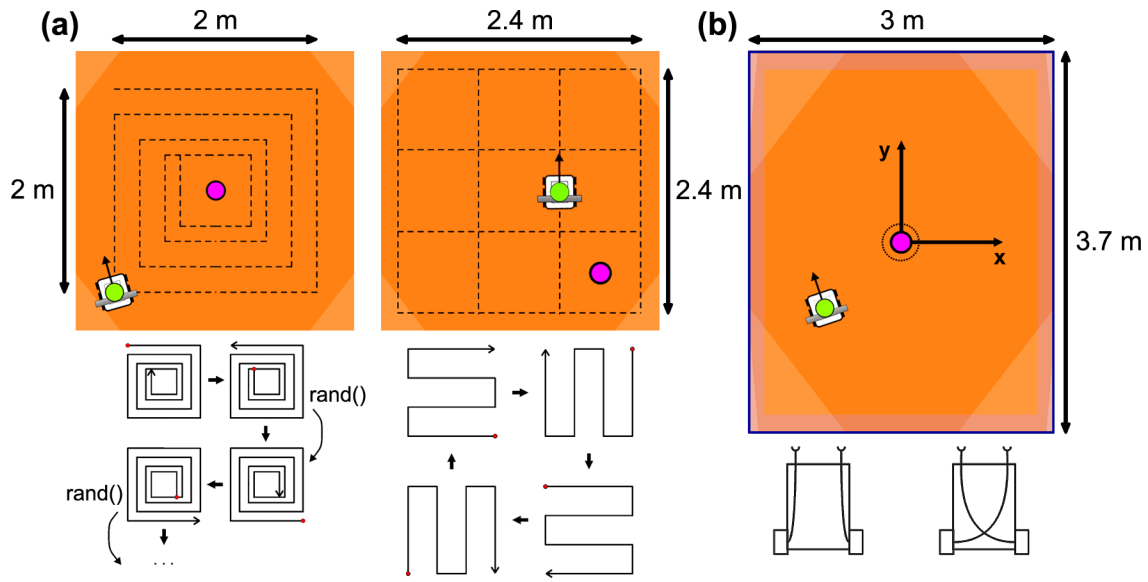


Fig. 7. Different data acquisition strategies performed for concentration mapping. (a) Outline of the predefined paths tested. (b) Sketch of the setup used for reactive gas source tracing experiments.

and rescue operations) if gas sensors could easily be added to existing mobile robots without the need to severely modify the behaviour. The same applies naturally if the target of a searching task is a gas source. Consequently, the concentration mapping algorithm was also tested with data obtained using *reactive gas source tracing* strategies (see Section 4.3).

4.1 Predefined Path – Rectangular Spiral

In this strategy, the robot followed a sequence of rectangular spirals around the location of the gas source. This path is sketched in Fig. 7(a). It consists of a sequence of inward and outward movements with a minimal distance to the centre of the source of 1 m, 0.75 m, 0.5 m, 0.35 m on the subsequent windings of the path. After an inward spiral is completed, the robot turns by 180° and drives back to the starting point of the inward spiral along the same path. Thus, points are passed equally often from opposite directions. In order to reduce self-induced disturbance of the gas distribution, the robot was driven at a constant speed of 5 cm/s along the straight lines. A constant speed was also found to enhance the gas source localisation capability [19, 5]. At the corners, the robot was rotated slowly ($10^\circ/\text{s}$) in order to minimise additional disturbance. One complete cycle including an inward and an outward spiral lasted about 25 minutes. These cycles were repeated with a randomly chosen starting corner and direction at the start of each trial. Note that the inspected area is not explored uniformly because of the need to avoid collisions with the beaker in the centre.

