

Creating Gas Concentration Gridmaps with a Mobile Robot

Achim Lilienthal

University of Tübingen, WSI,

D-72076 Tübingen, Germany

`lilien@informatik.uni-tuebingen.de`

Tom Duckett

University of Örebro, AASS,

S-70182 Örebro, Sweden

`tdt@tech.oru.se`

Abstract

This paper addresses the problem of mapping the features of a gas distribution by creating concentration gridmaps from the data collected by a mobile robot equipped with an electronic nose. By contrast to metric gridmaps extracted from sonar or laser range scans, a single measurement of the electronic nose provides information about a comparatively small area. To overcome this problem, a mapping technique is introduced that uses a Gaussian density 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. The stability of the mapped features and the capability to use concentration gridmaps to locate a gas source are also discussed.

1 Introduction

This paper addresses the problem of modelling gas distribution in indoor environments by a mobile robot equipped with an electronic nose, 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 distant gas source, especially in environments where it is impractical or uneconomical to install a fixed array of gas sensors. The 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 [6, 15, 4].

Gridmaps were originally introduced to mobile robotics in the early 1980s as a means of creating maps using wide-angle measurements from sonar sensors [12]. The basic idea is to represent the robot's surroundings 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 any object. In the suggested approach, the cells in the gridmap correspond to 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 robots equipped with gas sensors, discussed as follows.

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 of the mobile nose used here were estimated as $\tau_r \approx 1.8$ s and $\tau_d \approx 11.1$ s [7]. Thus, considerable integration of successive measurements is carried out by the sensors themselves. The impact of this effect on the concentration mapping is discussed in Section 2.2. In addition, because diffusion is a very slow transport mechanism for gases in general [13], the distribution of gas molecules in an environment that is not strongly ventilated tends to be dominated by turbulence and convection flow, typically resulting in a jagged pattern of temporally fluctuating eddies [9, 14].

To overcome these problems, a mapping technique is introduced that permits integration of many gas measurements over an extended period of time. Spatial integration of the point measurements is carried out by using a Gaussian density function to extrapolate on the measurements, by 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.

In order to build complete concentration gridmaps, the path of the robot should roughly cover the entire space, although a uniform exploration is not necessary. To increase spatial accuracy it is also advantageous to pass particular points from multiple directions.

The location estimates required for map building were obtained by the external, vision-based absolute positioning system W-CAPS [8], which is briefly described in Sec-

tion 3. However, the results are expected to apply to any mobile robot equipped with a suitably accurate on-board positioning system, e.g., by carrying out simultaneous localisation and mapping with other sensor systems [2].

2 Concentration Gridmaps

Gridmaps extracted from a sequence of measurements are able to represent time-constant features of the measured quantity. In rooms with a constant unidirectional airflow these structures should be plume-like ones [5]. This paper, however, presents investigations performed in an unventilated room. Recent experiments showed that the concentration profile in an unventilated environment is often relatively stable over time as well. These plume-like structures may be caused by constant air streams that occur as a consequence of spatial temperature differences [16, 11].

In order to create gridmaps the cells have to be updated multiple times. Gas sensor readings represent, however, just 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:

- Despite the jagged gas distribution of temporally fluctuating eddies [9, 14], 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 we are searching for.
- 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.

2.1 Building Concentration Gridmaps

Assuming that the raw sensor readings R_t represent the real concentration at the current location, the normalised readings r_t are convolved using the radially symmetric two dimensional Gaussian:

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

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

- First, 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 $\vec{\delta}_t^{(i,j)}$ to the grid cell's centre $\vec{x}^{(i,j)}$ is calculated as

$$\vec{\delta}_t^{(i,j)} = \vec{x}^{(i,j)} - \vec{x}_t. \quad (2)$$

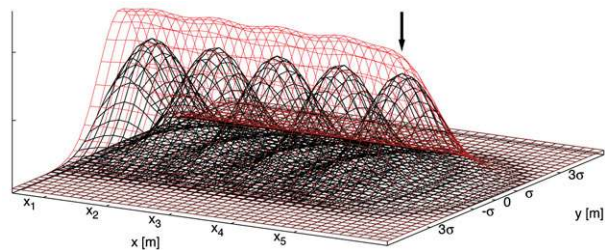


Figure 1: Sum of the likelihood function for a sequence of steps along a straight line.