4.2 Predefined Path – Sweeping Movement

In the second strategy, the robot performs a sweeping movement that encloses nine squares ($80\text{ cm} \times 80\text{ cm}$), as shown in Fig. 7 (a), providing nine possible locations in which to place the gas source (one is indicated in the figure). The sweeping movement was implemented as a trajectory consisting of the four segments shown in Fig. 7, which were executed repeatedly in the given sequence. Again, a constant speed of 5 cm/s was applied along the straight lines and the robot was rotated with a speed of $10^\circ/\text{s}$ at the corners. In contrast to the previous strategy, points on the path are not traversed equally often from opposite directions. Opposite edges of the nine squares, however, are passed from opposite directions in the course of each cycle. The trajectory also does not provide completely uniform exploration, because the outer edge of the explored area is traversed more often than the inner parts. In contrast to the spiral movement, however, the location of the source is not especially distinguished by the symmetry of the path.

4.3 Reactive Gas Source Tracing

As a sub-task of the general gas source localisation problem, gas source tracing is supposed to guide a gas-sensitive mobile system towards a source using its own measurements of the gas distribution. Under the assumption of a concentration field that exhibits smooth gradients, thus neglecting the effects of turbulence, gas source tracing can be accomplished by reactive gradient-following. The reactive strategies used in this work were based on Braitenberg vehicles [1], i.e., a direct sensor-motor coupling using the stereo architecture of the mobile nose. Two Braitenberg-type strategies were implemented, corresponding to uncrossed and crossed sensor-motor connections (see [17] for full details). Because inhibitory connections were used, maximum wheel speed results if the sensed concentration is low, which in turn implements a simple sort of exploration behaviour. Thus, the robot turns towards higher concentrations with uncrossed connections (“exploration and hill climbing”) while the robot turns away from them with crossed connections (“exploration and concentration peak avoidance”).

These experiments were carried out in a restricted area where the boundaries were implemented as a repulsive artificial potential field [14] that starts to be effective at a distance of 20 cm from the border. Both the inspected areas and the border region where the repulsive potential was effective are indicated in Fig. 7(b).

In contrast to the previous strategies, the trajectory is of course not known in

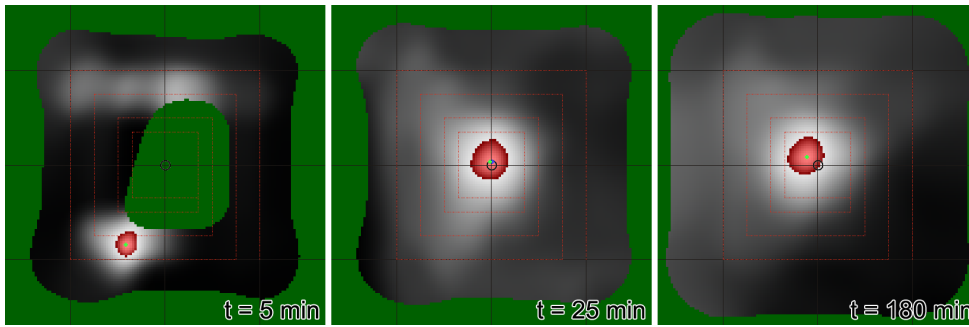


Fig. 8. Evolution of a concentration gridmap created by following a predefined path (rectangular spiral).

advance in the case of Braitenberg-type strategies. It is thus not guaranteed that the path complies with the mentioned conditions for concentration mapping. In reaction to a certain gas distribution, the resulting trajectory might not cover all regions and the robot might always approach some points from the same direction. It is also known that the “bouncing” behaviour of the robot due to the artificial potential fields (i.e., specular reflection) does not guarantee uniform coverage of an environment [6].

5 Results

5.1 Stability of the Mapped Structures

The concentration mapping algorithm was tested extensively using sensor data acquired with the Mark III mobile nose over a total of almost 70 hours of experiments and more than 5 kilometres of travel.