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

$$w_t^{(i,j)} = \begin{cases} f(\vec{\delta}_t^{(i,j)}) & : \delta_t^{(i,j)} \leq R_{co} \\ 0 & : \delta_t^{(i,j)} > R_{co} \end{cases} \quad (3)$$

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

$$W_t^{(i,j)} = \sum_t w_t^{(i,j)} \quad (4)$$

and the total sum of weighted readings

$$WR_t^{(i,j)} = \sum_t r_t w_t^{(i,j)} \quad (5)$$

are calculated using the normalised readings r_t obtained from the raw readings R_t as

$$r_t = \frac{R_t - R_{min}}{R_{max} - R_{min}}. \quad (6)$$

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

- 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. 1 shows the sum of the weighting functions for a sequence of steps along a straight line (with a step width of 2σ). The figure shows the last five steps and the current one indicated by an arrow. One can see that the readings collected at positions x_1 - x_5 were spread by the mapping process along the driven path. While the spreading along the path is approximately independent of the chosen parameter σ , the spreading orthogonal to the path is determined mainly by this variable. Note that the sum of the likelihood function assigns a strong weight along the path driven, because the actual sensor readings contain information especially about this path due to the implicit integration of successive measurements.

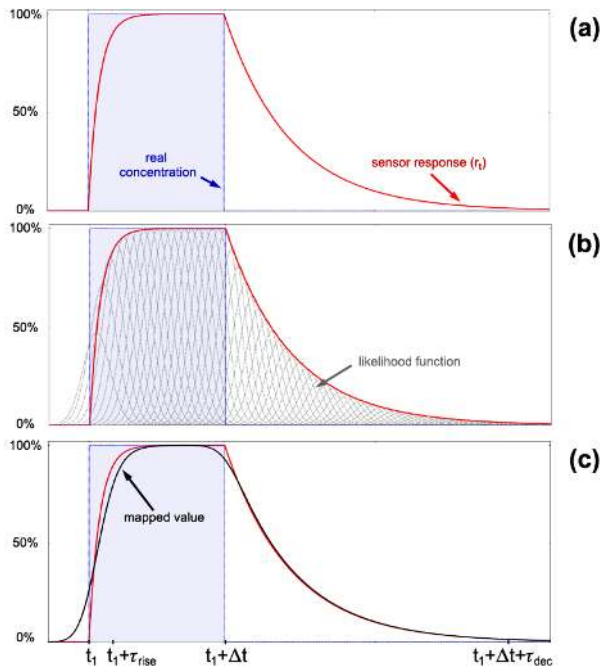


Figure 2: One-dimensional mapping of a rectangular step pulse with the technique described in eq.1-7.

2.2 Sensor Dynamics

It was mentioned that the sensor readings were assumed to represent the real concentration. As a consequence of this approximation 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. 2, which shows how a rectangular step pulse would be mapped by the described technique. In the upper part (a) the real distribution can be seen, which is a step pulse with a duration of $\Delta t = 10$ s. Here, the response of the mobile electronic nose used in this paper (the Örebro Mark III mobile nose [7]) is also shown. This curve was calculated using a first-order sensor model and the time constants $\tau_{rise} \approx 1.8$ s and $\tau_{dec} \approx 11.1$ s that were determined for this mobile nose by nonlinear least-squares fitting [7]. The middle part (b) of the figure shows the likelihood function (eq.1 with $\sigma = 1$ s) multiplied by the current reading. The readings were taken at a rate of 2 Hz in this example. Note that the applied value of σ corresponds to a distance of 5 cm if the described situation is considered to be caused by a robot that drives with a constant velocity of 5 cm/s through a 50 cm wide area of constant concentration. Finally in part (c) the normalised curve of the mapped values (eq. 7) is shown.