Due to the local character of single gas sensor measurements, it takes some time to build concentration gridmaps. In addition to spatial coverage, a certain amount of temporal averaging is also necessary in order to represent the time-constant structure of the gas distribution. Some examples showing the evolution of the gridmaps during these experiments can be seen in Figs. 8 and 9. Concentration values are indicated by shadings of grey (dark shadings correspond to low concentrations, and light shadings to high concentration), while the values higher than 90% of the maximum concentration value are shown with different shadings.

A relatively stable gas distribution was observed in almost all experiments. In order to evaluate the stability of the mapped distribution, we measured the distance of the grid cell with the maximum concentration value from the centre of the gas source. When this distance converged to a steady value, we




Fig. 9. Evolution of a concentration gridmap created with a reactive gas source tracing strategy (Braitenberg vehicle with uncrossed inhibitory connections).

considered the mapped distribution to be stable. In the experiments where the robot followed a predefined path, it was found that the mapped distribution stabilised when the robot had passed each point in the vicinity of the gas source at least two times from opposite directions. Thus, it took approximately 25 minutes in the trials where a rectangular spiral was driven, while stabilisation was observed after approximately 10 minutes in the sweeping trials. Rapid stabilisation after approximately 10 minutes was also observed with the reactive gas source tracing strategies due to the fact that areas with high concentration were actively explored. However, the location of the detected maximum changed more often because the Braitenberg vehicle has a tendency to get stuck in local concentration maxima.

In order to further demonstrate the consistency of the gridmaps over time, the data from one run were divided into three separate parts, as shown in Fig. 10, and separate gridmaps were created from the three different data sets. The results show that the gridmaps obtained are qualitatively very similar, and that the distance from the source to the concentration maximum reached the same level in all cases after roughly 25 minutes.

5.2 Suitability for Gas Source Localisation

In the case of a gas distribution controlled purely by diffusion, the location of a static, constantly evaporating gas source would correspond to the concentration maximum in the gridmap. This is rarely the case in real world environments due the effects of turbulence. However, for the indoor environment investigated, it was found that the concentration maximum could provide an approximate estimate of the source location in many cases, particularly when the shape of the distribution is roughly circular with a strong central peak.

Table 1 gives a summary of the results obtained in all trials, giving the percentage of time steps for which the distance from the source to the concentra-

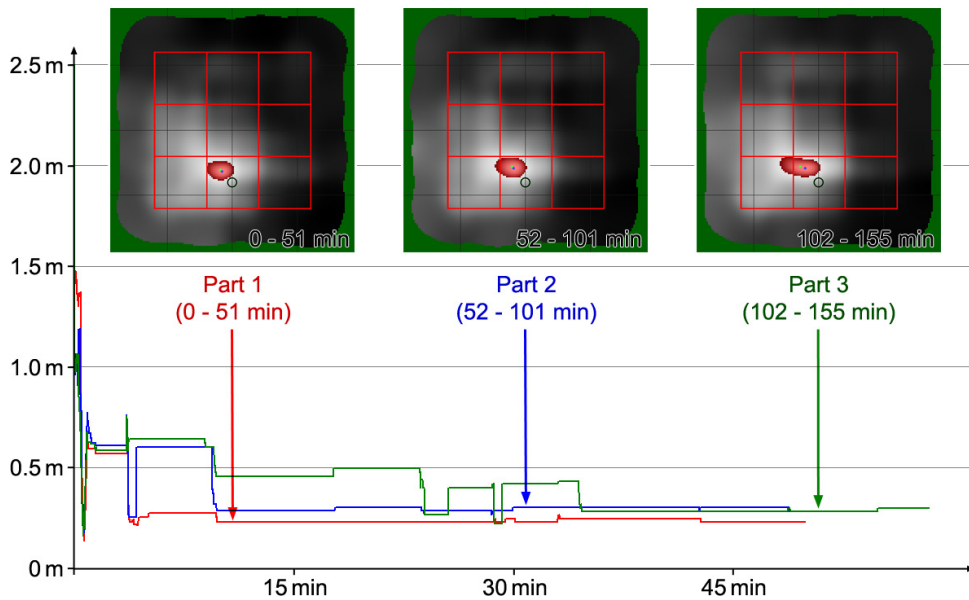


Fig. 10. Stability of the mapped structures. Three concentration gridmaps are shown created from the data of different parts of the same experiment. The graph shows the distance between the maximal concentration cell and the centre of the gas source. Each part corresponds to two completed cycles as defined in Section 4.2.