Comparing the real distribution with the course of the mapped values in Fig. 2(c) the asymmetrical shift as well as the blurring effect can be seen. This corruption is, how-

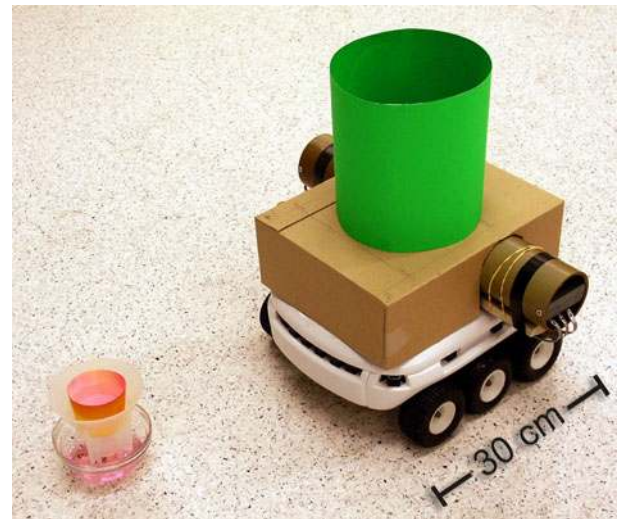


Figure 3: Koala Robot equipped with the Örebro Mark III mobile nose.

ever, not crucial. Due to the low speed of the robot, which never exceeded 5 cm/s during the experiments presented in this paper, the expected shift would be in the order of 10 cm at most. This effect is 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 passes a certain point from different directions. If this condition is fulfilled, the position of concentration maxima is represented closely by the described mapping process. The remaining effect of the blurred edges should be tolerable as long as the time-constant structures we are looking for can be distinguished in the maps created.

3 Experimental Setup

The experiments were performed with a Koala mobile robot (see Fig. 3) equipped with the Mark III mobile nose [7], comprising 6 tin oxide sensors manufactured by Figaro. This type of chemical sensor shows a decreasing resistance in the presence of deoxidising volatile chemicals in the surrounding air. 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 $8 \text{ m}^3/\text{h}$. The distance between the two sets of sensors was 40 cm.

To record the position of the robot the vision-based absolute positioning system W-CAPS [8] was applied, which tracks a distinctly coloured object mounted on top of the

robot (the green cardboard “hat” shown in Fig. 3). 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.

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, an odour 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. 3). 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.

4 Data Acquisition Strategy

Two different strategies were tested to collect concentration data. In one set of experiments the robot was driven along a *predefined path*, which was a rectangular spiral around the location of the odour source. The minimal distance to the centre of the source was 1 m, 0.75 m, 0.5 m, 0.35 m on the subsequent windings of the path. Along the straight lines a constant speed was applied because this was found to enhance the localisation capability [3, 11]. At the corners the robot was rotated slowly (10°/s) in order to minimise additional turbulence. A complete cycle including an inward and an outward phase lasted about 25 minutes. These cycles were repeated with a randomly chosen starting corner and direction at the start of each trial.

For a second set of experiments two different *reactive searching strategies* were applied in the manner of a Braitenberg vehicle [1]. Based on the stereo architecture of the mobile nose a direct sensor-motor coupling was implemented. Uncrossed as well as crossed inhibitory connections were used. In this way maximum wheel speed results if the sensed concentration is low, which in turn implements a simple sort of exploration behaviour. With uncrossed connections the robot turns toward higher concentrations (a behaviour that Braitenberg called *permanent love*) while the robot turns away from them with crossed connections (*exploring love*). Further details and the gas source localisation performance of such a “smelling Braitenberg vehicle” are discussed in [9].