tion maximum fell below some selected thresholds. This analysis is intended to show the localisation capability of the gridmaps after stabilisation, so the gridmaps obtained before the assumed stabilisation time were omitted from the calculation. The stabilisation time was assumed to be 30 minutes. The rightmost column of the table gives the radius of the circle around the maximum concentration cell that contains the centre of the gas source in 68.3% of the cases. A value of 68.3% corresponds to the expected fraction of counts of a normally distributed random variable with a distance lower than the width of the distribution σ . Although the observed gas distribution is obviously not normal, this value provides an indicator of the localisation accuracy obtained by the different strategies.

Type	< 0.25 m	< 0.5 m	< 0.75 m	< 1 m	> 68.3%
Spiral	60.5 %	90.4 %	90.4 %	90.4 %	≤ 27.5 cm
Sweeping	23.5 %	42.7 %	69.0 %	80.0 %	≤ 56.5 cm
Braitenberg	29.3 %	48.3 %	62.1 %	68.1 %	≤ 103.5 cm

Table 1

Percentage of time steps for which the distance between the maximum concentration grid cell and the centre of the gas source was below the specified thresholds. The values were calculated from 6 spiralling trials with a total length of 754 minutes, 8 sweeping trials (976 minutes), and 11 Braitenberg trials (1720 minutes). The covered experimental area was 2.4×2.4 m² (Spiral), 2.8×2.8 m² (Sweeping), and 3.4×4.1 m² (Braitenberg).

The best performance was obtained with the spiralling strategy. The lower performance of the sweeping strategy could be explained by the fact that points are not traversed equally often from opposite directions with this strategy, as discussed in Section 4. The same applies to the reactive strategies, which are also prone to uneven coverage of the experimental area due to their tendency to get trapped by local concentration maxima. These results confirm that uniform coverage of the environment is advantageous for building accurate concentration gridmaps. It should be noted that the experimental area was different for the different strategies as referenced in Table 1.

The statistics of the gas source localisation trials given in Table 1 have to be treated with care because they summarise the results of trials in which different airflow situations were encountered. Depending on the actual airflow situation, it *can* be meaningful to approximate the location of a gas source by the location of the maximum average concentration. This applies to conditions where a sharply peaked distribution with an approximately circular region of highest intensity was found, which yielded the best localisation results. On the other hand, the location of maximal concentration does *not* necessarily provide a reasonable estimate of the location of a gas source. Poorer localisation performance was obtained when the area of high concentration formed a plume-like structure, as can be seen in Fig. 9. Such a plume-like structure might indicate a stronger airflow, which would tend to increase the distance of the concentration maximum from the gas source. The maximum average concentration is also expected not to be suitable for gas source localisation if the airflow situation changes drastically during an experiment. Consequently, the results suggest that it is possible to determine the quality of the location estimates from the *shape* of the mapped distribution.

6 Conclusions and Future Work

This paper presents a new technique for modelling gas distributions by constructing concentration gridmaps with a mobile robot. The method is discussed in terms of its suitability regarding the slow response and recovery of the sensors. Results of almost 70 hours of mapping experiments with three different exploration strategies are presented and analysed with respect to the time needed to represent the time-invariant structures of the gas distribution. This was achieved more quickly with reactive control strategies, which was found to be a consequence of the fact that the two reactive behaviours applied increase the time the robot spends in regions of high concentration. If the robot was driven along a predefined path it took longer to determine a stable concentration profile. On the other hand, the structures found then remained comparatively stable while it took considerably longer for the maps to stabilise if the reactive control strategies were used.