5 Results

The introduced mapping algorithm depends on three parameters: the width of the gaussian σ , the cutoff radius R_{co} and the threshold W_{min} . The parameter σ has the greatest influence on the resulting gridmaps. Using $R_{co} = 3\sigma$ and $W_{min} = 1.0 \times$ (number of sensors) it was found that for small values of σ (≤ 10 cm) the spatial integration performed is not sufficient to bring out the time-invariant structure of the gas distribution [10]. Here, local variations along the path driven dominate the mapped distribution. Increasing σ causes these local maxima to be combined, and the time-constant structures to appear as contiguous spots. All of the subsequent gridmaps presented were created using a value of $\sigma = 15$ cm, which provides a reasonable compromise between the two effects.

5.1 Stability of the Mapped Features

Due to the local character of single gas sensor measurements, it takes some time to build a concentration gridmap. In addition to spatial coverage, a certain amount of temporal averaging is also necessary to represent the time-constant structure of the gas distribution.

In Fig. 4 the distance of the 90%-median (the median of the x- and y- coordinates of the area, which is defined by those cells with a value of at least 90% of the maximum) to the true centre of the gas source is plotted. The corresponding gridmaps were created with a cell size of 2.5×2.5 cm² using all six sensors and the parameters $\sigma = 15$ cm, $R_{co} = 3\sigma$, $W_{min} = 6.0$. Fig. 5 shows snapshots of these concentration maps that were created from data collected up to the time specified. Concentration values are indicated by shadings of grey (dark \rightarrow low, light \rightarrow high) while the values higher than 90% of the maximum are shown with different shadings (of red).

During the experiments where the robot moved along a predefined spiral path (section 4), it took approximately 25 minutes for the mapped structures to stabilise. This is indicated by the distance of the 90% median to the centre of the source which is shown in Fig. 4 (a). After a complete cycle (≈ 25 min) – meaning that the robot passed each point along its path two times – only small step-like variations remained. The transition between a preliminary mapped distribution, found after the robot has passed each point along its path just once, and the stable one is shown in example 1 and 2 in Fig. 5.

The presented results are approximately independent of the number of sensors used. Gridmaps extracted from data collected with 2 or 4 paired sensors reveal qualitatively the same features, thus indicating that it is not possible to accelerate the mapping process by using additional gas sensors mounted at approximately the same location. It is in-

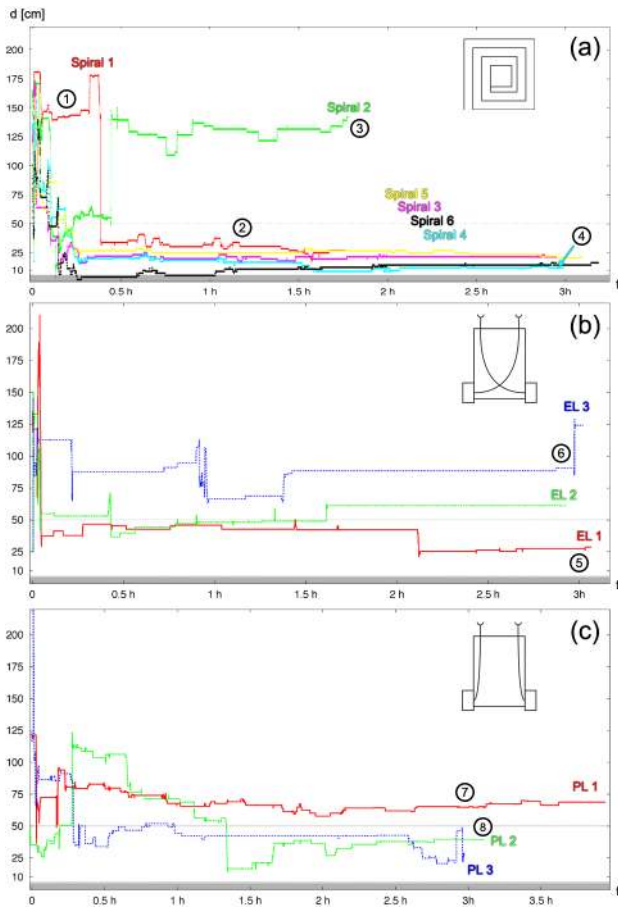


Figure 4: Distance of the 90% median to the centre of the gas source for 12 trials.

stead necessary to combine readings that are taken at different times in order to find the time-invariant structure of a particular distribution.