By comparing the location of the maximum concentration in the gridmap with the centre of the gas source, it was demonstrated that the location of the average maximum concentration can often be used to estimate the position of a source. Because the location of a gas source in a real-world environment does not necessarily correspond to the location of highest concentration, the suitability of this method for gas source localisation depends on the actual airflow situation in the inspected environment. Good localisation performance therefore requires also an analysis of the shape of the mapped distribution to distinguish different airflow scenarios.

It is assumed throughout the discussion in this work that stabilised concentration maps represent the true average distribution, possibly with a small distortion due to the memory effect of the sensors. This assumption is difficult to prove because there is no direct method for verifying the observed concentration field (“ground truth”) for the mapping experiments. There are, however, several indications of the consistency of the maps produced. This applies firstly to the applicability of concentration mapping for gas source localisation. The location of a gas source can be considered as some sort of ground truth about the real concentration field because it partly determines the gas distribution. Thus, the suitability for gas source localisation demonstrated in this work suggests that the obtained representation corresponds to the real physical gas distribution. Further indications are given by the self-consistency of the obtained concentration maps. This appears, for example, in the similarity of concentration maps that are created from different parts of the same trial. Fig. 10 shows a comparison of the concentration maps obtained in a sweeping experiment, where the data were divided into three parts each corresponding to two complete cycles shown in Fig. 7(a). This type of self-consistency could also be observed with data from different trials that were carried out in quick succession and using a different data acquisition strategy. A clear difference was also not discernible if separate concentration maps were created from the same trial using data from different sensors.

At present, only time-constant structures in the gas distribution were modelled by using temporal averaging. It would also be possible to model changing gas distributions by aging the measurements instead of averaging, so that older measurements gradually lose their weight. Other possible developments would include experimental comparisons of different exploration strategies for map building. Strategies based on the state of the map, e.g., by moving towards areas of high uncertainty, could also be considered. Future work could also include development of an actual source finding strategy based on the information about the peaks and plume-like structures extracted from these maps.

References

- [1] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press/Bradford Books, 1984.
- [2] C.T. David, J.S. Kennedy, J.S. Ludlow, and J.N. Perry. A Re-Appraisal of Insect Flight Towards a Point Source of Wind-Borne Odor. *Journal of Chemical Ecology*, 8:1207–1215, 1982.
- [3] T. Duckett. A Genetic Algorithm for Simultaneous Localization and Mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '2003)*, pages 434–439, Taipei, Taiwan, 2003.
- [4] A. Elfes and H. P. Moravec. High Resolution Maps from Wide Angle Sonar. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 1985)*, pages 116–121, 1985.
- [5] A. M. Farah and T. Duckett. Reactive Localisation of an Odour Source by a Learning Mobile Robot. In *Proceedings of the Second Swedish Workshop on Autonomous Robotics*, pages 29–38, Stockholm, Sweden, October 10–11 2002.
- [6] D. Gage. Randomized Search Strategies with Imperfect Sensors. In *Proceedings of SPIE Mobile Robots VIII*, volume 2058, pages 270–279, 1993.
- [7] A.T. Hayes, A. Martinoli, and R.M. Goodman. Distributed Odor Source Localization. *IEEE Sensors Journal, Special Issue on Electronic Nose Technologies*, 2(3):260–273, 2002. June.
- [8] J. O. Hinze. *Turbulence*. McGraw-Hill, New York, 1975.
- [9] H. Ishida, Y. Kagawa, T. Nakamoto, and T. Moriizumi. Odour-Source Localization in the Clean Room by an Autonomous Mobile Sensing System. *Sensors and Actuators B*, 33:115–121, 1996.
- [10] H. Ishida, A. Kobayashi, T. Nakamoto, and T. Moriizumi. Three-Dimensional Odor Compass. *IEEE Transactions on Robotics and Automation*, 15(2):251–257, April 1999.
- [11] H. Ishida, T. Nakamoto, and T. Moriizumi. Remote Sensing of Gas/Odor Source Location and Concentration Distribution Using Mobile System. *Sensors and Actuators B*, 49:52–57, 1998.
- [12] H. Ishida, K. Suetsugu, T. Nakamoto, and T. Moriizumi. Study of Autonomous Mobile Sensing System for Localization of Odor Source Using Gas Sensors and Anemometric Sensors. *Sensors and Actuators A*, 45:153–157, 1994.
- [13] H. Ishida, M. Tsuruno, K. Yoshikawa, and T. Moriizumi. Spherical Gas-Sensor Array for Three-Dimensional Plume Tracking. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, pages 369–374, Coimbra, Portugal, 2003.
- [14] O. Khatib. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 1985)*, pages 500–505, 1985.
- [15] A. Lilienthal and T. Duckett. A Stereo Electronic Nose for a Mobile Inspection Robot. In *Proceedings of the IEEE International Workshop*

- on *Robotic Sensing (ROSE 2003)*, Örebro, Sweden, 2003.
- [16] A. Lilienthal and T. Duckett. An Absolute Positioning System for 100 Euros. In *Proceedings of the IEEE International Workshop on Robotic Sensing (ROSE 2003)*, Örebro, Sweden, 2003.
 - [17] A. Lilienthal and T. Duckett. Experimental Analysis of Gas-Sensitive Braitenberg Vehicles. to appear to *RSJ Journal of Advanced Robotics*, 2004.
 - [18] Achim Lilienthal and Tom Duckett. Experimental Analysis of Smelling Braitenberg Vehicles. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, pages 375–380, Coimbra, Portugal, 2003.
 - [19] Achim Lilienthal, Andreas Zell, Michael R. Wandel, and Udo Weimar. Sensing Odour Sources in Indoor Environments Without a Constant Air-flow by a Mobile Robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2001)*, pages 4005–4010, 2001.
 - [20] T. Nakamoto, H. Ishida, and T. Moriizumi. A Sensing System for Odor Plumes. *Analytical Chem. News & Features*, 1:531–537, August 1999.
 - [21] R. A. Russell. Ant Trails - an Example for Robots to Follow? In *IEEE Int Conf. Robotics and Automation (ICRA 1999)*, pages 2698–2703, 1999.
 - [22] R. A. Russell. *Odour Sensing for Mobile Robots*. World Scientific, 1999.
 - [23] R. A. Russell, L. Kleeman, and S. Kennedy. Using Volatile Chemicals to Help Locate Targets in Complex Environments. In *Proceedings of the Australian Conference on Robotics and Automation*, pages 87–91, Melbourne, Aug 30- Sept 1 2000.
 - [24] R. A. Russell and A. H. Purnamadajaja. Odour and Airflow: Complementary Senses for a Humanoid Robot. In *IEEE Int Conf. Robotics and Automation (ICRA 2002)*, pages 1842–1847, 2002.
 - [25] R. A. Russell, D. Thiel, R. Deveza, and A. Mackay-Sim. A Robotic System to Locate Hazardous Chemical Leaks. In *IEEE Int Conf. Robotics and Automation (ICRA 1995)*, pages 556–561, 1995.
 - [26] R. A. Russell, D. Thiel, and A. Mackay-Sim. Sensing Odour Trails for Mobile Robot Navigation. In *IEEE Int Conf. Robotics and Automation (ICRA 1994)*, pages 2672–2677, 1994.
 - [27] T. Sharpe and B. Webb. Simulated and Situated Models of Chemical Trail Following in Ants. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S.W. Wilson, editors, *Proceedings of the 5th Conference on Simulation of Adaptive Behaviour*, pages 195–204, 1998.
 - [28] E. Stella, F. Musio, L. Vasanelli, and A. Distante. Goal-oriented Mobile Robot Navigation Using an Odour Sensor. In *Proceedings of the Intelligent Vehicles Symposium '95*, pages 147–151, 1995.