Slightly different results were obtained if the robot was controlled reactively as a Braitenberg vehicle, as explained in section 4. The distance of the 90% median to the centre of the gas source is shown in Fig. 4 either for crossed (b) and for uncrossed connections (c). For both cases stable structures could be determined more quickly, even though the average speed was lower (approx. 4 cm/s compared to 5 cm/s). On the other hand, the detected global maximum changed more often because the robot has a tendency to get stuck in local concentration maxima. An exception was found in the run “EL3” (see example 6 in Fig. 5). In this experiment the parameters of the applied strategy were chosen to enhance concentration peak avoidance. Consequently the robot failed to explore the complete available space and thus the area of maximum intensity could change drastically even after several hours if a new part of the room was explored. Concerning their suitability to create con-

centration gridmaps, no substantial difference could be determined between the two reactive strategies tested.

5.2 Gas Source Localisation

In the case of a gas distribution controlled purely by diffusion, the location of the gas source would correspond to the maximum in the concentration map. This assumption is, however, not fulfilled under realistic conditions due to the slow diffusion velocity of gases [13].

Nevertheless, the position of the 90% median can be used to estimate the location of the source in some cases. An exact agreement with the position of the gas source (indicated in Fig. 4 by the shaded area up to $R_{src} = 6$ cm) was observed only temporarily. On the other hand the error at the end of the experiment was < 50 cm in 8 out of 12 runs (see indicated examples 2,4,5 and 8) and < 75 cm in 10 out of 12 runs (see indicated examples 2,4,5,7 and 8). It is not guaranteed, however, to get a good estimate with this method.

Note that apart from the run “Spiral 2” (example 3) the concentration maps always showed a plume-like profile originating approximately from the location of the gas source (see example 2,4 – 8). Further experiments are necessary to investigate whether the asymmetric profile of the region of maximum concentration values provides a clue on the location of the gas source at its edge.

6 Outlook

This paper presents a new technique for modelling gas distributions by constructing concentration gridmaps with a mobile robot. It is discussed how the slow response and recovery of the metal oxide sensors affect these maps. Results of experiments carried out with three different exploration strategies are presented and analysed with respect to the time needed to represent the time-invariant structures of a particular gas distribution in the gridmaps created. This was achieved more quickly with reactive control strategies, which was found to be a consequence of the fact that the two different reactive behaviours applied increase the time the robot spends in regions of high concentration. However, the structures found in the gridmaps changed more often because the robot could temporarily get stuck in local concentration maxima.

If the robot was driven along a predefined path it took longer to determine a stable concentration profile compared to the reactive control strategies tested. On the other hand, the structures found then remained comparatively stable.

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

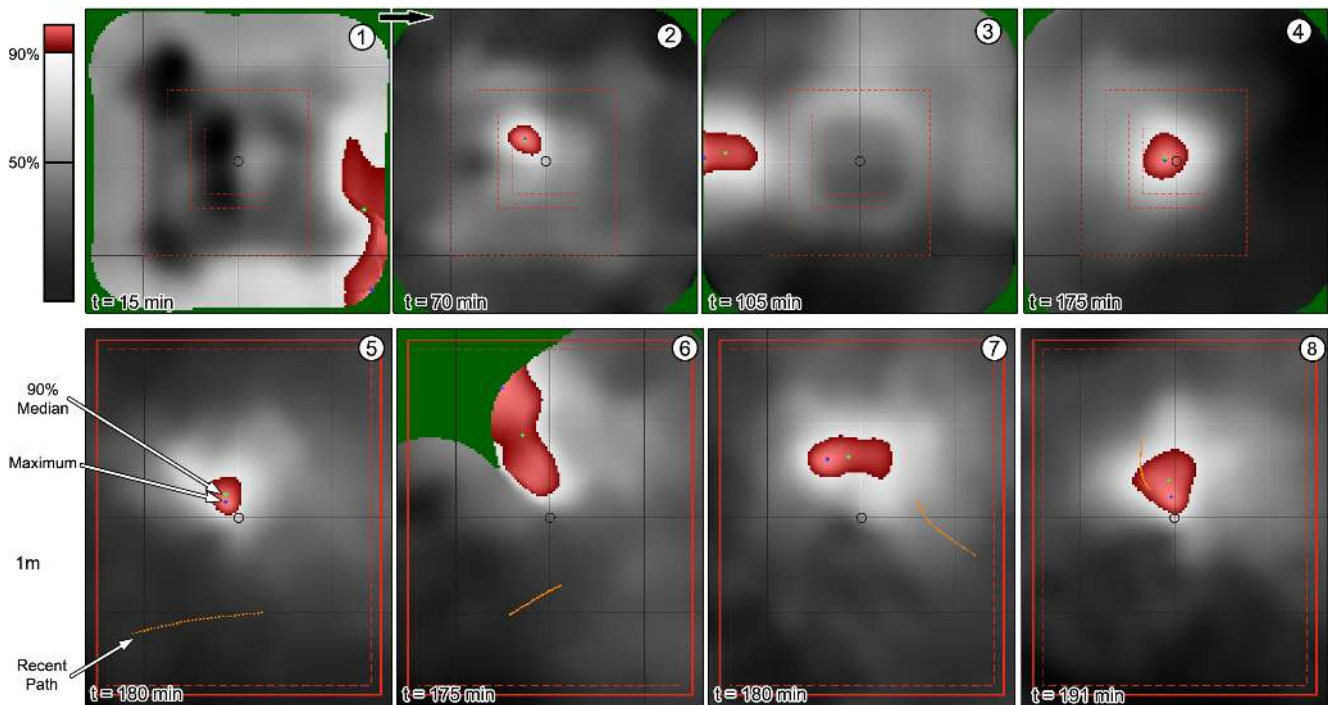


Figure 5: Concentration gridmaps created from data that were collected up to the time specified. Note that the point where the particular snapshot was taken is also referenced in Fig. 4.

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 plume-like structures extracted from these maps.

References

- [1] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press/Bradford Books, 1984.
- [2] T. Duckett. A genetic algorithm for simultaneous localization and mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'2003)*, Taipei, Taiwan, 2003.
- [3] 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.
- [4] A. Hayes, A. Martinoli, and R.M. Goodman. Swarm Robotic Odor Localization. In *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-01)*, volume 2, pages 1073–1078, Maui, Hawaii, USA, October 2001.
- [5] J. O. Hinze. *Turbulence*. McGraw-Hill, New York, 1975.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] A. Lilienthal and T. Duckett. Experimental Analysis of Smelling Braitenberg Vehicles. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, 2003.
- [10] A. Lilienthal and T. Duckett. Gas Source Localisation by Constructing Concentration Gridmaps with a Mobile Robot. In *Proceedings of the European Conference on Mobile Robots (ECMR 2003)*, Warszawa, Poland, 2003.
- [11] A. Lilienthal, M. R. Wandel, U. Weimar, and A. Zell. Experiences Using Gas Sensors on an Autonomous Mobile Robot. In *Proceedings of EUROBOT 2001, 4th European Workshop on Advanced Mobile Robots*, pages 1–8. IEEE Computer Press, 2001.
- [12] M. C. Martin and H. P. Moravec. Robot evidence grids. Technical Report CMU-RI-TR-96-06, The Robotics Institute, Carnegie Mellon University, 1996.
- [13] T. Nakamoto, H. Ishida, and T. Moriizumi. A Sensing System for Odor Plumes. *Analytical Chem. News & Features*, 1:531–537, August 1999.
- [14] R. A. Russell. *Odour Sensing for Mobile Robots*. World Scientific, 1999.
- [15] R. A. Russell, D. Thiel, R. Devezza, 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.
- [16] M. R. Wandel, A. Lilienthal, T. Duckett, U. Weimar, and A. Zell. Gas Distribution in Unventilated Indoor Environments Inspected by a Mobile Robot. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, 2003.